



DB2 UDB for iSeries - What's new?

Alison Butterill, IBM Canada



COMMON Belgium Congress
17 November 2004

© Copyright IBM Corporation, 2004. All Rights Reserved.
This publication may refer to products that are not currently available in your country. IBM makes no commitment to make available any products referred to herein.



DB2 UDB for iSeries Strategic Initiatives

Openness - Industry Standard Support


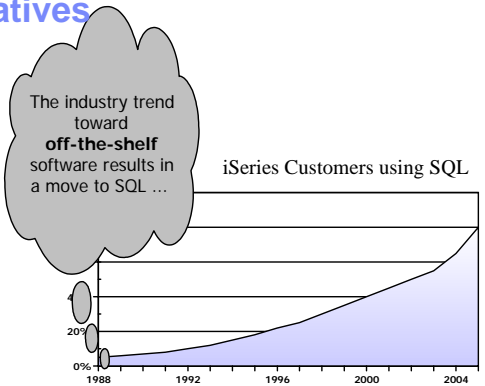
- Accomodate ISVs
- Portability/Compatibility
- Flexibility

Continued LEADERSHIP in database technologies

- Consistency across DB2 family
- Shared R & D across IBM Labs

Continued Leveraging of iSeries Strengths

- Availability
- Scalability
- Usability - Total Cost of Ownership
- Application Flexibility



DB2 UDB for iSeries Future Attractions

Application Flexibility & Portability

- Enhanced SQL Standards support
 - INTERSECT
 - GET DIAGNOSTICS
- Improved DB2 Family Compatibility
 - SEQUENCE Object
 - UTF encodings for Unicode
 - Column Encryption
- Native .NET Provider

Server Consolidation

- Database Migration Toolkits

OnDemand & Availability

- Online & Parallel Reorganize
- "Ragged" Save While Active
- Journal Enhancements

Performance

- SQL Query Engine enhancements
 - Star Join enhancements
 - Constraint Awareness
 - OnDemand Stats Generation
- Faster SQL Deletes
- Faster Stored Procedure Call Processing

Usability

- iSeries Navigator Enhancements
 - Reorg Table Manager
 - RI Constraint Manager
- Enhanced RPG SQL Precompiler

Application Development Enhancements

Application Development Enhancements - SQL & RPG

- **Enhanced ILE RPG SQL Precompiler**
 - ▶ V5 qualified subfield names support based on the RPG IV D-spec QUALIFIED keyword (only 1 level of qualification - arrays not supported)
 - ▶ Conditional precompilation: /If, /Define, /Include, and other directives
 - ▶ Support for Nested /Copy
 - ▶ Properly parse source members with sub procedures (eliminates out of sequence error)
 - ▶ Allow use of array data structures on blocked Fetch & Insert
 - ▶ Support for host structure in Update statements (cannot be array data structure)
 - ▶ Allow use of LIKEDS & LIKEREK (second optional parameter of LIKEREK not supported)
 - ▶ New COMPILEOPT parameter on CRTSQLRPGI to eliminate two-step compile process
 - ▶ Support for RPG logical (indicator) variable (application responsible for only allowing values of 0 or 1)
 - ▶ Some restrictions still remain...
 - Free-format SQL & support for qualified names (greater than 1 level)
 - Support for proper scoping of local variables (multi-pass)

Application Development Enhancements - .NET

- **Native .NET provider for iSeries**
 - ▶ Packaged in the next release of iSeries Access for Windows Client
 - ▶ Currently In Beta: www.iseries.ibm.com/access
- V5R3 Functionality Includes:
 - ▶ Connection Pooling
 - ▶ Stored Procedures
 - ▶ SQL and System Naming
 - ▶ LOB & Datalink Columns
 - ▶ Unicode
 - ▶ Isolation Level/Commitment Ctl
 - ▶ Tracing
 - ▶ Compression
- Restrictions:
 - ▶ Package Support (Ext Dynamic)
 - ▶ UDT Columns
 - ▶ Record Level Access
 - ▶ CMD/PGM call
 - ▶ Data Queues

Application Development Enhancements - .NET

- **DB2 Application Development Technology - DB2 LUW 8.2 "Stinger"**
 - ▶ Improved DB2 .NET Data Provider with better performance and support for DB2 UDB for iSeries servers - part of DB2 Connect
 - ▶ Enhanced DB2 add-ins for Microsoft Visual Studio
 - Integrate DB2 into Microsoft Development Tooling
 - Ease Creation of DB2 Objects:
 - Stored Procedures
 - Triggers
 - User Defined Functions



Application Development Enhancements

Java & Web Services

- **Support for Scrollable stored procedure result sets (all SQL interfaces)**
- **Toolbox JDBC Driver Enhancements**
 - ▶ Enhanced AutoCommit - isolation level other than *NONE can now be used
 - ▶ Enhanced XA (distributed transaction support)
- **Native JDBC Driver Enhancements**
 - ▶ Improved performance of character data processing
- **DB2 Web Services**
 - ▶ Actually available on Post GA V5R2 with WAS Express
 - ▶ Support for SQL Select, Insert, Update, & Delete statements along with Stored Procedure Calls
 - ▶ WOF component handles invocation and XML formatting out of results
 - ▶ More details at: ibm.com/servers/enabte/site/education/abstracts/db2web_abs.html

