

The Value of IBM's DB2 Utilities and Tools in 2011

A White Paper by Phil Grainger, Independent Consultant

January, 2011

Table of Contents

I. EXECUTIVE SUMMARY	3
II. ABOUT THIS WHITE PAPER.....	4
III. AVAILABILITY AND APPLICATION ACCESS TO DATA.....	5
IV. UTILITY PERFORMANCE AND TOTAL COST OF OWNERSHIP	7
V. STANDARDIZATION AND AUTOMATION	8
VI. CONTINUITY, INTEGRITY, AND RESILIENCY.....	10
VII. SUMMARY	10
ABOUT THE AUTHOR.....	11
ACKNOWLEDGEMENTS	11

I. Executive Summary

This paper examines how the IBM DB2 tools and utilities for z/OS have evolved and innovated in response of customer needs for support of critical DB2 for z/OS business applications. Topics covered in this paper include:

- How the IBM DB2 tools and utilities take advantage of the technology and features of System z and DB2 for z/OS; for example, supporting zIIPs, zAAP, IFLs and the latest versions of DB2 for z/OS and System z.
- How exponential data growth and complexity in applications have caused the DB2 tools and utilities to provide features and functions to not only respond to, but stay ahead of changing needs
- The cost of computing in today's IT environment. Customer demand increased performance and availability without increasing software or operating costs.

When DB2 was first developed, significant effort was directed at providing relational functionality and enhanced features and facilities. The DB2 utilities were largely only enhanced to keep pace with the changes in DB2 itself.

In the late 1980s/early 1990s, a fledgling independent software market started to emerge providing “add on” tools to complement DB2 processing (such as solutions to simplify the administration of schema management). It wasn’t long before these vendors identified a niche in providing complementary utilities beyond those initially provided as part of DB2 that would either provide additional functionality or enhanced performance.

However, in those early days, database administrators (DBAs) were sceptical and the feeling was that it was safer to rely on the IBM utilities rather than risk possible data loss or corruption that might impact the integrity of data and DB2 applications from using third-party DB2 utilities.

Over time, this “cost/benefit” balance changed. The need for additional functionality or, more often, the need for increased utility speeds when working with increasingly large amounts of data, meant that more DB2 administrators would begin using third-party utilities.

However, one of the biggest concerns in today’s 21st century is the need to reduce costs yet still remain competitive. Most companies today are faced with increasing pressures to contain or

reduce costs. New releases of DB2 deliver new features and functions that can be used to leverage additional value to businesses using DB2. If the adoption of these new features must be delayed due to late or selected support in third-party utilities, any possible advantages they may have provided are lost.

Therefore, each release of the DB2 Utilities Suite for z/OS has focused not only on lowering operating costs, but also on improving performance by reducing both CPU usage and elapsed time during execution, along with immediate support of new DB2 features.

For example, the DB2 Utilities led the DB2 utility market in 2006 with zIIP support that redirects index maintenance tasks to the zIIP processor to lower CPU costs. Further cost-saving innovations have continued with the addition of DB2 Sort for z/OS. DB2 Sort is a high-speed utility sort processor that also increases zIIP redirection. DB2 10 for z/OS has enhanced utilities such as RUNSTATS with zIIP support.

These additions and many other development investments have allowed the IBM utilities for DB2 to again become the utilities of choice and remove the need to purchase an additional set of DB2 utilities.

II. About this White Paper

The IBM DB2 for z/OS Utilities and utility management tools are the most comprehensive on the market today. By providing solutions to today's most complex business problems across many functional areas, they deliver the highest ROI for software tools investment.

DB2 for z/OS itself has a history of providing efficiency, resiliency, and growth for critical business applications. The IBM DB2 Utilities and utility management tools not only provide immediate support for new releases of DB2 for z/OS, but allow you to accelerate the return on the investment of migration and use of new features and releases. Staff will be able to leverage the new functions of every DB2 for z/OS release with more confidence and less time and error. New versions of DB2 applications can be moved into production faster with immediate out-of-the box savings and increased performance.

