Xen based GPU virtualization – dGPU prime

Julia Zhang

Xen Project

AMD

together we advance_

Background of dGPU prime

- Intention
  - OpenGL and Vulkan: dGPU prime allows an integrated GPU (iGPU) to display output data rendered by a discrete GPU(dGPU)

- Background
  - Virgl and Venus
    - Not support DMA
    - Not implement import/export

  - Xen VMs
    - Low performance: sync data depend on memcpy from guest to host via IOV

- Proposal
  - Implement dGPU prime on virgl and venus

dGPU Prime on Xen GPU Virtualization -- overall

- Platform
  - Ubuntu22.04
  - XEN hypervisor + QEMU
- GPU and driver
  - Integrated GPU(iGPU) + Virtio-GPU driver
  - Passthrough discrete GPU(dGPU) + AMDGPU driver

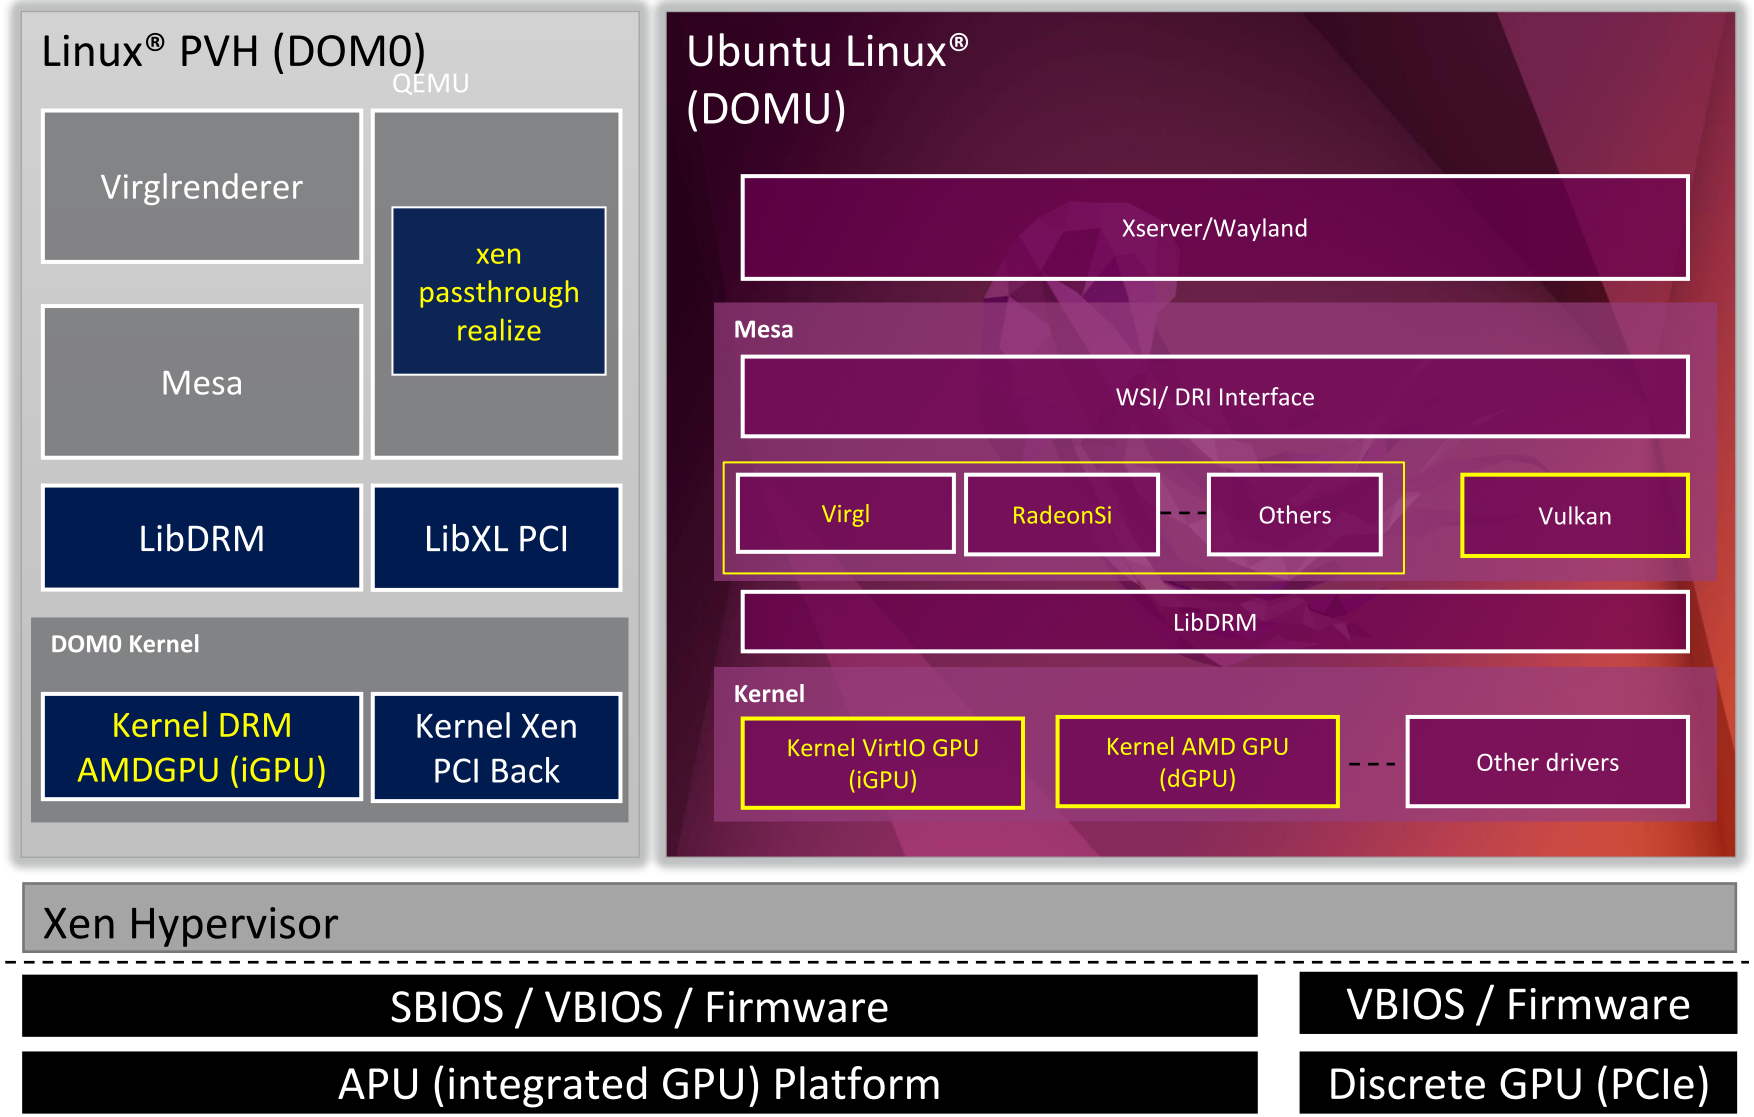

AMD together we advance_

dGPU Prime of Virgl(OpenGL)

- Create display buffer
  - Guest mesa side: send command to create host resource
  - Host virglrenderer side: create texture
- Import display buffer as linear buffer/blit dst
  - Get display buffer handle
  - Create linear buffer from display buffer handle
- Sync rendered data
  - Send command to copy data from guest IOV to host resource

