



KVMGT: a Full GPU Virtualization Solution

Oct, 2014

Jike Song, jike.song@intel.com

Zhiyuan Lv, zhiyuan.lv@intel.com

Kevin Tian, kevin.tian@intel.com



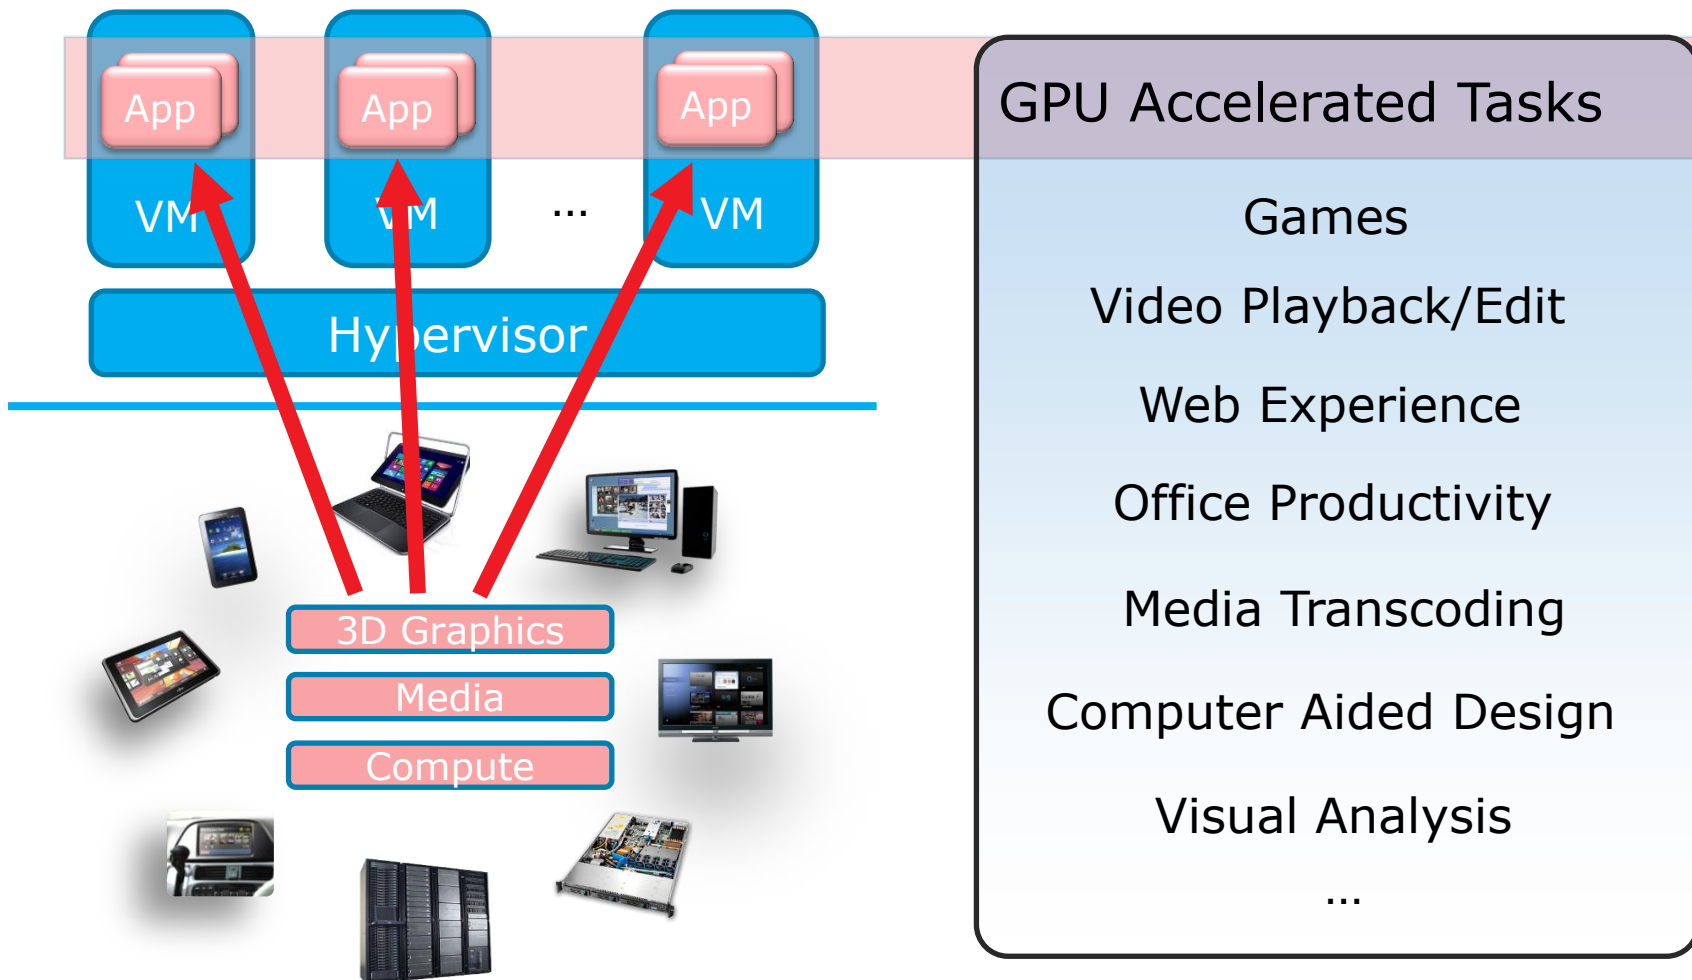


Agenda

- Background
- KVMGT
- Status
- Summary



GPU Virtualization Momentum



GPU virtualization becomes a fundamental requirement



Requirements of GPU Virtualization



Performance



Direct GPU acceleration



Capability



Consistent visual experience



Sharing



Multiple Virtual Machines



GPU Virtualization Approaches



API Forwarding	Direct Pass-through	Full GPU Virtualization
<p>Pros:</p> <ul style="list-style-type: none"> • Performance • Sharing <p>Cons:</p> <ul style="list-style-type: none"> • No media/GPGPU • Compatibility 	<p>Pros:</p> <ul style="list-style-type: none"> • Performance • Capability <p>Cons:</p> <ul style="list-style-type: none"> • No sharing 	<p>Pros:</p> <ul style="list-style-type: none"> • Performance • Capability • Sharing





Intel® GPU Virtualization Technology

- Intel® GVT-s for API Forwarding
 - ✓ Make existing API forwarding protocols running best on Intel® platform

- Intel® GVT-d for Direct Pass-through
 - ✓ Xen GPU pass-through upstreaming in progress

- Intel® GVT-g for Full GPU Virtualization
 - ✓ SW based approach with a good balance of performance, feature and sharing



Intel® GVT-g



- Full GPU virtualization
 - ✓ Run native graphics driver in VM



Full-featured vGPU

- Mediated pass-through
 - ✓ Pass-through performance critical operations
 - ✓ Trap-and-emulate privileged operations



