

APPLICATION NOTE

BUILDING A POWERFUL ORACLE 10g CLUSTER



Application Description

For decades, large database systems were synonymous with costly mainframes. Today, however, grid computing is enabling clusters of commodity servers to achieve enterprise class performance, availability, and scalability, as well as simplified installation, management, and provisioning through network consolidation, all at a fraction of the cost of traditional SMP systems. Two key enabling technologies making scalable database clustering a reality in the enterprise are Oracle's Real Application Clusters (RAC) and **RAC Accelerator**, SilverStorm Technologies' high performance cluster interconnect for Oracle 10g RAC.

Oracle RAC enables IT managers to create scalable, highly available single image databases running on clusters of commodity servers. By leveraging low-cost hardware, RAC improves return on investment and reduces total cost of ownership. With RAC, the key to maximizing performance becomes the cluster interconnect. InfiniBand, with its 10X bandwidth advantage and 10X latency reduction vs. Gigabit Ethernet, provides the highest performance cluster interconnect for RAC. For interconnect-intensive applications such as data warehousing, Oracle recommends only InfiniBand for RAC clusters of greater than eight nodes.¹

InfiniBand brings multiple benefits to database clustering and grid computing; however the most significant value-add of InfiniBand for database clustering is as a very high performance cluster interconnect. Developed jointly by SilverStorm and Oracle, **RAC Accelerator** is based on the Reliable Datagram Sockets (RDS) interconnect protocol. In addition to the bandwidth and latency advantages of InfiniBand, **RAC Accelerator** leverages the reliable delivery capabilities inherent in InfiniBand that offload end-to-end error checking to the InfiniBand fabric, freeing CPU cycles for application processing, thereby enabling processor scaling far greater than is possible with Ethernet.

SilverStorm's RDS-based **RAC Accelerator** provides the highest performance interconnect available² for Oracle 10g RAC along with the enterprise-class capabilities required by Oracle's most demanding RAC customers:

- **high bandwidth** – 3–4 times throughput gain over UDP Gigabit Ethernet measured running 10g
- **low latency** – ½ latency of UDP Gigabit Ethernet and IP over InfiniBand
- **low CPU utilization** – ½ CPU utilization of UDP Gigabit Ethernet
- **ease of use** – **RAC Accelerator** is easy to install and configure, and is automatically used by Oracle 10g RAC³ once installed; no special setup is required within RAC to take advantage of RDS
- **high reliability** – minimal changes to Oracle's existing UDP IPC library to support RDS and approximately 70% reduction in Oracle's RAC UDP module functionality; RDS successfully passed Oracle's internal RAC regression test suite less than two weeks following initial delivery
- **high availability** – RDS supports fail-over across HCA ports as well as across HCAs, supporting full infrastructure redundancy for enterprise data center configurations

"The future is now. Oracle today has thousands of customers successfully running their business applications using Oracle Real Application Clusters. Oracle believes that InfiniBand will be the key technology that enables us to keep scaling out to larger and larger clustered configurations. We are very pleased to partner with SilverStorm to bring InfiniBand technology to the commercial marketplace."

Juan Loaiza
Vice President, Systems Technology Group
Oracle Corporation

With advanced features and deployments today in Fortune 500 data centers, SilverStorm Technologies is the leader in high performance, high availability cluster solutions for Oracle 10g RAC. See SilverStorm's latest solutions in Oracle's Enterprise Technology Centers in Atlanta, GA or Reston, VA, or for more information, visit www.silverstorm.com or contact info@silverstorm.com.



Oracle Enterprise Technology Centers

¹ session S015, "Top 5 Performance Tips for Tuning a Data Warehouse on GRID" by Oracle Corporation, Oracle OpenWorld, September 2005

² Oracle CRTEST run over RDS in SilverStorm's lab, December, 2005

³ with Oracle 10g release 10.2.0.2