Application Development Enhancements

OLE DB, ODBC, CLI

- **Enhanced OLE DB provider**
 - ▶ SQL-only provider (IBMDASQL) including MTS Support
 - ▶ Support for commitment control
 - ▶ Support for SQL packages (ie, Extended Dynamic)
- **ODBC Driver Enhancements**
 - ▶ Enhanced MTS & XA Transaction Support
 - ▶ Enhanced AutoCommit
 - ▶ ODBC driver for Linux applications
- **CLI Improvements**
 - ▶ Column-wise binding to complement existing row-wise binding support
 - ▶ SQLGetInfo enhanced to return SQL_Keywords and Connection_Job_Name
 - ▶ New Connection Attributes for:
 - Setting minimum and maximum values for precision & scale
 - Unicode SQL statements (SQL_ATTR_UCS2 - also available on V5R2)
 - Hex literals designation

Application Development Enhancements

Cursor Sensitivity

- **Cursor sensitivity controls if database recognizes table changes (insert, update, delete) after a cursor has been opened**
 - ▶ **ASENSITIVE** - Default, DB2 can choose any setting. Updateable cursors always default to SENSITIVE
 - ▶ **INSENSITIVE** - Cursor will be insensitive to changes and behave as a read-only cursor by creating a temporary result file to process
 - ▶ **SENSITIVE** - Cursor implementation will have some sensitivity to changes. Prevents temporary copies and usage of DB2 SMP
- **Interface (get more details in Sensitivity article at: ibm.com/series/db2/sqe.html)**
 - ▶ **Embedded SQL: New keywords on DECLARE CURSOR & PREPARE ATTRIBUTES**
 - ▶ **OLE DB "Cursor Sensitivity" connection property (0-Asensitive/1- Insensitive/2-Sensitive)**
 - Static Cursors are always insensitive (same for ODBC)
 - Forward Only cursors can never be sensitive (same for ODBC)
 - ▶ **ODBC: SQL_ATTR_CURSOR_SENSITIVITY attribute on SQLSetStatementAttr (ODBC 3.0) or CURSORSENSITIVITY Connection keyword for ODBC 2.0**
 - ▶ **JDBC: "Cursor sensitivity" connection property (asensitive/insensitive/sensitive)**
 - Insensitive value only recognized with TYPE_FORWARD_ONLY result set type
 - Sensitive value only recognize with TYPE_SCROLL_SENSITIVE result set type

Application Development Enhancements

Improved SQL Diagnostics & Dynamic Cursor Control

- **Dynamic Cursor Control with Cursor Attributes on PREPARE**
 - ▶ Allows you to adjust the sensitivity, scrolling, lock holding, and result set attributes of an existing cursor
PREPARE s1 ATTRIBUTES :hv1 FROM :stmt
- **New GET DIAGNOSTICS statement**
 - ▶ Superset of all SQL error & diagnostic interfaces
 - ▶ Provides functionality & information similar to ODBC "SQLGet" functions like SQLGetConnectAttr & SQLGetStmtAttr
 - ▶ Examples:

Statement Info:

```
GET DIAGNOSTICS rcount = ROW_COUNT,  
                rcmd = COMMAND_FUNCTION,  
                rnbr = NUMBER,  
                rmore = MORE
```

Connection Info:

```
GET DIAGNOSTICS rcname = CONNECTION_NAME,  
                rcsts = CONNECTION_STATUS,  
                rnbr = DB2_PRODUCT_ID
```

DB2 & SQL Enhancements

New Column Types & Attributes

- Binary & VarBinary data types
- Large Decimals - up to 63 digits precision
- New Unicode support: UTF-8 & UTF-16 encodings
 - ▶ Enhances performance in unicode-based environments like Java by eliminating EBCDIC to Unicode conversion
 - ▶ Examples:
 - UTF-8: name CHAR(10) CCSID(1208)
 - UTF-16: name GRAPHIC(10) CCSID(1200)
 - UCS-2: name GRAPHIC(10) CCSID(13488)

New functions for Column Encryption/Decryption

- Encrypt & Decrypt SQL scalar functions
 - ▶ Requires the IBM Cryptographic Access Provider 128-bit product
 - ▶ Column Requirements:
 - Data Type Requirements: BINARY/VARBINARY, CHAR/VARCHAR FOR BIT DATA, BLOB, and DDS CHAR/VARCHAR with CCSID(65535)
 - Length Requirements:
 - Extra 8 bytes & total length must be rounded to 8-byte boundary (*replace 8 with 16, if BLOB or double-byte CCSID*)
 - 32-byte hint can optionally be stored with encrypted value
 - Example: 6-byte employee id with no hint needs to be stored in a VARBINARY(16)

```
CREATE TABLE emp(
  id VARCHAR(16) FOR BIT DATA,
  name VARCHAR(50))
```

```
SET ENCRYPTION PASSWORD = 'protect'
```

```
INSERT INTO emp VALUES(ENCRYPT('112233'), 'BOB SANDERS' )
```

```
SELECT DECRYPT_CHAR(id), name FROM emp
```

