



IBM Z Table Accelerator

Optimize the performance of mission-critical IBM Z applications while reducing the associated total cost of ownership of your IT environment

IBM Z Table Accelerator (IZTA) is an IBM mainframe real-time high performance in-memory table manager. As a mainframe optimization solution, IZTA is ideal for organizations that need to get every ounce of power from their mainframe systems to maximize performance and transaction throughput, while minimizing system resource usage at the application level. It allows for more powerful and efficient applications while lowering mainframe Total Cost of Ownership (TCO) and reducing operational spend. It can also help to increase the timeframe between upgrade capital expense.

In-Memory Table Manager

As an off-the-shelf solution, it is perfect for those organizations that have a need to focus their development efforts more on revenue-generating business activity, and less on supporting in-house developed table management utilities.

Db2 Accelerator

Overall Db2 database performance can be improved by optimization at the application level. Db2 applications with very high transaction processing rates can be optimized to use fewer resources (CPU and I/O) while improving application performance. Applied to several resource-intensive applications will dramatically improve overall Db2 performance.

VSAM and IMS Database Accelerator

Overall VSAM or IMS performance can be improved by optimization at the application level. Applications with very high transaction processing rates can be optimized to use fewer resources (CPU and I/O) while improving application performance. Applied to several resource intensive applications will dramatically improve overall VSAM or IMS performance

Batch Application Accelerator

Batch applications with very high transaction throughput rates can be optimized to use fewer resources (CPU and I/O) while improving application performance. Batch run times can be sharply reduced.

OLTP Accelerator

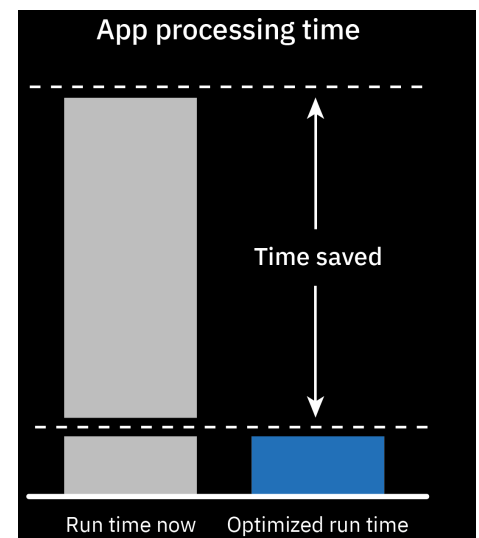
OLTP applications with very high transaction throughput rates can be optimized to use fewer resources (CPU and I/O) while improving application performance. The time needed to complete each transaction can be reduced significantly, dramatically improving the overall system performance.

Query Accelerator

Systems that run both long-run queries (like those used for analytics) and short-run queries (like those used in transaction processing) often need two solutions. IZTA is the ideal solution to improve access time for short-run queries.

Agile Market-Reactive Application Solution

Business rules are processed ultra-fast when they are embedded within rules processing applications. But updating rules requires application recompiles. IZTA applications can run almost as fast as embedded-rules applications, but with rules externalized they can be updated in hours or days rather than weeks or months.



Without IZTA
With IZTA

Figure 1

IBM Z Table Accelerator

How IBM Z Table Accelerator (IZTA) Works

Generally, data access follows the 80 to 20 rule—80% of your data is accessed only 20% of the time, and 20% of your data is accessed 80% of the time. The 20% of your data that you are accessing 80% of the time is the reference data that you access many times for every business transaction.

This is typically where costs get eaten up and represents a significant opportunity to increase performance and decrease operating costs—considerably. You can use IBM IZTA to do just this. Copying some of the most often accessed data into IZTA high performance in-memory tables allows access to that data up to 100 times faster than is possible using other methods.

As shown in Figure 2, the data that is used most often during processing is accessed from IZTA high-performance in-memory tables, while the rest of the data is accessed from DASD using current techniques. In this way, applications can run much faster—and the difference can be astounding.

