PRODUCT SHEET
CA Datacom/DB

CA Datacom®/DB Version 14.0

CA Datacom®/DB for z/OS is a high-performance relational database management system (RDBMS) for mainframe operating systems. CA Datacom/DB provides a central data repository for your enterprise-class applications that demand the highest service-level requirements.

Overview

Corporate database management systems must be highly scalable to handle the ever-increasing demand to support additional workload requirements. They must also be flexible enough to support existing applications while deploying new web-enabled technologies without requiring costly rewrites and conversions. CA Datacom/DB provides award-winning transaction processing performance for both existing and new workloads, helping you to consolidate business functions on a central data repository as needed by your business.

Business value

By exploiting the geometric growth of mainframe memory capacities and processor speeds, the CA Datacom/DB RDBMS is scalable to manage extremely high-volume workloads. The database engine also takes advantage of new hardware technologies with improved memory optimization techniques at each new release. The ability to scale to handle additional workloads, which is due to highly efficient CA Datacom/DB resource utilization techniques, has allowed some customers to delay or cancel plans for hardware upgrades.

Features

Next-Generation Mainframe Management

CA Datacom/DB has adopted key features that are designed to simplify your use of CA Datacom/DB and help your staff to install, configure and maintain it more quickly and effectively.

- **CA Mainframe Software Manager™:** CA Mainframe Software Manager (CA MSM) automates CA Datacom/DB acquisition, installation, deployment, configuration and maintenance and removes SMP/E complexities. CA MSM saves time and resources when compared with the manual steps used to perform these tasks.
  - The Software Acquisition Service helps you to more easily move product installation packages and maintenance from CA Support Online directly to your mainframe environment and prepare them for installation.
  - The Software Installation Service standardizes CA Datacom/DB installation, which includes a new, streamlined Electronic Software Delivery (ESD) method that allows CA Datacom/DB to be installed using standard utilities. This service also provides standardized SMP/E product installation and maintenance via APARs and PTFs, and simplifies SMP/E processing through an intuitive graphical user interface and an intelligent Installation Wizard.
  - The Software Deployment Service helps you to more easily deploy CA Datacom/DB in your mainframe environment.
  - The Software Configuration Service helps you to more easily configure CA Datacom/DB if it has been acquired, installed, maintained and deployed using CA MSM.
- **Installation Verification Program (IVP) and Execution Verification Program (EVP):** As part of qualification for inclusion in the set of CA Technologies mainframe products released every May, CA Datacom/DB has passed stringent tests performed through the IVP and EVP to find and resolve interoperability problems prior to release. These programs are an extension of our ongoing interoperability certification initiative launched in May 2009.
- **Best Practices Guide:** This guide provides information on CA Datacom/DB installation, initial configuration and deployment to shorten the learning curve for staff that is responsible for the installation and management of this product.
- **Health Checker:** The Health Checker provides CA Datacom/DB Health Checks that execute under the IBM Health Checker for z/OS.

- **Serviceability**: To provide additional levels of information and consistency for problem determination, CA Datacom/DB has adopted the following serviceability enhancements:
  - Product description information: CA Datacom/DB provides a Product Description module that contains the product name, version/release number and other useful information. To determine which CA Technologies products are installed within each z/OS environment, you can access and report on this module using the CA Common Services EXAMINE facility.
  - Module identification: Each object deck comprising CA Datacom/DB contains a copyright and module identifier to assist with error analysis and debugging.
  - Common dump title formatting: When applicable, CA Datacom/DB recovery routines use standard formatting of dump titles to assist with problem identification.
  - Integration with CA OPS/MVS® Event Management and Automation System State Manager: CA Datacom/DB integrates directly with CA OPS/MVS to indicate programmatically its state and readiness in your environment.
- **zIIP Specialty Processor Exploitation (z/OS)**: This feature provides support for offloading processing cycles from general processors to zIIP processors, increasing overall CPU throughput at a lower operational cost. A portion of all database requests serviced by CA Datacom/DB Multi-User Facility (MUF) is eligible to be offloaded to the zIIP processor (see Figure A).

FIGURE A.

zIIP specialty processor support

For a given transaction, CA Datacom/DB exploits IBM's zIIP processor by establishing a TCB owning task and spawning SRB sub-tasks, which are zIIP-eligible. If the IBM Workload Manager (WLM) determines that a zIIP is available, then the SRB sub-tasks are dispatched to the assigned zIIP. If no zIIP is available, then the SRBs are re-routed back to the General Processor (GP).