New functions for Column Encryption/Decryption

- **Encrypt & Decrypt SQL scalar functions**
 - ▶ **Password management**
 - When connected to remote databases, password is sent in the clear may consider using communication encryption (eg, IPSEC or SSL)
 - Best option would be store the password in OS/400 validation list
 - Encryption password would be stored in the encrypted data value and the list entry name could be table name or some other identifier
 - Validation list suggestions:
 - ◆ Make Public authority *EXCL
 - ◆ Programs would need authority to access validation list, best to use program adopted authority
 - Password Hint (if stored) retrieved with GETHINT function
SELECT GETHINT(id) FROM...
 - ▶ **Native programs could utilize by defining SQL View containing encrypt/decrypt and then opening SQL View as a logical file**

Sequence Object

- **Another DB2 construct that supports the automatic generation of column values**
 - ▶ Viewed as a superset of V5R2 identity columns
 - ▶ Generated values easily shared across tables
 - ▶ Can create constant sequence to be used as Global DB2 variables
- **Example:**

```
CREATE SEQUENCE order_seq
  START WITH 1 INCREMENT BY 1 NO MAX VALUE

INSERT INTO orders(ordnum,custnum)
  VALUES (NEXT VALUE FOR order_seq, 123)

VALUES NEXT VALUE FOR order_seq INTO :hostvar

UPDATE orders SET ordnum = :hostvar
  WHERE custnum = 123
```

Sequence Object

- Sequence values can be changed & altered with ALTER SEQUENCE statement
- Sequence values can be used to generate non-numeric key
CREATE SEQUENCE s START WITH 1001; ... ID='N'||CAST(NEXT VALUE FOR s AS CHAR(4))
- Customizable Sequence Attributes
 - ▶ START WITH & INCREMENT BY
 - ▶ MINVALUE & MAXVALUE
 - ▶ CYCLE & NO CYCLE
 - ▶ CACHE & NO CACHE - To improve performance, DB2 allocates a block of sequence values at the job/connection level.
 - ▶ ORDER & NO ORDER - ORDER ensures that values are returned in the actual order that they are requested independent of the job/connection. NO ORDER is the default. ORDER also disables caching.

```
CREATE SEQUENCE s1 CACHE 20 ORDER
Job1: NEXT VALUE FOR s1 => VALUE = 1
Job 2: NEXT VALUE FOR s1 => VALUE = 2
Job 1: NEXT VALUE FOR s1 => VALUE = 3
Job 1: NEXT VALUE FOR s1 => VALUE = 4
Job 2: NEXT VALUE FOR s1 => VALUE = 5
```

```
CREATE SEQUENCE s1 CACHE 20 NO ORDER
Job1: NEXT VALUE FOR s1 => VALUE = 1
Job 2: NEXT VALUE FOR s1 => VALUE = 21
Job 1: NEXT VALUE FOR s1 => VALUE = 2
Job 1: NEXT VALUE FOR s1 => VALUE = 3
Job 2: NEXT VALUE FOR s1 => VALUE = 22
```

INSERT & SELECT enhancements

- INSERT row expression values for blocked inserts
INSERT INTO table1 VALUES (1,'TESTING'), (2,'ADMINISTRATION')
- EXCEPT & INTERSECT operators
 - ▶ Return all rows that are in t1, but not t2
(SELECT cusnum FROM orders2003) EXCEPT DISTINCT
(SELECT cusnum FROM orders2004)
 - ▶ All rows that exist in t1 & t2
(SELECT cusnum FROM orders2003) INTERSECT DISTINCT
(SELECT cusnum FROM orders2004)
- Named Column Joins
SELECT * FROM t1 INNER JOIN t2 USING (c1, c2, c3)
- Lateral Correlation
SELECT * FROM t1, TABLE(SELECT * FROM t2 WHERE t1.c1=t2.c1) AS t3

Stored Procedure Enhancements

- **Faster call processing**
 - ▶ Up to 50% performance improvement on repeated calls within a connection
 - ▶ Improved processing of long SQL procedure names
 - ▶ Improved procedure resolution with unqualified procedure calls & *SYS naming
- **External Procedure Enhancements**
 - ▶ Support for Service Program Objects
 - ▶ Improved external program marking
- **Improved SQL Procedure (function & triggers) with more efficient code generation**
- **DB2 Family Compatability:**
 - ▶ WITH HOLD clause for FOR loop

```
FOR v1 AS c1 CURSOR WITH HOLD FOR
  SELECT firstname, lastname FROM emp
DO
  SET fullname = firstname || ' ' || lastname;

  INSERT INTO namelist VALUES (fullname);
  ...
END FOR;
```


Stored Procedure Enhancements

- **Result Set Enhancements**
 - ▶ Scrollable stored procedure result sets (cursor result sets only)
 - ▶ Result Set Consumer Control
- ```
RETURN TO CLIENT
DECLARE c1 CURSOR FOR WITH RETURN TO CLIENT SELECT * FROM t1

SET RESULT SETS WITH RETURN TO CLIENT FOR CURSOR x1

RETURN TO CALLER
DECLARE c1 CURSOR FOR WITH RETURN TO CALLER SELECT * FROM t1

SET RESULT SETS WITH RETURN TO CALLER FOR ARRAY :array1 FOR :hv1 ROWS
```



