



# TABEX4 JTC sync

MAKE THE MOST  
OF YOUR TABLE DATA

**BOI** BETTER  
ORGANIZED  
INFORMATION

## TABEX4 JAVA TABLE CACHE sync (TABEX4 JTC sync)

### MAKE THE MOST OF YOUR TABLE DATA

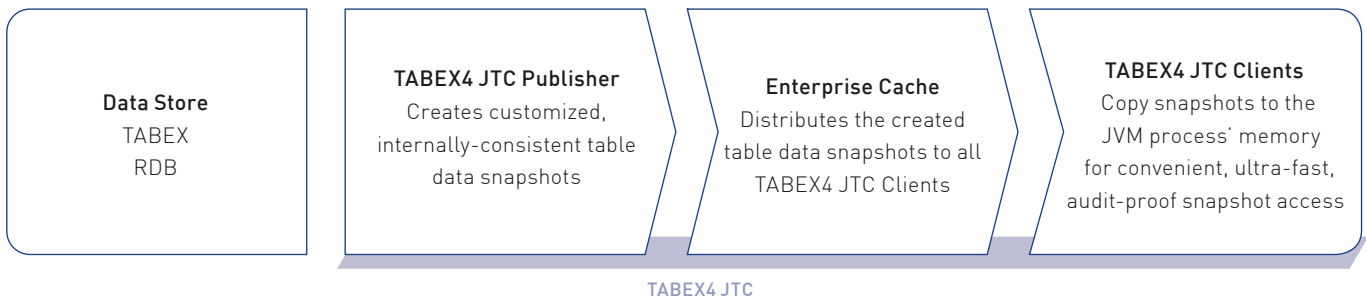
Are your servers challenged by the table data queries of your business applications?

BOI's TABEX4 JAVA TABLE CACHE sync (TABEX4 JTC sync) offers the solution:

- Create table data snapshots of your table data sources – audit-proof and perfectly tailored to your company's requirements.
- Replicate those snapshots within your data centers or across your worldwide network to an unlimited number of TABEX4 JTC Clients – easily and conveniently by using your existing infrastructure.
- Benefit from the ultra-fast in-memory access to your table data stored in the TABEX4 JTC Clients - user-friendly and ultra-fast, using modern Java technology.
- Take advantage of TABEX4 JTC's high-level security options: With TABEX4 JTC sync, your IT architecture supports table data access and distribution – automated and audit-proof.

### TABEX4 JTC sync at a glance

TABEX4 JTC sync, a product of BOI Software GmbH (BOI), delivers your table data to your Java business applications. TABEX4 JTC sync is linearly scalable and provides table data for ultra-fast Java read access - internally consistent, versioned, and ultimately traceable. TABEX4 JTC sync consists of three components: the TABEX4 JTC Publisher, an enterprise caching system of your choice, and the TABEX4 JTC Clients. TABEX4 JTC Publisher collects table data from various sources and uses the object distribution system to publish these data to the TABEX4 JTC Clients. TABEX4 JTC Client is a small pure-Java library linked into your Java business applications. It copies the tables from the object distribution system into the JVM process' memory and provides a concise API to access them as in-memory indexed tables.



TABEX4 JTC sync is a highly focused product. It is designed to provide, distribute, and access table data – ultra-fast and perfectly tailored to your operational requirements

In a TABEX4 JTC deployment, you retain full control over all aspects of table data management including the processes, policies and tools you use to collect and curate table data.

With TABEX4 JTC sync, you will benefit from your investment into an enterprise caching system in more than one way: It allows you to use any enterprise-class, key/value distributed object cache supporting PUT and GET semantics.

## TABEX4 JTC sync: Benefit from unique features

### TABEX4 JTC sync makes use of your existing data organization

As data stores serve TABEX Databases, SHS dataspace as well as relational databases, both from mainframe and non-mainframe environments. TABEX4 JTC sync makes use of your existing RDB-tables, TABEX4 Tables, Views and AI Tables. TABEX4 JTC sync also supports TABEX4 table versions.

### TABEX4 JTC sync ensures transactionally consistent table data access

TABEX4 JTC sync creates internally consistent snapshots of your table data using the TABEX4 JTC Publisher. You control the content of each snapshot by configuring TABEX4 JTC Publisher with table- or index-generating queries. These queries can span data sources on a variety of operating systems and database platforms. TABEX4 JTC sync swiftly distributes these snapshots to the TABEX4 JTC Clients. Every snapshot is already internally consistent when it arrives at the client. Since TABEX4 JTC sync queries run against a single snapshot, there is no need to hold database locks of any kind, and there is no performance penalty when a snapshot is accessed over an extended period of time by TABEX4 JTC Clients.