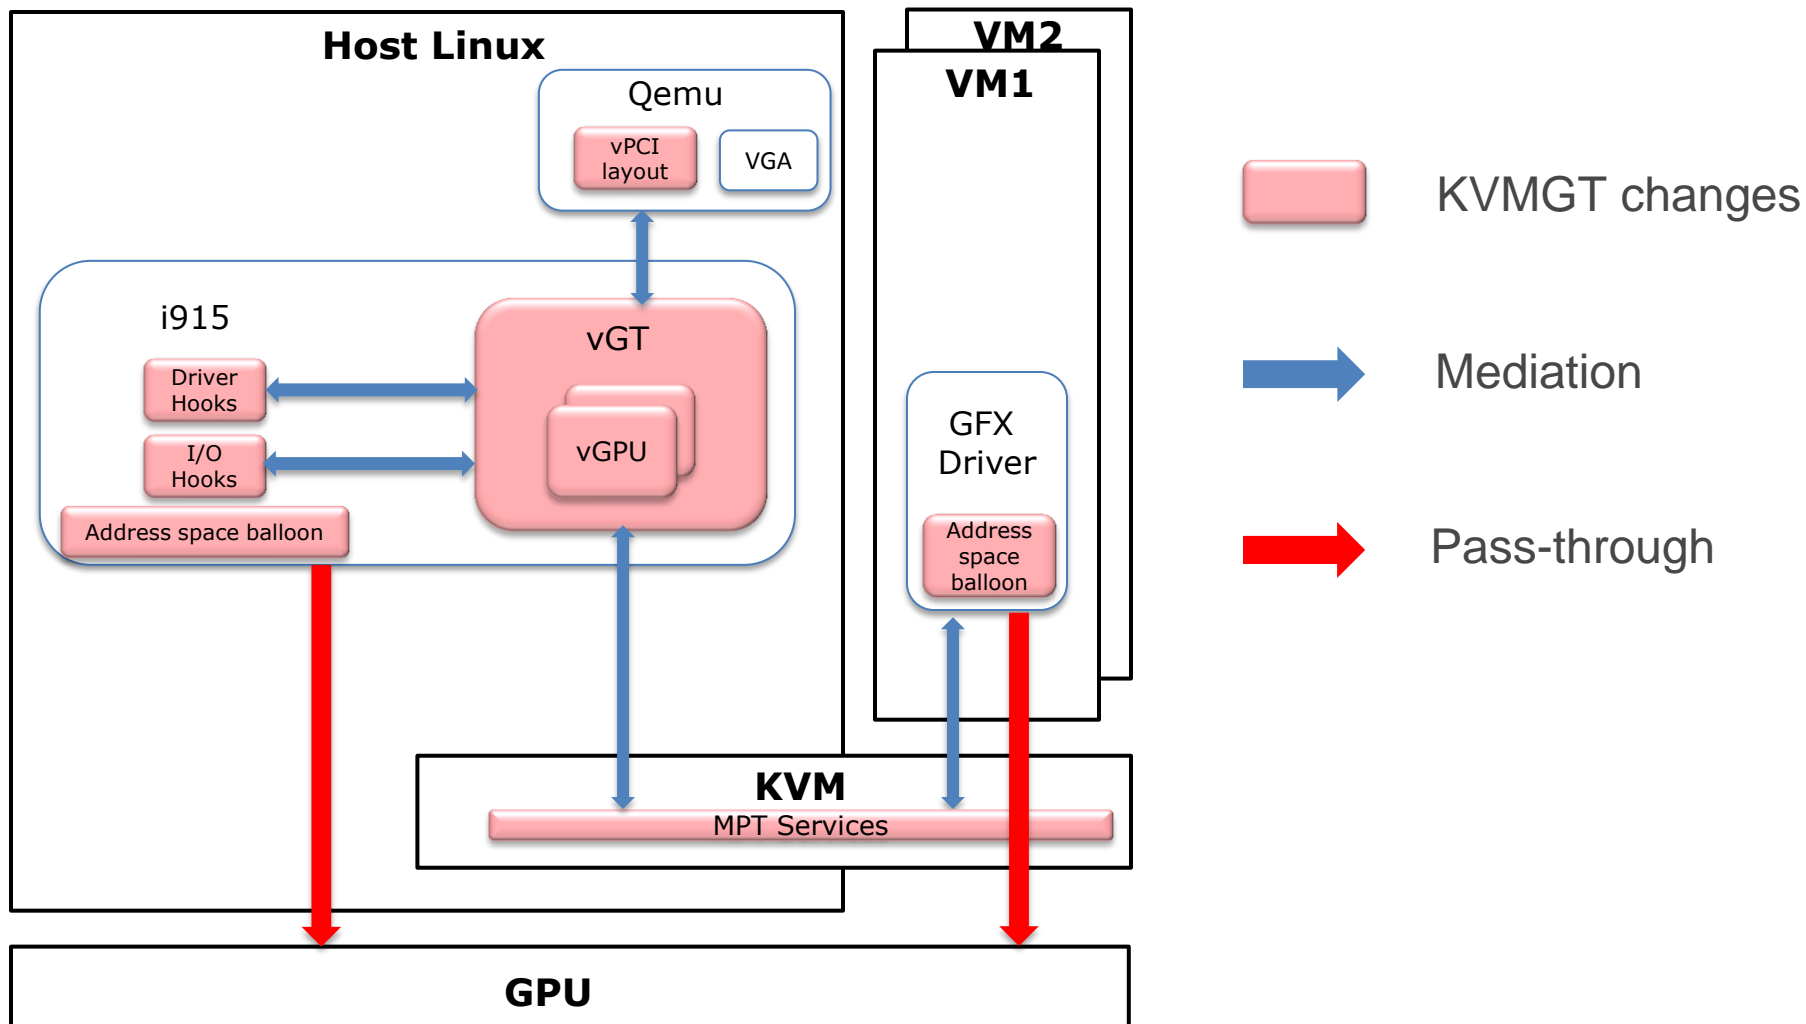
Close-to-native performance

Flexible sharing among VMs

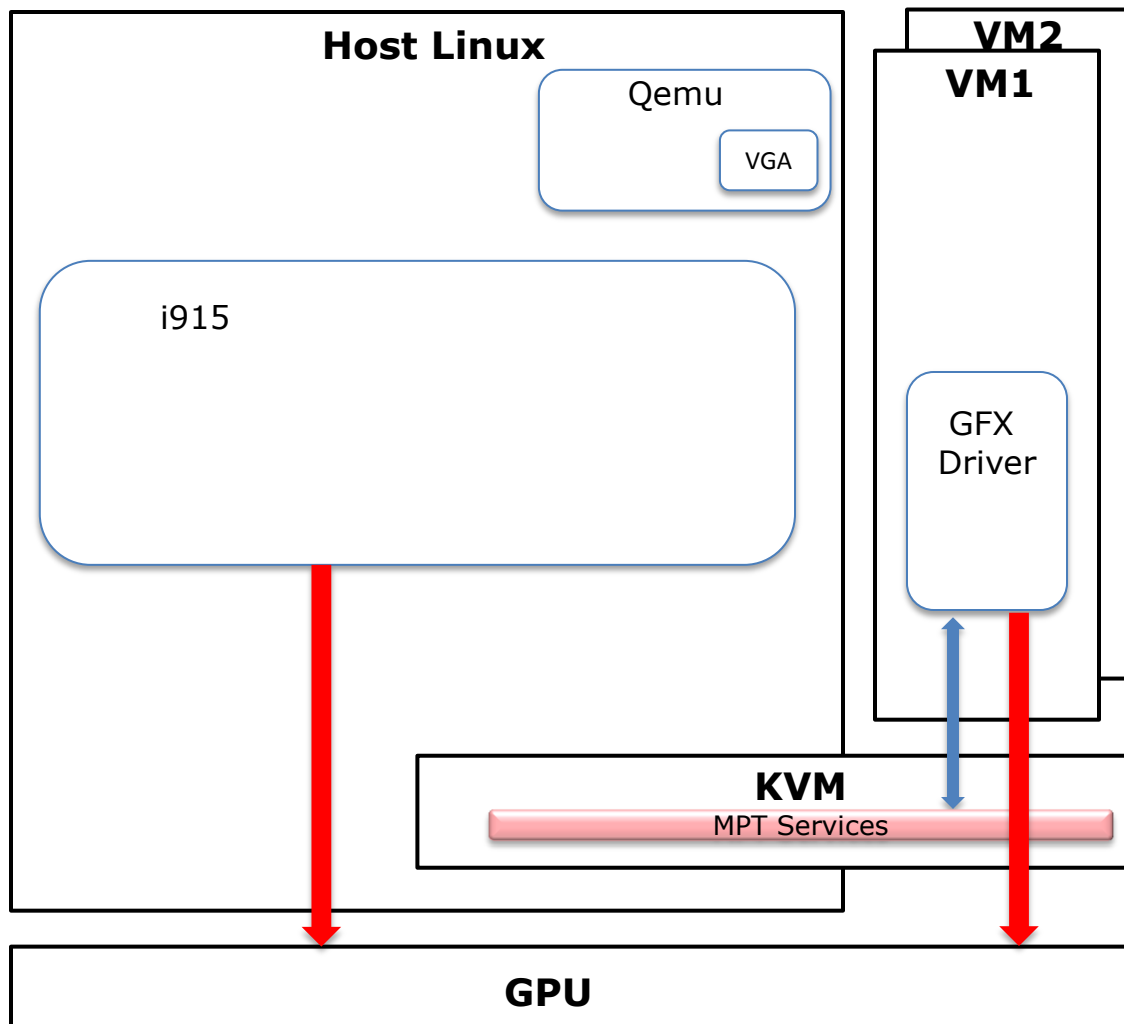
- KVMGT is the open source implementation of Intel® GVT-g for KVM
 - ✓ Support Intel® Processor Graphics built into 4th generation Intel® Core™ processors
 - ✓ vGT device model is GPL/MIT dual-licensed and VMM agnostic



