



IBM DB2 11 Tools and Utilities: *Delivering Timely Value to Your Business*

By
Craig S. Mullins

President & Principal Consultant
Mullins Consulting, Inc.

<http://www.mullinsconsultinginc.com>

The tools and utilities used to manage and administer your DB2 environment can be just as important for delivering availability and performance as the database management system itself.

However, tools and utilities are not always valued accordingly and therefore may not be managed appropriately and with the proper due diligence. This white paper examines the importance of DB2 tools and utilities and why it is essential that they are kept in sync with new versions of DB2 for z/OS.



www.craigsmullins.com

Mullins Database Research
A division of Mullins Consulting, Inc.

15 Coventry Court
Sugar Land, TX 77479

tel 281-494-6153



Executive Summary

Many companies assume that their database tools vendors provide accurate and timely support for each and every release of DB2 for z/OS. System z customers plan their migration to a new release based on their business needs. They want to be able to take advantage of any features that are important to achieving business objectives without having to wait for a vendor to catch up and support their DB2 version level. There are many implications to a customer and their business with the lack of timely DB2 version support.

This white paper addresses the latest features of DB2 11 for z/OS and how IBM DB2 Tools provide support that is available on day one of General Availability. This paper also discusses the ongoing investment being made by IBM in the DB2 Tools market to re-assure customers that they will obtain the highest degree of ROI when selecting IBM as their DB2 Tools vendor.

Modern enterprise applications rely on DB2 for z/OS to manage and maintain the most important mission critical data required by the business. But the core DBMS does not provide all of the requisite capabilities for assuring robust performance, round-the-clock availability, and scalable applications. DB2 tools and utilities are also needed to deliver the full functionality required of enterprise systems. Without DB2 tools and utilities performance will suffer, availability will decline, and business objectives will not be met.

Indeed, DB2 Tools and Utilities are an essential component for reducing cost and improving the performance, availability, and governance of your DB2 databases and applications. Fortunately, IBM offers a full complement of robust and affordable tools and utilities for managing and administering DB2 for z/OS.

Are Your DB2 Tools and Utilities Up-to-Date?

Given the importance of DB2 tools and utilities, it stands to reason that they should be able to work with the latest features and functionality of DB2 for z/OS. After all, if your company has spent the time and effort to upgrade to a new version of DB2 there must have been a reason for doing so. Usually this reason involves achieving higher performance or using a new feature only available in the new version. Succinctly stated, IBM System z customers plan their migration to a new DB2 version based on business needs.

If your organization is using DB2 11 for z/OS in production, then you will want to take advantage of any and all features delivered by DB2 11. Implementing new database functionality can be a competitive advantage and wise organizations will want to adopt new features at their own pace to achieve important business objectives. And they do not want – and should not have – to wait for a vendor to catch up and support their DB2 version level.

And that brings us to an important criterion that must be understood before examining a vendor's ability to support a new DB2 version. Because many DB2 tools and utilities vendors have had difficulty keeping up to date with new DB2 versions they have developed a nomenclature for gauging their level of support that can be somewhat confusing. The vendors differentiate between supporting and exploiting a new DB2 version or release. Software that **supports** a new release will continue to function the same as before DB2 was upgraded, but with no new capabilities. So, if a DB2 tool, is touted to support a new ver-

sion of DB2, it can provide all of the services it did for the past release, as long as none of the new features of the new version of DB2 are used. A DB2 tool that **exploits** a new version provides the requisite functionality to operate on the new features of the new version of DB2.

DB2 11 for z/OS was released for general availability on October 25, 2013

To understand the difference between supporting and exploiting, let's take a look at an example. IBM added support for transparent archiving in DB2 11 for z/OS. A tool can support DB2 11 without operating on archived tables, but it must be able to access, manage and operate on archived tables in order to exploit DB2 11. So a vendor can say they support DB2 11, but that does not mean what you might think it means. Not understanding this nuance can cause problems. Be sure to understand the difference between supporting and exploiting a new version, as well as your vendor's exploitation schedule, before proceeding with your DB2 upgrade and any DB2 tool implementation. If you want to use new DB2 features you need to ensure that your tools and utilities exploit the new version... supporting the new version is insufficient.

