

CUTTING-EDGE PERFORMANCE WITH SCALABILITY AND FLEXIBILITY TO MEET SHIFTING NEEDS

VDI powered by AMD Multiuser GPU (MxGPU) technology delivers the consistent performance required by fluctuating project teams while optimizing resource usage.

The media and entertainment industry is on the forefront of delivering innovative, visually compelling still images and video for a variety of purposes from print production to game and interface design, photography, CGI, movies, and much more. These place a premium on performance to allow faster iteration and shorter rendering/processing cycles among teams of internal and/or third-party users that often vary by project and include members in varying locations. The need to deliver high-end GPU acceleration to geographically diverse, project-specific teams poses unique challenges for deploying and managing IT resources.

AMD Multiuser GPU (MxGPU) technology offers the following key benefits to the media and entertainment industry:

- 1. Mobility & Collaboration:** Users can access their virtual desktops and allotted GPU resources on virtually any device and have full OpenCL™ and GPU acceleration from virtually any location. Migrating data to the data-center simplifies version control, helps boost productivity, and fosters collaboration among distributed teams.
- 2. Scalability and Flexibility:** Adding and removing users is as easy as creating or removing accounts and allocating resources. Users needing access to different applications and/or GPU performance when shifting projects can receive the resources they need within minutes.
- 3. Performance:** Hardware-based virtualization brings workstation-grade 2D/3D graphics acceleration to the datacenter using the Single Root I/O Virtualization (SR-IOV) PCIe® virtualization standard. This eliminates proprietary and complex software from the hypervisor and allows each VM to use native AMD drivers with 100% compatibility and access to all GPU graphics and compute functions on the server. Each GPU can support 1 to 16 users and requires no profiles. The consistent performance ensured by hardware-based virtualization also simplifies resource allocation and planning.
- 4. Cost Effectiveness:** IT departments no longer need to procure, support, repair and upgrade individual workstations with multiple hardware, OS, and application configurations, especially among ever-shifting project teams.

Users simply log in to their virtual desktops at virtually any time from virtually any device and receive full workstation performance, GPU acceleration, and application/data access based on their credentials and assigned resources. Persistent desktops can be used for internal parties with consistent needs. Third parties or users with widely varying needs can access pooled desktops.



Empowers Mobility and Collaboration

Users have full access to applications and data at virtually any time, from virtually any location, on virtually any device, including thin and zero clients. Each user accesses only the applications and data they need, enhancing security and optimizing resource utilization.

Users transmit commands to the virtual machines and receive fully rendered pixels at full resolution and with full graphics performance. Centralized data storage and processing reduces the need for time-consuming file transfers and version control among multiple devices.



Allows Scalable, Flexible Resource Usage

Adding a new user is as fast and easy as assigning a virtual desktop with appropriate application and access privileges. The user is productive in minutes on any device, with no need to purchase and maintain an expensive workstation. Persistent desktops allow users to pick up where they left off at any time, while pooled desktops maximize user density by sharing resources, such as among different shifts.

The hardware virtualization implemented in AMD MxGPU enables workstation-grade 2D/3D graphics performance and provides fast, accurate resource monitoring and metrics to facilitate planning to meet future needs.



Performance at the Speed of Creativity

AMD Multiuser GPU technology enables consistent AMD FirePro GPU acceleration that can meet or exceed traditional workstation graphics cards, allowing modelers, designers, animators, and others to run everything from standard office

AMD RADEON PRO

Pure Datacenter Graphics

applications to streaming video and high-end 3D applications with complete fidelity. Hardware-based virtualization assures performance through dedicated frame buffers while helping ensure total compatibility by using native drivers. This helps shorten iterations between drafts and speed final renderings.

Lowers TCO and Maximizes ROI

AMD Multiuser GPU technology can help lower costs across the board compared to traditional workstations or laptops. Initial deployment is fast and easy, and future upgrades are as easy as purchasing additional servers and access points. Firms can assign resources to internal and third-party users within minutes, thereby optimizing resource usage. Further, there are no additional hardware costs to use AMD Multiuser GPU technology beyond purchasing the hardware itself.

Peace of Mind

- Three-year limited product repair/replacement warranty
- Direct toll-free phone (US, Canada) and global email access to dedicated technical support team
- Advanced parts replacement option

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

Specifications

S7100X



S7150X2



V340



Max. Virtual Machines	16	16	32
Max. Power	100W	265W	<300W
Form Factor	PCIe® 3 MXM 3.1	Full height & length PCIe® 3x16	Full height & length PCIe® 3x16
Cooling	Passive	Passive	Passive
Interface	256 bit	256 bit	256 bit
Memory	8GB GDDR5	16GB GDDR5	32GB HBM2
ECC Memory	supported	supported	supported
API Support	DirectX® 11.1, OpenGL® 4.4 and OpenCL™ 2.0		
OS Support	Microsoft® Windows 10, Windows® 7, Windows® Server 2016, Windows® Server 2008 R2 (64-bit only)		
Hypervisor Supt.	VMware® ESXi™ 6.5, 6.0, Citrix® XenServer® 7.4+		
Remote Vis. Supt.	VMware® Horizon® View 7.0+, Citrix® XenDesktop® 7.15+, Citrix® XenApp® 7.15+		

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.

© 2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Linux is a registered trademark of Linux Torvalds. OpenCL is a trademark of Apple Inc. used by permission by Khronos. 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.

