Systems/Enterprise:

# Tangosol Demos Data Grid Speed, Adds Spring Support

Tangosol announced that they, along with Intel and IBM, have completed Grid-based benchmarks that demonstrate the linear scalability of parallel processing of data using Tangosol Coherence Data Grid running on a computing grid of 100 IBM servers based on the Intel Xeon processor 5100 series.

A U.S.-based investment bank, a public company with more than $600 billion in assets under management, initiated this benchmark test to establish a performance and scalability baseline for Coherence. This Tangosol customer has been using Coherence since mid-2005. Coherence is the distributed in-memory data grid backbone for several trading and mission-critical applications at the company.

The requirements for the benchmark identified by the bank included:

- reliability and high availability;

- linear scalability (scale-out);

- very high performance (scale to the bandwidth of the network);

- ease of use and extensibility.

The benchmark results showed that Coherence enables linear scaling of parallel processing across the grid as the grid (and the data set) increases in size and is limited only by aggregate CPU cycles across the grid. In the test, Coherence was able to linearly scale from 2 million aggregations with two servers to more than 60 million aggregations across the 96 servers. This 30 times increase in processing throughput was achieved with only one tenth of a second increase in processing time, or 1.2 seconds compared to 1.1 seconds. Additionally, the tests demonstrated that the data grid storage capacity increases linearly as additional resources are added to the grid and is limited only by the amount of RAM available to the data grid.

Coherence™ Data Grid Throughput
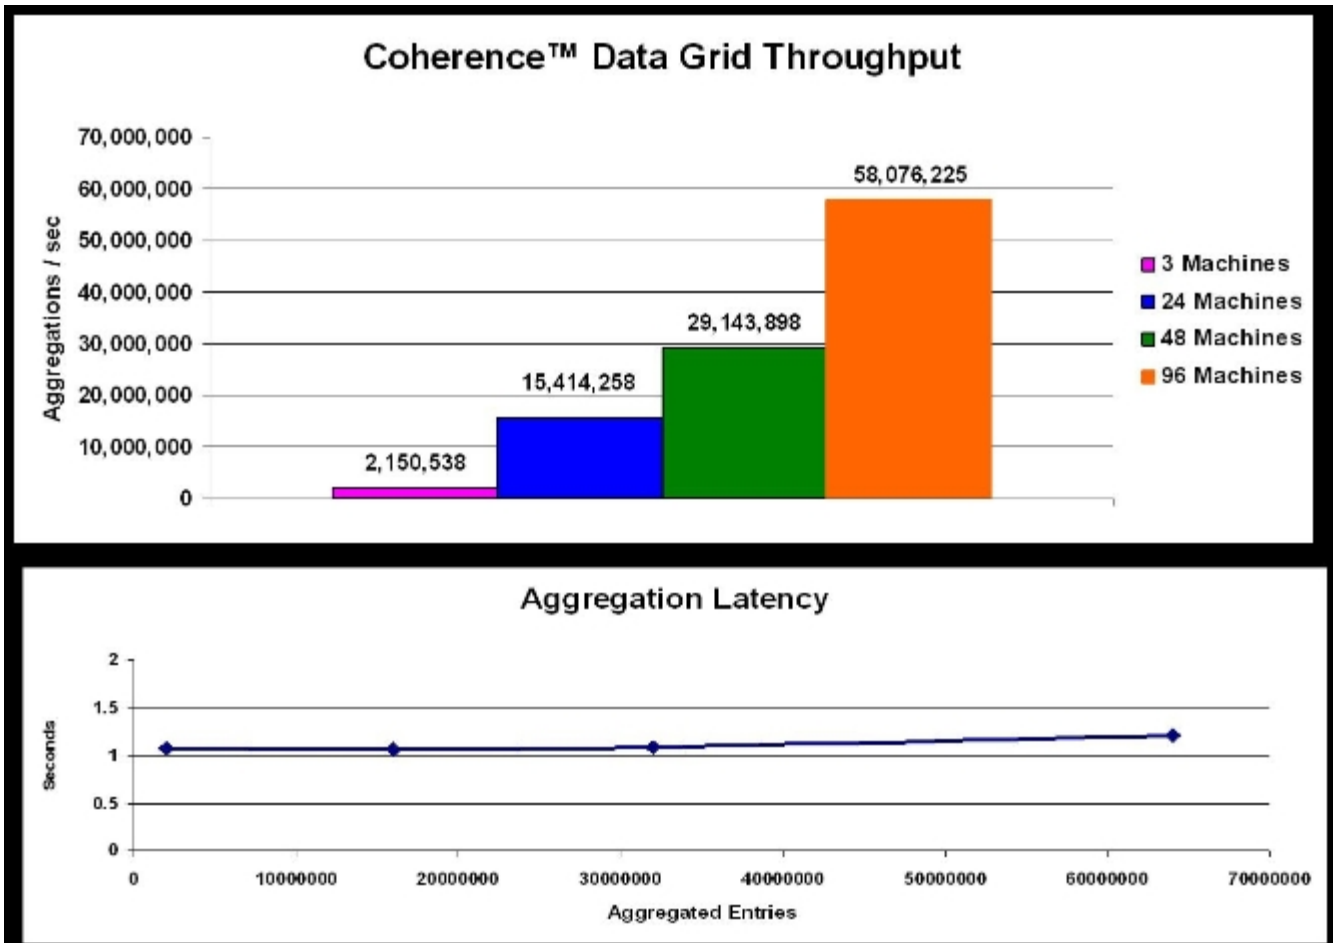
The tests were on a 100-server grid, with dual 3.6 GHz CPUs, four Intel dual-core Intel Xeon 5100 processors, and a 1 Gb network with 4 Gb interconnects between blade centers.