There are many implications to a customer and their business with the lack of timely DB2 version exploitation. Application programmers may start to code programs using new features before the tools, and therefore the DBA group, is capable of supporting those features. This can result in a negative application impact affecting availability, retention of customers, reputation and revenue.

Furthermore, DB2 shops do not want to wait to plan their migration to the latest version of DB2 based on their tool and utilities vendor's time schedule. And the problem is exacerbated if more than one vendor supplies their tools and utilities for DB2 because now the schedule is impacted by multiple parties, each with different schedules for fully exploiting new versions of DB2. The longer the delay, the greater the negative impact on your organization's implementation schedule, which can impact application availability and corporate profitability.

Organizations using DB2 for z/OS do not want any loss of continuity or business resiliency from a ven-

dor that provides partial and/or late support for their critical business applications.

The Advantage of IBM's Expertise

Although there are quite a few vendors supplying DB2 for z/OS tools and utilities available in the marketplace, IBM should be at the top of your list when evaluating tools. IBM has the most expertise with the platform. IBM makes DB2 as well as the tools that manage it. This experience and knowledge helps to create robust management tools. And IBM understands the mainframe better than any other vendor... after all, IBM engineers and sells mainframes, too. In fact, IBM invented both mainframe technology, and the relational model upon which DB2 is based. So it should come as no surprise that IBM can offer synergy and exploitation of the System z architecture faster and more safely than other vendors.

IBM's investment in DB2 Tools is substantial and ongoing. Whereas some mainframe tools vendors have significantly reduced (or stopped) introducing new offerings and functionality, IBM has increased its investment in DB2 tools, with new offerings and functionality being delivered on a regular basis. IBM continues to make significant investments in product development, technical support, migration and implementation, and customer and industry partnerships. Your organization should consider deploying IBM's DB2 Tools and Utilities to reduce the cost, time, and effort involved in building and maintaining efficient DB2 databases and applications.

To verify this claim, let's take a look at IBM's exploitation of DB2 11 for z/OS in their tools and utilities.

IBM Tools & Utilities Support & Exploit DB2 11 for z/OS

'DB2 11 for z/OS became generally available on October 25, 2013. The broad array of new features and the promise of improved performance should cause customers to quickly migrate to this new version of DB2. Perhaps more importantly though, on October 1ST, IBM announced more than 25 new product releases that were upgraded to exploit new features and functionality... and delivered them for general availability on the same day as DB2 for z/OS, October 25TH.

If we examine some of the more prominent new features of DB2 11 and take a look at how IBM's tools and utilities exploit the new functionality it becomes evident just how much work was put into IBM's products.

Simplification of DB2 Migration

Perhaps the first thing that an organization upgrading to a new version of DB2 should look for is ways to improve the migration process. DB2 11 delivers out-of-the-box help here with a DB2 Catalog upgrade process that is up to 16 times faster than in previous versions.

It is a good idea to look for tools that aid in the migration effort, too. The **IBM DB2 Cloning Tool** can be used to simplify the migration process. You can use it to clone just the catalog and directory objects if required for migration test purposes, And you can use the Cloning Tool to clone one subsystem to another, including the translation of VCAT names from source to target. This is important if you want to clone within the same LPAR because there are log records that contain the VCAT name. Previously, the cloning process could fail when performing log apply to an arbitrary point in time. But as of DB2 11, the IBM DB2 Cloning Tool supports VCAT name translation during log apply processing on RESTORE SYSTEM LOG ONLY.

Another particularly thorny issue when migrating to a new version of DB2 is testing out your applications afterward to ensure that things are still working. You can use **IBM Optim Workload Replay** to replay production workloads to ensure that functional errors or scalability constraints are detected before hitting production.

Superior DB2 Performance

Performance and CPU reduction are additional driving factors that cause organizations to migrate to a new version of DB2. And as with most of the recent new DB2 versions, IBM boasts of performance improvements that can be achieved by migrating to DB2 11. Claims for DB2 11 out-of-the-box performance improvement range from 10 percent to 40 percent for different types of query workloads: up to 10 percent for complex OLTP and update intensive batch; up to 40 percent for queries. As usual, your actual mileage may vary. It all depends upon things like the query itself, number of columns requested, number of partitions that must be accessed, indexing, whether REBIND was performed, and so on.

So even though performance should improve when you migrate to DB2 11, there are no hard-and-fast percentages that you should expect to achieve because every shop is different. The standard operating procedure of rebinding to achieve the best results still applies.

