

Motivation and Goal

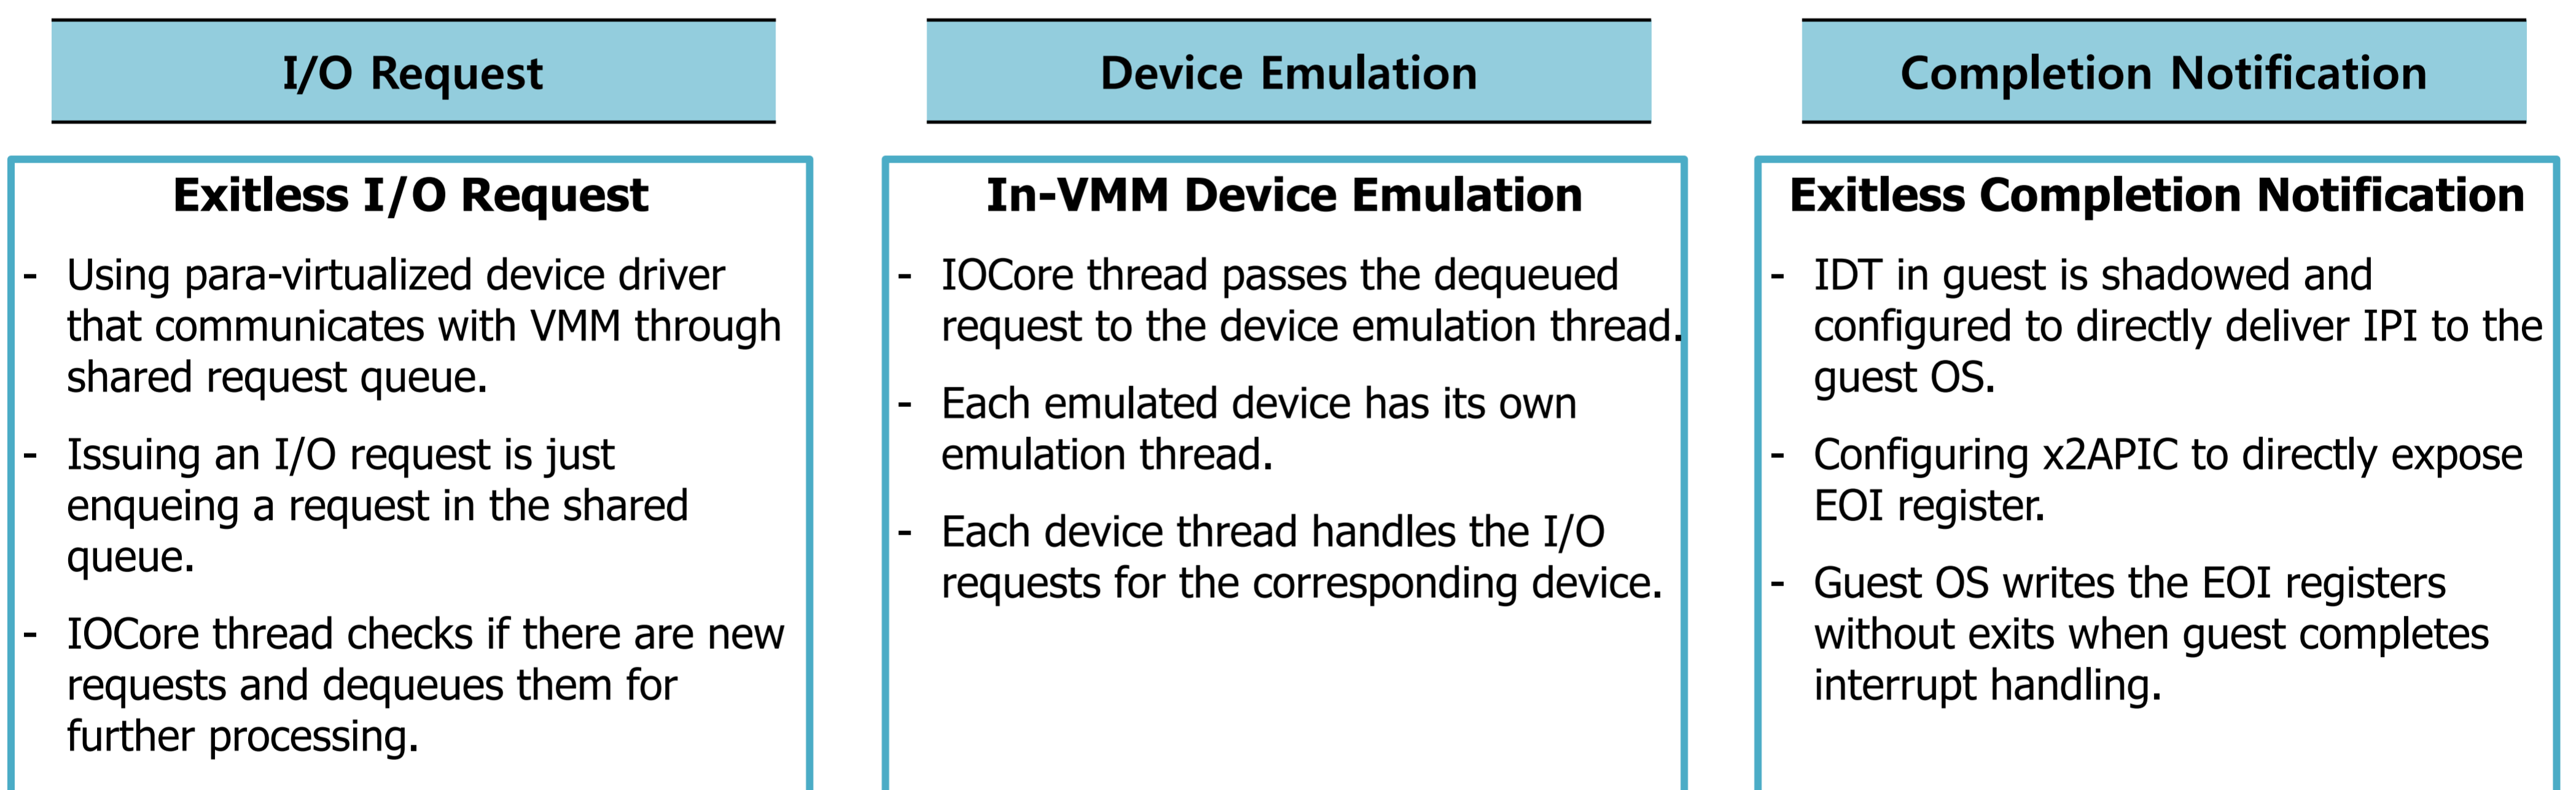
I/O workloads in VM suffer performance degradation due to virtualization overhead.

