



EMBARCADERO
TECHNOLOGIES®

A List of DB2 Top Ten Lists

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

Craig S. Mullins
Director, Product Strategy

Embarcadero Technologies
100 California Street
San Francisco, CA 94111-4517
<http://www.embarcadero.com>

Craig.Mullins@embarcadero.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 statement**
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**



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



EMBARCADERO
TECHNOLOGIES®

Top Ten New 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**



NPI versus DPSI

EXP_DATE	CUSTNO	CUSTNAME ...	
07-11-2004	500	ACME RENTAL	100 175 500 600
07-19-2004	100	MULLINS AND ASSOC.	
07-25-2004	600	BMC SOFTWARE	
07-31-2004	175	SUPER BANK	
08-01-2004	400	SPUMCO	200 333 400 900
08-16-2004	333	SIGH BASE CORP.	
08-27-2004	200	BETH-ANN INC.	
08-28-2004	900	DIAL-A-DBA	
09-01-2004	800	VAN THE MAN	300 715 800 950
09-02-2004	715	DBAZINE.com	
09-04-2004	300	RENT-IT, INC.	
09-10-2004	950	VANDALAY INDUST.	

- 100
- 175
- 200
- 300
- 333
- 400
- 500
- 600
- 715
- 800
- 900
- 950

NPSI on CUSTNO

DPSI on CUSTNO



Recursive SQL Example

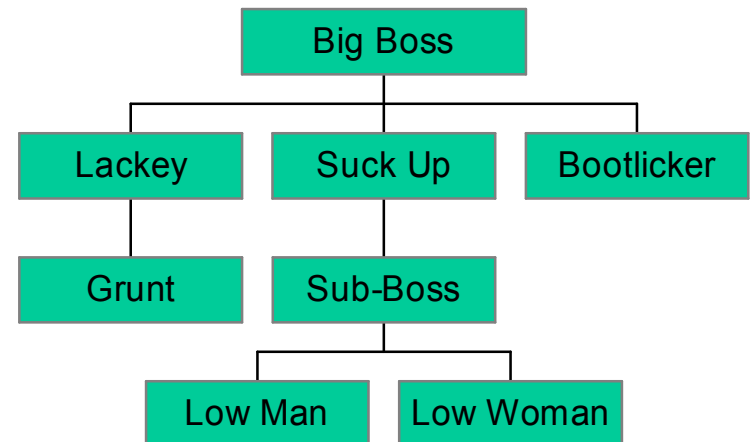
```
CREATE TABLE ORG_CHART
(MGR_ID    SMALLINT,
EMP_ID    SMALLINT,
EMP_NAME  CHAR(20));
```

```
WITH EXPL (MGR_ID, EMP_ID, EMP_NAME) AS
(
SELECT ROOT.MGR_ID, ROOT.EMP_ID, ROOT.EMP_NAME
FROM ORG_CHART ROOT
WHERE ROOT.MGR_ID = -1

UNION ALL


SELECT CHILD.MGR_ID, CHILD.EMP_ID, CHILD.EMP_NAME
FROM EXPL PARENT, ORG_CHART CHILD
WHERE PARENT.EMP_ID = CHILD.MGR_ID
)
SELECT DISTINCT MGR_ID, EMP_ID, EMP_NAME
FROM EXPL
ORDER BY MGR_ID, EMP_ID;
```

MGR ID	EMP ID	EMP NAME
-1	1	BIG BOSS
1	2	LACKEY
1	3	LIL BOSS
1	4	BOOTLICHER
2	5	GRUNT
3	6	TEAM LEAD
6	7	LOW MAN
6	8	SCRUB



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)




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)**



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**



Top Ten 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 BP0 only performs OK
9. PIECESIZE matches up IX and TS partitions
10. It depends!



Top Ten Most Under-utilized Features of DB2

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



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)





EMBARCADERO
TECHNOLOGIES®

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+



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)



1-3, 8



4-6



EMBARCADERO
TECHNOLOGI



Buffer Pool Thresholds

Variable Thresholds

VPXPSEQT: Assisting Parallel Sequential
(0% of VPPSEQT)

VPPSEQT: Parallel Sequential
(50% of VPSEQT)

Prefetch Disabled
(90%)

Static Thresholds


Data Manager Critical
(95%)

VDWQT: Vertical Deferred Write
(10%)

DWQT: Deferred Write (50%)

VPSEQT: Sequential Steal (80%)

Immediate Write
(97.5%)




Buffer Pool Guidelines

- **Do not default everything to BP0**
 - Explicitly specify a buffer pool for every TS and IX
- **Ideas:**
 - BP0 – system and catalog objects
 - BP1 – most table spaces
 - BP2 – most indexes
 - BP3 – most lookup tables
 - BP4 thru BP6 – reserve for tuning
 - BP7 – sort work
 - BP8+ – specialized processing
 - there is no “silver bullet” approach



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 DBA Tools

1. **Data Modeling/ER**
2. **Change Manager**
3. **SQL Performance Tuning**
4. **System Performance Monitor**
5. **Database Audit / Analysis**
6. **Catalog Visibility**
7. **Table Editor**
8. **Database Structure Analysis**
9. **Explain Plan Analysis**
10. **Database & Application Testing**




Top Ten Sources for DB2 Information

1. IBM manuals
2. IDUG
3. Local DB2 user groups
4. DB2 Magazine(s)
5. Vendor web sites
6. Web portals
(DBAzone.com, DB2portal.com, TDAN.com)
7. DB2 books
8. IBM DB2 DeveloperWorks
(<http://www.ibm.com/developerworks/db2>)
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



Covers Version 8



Contact Information


http://www.craigsmullins.com/dba_book.htm



Database Administration

Practices and Procedures

Craig Mullins



<http://www.craigsmullins.com/cm-book.htm>

- Craig S. Mullins
- *Director, Product Strategy*
- Embarcadero Technologies
100 California Street
San Francisco, CA 94111-4517
- <http://www.embarcadero.com>
- Craig.Mullins@embarcadero.com



EMBARCADERO
TECHNOLOGIES®

Check out the new DB2 for z/OS Portal

DB2PORTAL.com

<http://www.DB2portal.com>

A site documenting mainframe DB2 information for mainframe DB2 users.

DB2 topics

- [DB2 GENERAL INTEREST](#)
- [DB2 CERTIFICATION](#)
- [DB2 PERFORMANCE](#)
- [DB2 DATA SHARING](#)
- [DB2 VERSION INFORMATION](#)
- [DB2 WEB AND eBUSINESS](#)
- [DB2 & CICS](#)
- [DB2 DATABASE ADMINISTRATION](#)
- [DB2 TOOLS AND UTILITIES](#)
- [DB2 BLOG](#)
- [DB2 & ERP](#)
- [DB2 STORED PROCEDURES, TRIGGERS & FUNCTIONS](#)
- [DB2 CODING & DEVELOPMENT](#)
- [DB2 SECURITY](#)
- [DB2 DATA WAREHOUSING & BUSINESS INTELLIGENCE](#)
- [DB2 CROSS PLATFORM](#)
- [DB2 BOOKS](#)
- [DB2 SITES](#)



EMBARCADERO
TECHNOLOGIES®