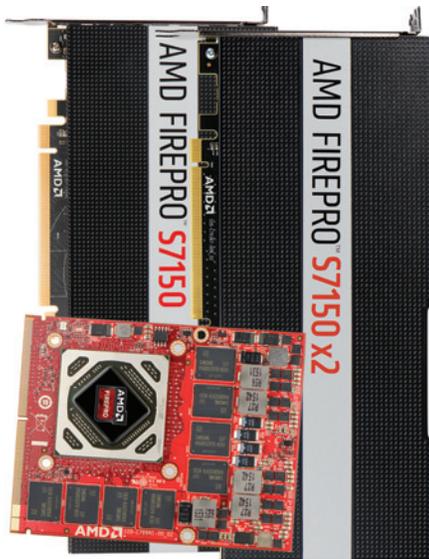


# Quality of Service in Virtualized Graphics Deployments Using AMD MxGPU Technology

## AMD FirePro™ S7100x, S7150, and S7150x2 Specifications

- **Max. Power:** 100W (S7100x), 150W (S7150), 265W (S7150x2)
- **Form Factor:** PCIe® MXM 3.1 (S7100x), or full height/length PCIe x16 (S7150/S7150x2)
- **Cooling:** Passive (active available for S7150)
- **RAM:** 8GB (S7100x/S7150) or 16GB GDDR5 (S7150x2)
- **Interface:** 256-bit
- **Performance:** 3.77 TFLOPS single-precision and 250 GFLOPS double-precision peak floating point performance (S7150). 7.54 TFLOPS single-precision and 500 GFLOPS double-precision peak floating-point performance (S7150x2).
- **OS:** Microsoft® Windows® 10, 8.1, 7, and Linux® (32- or 64-bit)
- **Compatible Virtualization Technologies:** Citrix® & VMware®
- **Concurrent Users:**  
Up to 16 (S7100x/S7150)  
Up to 32 (S7150x2)



*AMD MxGPU graphics virtualization technology delivers leading Quality of Service (QoS) with performance that is up to nine times more consistent than the competition<sup>2</sup>.*

## QUALITY OF SERVICE?

Predictable, reliable infrastructure is paramount in an enterprise environment. Traditional workspaces ensure predictability using physically isolated personal computers on each user's desk. Modern virtualized deployments have more dependencies for predictability and reliability. For example, the activity of other users can directly impact your experience, or Quality of Service (QoS).

## GRAPHICS IN THE CLOUD

Public and private cloud infrastructures are increasingly adopting dedicated graphics hardware using Graphics Processing Units (GPUs). More and more Independent software vendors (ISVs) are integrating GPU acceleration into their products. Users of applications such as Microsoft PowerPoint®, Autodesk Revit™, and YouTube can use GPUs to provide a richer experience. This trend is causing enterprises to start paying attention to graphics QoS.

## NOISY NEIGHBORS

Imagine an architect who needs to render a kitchen design using different cabinets. She starts the rendering, but notices unusually slow performance. Investigation reveals that a nearby coworker began a similar render at almost the same moment. Nobody could predict this timing, and neither user can afford degraded performance.

## UNTRUSTED NEIGHBORS

Many enterprises have at least one user who responds to spam and clicks flashing ads despite warnings. These users can be more likely to infect enterprise infrastructure with malware that can degrade performance as it spreads from virtual machine to virtual machine.

## AN ELEGANT SOLUTION: AMD MxGPU

AMD Multi-User GPU (MxGPU) technology is designed to deliver predictable and consistent performance irrespective of noisy and untrusted neighbors. Hardware-enforced time-slicing and memory isolation divide and isolate resources to prevent mingling. AMD FirePro S7100-series graphics accelerators deliver world-class QoS.

## HOW WELL DOES IT WORK?

AMD performed lab tests to compare the QoS delivered by the AMD FirePro S7150 x2 versus the NVIDIA Tesla M60. SPECcap® for Solidworks™—a real-world benchmark—was instantiated simultaneously across 16 virtual machines in a vSphere 6.5 environment<sup>1</sup>. Figure 1 (next page) displays the resulting normalized graphics composite scores and demonstrates the consistency and predictability of AMD MxGPU technology.

Radeon Pro graphics include professional graphics products such as Radeon Pro WX Series, Radeon Pro SSG, Radeon Pro Duo as well as legacy product solutions such as AMD FirePro, ATI FirePro, ATI FireMV, ATI FireGL and ATIFireStream. Learn more at: <http://www.amd.com/en-gb/products/graphics/workstation/radeon-pro>.

