



GPU Compute Virtualization with VirtIO

Huang Rui <ray.huang@amd.com>

Honglei Huang <honglei1.huang@amd.com>



AMD  together we advance_


Team members

- Ray Huang - China 
 - Kernel (GPU, Xen)/QEMU/Xen
- Julia Zhang - China 
 - Mesa 3D (OpenGL, Vulkan)/Virglrenderer/QEMU
- Honglei Huang - China 
 - ROCm/Thunk/Virglrenderer/QEMU
- Jiqian Chen - China 
 - Kernel (GPU, Xen)/QEMU/Xen
- Penny Zheng (*new*) - China 
 - Xen/QEMU/Kernel (GPU, Xen)
- Trigger Huang (*new*) - China 
 - Mesa 3D (OpenGL, Vulkan)/Kernel (GPU)/Virglrenderer
- Pierre-Eric Pelloux-Prayer - France 
 - Mesa 3D (OpenGL, Vulkan)/Virglrenderer/QEMU
- Xenia Ragiadakou (Stefano's team) - Greece 
 - Xen/QEMU/Kernel (Xen, GPU)
- Leo Liu - Canada 
 - Mesa Multimedia
- Boyuan Zhang - Canada 
 - Mesa Multimedia
- Lingshan Zhu (*new*) - China 
 - Kernel (VirtIO, AMDKFD)/QEMU/ROCm
- Wei Zhao (*new*) - China 
 - Mesa 3D (OpenGL, Vulkan)/Virglrenderer
- Dmitry Osipenko - Russia 
 - Mesa 3D/Virglrenderer/QEMU/Kernel VirtIO



Last year

- Xen Project Summit 2023 - June
 - Kickoff GPU Virtualization design for Xen
 - Enable and implement traditional 3D Graphic and Multimedia hardware acceleration based on Virgl and Venus - **first time to enable Vulkan 3D on Xen**
 - Design and implement dGPU passthrough on PVH dom0
- X.Org Developers Conference 2023 - Oct
 - Continue implementing OpenGL and Vulkan 3D libraries support on Xen
 - Continue upstream of dGPU passthrough
 - Introduce the prototype of virtio native context - **OpenGL/Vulkan enhancement**



