TERRACOTTA DB
NEXT-GENERATION IN-MEMORY DATA MANAGEMENT

Support hybrid cache, store and compute workloads on a single architecture

Analysts predict the digital universe is doubling in size every two years. By 2020, we will create and copy 44 zettabytes, or 44 trillion gigabytes, a year. How will you be able to process, search, analyze, interpret and consume that data so it's accessible—and actionable—by the right people, in the right way and at the right time? Your solution: Terracotta DB.

## The fastest platform for your biggest data challenges

Terracotta DB is a comprehensive, distributed in-memory data management platform—offering a unified architecture for all your operational and analytical data caching, storage and computing needs. Terracotta DB has the speed, performance and availability your business needs to meet massive data storage and processing demands, all within a single operational store.

Unlike other data management systems, Terracotta DB can be used as a "translytical" data store. That means you can manage transactional and analytical information on one system—and faster than on any other solution. Terracotta DB manages 10 to 100 times more data in-memory than data grids using the same number of commodity servers. This next-generation platform radically reduces data management costs by using one of the most powerful query and computation capabilities in its class, based on Java® Streams, which enables developers to meet the demands of modern real-time applications.

Key benefits

- Rapidly store, analyze and exploit ever-expanding volumes of data

- Efficiently control the growth, complexity and cost of your systems and data management

- Optimize application performance across diverse workloads

- Provide flexible and highly scalable capacity for total resilience without complexity

- Make best use of your skilled resources by doing more with less

- Automatic permanent storage of in-memory data with ultra-fast recovery upon server restarts

TERRACOTTA DB
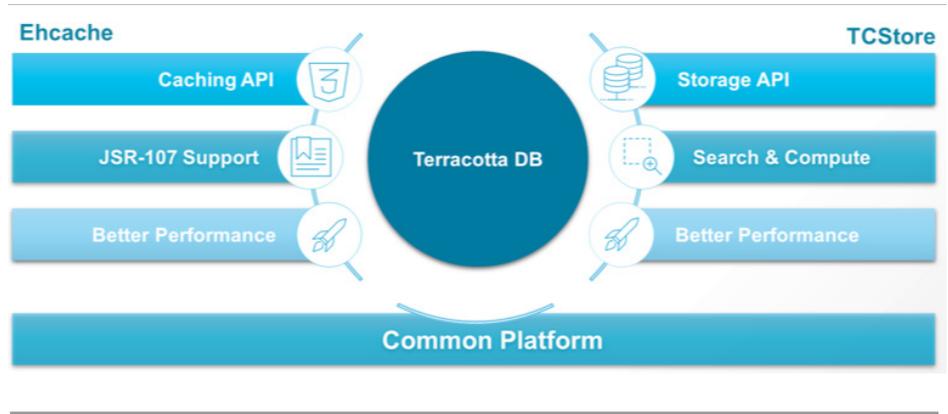
A new TCStore API and improved JSR-107 standard-based Ehcache API work seamlessly side by side and can be used independently—Ehcache for distributed, in-memory caching and TCStore for in-memory storage, compute & search.

Architected for performance, scalability and reliability

Next-generation Terracotta Server
Terracotta Server provides the distributed data platform for Terracotta products. A Terracotta Server Array can vary from a single server to a basic two-server tandem for high availability, to a multi-server array for configurable scale, high-performance and deep failover coverage. The server offers:

- Simple configuration and deployment options for scaling up and/or out

- High availability with instant failover for continuous uptime and services

- Persistent application state with automatic permanent storage of all current shared in-memory data

- Fast restart for enterprise-ready crash resilience, keeping a consistent, real-time record of in-memory data on disk so all data remains available after a shutdown

Terracotta DB is able to connect to a wide variety of enterprise applications via the Software AG webMethods adapter for Terracotta DB and is fully integrated, using native connectivity to Apama for advanced stream processing capabilities and MashZone NextGen for real-time data visualization.

Features

Microsecond response times
By performing all storage, analysis and compute functions in-memory, Terracotta DB reduces response times from milliseconds to microseconds—that's extremely rapid response and unbeatable performance compared to any other spinning disk or solid-state-drive based solution. In addition, Terracotta DB understands the data that's stored, which massively increases the speed of analysis. It can quickly perform calculations for aggregating data within the Terracotta database, rather than on the client system. As a result, it doesn't waste network and application resources by needlessly returning and analyzing a lot of irrelevant data.
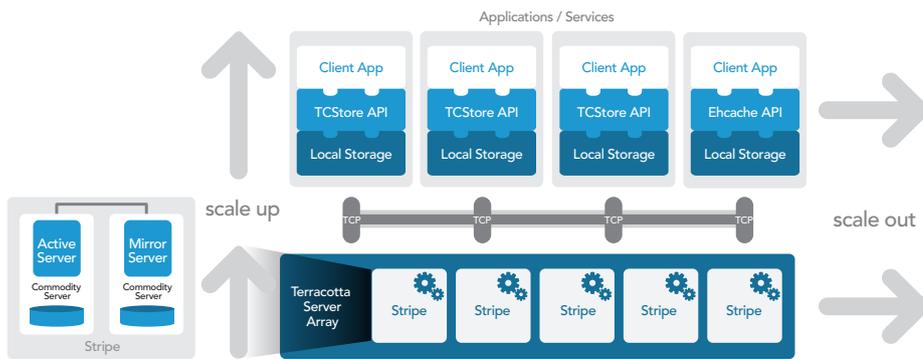
Terracotta Ehcache can offer 99.999% availability on high-velocity data at scale. Scale up by adding more memory. Scale out by adding more commodity servers.

Intelligent workload distribution

As a translytical database, Terracotta DB removes the need for separate databases to perform specific tasks, saving you time on performance/load balancing and splitting data. The Terracotta Server intelligently distributes workload so it can perform sub-aggregations on different servers for even faster response times. By enabling the creation of a single, comprehensive set of data, Terracotta DB avoids the need to partition/shard the database, so even large queries can be performed easily and efficiently, without the need for another database purely for analytics.

Virtually unlimited scalability

To scale up, add more RAM to existing servers. Or scale out to increase processing power by adding more servers, spreading your data and workload over multiple machines in cluster. All RAM is treated collectively and where it is physically located is completely transparent. You can store more data in-memory or create a bigger cache to improve performance since more data can be stored locally near the application and doesn't need to be retrieved from the database.

High availability & resilience

Terracotta DB operates on a highly resilient distributed data platform with bullet-proof availability. Even though it acts as one pool of data, it can be spread across many servers since every node in the cluster has the same role and every node can service any request. This ensures no single point of failure. For complete redundancy, Terracotta DB can ensure data is fully mirrored. If you require extreme redundancy, Terracotta DB supports multiple mirrors which can be located across distributed data centres.

Take the next step

To see how Terracotta DB can transform your digital business, talk to your local Software AG representative or visit www.softwareag.com.

SAG_Terracotta DB_FactSheet_Oct17

software AG