KVMGT Architecture



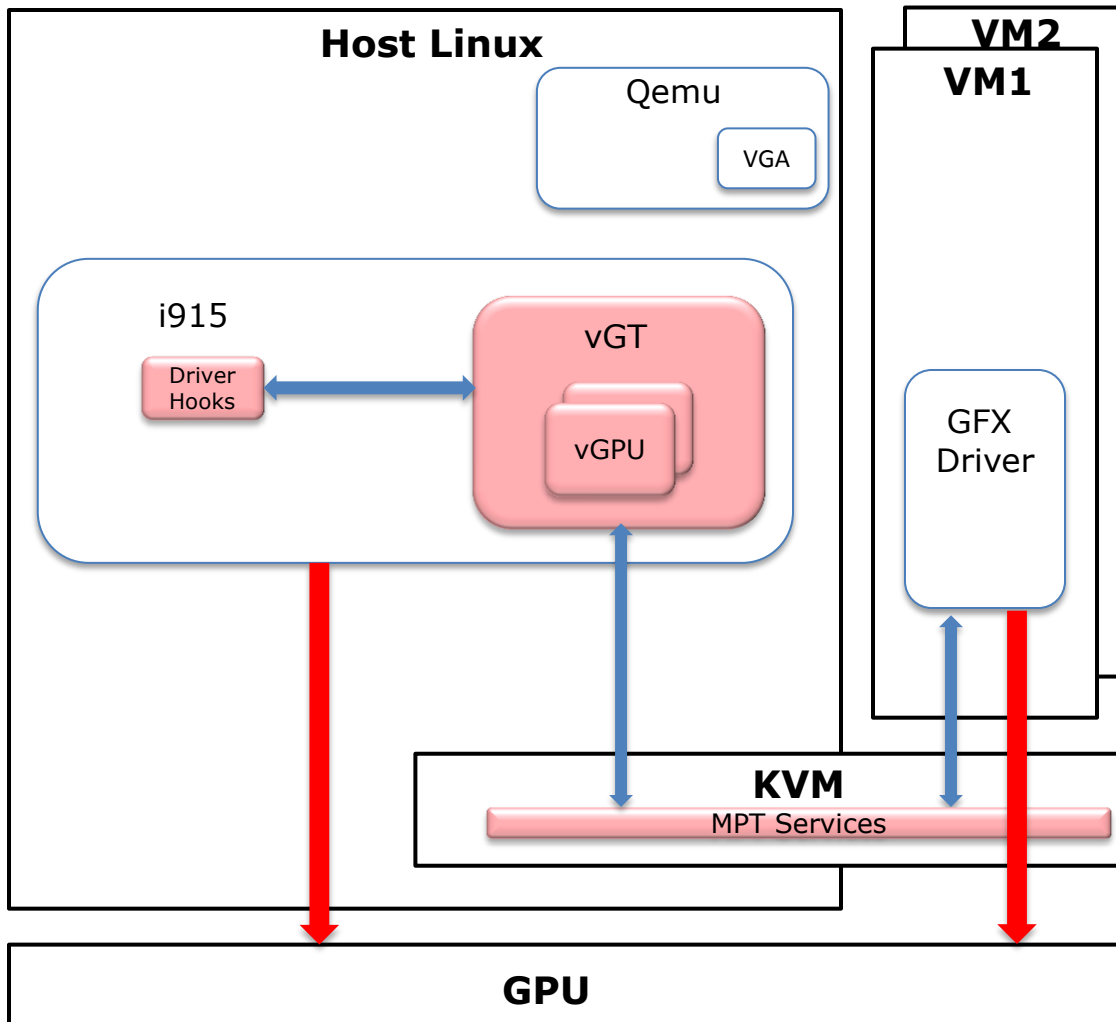
Mediated Pass-Through (MPT) Services



- In-kernel device model framework
 - ✓ Selectively pass-through
 - ✓ Trap-and-emulation