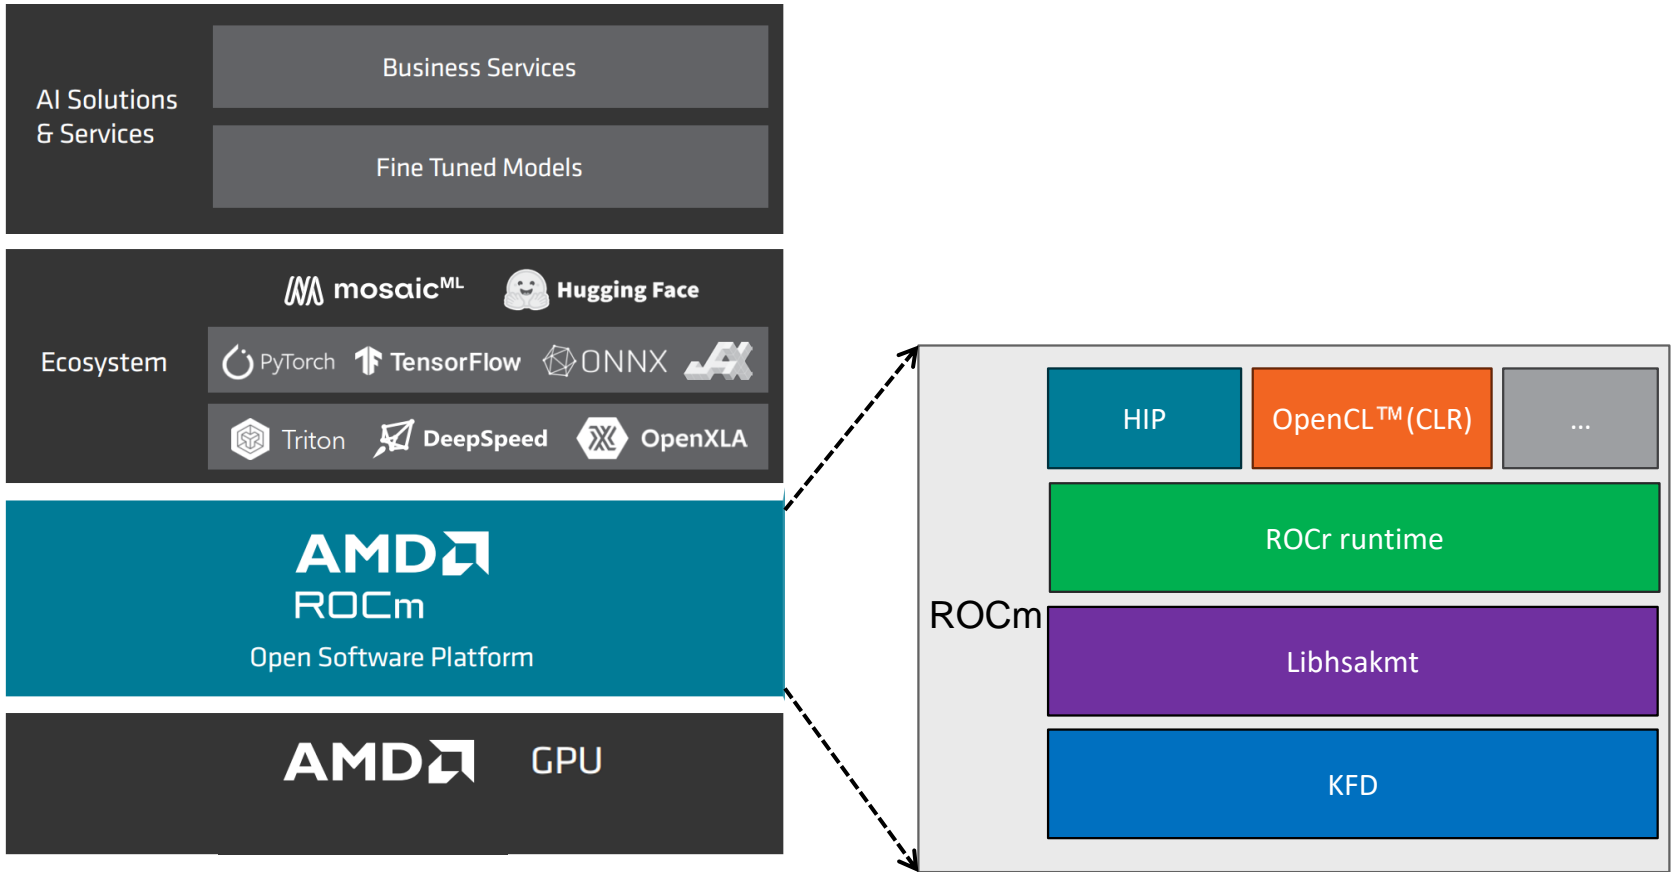
Updates in 2024

- GPU Para-Virtualization
 - Introduce VirtIO native context solution for 3D Graphic and Multimedia hardware acceleration
 - **Prototype AMD ROCm native context solution to support OpenCL™ on VirtIO GPU - coming**
 - Finish PCIe Passthrough on Xen PVH dom0
 - To support VirtIO GPU on PVH guest domU - **ongoing**
 - To support VirtIO GPU on Xen Hyperlaunch - **ongoing**



