Updates of VirtIO GPU support on Xen

Ray Huang

X.Org Developer’s Conference 2025

AMD
together we advance_

[Public]

Team members

Ray Huang
- Kernel (GPU,Xen)/QEMU/Xen

Julia Zhang
- Mesa 3D (OpenGL, Vulkan)/Virglrenderer/QEMU

Honglei Huang
- ROCm/Thunk/Virglrenderer/QEMU

Jiqian Chen
- Kernel (GPU,Xen)/QEMU/Xen

Penny Zheng
- Xen/QEMU/Kernel (GPU,Xen)

Trigger Huang
- Mesa 3D (OpenGL,Vulkan)/Kernel (GPU)/Virglrenderer

Ke Chen
- Mesa 3D (OpenGL,Vulkan)/Virglrenderer/Weston

Lingshan Zhu – Beijing
- Kernel (VirtIO,AMDKFD)/QEMU/ROCm

Wei Zhao
- Mesa 3D (OpenGL,Vulkan)/Virglrenderer

Yiru Ma
- Program Coordination and Management

Shuaidong Chen
- Kernel (GPU,Xen)/QEMU/Xen

Junhua Shen
- ROCm/Thunk/Virglrenderer/QEMU

Yuling Li
- Mesa 3D (OpenGL,Vulkan)/Virglrenderer

Jiming Liu
- Mesa 3D (OpenGL,Vulkan)/Virglrenderer

AMD
together we advance_

[Public]

Background

- Para-Virtualization basic functions are supported
  - VirtIO Native Context for Graphics and Compute ROCm
  - However current basic function is not enough for automotive requirement

- Introduce more advanced features like below
  - GPU Recovery on Xen (DOM0 & DOMU) - *Developed by Trigger/Wei*
  - Perfetto PPS producer on Xen DOMU - *Developed by Julia*

AMD
together we advance_

[Public]

Use Case of GPU Recovery

- GPU Reset
  - Mode1 on dGPU
  - Mode2 on APU
- Current GPU scheduler design
  - All contexts will be killed in the ring while the current job is timeout
- Use case 1
  - Context A is guilty, but context B, C, D will be killed if A, B, C, and D are in the same ring
- Use case 2
  - Context A, B, C are in the same ring, and context A is guilty, but context B, C will be killed. Only context D in another ring can be survived