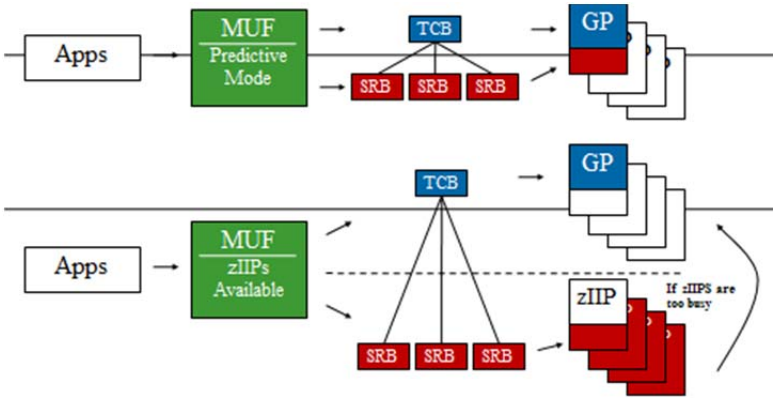

What's new in CA Datacom/DB Version 14.0

CA Datacom/DB Version 14.0 includes a number of new features, with many that focus on lowering your Total Cost of Ownership (TCO) through additional performance enhancements and expanded support for the IBM zIIP specialty processor. Other features are designed to improve usability with a more intuitive approach to managing a CA Datacom/DB environment and to assist training new personnel. New features for CA Datacom/DB Version 14.0 include:

- **zIIP Specialty Processor Exploitation (z/OS):** CA Datacom/DB Version 14.0 provides zIIP enhancements that optimize the zIIP architecture, deliver more efficient zIIP utilization and make additional CA Datacom workload eligible for execution on a zIIP processor. zIIP eligibility is increased by about 50% and early testers of CA Datacom/DB Version 14.0 zIIP processing have seen offloads to zIIP approaching 48% of the CPU consumption of the MUF.
- **Encryption:** CA Datacom Version 14.0 provides encryption for data at rest at the table level to satisfy user requirements that are driven by federal regulations (PCI) and the desire to protect sensitive information such as social security and credit card numbers. CA Datacom data encryption allows the user to select encryption on a table-by-table basis. Tables selected for encryption are stored by CA Datacom/DB on DASD as encrypted images using the IBM z/OS encryption macros and encryption processor.
- **SIMPLIFY feature simplifies configuration:** CA Datacom Version 14.0 delivers improved system set up and configuration including eliminating the SVC requirement and adding more error protection. CA Datacom allows clients to select a simplify feature that greatly reduces the risk of accidental damage to the CA Datacom system components by users, for example, by running a DBUTLTY INIT LXX against an executing MUF. Simplify is also meant to reduce or eliminate a lot of technical decisions needed to run MUF. There are numerous changes, simplifications and protections encompassing this feature including:
  - CA Datacom Data Set Protection. For example, a data set being updated in one environment is protected from concurrent initialization from another environment.
  - Directory (CXX) concurrent update requirement and protection. With MUF enabled, CXX is protected from nearly all risks to its safety from inappropriate use and, at startup, avoids frequent protect/release cycles thus saving substantial CPU and I/O resources.
- **Buffers:** As with the IXX buffers in CA Datacom/DB Version 12.0, the DXX buffers can now be optionally moved from 31-bit storage into 64-bit storage. In Version 14.0, MUF supports up to 99 additional buffer pools of each type: IXX, DXX, and DATA. The additional MUF buffer pools provide the capability to manage all large high-volume MUF instances to execute in a way that prioritizes memory usage, enabling you to speed the most important

databases over less important ones. Large-page support is now provided for index buffers (IXX and DXX) located in 64-bit storage.

- **Performance:** CA Datacom Version 14.0 introduces new performance improvements such as automated data collection and reporting for better performance tuning and as the foundation for more self-automated tuning in future releases, more index and buffer pools, additional 64-bit memory usage, better memory management and log file management.

- **Reorganization:** CA Datacom Version 14.0 provides greater reorganization advancements including "patent pending" techniques that improve the planning and execution of data reorganization activities. Included are new efficiency reports that show whether reorganization is currently needed. This feature also includes new online reorganization options that optimize the reorganization activities based on free space available.
  - New OLREORG AREA Entity-type SPACE_MNGMNT Attributes
  - New OLREORG capability set
  - New DATANE and DATASP reports
- **Miscellaneous enhancements:** There are many other new features included in the CA Datacom/DB Version 14.0 product delivery. Refer to the *CA Datacom Release Notes Version 14.0* guide for further details.