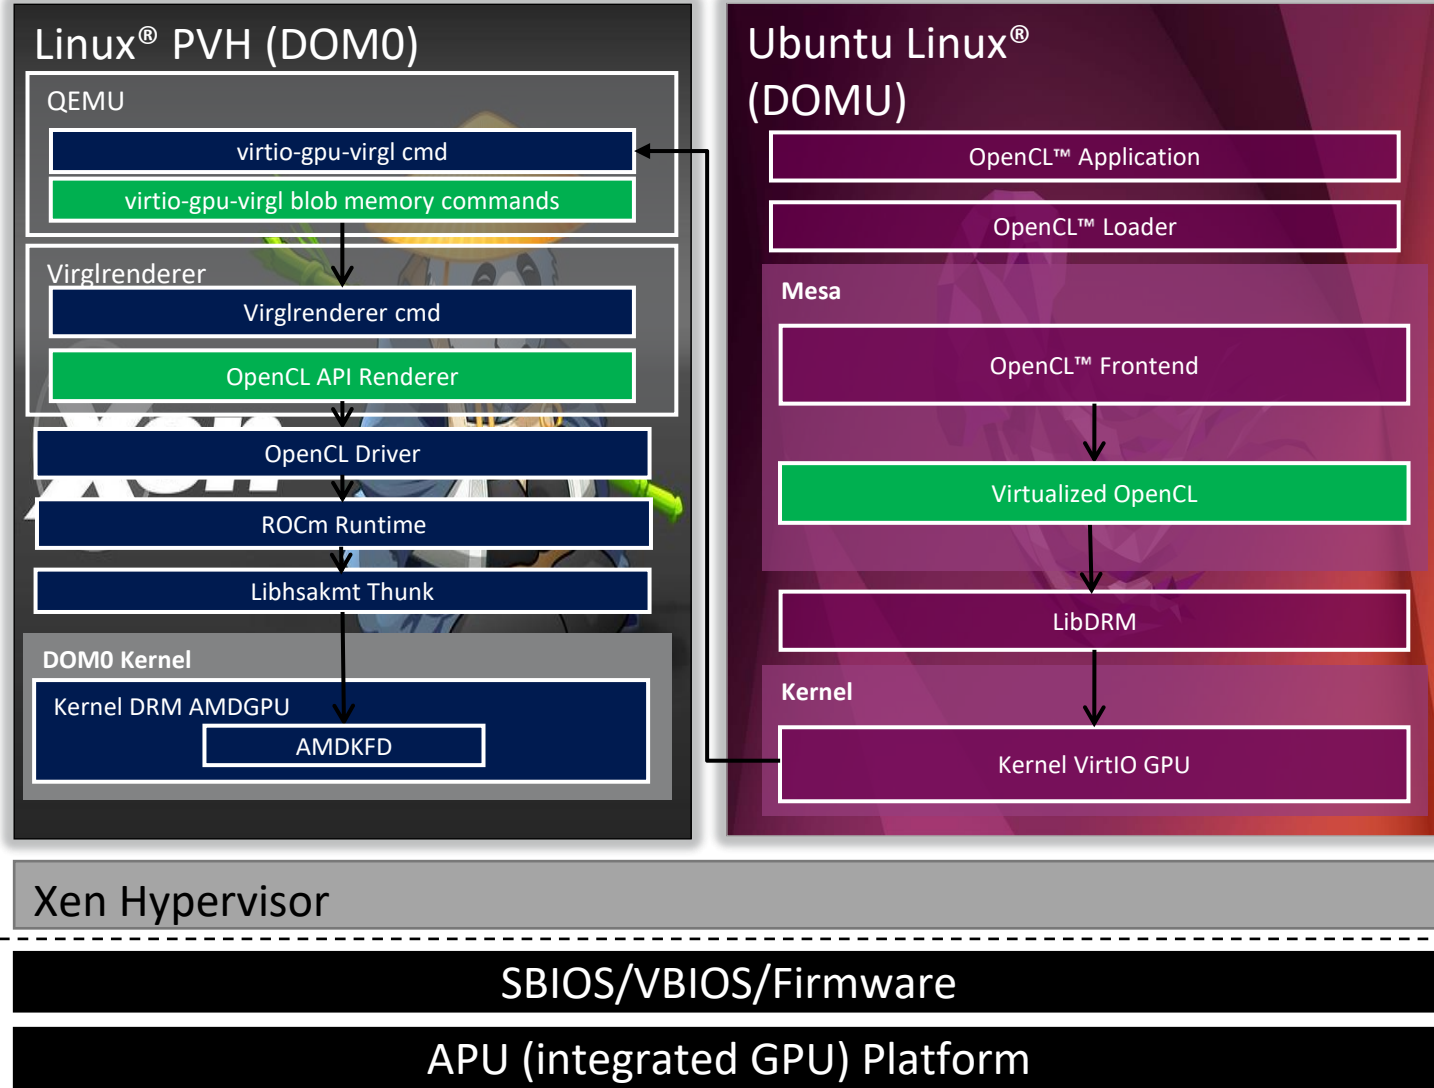
What's ROCm

- An open-source software platform for compute/AI on AMD GPU series
 - Provide multiple interface support for popular computing frameworks like HIP & OpenCL™
 - Mainly includes CLR (OpenCL™ runtime), HIP, ROCr runtime, libhsakmt, KFD kernel driver.
- OpenCL™ over ROCm
 - Support OpenCL protocol based on ROCm compute stack on AMD GPU series