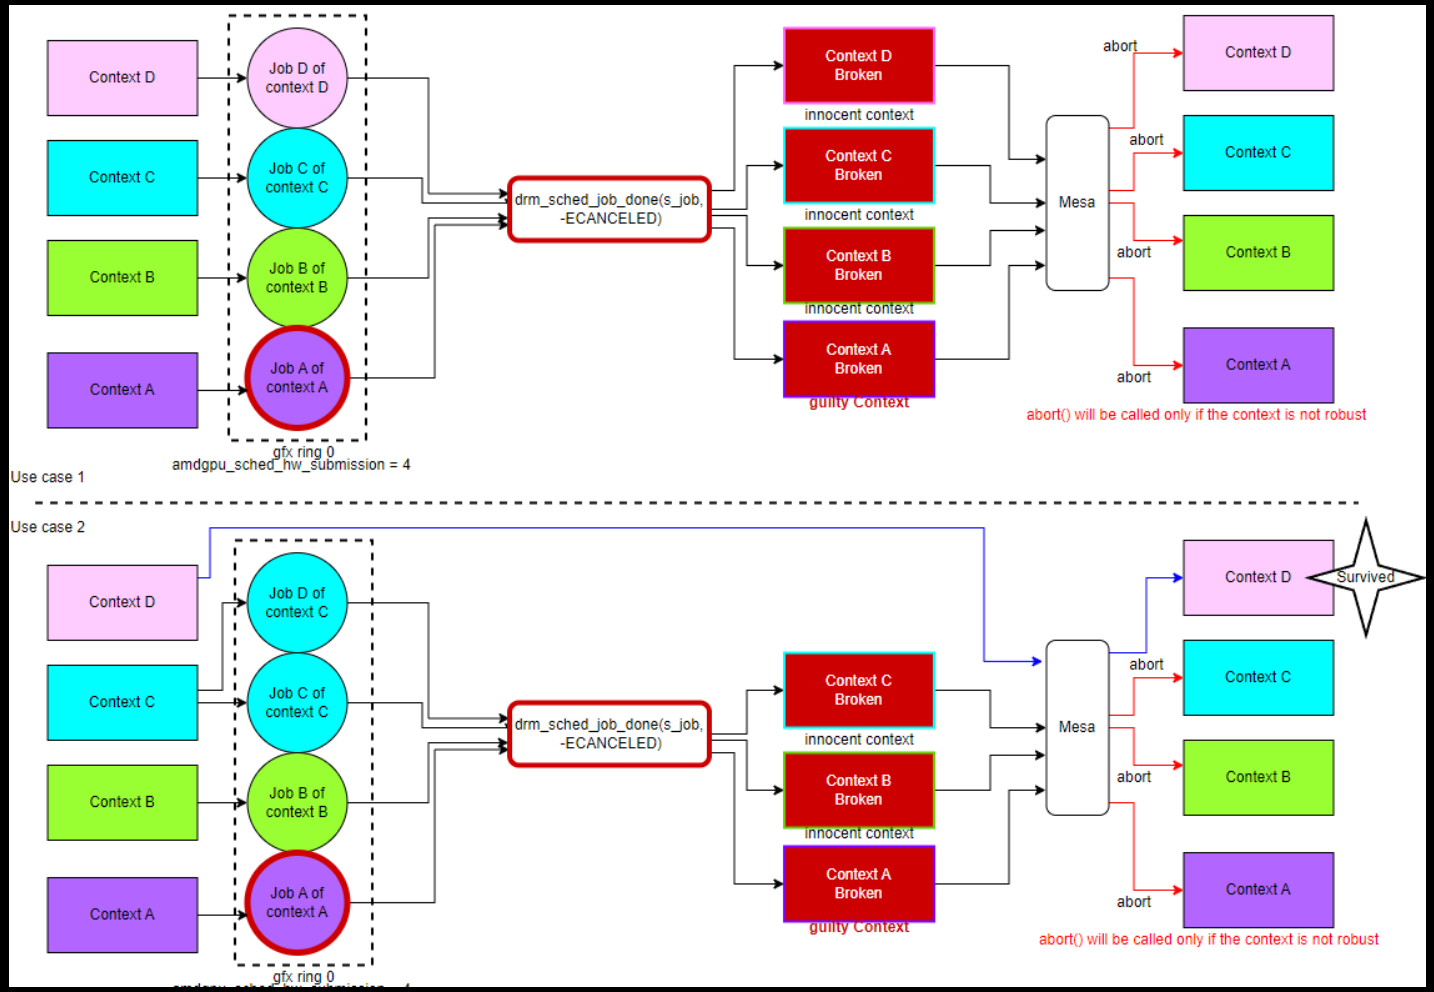

AMD
together we advance_

[Public]

GPU Recovery on Weston compositor

- GPU recovery on User Space
  - Introduce Robust Context Support in weston compositor
  - Make sure weston desktop is always valid during the GPU recovery

- Merge Request
  - https://gitlab.freedesktop.org/wayland/weston/-/merge_requests/1662

- Workflow
  - Create context with GL_EXT_robustness for gl renderer
  - Check whether it is requested by a GPU reset
  - If yes, destroy and re-create GL and EGL contexts and resources
  - Restore the surfaces of client applications