Other Key Features

- **High-volume workload processing:** Many IT organizations depend on CA Datacom/DB as an industrial-strength database management system for their core business applications.
- **Best-in-the-industry OLTP throughput:** CA Datacom/DB customers have registered award-winning online transaction processing rates.
- **New advances in hardware:** Significant throughput gains have been achieved by exploiting the latest performance improvements in processor power and channel/device enhancements.
- **Memory management:** Increased task/buffer limits and exploitation of larger virtual memory pools have contributed to vast improvements in multi-task processing power.
- **Symmetrical multiprocessing:** The high-performance database engine uses multitasking and multithreading to fully utilize n-Way processors.
- **Parallel variable log:** This feature provides a notable performance boost for sites that have a significant number of database updates (add/delete/update).
- **Row-level locking:** This CA Datacom/DB characteristic limits the scope of exclusive update control to a single row.

- **Continuous availability:** CA Datacom/DB includes a number of features that have enabled organizations to access their critical data around the clock.
- **Online data reorganization:** A critical component in our 24-hour-a-day, 7-day-a-week availability strategy is to deliver data area reorganizations without also requiring a scheduled outage. This feature is designed to provide the data center with the ability to selectively reorganize user data areas while allowing the data to remain online for real-time processing. This function has a minimal resource footprint and can be stopped and started as needed. The Online Reorg facility offers the user various parameter selections to target the number of data blocks to be reorganized. Many competitive database management systems implement online reorganization techniques that require significant extra data space and move all rows regardless of their order in the data area. But the CA Datacom/DB Online Reorg process works within the existing data area space and only moves those rows that are not in native sequence. Although it is generally recommended to run this process during non-peak time periods, the online reorganization function can be stopped and started as needed. Aspects of this feature are protected under U.S. patent # 7,519,637.
- **Shadow MUF (z/OS):** IT organizations are ever vigilant to ensure that when unscheduled outages occur, systems are brought back online as soon as possible. This feature is designed to directly address MUF failures and LPAR outages. A Shadow MUF is initiated in an address space on the same LPAR or a different LPAR (recommended) within a Parallel Sysplex. The Shadow MUF watches the primary MUF and mimics dataset opens. If the primary MUF (or its LPAR) fails, the Shadow automatically starts a takeover process to assume the identity of the primary MUF, clean up any in-process transactions and begin processing new requests. The Shadow MUF utilizes very few resources while in Shadow mode and has no performance affect on the primary MUF. This feature is designed to provide a significant reduction in the amount of downtime that occurs during a system failure. The Shadow MUF can also be coupled with the data sharing functionality of CA Datacom/DB to provide 24 hours a day, 7 days a week access during maintenance windows for multi-LPAR systems that deploy a rolling IPL process (see Figure B). Aspects of the Shadow MUF feature are protected under U.S. patent # 7,487,188.

FIGURE B.

Shadow MUF Feature

The Shadow MUF feature is designed to provide very efficient failover support for CA Datacom/DB during unscheduled system outages and provide continuous service during a rolling IPL maintenance window.