## Richer SQL function set

|               |                                                                                    |                                                                                          |
|---------------|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| REPLACE       | Returns a string where a given string is replaced with another string.             | <b>SELECT REPLACE('ABCXYZ','ABC','123') FROM t1</b><br>returns '123XYZ'.                 |
| EXTRACT       | Returns the specified portion of a datetime value.                                 | <b>SELECT EXTRACT(YEAR FROM datecol) FROM t1</b><br>returns the year from the date.      |
| INSERT        | Returns a string where one substring is deleted and another substring is inserted. | <b>SELECT INSERT('INSERTING',1,3,'XX') FROM t1</b><br>returns 'XXERTING'                 |
| REPEAT        | Returns a string composed of another string repeated n times                       | <b>SELECT REPEAT('ABC', 2) FROM t1</b><br>returns 'ABCABC'.                              |
| DAYNAME       | Returns the name of the day of the week.                                           | <b>SELECT DAYNAME(datecol) FROM t1</b><br>returns the name of the day from the date      |
| MONTHNAME     | Returns the month name from a datetime value.                                      | <b>SELECT MONTHNAME(datecol) FROM t1</b><br>returns the name of the month from the date. |
| RIGHT         | Returns the rightmost N characters from a string.                                  | <b>SELECT RIGHT('ABCXYZ',3) FROM t1</b><br>returns 'XYZ'.                                |
| TIMESTAMP_ISO | Returns a timestamp based on a date, time, or timestamp argument.                  | <b>SELECT TIMESTAMP_ISO(datecol) FROM t1</b><br>returns the timestamp from datecol.      |

- **Other new functions include:** MULTIPLY\_ALT, BIT\_LENGTH, OCTET\_LENGTH, DATABASE

## Miscellaneous DB2 Enhancements

- **Limit Improvements**
  - ▶ 256 Tables in a View
  - ▶ 1.7 TB table size limit
- **CPYFRMIMPF support for BLOBs (V5R2 PTF available: SI08237)**
- **DATE, TIME, & TIMESTAMP literals**
- **LOCAL TIME & LOCAL TIMESTAMP special registers**
- **Improved robustness of Cross-Reference Files**
- **Easier database creation script management with RUNSQLSTM**
  - ▶ Severity level of 20 now causes DROP errors to be ignored

## *Performance Enhancements*

### **DB2 Performance Enhancements**

- **Enhanced SQE Capabilities**
- **Caching of Deterministic UDF Calls - SQE Only**
- **Faster Stored Procedure Call Processing**
- **Improved SQL Function & Procedure Expressions**
- **Faster DELETE without WHERE clause statements**
  - ▶ CLRPFM used when used with "No Commit" isolation level
  - ▶ Special ALTER TABLE used when run with Isolation Level/Commit
- **Enhanced Database Monitor**
  - ▶ Ability to suppress monitor records for short-running SQL requests
  - ▶ Filter to eliminate DB2 generated SQL statements
  - ▶ More diagnostic details on events that force Full Open/ODP Creation
- **Parallel Reorganize**

## SQE Enhancements

- **Elimination of the following V5R2 restrictions:**

- ▶ View References
- ▶ UNION
- ▶ Subqueries
- ▶ Derived Tables & Common Table expressions
- ▶ Update & Delete-capable statements

- ▶ **Remaining SQE restrictions:**

- LIKE Predicate
- LOB columns
- Sort sequences
- ALWCPYDTA(\*NO) & SENSITIVE Cursors
- Logical File References
- Select/Omit Logical Files
- Non-SQL interface

## SQE Enhancements

- **Star Join Recognition (ie, Lookahead Predicate Generation)**

- **Check Constraint Awareness**

- ▶ SQE Stats Manager can use constraints to improve accuracy of Filter Factors & Cardinality
- ▶ SQE Optimizer will rewrite query to avoid execution when input data is out of range

Constraint Def: CHECK(col1 BETWEEN 1 AND 100)  
Original Query: SELECT \* FROM t2 WHERE col1=:hv