	I/O Request Cost	Device Emulation Cost	Completion Notification Cost
Direct Cost	VM-to-VMM world switching (Issue I/O requests to VMM)	Domain or user-kernel mode switching (Emulate requests in a separated domain or process)	VM-to-VMM world switching (Inject interrupts to guest and signal the completion of interrupt)
Indirect Cost	Cache pollution TLB flush	Cache pollution TLB flush	Cache pollution TLB flush
Synchronous Cost	ALL processes on a VM are stopped.		A process on a CPU is stopped.

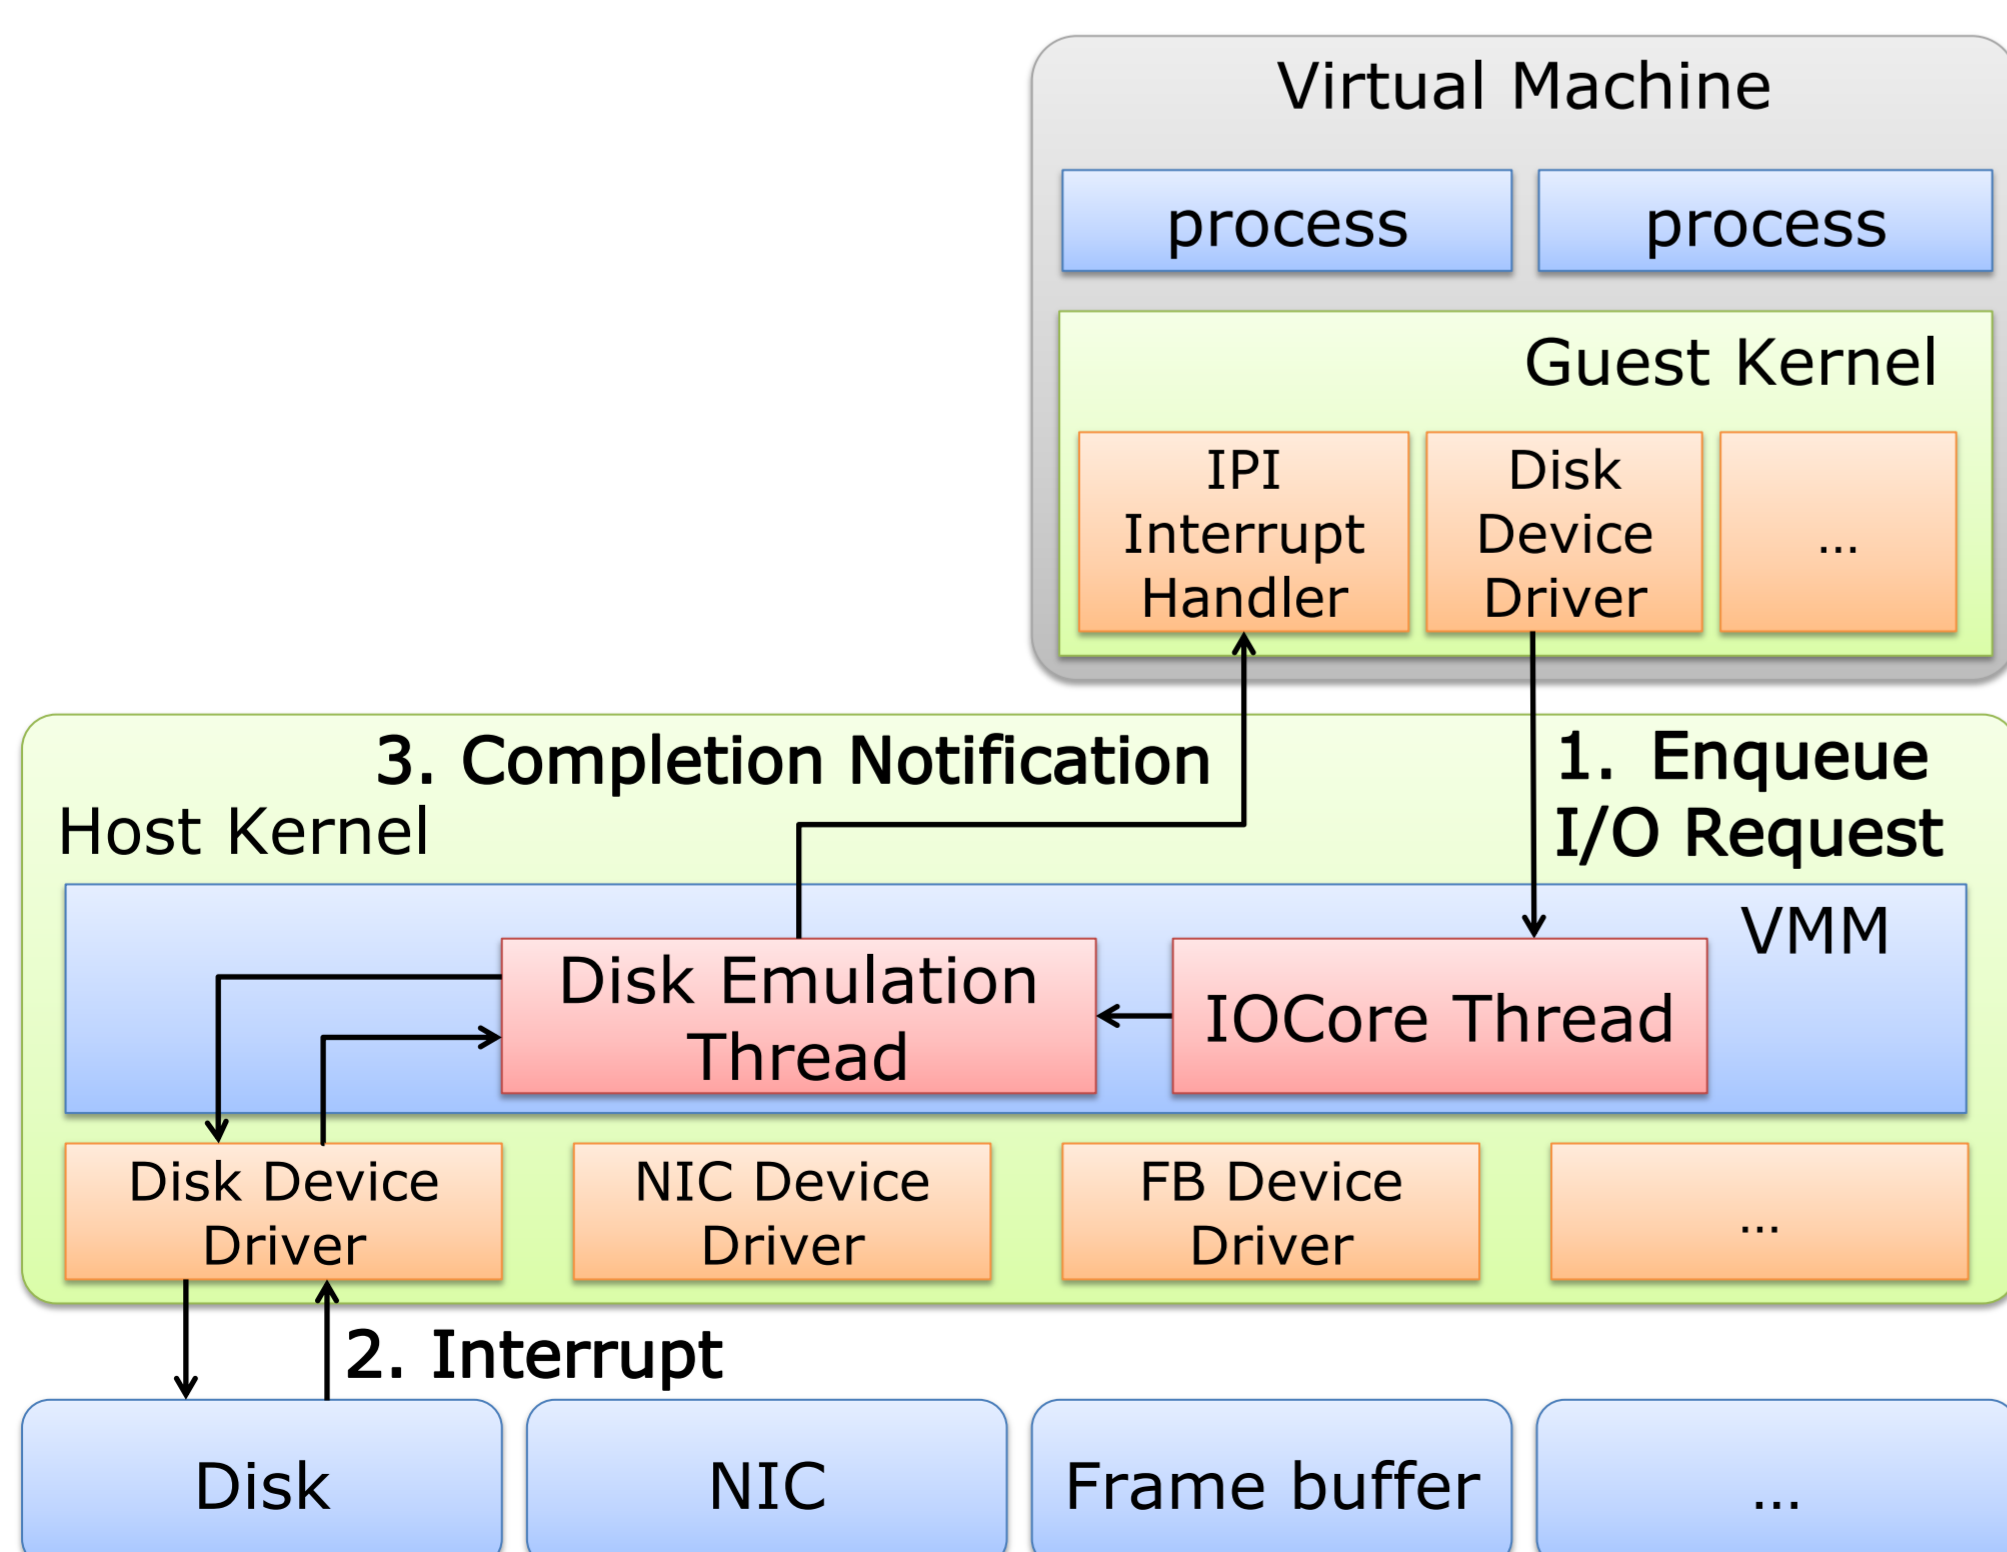
Goal : Accelerating I/O performance in VM by reducing the cost induced by exits