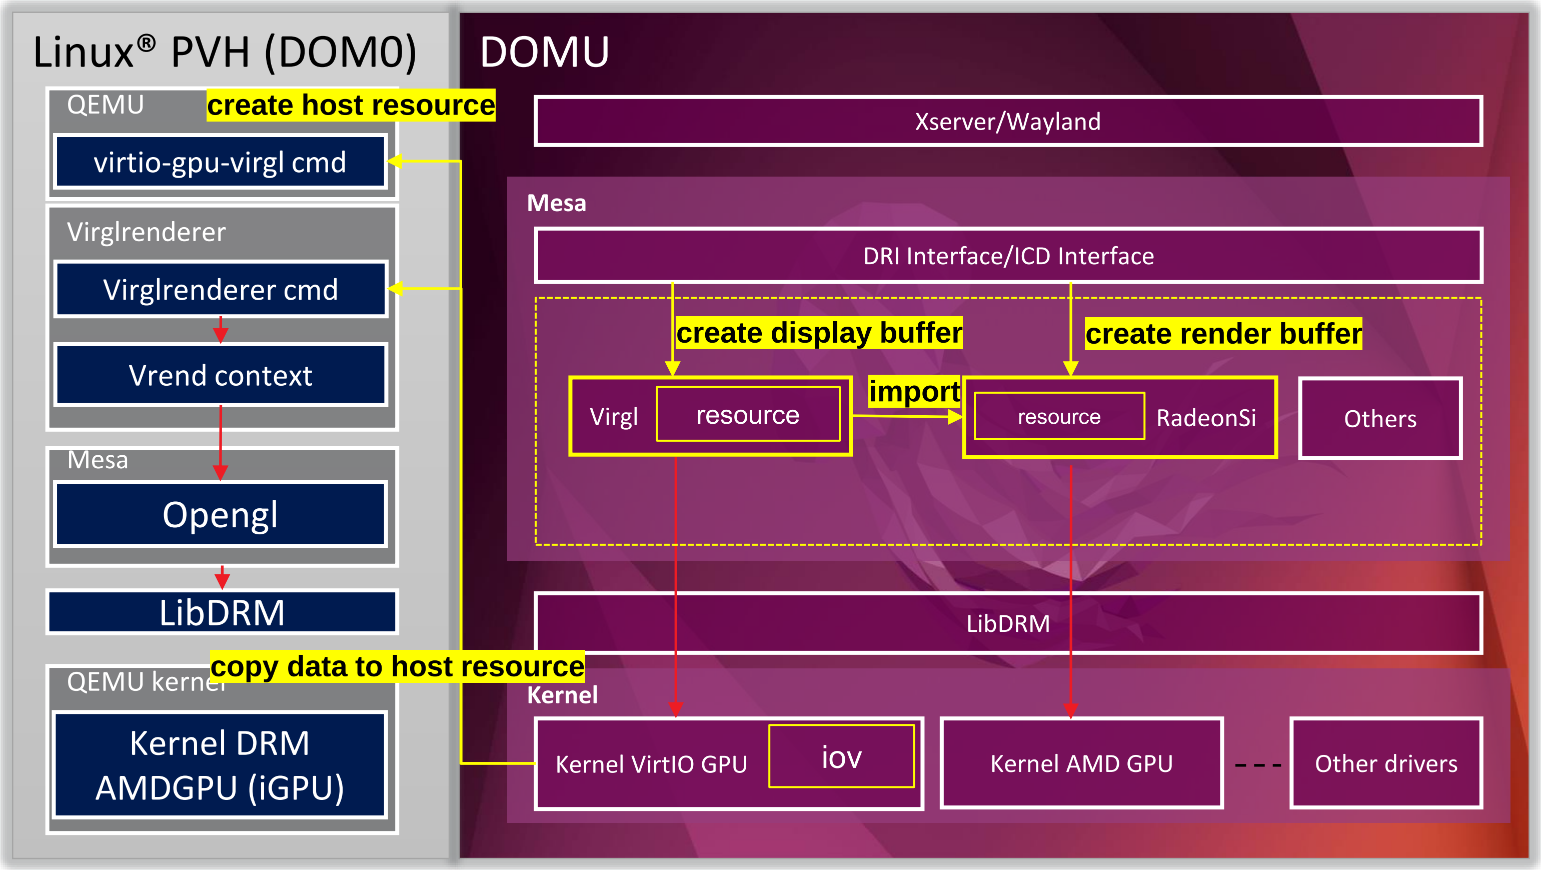

AMD together we advance_

dGPU Prime of Venus(Vulkan)

- Create display buffer
  - Guest mesa side: GBM create display buffer→ send command to create host resource
  - Host virglrenderer side: import host resource as pipe resource → create texture
- Import display buffer as linear buffer(blit dst)
  - Import gbm_bo as foreign mem
  - Bind linear buffer of swapchain with foreign mem
- Sync rendered data
  - Send command to copy data from guest IOV to host resource