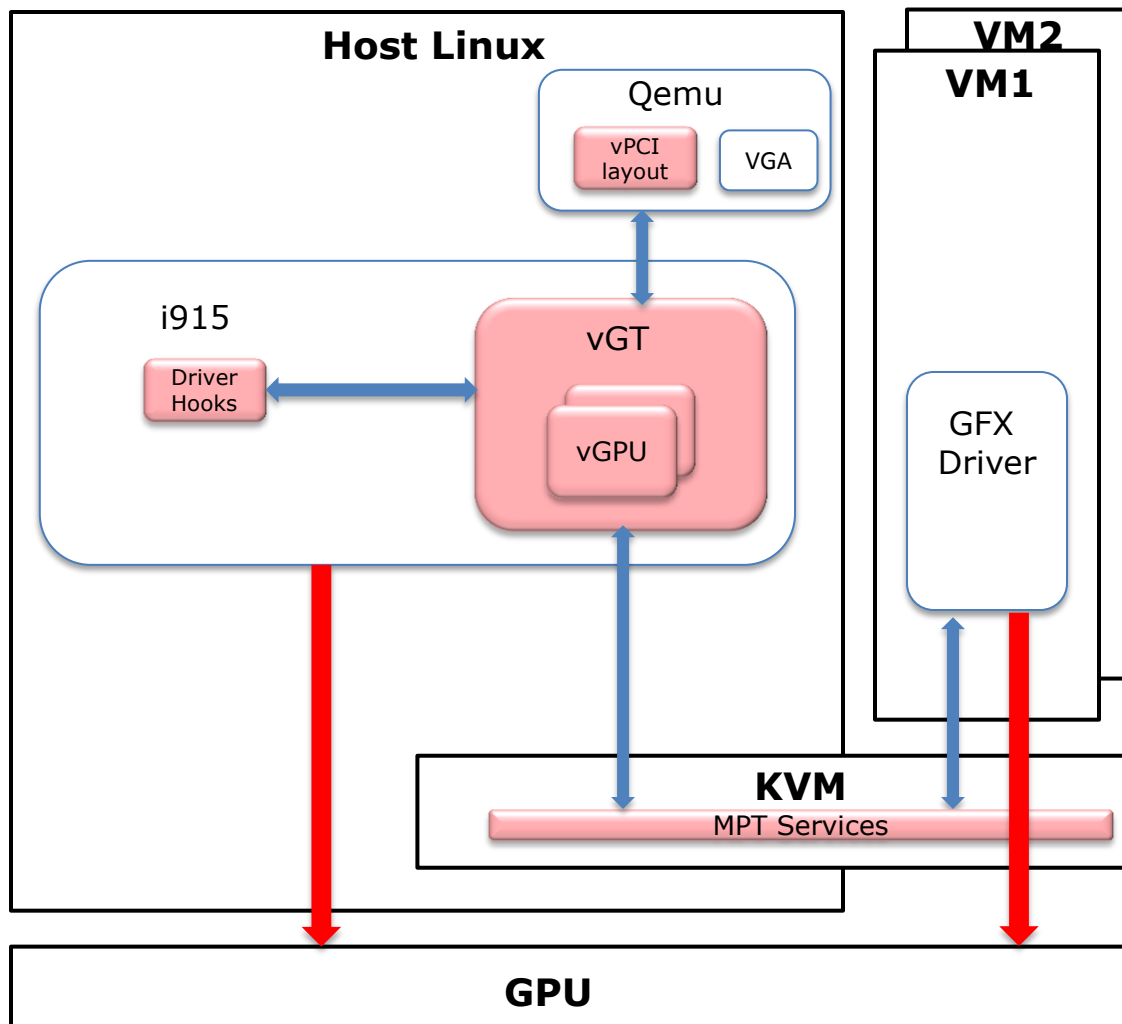
vGT



- Part of host i915 driver
- Maintain 1 vGPU instance for each VM
- vGPU scheduling in a time-shared manner