### TABEX4 JTC sync distributes your data to an unlimited number of TABEX4 JTC Clients

In a TABEX4 JTC sync deployment, the object distribution system is the component that actively delivers scalability — TABEX4 JTC Publisher and TABEX4 JTC Client plug directly into Hazelcast, Infinispan, and Terracotta, or any other key/value distributed object cache supporting PUT and GET semantics. Information flows through the object distribution system from the TABEX4 JTC Publisher to the TABEX4 JTC Client, not vice-versa. Thus, there is no need for distributed transactions or locks, and there is no limit on the number of TABEX4 JTC Clients that can receive and cache your table data snapshots.

### TABEX4 JTC sync guarantees ultra-fast read access on your data

Because TABEX4 JTC sync brings the entire table data to your business application, query latencies are not affected by the shortcomings of the JDBC API, by runtime query compilation, by network traffic, or even by inter-process communication. The business application query goes directly to the data in the business application's JVM—there is no thread context switching and no network traffic. TABEX4 JTC Client queries are up to 550 times faster than direct database access via JDBC.

### TABEX4 JTC sync simplifies data searches

In addition to the ultra-fast key access method to search your in-memory table rows, the TABEX4 JTC Client offers convenient Java methods to search multiple columns with various search operators.

### TABEX4 JTC sync facilitates the use of client-side indices

TABEX4 JTC Publisher can create and distribute indexes: you configure which indexes TABEX4 JTC sync should create for each table. TABEX4 JTC Client will make the data tables, and their corresponding index tables, available for processing in-memory query requests.

### TABEX4 JTC sync puts your tables under technical version control

Every TABEX4 JTC Client can hold multiple consecutive TABEX4 JTC Snapshots. Thus, your business applications can access various TABEX4 JTC Snapshots simultaneously. In data sources with technical validity periods (validity data from - to), TABEX4 JTC Client will automatically use the table version that was valid for the requested date. If you could not use technical validities in your tables so far, you can now implement them with TABEX4 JTC sync.

### TABEX4 JTC sync offers multiple data spaces and multi-tenancy

TABEX4 JTC allows you to create independent data spaces within the same snapshot. Each data space can hold similar table and index structures with different data content.

### TABEX4 JTC sync is audit-proof

Every TABEX4 JTC snapshot has a unique ID. TABEX4 JTC Publisher can automatically archive the full content of each TABEX4 JTC snapshot. Every business application can record the used snapshot ID of every access. Every TABEX4 JTC query result set is based on data from a single TABEX4 JTC snapshot. The TABEX4 JTC Client API enables your business application to access the snapshot ID that corresponds to any TABEX4 JTC query result set.

# TABEX4 JTC sync

MAKE THE MOST  
OF YOUR TABLE DATA

## Conclusion

With TABEX4 JTC, you will maximize the profit of your company's table data. TABEX4 JTC sync is the ideal extension of your TABEX4 installation for your Java applications.

TABEX4 JTC Publisher supports you in creating perfectly tailored, internally consistent snapshots of your table data from virtually any data source or combination of data sources. You choose an enterprise caching system to distribute your snapshots. TABEX4 JTC Client provides your Java business applications with audit-proof, ultra-fast, in-memory, and thread-safe access on your replicated TABEX4 JTC Snapshots.

Together, the three components of the TABEX4 JTC system provide table data access to your business applications that is scalable, convenient, ultra-fast, audit-proof and low-latency.

In a TABEX4 JTC deployment, you retain full control over all other aspects of table data management including the processes, policies, and tools you use to collect and curate table data.

## Specifications

### TABEX4 JTC Publisher

- Operating systems: AIX, Unix, Linux, Windows, z/OS, BS2000, VSE
- Data source compatibility: TABEX, DB2, Oracle, MySQL, PostgreSQL
- Web-based Frontend

### TABEX4 JTC Client

- Pure Java
- Requires Java version 1.6 or newer
- No third-party Java dependencies

### Object Distribution System

- Hazelcast
- Infinispan
- Terracotta
- Any enterprise caching system that supports GET and PUT operations

All names and designations may be trademarks of their respective owners.