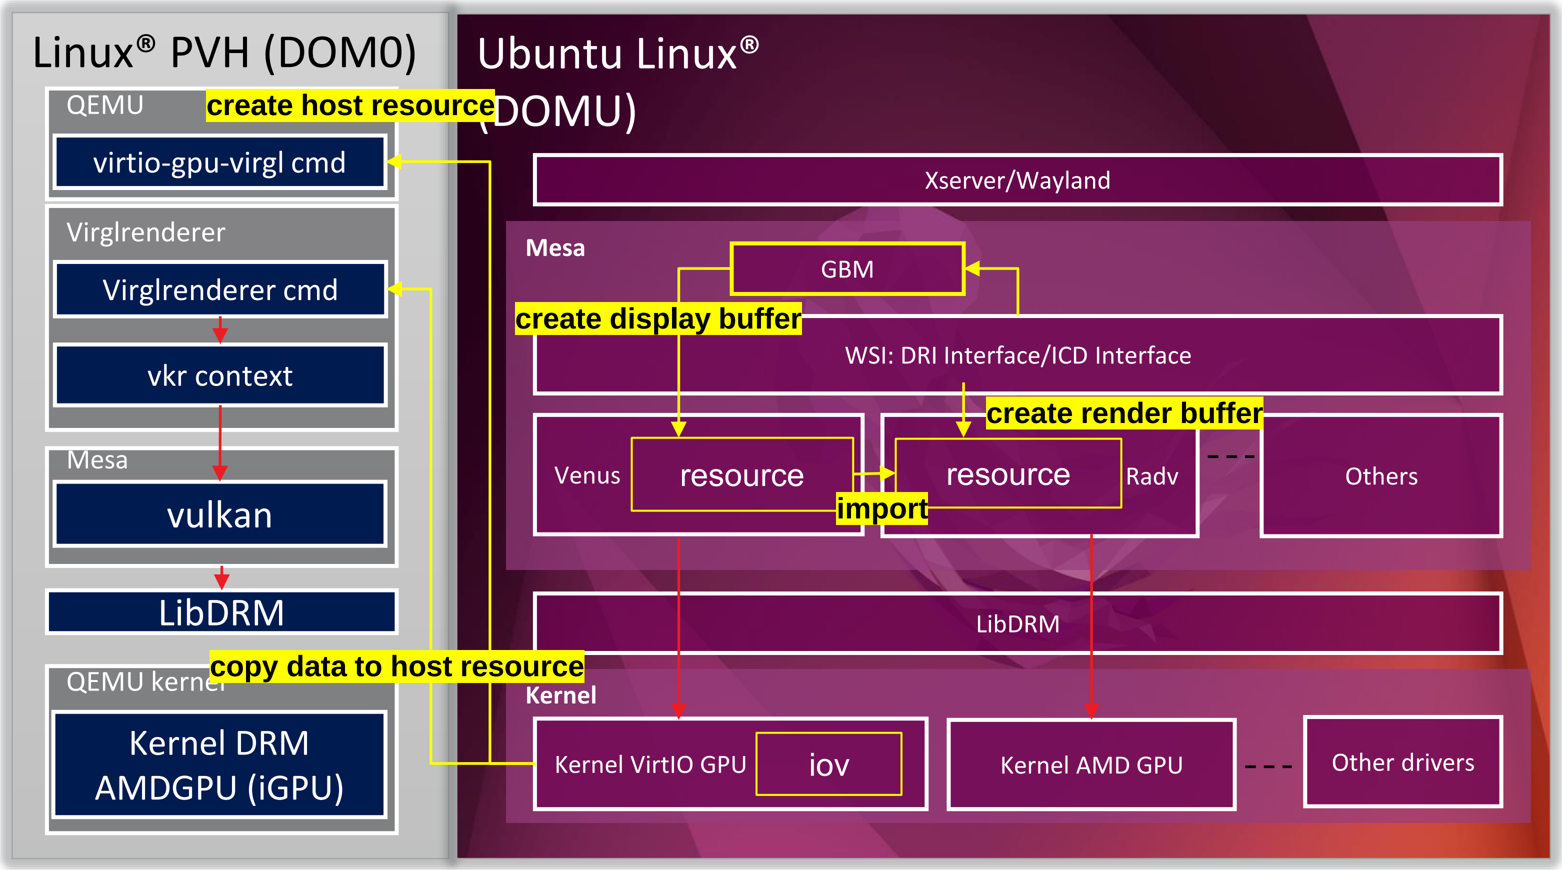

AMD
together we advance_

Improvement of dGPU prime

- Query stride
  - Existing logic: align virgl resource stride to 256
  - Our design: send command to host to query buffer stride using GBM
- Improve performance
  - Existing logic: send command every time swap buffer → Low benchmark scores
  - Our design
    - Create blob resource for linear buffer
    - Map host blob resource to guest → Guest mesa blit data to host pipe resource directly