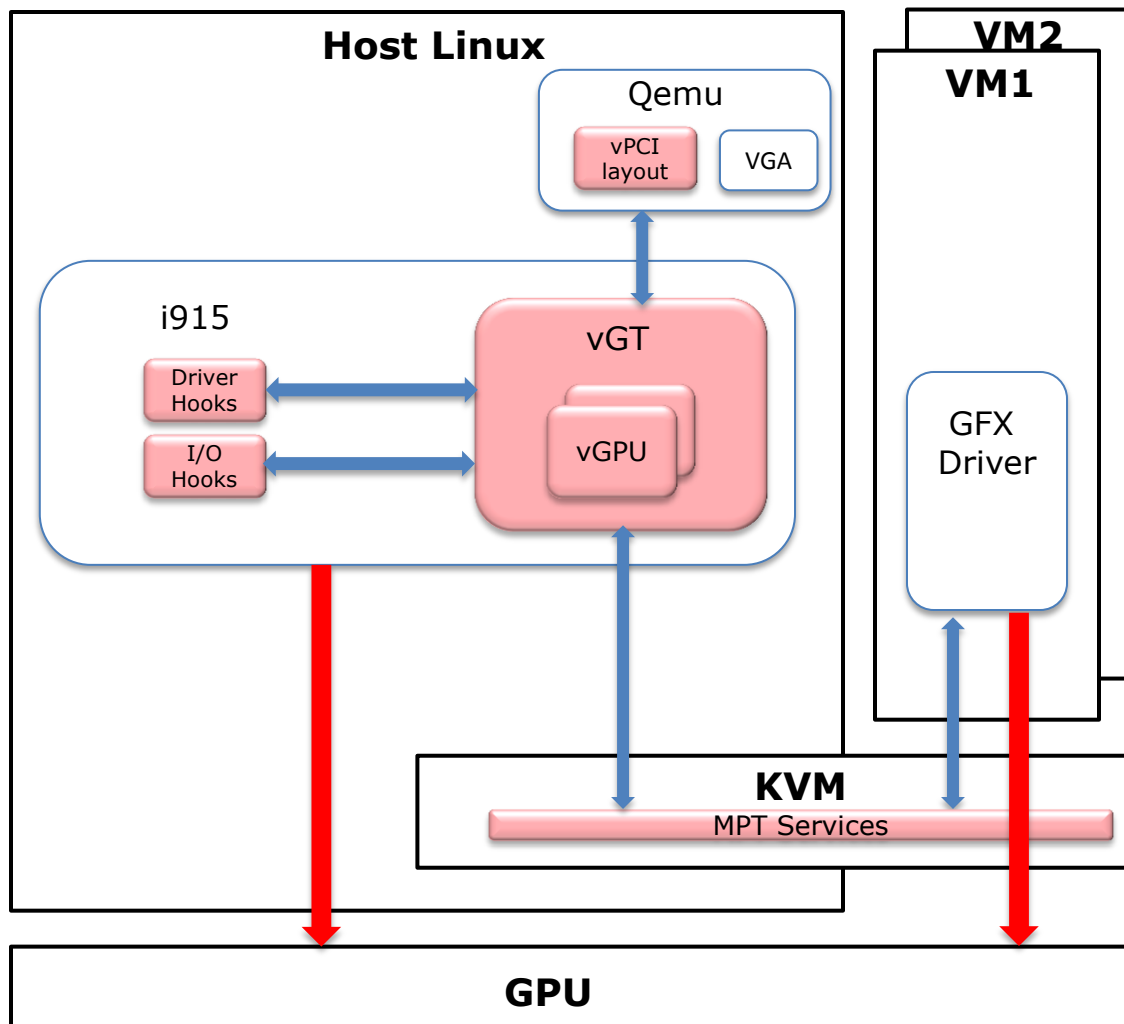
Qemu Changes



- Present vGPU in the virtual PCI hierarchy
- Request vGPU instance creation/destroy