*Rewritten Query:*

*SELECT \* FROM t2 WHERE (:hv BETWEEN 1 AND 100) AND col1 = :hv*

- **RI Constraint Awareness**

- ▶ SQE Optimizer can rewrite query to eliminate join combinations

- **Immediate on-demand statistics generation**

- **Result Set Caching**

- **Customized Visual Explain for SQE access methods**

## Performance - New QAQQINI Options

| Option                     | Description                                                                                             | Possible values                          |
|----------------------------|---------------------------------------------------------------------------------------------------------|------------------------------------------|
| DATABASE_MONITOR_THRESHOLD | Allows only SQL statements with estimated runtime exceeding the threshold to be captured by the monitor | Integer, <b><u>2147483647</u></b> secs   |
| SQL_DBMON_OUTPUT           | Controls the types of SQL statements collected by the monitor based on the requestor                    | <b><u>*USER</u></b> , *ALL, *SYSTEM      |
| SQL_STMT_COMPRESS_MAX*     | Allows the user to adjust background access plan compression when using SQL packages                    | Integer(1-255, <b><u>2</u></b> )         |
| IGNORE_DERIVED_INDEX       | Allows SQE to process SQL statement even when an unsupported index type exists over the table(s)        | <b><u>*NO</u></b> , *YES                 |
| SQL_FAST_DELETE_COUNT *    | Allows user to control when & how V5R3 SQL Fast Delete support is used                                  | *NONE, <b><u>*OPTIMIZE</u></b> , Integer |
| CACHE_RESULTS *            | Allows SQE queries to use cached results sets from previously run queries                               | <b><u>*SYSTEM</u></b> , *JOB, *NONE      |

\* - Only available on V5R3, no PTFs for prior releases

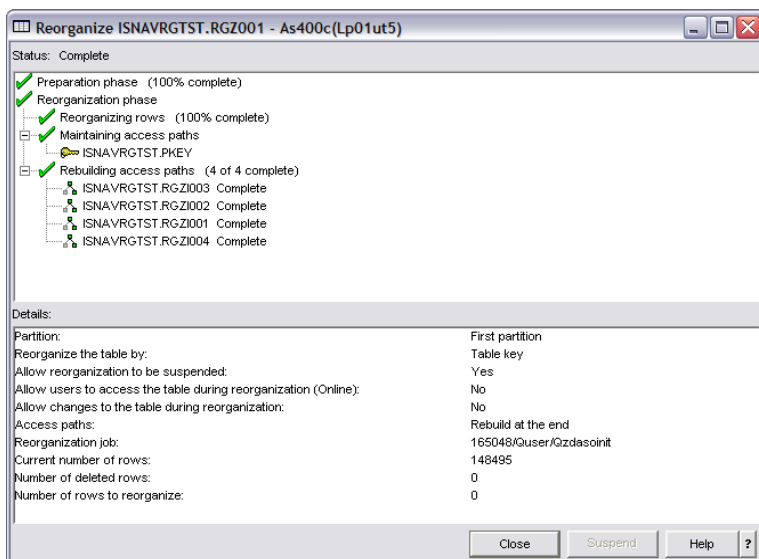
## OnDemand & Availability Enhancements

## Journal Enhancements

- ▶ "Ragged" Save While Active - SAVACTWAIT(\*NOCMTBDY)
- ▶ New CHGJRNOBJ command to adjust journal attributes on the fly
- ▶ New journal sequence maximum - \*MAXOPT3
- ▶ FRCAPPPTH(\*YES) & JRNSTATE(\*INACTIVE) settings are now ignored
- ▶ RCVJRNE now supports 1024 journal receivers
- ▶ NEW Defaults for Journal Commands & Settings
  - ▶ CRTJRRCV: THRESHOLD default changes from \*NONE to **1.5 GB**
  - ▶ CRTJRN: MNGRCV default changes from \*USER to **\*SYSTEM**
  - ▶ APYJRNCHG/RMVJRNCHG: CMTBDY default changed to \*YES
  - ▶ AUDIT Journal: Uses **RCVSIZOPT(\*MAXOPT1)**
  - ▶ SMAPP(EDTRCYAP): \*SYSDFT drops from 90 to **60 minutes**
- ▶ New Journal Minimal Data Checksum to increase data integrity of journal apply
- ▶ Journal Performance Improvements
  - ▶ Faster long-running ROLLBACKs
  - ▶ Remote Journal Super Bundling
  - ▶ RCVJRNE Performance Improved 15-20%

## Online & Parallel Reorganize

- New Reorganize capabilities activated with new parameter ALWCANCEL(\*YES)
  - ALWCANCEL(\*YES) requires file to be journaled
  - New parameter, LOCK, controls the concurrent access
  - If Exclusive lock not requested, then row order may be different & space may not be reclaimed
- Parallel capabilities rely on DB2 SMP licensed feature being installed & activated
- New Index Rebuild parameter RBDACCPH
  - RI & Unique indexes always maintained



# Online & Parallel Reorg Comparison

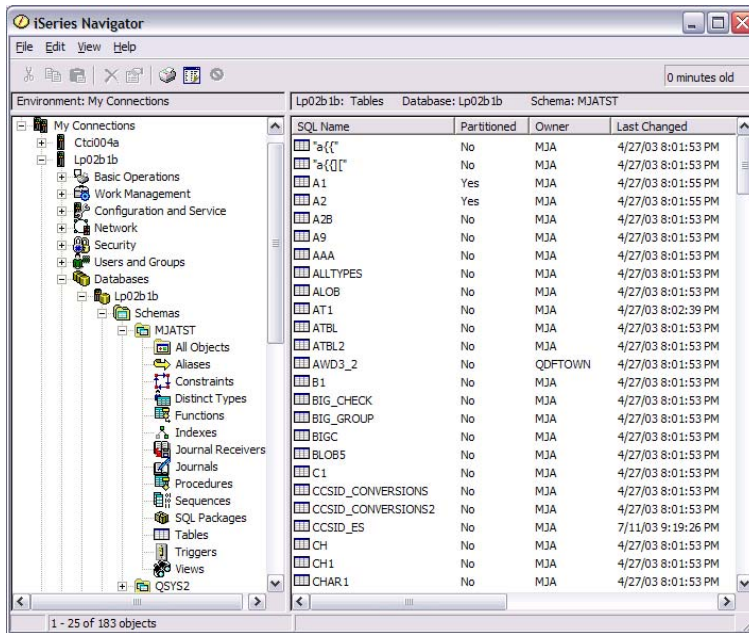
ALWCANCEL(\*NO)