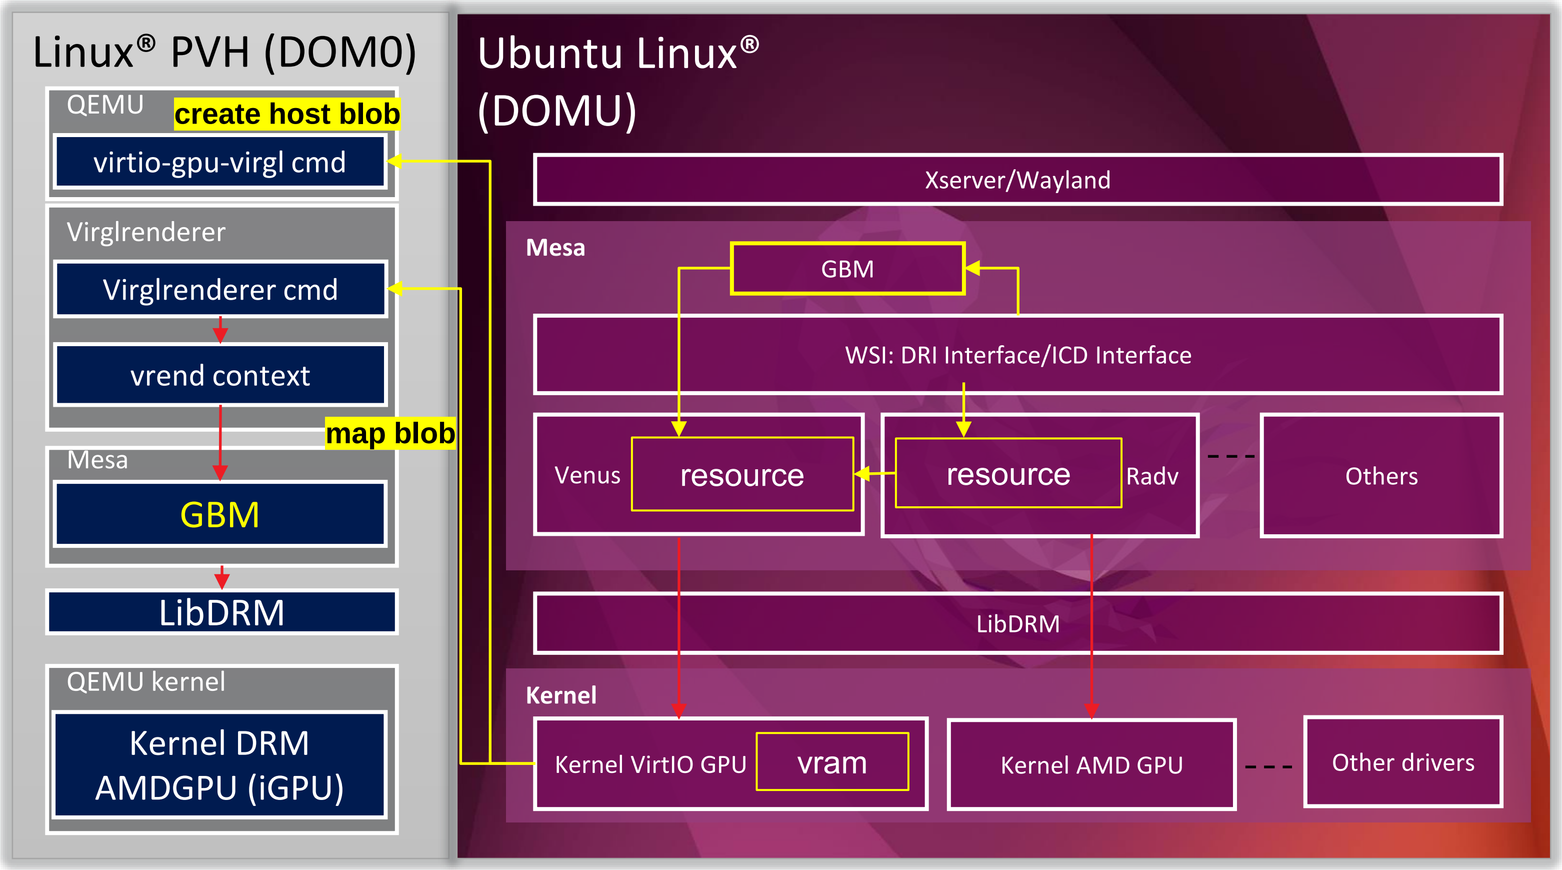

Improvement of dGPU prime

- Native context: virtio-GPU + passthrough GPU

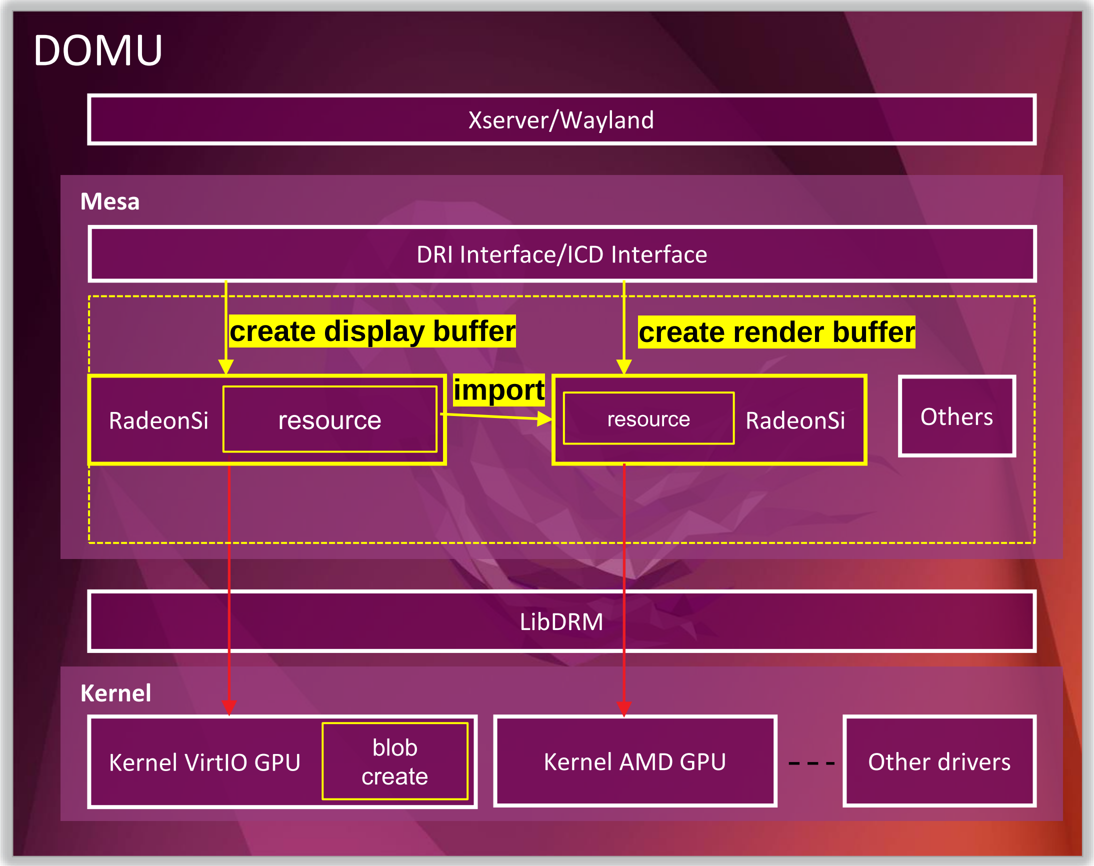

- Two virtio-GPU

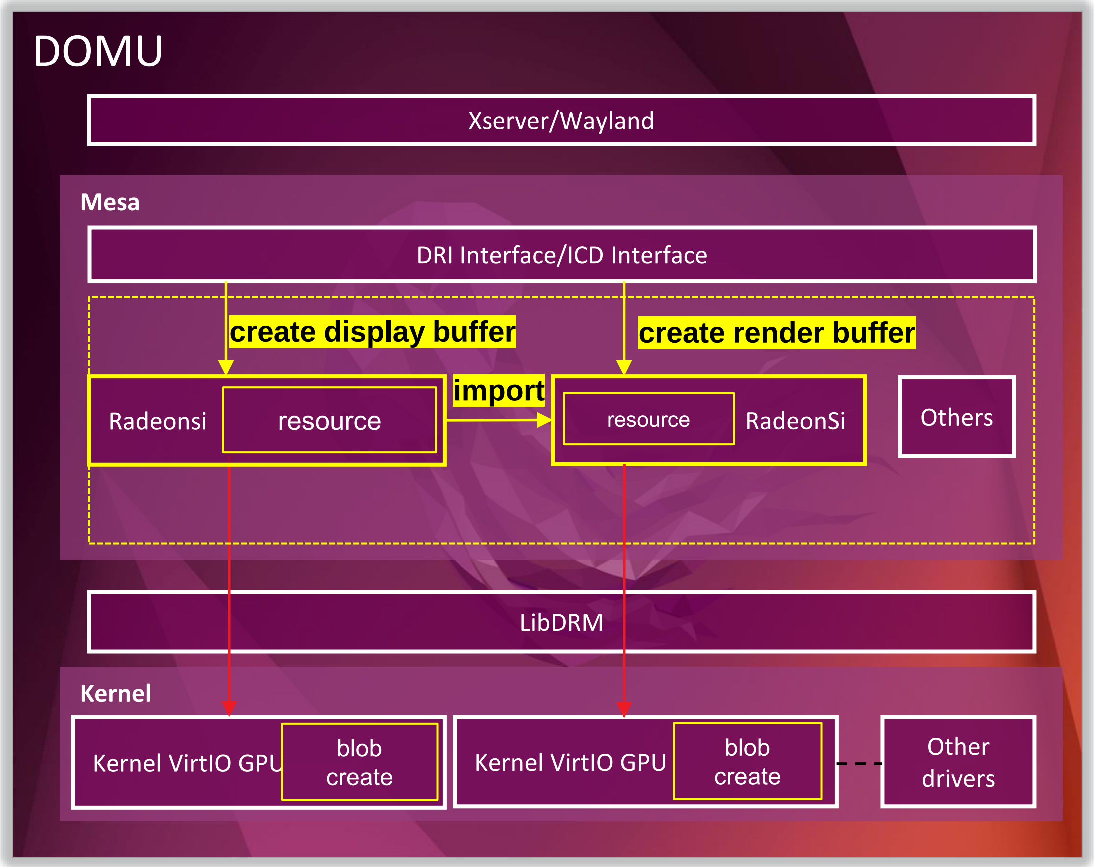

AMD
together we advance_

dGPU prime Comparison between virgl and virgl + blob

