




AMD FirePro™ S9150 Server GPU

The world's most powerful server GPU
for high-performance computing¹

Key Features:

- Passively cooled solution for server environments
- AMD Graphics Core Next architecture
- 2,816 stream processors (44 compute units)
- 5.07 TFLOPS peak single precision floating point
- 2.53 TFLOPS peak double precision floating point
- 1/2 rate double precision
- Error correcting code memory support (external only)
- 16GB ultrafast GDDR5 memory
- 512-bit memory interface
- Up to 320GB/s memory bandwidth
- 235W maximum power consumption
- Support for SMBus temperature reporting at boot up
- AMD PowerTune technology²
- AMD STREAM technology³
- OpenCL™, DirectX® and OpenGL support
- PCIe® x16 bus interface, PCIe 3.0 compliant
- Full height/full length dual-slot form factor
- Headless display support
- Microsoft® Windows® 7, Windows 8.1 and Linux OS support (64- and 32-bit)
- Microsoft Windows Server® 2012 R2 and Windows Server® 2008 R2 SP1 support
- FCC, CE, C-Tick, BSMI, KCC, ULVCCI, RoHS and WEEE compliance
- Designed, built and tested by AMD
- Planned minimum three-year life cycle
- Limited three-year warranty



The most compute-intensive workloads and complex calculations are no challenge for the AMD FirePro S9150 server GPU. With support for OpenCL™ 1.2, 16GB GDDR5 memory, up to 2.53 TFLOPS of peak double-precision and up to 10.8 GFLOPS-per-watt peak double-precision performance, the choice is clear. AMD FirePro S9150 GPUs deliver unparalleled compute performance and performance-per-watt.

Industry-leading GPU Compute Performance¹

As the first server GPU to feature 1/2 rate double precision⁴ and break the 2.0 TFLOPS double-precision barrier,⁵ the AMD FirePro S9150 brings up to 5.07 TFLOPS of peak single-precision and up to 2.53 TFLOPS of peak double-precision compute performance.

Industry-leading Performance-per-Watt¹

A dual PCIe® slot form factor consuming 235W at maximum power, the AMD FirePro S9150 delivers up to 21.6 GFLOPS-per-watt of single-precision floating point performance, and up to 10.8 GFLOPS-per-watt of double-precision performance – up to 77% more than the competition¹. With AMD FirePro S9150 server GPUs, supercomputers and computing clusters can achieve massive compute performance and processor density within an eco-conscious power budget.

Industry-leading Memory Configuration⁶

The AMD FirePro S9150 is the industry's first server GPU with 16GB ultra-fast GDDR5 onboard memory – 33% more than the competition.⁷ It also features a 512-bit memory interface for up to 320GB/s of memory bandwidth, helping to improve overall workload speed and system responsiveness.

AMD FirePro™ S9150 High-Performance Server GPU ▲

FEATURES	BENEFITS
AMD Graphics Core Next (GCN) Architecture	The first GPU architecture designed with compute workloads in mind. Engineered for high utilization, high throughput and multitasking.
5.07 TFLOPS of Peak Single-precision Floating Point Performance	Helps speed up time required to complete single-precision operations used within Video Effects and Rendering, Signal Processing, Transcoding and Digital Rendering applications where high performance takes precedence over accuracy.
2.53 TFLOPS of Peak Double-Precision Floating Point Performance	Helps speed up time required to complete double-precision operations used within Computational Fluid Dynamics, Structural Mechanics, Reservoir Simulation and Aerodynamics applications, where numerical precision is mission critical.
½ Rate Double Precision	Unlike competing server cards that are not optimized for double precision, AMD FirePro S9150 server GPUs offer the most double-precision performance ⁵ , completing compute-intensive tasks faster than ever before.
Error Correcting Code (ECC) Memory Support	Helps ensure the accuracy of computations by correcting any single or double bit error as a result of naturally occurring background radiation. External support only.
Multi-GPU Support	Combine more than 10 AMD FirePro S9150 server GPUs in a single system and leverage the combined processing power to tackle the most compute-intensive workloads.
16GB GDDR5 Memory	Helps to accelerate memory-intensive applications and process larger and more computationally complex workloads with ease.
AMD PowerTune Technology ²	An intelligent system that performs real-time analysis of applications that utilize a GPU. In the event that an application is not making the most of the power available to the GPU, AMD PowerTune can improve that application's performance by raising the GPU's clock speed by up to 30% – automatically.
AMD STREAM Technology ³	Powers the ecosystem that enables AMD FirePro™ S-Series server cards to be used for compute-intensive workflows leveraging the massively parallel processing power of AMD GPUs, and accelerate many applications beyond just graphics.
Ready to Support OpenCL™ 2.0 ⁸	The AMD FirePro S9150 server GPU is expected to support OpenCL™ 2.0, allowing developers to take advantage of new features that give GPUs more freedom to do the work they are designed to do.



For more information, please visit
www.amd.com/firepro

1. AMD FirePro™ S9150 max power is 235W and delivers up to 2.53 TFLOPS peak double and up to 5.07 peak single precision floating point performance. Nvidia's highest performing server cards in the market as of June 2014 are the Tesla K40, max power of 235W, with up to 1.43 TFLOPS peak double and up to 4.29 peak single, and the K10, max power 225W, with up to 4.58 TFLOPS peak single and 190 GFLOPS peak double precision. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-97.

2. AMD PowerTune technology is offered by certain AMD FirePro™ graphics products and is designed to intelligently manage GPU power consumption in response to certain GPU load conditions. Not all products feature all technologies – check with your component or system manufacturer for specific model capabilities.

3. AMD STREAM technology is a set of features offered with select AMD FirePro graphics cards for the acceleration of compute-intensive workflows. Not all products have all features and full enablement of some capabilities may require complementary software. Check with your system manufacturer for specific capabilities and supported technologies.

4. AMD FirePro™ S9150 supports ½ rate double precision; 5.07 TFLOPS peak single precision divided by two is 2.53 TFLOPS peak double. Nvidia Tesla K40 supports 1/3 rate double precision, Nvidia's highest rate for server GPUs as of June 2014; 4.29 TFLOPS single precision divided by three is 1.43 TFLOPS double. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-96.

5. AMD FirePro™ S9150 delivers up to 2.53 TFLOPS peak double precision floating point performance, and Nvidia's highest server GPU in the market as of June 2014 is the Tesla K40 with up to 1.43 TFLOPS peak double precision. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-95.

6. AMD FirePro™ S9150 features 16GB GDDR5 memory, a 512-bit memory interface and up to 320 GB/s memory bandwidth, and Nvidia's highest performance server GPU in the market as of June 2014 is the Tesla K40 with 12GB GDDR5 memory, a 384-bit memory interface and up to 288 GB/s memory bandwidth. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-100.

7. AMD FirePro™ S9150 features 16GB GDDR5 memory, and Nvidia's highest performance server GPU in the market as of June 2014 is the Tesla K40 with 12GB GDDR5 memory. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-98.

8. AMD plans to release OpenCL 2.0 drivers for AMD FirePro graphics cards in Q4 2014; conformance testing is planned at that time.