This paper looks at the some of the latest functionality provided by:

- The LOAD and UNLOAD utilities--used for populating DB2 tables and for unloading relational data
- DB2 High Performance Unload--used for unloading and extracting data for movement across enterprise systems or for in-place reorganizations
- The COPY and RECOVER utilities--used for securing backups and performing recoveries of DB2 objects
- The REORG utility--used to ensure that DB2 data is optimally organized in DB2 page sets
- The RUNSTATS utility--used to collect statistics on DB2 data enabling the DB2 Optimizer to make accurate cost-based access path decisions when executing SQL

- DB2 Sort for z/OS--provides enhanced sorting performance for DB2 utility processing
- DB2 Automation Tool--provides the ability to automate conditional or routine maintenance tasks based on user-defined criteria
- DB2 Utilities Enhancement Tool--manages thread activity and provides additional utility features and the ability to set company-wide standards for executing DB2 syntax
- DB2 Cloning Tool--provides for the replication of DB2 subsystems

This paper examines how IBM has responded to the needs of increased performance, while decreasing the use of system and staff resource. We discuss:

- Availability and application access to data
 - Fastest possible retrieval of information
 - Reducing or eliminating the amount of down time required for maintenance
- Utility performance and total cost of ownership (TCO)
 - Meeting or exceeding Service Level Agreements (SLAs) and implications for chargeback
 - Reducing CPU and elapsed time to achieve lowest possible TCO
- Automation and Standardization
 - Reducing the need for repeated tasks, thus reducing manual effort and possibilities for error
 - Ensuring consistency
- Continuity, Integrity and Resiliency
 - Ensuring data integrity during utility execution
 - Providing Day-1 support of new versions of DB2 for z/OS

III. Availability and Application Access to Data

This first section discusses some of the recent enhancements to the IBM DB2 Utilities Suite that ensure the fastest possible retrieval methods of information from DB2, as well as reductions or elimination in the amount of down time required for maintenance.

In many ways, the most cost effective and least disruptive utility to run is the one you don't actually NEED to run. IBM continues to enhance the online execution of utilities and the customizable automation of the associated DB2 tools, so that together they deliver the highest degree of availability suited to your DB2 for z/OS application needs.

DB2 9 made great strides with more online utility functionality and this investment has continued with DB2 10. Running utility functions online while applications are still processing the underlying data provides greater savings in utility elapsed time.

Large Objects (LOBs) that are typically used to store unstructured data can also be a concern to data accessibility and application availability. They introduced problems in the maintenance of DB2 range-partitioned table spaces. For example, in DB2 9 it was not possible to rebalance partitions if the table contained a LOB column definition. DB2 could not move the LOB data between partitions. DB2 10 resolves this issue by supporting the rebalancing of table space partitions, resulting in much simpler management of these table spaces—even where there is an unpredictable spread of data across the chosen key ranges. This enhancement also allows for manual alteration (followed by a REORG of partition limit keys if needed).

DB2 10 has also introduced the ability to reorganize LOB table spaces with full availability with a SHRLEVEL CHANGE REORG. This means that large object table spaces can now also be reorganized without having to take the data offline for the duration of the REORG. Other enhancements to the REORG utility allow partitioned table spaces that contain LOB columns to also be reorganized, even if data needs to be moved between partitions. These changes add to the usability and more importantly, the availability of managing large objects.

DB2 10 for z/OS allows image copies to be created with transaction consistency. Regardless of whether they are DB2 standard image copies, or copies taken using flashcopy technology, DB2 now ensures that any in-flight activity at the time of the copy is removed from the final copy. An image copy therefore contains only data from committed transactions and can be regarded as consistent.

The IBM DB2 Automation Tool for z/OS allows you to set up reoccurring utility jobs for conditional and routine DB2 maintenance tasks. The Automation Tool automates common utility tasks and allows you to define criteria for DB2 objects to determine if and when utility maintenance should occur, saving valuable system and staff resources. By taking advantage of the Automation Tool to create an image copy flashcopy data set in DB2 10, you can automate the copy and ensure you are getting the highest degree of availability possible.