Test Cases (Unit: FPS)	Histogram	Bare-Metal - Native (RadeonSi)	DOMU - virgl	DOMU – Virgl + blob
glMark2	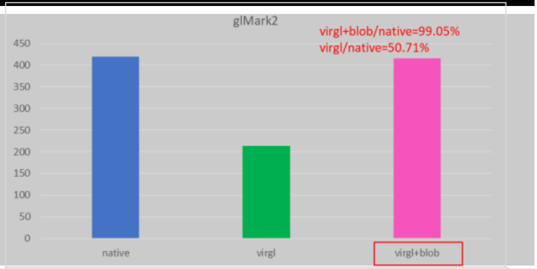 glMark2 virgl+blob/native=99.05% virgl/native=50.71% 	420	213	416 (195.31%)
Madmax	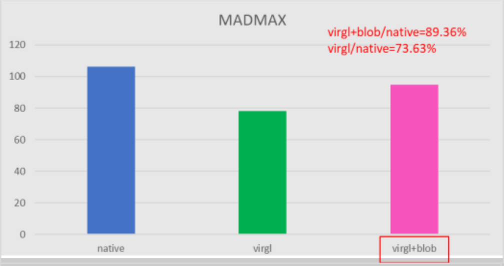 MADMAX virgl+blob/native=89.36% virgl/native=73.63% 	106.2	78.2	94.9 (121.36%)
Unigine Valley	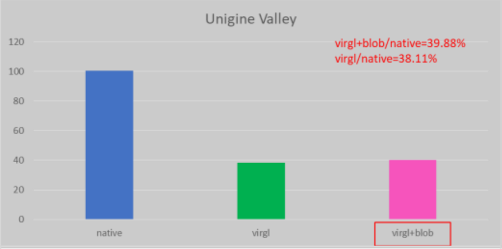 Unigine Valley virgl+blob/native=39.88% virgl/native=38.11% 	100.5	38.3	40.1 (104.70%)
GFXBench (gl_Manhattan)	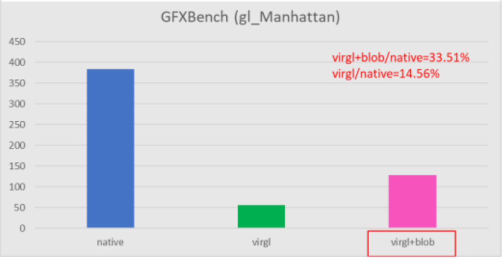 GFXBench (gl_Manhattan) virgl+blob/native=33.51% virgl/native=14.56% 	383.2	55.8	128.4 (230.11%)