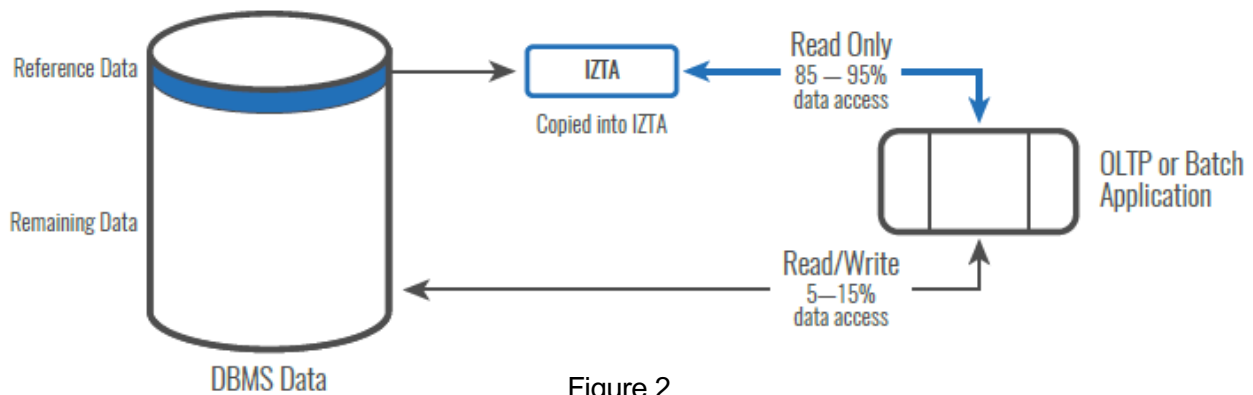


Figure 2

Also, as part of transaction processing, temporary files are often created for the purpose of sorting: using temporary in-memory tables instead can save CPU resource usage and processing time for every business transaction.

IZTA Code Path

IZTA high-performance in-memory tables are accessed using the shortest possible code path—as close as possible to the calling application. Accessing buffered data uses the same code path that is used for accessing data from disk; it reduces I/O, but still uses the same code path. Figure 3 (top) shows an approximation of the Db2 code path taken for accessing data on disk and from buffers.

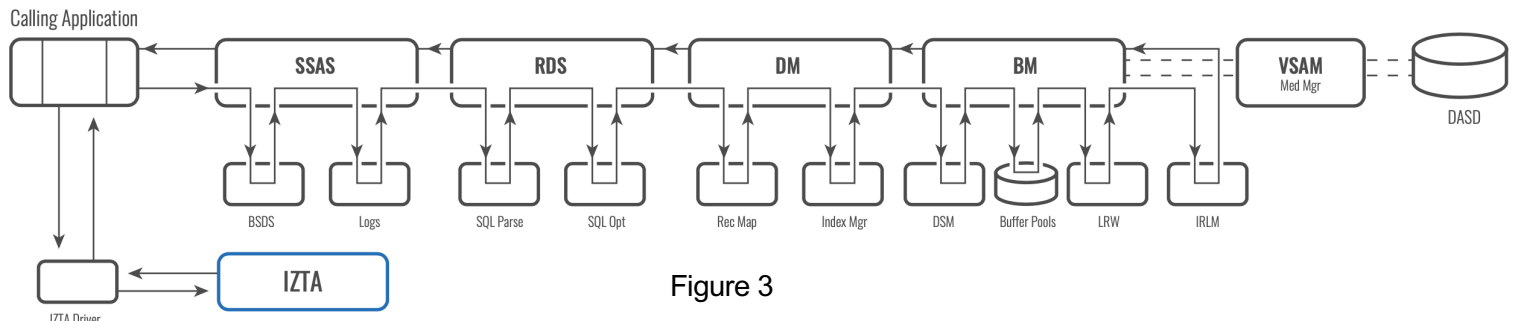


Figure 3

Figure 3 (bottom) shows the code path taken to access data from IZTA high-performance in-memory tables. This is why accessing just a small amount of often-accessed data using IZTA can make such a dramatic improvement to an application's processing speed.



IBM Z Table Accelerator

VTS (Virtual Table Share)

VTS augments the IZTA core product by providing the capability to share table data. Using IZTA, every region loads tables into a local table space region. After augmenting the core product with VTS, all regions can access the same data from a single shared region. Applications do not need to load data into their local table space, and system resource usage is optimized.

CICS and IMS interfaces

Besides the default batch interface, there are two optional interfaces available. An IMS/TM interface provides compatibility for increased throughput IMS online applications, and a CICS interface provides compatibility for CPU-cycle and I/O-operation efficient CICS online applications.

IZTA features

Accessing IZTA

IZTA is accessed via its API—a single call, and can be used in batch, CICS, IMS and other operating environments. This allows organizations to retain and optimize existing applications. A swap of I/O calls for IZTA API calls will obtain up to a 100x improvement in application performance.

z/OS version compatibility

IZTA operates on all version levels of the z/OS operating system.