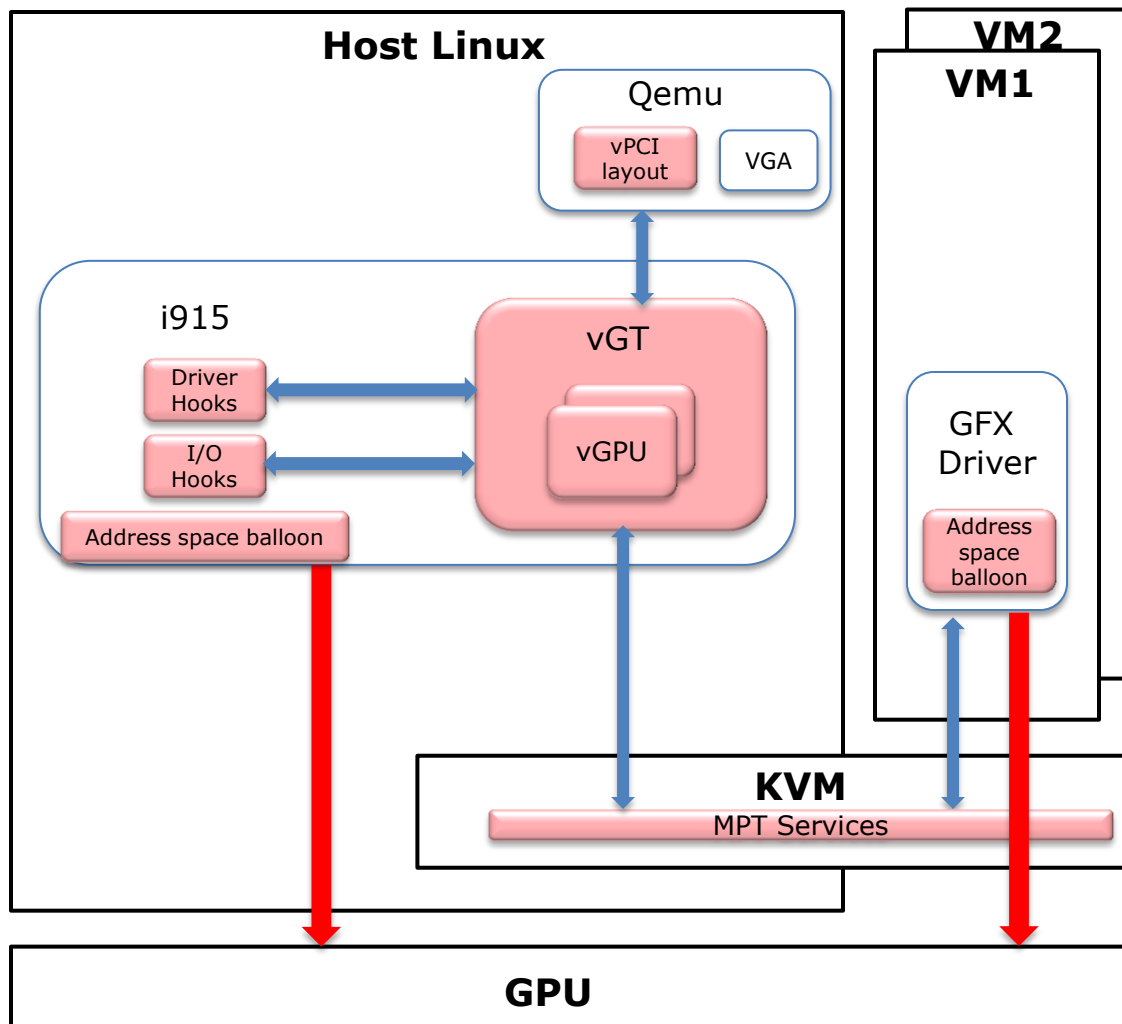
Host i915 Mediation



- Host i915 accesses are also mediated
 - ✓ Hook in i915 I/O interfaces
- Centralized GPU sharing management
 - ✓ For both host and VMs



Address Space Ballooning



- Graphics memory partitioning
 - ✓ Guest view vs. host view
- Address space ballooning
 - ✓ Consistent guest/host views
 - ✓ No need for address translation



Implementation Status


- Basic features complete
 - ✓ Linux guest works steadily with PPGTT(Per-Process Graphics Translation Table) disabled
- MPT services
 - ✓ Trap via kvm_io_device, similar to the in-kernel irqchip
 - ✓ Pass-through via a special in-kernel memslot
- vGT and i915 driver
 - ✓ Start upstreaming patches to i915 community
- Qemu
 - ✓ Provide MCH/PCH emulation similar to GPU pass-through




Performance



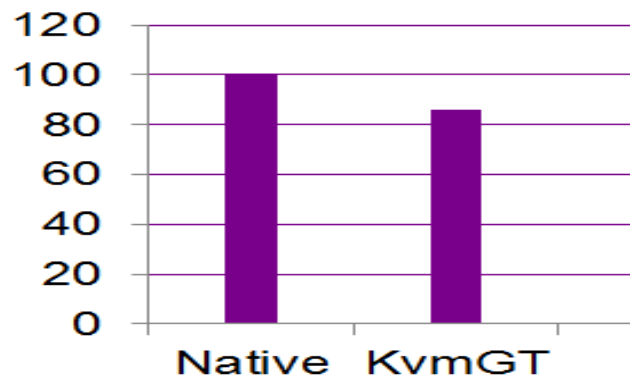
lightsmark



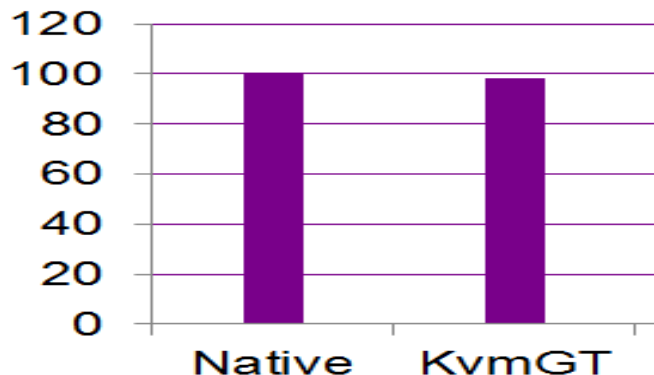
nexuiz



urbanterror



warsow



Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks



TODO

- Implement shadow PPGTT
 - ✓ Is a must to support Windows guest

- Clean up and stabilize the prototype code

- Involve with community for architectural discussion

- Work with community for upstreaming





Summary

- Intel ® GVT-g provides a good balance of performance, feature and sharing capability
 - ✓ Minimum impact on performance critical operations
 - ✓ Running native graphics driver in VM
 - ✓ Multiplexing capability

- KVMGT implements Intel ® GVT-g for KVM, with >80% of native 3D performance
 - ✓ RFC code is coming

- Community contribution is highly appreciated!



Notices and Disclaimers



1. © 2014 Intel Corporation.
2. Intel, the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.
3. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.
4. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.
5. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
6. Configurations: {describe config + what test used + who did testing}. For more complete information visit <http://www.intel.com/benchmarks>
7. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks .
8. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.
9. Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
10. Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer.
11. No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.
12. Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, updated BIOS, and virtual machine monitor (VMM). Functionality, performance or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Check with your system manufacturer. Learn more at <http://www.intel.com/go/virtualization>.
13. KVM Remote Control (Keyboard, Video, Mouse) is only available with Intel® Core™ i5 vPro™ and Core™ i7 vPro™ processors with integrated graphics and Intel® Active Management technology activated. Discrete graphics are not supported.