ALWCANCEL(\*YES)

|                                                                      | KEYFILE (*NONE)                      | KEYFILE (*FILE or keyfile)           | KEYFILE (*RPLDLTRCD)                                     | KEYFILE (*NONE)                                          | KEYFILE (*FILE or keyfile)                               |
|----------------------------------------------------------------------|--------------------------------------|--------------------------------------|----------------------------------------------------------|----------------------------------------------------------|----------------------------------------------------------|
| Cancel and restart                                                   | No                                   | No                                   | Yes                                                      | Yes                                                      | Yes                                                      |
| Concurrent Access                                                    | No                                   | No                                   | Yes                                                      | Yes                                                      | Yes                                                      |
| Parallel processing                                                  | Only index rebuilds                  | Only index rebuilds                  | Data movement and index rebuilds                         | Data movement and index rebuilds                         | Data movement and index rebuilds                         |
| Non-parallel performance                                             | Very fast                            | Fast                                 | Very fast                                                | Slower                                                   | Slowest                                                  |
| Temporary storage                                                    | Double data storage                  | Double data storage                  | Journal receiver storage                                 | Journal receiver storage                                 | Journal receiver storage                                 |
| LIFO KEYFILE index processing                                        | N/A                                  | Duplicates reversed                  | N/A                                                      | N/A                                                      | Duplicate ordering preserved                             |
| Index processing (non-KEYFILE)                                       | Synchronous or asynchronous rebuilds | Synchronous or asynchronous rebuilds | Maintain indexes or synchronous or asynchronous rebuilds | Maintain indexes or synchronous or asynchronous rebuilds | Maintain indexes or synchronous or asynchronous rebuilds |
| Final row position exact                                             | Yes                                  | Yes                                  | Only if LOCK(*EXCL) and not restarted                    | Only if LOCK(*EXCL) and not restarted                    | Only if LOCK(*EXCL) and not restarted                    |
| Amount of CPU & I/O used                                             | Smallest                             | Next smallest                        | Smallest                                                 | More                                                     | Most                                                     |
| Variable length segment reorganize                                   | Good                                 | Good                                 | Worse                                                    | Worse                                                    | Worse                                                    |
| Allows referential integrity parents and FILE LINK CONTROL DataLinks | Yes                                  | Yes                                  | No                                                       | No                                                       | No                                                       |
| Allows QTEMP & Database Cross Ref Files                              | Yes                                  | Yes                                  | No                                                       | No                                                       | No                                                       |
| HABP replication cost                                                | Minimal - one journal entry          | Minimal - one journal entry          | More - journal entires for all rows moved                | Most - journal entires for all rows moved                | Most - journal entires for all rows moved                |

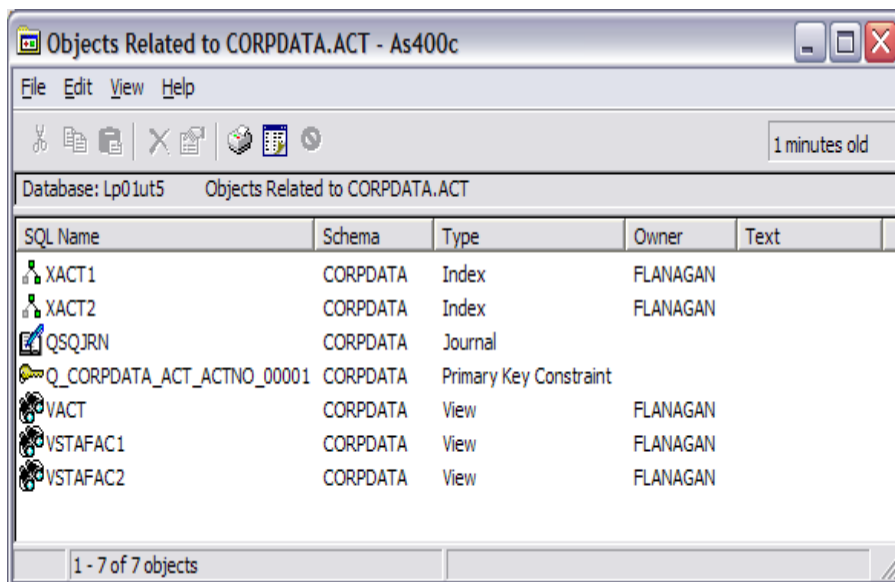
## Tools & Utilities

## iSeries Navigator Enhancements



- INDEX ANALYZER
- DB2 Object Type Folders
- Show Related (graphical DSPDBR)
- Constraint Management Interface (graphical WRKPFCSST)
- Reorganize Table Manager
- SQE Aware Visual Explain
- Run SQL Enhancements
  - Font Chooser
  - Integrated support for Graphical Debugger
  - Sortable Explainable Stmtms
  - Run & Explain support for Stored Proc calls
- New Show Command
- Improved Multi-tasking Support