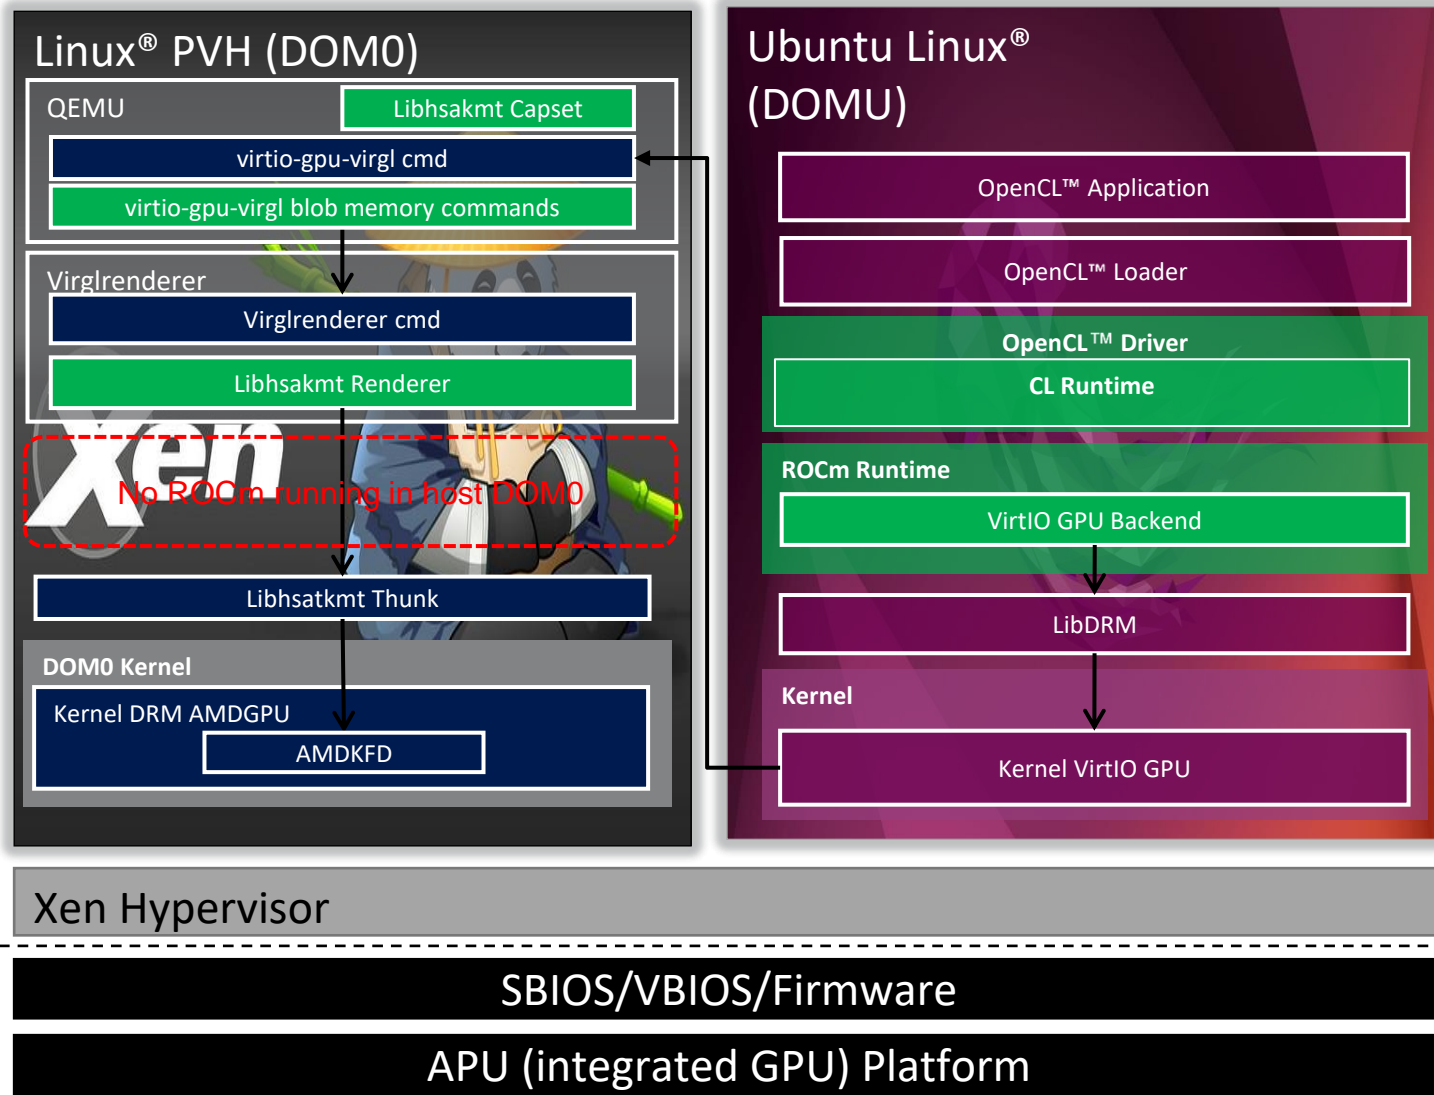
Virtualized OpenCL™ (VirCL)