From a DB2 tools perspective, **OMEGAMON Performance Expert** for DB2 can be used to measure the performance improvements you accrue as you move from DB2 10 to DB2 11, and as you REBIND your programs in DB2 11. Using OMEGAMON Performance Expert for DB2 you can trace and uncover just about everything you might want to know about a DB2 thread including elapsed time, CPU time, zIIP/zAAP usage, wait time, lock and latch activities, buffer pool usage, parallel access, and more. The new version of OMEGAMON can identify threads of autonomous procedures (native SQL PL) in relation to their calling “parents,” making stored procedure development and tuning easier. And OMEGAMON also boasts a new feature to export performance data to a spreadsheet making it easier to analyze the historical performance data collected.

Expanded RBAs and LRSNs

One of the bigger impacts of DB2 11 for z/OS is extended log record addressing in the form of larger Relative Byte Addresses (RBAs) and Log Record Sequence Numbers (LRSNs). DB2 11 expands the RBA/LRSN from its previous size of 6 bytes to 10 bytes, which increases addressability to a yottabyte. That should help out those DB2 shops that are approaching (or who have already hit) the maximum that a 6 byte RBA could store. By expanding its size your organization can avoid the outage that is required if the amount of log records accumulated exhausts the capability of DB2 to create new RBAs or LRSNs. Moving to the new extended log record addressing requires converting your BSDSs.

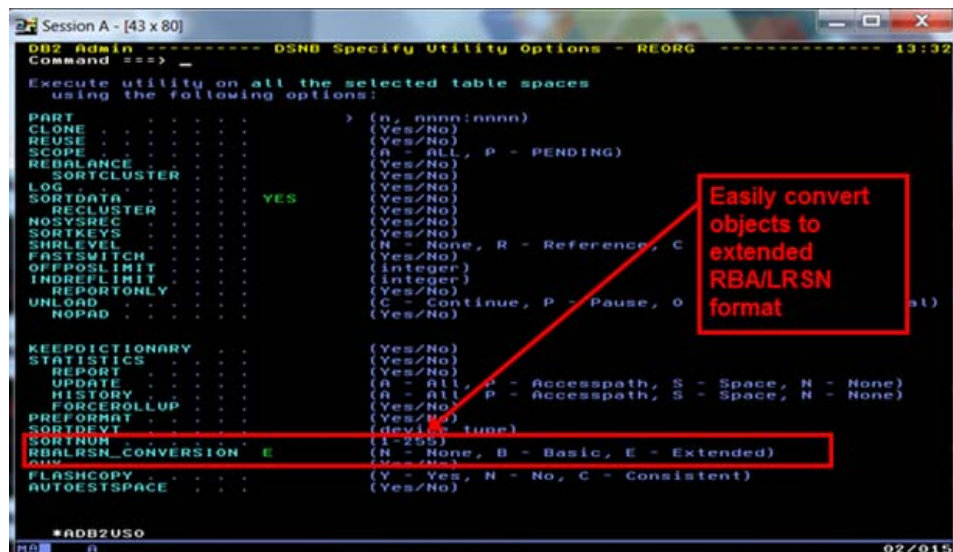


Figure 1. DB2 Admin Tool Supports Conversion of RBA/LRSN

As you might imagine, a larger size for log RBAs and LRSNs has a significant impact on tools and utilities that interact with DB2 logs and page formats. One particular area worth discussing is the automated mapping tables and table management capabilities of the **REORG** utility. The mapping tables required for an online REORG previously required a 6 byte RBA, which has of course been expanded to 10 bytes for DB2 11. As of DB2 11 NFM the REORG utility will no longer support the old format mapping tables. However, if you do not supply a mapping table, the REORG utility can now create one for you and then automatically drop it when the REORG concludes. Even better, if you get to NFM in DB2 11 and try to use REORG with an old format mapping table, the utility will recognize the old format and not use it. Instead a new mapping

table using the appropriate format with 10 byte RBAs will be created and used instead. IBM is to be lauded for this functionality because it saves a lot of work for organizations with hundreds or thousands of REORG jobs.

The IBM **Log Analysis Tool** has also been modified to support the new 10 byte RBA/LRSN values. After all, what good would such a tool be for a DB2 11 subsystem that has been converted to use the extended values if the tool only supported 6 byte RBAs/LRSNs?

