



HOW TO SELL THE NEW

# AMD FirePro™ S9150

*The World's Most Powerful Server GPU for High-Performance Computing<sup>1</sup>*

## Why it's great

The AMD FirePro™ S9150 server GPU is designed for server environments to accelerate graphics intensive applications and process the most compute-intensive workloads found in:

- Academic and Government Clusters
- Oil and Gas Industries
- Deep Neural Networks
- Virtualization

...and more



## Sell it in 5 seconds

- Designed for **large-scale multi-GPU support and unmatched compute performance**
- The first Server GPU to support enhanced double precision and break the 2.0 TFLOPS double precision barrier<sup>1,2</sup>
- Delivers up to **10.8 GFLOPS** double-precision performance-per-watt – **up to 77% more performance-per-watt than the competition<sup>1</sup>**
- Support for open standards including OpenCL 2.0



## Sell it in 60 seconds

- Based on the AMD Graphics Core Next (GCN)GPU architecture, it is capable of delivering upto 2.53 TFLOPS of peak double-precision floating-point performance
- The AMD FirePro™ S9150 offers the highest peak double-precision performance – **delivering up to 77% more than the competition<sup>1</sup>**
- Features a 512-bit memory interface and up to 320GB/s of memory bandwidth, **helping to improve overall workload speed and system responsiveness**
- Offers the **highest performance-per-watt**, up to 10.8 GFLOPS-per-watt double-precision performance
- Features 16GB of ultrafast GDDR5 memory – **33% more than the competition<sup>3</sup> – which can help accelerate memory-intensive applications and process larger and more computationally complex workloads**





HOW TO SELL

THE WORLD'S MOST POWERFUL SINGLE-GPU SERVER CARD FOR HIGH-PERFORMANCE COMPUTING

Why it's great

<b>5.07 TFLOPS of peak single-precision floating-point performance</b>	Helps <b>speed up time required</b> to complete single-precision floating-point operations used within simulations, video enhancement, signal processing, video transcoding, and digital rendering applications <b>where high performance takes precedence over accuracy.</b>
<b>2.53 TFLOPS of peak double-precision floating-point performance</b>	Helps <b>speed up time required</b> to complete double-precision floating-point operations used within computational fluid dynamics, structural mechanics, reservoir simulation, and aerodynamics applications <b>where numerical precision is mission critical.</b>
<b>16GB GDDR5 memory</b>	Equipped with 16GB of GDDR5 memory, the <b>S9150 can accelerate memory-intensive applications and process larger and more computationally complex workflows with ease.</b>
<b>Error correcting code (ECC) memory support</b>	Helps <b>ensure the accuracy of your computations by correcting any single or double bit error</b> as a result of naturally occurring background radiation; supported on GDDR5 memory only.
<b>AMD STREAM technology<sup>4</sup></b>	Powers the ecosystem that enables AMD FirePro™ S-series server cards to be used for compute-intensive <b>workflows leveraging the massively parallel processing power of AMD GPUs, and to accelerate many applications beyond just graphics.</b>
<b>Supports Open Standards</b>	The AMD FirePro S9150 server GPU supports 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. The S9150 is expected to support OpenMP 4.0, further solidifying AMD's commitment to open standards. <sup>5</sup>

How they stack up<sup>5</sup>

	NVIDIA TESLA K40	AMD FIREPRO™ S9150	AMD ADVANTAGE
<b>OpenCL™ Support</b>	1.1	<b>2.0</b>	Yes
<b>Peak Single Precision</b>	4.29 TFLOPS	<b>5.07 TFLOPS</b>	18%
<b>Peak Double Precision</b>	1.43 TFLOPS	<b>2.53 TFLOPS</b>	77%
<b>Performance/watt DFPF</b>	up to 6.09 GFLOPS/W	<b>up to 10.8 GFLOPS/W</b>	Yes
<b>Performance/watt SPFP</b>	up to 18.3 GFLOPS/W	<b>up to 21.6 GFLOPS/W</b>	Yes
<b>Performance/\$DFPF*</b>	.286 GFLOPS/dollar	<b>.603 GFLOPS/dollar</b>	Yes
<b>Double Precision Rate</b>	1/3	<b>1/2</b>	Yes
<b>Memory Size &amp; Bandwidth</b>	12GB GDDR5 and 288GB/s	<b>16GB GDDR5 and 320GB/s</b>	Yes
<b>Maximum Power</b>	235W	<b>235W</b>	-

NVIDIA Tesla K40 data source: <http://www.nvidia.com/object/tesla-servers.html>  
 \*Performance-per-dollar based on AMD FirePro S9150 \$4199 MSRP and \$5000 street price for K40

For more information, please visit [in.amdfireprohub.com/resources/partners/hp/](http://in.amdfireprohub.com/resources/partners/hp/)

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 single GPU server card in the market as of January 2015 is the Tesla K40, max power of 235W, with up to 1.43 TFLOPS peak double and up to 4.29 TFLOPS peak single-precision. Visit <http://www.nvidia.com/object/tesla-servers.html> for Nvidia product specs. FP-97
2. AMD FirePro™ S9150 delivers up to 2.53 TFLOPS peak double-precision floating-point performance, and Nvidia's highest single GPU server card in the market as of January 2015 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-97
3. 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 single GPU server card in the market as of January 2015 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
4. 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.
5. AMD plans to release OpenMP 4.0 drivers for AMD FirePro graphics cards in Q2 2015; conformance testing planned at that time.

