egenera®

# Benchmark

## Clustered Database Performance and Scalability on the Egenera® BladeFrame® System

ORACLE CERTIFIED PARTNER

## Summary

Performance testing on the Egenera BladeFrame system has demonstrated that the platform is capable of delivering high throughput from multiple servers using Oracle Real Application Clusters (RAC) database software. Analysis using Oracle's Swingbench demonstration tool and the Calling Circle schema has shown very high transactions-per-minute performance from single-node implementations with dual-core, 4-socket SMP servers based on Intel and AMD architectures running a 64-bit-extension Linux operating system. Furthermore, results demonstrated 92 percent scalability on either server type up to at least 10 servers.

The BladeFrame's architecture naturally provides a host of benefits over other platforms in terms of manageability, server consolidation and high availability for Oracle RAC.

## BladeFrame Overview

The Egenera BladeFrame system is purpose-built to address the problems of large datacenters. Remote management, server consolidation, reduced wiring and switches, and high availability are some of the physical problems solved.

The system benefits of the hardware architecture provide for superior performance. The diskless, bladed form factor permits very-high-powered chips to be used. For example, an Egenera Processing Blade™ with 4-way, dual-core processors at 2.66 GHz or 2.4 GHz with 32 GB of memory provides the kind of processing power previously found only on large servers. Moreover, the BladeFrame backplane interconnect provides inter-server communications at approximately 2.5 Gb per second. These features result in very high throughput.

## Clustered Database

Oracle's Real Application Clusters (RAC) software provides a clustered or grid-oriented database platform for use with multiple databases. The platform enables multiple servers to access any of the databases running in the cluster. All servers operate independently, yet are interconnected to allow access to the same disk data simultaneously. RAC software provides the clustering elements for inter-server communication and efficient sharing of in-memory blocks, called "Cache Fusion."

## Problem Description

Every database deployment has unique requirements and may be implemented differently. In a clustered implementation, certain considerations must be made to ensure that the problem performs well in a parallel environment. The problem selected for this work has proven to be scalable.

The Calling Circle application has been developed in SQL, PL/SQL and Java by Michael Hallas and Dominic Giles of the Oracle Solutions Group. The Calling Circle application generates a heavy database workload and allows demonstration of failover/failback and load-balancing capabilities.

The Calling Circle application represents a self-service OLTP application that models the customers of a telecommunications company registering, updating and inquiring on a calling circle of their most frequently called numbers in order to receive discounted call pricing.

The Calling Circle application was originally developed to include aspects known to cause scalability challenges with RAC. For example, the database features several keys generated from sequence numbers, resulting in contention for right-growing indexes. The application also maintains activity counters and a history of changes to each calling circle, further increasing the proportion of insert and update statements.

The goal of the Calling Circle application is to simulate a randomized workload of customer transactions and to measure transaction throughput and response times. The test workload is specifically designed to perform a high proportion of database changes to reveal any contention issues. In fact, approximately 97 percent of customer transactions cause at least one database update, with well over three-quarters performing two or more updates.

A second problem was run to provide information on I/O throughput. The Swingbench team added a Sales History problem to represent large-data-warehouse reporting environments. A few simulated users run very large reports to stress the I/O subsystem. Scaling this problem out to multiple nodes provides an environment with high total throughput and I/Os per second.

## Performance and Scalability

A series of tests was run with single servers as a baseline for the scalability runs. These were performed with Processing Blades based on 4-way, dual-core 2.66 GHz Intel technology and 4-way, dual-core 2.4 GHz AMD technology. Red Hat Enterprise Linux 4 Update 3 for 64-bit extensions was used for the test. A set of 12 75-GB LUNs cut from 120 individual SAN spindles was used to store the database. Oracle's Automatic Storage Manager (ASM) was used to provide access to the storage disks.

Once the optimium set of users for the problem was determined, the tests were re-run for two, four, six, eight and 10 servers. The scalibility of the problem across these sets of nodes was found to be excellent with 92 percent at 10 nodes.

**Figure 1**
**Oracle 10*g* Relative Scalability**

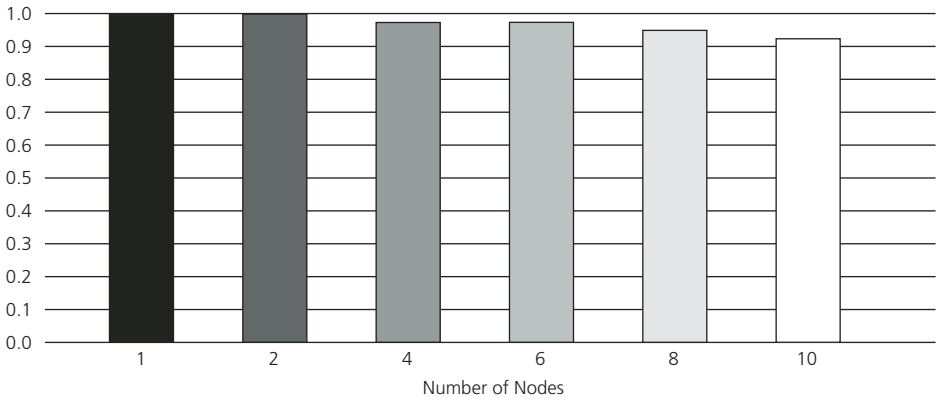


The 10-node cluster ran nearly 8,500 transactions per second. (Versions of Swingbench prior to 2.2 used a different metric for transactions per second, so this data is not directly comparable to early calculations.) Cache Fusion traffic for these runs, based on CPU utilization, shows very high utilization with internode throughput of 8 Mb per second.

Since the Calling Circle problem is primarily CPU-intensive, a second problem was run to demonstrate total system throughput. Twelve nodes were run simultaneously to create a large I/O load on the BladeFrame subsystem. The load reached approximately 200 Mb per second and nearly 20,000 I/O per second. In these tests, the storage configuration (120 available spindles) limited total system performance. Results confirm that Oracle RAC I/O requirements will be met if the SAN subsystem is large enough. This performance is a fraction of the total throughput and IOPS performance of which the BladeFrame is capable when storage is configured optimally.

## Conclusions

The Egenera BladeFrame has demonstrated its ability to provide an excellent platform for Oracle RAC. The tight intergration of the hardware provides a very-high-performance system that also offers many benefits for the Grid environment. Both CPU and I/O tests support the BladeFrame's ability to meet user needs for Oracle RAC database solutions.

egenera®

Corporate Headquarters
Egenera, Inc.
165 Forest Street
Marlboro, MA 01752
U.S.A.
Phone: 508-858-2600
Fax: 508-481-3114
www.egenera.com

European Headquarters
Egenera Ltd.
FSC Campus
The Boulevard
Cain Road
Bracknell, Berkshire RG12 1HH
Phone: +44 (0)1344 475237
Fax: +44 (0)1344 475775
www.egenera.com

Asia Pacific Headquarters
Egenera (Hong Kong) Limited
Suite 1903
Central Plaza
18 Harbour Road,
Wanchai, Hong Kong
Phone: 011 852-2877-9101
Fax: 011 852-2877-8611
www.egenera.com