IBM DB2 utilities and tools

A leader in data management solutions

Highlights

- An industry leader in DB2 for z/OS support from early release to general availability, continuous delivery and beyond
- A focus on improving efficiency and availability while minimizing application impact
- Responsive to the needs of today while helping grow competitive advantage tomorrow
- Provide a clear choice to help reduce software and vendor costs

The choice is clear

An organization's data is its core business asset. The amount and types of data include a combination of structured, unstructured and replicated data, and continue to grow and change over time. Continual demands are placed on the accessibility, availability and performance of the data to ensure that business goals are met in the constantly changing global economy. To keep pace with these demands, organizations must find solutions to effectively manage this wealth of data.

In response, IBM offers a comprehensive set of utilities and tools to assist with the management of IBM® DB2® for z/OS® data. They can help you more efficiently perform large-scale operations across your enterprise.

However, IBM DB2 utilities and tools do more than improve the efficiency of DB2 systems and data. The DB2 utilities are essential in exploiting new capabilities in DB2, which are growing in complexity and size. As these systems evolve, there are more stringent demands for availability and efficiency. This evolution, in turn, places more demands on even the most experienced systems administrators.

IBM DB2 utilities and tools can help you reduce your total cost of ownership (TCO) by maximizing business performance and IT staff efficiency. IBM continues to invest in offering one of the most complete utilities and tools management portfolios in the market with your choice of standalone products or integrated solutions to help drive down software costs.

The key to enabling new DB2 functions

IBM DB2 utilities and tools provide comprehensive support for all aspects of DB2 for z/OS. You can take advantage of new features more quickly and on your time schedule without waiting for vendor support. IBM works to help ensure new DB2 functions have the necessary IBM DB2 utility support. At the same time, IBM invests in the tooling area to provide the DB2 support you need.

DB2 for z/OS customers must be able to test new functions in DB2 as part of the Early Support Program (ESP). That's why IBM makes both the DB2 utilities *and* tools available to customers in the ESP program. The results are overwhelmingly positive, particularly with the ability to use IBM OMEGAMON® XE for DB2 Performance Expert (OMPE) reporting from IBM and non-IBM tools customers.

"This was our first time using OMPE reports and I was very glad for the capabilities it provided."

— DB2 for z/OS ESP participant

DB2 12 and the REORG utility

The IBM REORG utility changes and improves with each release of DB2 and DB2 12 for z/OS is no exception. However, there are significant changes being made by IBM in how the REORG utility is used. Over the past several DB2 releases, fewer executions of the REORG utility are needed to regain clustering and enhance application performance because other improvements are changing its original use. These improvements include recent DB2 optimizer and I/O and hardware upgrades, better caching, and greater exploitation of memory by DB2 for enhanced performance. An example in DB2 12 is contiguous buffer pools support.

DB2 is increasingly using the REORG utility to deliver new functions without impacting application availability. IBM has chosen this path over several years and DB2 12 has more examples of this usage, making it a key differentiator for the REORG utility. With DB2 12, you now have the ability to:

- Insert a new partition in the middle of a partitioning range
- Convert to partition by range (PBR) with relative page numbering (RPN)
- Exploit compression of a binary large object (LOB)
- Defer column-level alter operations

These enhancements only take effect after the next execution of the DB2 REORG utility. In DB2 12, you can also reorganize a single partition-by-growth (PBG) table space and have it automatically overflow to a new partition if the data no longer fits after the REORG operation. Many of the scalability and availability benefits in DB2 12, which take advantage of greater scalability and application availability, use the DB2 12 REORG utility.

More REORG performance enhancements

There are many new enhancements beyond DB2 12 support in the REORG utility, including greater IBM z Systems Integrated Information Processor (zIIP)-eligible offload capabilities. During the REORG RELOAD phase, there is a 17 percent CPU reduction in the general processor and a 57 percent offload to zIIP engines. Every phase of the REORG utility is now zIIP-eligible. The IBM DB2 Sort for z/OS tool also added more zIIP capabilities for REORG. This feature provides block-level operations for variable length record sorts, which sort data in the REORG utility. Customers running the REORG utility with DB2 Sort and zIIP engines may see the following reductions:

- Up to 70 percent reduction of sort CPU usage
- Up to 58 percent reduction of utility CPU usage
- Up to 54 percent reduction in elapsed time

IBM has enhanced the REORG utility for specific application requests. For example, if you use active-active or replication products with the REORG utility, you can now perform a REORG operation on READ ONLY tables. A usability feature was added for SAP applications for REORG and its drain process. If it fails during the drain process, the REORG task can retry the process and displays claimer information on each drain failure, enabling SAP applications to take remedial action before the REORG utility fails. This availability enhancement benefits applications such as SAP. Another specific SAP enhancement is the REORG operation no longer fails if, during a REORG on catalog objects, no recovery base is found.