One of the challenges to availability is changing the database structures without affecting application availability. Adding columns to tables, adding indexes to tables, and changing security can all be performed with minimal impact to existing applications and therefore minimal outage for the business. Online schema evolution addressed many types of structure changes.

Unfortunately, many changes have historically required tables to be unloaded, dropped, recreated and reloaded, as well as requiring other utility operations. These types of changes introduced a period of outage to the application and also an element of risk. The possibility also existed that a change needed to be reversed, which carries the same outage and risk and disrupts the application and business.

With DB2 10, IBM introduced a concept known as ***Deferred Schema Evolution***, where certain changes to physical structures of DB2 objects can be implemented with minimal impact to applications through a two-phase approach:

1. The change becomes a “pending alter” described in the DB2 catalog.
2. These pending alters are implemented by the IBM REORG utility at a later date. The REORG utility can be executed online (SHRLEVEL CHANGE) so that availability of active applications is unaffected.

Pending changes can be removed quickly and simply with the use of the ALTER statement.

Deferred Schema Evolution removes the risk of dropping and recreating DB2 objects. It also mitigates the risk and impact of taking outages, because all changes can be implemented with an online REORG.

IV. Utility Performance and Total Cost of Ownership

This section considers how the IBM DB2 Utilities are being continuously enhanced to ensure that Service Level Agreements can be met or exceeded, and helps reduce CPU usage and elapsed times to achieve the lowest possible TCO.

The cost of sorting data and keys is one of the most significant contributors to CPU usage in utilities, especially REORG, LOAD, REBUILD INDEX and RUNSTATS.

Enhancements to DFSORT and DB2 have resulted in significant portions of utility sort processing CPU becoming eligible for redirection to a System z Integrated Information Processor (zIIP). This lowers the utility sort processing time, as well as yielding reduction in CPU usage. Redirecting CPU usage to a zIIP means lowering overall central processor usage, which can lower software licensing costs and reduce the overall TCO.

IBM DB2 Sort for z/OS is designed for customers who want to further maximize their batch window availability and use of system resources during utility sort processing. The IBM DB2 Utilities Suite for z/OS takes full advantage of DB2 Sort to significantly reduce elapsed time and CPU usage during DB2 utility sort processing. During the execution of utility processing, DB2 Sort monitors and adjusts the allocation of system resources to optimize CPU processing, I/O performance, and memory usage and provides zIIP redirection of utility sort processing. DB2 Sort also helps to determine the availability of system resources and choose the appropriate amount of storage to allocate. It employs memory objects and data spaces to avoid disk access whenever possible, reducing elapsed time and optimizing the use of intermediate work space.

By choosing the most effective use of resources for each sort, rather than trying to use a “one size fits all” approach, DB2 Sort can increase the number of simultaneous sorts that can be run. This lowers utility elapsed times without necessarily increasing other resource utilization.

DB2 Sort also optimizes specific I/O devices to utilize the best transfer techniques and monitors the I/O transfer rates so that if performance objectives are not being met, it can adjust the use of devices to balance the I/O load and achieve the best elapsed time.

In addition, IBM DB2 High Performance Unload for z/OS offers performance options for extracting large amounts of data across the enterprise. Because HPU works outside of the DB2 address space, it does not compete for the same DB2 resources and typically maximizes performance by reducing processor usage and improves availability.

With DB2 10 for z/OS, the RUNSTATS utility takes advantage of possible zIIP redirection to lower its CPU consumption. RUNSTATS has also been enhanced in DB2 10 with improved sampling options. This allows the utility to collect statistics at a page sampling level instead of reading every row of a table. This enhancement makes a dramatic improvement in the I/O costs of collecting statistics and helps improve RUNSTATS performance.

Taking secure backups of DB2 data is something that every DB2 installation must plan for – just in case a recovery of the data is ever required. DB2 10 for z/OS and the DB2 Utilities Suite have increased its support of flashcopy technology to improve performance. By using the COPY or LOAD utilities with the SHRLEVEL CHANGE option, it is now possible to create consistent flashcopy image copies with full data availability.