Key Techniques

Reducing the I/O virtualization overhead by exploiting multicore architecture

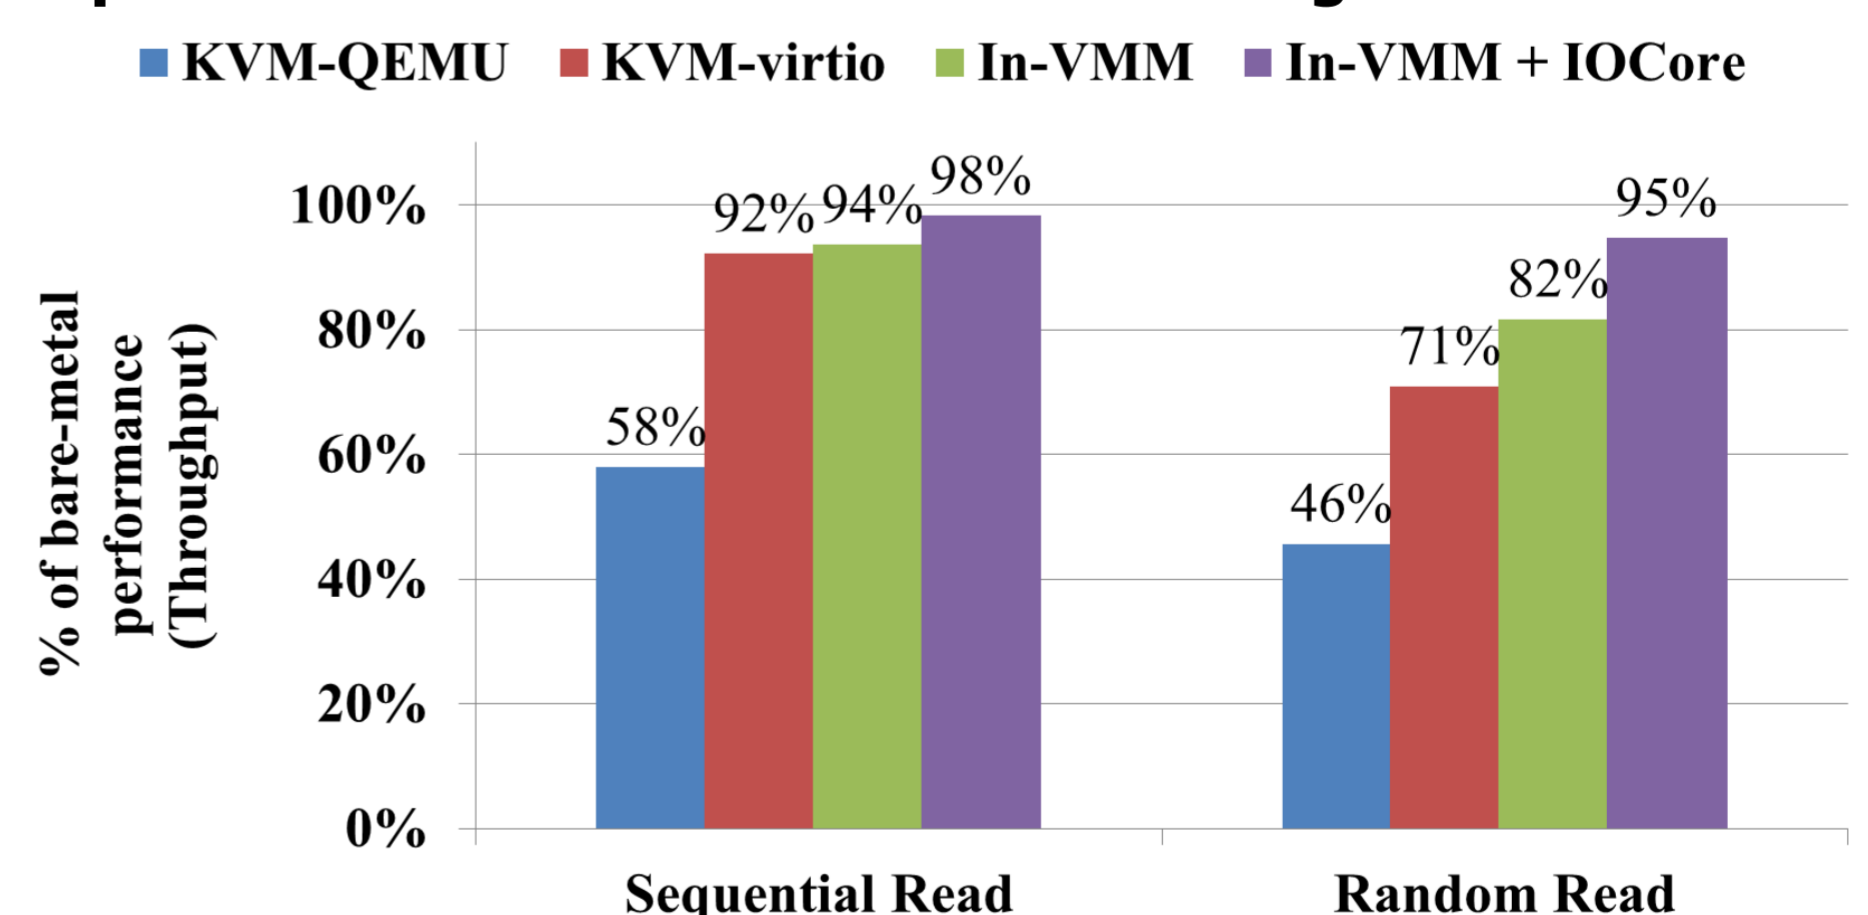


Architecture



Experimental Results

Sequential reads and random reads on high-end SSD



Our approach highly close to the bare-metal performance
Seq. Reads by 98% and Rnd. Reads by 95%

In-VMM device emulation benefit: 2~11%

Exitless I/O Request benefit: 4~13%

Considering our prototype is **not optimized** and **exitless completion notification** is not implemented yet, the performance result is quite encouraging