

VMware vFabric SQLFire

Fast, Scalable SQL Data in The Cloud

REDUCED LATENCY FOR SQL APPLICATIONS

Memory-based data management accelerates application performance, completely eliminating many disk and network latencies.

STANDARD SQL SYNTAX AND TOOLS

Database administrators can use and adapt SQLFire using their current skill sets and tools.

SCALES EASILY TO MEET THE HIGHEST DEMANDS

Scales memory out or back instantaneously to meet changes in loads or resource availability.

HIGH AVAILABILITY AND DISASTER RECOVERY

Ensures continuous availability within or across data centers, and supports disaster recovery with granularity to the level of individual tables.

IDEAL FOR HIGH TRANSACTION RATES

SQLFire is ideal for applications such as large transaction-oriented Web sites where disk and network overheads choke delivery of many small data items.

MEMORY ORIENTED AND CLOUD OPTIMIZED

Use of nonproprietary hardware offers an economical way to achieve high database performance at extremely large scale.

Overview

New VMware SQLFire replaces or augments disk-oriented database architectures with data structures and indexes optimized for fast main memory, with options for write-through or write-behind to disk. Built on the proven vFabric GemFire distributed data management platform, SQLFire is ideal for primary data stores that require high transaction rates, continuous availability, and support by database programming staff without specialized coding assistance. Deployed as the front-end data-management layer for an Oracle® or other RDBMS, SQLFire delivers substantially better performance and scalability for modest investments, and retains full compatibility with current DBA skills and tools.

VMware SQLFire Highlights

- **Low latency:** Memory-based data management maintains consistently high application performance by eliminating lookup, read/write, and network round-trip latencies.
- **Extreme write performance:** Memory-speed write performance is ideal for large-scale databases with high transaction volumes and demanding Service Level Agreements.
- **Simplified scale-out and scale-back:** Repartitions, replicates, and balances data across independent nodes to accommodate shifting loads or new resources.
- **Standard SQL interface:** Integrates fast, memory-based data management into applications using current database administration skills and tools.
- **Flexible HA and DR options:** Ensures continuous availability within or across data centers, with partial or full, synchronous or asynchronous writes to disk to meet disaster recovery or regulatory requirements.

VMware SQLFire

SQLFire enables dynamic horizontal scaling by creating a shared pool of memory across multiple standalone physical devices or x86 rack/blade hardware, and then growing or the shrinking the pool as demand changes.

Replicated or partitioned tables can be managed in memory alone, or in memory and on disk. Application logic can be routed to the data location, for increased performance from parallel execution.

Used as either the primary data store or a front-end data-management layer for one or more existing databases, SQLFire assures continuous availability for data within and across data centers. Any table can be configured to be replicated or partitioned into one or more redundant copies.

Built on a foundation of vFabric GemFire, SQLFire benefits from years of testing and reliability of its distribution sub-system, and adds a more sophisticated SQL query engine that compiles a query plan into byte codes, and a much more sophisticated cost-based optimizer. Anyone with relational database experience will find the SQLFire configuration and deployment model simple and intuitive to use and adapt. Unlike many popular data grids, SQLFire offers native persistence and recovery capabilities, and can be used as a distributed data store.

Its use of established standards such as SQL, JDBC and ADO.NET makes SQLFire straightforward to adopt in existing applications that use relational databases. Configuration and deployment is much simpler, and the product works effectively with a large eco-system of compatible products and frameworks, including object relational mapping tools (Hibernate, NHibernate, etc), schema-editing and database management tools (SquirrelSQL), database replication and change data capture products, Spring JDBC, and many more.

Applications using the standard SQL syntax supported by SQLFire can easily migrate to and from other relational databases, for flexibility and future-proofing as well as unparalleled performance.

