Hardware-Accelerated Graphics in Secure Multi-Tenant Environments

What is the current status, and what is blocking it?

Demi Marie Obenour
Invisible Things Lab

AMD GPU virtualization

- If you are using enterprise GPUs on VMware, life is great
  - AMD MxGPU is fully supported by AMD and VMware
  - Appears to be implemented in hardware, with little mediation in software
- But this solution is not an option for many
  - Both MxGPU drivers and the VMware hypervisor are proprietary
  - VMware licensing and MxGPU-capable hardware is costly
  - No support for other hypervisors whatsoever!
    - Public cloud offerings excluded
- LibVF.io has a libre offering, but it only supports ancient hardware

NVIDIA GPU virtualization

- Official: NVIDIA GRID
  - Requires enterprise GPU and per-user licensing fees
  - Only supported by proprietary NVIDIA driver
- Unofficial: LibVF.io
  - Supports inexpensive commodity GPUs
  - Still requires proprietary drivers :(
- Pre-Ampere: primarily implemented in software
  - Large attack surface
  - Not suitable for high security solutions such as Qubes OS
- Ampere and later: SR-IOV, but likely a lot of software mediation
  - Host driver must be at least as recent as guest driver
  - IOCTLs must pass through host

Intel GPU virtualization

- Intel claims to support SR-IOV
- Unfortunately, zero Linux driver code is available
  - Windows drivers may support it, but that is not helpful to those who do not run Windows.
- Intel claims to have SR-IOV support for their Linux drivers
  - Intel developers: Please send patches!

Windows vGPU

- RemoteFX deprecated, GPU-PV and GPU-P are its replacements
- Both support commodity hardware
- GPU-PV (vGPU) is just a wrapper for kernel-mode DirectX driver :(
  - Same attack surface as exposed to userspace code
  - Much of the security benefit of virtualization is lost
- Microsoft warns that vGPU may increase attack surface in Windows Sandbox
  - While any solution would increase attack surface somewhat, an explicit warning implies that Microsoft is less than confident about its security
- Not enough information about GPU-P

Conclusion

- NVIDIA continues to be a non-starter for some applications
  - Proprietary driver requirements
  - Software-based mediation has large attack surface
- AMD is best if one has $$$
  - Only public clouds and VMware supported
  - Neither are suitable for most high-assurance solutions!
- Intel: Please release SR-IOV capable Linux driver code!
  - Or at least specifications from which others could do the same