

A List of
DB2 Top Ten Lists

*In which we ponder numerous DB2 topics for
learning and amusement*

Craig S. Mullins

Corporate Technologist
NEON Enterprise Software
<http://www.neonesoft.com>

<http://www.CraigSMullins.com>



The Top Ten Lists



And now, from the home office in Sugar Land, Texas... a series of DB2 Top Ten lists about various topics ranging across the following subjects:

- Performance
- Coding
- Design
- Administration
- Management
- Features
- Tools

Top Ten SQLCODEs to Memorize



1. 000 / +100 successful / "no more rows"
2. -904 resource unavailable
3. -818 timestamp mismatch
4. -101 SQL statement too complex
5. -104 illegal symbol in SQL stmt.
6. -530 / -532 RI constraint violation
7. -803 unique violation (duplicate data)
8. -913 deadlock or timeout
9. -922 authorization failure
10. -805 program not found in plan



TAKE CONTROL

Top Ten DB2 V7* Features

* or V6 refresh



1. Real Time Stats
2. Scrollable Cursors
3. SQL Procedure Language
4. Declared Temporary Tables
5. Identity Columns
6. Limited FETCH
7. Stored Procedure Builder
8. Historical Statistics
9. External SAVEPOINTS
10. Deferred Data Set Creation



TAKE CONTROL

Top Ten DB2 V8 Features

1. 2M SQL Limit
2. Partitioning changes
up to 4096 Partitions, table-based partitioning, clustering separation
3. Stage 1 for Unlike Data Types
4. Data Partitioned Secondary Indexes
5. Sequences
6. Materialized Query Tables
7. Multi-Row FETCH and INSERT
8. Dynamic Scrollable Cursors
9. Recursive SQL
10. Online Schema Change



TAKE CONTROL

Top Ten DB2 9 for z/OS Features



1. Universal Table Spaces
and no more creation of simple table spaces (still supported if you have them tho')
2. Rows arranged for variable data
3. Index on Expressions
4. ORDER BY and FETCH FIRST on Subselects
5. Clone Tables
6. Not Logged Table Space ...but beware
7. BINARY / VARBINARY ...at last, a true binary data type!
8. Native SQL Procedure Language ...no more C compiler
9. SELECT from UPDATE, DELETE, MERGE
10. Implicitly Hidden Columns
11. And, OK, I guess I have to include it, pureXML

Top Ten Significant Features of DB2's First 20+ Years

1. Packages (V2.3)
2. Data Sharing (V4)
3. Referential Integrity (V2.3)
4. Type 2 Indexes (V4)
5. Segmented Table Spaces (V2.3)
6. Triggers and UDFs (V6)
7. Stored Procedures (V4)
8. Multiple Buffer Pools (V3...)
9. Breaking many limits (V8)
10. DATE / TIME data types (V1.3)



TAKE CONTROL

Top Ten Most Common DB2 Performance Problems



1. PEBCAK
2. Poorly coded SQL
3. Improper indexing
4. Bad program design
5. Bachelor programming syndrome
6. Improperly defined buffer pools
7. Index / table space needs to be reorganized
8. Improperly designed database structures
9. Copied code syndrome
10. RUNSTATS not up-to-date (or not even run)



TAKE CONTROL

Top Ten Steps to Proper Indexing

1. Index by workload, not by object
2. Build indexes based on predicates
3. Index most-heavily used queries
4. Index important queries
5. Index to avoid sorting (**GROUP BY, ORDER BY**)
6. Create indexes for uniqueness (**PK, U**)
7. Create indexes for foreign keys
8. Consider adding columns for IXO access
9. Don't arbitrarily limit number of indexes
10. Be aware of I/U/D implications



Top Ten Most Common Physical DB2 Database Design Mistakes



1. Relying on the defaults
2. Not basing the physical on a logical model
3. Over-relying on logical design
4. Normalization problems
(Over-normalized or too denormalized)
5. Not enough indexes
6. Indexing by table, not by workload
7. Too much (or not enough) free space
8. Failing to plan for data purging or archiving
9. Failure to share data
(not Data Sharing, but sharing data!)
10. Kludging



TAKE CONTROL

Top Ten Most Common Misunderstandings About DB2




1. "There's a problem with DB2!"
2. Using nulls can save space
3. DB2 is a "database"
4. DB2 is self-managing!
5. SQL is simple to learn and code (*properly*)
6. If it uses an index it doesn't need ORDER BY
7. Extents don't matter anymore
8. Using BPO only performs OK
9. PIECESIZE matches up IX and TS partitions
10. It depends!



TAKE CONTROL

Top Ten Most Under-Utilized Features of DB2

1. Real Time Stats
2. CASE statements
3. Table Expressions
4. DISTINCT Types
5. LOBs
6. User-Defined Functions
7. Triggers
8. Recursive SQL
9. Date/Time Data Types
10. Dynamic SQL ↓



Top Ten Extinct* DB2 Features

* or soon to be extinct



1. Type 1 indexes

2. The RCT



3. Host variables w/o a colon



4. SROD



5. Data set passwords



6. Simple table spaces



7. Manual stored procedure registration

8. Non-DRDA distribution



9. Hiperpools (and VPs in data spaces)

10. Denormalization (because of MQTs?)



TAKE CONTROL

Top Ten DB2 Annoyances

1. Changing the SQL Terminator for Triggers
2. No EXPLAIN parameter for CREATE TRIGGER
3. Fumbling thru the SQL Reference for Syntax
(specifically for SELECT)
4. SQL examples are too simple in the manuals
5. Lack of 100% Platform Compatibility
6. Managing Tables with LOBs
7. DSNZPARM documentation
8. Utilities cost extra
9. The database object is strangely implemented
10. It is so good that people take it for granted!