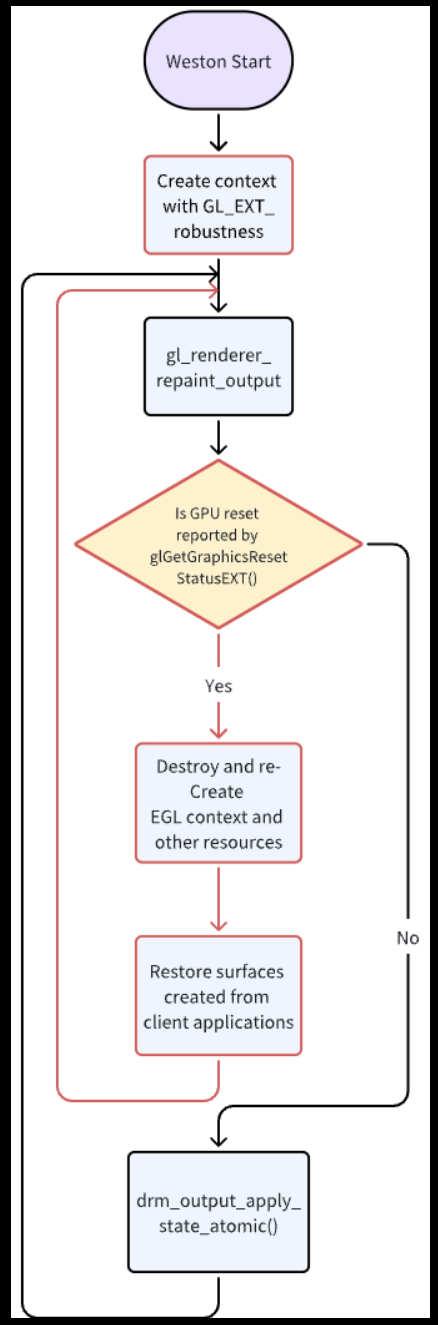

AMD
together we advance_

[Public]

GPU Recovery for VirtIO GPU

- GPU recovery on virtualization
  - QEMU can be treated as an OpenGL application
  - There are three contexts in QEMU and Virglrenderer and create them with robustness
    - sdl2_window_create created by sdl2_display_init()
    - virgl_create_context by vrend_renderer_init()
    - virgl_egl_create_context by vrend_renderer_init()
  - Re-create above three contexts
  - Re-store QEMU application with robust contexts