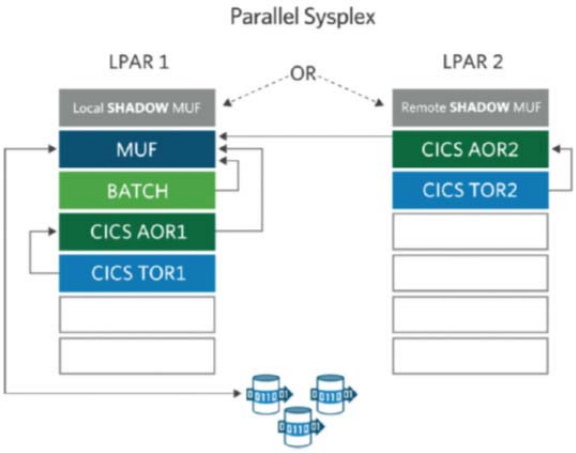

- **Online index defragmentation:** This continuous availability feature is designed to allow you to optimize your database indexes without delaying or interrupting normal real-time data access. The process combines partially used blocks into single blocks and returns the emptied blocks to the free-space pool. In addition, the newly combined index blocks allow more information to be kept in the index buffers, providing better buffer reuse and improved performance. By regularly running the index defrag, the index space remains constant and avoids the expansion of index levels, which also helps to maintain optimum performance. The index defrag function can be executed at any time and can be stopped and started as needed.
- **Open architecture:** Data stored in CA Datacom/DB is accessible from a variety of open system platforms using industry-standard APIs. CA Datacom/DB has a number of features that have enabled organizations access to their critical data around the clock.
  - Database connectivity: Using CA Datacom® Server, access to data is supported through standard interfaces, such as JDBC, ODBC, OLE/DB, ADO and CGI.
  - Platforms: Support is provided for the most common distributed and Web platforms, including Windows, AIX, Solaris, HP/UX, Linux, Linux on System z and other platforms, and z/OS UNIX.

  - Multiple database access methods: Application programmers can match their processing requirements and program logic with the most appropriate database access method.
  - SQL access: Supports industry-standard SQL command access for host and remote applications.
  - Navigational access: Includes a high-performance record-at-a-time (RAAT) API that allows programmers to quickly retrieve a single data row or index value with very low overhead. In addition, the set-at-a-time (SAAT) API provides the same highly efficient processing for a set of rows on a single table.
  - Mixed mode: Supports any mixture of SQL, RAAT and SAAT commands in the same program with full logical unit-of-work integrity. This enables programmers to select the most efficient access command for any task.

- **Ease of use:** Database administration and user interface facilities are designed to bring new IT personnel up to speed with minimal effort.
- **Installation and configuration:** The CA Datacom/DB installation and upgrade process has been streamlined to help provide for a simple implementation. Facilities are provided to support mixed releases in the same environment, as well as easy fallback in case of a problem during or after the upgrade.
- **One system utility:** CA Datacom/DB system administration can be managed through a single batch utility, called DBUTLTY. This helps simplify database administration training efforts.
- **SQL functionality:** The SQL language component provides support for data definition, access and security using industry-standard SQL syntax.
  - Performance: Improved statement cache, join method processing and MIN/MAX optimization have contributed to measurable performance gains. When all columns are in a traversal key, index-only processing will occur, which cuts down on I/O significantly.
  - SQL standards: Supports both ANSI and FIPS standards.
  - SQL extensions: Supports syntax extensions to enhance its power and flexibility.

- **Support for non-relational data:** Customer data can be easily converted from IBM's VSAM and DL/1 or Cincom System's TOTAL into CA Datacom/DB data tables. Existing applications continue to execute without change with the converted data and developers can take advantage of the powerful CA Datacom features to extend these applications and create new ones.
  - Object-level transparency: Location of the transformed data is completely transparent to the application program. No program recompiles or re-links are required.
  - Full relational database support: Non-relational data converted to CA Datacom/DB data will unleash full relational DBMS features and functions. The SQL command CREATE DATACOM VIEW allows SQL access to the following legacy data formats:
    - Array elements
    - Redefined fields
    - Compound fields
- **Datadictionary:** An integrated data dictionary helps provide that data descriptions are consistent with data access and storage requirements.
- **Dynamic system tables:** These tables contain the primary set of system management metrics, including updatable startup parameters, performance management settings and transaction status data. This set of tables resides in the MUF address space and is accessible using standard SQL statements. Since the most current statistics are stored in memory, this is a highly efficient method to monitor the CA Datacom/DB system from any application platform.
- **CA SYSVIEW® Performance Management CA Datacom® option:** This performance management tool combines a number of CA Datacom/DB metrics into simple customizable displays for quick problem analysis and provides real-time monitoring of system activity.
- **Table partitioning:** Partitioned tables allow you to store the rows of a single table in multiple datasets. This delivers enhanced data availability, expanded database administration flexibility, extended application support, performance gains and 24x7 operational capabilities.
  - Performance: Tremendous performance gains can be achieved by identifying heavily-used data partitions that can be processed as memory resident.
  - Concurrent processing: Partitioning allows multiple, concurrent DBUTLTY data area processes that can reduce requirements for system downtime.
  - Partition level processing: Utility processing is performed only on the partitions that are needed. Other partitions are not affected and can be available for application processing.