"These tests indicate that companies using Coherence Data Grid can achieve unprecedented performance on industry-standard hardware, for operational and cost advantages," explained Cameron Purdy, CEO of Tangosol. "Because Coherence is so finely optimized, it fully utilizes all resources by moving the data closer to the applications -- in memory -- where it can be available instantaneously as demand grows."

"This benchmark of Tangosol Coherence Data Grid successfully demonstrates predictable scalability," noted Elliot Garbus, general manager of Intel's Developer Relations Division. "As the industry moves from dual-core to quad-core processors like the new Quad-Core Intel Xeon processor 5300 series for data grid servers and the new Intel Core 2 Extreme quad-core processors for data grid clients, enterprise customers will be able to achieve even greater performance and energy efficiency."

"This internal benchmark proves that the improvements to today's computing hardware combined with capabilities of products like Tangosol Coherence both validate and greatly enhance the power of grid computing," said Tamara Crawford, program director for IBM's Grid Ecosystem Strategy. "IBM is pleased to partner with companies, like Tangosol, to demonstrate the benefits of grid technology, to extend the scope and capabilities of IBM Grid offerings and to speed enterprise adoption of grid computing."

Tangosol also performed parallel aggregation and processing tests across the grid, which demonstrated a linear scale model allowing for such processing to take less time as the grid grows, or to process more data as the grid grows.

In a separate announcement, Interface21, the provider of the open source Spring application Framework, and Tangosol, introduced Coherence Data Grid for Spring, bringing enterprise-class data availability to the Spring Framework. With Coherence Data Grid for Spring, Spring applications can directly tap into the power and scalability of the data grid to meet increasing information access demands.

Coherence Data Grid manages data in memory across data grids, offering scalable performance that enables true linear growth; automatic, dynamic and transparent reliability of data; and reduced latency and rapid data access for customers, partners and employees.

Coherence Data Grid for Spring will enable standalone Spring applications to utilize the resiliency of a Coherence Data Grid. Spring applications may then transparently publish messages and share state in a recoverable, scalable, and high-performance manner. Spring application events are no longer constrained to the scope of a single JVM, and can now work together, scaling out data and processing capacity across a data grid.

An open source project since 2003, Spring is an open source Java/JEE application framework, with well over one million downloads. It simplifies enterprise Java development and testing and makes existing technologies easier to use. Spring Framework was named as a strong performer in open source projects in the The Forrester Wave Project Summary, Q2 2006.

"Spring Framework is an open source Java application framework that is gaining rapid popularity among Java developers," noted Michael Goulde, senior analyst, Forrester. "Spring Framework's high degree of automation, simplification of development, and simplified testing make it a good choice for Java developers looking to simplify and accelerate their development process when building enterprise applications."

"The modular, flexible characteristics of Spring combined with the proven performance, reliability, and scalability of Tangosol's Coherence Data Grid will permit standalone Spring applications to share and scale out data processing," noted Rod Johnson, founder

of the Spring Framework and CEO of Interface21. "Most of the world's top investment banks have selected the Spring Framework, and Interface21 and Tangosol share a growing customer base. Clustering through Coherence Data Grid for Spring will provide the added benefits of clustered reliability and linear scalability, which are so critical to their operations and growth."

Coherence Data Grid for Spring will feature a new type of Spring "Bean" component, the Spring Data Grid Bean. Spring Beans may be automatically and transparently managed in highly available data grids built on top of Coherence. This allows Spring Beans to break free of single-server environments and be scalable across a data grid. Through the use of innovative Coherence in-memory data management technology, Data Grid Beans appear and behave as if they are traditional Spring Beans, and may be may be queried, indexed, aggregated, and updated in parallel across a data grid. As data grid capacity increases, so too does the capacity of each Spring application.

In addition, Spring applications using Data Grid Beans may fail and recover, or alternatively be restarted, without the need to manually reconstruct the state of their said beans (from databases or disk), thus improving the availability and resilience of Spring applications.

Lastly, based on Coherence data management technologies, Data Grid beans used in Spring may be transparently accessed from other non-Java platforms, without the need for complicated and intrusive bridge/messaging/xml-based technologies.

Coherence Data Grid for Spring will be available for developers this year and generally available in the first quarter of 2007.