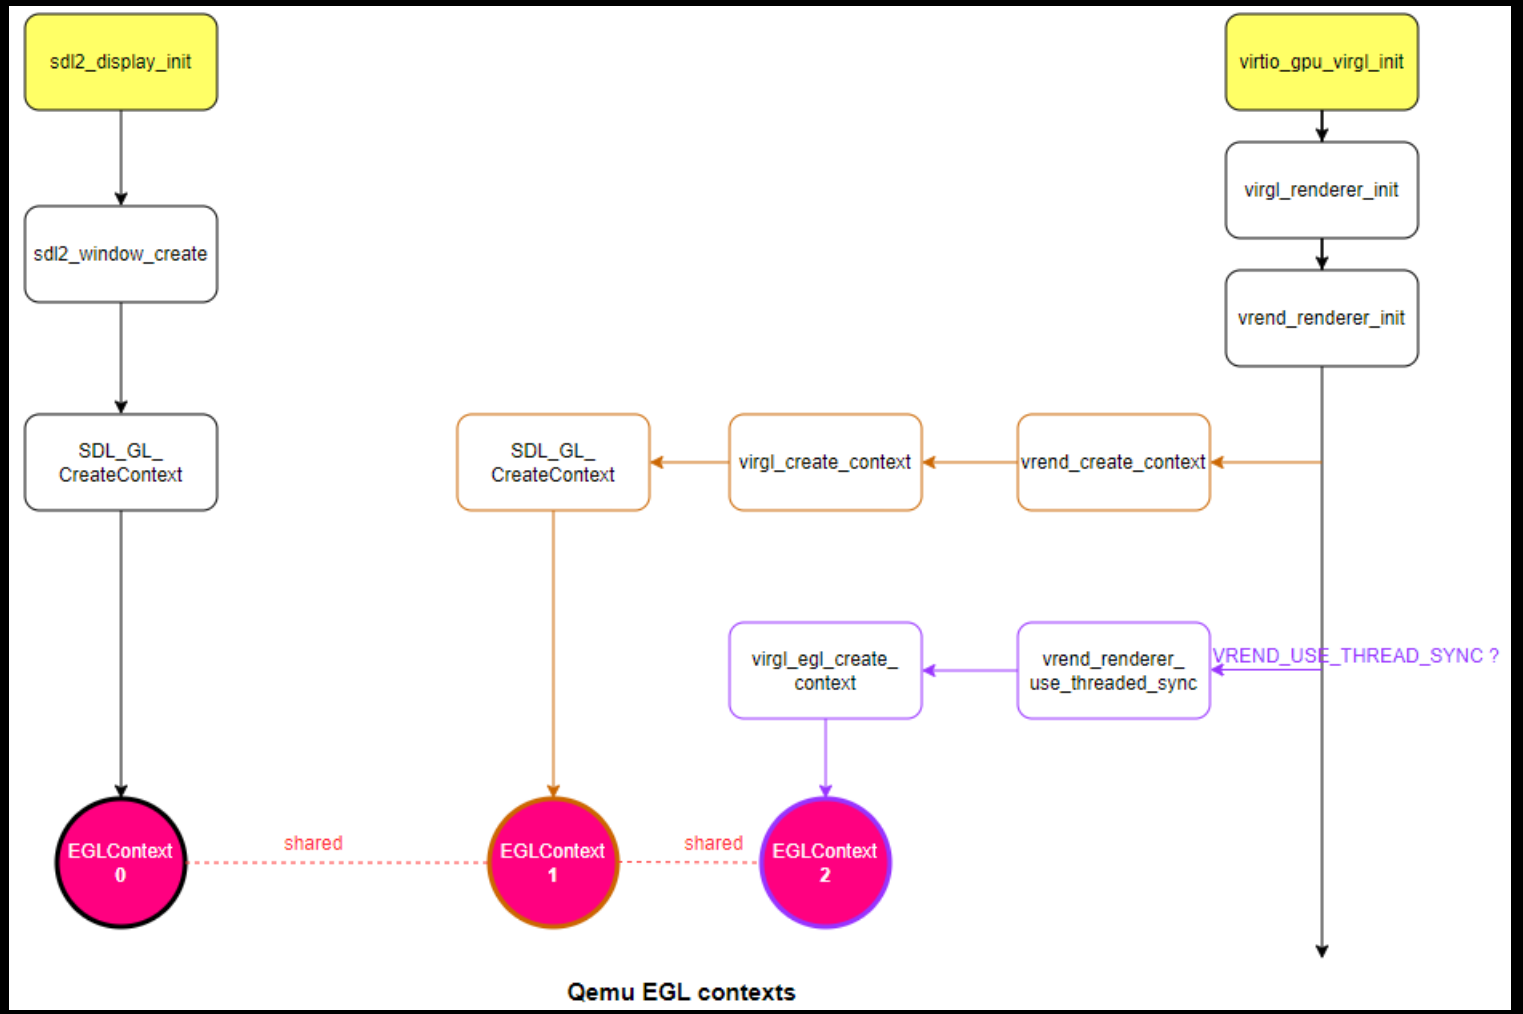

AMD
together we advance_

[Public]

Perfetto - Performance Counter

- Merge Request
  - https://gitlab.freedesktop.org/mesa/mesa/-/merge_requests/34741
- pps-producer (Per-process producer) is a systemwide daemon of mesa that can collect global performance counters
- Dataflow
  - PPS driver setup data source, send commands like init performance counter, start collecting counter value and dump counter value
  - Create the Mesa interface class to handle commands from PPS driver and then call into kernel driver
  - Calculate result for example get weighted sum of several counters and pass it to Perfetto to display it

