



SMARTCACHE
Any Speed. Any Scale.

The SMARTCACHE logo features a stylized icon of four parallel diagonal lines above the word "SMARTCACHE" in a bold, white, sans-serif font. Below it, the tagline "Any Speed. Any Scale." is written in a blue, sans-serif font.

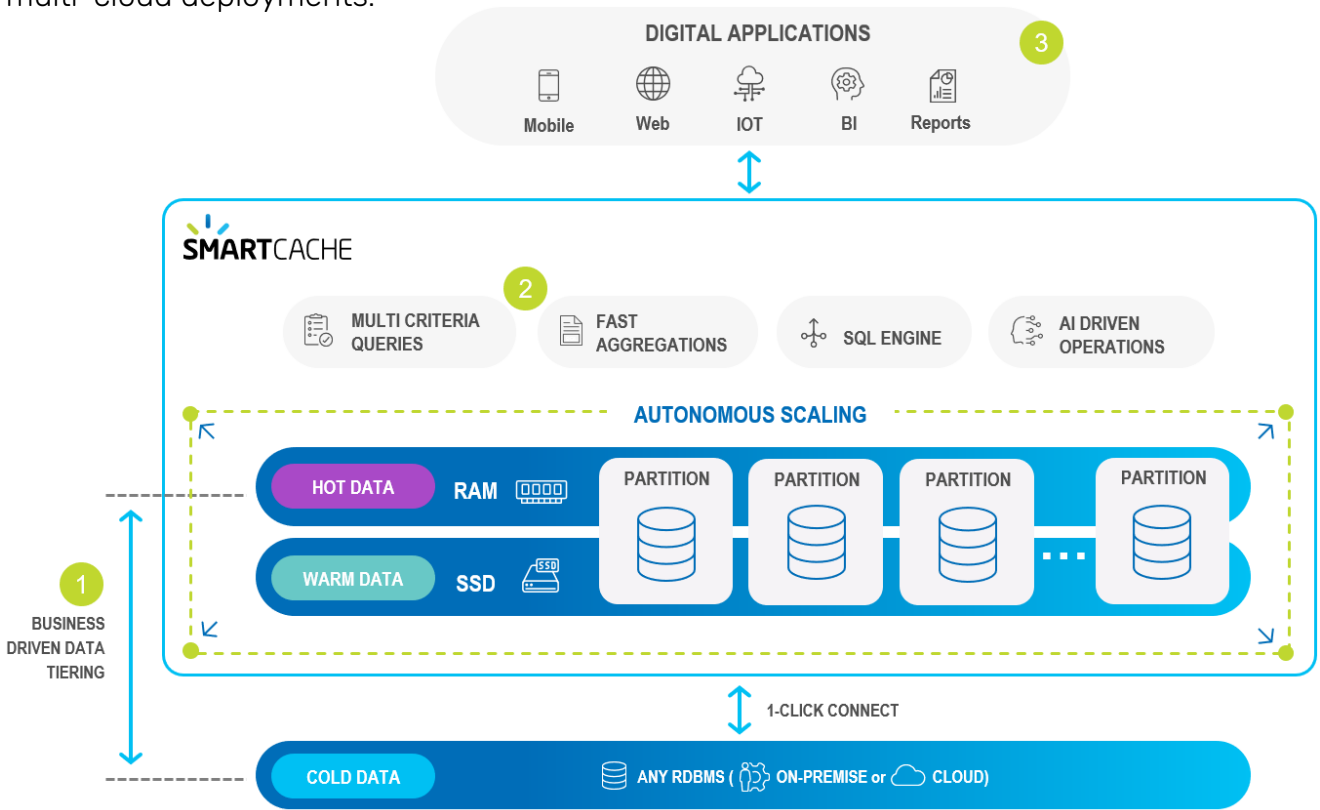
GigaSpaces Smart Cache provides a powerful solution for ultra-fast transactional processing and task execution, offering high availability, reliability and scale.

Smart Cache improves the performance, scalability, and reliability of applications that require high-speed access to frequently accessed data. For organizations that depend on accurate, near real time responses and require full persistence and ACID compliance, Smart Cache accelerates digital applications at the speed and scale of your business, across any environment.

