



Craig S. Mullins

[Return to Home Page](#)

Dec 2005 / Jan 2006



zData Perspectives

by Craig S. Mullins

Top Five DB2 V8 Transition Do's and Don'ts

As more shops embark on the transition from DB2 V7 to V8, the time is right to review some of the biggest things to “keep an eye on” as you migrate your DB2 subsystems. DB2 V8 is indeed a very large new version of DB2 with a substantial number of changes. With that in mind, please understand that this treatment will not be comprehensive; but we will touch on some of the biggest transition considerations.

Number 5: Do Your Homework: The single biggest cause of migration and post-migrations problems is lack of upfront study and planning. Be sure to budget adequate time and resources to succeed with your migration. This includes training,

time to read the “what’s new” documentation, and preparation of a thorough migration plan.

Be sure to know the pre-requisites and prepare your system for V8. This means knowing things like you can only migrate from V7, you must be running on a zSeries box in 64-bit mode, and you need at least z/OS R3 (higher for certain features). Furthermore, you should understand the new modes of DB2 V8 (COMPAT, ENFM, and NFM), including how to migrate through the modes, what features are supported in each mode, and whether or not you can fallback in each mode. Your migration plan should include an estimate for how long you will run in each mode, and how you will introduce the modes across your test, QA, and production systems.

Also, be sure to run the health check job, DSNTIJPM, provided by IBM to help you identify unsupported objects. If you want to prepare early, then use job DSNTIJP8 from V7. It queries the catalog and identifies objects that will cause a migration failure.

Number 4: Bone up on CCSID Issues: DB2 V8 requires you to provide a DSNHDECP module that specifies valid, non-zero CCSIDs for single-byte character sets (SBCS) for both EBCDIC and ASCII. If you don’t know what a CCSID is, it stands for Coded Character Set Identifier, and it is a unique value that identifies a code page.

The CCSID should be consistent across members of a data sharing group. Additionally, your terminal emulators should have a compatible code page – DB2 provides automatic conversion for remote access when needed. If your CCSID is not correct, character conversion produces incorrect results.

The DSNTIJPM health check job we discussed earlier will check for various CCSID problems, but there can be others. Be sure to check work stations and emulator software for other CCSIDs.

Number 3: Be Ready for Unicode: Unicode is an issue for V8 because the DB2 catalog is converted to Unicode in DB2 and most internal DB2 processing is done in Unicode. This includes literal parsing, utility parsing, program preparation, optimization, authorization, internal names, special registers, and tracing. So you will need to prepare your environment to support Unicode. But don’t worry too much, you do NOT have to convert your application data to Unicode (of course, you can if you wish).

Unicode provides a unique number for every character, no matter what the platform, no matter what the program, no matter what the language. Unicode is a modern standard for character encoding; it is the foundation for the globalization of data storage and access because it handles just about any symbol you will encounter.

Number 2: Be Prepared for Performance Issues: Just running DB2 V8 means there is additional work to be performed, due to factors such as: increased code size, 64-bit address translation, and increased module linkage costs. IBM's stated performance objective is less than 10% average CPU regression — when comparing V7 and V8 subsystem with the same size buffer pools, hardware, workload, and so on.

Real storage requirements can also increase by at least ten percent as compared to V7. On migration to V8, buffer pools move above the 2GB bar – and the new limit for virtual buffer pools is 1 TB (instead of 1.6 GB). Additionally, the EDM pool, RID pool, SORT pool, and compression dictionaries are stored above the 2Gb bar; the IRLM is moved above the 2GB bar, too (with PC=NO no longer supported). So be ready to deal with increased memory usage requirements.

Of course, once you get to NFM you can take advantage of new features to improve performance, such as more Stage 1 predicates, MQTs with AQR, multi-row INSERT and FETCH, distribution stats, backwards index scan, volatile tables, and so on. So, IBM taketh away, but they also giveth!

Number 1: Know Those New Features: Of course, the biggest reason to migrate to V8 is to take advantage of the new features. There have been numerous articles written about new features, so I won't try to list them all here. Most tool vendors are introducing phased support for V8, so keep track of your vendors to make sure that they support the new V8 features when you need them. Features that seem to be later in the support cycle include support for 4,096 partitions, multi-level security, sequence objects, and MQTs.

From [zJournal](#), Dec 2005 / Jan 2006

.

© 2006, 2005 Craig S. Mullins, All rights reserved.

[Home](#).