## iSeries Navigator - Show Related



## iSeries Navigator - RI Constraint Management

Edit Check Pending Constraint ISNAVCPTST.CHK1CST - As400c(Lp01ut5)

File Options Actions Rows Help

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         | <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

## Rational XDE Modeler - iSeries Support Coming Soon

Modeling - Rational XDE

File Edit Navigate Search Project ClearCase Run PurifyPlus Modeling Diagram Format Tools Window Help

«database»  
**MJA2**

«table»  
**Table1**  
 NN Column1 : SMALLINT = 5  
 NN Column2 : TIMESTAMP

Model Explorer  
 Element Type  
 (Miscellaneous Models) DB2 iSeries Data Mod  
 MJA2  
 Table1  
 NN «column» Column1  
 NN «column» Column2  
 «unique constraint» Tab\_Constraint  
 «unique constraint» Tab\_Constraint  
 MJA2

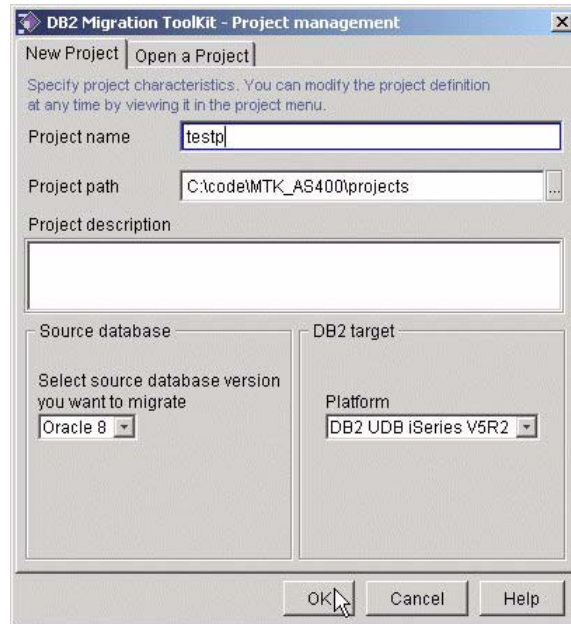
Properties  
 Storage Modeling Diagram <Diagram>  
 Property Value  
 Print  
 Dimension 0 - Inches  
 MarginBottom 1  
 MarginLeft 1  
 MarginRight 1  
 MarginTop 1  
 Orientation 0 - Portrait  
 PaperSize 0 - Letter (8.5 x 11...  
 ShowFooter True  
 ShowHeader True

Ready

## DB2 MTK Enhancements

### ■ Migration Toolkit Enhancements

- ▶ Support for V5R3 enhancements, Sequence objects big improvement for Oracle migrations
- ▶ Re-engineering of existing SQL Server & Sybase Migration Tool into the "Oracle" MTK architecture and interface
- ▶ Informix Database Support
- ▶ Free Download at:  
[ibm.com/servers/enable/site/db2/porting.html](http://ibm.com/servers/enable/site/db2/porting.html)



## SWG Utilities for DB2 UDB for iSeries

- DB2 UDB Extenders for iSeries
  - ▶ Support for XML Schemas & NOT FENCED UDFs
  - ▶ Leverages XML Toolkit for iSeries XML parser
  - ▶ Decomposition of XML now supports non-unique XML elements & attributes
  - ▶ UDFs now support language encoding parameter
- QMF for Websphere / Distributed Edition 8.1
  - ▶ Follow on to QMF for Windows
  - ▶ Newly GA'd- <http://ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&apname=iSource&supplier=897&letternum=ENUS204-028>
- DB2 Information Integrator
- DB2 Development Center
  - ▶ Support for Java & External Stored Procedures
- DB2 OLAP Server
- DB2 Web Query Tool (Websphere Based)
- DB2 Table Editor (graphical STRDFU)
- DB2 Data Propagator

## Positioning Partitioned & Materialized Query Tables

- **Partitioned Tables**
  - ▶ **ONLY** should be used in cases where single table limit of 1.7 TB or 4.2 billion rows will be exceeded
  - ▶ Partition tables should not be used to improve performance
    - Limited optimizer awareness of partitions, especially the CQE query optimizer
    - Fast delete of rows in a partition is supported
  - ▶ If you are considering using partitioned tables in V5R3, then contact the DB2 Solutions Enablement Team first
  
- **Technology Preview: Materialized Query Tables (Automatic Summary Tables)**
  - ▶ Only creation of MQTs supported
  - ▶ Query Optimizer is not aware of MQTs - MQTs will not be used by optimizer to improve query performance until future releases
    - Can manually query the MQTs