Db2 version compatibility

IZTA supports all version levels of Db2, and is designed to work in a Db2 stored procedure environment with multiple TCBS.

IMS version compatibility

IZTA/IMS supports all version levels of IMS. ISPF version compatibility. The optional IZTA/ISPF interface supports all version levels of ISPF.

Sysplex support

IZTA is fully capable of running in a Sysplex environment. Read-Only tables are fully supported across LPARs—updates to Read-Only tables are managed using Linear Data Set mappings. Read/Write tables can be managed using Db2 stored procedures. Datasets are fully supported across all MVS images in a GRS configuration.

CICSPLEX support

IZTA is CICSPLEX compliant, but applications using IZTA may require transaction affinities. This limitation can be eliminated by using the optional IZTA VTS interface—this provides support for multiple regions accessing a shared Read/Write table. A transaction running in any region can create, update and access any table in a common shared region.

64-bit support

IZTA is a 31-bit application; it does not use virtual storage above the 2G bar. z/OS fully supports 31-bit applications.

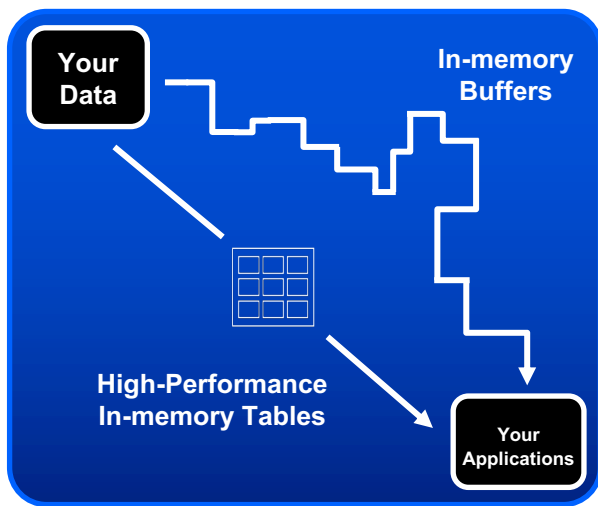


Figure 4

Other Features

IBM Z Table Accelerator (IZTA) uses the shortest possible code path to execute your application's table calls as shown in Figure 4. It is a robust and mature in-memory table manager that has features that your Data Base Management System (DBMS) does not have, including:

- In-memory table management that complements DBMS
- Date sensitive processing
- Indirect table access
- Flexible high-performance indexing
- Version control (up to 9 table generations)
- Multitasking – IZTA is fully re-entrant
- In-memory repository for reference data
- Ideal holding area for temporary data
- Automatic table load and unload
- Multiple high-performance search methods
- Dynamic run-time table expansion
- Dynamic table reorganization
- Dynamic index creation and modification
- Dynamic run-time alternate views
- TSO/ISPF, CICS TS, IMS TM, Db2 SPAS compatibility • C, C++, COBOL, PL/1, Assembler, Fortran accessibility • Data access 1000x faster than I/O
- Data access 10x faster than DBMS buffer pooling
- Compatible with Db2, VSAM, BDAM, QSAM
- Compatible with MVS, IMS
- RACF compatibility



© Copyright IBM Corporation 2020
 IBM Corporation
 New Orchard Road
 Armonk, NY 10504
 U.S.A.
 02/20

IBM, ibm.com, IBM logo, LinuxONE, zEC12, z13, z14, z15 and zEnterprise are trademarks or registered trademarks of the International Business Machines Corporation.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

RStudio®, the RStudio logo and Shiny® are registered trademarks of RStudio, Inc.

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

Zowe™, the Zowe™ logo and the Open Mainframe Project™ are trademarks of The Linux Foundation.

IBM IT Economics Disclaimer

IBM IT Economics Studies are presented as illustrations of how clients can estimate their associated costs for IBM products and non-IBM products. Study results are projections using a total cost of ownership analysis, including costs such as hardware, software, network, storage, facilities and other costs over multiple years. Most costs and performance characteristics are provided by you, the client.

If costs and performance characteristics are obtained from a supplier, it is based on published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the vendor of those products.

Actual costs and performance characteristics will vary by client. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the data used in the study.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

IBM IT Economics Study content is provided "AS IS" without warranty of any kind. Any reliance by you on the study and any results from the study is at your sole risk and will not create any liability or obligation for IBM.

Learn more:

www.ibm.biz/IZTAinfo

www.ibm.biz/IZTAAnnounce