

AMD RADEON PRO

Pure Datacenter Graphics

SOLUTION BRIEF | AMD MULTIUSER GPU

MANUFACTURING

Cover image provided courtesy of Dassault Systèmes.



SECURE, SCALABLE COLLABORATION WITH PREDICTABLE GPU ACCELERATION AND LOW TCO

VDI powered by AMD Multiuser GPU (MxGPU) technology impacts IT management overhead and empowers collaboration while helping safeguard sensitive data.

Staying competitive in today's economy means shorter product innovation time and development cycles and working with remote employees and external parties around the clock. It means finding new ways to do more with less, including replacing manual tasks with powerful 3D applications and the powerful computers needed to run them. It also means ensuring that authorized users can access current data while also protecting that data. It can even mean supporting a variety of hardware and operating systems.

AMD Multiuser GPU (MxGPU) technology offers the following key benefits to manufacturers:

- 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 datacenter simplifies version control, helps boost productivity, and can foster collaboration among local and remote teams.
- 2. Security:** Centralized data storage facilitates backup and other protective measures while eliminating the possibly catastrophic losses caused by loss or theft of that data.
- 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 while also removing potentially vulnerable abstraction layers. Each VM uses 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.
- 4. Simplicity:** IT departments no longer need to support individual workstations with multiple hardware, OS, and application configurations. Users simply log in to their virtual desktops at virtually any time from virtually any device and receive full workstation performance and GPU acceleration.



Empowers Mobility and Collaboration

Replacing an individual workstation with an access portal means that users have full access to application and data at virtually any time, from virtually any location, on virtually any device, including thin and zero clients. Giving every user the same OS and application environment helps ensure compatibility.

Users transmit commands to the virtual machines and receive fully rendered pixels at full resolution and with full graphics performance in return. Storing and processing data in the datacenter reduces the time-consuming need to transfer and track multiple copies and versions across multiple devices.



Safeguards Sensitive Data

Storing and processing very sensitive CAD and PLM data in the datacenter reduces the need for local copies while improving version control and helping ensure proper backup and archival. It also protects data against situations such as losing a laptop loaded with sensitive data, data loss caused by a virus or hardware failure on a local workstation, and unauthorized access or theft. Moving applications and compute to the datacenter also helps manage software licensing and eliminate piracy or installing unauthorized plug-ins that can jeopardize compatibility among multiple users.

End users access the VDI environment using their individual network credentials. The IT department always retains control over who can access which applications and data.



Flexible, Scalable Performance Supports Future Growth

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.

AMD RADEON PRO

Pure Datacenter Graphics

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.

IT Management Efficiency

MxGPU-based VDI saves time and money by freeing IT from continually procuring, maintaining, and repairing individual workstations. System maintenance and upgrades takes place in the datacenter, with all affected users seeing near-immediate results. Hardware-based GPU virtualization also means that IT can assign the exact resources needed by any user with no need to install special drivers or management software in the hypervisor or virtual machines. Instead, each machine uses native workstation AMD drivers. Nearly every application that runs on a workstation will run in an MxGPU VDI deployment.

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 Linus 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.

