Hardware Support for Virtualization

Guillaume Urvoy-Keller

January 3, 2019

Source documents

- https://www.intel.com/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html
- https://en.wikipedia.org/wiki/X86_virtualization
- https://www.hardwaresecrets.com/everything-you-need-to-know-about-the-intel-virtualization-technology/3/

Intel and AMD

- x86 ISA common to Intel and AMD processors
- Intel Virtualization Technology (Intel VT)
- AMD Virtualization (AMD-V)
- Introduced in 2006...
- ....new functionalities introduced gradually.

Intel VT portfolio

A set of technologies that hypervisors might use:

CPU virtualization

- CPU abstraction, i.e. let guest VM, and especially guest OS, do as much as it can without support of hypervisor
  $\rightarrow$ introduction of new instructions in x86 ISA.
- Live migration over different CPU generations
  $\rightarrow$ When migration occurs, the status of execution on the processor has to be migrated and it might need adaptation

Intel VT CPU virtualization

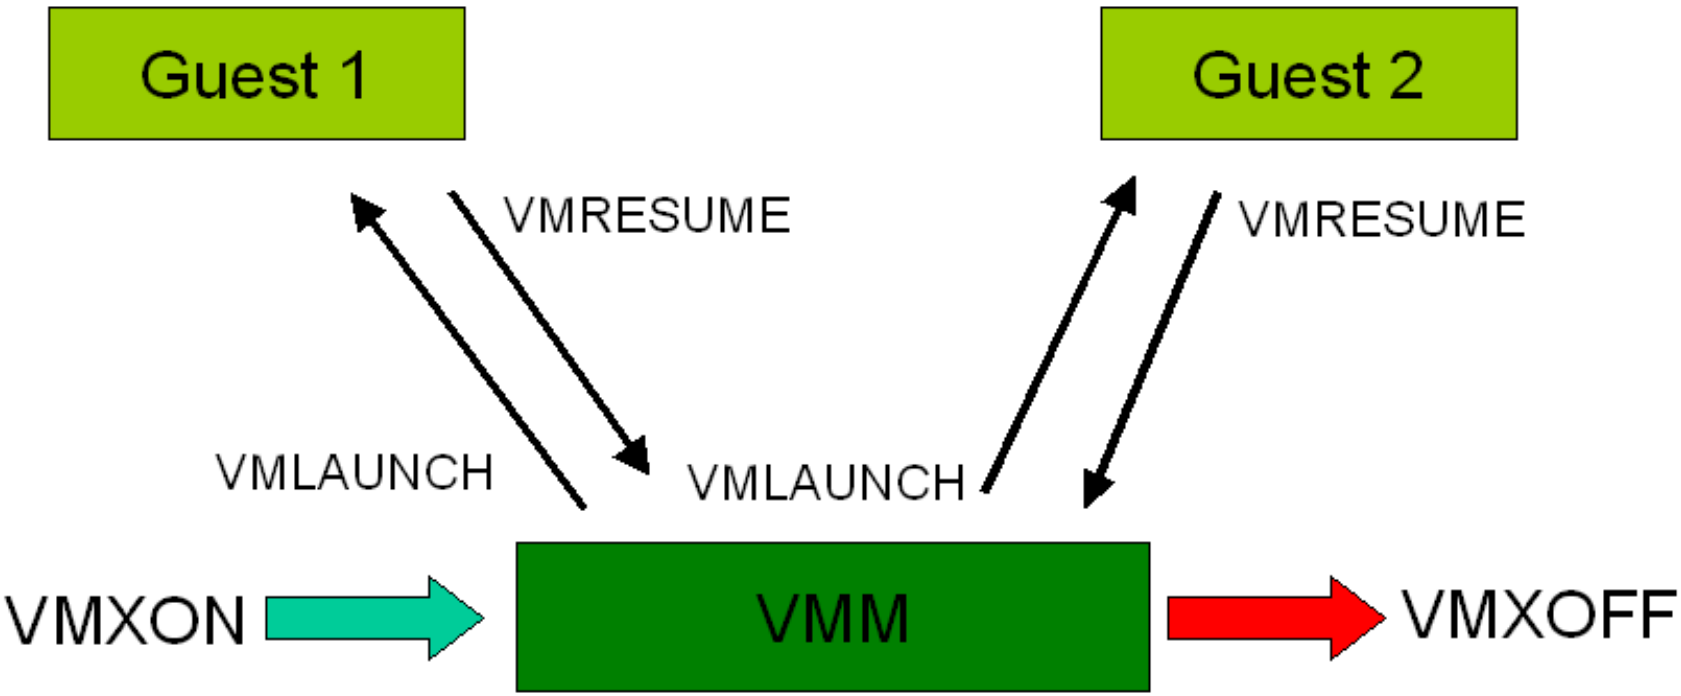

Intel VT CPU virtualization

These instructions, esp. VMXON and VMXOFF explains the following schema with the "non-real" ring -1 where the hypervisor sits :