## Footnotes

### 1. Server configured as:

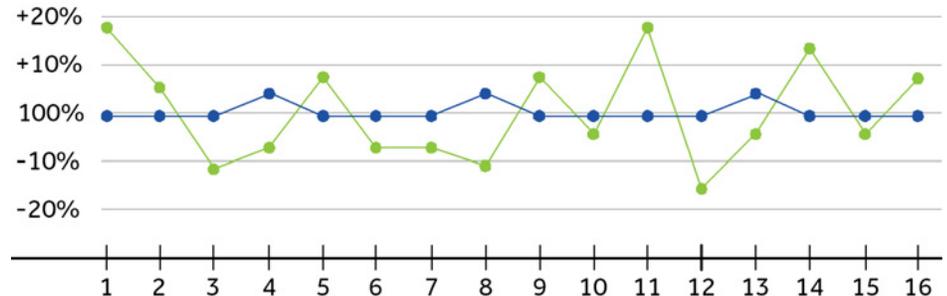
- **Host:** VMware ESXi 6.5.0 with host driver Radeon Pro VMware vSphere Driver– Revision 1.02/NVIDIA VMware ESXi 6.5 Host Driver 367.106
- **Virtual Machines:** Intel® Xeon® E5-2698 v4 @2.20 GHz, 4 Cores, 4 Logical Processors, 8GB Physical Memory, 120 GB Storage, Windows 7 Enterprise 64-bit SP1 Build 7601 with Radeon Pro Software 17.30 driver/NVIDIA GRID 370.12 guest graphics driver

### 2. Testing conducted by AMD Performance Labs as of August 1st under following host and virtual machine configuration:

**Host:** VMware ESXi 6.5.0 with host driver Radeon Pro VMware vSphere Driver– Revision 1.02/NVIDIA VMware ESXi 6.5 Host Driver 367.106. **Virtual Machines:** Intel® Xeon® E5-2698 v4 @2.20 GHz, 4 Cores, 4 Logical Processors, 8GB Physical Memory, 120 GB Storage, Windows 7 Enterprise 64-bit SP1 Build 7601 with Radeon Pro Software 17.30 driver/NVIDIA GRID 370.12 guest graphics driver. **Benchmark:** SPECapc Solidworks simultaneously on 16 virtual machines NVIDIA GRID vGPU configuration: Each VM has 1 vGPU assigned from NVIDIA GRID M60-0Q profile (16 vGPU on a physical GPU). AMD MxGPU configuration: Each VM has 1 vGPU assigned from 16 VFs configuration (16 vGPU on a physical GPU). **Raw, normalized, AMD scores across 16 VMs:** 0.994174757 0.994174757 0.994174757 1.025242718 0.994174757 0.994174757 0.994174757 1.025242718 0.994174757 0.994174757 0.994174757 1.025242718 0.994174757 0.994174757 0.994174757

**Raw, normalized, NVIDIA scores across 16 VMs:** 1.181981982 1.066666667 0.893693694 0.922522523 1.066666667 0.922522523 0.922522523 0.893693694 1.066666667 0.951351351 1.181981982 0.836036036 0.951351351 1.124324324 0.951351351 1.066666667. **FirePro™ S7150x2 graphics standard deviation:** 0.00390. **NVIDIA Tesla M60 graphics standard deviation:** 0.03618 **Performance Differential:** 0.03618/0.00390 = ~9.27x higher standard deviation on NVIDIA Tesla M60. **FirePro™ S7150x2 graphics performance predictability (inverse of standard deviation):** 1/0.00390 = 256.41. **NVIDIA Tesla M60 graphics performance predictability (inverse of standard deviation):** 1/0.03618 = 27.64. PC and server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers. RPW-177

To further quantify QoS, we used the reciprocal of the statistical standard deviation of the 16 virtual machine graphics scores to measure predictability. This approach reveals that the AMD S7150 x2 MxGPU delivers performance that is nine times more stable<sup>2</sup> than the NVIDIA Tesla M60, as shown in Figure 2, below



**Figure 1: Performance stability across 16 virtual machines for AMD MxGPU (blue) and NVIDIA Tesla (green). Values normalized to make the mean 100%. Flatter curve nearer to 100% is better.**



**Figure 2: Performance predictability<sup>2</sup> (higher is better). Numbers are derived from a reciprocal of the raw performance numbers used in the standard deviation calculation shown in Figure 1.**

The performance variability on the AMD FirePro S7150 x2 with 16 simultaneous workloads is in line with expected observations on 16 identical, isolated physical machines. AMD MxGPU technology replaces noisy and/or untrusted neighbors with good neighbors.

## PERIPHERAL BENEFITS FOR YOUR BUSINESS

The stable and predictable QoS delivered by AMD MxGPU technology offers the following compelling benefits to your business:

- **Fewer IT Support Calls:** Effective use of virtualized infrastructure should reduce IT support and not clog the lines with performance-related complaints—and especially not complaints about performance varying for no apparent reason.
- **Better Business Planning:** Your enterprise's Program Management Office likely isn't fond of committing to projects while there is a significant risk computational resources can fail to perform as expected at any given moment of any given day.



# RADEON™ PRO

For more information, please visit <http://www.amd.com/mxgpu>

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of non-infringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

© 2017 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, FirePro and combinations thereof are trademarks of Advanced Micro Devices, Inc. Linux is a registered trademark of Linus Torvalds. OpenCL is a trademark of Apple Inc. used by permission by Khronos. OpenGL is a registered trademark of Khronos Group, Inc. PCIe is a registered trademark of PCI-SIG Corporation. Microsoft, DirectX and Windows are registered trademarks of Microsoft Corporation in the U.S. and/or other jurisdictions. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.