"The IBM DB2 utilities zIIP offload percentage of DB2 12 CPU usage is much better than DB2 11."

— Mateusz Książek, PKO Bank Polski

DB2 12 and the LOAD and UNLOAD utilities

As with the zIIP offload in the REORG utility, the LOAD utility increased its zIIP capabilities in DB2 12 with up to 90 percent of the LOAD utility processing being zIIP-eligible. Other performance enhancements include:

- LOAD PART REPLACE can skip the non-partitioned index (NPI) scan if the PBR partition is empty. Prior to this feature, a customer running a LOAD COPYDICTIONARY required two hours. Now, by implementing this *test input*, there's a potential 99 percent reduction in CPU usage and 98 percent reduction in elapsed time.
- An important availability option was added for the LOAD utility and is supported in LOAD REPLACE or LOAD RESUME SHRLVEL REFERENCE and the prevalidation of input data to the LOAD utility without impacting the target table. Previously, a typical usage example might consist of a reference table that is read by applications throughout the day, and at midnight the data is refreshed. That refresh of the data at midnight discards the existing data and replaces it with the new data. This process is the definition of the LOAD REPLACE function.
  LOAD REPLACE also implied that there was no access to the table while the LOAD was running. If the LOAD runs for 30 minutes, there would be 30 minutes during which time your DB2 applications are not available for use.
  The new enhancement allows those applications to continue to read the data in the table while the data is being replaced. LOAD REPLACE now creates a copy or a shadow of the target table. Then LOAD is redirected to load into the shadow copy. Once finished, post-processing terminates for both the shadow and the original page sets. A swap and rename of the data sets and final clean up then occurs. The result is a lengthy production application outage is reduced to mere seconds. This enhancement works for LOAD RESUME and can also be used to validate data without impacting the real table. Prior to this enhancement, a failure of the LOAD REPLACE/RESUME utility left the data in a recovery situation. With this enhancement, a potential LOAD failure avoids the recovery, as the target table is not affected.
- LOAD SHRLEVEL CHANGE PARALLEL support for PBGs was enhanced in DB2 12. This support was available in version 11 of DB2 and extended in DB2 12 to cover PBGs, including support for partition growth. There's a 76 percent elapsed time reduction from the previous version.

- LOAD RESUME BACKOUT YES allows you to avoid a RECOVERY PENDING situation, and therefore a recovery operation, should the LOAD utility fail. This avoids the need to perform any recovery operation.
- The LOAD utility can now maintain MAXASSIGNEDVAL for identity columns. LOAD REPLACE now maintains the value at the start of the LOAD operation. User-supplied values are also maintained in MAXASSIGNEDVAL. This improvement helps sequences and identity columns in the LOAD utility. By maintaining MAXASSIGNEDVAL, which is basically control information in the DB2 catalog, it helps minimize application failures after loading data into a table.
- REGISTER NO can effectively eliminate data-sharing overhead for unloading data. Previously, customers performing an UNLOAD operation, incurred page registration situations that flooded their coupling facility in a data-sharing environment. This feature eliminates the situation and provides a unique competitive advantage over other DB2 utility vendors and is also extended to the RUNSTATS utility in DB2 12.

"The RUNSTATS enhancement with profiles, inline stats and the DB2 optimizer update, completes the picture for us. We are extremely satisfied."

— Walter Janißen, ITERGO

DB2 12 RUNSTATS enhancements

COLGROUP processing has been improved for RUNSTATS to significantly reduce CPU usage and elapsed time.

DB2 12 delivers major improvements in statistics management, enabling the use of statistics profiles to help improve and simplify DB2 self-management.