Top Ten SQL Mistakes

1. Syntax
2. The “flat file” mentality
3. Ignorance of *New* Features
(such as CASE and table expressions)
4. Fear factor
5. Copied code syndrome
6. Not coding for performance
(ignorance of Stage 1/Stage 2, indexing, etc.)
7. Too many columns! .
8. Not running the most efficient SQL statement
9. Improper “existence” checking
10. The Never-Ending Story!



What's Wrong With This SQL?

```
SELECT LAST_NAME, FIRST_NAME,  
       JOB_CODE, DEPT, PHONENO  
FROM   DSN8810.EMP  
WHERE  JOB_CODE = 'A' ←  
AND    DEPT = 'MIS'; ←
```



Existence Checking Options

```
SELECT 1
FROM SYSIBM.SYSDUMMY A
WHERE EXISTS (SELECT 1
              FROM DSN8810.EMP B
              WHERE LASTNAME = 'JONES'
              AND A.IBMREQD = A.IBMREQD);
```

```
SELECT 1
FROM DSN8810.EMP
WHERE LASTNAME = 'JONES'
FETCH FIRST 1 ROW ONLY;
```

Version 7+

<http://www.dbazine.com/db2/db2-mfarticles/custard1/view?searchterm=randy%20custard>

TAKE CONTROL




Top Ten Buffer Pool Tuning Steps



1. Do not use one large BP0 - *spread the wealth!*
2. Use BP0 for system objects only
3. Separate BP for indexes and table spaces
4. Set DWQT to enable trickle writing
5. Separate random and sequential
6. Use VPSEQT to control sequential usage
(increase for sequential, decrease for random)
7. "Peg" (small) frequently used tables in memory
8. Assign DSNDB07 to BP7 - *tune it for sorting*
9. Do not undersize hiperpool if you use them
(ROT: setup HP to be 2x - 3x the size of the VP it backs up)
10. Consider dedicated buffer pools (for "special" table spaces)

TAKE CONTROL

Top Ten DBA Excuses

1. It depends.
2. RTFM
3. "Did you fill out the form?"
4. "I'm busy."
5. "It's working as designed - leave me alone."
6. IBM says... 
7. You couldn't possibly understand why...
8. "That's what they said in class."
9. "Our standards say we do it this way."
10. "Because I'm the DBA, that's why!"



Top Ten Programmer Excuses

1. "There's something wrong with DB2!"
2. "But I copied that from another program."
3. "It worked yesterday."
4. "Isn't there something you can do to make it work?"
5. "But I can do that better in C; Java; etc."
6. "It works that way in Oracle; Access; etc."
7. "It's too late in the project to re-write that."
8. "But I heard somewhere it works this way."
9. "Why do I have to BIND every time?"
10. "DB2 is a hog."



Top Ten Management Excuses

1. "We're over-budget."
2. "The project is under-funded."
3. "Work smarter, not harder."
4. "You better work overtime on that."
5. "This comes from upper-level management."
6. "We're running behind schedule on this."
7. "You can't be out of the office that long."
8. "I read somewhere that isn't how it works."
9. "When I was a DBA/programmer/etc. ..."
10. "That is no longer strategic."



Top Ten Database Trends

1. From Many to "3"
2. Open Source
3. The Giant Sucking Sound
4. Complexity
5. Heterogeneity
6. Autonomic/Self-managing
7. **Marketing**
8. The Checkbox Wars
9. From VLDB to VHDB
10. Application Centricity



Top Ten Types of DB2 Tools

1. Change Manager
2. BIND and Access Path Analysis
3. Catalog Visibility
4. SQL Performance Monitor
5. System Performance Monitor
6. DBA Automation / Utility Automation
7. Table Editor
8. Recovery & Log Analysis
9. Database Structure Analysis
10. Application Restart Control





Top Ten Sources for DB2 Information

1. IBM manuals
2. IDUG
3. Local DB2 user groups
4. DB2 Magazine(s)
5. Vendor web sites/webinars
(neonesoft.com, etc.)
6. DB2 books
7. Web portals
(DB2portal.com, DBAzone.com, SearchDataManagement.com)
8. IBM DB2 Developer's Domain
9. Consultant web sites
10. Your co-workers!



Top Ten Books for DB2 Professionals

1. DB2 Developer's Guide
2. DB2 Developer's Guide
3. DB2 Developer's Guide
4. DB2 Developer's Guide
5. DB2 Developer's Guide
6. DB2 Developer's Guide
7. DB2 Developer's Guide
8. DB2 Developer's Guide
9. DB2 Developer's Guide
10. DB2 Developer's Guide







Craig S. Mullins
NEON Enterprise Software, Inc.
craig.mullins@neonesoft.com

<http://www.craigsmullins.com>


TAKE CONTROL

Available Now



DB2 Developer's Guide, 5ed

www.craigsmullins.com/cm-book.htm



DBA: Practices & Procedures

www.craigsmullins.com/dba_book.htm



Intelligent Solutions for Enterprise Data.