

November 12, 2012



# AMD Introduces Industry's Most Powerful Server Graphics Card

## AMD FirePro S10000 Is the World's First Professional Graphics Card to Exceed One TeraFLOPS of Peak Double Precision Performance and Unparalleled Single Precision Performance

SUNNYVALE, CA -- (Marketwire) -- 11/12/12 -- [AMD](#) (NYSE: AMD) today launched the [AMD FirePro™ S10000](#), the industry's most powerful server graphics card, designed for high-performance computing (HPC) workloads and graphics intensive applications. The AMD FirePro S10000 is the first professional-grade card to exceed one teraFLOPS (TFLOPS) of double-precision floating-point performance, helping to ensure optimal efficiency for HPC calculations(1). It is also the first ultra high-end card that brings an unprecedented 5.91 TFLOPS of peak single-precision and 1.48 TFLOPS of double-precision floating-point calculations(2). This performance ensures the fastest possible data processing speeds for professionals working with large amounts of information. In addition to HPC, the FirePro S10000 is also ideal for virtual desktop infrastructure (VDI) and workstation graphics deployments.

"The demands placed on servers by compute and graphics-intensive workloads continues to grow exponentially as professionals work with larger data sets to design and engineer new products and services," said David Cummings, senior director and general manager, Professional Graphics, AMD. "The AMD FireProS10000, equipped with our Graphics Core Next Architecture, enables server graphics to play a dual role in providing both compute and graphics horsepower simultaneously. This is executed without compromising performance for users while helping reduce the total cost of ownership for IT managers."

Equipped with AMD next-generation [Graphics Core Next Architecture](#), the FirePro S10000 brings high performance computing and visualization to a variety of disciplines such as finance, oil exploration, aeronautics, automotive design and engineering, geophysics, life sciences, medicine and defense. With dual GPUs at work, professionals can experience high throughput, low latency transfers allowing for quick compute of complex calculations requiring high accuracy.

### *Responding to IT Manager Needs*

With two powerful GPUs in one dual-slot card, the FirePro S10000 enables high GPU density in the data center for VDI and helps increase overall processing performance. This makes it ideal for IT managers considering GPUs to sustain compute and facilitate graphics intensive workloads. Two on-board GPUs can help IT managers reap significant cost savings, replacing the need to purchase two single ultra-high-end graphics cards, and can help reduce total cost of ownership (TCO) due to lower power and cooling expenses.

### *Key Features of AMD FirePro™ S10000 Server Graphics*

- Compute Performance: The AMD FirePro S10000 is the most powerful dual-GPU server graphics card ever created, delivering up to 1.3 times the single precision and up to 7.8 times peak double-precision floating-point performance of the competition's comparable dual-GPU product. It also boasts an unprecedented 1.48 TFLOPS of peak double-precision floating-point performance(3);
- Increased Performance-Per-Watt: The AMD FirePro S10000 delivers the highest peak double-precision performance-per-watt -- 3.94 gigaFLOPS -- up to 4.7 times more than the competition's comparable dual-GPU product(3);
- High Memory Bandwidth: Equipped with a 6GB GDDR5 frame buffer and a 384-bit interface, the AMD FirePro S10000 delivers up to 1.5 times the memory bandwidth of the comparable competing dual-GPU solution(4);
- DirectGMA Support: This feature removes CPU bandwidth and latency bottlenecks, optimizing communication between both GPUs. This also enables P2P data transfers between devices on the bus and the GPU, completely bypassing any need to traverse the host's main memory, utilize the CPU, or incur additional redundant transfers over PCI Express®, resulting in high throughput low-latency transfers which allow for quick compute of complex calculations requiring high accuracy;
- OpenCL™ Support: OpenCL has become the compute programming language of choice among developers looking to take full advantage of the combined parallel processing capabilities of the FirePro S10000. This has accelerated computer-aided design (CAD), computer-aided engineering (CAE), and media and entertainment (M&E) software, changing the way professionals work thanks to performance and functionality improvements.

Please visit AMD at [SC12](#), booth #2019, to see the AMD FirePro S10000 power the latest in graphics technology.

### *Supporting Resources*

- Learn more about [AMD Firepro Server Graphics](#)
- Learn more about [AMD Professional Graphics](#)
- Learn more about [GCN Architecture](#)
- Follow AMD professional graphics news on Twitter at [@AMDFirePro](#)
- Become a fan of AMD technology on [Facebook](#)

### *About AMD*

AMD (NYSE: AMD) is a semiconductor design innovator leading the next era of vivid digital experiences with its ground-breaking AMD Accelerated Processing Units (APUs) that power a wide range of computing devices. AMD's server computing products are focused on driving industry-leading cloud computing and virtualization environments. AMD's superior graphics technologies are found in a variety of solutions ranging from game consoles, PCs to supercomputers. For more information, visit <http://www.amd.com>.

*AMD, the AMD Arrow logo, AMD FirePro, Radeon, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. OpenCL is a trademark of Apple, Inc. and used by permission of Khronos. Other names are for informational purposes only and may be trademarks of their respective owners.*

(1) AMD FirePro™ S10000 delivers 1.48 TFLOPS peak dual-precision floating-point performance -- no other AMD or Nvidia server graphics card has exceeded one TFLOPS before. Prior to launch of AMD FirePro S10000, AMD's highest performing graphics card for double precision was the AMD FirePro™ W9000 with 1.0 TFLOPS, and Nvidia's highest

performing card in the market as of Oct. 31, 2012 is the Tesla M2090 with 665 GFLOPS double precision. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-71

(2) AMD FirePro™ S10000 delivers up to 5.91 TFLOPS of peak single-precision and 1.48 TFLOPS of peak double-precision floating-point performance, compared to Nvidia Tesla K10 that is capable of up to 4.58 TFLOPS of peak single-precision and 190 GFLOPS peak double-precision floating-point performance. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. Comparison as of 10/31/12. FP-65

(3) Based on comparison as of Oct. 31, 2012 between AMD FirePro S10000 with 1.48 TFLOPS peak double precision, max board power of 375W and 3.94 GFLOPS double-precision performance-per-watt. Compared to the dual-GPU Nvidia Tesla K10 with 190 GFLOPS peak double precision, 225W max board power, and less than 1 GFLOPS per watt double-precision performance (0.84). Nvidia Tesla K10 product specs found here [http://www.nvidia.com/content/PDF/kepler/Tesla\\_K10\\_BD-06280-001\\_v05.pdf](http://www.nvidia.com/content/PDF/kepler/Tesla_K10_BD-06280-001_v05.pdf). FP-68

(4) AMD FirePro™ S10000 features 480 GB/sec memory bandwidth compared to Nvidia Tesla K10 with 320 GB/sec memory bandwidth. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-66

[Add to Digg](#) [Bookmark with del.icio.us](#) [Add to Newsvine](#)

Contact:  
Dave Erskine  
AMD Public Relations  
(289) 695-0903  
[dave.erskine@amd.com](mailto:dave.erskine@amd.com)

Matthew Kanas  
Edelman for AMD  
(416) 849-3324  
[matthew.kanas@edelman.com](mailto:matthew.kanas@edelman.com)

Source: Advanced Micro Devices