OMEGAMON for DB2 Performance Expert and **DB2 Administration Tool** have also been modified to support the extended log RBA/LRSN. OMEGAMON has been augmented to use the 10 byte values in all of its modes including batch, real time UIs, near-term history and the performance database. And also, the IBM DB2 Admin Tool has been enhanced to support conversion of database objects to longer RBA/LRSN values (see Figure 1).

Specialty Processor Support

IBM has introduced several different types of specialty processors over the past few years. The basic idea of a specialty processor is that it augments the main general purpose CPUs and specific workloads are shuttled to the

specialty processors. The primary benefit of specialty processors is that workload that runs on them is not subject to IBM or independent software vendor (ISV) licensing charges. And, over time, more software is being enabled to exploit the capabilities of these processors.

The primary type of specialty processor that is relevant for DB2 workload is the zIIP, or Integrated Information Processor. It is used for processing specific types of distributed database workloads, but its usage has expanded in both DB2 and IBM's DB2 tools and utilities.

Several of the IBM DB2 utilities have had their ability to exploit zIIP processors improved for DB2 11. The **REORG**, **LOAD**, and **RUNSTATS** utilities can now offload more work to the zIIP:

- Up to 81% zIIP-eligible CPU with RUNSTATS COLGROUP
- Up to 40% zIIP-eligible CPU in REORG and LOAD with inline distribution stats

Additionally, if you are using LOAD REPLACE to clear out partitions (using dummy input) the process of fixing the NPIs has been zIIP-enabled. Tests have shown that as much as 98% to 99% of this work can be offloaded to the zIIP.

Other tools have also been augmented to improve their ability to leverage specialty processors. Both the **DB2 Query Monitor** for z/OS and the **Log Analysis Tool** for z/OS deliver enhanced zIIP support, which can reduce the cost of monitoring DB2 and analyzing log records.

The IBM DB2 Utilities and DB2 11 for z/OS

IBM clearly expended a great amount of effort into ensuring that its suite of DB2 utilities were upgraded to support DB2 11 for z/OS. But the updated utilities suite not only exploits new DB2 functionality, but also delivers new enhancements that:

- Reduce CPU usage
- Improve performance for utilities reducing elapsed time and minimizing resource consumption
- Maximize availability to further reduce application impact when running utilities
- Remove many previously existing constraints and limitations
- Simplify data management making it easier for DBAs to manage DB2 systems.

Let's take a moment to walk through some of the improvements made to the IBM utilities for DB2.

REORG

The **REORG** utility now offers a new option to defer the shadow index build until all the keys have been passed through the sort.

Why is that important? Well, let's consider the way it used to be done. For a partition-level REORG with non-partitioned secondary indexes (NPIs), shadow NPI processing consumed a significant portion of the processing cost. For example, if you have 10 partitions, and you want to reorganize partition 5, REORG would extract the index keys from the NPI for the partitions that are not being reorganized. The utility would then build the NPI with those index keys and then the data for partition 5 was reorganized. The index keys from partition 5 would be extracted and then inserted into the shadow NPI.

REORG in DB2 11 avoids that index key insert processing by not building the NPI first and then inserting the additional keys. Instead, REORG extracts the index keys for the partitions that are not being REORGed, sorts them, combines them with index keys for the partitions being reorganized, merges them together and builds the index shadow NPI in one step. This can be much more efficient and less costly.

This new capability was tested at one customer site using a 250 partition table space where 100 of those partitions were being reorganized. With this enhancement, REORG elapsed time was reduced by approximately 55%. But CPU usage increased because REORG had to pass all those additional index keys through the sort. Fortunately, there was a solution to reduce the CPU – the **IBM DB2 Sort** product. By adding DB2 Sort to the process the CPU cost went from a net increase to a net decrease, and elapsed time decreased even further, from 55% to 65%.

But this new REORG technique is not always optimal. For example, consider a billion row table space with one partition having only 5 rows in it. The new algorithm is not a fit for this scenario because REORG would need to pass one billion index keys through sort just to save the cost of inserting 5 index keys into the shadow NPI. Fortunately, there is a new REORG parameter and a DSNZPARM where you can enable it always, disable it or let the REORG utility decide whether to use this new option or not (this is the default).