IBM removed the restrictions on using statistics profiles by now allowing their use in inline statistics in version 12 of DB2. In DB2 11, the DB2 optimizer could update the catalog with information about missing statistics it needed. A tool, such as IBM DB2 Query Workload Tuner for z/OS, could take that information and update the statistics profile for the table.

In DB2 12, DB2 updates the statistics profile directly. This procedure includes not missing statistics required by the DB2 optimizer. It also includes support for DDL changes, such as dropping or renaming an index, dropping or renaming a column, or renaming a table. DB2 12 now automatically updates the relevant statistics for the profile for all relevant actions. Furthermore, clean-up is improved with an automatic removal of unneeded distribution statistics when using the DB2 12 RUNSTATS USE PROFILE option.

When a profile has been updated, it can be detected by the IBM DB2 Automation Tool for z/OS, which then completes the cycle and helps provide greater DB2 self-management capabilities. The Automation Tool can generate the necessary RUNSTATS job to gather the correct statistics that the DB2 optimizer needs. This DB2 12 statistics enhancement fits the infrastructure delivered in the IBM tools portfolio to move from simple automation of tasks to true DB2 self-management. Some complex decisions are made to eliminate unnecessary overhead and enable the running of DB2 systems to be more efficient and minimize database administrator (DBA) involvement.

Other new capabilities in the DB2 12 RUNSTATS utility include:

- A new option to help determine if you want to invalidate the statements in the statement cache. Prior to DB2 12, there was no choice. This upgrade is a unique feature for IBM as other vendors have not been making investment in supporting this feature natively or in vendor-specific versions of statistics.
  Many customers use dynamic SQL in their DB2 applications. DB2 caches these statements in memory. Prior to DB2 12 RUNSTATS, a customer running RUNSTATS in this situation would invalidate these in-memory cached SQL statements, resulting in potential costs associated with performing operations such as a PREPARE.
  There was also no guarantee that you would get the same access path from the DB2 optimizer. This new feature in DB2 RUNSTATS allows you to invalidate these statements in the cache.
- There is also support for inline statistics for LOAD PARALLEL in DB2 12.
- LISTDEF performance improvements were added to DB2 12 RUNSTATS. You can now avoid repeated index processing for a list of parts.
- There's an improved CLUSTERRATIO formula in DB2 12 RUNSTATS.
- A new REGISTER NO option is provided in DB2 12 for RUNSTATS SHRLEVEL CHANGE. This option helps eliminate data-sharing overhead by avoiding the coupling facility page registration.
- There's a performance improvement for COLGROUP in DB2 12 RUNSTATS. DB2 gathers statistics from the index if it finds a match in the COLGROUP columns. The ability to group columns for gathering statistics is more efficient than the previous method of scanning the entire table. RUNSTATS can gather the needed information from only the index defined for the columns. Testing by the IBM DB2 lab shows a possible 25 percent reduction in CPU usage and a 15 percent reduction in elapsed time.
- The COUNT value for FREQVAL is automatically determined in the DB2 12 RUNSTATS.

DB2 12 improvements for backup and recovery tasks

There were significant improvements added to the DB2 12 backup and recovery capabilities:

- The default for the RECOVER utility is PARALLEL. In early IBM testing, this improved elapsed time by as much as 35 percent.
- There's added support for point-in-time (PIT) recovery prior to materializing REORGs for PBG tablespaces.
- There's a RECOVER SCOPE UPDATED option for PIT recovery. You can skip the recovery process for non-updated page sets. This option helps avoid unnecessary PIT recoveries. For example, if you have two tables that you needed to recover three hours ago, the PIT RECOVER operation can now determine which of those recovery operations is not needed. This saves system resources from being used on unnecessary processing.
- You have the option to physically delete data sets with the DB2 12 MODIFY RECOVERY DELETEDS option.

The MODIFY COPY NOCOPYPEND option avoids leaving objects in a copy pending situation. This option is particularly useful for system-level backup and recovery environments.