## Additional Information

- **DB2 UDB for iSeries home page** - <http://www.iseries.ibm.com/db2>
- **Newsgroups**
  - ▶ USENET: comp.sys.ibm.as400.misc, comp.databases.ibm-db2
  - ▶ iSeries Network (NEWS/400 Magazine) SQL & DB2 Forum - <http://www.iseriesnetwork.com/Forums/main.cfm?CFApp=59>
- **Education Resources - Classroom & Online**
  - ▶ <http://www.iseries.ibm.com/db2/gettingstarted.html>
  - ▶ <http://ibm.com/servers/enable/education/i/ad/db2/recentindex1.html>
- **DB2 UDB for iSeries Publications**
  - ▶ Online Manuals: <http://www.iseries.ibm.com/db2/books.htm>
  - ▶ Porting Help: <http://ibm.com/servers/enable/site/db2/porting.html>
  - ▶ DB2 UDB for iSeries Redbooks (<http://ibm.com/redbooks>)
    - [Stored Procedures, Triggers, and UDFs on DB2 UDB for iSeries](#) (SG24-6503)
    - [Preparing for the V5R2 SQL Query Engine](#) (SG24-6598)
    - [DB2 UDB for AS/400 Object Relational Support](#) (SG24-5409)
  - ▶ **SQL/400 Developer's Guide** by Paul Conte & Mike Cravitz
    - <http://as400network.com/str/books/Uniquebook2.cfm?NextBook=183>



# DB2 UDB Family Certifications

DB2

## Certified Database Associate - DB2 UDB Family (Test 700)

- **Website:** [ibm.com/certify/certs/dbdaudv81.shtml](http://ibm.com/certify/certs/dbdaudv81.shtml)
- **Education Resources:** [ibm.com/certify/tests/edu700.shtml](http://ibm.com/certify/tests/edu700.shtml)
- **Online Tutorial:**  
[www7b.boulder.ibm.com/dmdd/library/tutorials/db2cert/db2cert\\_V8\\_tut.html](http://www7b.boulder.ibm.com/dmdd/library/tutorials/db2cert/db2cert_V8_tut.html)

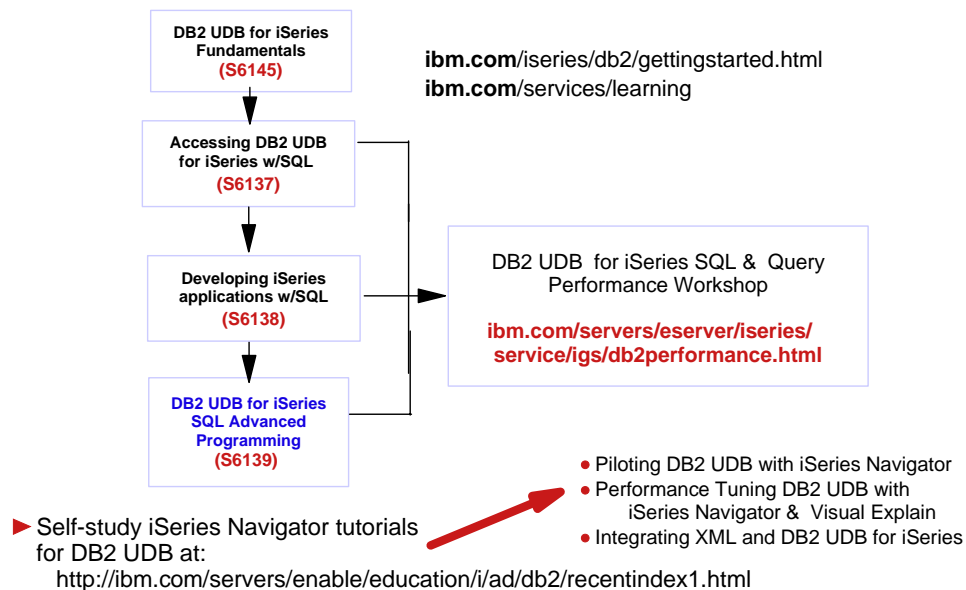
## Certified Application Developer - DB2 UDB Family (Test 703)

- **Website:** [ibm.com/certify/certs/dbapudv81.shtml](http://ibm.com/certify/certs/dbapudv81.shtml)
- **Education Resources:** [ibm.com/certify/tests/edu703.shtml](http://ibm.com/certify/tests/edu703.shtml)

➤ **Sample Tests:** [certify.torolab.ibm.com/ice](http://certify.torolab.ibm.com/ice)

➤ Exams were refreshed & updated for DB2 UDB for iSeries

## Education Roadmap



## Trademarks & Disclaimers

© IBM Corporation 1994-2004. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

|                   |            |
|-------------------|------------|
| AS/400            | IBM (logo) |
| AS/400e           | iSeries    |
| e (logo) business | OS/400     |
| IBM               |            |

Lotus, Freelance Graphics, and Word Pro are registered trademarks of Lotus Development Corporation and/or IBM Corporation. Domino is a trademark of Lotus Development Corporation and/or IBM Corporation.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.  
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.  
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.  
ActionMedia, LANDesk, MMX, Pentium and ProShare 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.  
SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.  
Other company, product and service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information in this presentation concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information in this presentation addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.