


Lava Distributed SQL

Tesseract Digital Information (Pty) Ltd

Legacy Mode

On the surface, Lava SQL is a conventional client-server relational SQL database.




Major functionality includes :

- Comprehensive SQL capability
- The most powerful relational engine available, including full support for logical schema layout (conditional relationships)
- The most powerful stored procedure language available - a procedural language fully capable of general purpose programming, but tightly integrated with the database
- Hot-standby server (replication server)
- Distributed (satellite) server
- Comprehensive and lightning-fast backup and restore
- Slice backup and restore (relational database slices)

The Lava database engine is probably the fastest available, and one of the best in terms of data capacity. On a moderate server, the following specifications may be achieved :

- Sustained insertion rate of 20,000 rows per second or more
- Row retrieval of over 30,000 rows per second
- Data tables of 500 million rows or more



Distributed Mode

In addition to the standard database capabilities, Lava Distributed SQL has comprehensive distribution capability - including a fully distributed client. This functionality provides a new dimension in database applications.

- Application design and implementation is identical to a single-user application
- Processing performed on the client workstation reduces server load
- Responsiveness of client application is vastly improved as the majority of queries are answered directly without referring to the server
- Workgroup capability in client applications may be achieved without any special programming
- Distributed database efficiency is such that modern notebooks have sufficient power to run in distributed mode

*Lava Distributed SQL is a product developed by
Tesseract Digital Information (Pty) Ltd*

www.lava-sql.com
support@qubik.net