Additionally, REORG utility improvements have optimized drain acquisition and thereby gains some relief from the impact of the SWITCH phase. The general idea here is that REORG will prevent new claims on all target partitions while it is waiting for drains. The result will be faster drain acquisition for partition-level REORGs. This is an automatic gain in DB2 11 and does not require specifying any new parameters. The drain improvement coupled with an up to 90% performance improvement for the SWITCH phase itself, both contribute to further minimizing application impact from REORG.

Another improvement when reorganizing a subset of partitions is the new DRAIN_ALLPARTS option. How does this work? Let's learn using an example. Consider a table space with 10 partitions, but you only want to reorganize one of them. This requires a drain on the partition (say partition #2) and after draining the data, the NPI must be drained as well. A deadlock can occur if an application is accessing another partition (say partition #5) which requires a claim on both partition #5 and the NPI. To reorganize partition #2 REORG must ask for a drain on the NPI and that application is going to have to commit and release its claim and then the REORG can begin processing. A deadlock can occur in this situation as the application waits to get a claim on partition 2 and the REORG holds a drain on it; and REORG is waiting for the application to release its claim on the NPI.

The solution is to use the DRAIN_ALLPARTS option, which indicates that the entire table space should be drained when reorganizing a single partition. Once the table space is drained, REORG can easily obtain a drain on the NPI and avoid the deadlock. As soon REORG gets the drain, the drain on the table space is released keeping just the drain on the one partition that is being reorganized.

Yet another new feature of REORG is the ability to repartition data online (SHRLEVEL CHANGE) with the REBALANCE parameter of the REORG utility. In DB2 10 and prior, rebalancing was only possible when specifying SHRLEVEL REFERENCE. Rebalancing data online complements another new DB2 11 feature of REORG, online ALTER of table space limit key values. With DB2 11, you can use ALTER to change partition limit key values and that change will become a deferred ALTER. It will take place the next time an online REORG is run.¹

REORG REBALANCE comes with improved resiliency in DB2 11 because it will no longer fail if it run out of rows to put into remaining partitions. As of

DB2 11, REORG will simply choose a limit key value for the remaining partitions. Additional enhancements for REORG REBALANCE include building compression dictionaries for all partitions and a new SORTCLUSTER option to sort data in clustering as well as partitioning order. This helps to avoid the AREO* condition which occurred when the partitioning key was not a superset of the clustering key.

There are quite a few additional improvements made to the REORG utility for DB2 11 including:

- The ability of REORG to physically empty PBG partitions (controlled by a new DSNZPARM named REORG_DROP_PBG_PARTS)
- Automated mapping table handling (as discussed earlier in this paper)
- The ability to reorganize without sorting data for those REORGs being performed for reasons other than to regain clustering of data
- Faster partition-level inline image copy during REORG
- New default parameters that better match industry best practices
- Improved REORG LOB processing
- Improved REORG LISTDEF processing

RUNSTATS

The **RUNSTATS** utility has been enhanced for DB2 11, too. As mentioned earlier, more RUNSTATS functionality has been made zIIP eligible in DB2 11, which will help cost containment for organizations using zIIP processors.

A new RUNSTATS option – RESET – can be used to reset all of the catalog statistics for a table. RUNSTATS RESET is an ideal way to remove stale statistics from the DB2 Catalog that could be negatively impacting the DB2 Optimizer.

Another nice statistics-related feature of DB2 11 is that the DB2 Optimizer can determine if there are missing statistics. The Optimizer knows what statistics are available and any missing statistics are populated by DB2 into a new, but unusable, DB2 Catalog table. However, **Optim Query Workload Tuner** can access and use the information from the new catalog table, interpret the information and update a statistics profile for the table. This profile can then be used by the **DB2 Automation Tool** which will detect that the profile has been updated and build the correct RUNSTATS job to go out and obtain the missing statistics that the DB2 Optimizer needs. With this integrated functionality among the DB2 Optimizer, Query Workload Tuner and the Au-

¹ But only if your partitioned table space is table-partitioned, not index-partitioned.

tomation Tool you can better assure that DB2 is working with the optimal set of statistics, and thereby improve the overall performance of your DB2 applications. And as with REORG, LISTDEF processing has been improved.