Accessing DB2 data in high volume, transaction-based environments has always been a trade off between: 1) providing an indexing strategy that provides quick access to the data without overloading the DB2 tables; and 2) providing so many indexes that the index maintenance involved in changes to the data itself causes performance problems.

DB2 10 introduces a new organization type for table spaces known as HASH ACCESS, which enables rapid access to a single row, using a primary key (known as a hash key) instead of the traditional traversal of an index. The REORG utility in DB2 10 allows automatic calculation of the primary hash space size when executed with the AUTOESTSPACE keyword. During REORG, DB2 will determine the optimum size of the hash space from real-time statistics. This hash space is then used when the data is subsequently reloaded, removing the need to manually size this critical table space area. To improve the performance of loading data into a hash access table, the IBM DB2 Utilities Enhancement Tool provides the ability in DB2 10 to pre-sort the data in hash key sequence. This results in significant reductions in elapsed time and I/O costs.

V. Standardization and Automation

The importance of standardization and automation when running DB2 utilities is closely related to performance, availability and maximizing the use of resources to reduce errors and costs. IBM continues to invest in the DB2 Utilities Suite for z/OS for each DB2 release and to deliver additional non-release specific performance and availability features. It is when the DB2 Utilities are combined with the IBM DB2 utility management tools, such as DB2 Automation Tool for z/OS and the DB2 Utilities Enhancement Tool for z/OS that maximizes the efficiency of utility processing to yield the greatest TCO.

Most DB2 installations have hundreds, if not thousands, of utility job streams. It is imperative that the need for repeated tasks is minimized, reducing manual effort and possibilities of error.

It is also critically important to ensure consistency across the entire installation to allow consistent processing. The utilities should be run **only** when they are needed; therefore, a mechanism is required to control how the utilities are created with the desired settings and syntax, as well as scheduling when they need to be run.

The IBM DB2 Automation Tool for z/OS can automate DB2 utility tasks and reduces the need to build JCL manually. While helping to reduce the use of system and people resources, it creates re-useable actions that perform maintenance that is based on user-defined conditions. It also provides object profiles, which are lists of table spaces and indexes to be included for processing by utility profiles. Wildcards make the selection process easy and error-free.

Profiles can be created for the object, job, and condition, thereby reducing manual effort and increasing user productivity. This is especially helpful for REORG and COPY utilities, where triggering their execution can be determined based on user-defined conditions, reducing the manual creation of these jobs that can often be a source of errors.

The online capabilities of the REORG utility increases application availability, but they do not provide the complete solution. If a utility, such as REORG, encounters a DB2 thread running in a production environment, the only recourse is for the utility to wait, and possibly fail. Using the IBM DB2 Utilities Enhancement Tool for z/OS, DBAs can cancel conflicting application threads that would otherwise block critical utilities. Blocking these threads clears the way for batch applications and utilities to run. However, it is equally important not to inadvertently cancel threads that another application or utility is using. The DB2 Utilities Enhancement Tool provides numerous types of thread-filtering criteria to help specify precisely the threads to display or cancel.

The RUNSTATS utility has been enhanced in DB2 10 to increase the automation of collecting accurate statistics. Profiles can be created during the execution of the RUNSTATS utility, saving the current options that are specified, or DB2 can create a new profile by using RUNSTATS statistics that are already collected and held in the DB2 catalog.

The RUNSTATS utility supports collecting statistics while applications are accessing the data (using the SHRLEVEL CHANGE keywords). Try using the DB2 10 RUNSTATS with the Utility Syntax Monitor that is provided with IBM DB2 Utilities Enhancement Tool. This feature increases the automation and standardization of ensuring that DB2 utility syntax is followed according to company policies. The Utility Syntax Monitor can assist IT staff by enforcing company standards and policies for processing DB2 utilities. By defining rules within an intercept Policy, IT staff can control the authority to run IBM DB2 utilities, as well as enforce which parameters are used with specific DB2 utilities and objects. DB2 Utilities can also be purposely failed if they violate the rules that are defined within the policy.

The DB2 Utilities Enhancement Tool helps ensure that threads are immediately cancelled before the switch phase for DB2 10 online REORGs. Previously, many shops avoided using online REORG because they could not obtain the necessary drain to process writers in a timely manner.