  - Continuous access: If a partition is unavailable, processing can continue against other available partitions of the same table.
  - Use existing applications: Existing applications will run unchanged when converting from non-partitioned tables to partitioned tables.
  - New applications: New applications can be created to access a specific partition or group of partitions to enhance the performance or flexibility of data access and availability.

- **Data sharing:** With this feature, up to seven MUFs (in a MUFplex running within an IBM Parallel Sysplex) can reside on different z/OS operating system images with full access to the same data. Some customers use this capability in low-volume workload environments to provide fault tolerance in case one of the images experiences a failure.
- **Exploitation of z/OS Resource Recovery Services (RRS) and Two-Phase Commit:** CA Datacom/DB delivers support of the z/OS RRS feature that provides data integrity for applications accessing more than one data resource, such as another CA Datacom/DB system, CA IDMS™/DB, IBM DB2 or WebSphere. This service covers multiple instances of the same resource.
  - Remote Resource Manager (RRM): RRS provides full two-phase commit processing for all concerned data resources on a given z/OS image (LPAR). RRM extends RRS support for CA Datacom applications across multiple LPARs.
  - No changes to user applications: Two-phase commit support is available to applications using multiple MUFs or a single MUF associated with other RRS resources. No user application changes are required to implement this feature.
- **Multi-MUF CICS applications:** With CA Datacom® CICS Services, a CICS application using set-at-a-time or record-at-a-time database access commands can connect to multiple MUF environments at the same time. This provides significant flexibility for installations that run multiple CA Datacom/DB production environments.
- **Exploitation of 64-bit storage (z/OS):** The Memory Resident Data Facility provides the ability to exploit 64-bit storage, where available, in order to reduce I/Os and improve transaction throughput.
- **Historical database (z/OS):** This system database holds various types of system utilization data that assists the database administrator in managing and tuning the CA Datacom environment. The first table delivered for this database captures information that assists the DBA in making decisions regarding the frequency and necessity for reorganization of user data areas. By monitoring these sequential processing metrics, some scheduled maintenance windows may be diminished or entirely eliminated. Starting with CA Datacom/DB Version 12.0, a new table provides a history of recovery file generations to

assist in disaster recovery situations. The historical database provides the basis for autonomic processing.

- **Change Data Capture (z/OS):** The Change Data Capture (CDC) feature allows committed database changes to be immediately stored and accessible by a user application from all standard APIs. It is intended for users that must replicate CA Datacom/DB database updates to another data repository in near real-time with very minimal overhead in the MUF. In the past, customers with data replication requirements used the Recovery File (RXX) and the READRXX routine to accomplish the replication task. This approach does not always satisfy stringent, real-time data requirements. In other cases, users have implemented a real-time synchronous replication process by using an SQL procedure, which could provide a negative MUF performance impact and would prevent data replication if the other data source was not available. The CDC method was designed to remove constraints that were present in other methods.

Benefits

This product release is deploying a number of technology advancements that provide an immediate impact on how CA Datacom/DB is deployed and managed in your environment. They include:

- CA Mainframe Software Manager simplifies your CA Technologies software acquisition, installation, deployment, configuration and maintenance tasks for CA Datacom, as well as all mainframe products from CA Technologies in your software portfolio.
- zIIP specialty processor support, which offloads work from the general purpose (GP) processor to the zIIP processor. This allows you to add more workload using existing resources and to delay planned MIPS upgrades, while lowering your net cost per business transaction.
- Other enhancements that improve performance provide additional functionality to perform system tasks while database access remains available and reduce the total cost of ownership by adding new product capabilities with the base product.

As we deliver each release of the product, we continue to add new ease-of-use features designed to shorten the training cycle for new database administration personnel.

Delivery approach

CA Services provides a portfolio of mainframe services delivered through CA Technologies internal staff and a network of established partners chosen to help you achieve a successful deployment and get the desired business results as quickly as possible. Our standard service offerings are designed to speed deployment and accelerate the learning curve for your staff. CA Technologies field-proven mainframe best practices and training help you lower risk, improve use/adoption and ultimately align the product configuration to your business requirements.

The CA Technologies advantage

CA Technologies has over 30 years of recognized expertise in robust, reliable, scalable, and secure enterprise-class IT management software. CA Datacom/DB is a key component of the Next-Generation Mainframe Management initiative from CA Technologies to change the way the mainframe is managed forever by helping you maximize the value of our mainframe products and by providing a simplified experience and innovative solutions that deliver value quickly and flexibly.

 CS2277_0412