- AMD ROCm on host DOM0
 - Support OpenCL™ over ROCm on host
- API Forward for OpenCL™
 - Introduce Virtualized OpenCL in Mesa driver to talk with VirtIO GPU
 - Leverage blob memory
 - Add OpenCL™ renderer in virglrenderer to talk with OpenCL™

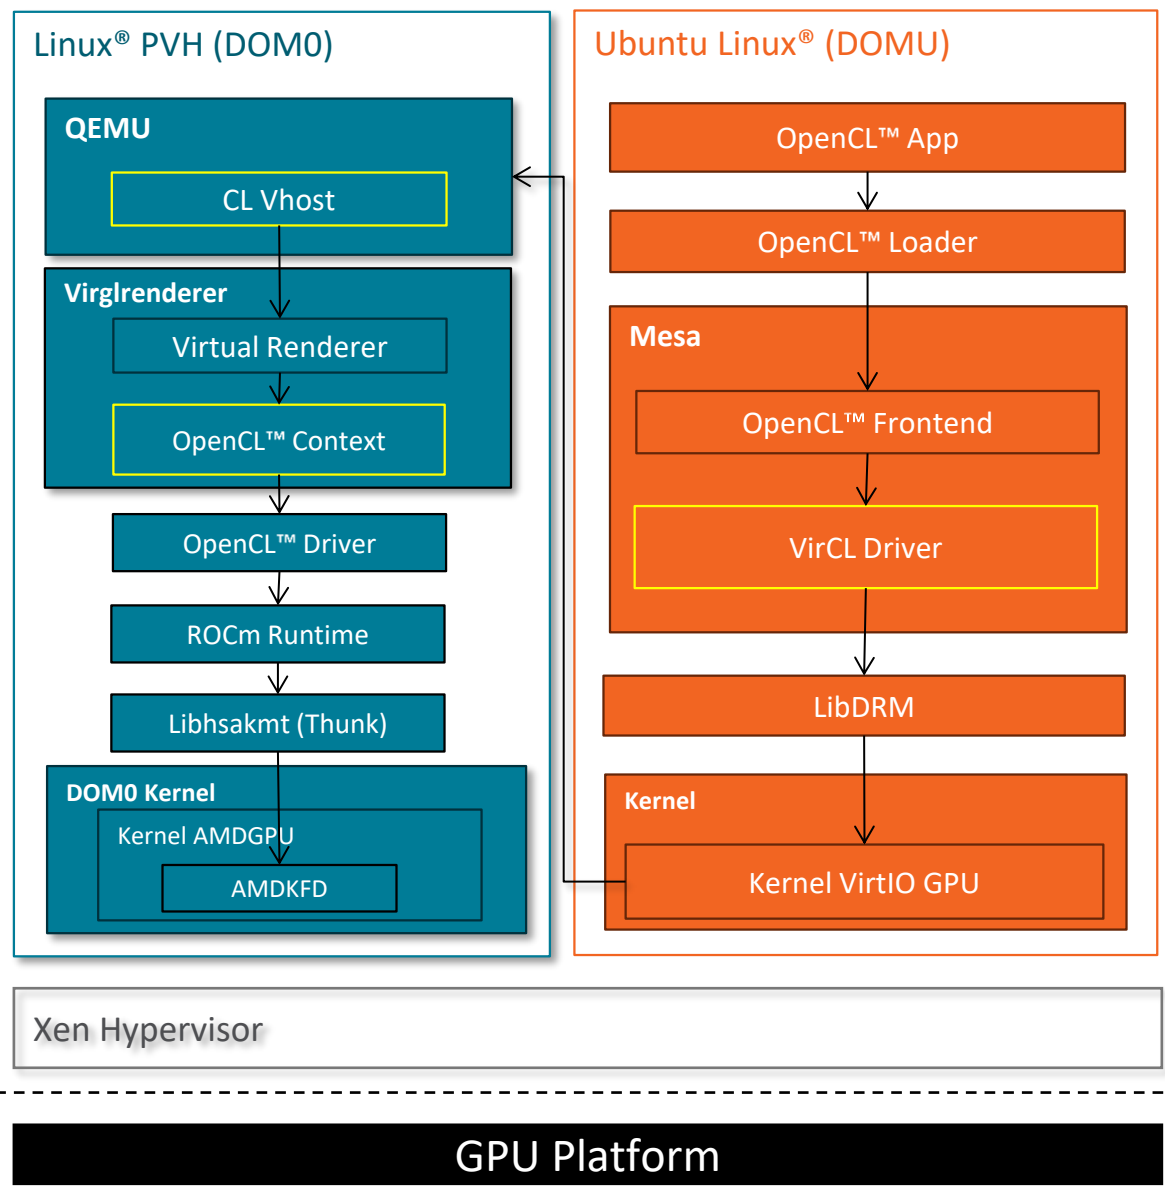


