

Contacts:

Investor Relations

Nancy Fazioli
650.224.8291
ir@fusionio.com

Media Relations

Robert Brumfield
917.224.7769
bbrumfield@fusionio.com

**Fusion-io and Lawrence Livermore National Laboratory
Traverse a 68 Billion Node Graph**

Powerful Algorithm and 12 TB of Fusion ioMemory Break Previous LLNL Record

SEATTLE – Nov. 15, 2011 – Fusion-io (NYSE: FIO) today announced that Lawrence Livermore National Laboratory recently broke its previous graph size record set in the June 2011 Graph500 competition with a single node server containing 12 TB of ioMemory technology from Fusion-io.

The LLNL system, called “Leviathan,” is based on a four-socket 40-core Intel Xeon 7500 processor, nine Fusion ioDrive Duos, and one ioDrive. By storing the graph in Fusion-io’s direct-attached high performance ioMemory technology, the LLNL algorithm could use a single computer to process a graph with 68,719,000,000 nodes (known as “scale 36” meaning 2^{36} nodes) – four times the size previously attained. The Fusion ioDrives provided low latency access to graph nodes and edges, enabling a unique alternative for data intensive computing. In contrast, most other Graph500 submissions rely on expensive supercomputers with many hundreds to thousands of compute nodes.

“As a multidisciplinary national security R&D laboratory that deploys some of the world’s most powerful computing systems, we must continually push the state of the art in such areas as data intensive computing, which is increasingly important to our mission work,” says Maya Gokhale of LLNL. “To advance the data-intensive computing capabilities we need for applications in cyber security and informatics, we work with industry leaders such as Fusion-io. The innovative use of technology has allowed us to solve large data intensive problems on commodity hardware.”

The Fusion-io solution also provides significant scalability, as LLNL demonstrated a second scale 36 result with Fusion ioMemory, using 64 compute nodes, each integrated with two Fusion-io ioDrive Duos on the Hyperion Data Intensive Testbed. The Hyperion Data Intensive Testbed is a LLNL and Fusion-io collaboration to explore data-intensive computing systems that leverage low latency flash memory. The single node Leviathan system achieved 52.796 million traversed edges per second (TEPS), the Graph500 speed metric, while the Hyperion DIT attained more than an order of magnitude higher TEPS.

“Lawrence Livermore’s achievement is significant not only for its impressive results, but also due to the impact this extremely efficient system could have on the computing industry as a whole,” said Neil Carson, Fusion-io Chief Technology Officer. “When you decouple performance from scale, customers no longer have to pay twice for performance in their server and their SAN. Given the energy and costs associated with building out and running a supercomputer, LLNL’s benchmarks point to vast savings while achieving the top tier performance required by leading HPC research. LLNL was one of the first to adopt Fusion-io technology, and we congratulate their visionary team for beating their previous results with Fusion ioMemory and their powerful algorithm on a single node.”

Powered by Fusion’s Virtual Storage Layer (VSL) software, Fusion ioMemory is integrated within the server to offer applications and databases low latencies to deliver advanced performance and scalability, as well as enterprise reliability. Fusion ioMemory is integrated into a number of HPC solutions at Supercomputing 2011 (SC 2011) in Seattle taking place November 14 - 17. For details on Fusion-io Alliance demonstrations at SC 2011, go to www.fusionio.com/blog.

Visit Fusion-io in booth 5607 at SC 2011 to learn more, or visit www.fusionio.com. Follow Fusion-io on Twitter at www.twitter.com/fusionio or www.twitter.com/fusionioUK and on Facebook at www.facebook.com/fusionio.

About Fusion-io

Fusion-io has pioneered a next generation storage memory platform for shared data decentralization that significantly improves the processing capabilities within a datacenter by relocating process-critical, or “active”, data from centralized storage to the server where it is being processed, a methodology referred to as data decentralization. Fusion’s integrated hardware and software solutions leverage non-volatile memory to significantly increase datacenter efficiency and offers enterprise grade performance, reliability, availability and manageability. Fusion’s data decentralization platform can transform legacy architectures into next generation datacenters and allows enterprises to consolidate or significantly reduce complex and expensive high performance storage, high performance networking and memory-rich servers. Fusion’s platform enables enterprises to increase the utilization, performance and efficiency of their datacenter resources and extract greater value from their information assets.

Note on Forward-looking Statements

Certain statements in this release may constitute “forward-looking statements” within the meaning of Section 21E of the Securities Exchange Act of 1934 and Section 27A of the Securities Act of 1933. These statements are based on current expectations and assumptions regarding future events and business performance and involve certain risks and uncertainties that could cause actual results to differ materially from those contained, anticipated, or implied in any forward-looking statement, including the risks set forth in the registration statements and reports that Fusion-io files with the U.S. Securities and Exchange Commission, which are available on the Investor Relations section of our website at www.fusionio.com. You should not rely upon forward-looking statements as predictions of future events. Although we believe that the expectations reflected in the forward-looking statements are reasonable, we cannot guarantee that the future results, levels of activity, performance or events and circumstances reflected in the forward-looking statements will be achieved or will occur. Fusion-io undertakes no obligation to update publicly any forward-looking statement for any reason after the date of this press release.

###