IBM has added enhanced FlashCopy support in DB2 12, including support for multiple copy pools, better messaging for system-level backups and better template support. The REORG utility also better handles any FlashCopy failures. Other FlashCopy DB2 12 improvements in utilities include:

- MGMTCLAS and STORCLAS support for FlashCopy templates
- New COPY_FASTREPLICATION ZPARM was introduced. The default option is still PREFERRED, but it now supports REQUIRED.
- FlashCopys created by the COPY utility in DB2 12 maintain COPYPAGESF in SYSCOPY. Previously, subsequent COPYTOCOPY processing acquired a read claim on the page set to determine the COPYPAGESF value.
- IBM Data Facility Storage Management Subsystems (DFSMS) Hierarchical Storage Manager (hsm) messages are included in the DB2 12 utility job output for BACKUP and RESTORE. This inclusion is an indication of how IBM responds to customer requirements and was requested by DB2 SAP customers with large DB2 systems. It also demonstrates the collaboration between z/OS and DFSMShsm development groups to deliver an important new function.

Why IBM DB2 utilities

DB2 utilities go far beyond providing core DB2 12 support with enhancements that:

- Reduce processor usage
- Improve performance through elapsed-time and resource-consumption reductions
- Maximize availability to reduce application impact when running utilities
- Remove constraints and streamline the management of complex DB2 for z/OS systems

---

"We were the requester for most of the new IBM DB2 utilities enhancements and we're very satisfied with the IBM DB2 12 delivery of them."

— IT manager of a worldwide global financial institution

Self-managing DB2 systems

Professionals who manage a DB2 for z/OS environment have some common concerns:

- Increasing database size and complexity
- Exponential volumes of data and applications to manage
- Fewer experienced workers available

Each new release of DB2 introduces new features and functions. However, the scope and range of DB2 applications is moving into different areas that often require more time and effort. IBM is addressing these challenges in a number of ways. IBM Data Server Manager (DSM) helps administer, monitor, manage and optimize the performance of IBM data management platforms across the enterprise. It provides DBAs and other IT staff with the information they need to manage performance proactively and helps minimize the impact these issues have on your business.

DSM is designed to provide a single, simplified interface to manage DB2 systems and databases, as shown in Figure 1.

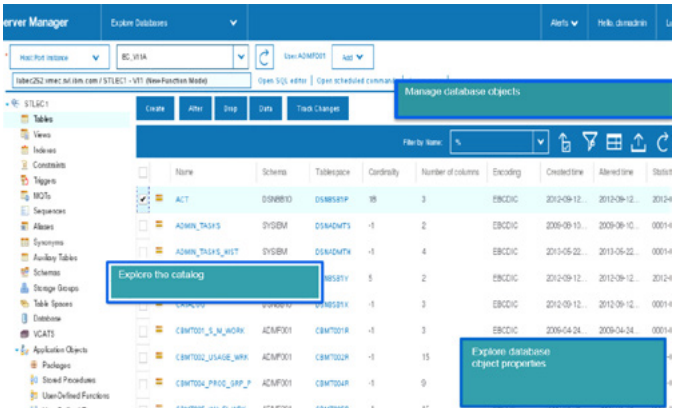

Figure 1: Screen capture of the DSM for z/OS user interface

Data Server Manager for z/OS is a zero-install web-based tool. Some of the base features available include:

- Connect to DB2 for z/OS and navigate database objects, viewing details and links to related objects and object dependencies
- Browse and edit data
- Perform basic database object operations, such as creation of tables, indexes, constraints and tablespaces; dropping of tables, indexes and constraints; altering tables
- Show system privileges from the perspective of group or user, role, or objects

The combined use of DSM and the IBM DB2 Utilities Solution Pack for z/OS moves utility management from autonomation to autonomics. Autonomics takes automation to a different level of operational efficiency. There are four levels of autonomics computing: automation, adaptive, aware and self-managing. One of the key advantages of DSM is the ability to consolidate and help simplify information from various sources. It helps simplify the presentation of complex information. The DB2 Utilities Solution Pack delivers smarter infrastructure that adds intelligence to tasks such as:

- Collection and analysis of data
- Informed decision making based on analysis of collected data
- Execution of decisions during defined time periods
- Integration of tools and DB2 engine functions

Figure 2 shows a screen capture of the user interface for the integration of DSM into the DB2 Utilities Solution Pack.