dGPU prime Comparison between venus and venus + blob

Test Cases (Unit: FPS)	Histogram	Bare-Metal - Native (Radv)	DOMU - Venus	DOMU – Venus + blob
Rise Of The Tomb Raider	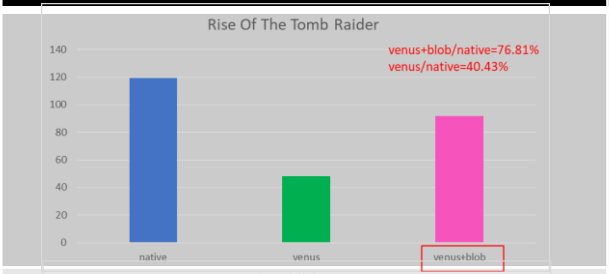 Rise Of The Tomb Raider venus+blob/native=76.81% venus/native=40.43% 	119.18	48.18	91.54 (190.0%)
3DMark Wildlife Extreme	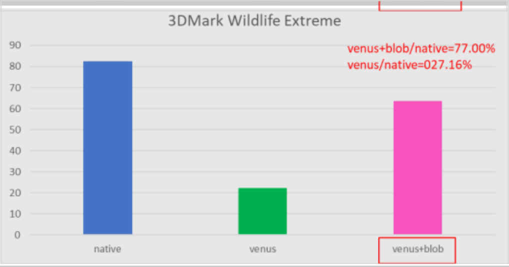 3DMark Wildlife Extreme venus+blob/native=77.00% venus/native=027.16% 	82.47	22.4	63.5 (283.4%)
GFXBench (vulkan_5_high)	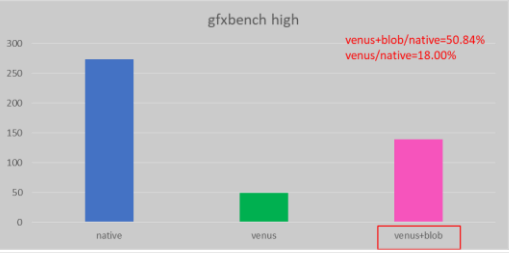 gfxbench high venus+blob/native=50.84% venus/native=18.00% 	273.4	49.2	139 (282.5%)
GFXBench (vulkan_5_normal)	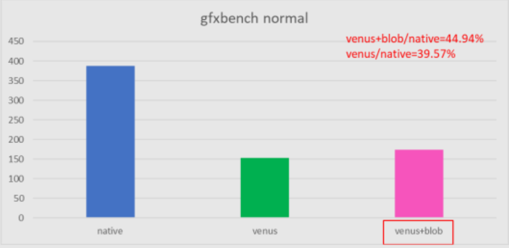 gfxbench normal venus+blob/native=44.94% venus/native=39.57% 	387.63	153.4	174.3 (113.6%)

References

- Repositories
  - Virglrenderer - https://gitlab.freedesktop.org/Julia/virglrenderer/-/commits/upstream-dGPU_prime
  - Mesa - https://gitlab.freedesktop.org/Julia/mesa/-/commits/upstream-dGPU_prime_virgl

- Upstream is in progress
  - Mesa - https://gitlab.freedesktop.org/mesa/mesa/-/merge_requests/23896
  - Virglrenderer - https://gitlab.freedesktop.org/virgl/virglrenderer/-/merge_requests/1268

Thanks

- Bob and Daniel from Collabora (GBM bo create)
- Ray from AMD(Enable blob memory on XEN)
- Pierre Eric from AMD(Native context for amdgpu)

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.

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