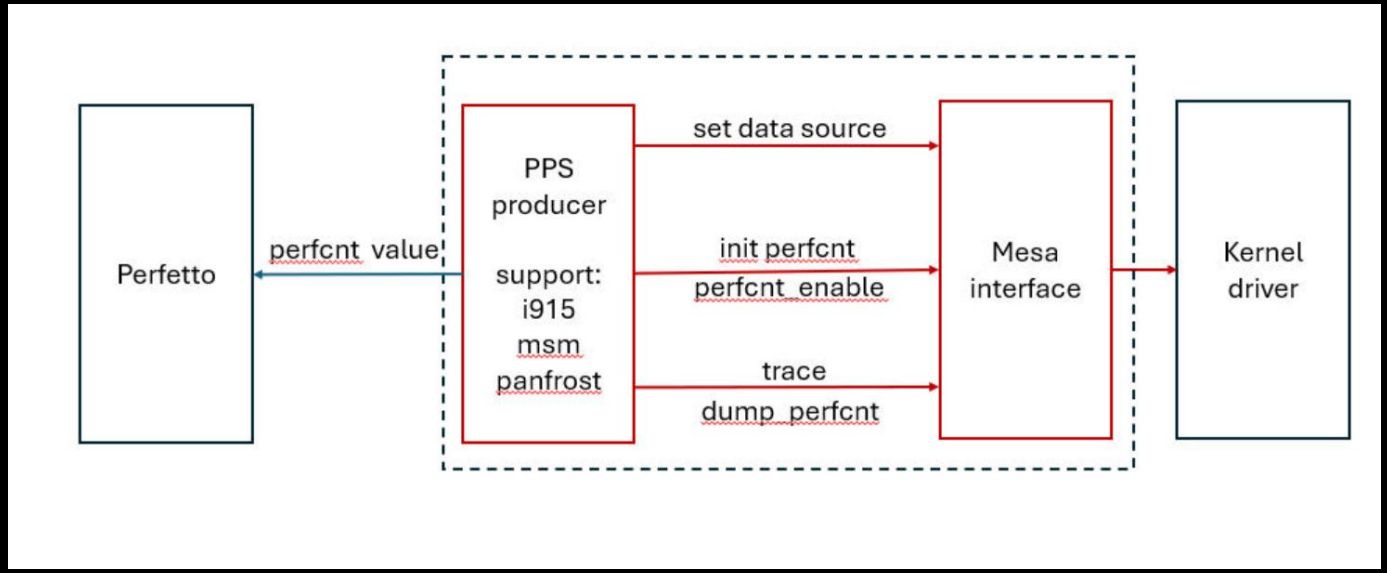

AMD
together we advance_

[Public]

Perfetto - Screenshot of Performance Counter

AMD
together we advance_

[Public]

The Best is Yet to Come

- Multiplane Overlay Virtualization
- HIP in ROCm stack for Virtualization
- 3D support for Windows Guest DOMU ☺

AMD
together we advance_

[Public]

References

- X.Org Developer’s Conference 2024
  - GPU Compute Virtualization with VirtIO
  - https://indico.freedesktop.org/event/6/contributions/285/
- Xen Summit 2024
  - GPU Para-Virtualization on Xen
  - https://xenprojectsummit2024.sched.com/event/1bCFX
- X.Org Developer’s Conference 2023
  - Xen based GPU virtualization - VirtIO/Passthrough
  - https://indico.freedesktop.org/event/4/contributions/216/
  - dGPU prime on VM
  - https://indico.freedesktop.org/event/4/contributions/189/
- Xen Summit 2023
  - VirtIO GPU and Passthrough GPU Support for Xen
  - https://xen2023.sched.com/event/1LKln

AMD
together we advance_

Thank You and Q&A

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.

© 2025 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