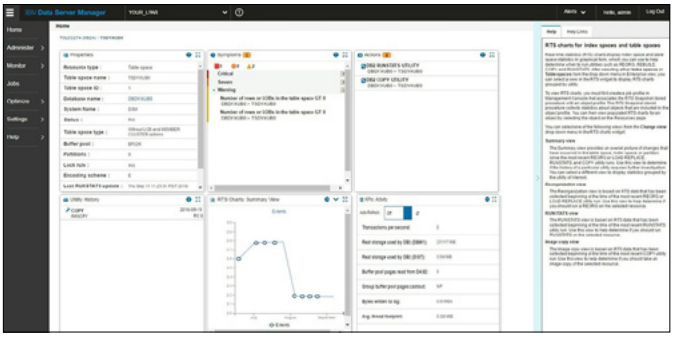

Figure 2: Screen capture showing the integration of IBM DSM and the DB2 Utilities Solution Pack

From the DSM dashboard, there are various views that show all relevant information on utility-based operations. You can view symptoms and the actions on subsystems and objects, and upcoming autonomic user-configured maintenance windows with scheduled actions. The analysis and identification of potential problems with recommended solutions can easily be viewed and changed from the dashboard.

Comparing costs

One of the actions that is enlightening is to compare the cost of utilities. Independent software vendors (ISVs), such as BMC and CA, require the IBM DB2 Utilities Suite to provide DB2 functions they do not support, either based on the function itself or timeliness of their delivery of the feature. A customer must purchase two sets of utilities, the IBM DB2 Utilities Suite and one or more of the ISV DB2 utilities to help ensure DB2 consistency and reliability. Another factor to consider is that the ISV DB2 utilities are not sold together, but separately. For example, REORG, LOAD and others are separate ISV utility products.

One of the claims made by DB2 utility ISVs is performance advantages over the ISV DB2 utilities when compared to comparable IBM DB2 utilities and tools. IBM encourages you to evaluate those performance claims for your environment. While making that evaluation, there should be consideration and factors that help you determine the cost of being current with the latest version of DB2 for z/OS.

Cost of being current on DB2

The most recent version of DB2, DB2 12, provides significant performance benefits that can be realized in MIPS reductions, including:

- Improvements in support of analytics workloads, including optimized performance and improved zIIP offload in support of in-transaction analytics. All parallel child task processing is now zIIP-eligible.
- Up to 25 percent CPU usage improvement for query workloads and up to 10 percent online transaction processing (OLTP) CPU savings[1] with larger memory and activation of memory exploitation features. CPU usage reductions and performance improvements for certain OLTP, as well as select query workloads.
- Expanded in-memory processing for greater performance improvement and emerging use case support, such as synergy with latest IBM z Systems™ processors, which offer significantly expanded memory capacity
- High-volume insert performance equaling two times throughput increase[2] for concurrent sequential insert without clustering
- Faster insights with up to 50 percent elapsed query time improvement.[3]

A study by the IBM Competitive Project Office performed a conservative cost estimation of potential savings realized from greater zIIP exploitation when a customer moves to DB2 12 for z/OS. The study concluded:

- Query performance may yield 100 percent zIIP eligibility for parallel query child tasks.
- The REORG utility may yield up to an additional 17 percent of zIIP offload, which makes almost all of the REORG utility zIIP-eligible.
- Up to 90 percent additional zIIP offload is achieved from the RELOAD phase of the LOAD utility.

The CPO team provided the following cost estimates:

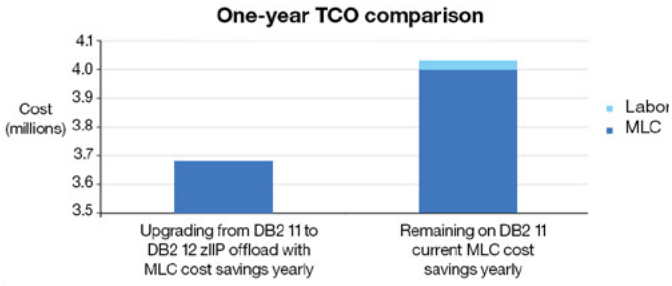

Figure 3: One-year TCO based on upgrading from DB2 11 to DB2 12

This estimate was performed on a 2,900 MIPS customer system running DB2 11 for z/OS. There was an assumption that daytime, sub-capacity peak queries would benefit from greater zIIP exploitation. There was also an assumption that the REORG and LOAD utilities would benefit from greater zIIP exploitation during nighttime batch sub-capacity peaks. A modest 8 percent reduction in sub-capacity peak was calculated into the equation. There was also an assumption of one-quarter time from an IT professional to maintain duplicate sets of utilities and tools. The cost comparison clearly shows a reduction for being current with the latest version of DB2, DB2 12.