ROCm with VirtIO Native Context

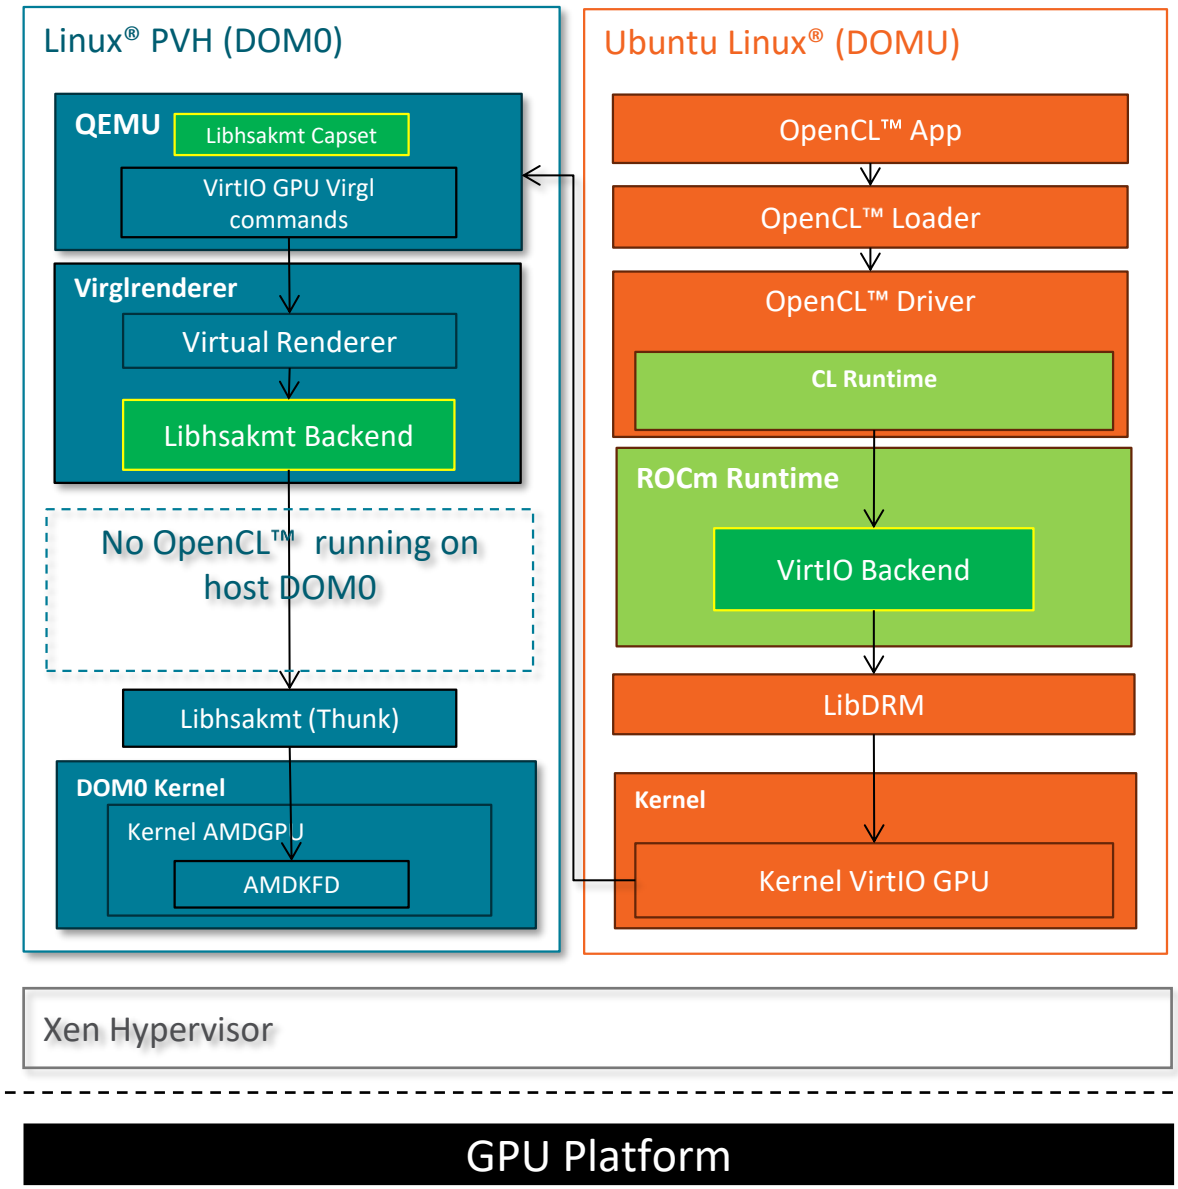
- AMD ROCm on guest DOMU - **coming**
- Support OpenCL™ over ROCm for virtualization
- Also inspired by VirtIO native context on graphic design
- API Forward for Libhsakmt (Thunk)
 - Introduce VirtIO GPU backend in ROCm runtime and OpenCL™ runtime
 - Add libhsakmt capacity in QEMU
 - Leverage blob memory
 - Add libhsakmt renderer in virglrenderer
- Kickoff upstream in Freedesktop
 - https://gitlab.freedesktop.org/virgl/virglrenderer/-/merge_requests/1370