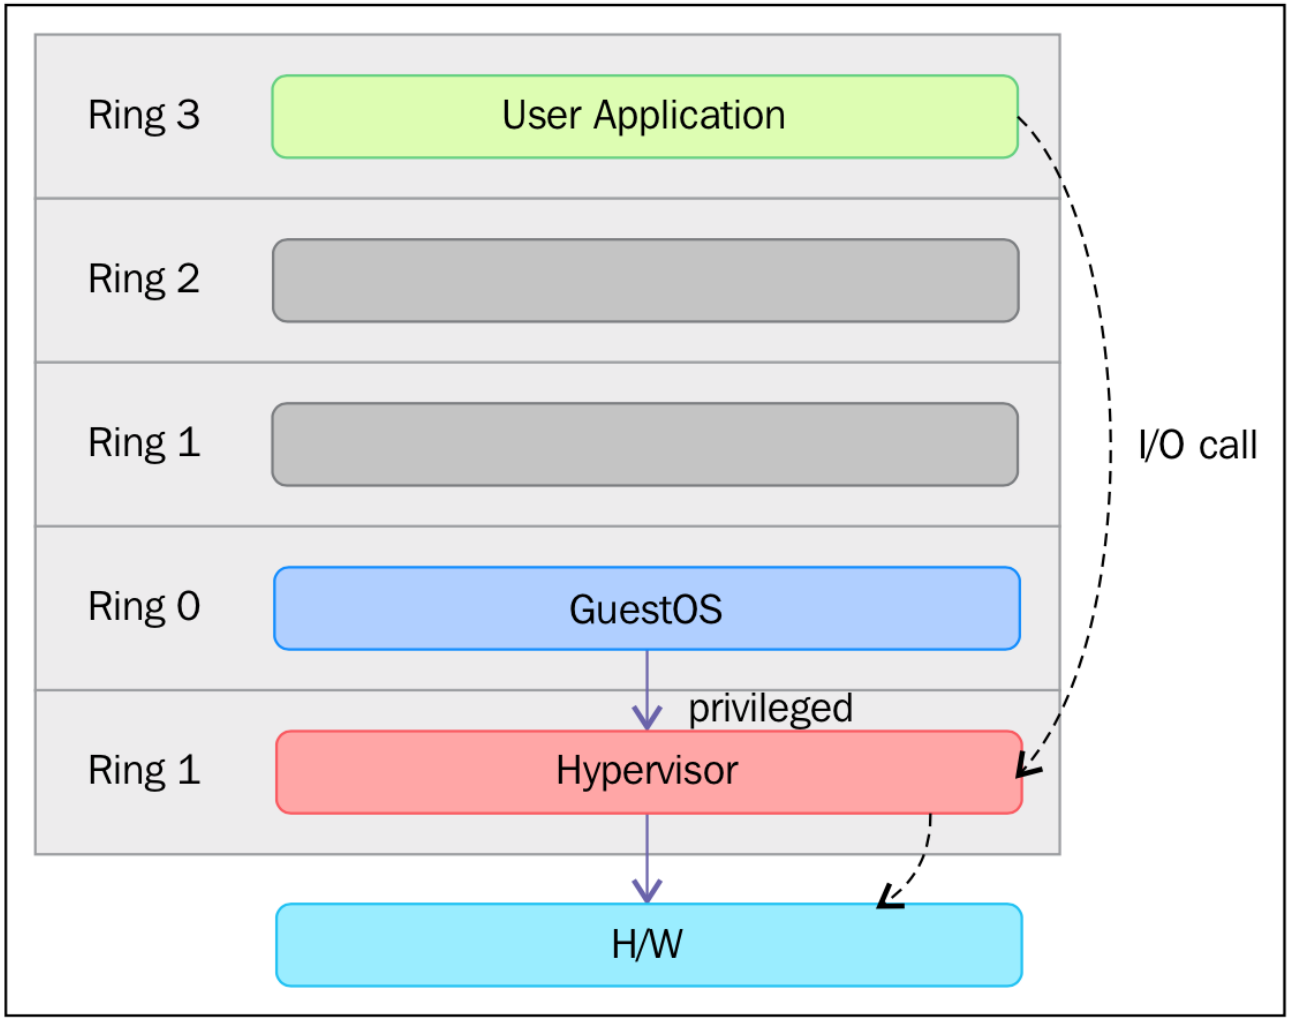

Intel VT portfolio - cont'd

Memory Virtualization

- Non virtualized case: there exists already virtual pages (seen by programs) and real pages seen by OS.
- In virtualized cases, two levels of traductions:
  - Guest program $\rightarrow$ guest OS
  - Guest OS $\rightarrow$ host OS
- Intel offers extended page tables (EPT)
  - known under the name *SLAT, Second-Level Address Translation* in AMD

Intel Extended Page Tables

- Extended Page Tables (EPT) is an Intel x86 virtualization technology for the memory management unit (MMU), ie circuit in charge of translating virtual memory (seen by program) to real memory (on chipset)
- EPT support is found in Intel's Core i3, Core i5, Core i7 and Core i9 CPUs....
- From wikipedia: "According to a VMware evaluation paper: "EPT provides performance gains of up to 48% for MMU-intensive benchmarks and up to 600% for MMU-intensive microbenchmarks", although it can actually cause code to run slower than a software implementation in some corner cases [9]."

[9] : `http://www.vmware.com/pdf/Perf_ESX_Intel-EPT-eval.pdf`

Intel VT portfolio - cont'd

I/O Virtualization

- E.g. enable :
  - packet processing offloading to network adapters (to compute checksums - a known performance issue at high speed)
  - Direct disk I/O
- Set of technologies:
  - Intel® Virtualization Technology for Directed I/O (VT-d)
  - Virtual Machine Device Queues (VMDQ),
  - Single Root I/O Virtualization (SR-IOV, a PCI-SIG standard)
  - Intel® Data Direct I/O Technology (Intel® DDIO) enhancements

Intel VT I/O Virtualization: example of DMA

- Guest OS would like to do DMA for its programs
- Problem: guest OS does not see physical pages (only hypervisor does)
- Use of I/O MMU $\rightarrow$ maps guest-physical address to host-physical addresses