The REORG enhancement in DB2 10 for z/OS allows you to cancel the writer prior to the switch phase. This can significantly reduce the outage and make online REORG possible to applications. Previously, the REORG was just taken offline. Using the Utility Syntax Monitor, you can ensure that your online REORGs are always run with the FORCE option.

VI. Continuity, Integrity, and Resiliency

This section discusses some of the utility processing factors that are not as straightforward to quantify and measure, such as Service Level Agreement performance measurements, business continuity, and data integrity. These factors are as important, if not more so, than others discussed thus far.

The DB2 Utilities provide day-one support for new releases of DB2 for z/OS. No lapse occurs in the utilities support for customers who are migrating to new releases of DB2. This is important for business resiliency.

DB2 for z/OS is constantly changing. Change and how it affects business resiliency is also critical when DB2 for z/OS delivers an enhancement or corrective action through the maintenance stream. The DB2 Utilities are able to ensure that any metadata changes caused by the DB2 for z/OS maintenance stream do not affect the data integrity or DB2 application failure that might be caused if using an Independent Software Vendor's (ISV) utilities. This is also an important consideration in a recovery situation, where the data has changed and the recovery point is prior to the change being implemented. The metadata must match or the recovery is not possible. The close tie of the DB2 Utilities to the DB2 z/OS engine ensures the consistency of data.

Business resiliency and continuity is also extended to the DB2 Tools that IBM delivers. No other vendor has the capabilities that are found in IBM's DB2 Cloning Tool for z/OS. The DB2 Cloning Tool can create a copy of sets of DB2 objects (cloning objects). By offering advanced, highly efficient functionality, it makes it simpler and quicker to clone a DB2 subsystem or a DB2 table space. The tool can automate the cloning process to provide usable DB2 clones within minutes, thus boosting efficiency and resiliency and freeing database administrators from performing time-consuming multi-step tasks.

DB2 Sort for z/OS also delivers features that increase the resiliency and business continuity when running DB2 Utilities. By using DB2 Sort, costs and interruption to DB2 applications can be reduced.. DB2 Sort can help to ensure that utilities are run the first time, because it dynamically allocates and adjusts the amount of temporary work space that is needed to run the utility. The result is a well-allocated set of resources.

VII. Summary

The IBM DB2 for z/OS Utilities Suite and DB2 utility management tools are the most comprehensive on the market. They are unique, providing critical solutions to today's most

complex business problems across many functional areas to lower the TCO of your software investment.

The IBM utilities and DB2 tools have a robust history of providing efficiency, resiliency, and growth for critical business applications--with a continued focus on investment for changing business demands.

The IBM DB2 utilities and utility management tools not only provide day-one support for new releases of DB2 for z/OS, but accelerate your time to value for migration and use of new releases.

About the Author

Phil Grainger is Consulting Director at Grainger Database Solutions Ltd – an independent provider of skills relating to IBM's DB2 for z/OS.

Phil began his career with DB2 back in 1987 as a Database Administrator for one of the largest European users of DB2 at that time. Since then, he has worked with every release of DB2 from Version 1.2 through DB2 10 for z/OS.

Phil has also held positions as DB2 tools Product Manager for a number of independent software vendors. He is active in many user groups and is a popular speaker, due to his in-depth DB2 knowledge and his recognized presentation skills.

Phil's detailed knowledge of DB2, together with his product management background and education/presentation abilities make him an ideal choice as the author of this white paper.

Acknowledgements

I would like to thank the following people at IBM for their assistance and expertise in providing invaluable information for this white paper. Without their collective help, this would have been a much more difficult task.

Craig Friske, IBM DB2 Utilities Software Developer

Calene Janacek, IBM Product Marketing Manager, DB2 Tools and Utilities

Roger Miller, IBM DB2 for z/OS Evangelist, Designer and Architect

Haakon Roberts, IBM Distinguished Engineer – DB2 for z/OS Development

Tom Vogel, IBM Senior Product Marketing Manager, DB2 Tools