Key Business Benefits

- ✓ Enhances customer experience with millisecond response time
- ✓ Balances cost and performance with smart data management
- ✓ Faster time-to-market with easy deployment and integration
- ✓ Always-on high availability guarantees business continuity
- ✓ Maintains business and service levels for applications under any load and demand

Smart Cache offers **unique capabilities** that are not available in simple caching solutions, such as full SQL compatibility, multi-criteria queries, dynamic server-side processing, full data integrity, policy-driven data tiering, seamless database integration, and hybrid and multi-cloud deployments.



Smart Cache Architecture

- 1 Data is stored in multiple data storage tiers - hot (in-memory), warm (SSD).
- 2 High performance multi-criteria queries, fast aggregations on large datasets, and high concurrency.
- 3 Digital applications connect and trigger SQL or client SDK queries with millisecond performance.

Key Features



High Availability

Smart Cache is an always-on, ACID-compliant and secured system, with disaster recovery ensured through efficient data replication (Active-Active) between sites, including hybrid environments.




Storage Tiering

Data is stored in multiple data **storage tiers** – hot (in-memory), and warm (SSD). User-defined business rules place the most important data - the most accessed data, or the data that should be retrieved the fastest, or more likely to be used - in hot storage. Other data that is less frequently accessed can be kept in warm storage.



Persistence

Smart Cache data persistence handles the fast cache recovery of data from the persistence layer (source) and handles changes made within the space delegation to the persistence layer (endpoint). The GigaSpaces mirror service (write-behind) provides reliable asynchronous persistence that ensures that data will not be lost in the event of a failure.



Scaling Up and Out

Smart Cache delivers a scaling process that can be performed both horizontally and vertically, enabling the increase or reduction of the number of partitions allocated to a service.

Deployment

Smart Cache can be deployed on any public, private cloud or hybrid environment, offering:

- Easy deployment and provisioning across any scale, on cloud, on-premises and hybrid environments.
- Hot rolling and automated deployment, upgrading, monitoring, scaling and recovery.
- Support of elastic orchestration systems such as Kubernetes, Red Hat OpenShift among others.
- Efficient and cost-effective data replication across regions/clouds or between on-premises and cloud deployments.



Enables Speedy Complex Queries

Supports an unlimited number of advanced indexes with SQL-99 compliance, only creating additional indexing data structures on one instance of data, enabling complex queries to be performed with optimal speed and minimal memory footprint.



Supports Multiple Data Models

Multi-model cache supporting structured, semi-structured and unstructured data models in a single solution without fragmenting the data to multiple different data stores.



Operational Reporting

Easily integrates with operational reporting and data integration tools.



Ensure Optimal Customer Experiences at Peak Data Volumes & High Concurrency



Native Multi-Criteria Queries

Supports an unlimited number of diverse indexes including nested objects, collections, compound index, geo-spatial and full text search for faster performance and smaller footprint.



Never Lose Writes

GigaSpaces provides strong consistency which ensures data integrity and data freshness, unlike other caching solutions which support eventual consistency between nodes.



Business Policy-Driven

Optimizes performance and cost with advanced business-driven policies that place the data in fast access RAM, or in SSDs, depending on the how quickly the data should be retrieved.



Fastest In-Memory Performance

Data processing is distributed on the server side in the same memory space as the data to accelerate performance and reduce networking and serialization overhead.

About GigaSpaces

GigaSpaces is building on its in-memory computing and operational data store technologies to offer one of the market's first Operational Data Hubs, an out-of-the-box unified real-time data platform that empowers organizations' digital transformation, by unlocking data that resides in legacy and cloud systems. Whether you need to accelerate one application with cache, or modernize your architecture,

GigaSpaces can future-proof your investment. Never before has it been this straightforward to accelerate API-powered digital applications and democratize access to data across your enterprise. Smart DIH is part of the GigaSpaces Smart Suite of products, alongside the award-winning Smart Cache solution. GigaSpaces offices are located in the US, Europe and Israel, with partners such as Capgemini and Cognizant around the globe, serving customers such as Morgan Stanley, Bank of America, CSX, Goldman Sachs, Societe Generale, Credit Agricole, American Airlines, Avanza Bank, Avaya, CLSA, and UBS.

For more information visit www.gigaspace.com.