While discussing RUNSTATS, it is worth mentioning that inline statistics gathering in the REORG and LOAD utilities has been improved. In DB2 10 and earlier releases there were certain types of statistics that could not be gathered using inline processing, but this situation has been remedied in DB2 11, bringing inline statistics gathering up to par (in terms of functionality) with running an independent RUNSTATS. This enhancement includes histogram statistics, distribution statistics, and statistics collection on NPIs.

COPY and RECOVER

More robust backup and recovery features have been added to the IBM DB2 utilities, too. The first new capability is that the **COPY** utility can be run concurrently with an online REORG. Of course, contention arises if the REORG is in the middle of a SWITCH phase or trying to get a drain. But COPY SHRLEVEL CHANGE is now compatible with the unload, reload and log phases of the REORG utility.

An additional improvement for DB2 11 is that incremental image copy and COPY CHANGELIMIT no longer allocate empty image copy data sets.

For the **RECOVER** utility, numerous point-in-time recovery restrictions have been removed for LOB, XML, and partition-by-range Universal table spaces.

These restrictions involve dealing with deferred ALTERs, but also immediate ALTERs. The enhancements allow point-in-time recovery to a point before the materializing REORG, but will place the object into a hard REORG PENDING status. So a REORG is required immediately after the recovery to materialize the changes again. But the benefit is that you can now recover without worrying too much about the ALTERs that occurred during the recovery timeline.

Keep in mind, though, that not all point-in-time recovery restrictions have been lifted. PIT recovery restrictions are still in place for table space conversion and PBG table space deferred ALTERs. Additionally, two new restrictions have been added: PBG partition pruning and online DROP COLUMN.

LOAD and UNLOAD

IBM has made some important improvements to the **LOAD** and **UNLOAD** utilities. Parallel processing for data conversion is now possible when running LOAD SHRLEVEL NONE from a single input SYSREC data set (see Figure 2). This can be enabled using the PARALLEL option of the LOAD utility and it can reduce elapsed time especially when complex data types are being processed.

For LOAD SHRLEVEL CHANGE, not only can the conversion process be parallelized, but also the inserts to the target page sets. Elapsed time reduction in this case can be as high as 80% when using the new PARALLEL option.

Additionally, the IBM **DB2 Cross Loader** product has been augmented to support XML data.