VirCL



CL Native Context



Xen Hypervisor

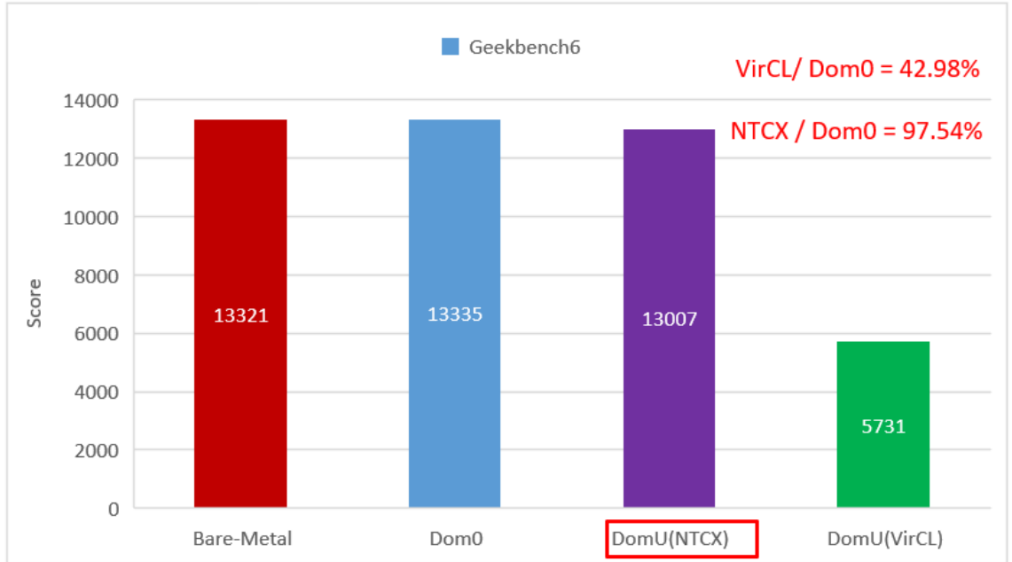
GPU Platform

Xen Hypervisor

GPU Platform

OpenCL™ Performance Preview in Virtualization

- OpenCL™ Comparison between Bare-metal, DOM0, VirCL (DOMU), and Native Context (DOMU)



| Test Cases (Unit: Score) | Bare-Metal - Native | DOM0 - Native | DOMU - VirCL | DOMU - OpenCL Native Context |
|--------------------------|---------------------|---------------|---------------|------------------------------|
| GeekBench6 | 13321 | 13335 | 5731 (42.98%) | 13007 (97.54%) |

The Best is Yet to Come

- Continue upstream for the whole solution - ongoing
- Support multiple processes on AMD KFD for native context
- **Support HIP in ROCm stack for virtualization**
- Support leading AI Frameworks over ROCm for virtualization



AMD 
ROCm

 TensorFlow


 PyTorch

DRIVING THE FUTURE OF IN-VEHICLE EXPERIENCE (IVX)



Passenger needs evolve with every generation, driving vehicle advancement to provide **more capability and functionality.**

70
 BILLION HOURS

Americans collectively spent 70 billion hours each year behind the wheel*

CHANGE IS COMING
 Work environments are more dynamic than ever, placing greater expectations on technology to keep up with the demands of work and life.



DRIVER EXPERIENCE

Drivers face many distractions on the road. Advanced in-vehicle features are crucial for keeping drivers focused.



LARGE HIGH-FIDELITY DISPLAYS
for easy viewing and control





AR NAVIGATION
for accurate wayfinding



SAFETY FUNCTION ALERTS
to improve driver habits and behavior

References



- Hardware
 - AMD Ryzen™ Embedded V2000 Series
 - AMD Radeon™ RX 6000 Series GPUs
 - **AMD Ryzen™ 8000 Series (ongoing)**
- VirtIO GPU and Passthrough GPU Support for Xen
 - Xen Project Summit 2023
 - <https://xen2023.sched.com/event/1LKIn>
- Xen based GPU virtualization - VirtIO/Passthrough
 - X.Org Developer's Conference 2023
 - <https://indico.freedesktop.org/event/4/contributions/216/>
- dGPU prime on VM
 - X.Org Developer's Conference 2023
 - <https://indico.freedesktop.org/event/4/contributions/189/>
- GPU Para-Virtualization on Xen
 - Xen Project Summit 2024
 - <https://xenprojectsummit2024.sched.com/event/1bCFX>

Demo, Q&A, and Thank You





Disclaimer:

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2024 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, Radeon, Ryzen and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. Linux is a trademark of Linus Torvalds and OpenCL is a trademark of Apple Inc. Windows and DirectX are the registered trademarks of Microsoft Corporation in the US and other jurisdictions.

AMD 