Cost on the z/OS platform

Consider also the cost difference of software from IBM versus software from an ISV running on the z/OS platform. The 2016 11th annual mainframe survey from one of the ISVs questioned over 1,200 executives and technical professionals on their perspectives and experiences on the mainframe and their business. Results of the 2016 survey found 89 percent of respondents predict long-term viability of the platform with 58 percent indicating that they would take advantage of the mainframe and its capabilities and add a new workload. A cost comparison was done based on those results.

The IBM CPO evaluated the cost difference in a growing workload on the mainframe and its effect on the price of software. If you consider adding a new logical partition (LPAR) for the purpose of supporting a new workload, there are significant cost considerations. For example, BMC licenses their software products on the entire central processor complex (CPC) frame, not just the LPARS where the software is running. IBM used the example of a new non-DB2 workload, an IBM WebSphere® Application Server (WAS) on this new LPAR. And to illustrate, IBM considered the cost of the DB2 utilities, IBM versus BMC's newest version of the utilities, Next Generation Technology (NGT).

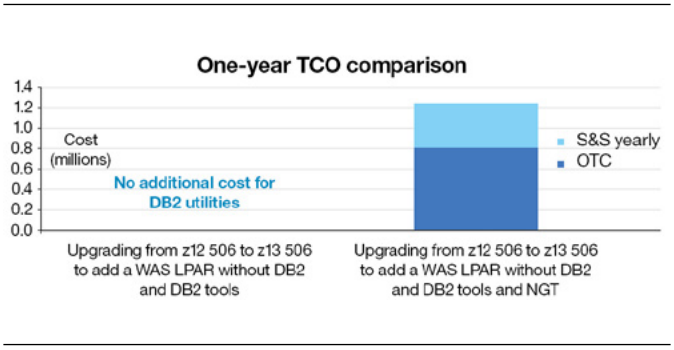

Figure 4: Cost comparsion of software on z/OS: IBM versus BMC

This graph shows the difference in how an ISV prices software versus how IBM prices software. In this case, a new non-DB2 workload was added, which, on the ISV side, raises the cost of exisitng DB2 software on the entire CPC. BMC as other ISVs, prices software on the entire CPC frame. IBM prices it on the LPAR level. The new LPAR in this example did not require the DB2 utilities—neither for IBM or BMC. Because BMC prices software for the entire CPC frame, a customer incurs the additional cost of the BMC NGT utilities, whether used or not, when purchasing software from BMC. There's no additional cost for the DB2 utilities since IBM does not license its software at the CPC level—a *substantial* cost differentiation.

DB2 for z/OS continuous delivery

Starting with DB2 12 for z/OS, customers see a single DB2 maintenance stream, with consumable new functions being released to respond to growing changes and market needs. There will be faster delivery of new features when ready. IBM utilities support new core functions in DB2. It's inherent to the reliability and business continuity of DB2 for the utilities to support new features in its continuous delivery. At the same time, IBM invests in the tools area to help ensure you have the DB2 support needed. There's no guesswork in DB2 support with IBM DB2 utilities and tools.

"IBM is the only vendor I trust to touch my data"

— IT manager of a North American financial institution

Conclusion

The DB2 Utilities Suite helps ensure almost immediate support for all core functions of DB2 for z/OS. IBM is also committed to delivering increased savings and performance.

The IBM DB2 for z/OS utilities and tools are among the most comprehensive on the market. They have the ability to provide critical solutions to today's most complex business problems that span many functional areas to help deliver the highest return on investment from your software.

For more information

To learn more about IBM DB2 utilities and tools, contact your IBM representative, or visit: **ibm.com**/analytics/us/en/technology/db2/zos/tools

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
February 2017

1 IBM Competitive Project Office

2 IBM Competitive Project Office

3 IBM Competitive Project Office

Please Recycle

IMS14510-USEN-00