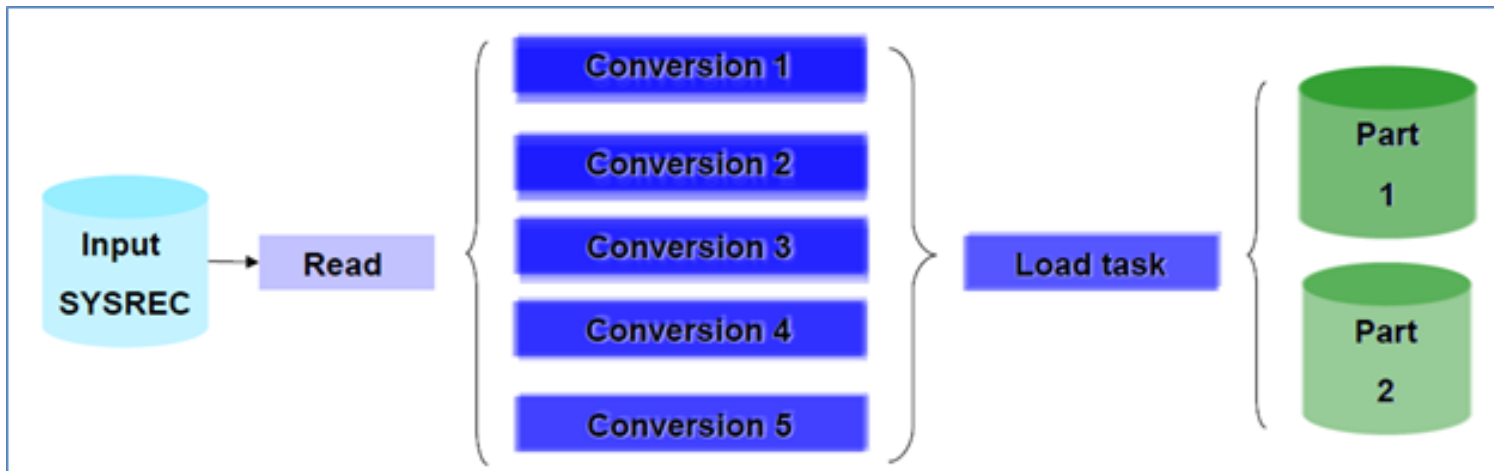


Figure 2. More Parallel LOAD Capabilities

Excluding NULL KEYS in DB2 Indexes

Another interesting new feature of DB2 11 is the ability to create indexes that skip keys where every key column is NULL. This is accomplished using the new option EXCLUDE NULL KEYS. There are, of course, exceptions whereby this clause is not supported, including unique indexes, indexes built on expressions, XML specifications, among others.

At any rate, the IBM DB2 utilities can greatly benefit from operating on such indexes. Benchmarking has shown CPU improvement for utilities with EXCLUDE NULL KEYS as follows:

- **LOAD** up to 12%
- **REORG** up to 25%
- **REORG INDEX** up to 88%
- **REBUILD INDEX** up to 72%
- **CHECK INDEX** up to 79%
- **RUNSTATS** up to 90%

And the IBM **DB2 Admin Tool** has been enhanced to support indexes that exclude NULL keys.

Database Schema Change

IBM continually looks for opportunities to enable online schema changes in new versions of DB2 and DB2 11 offers up some nice functionality in this realm. Perhaps the most interesting new capability is DROP COLUMN. Dropping a column from an existing table has always been a difficult task requiring dropping and recreating the table (and all related objects and security), so many DBAs just left unused and unneeded columns in the table. This can cause confusion and data integrity issues if the columns are used by programs and end users. Now, DROP COLUMN can be used (as long as the table is in a UTS). It is important to keep in mind that DROP COLUMN is implemented through an IBM

online REORG. Of course, there are some other restrictions on its use, but this capability may help many DBAs clean up unused columns in DB2 tables.

And IBM is exploiting this new capability right out of the gate with **DB2 Admin Tool** being able to exploit the new DROP COLUMN capability. The DB2 Admin Tool also supports other DB2 11 online schema change and utility features including:

- Online altering of limit keys
- Improved recovery for deferred schema changes
- Performance improvements for the LOAD utility, BIND, and DDL
- Automatic mapping of online REORG table
- Improved access path reuse with warnings
- Support for arrays
- Support for user-defined GLOBAL variables, array variables, and temporal special registers

The DB2 Admin Tool has also been augmented to support multi-target change management. This capability makes it easier for DBAs to distribute and track changes to each of many environments. Ensuring changes are propagated appropriately to multiple environments has been problematic for DBAs as they attempt to manage different object naming conventions, authids, and ownership issues.

With multi-target change management, the DBA can register a multi-target change once on a central system and then easily distribute the change to individual targets or a group of targets. Options are provided for masking at each target to adjust for naming differences. Changes are made automatically via DRDA. Each target's status and the overall status is reflected on the central system. This gives the DBA a single view of where each change has been deployed and applied.

The Bottom Line

DB2 tools and utilities provide valuable functionality for managing the performance, availability and operations of your DB2 for z/OS environment. Too often these solutions are viewed as afterthoughts instead of as vital components of your DB2 infrastructure. It is important that DB2 shops work with tools and utilities that are in sync with the DB2 release cycle, but also that are being modified regularly for new functionality and capabilities. Some vendors are better than others at providing both. And as we have learned in this white paper, IBM is providing both rapid support for new features of DB2 as well as new tool functionality. So if you want to be able to exploit features and functionality in DB2 without concern about whether your vendor tools have the necessary support, it is a good idea to consider using IBM's tool and utilities.

Published by:



Mullins Consulting, Inc.
15 Coventry Court
Sugar Land, TX 77479
Telephone: 281 494 6153
www.MullinsConsultingInc.com



About Mullins Consulting, Inc.

Mullins Consulting, Inc. is a database research and consulting company specializing in database performance, database administration, and database tools. The company was founded by Craig S. Mullins, a data management strategist, researcher, and consultant. He has nearly three decades of experience in all facets of database systems development including developing and teaching DB2 and SQL classes, systems analysis and design, database administration and system administration, and data analysis and modeling. He has worked with DB2 for z/OS since Version 1 and has experience working with other database technology including Microsoft SQL Server and IMS.

