


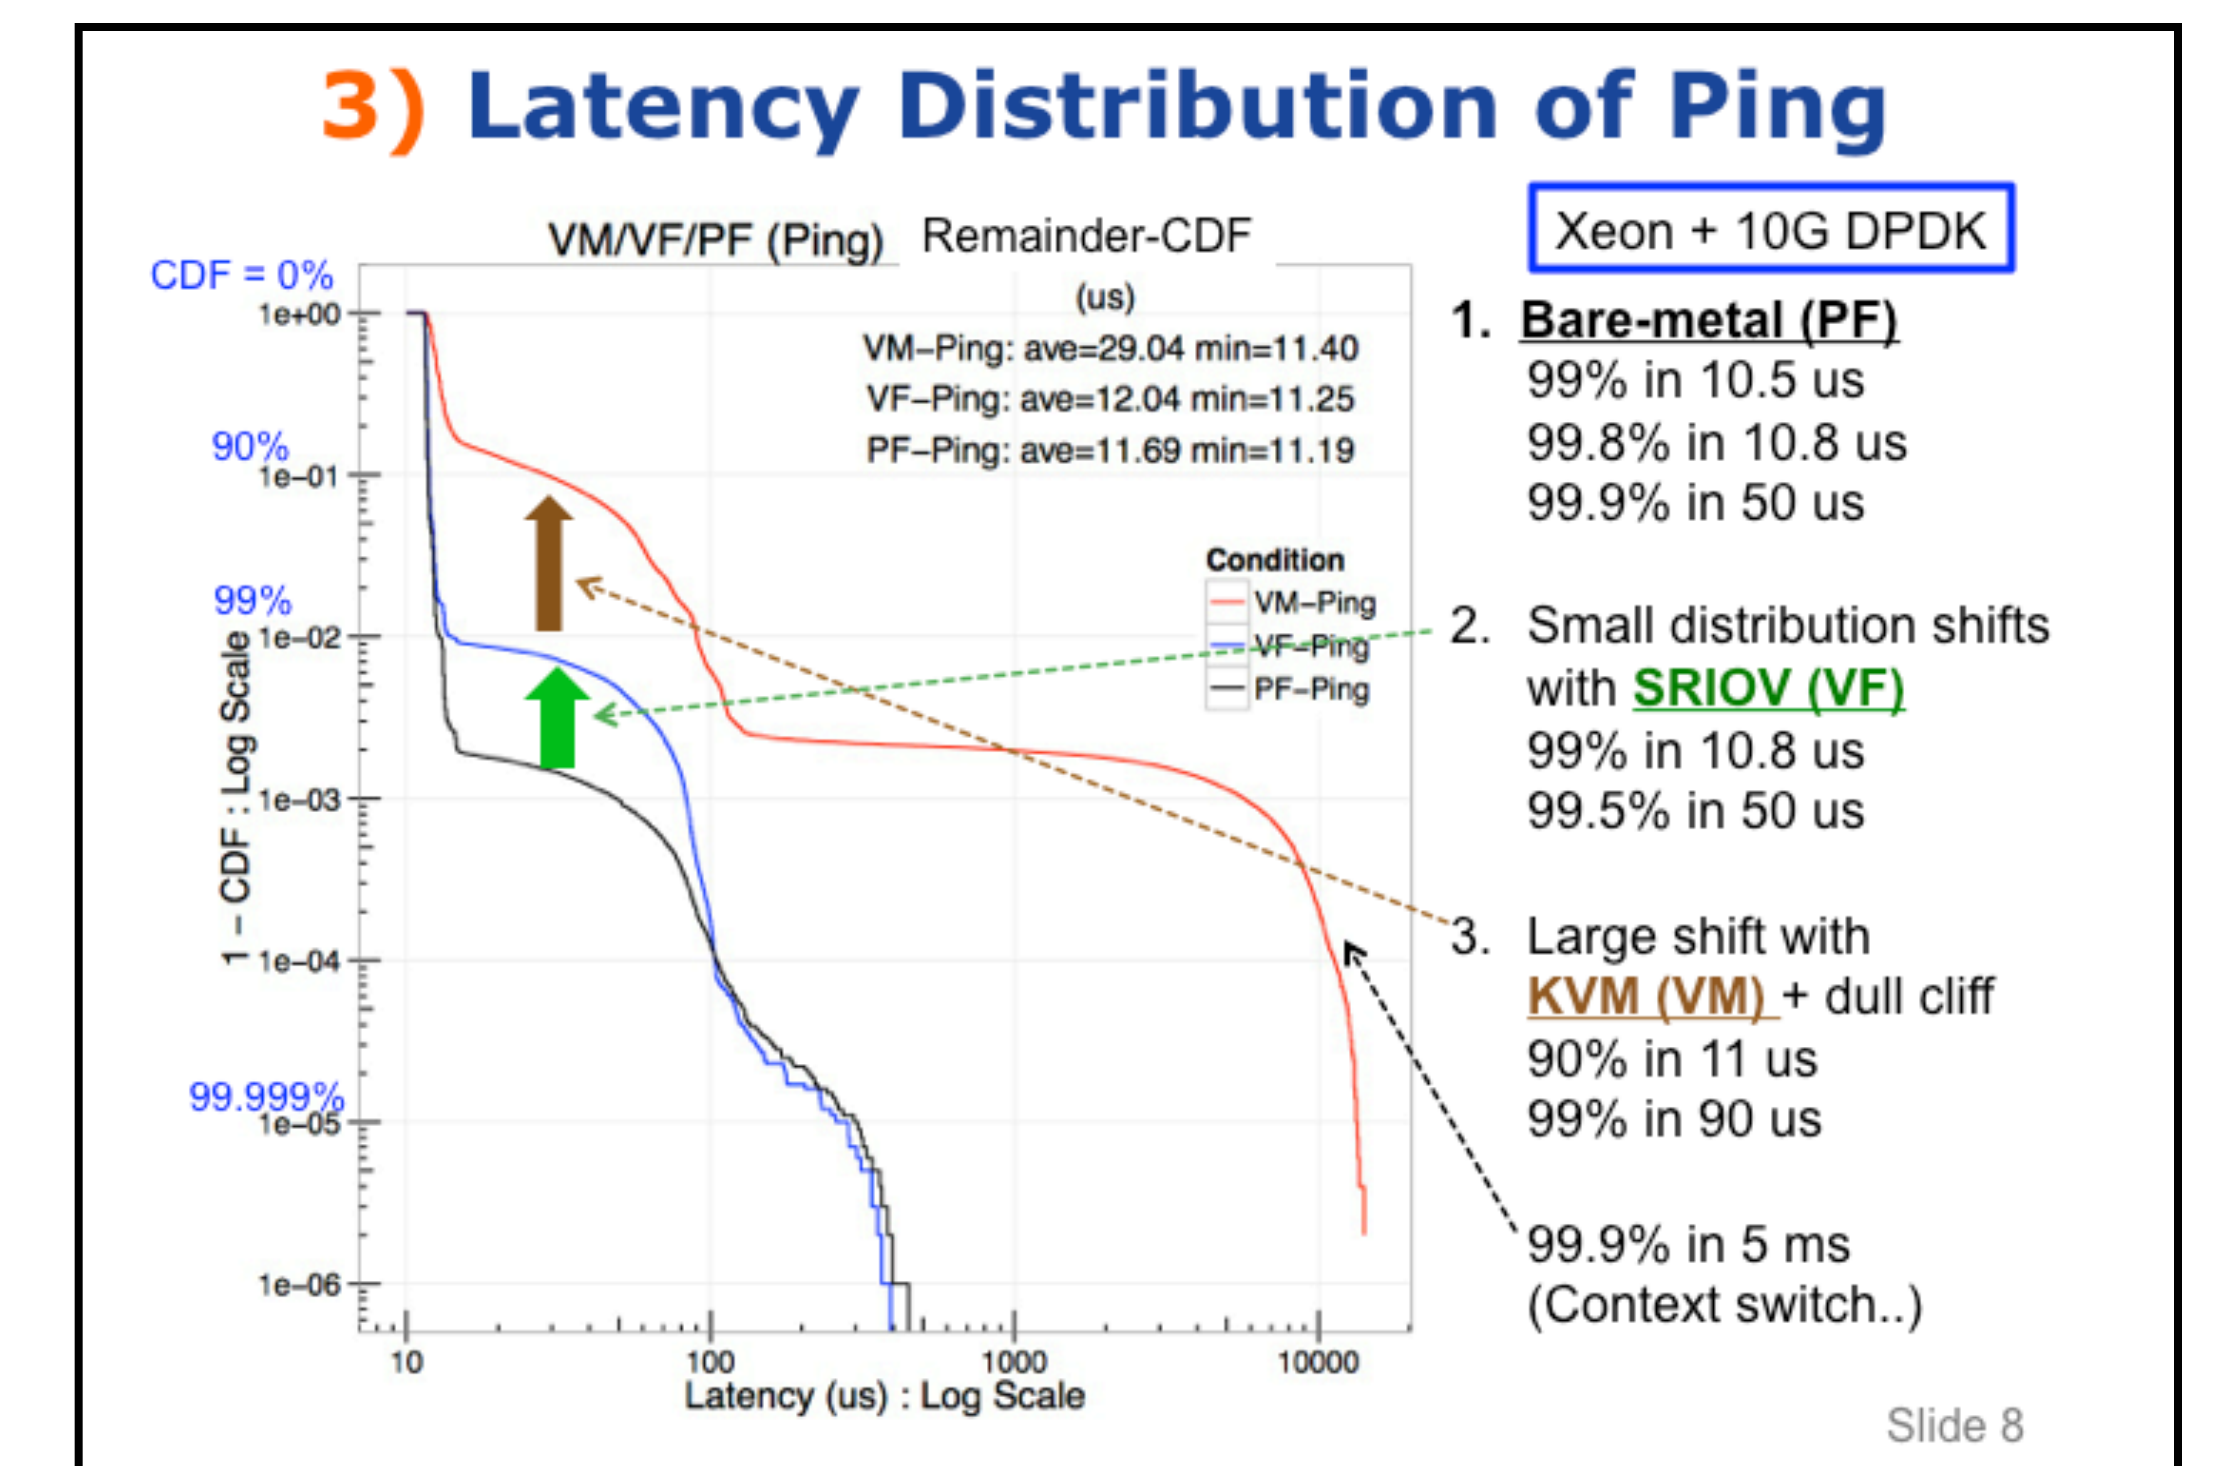
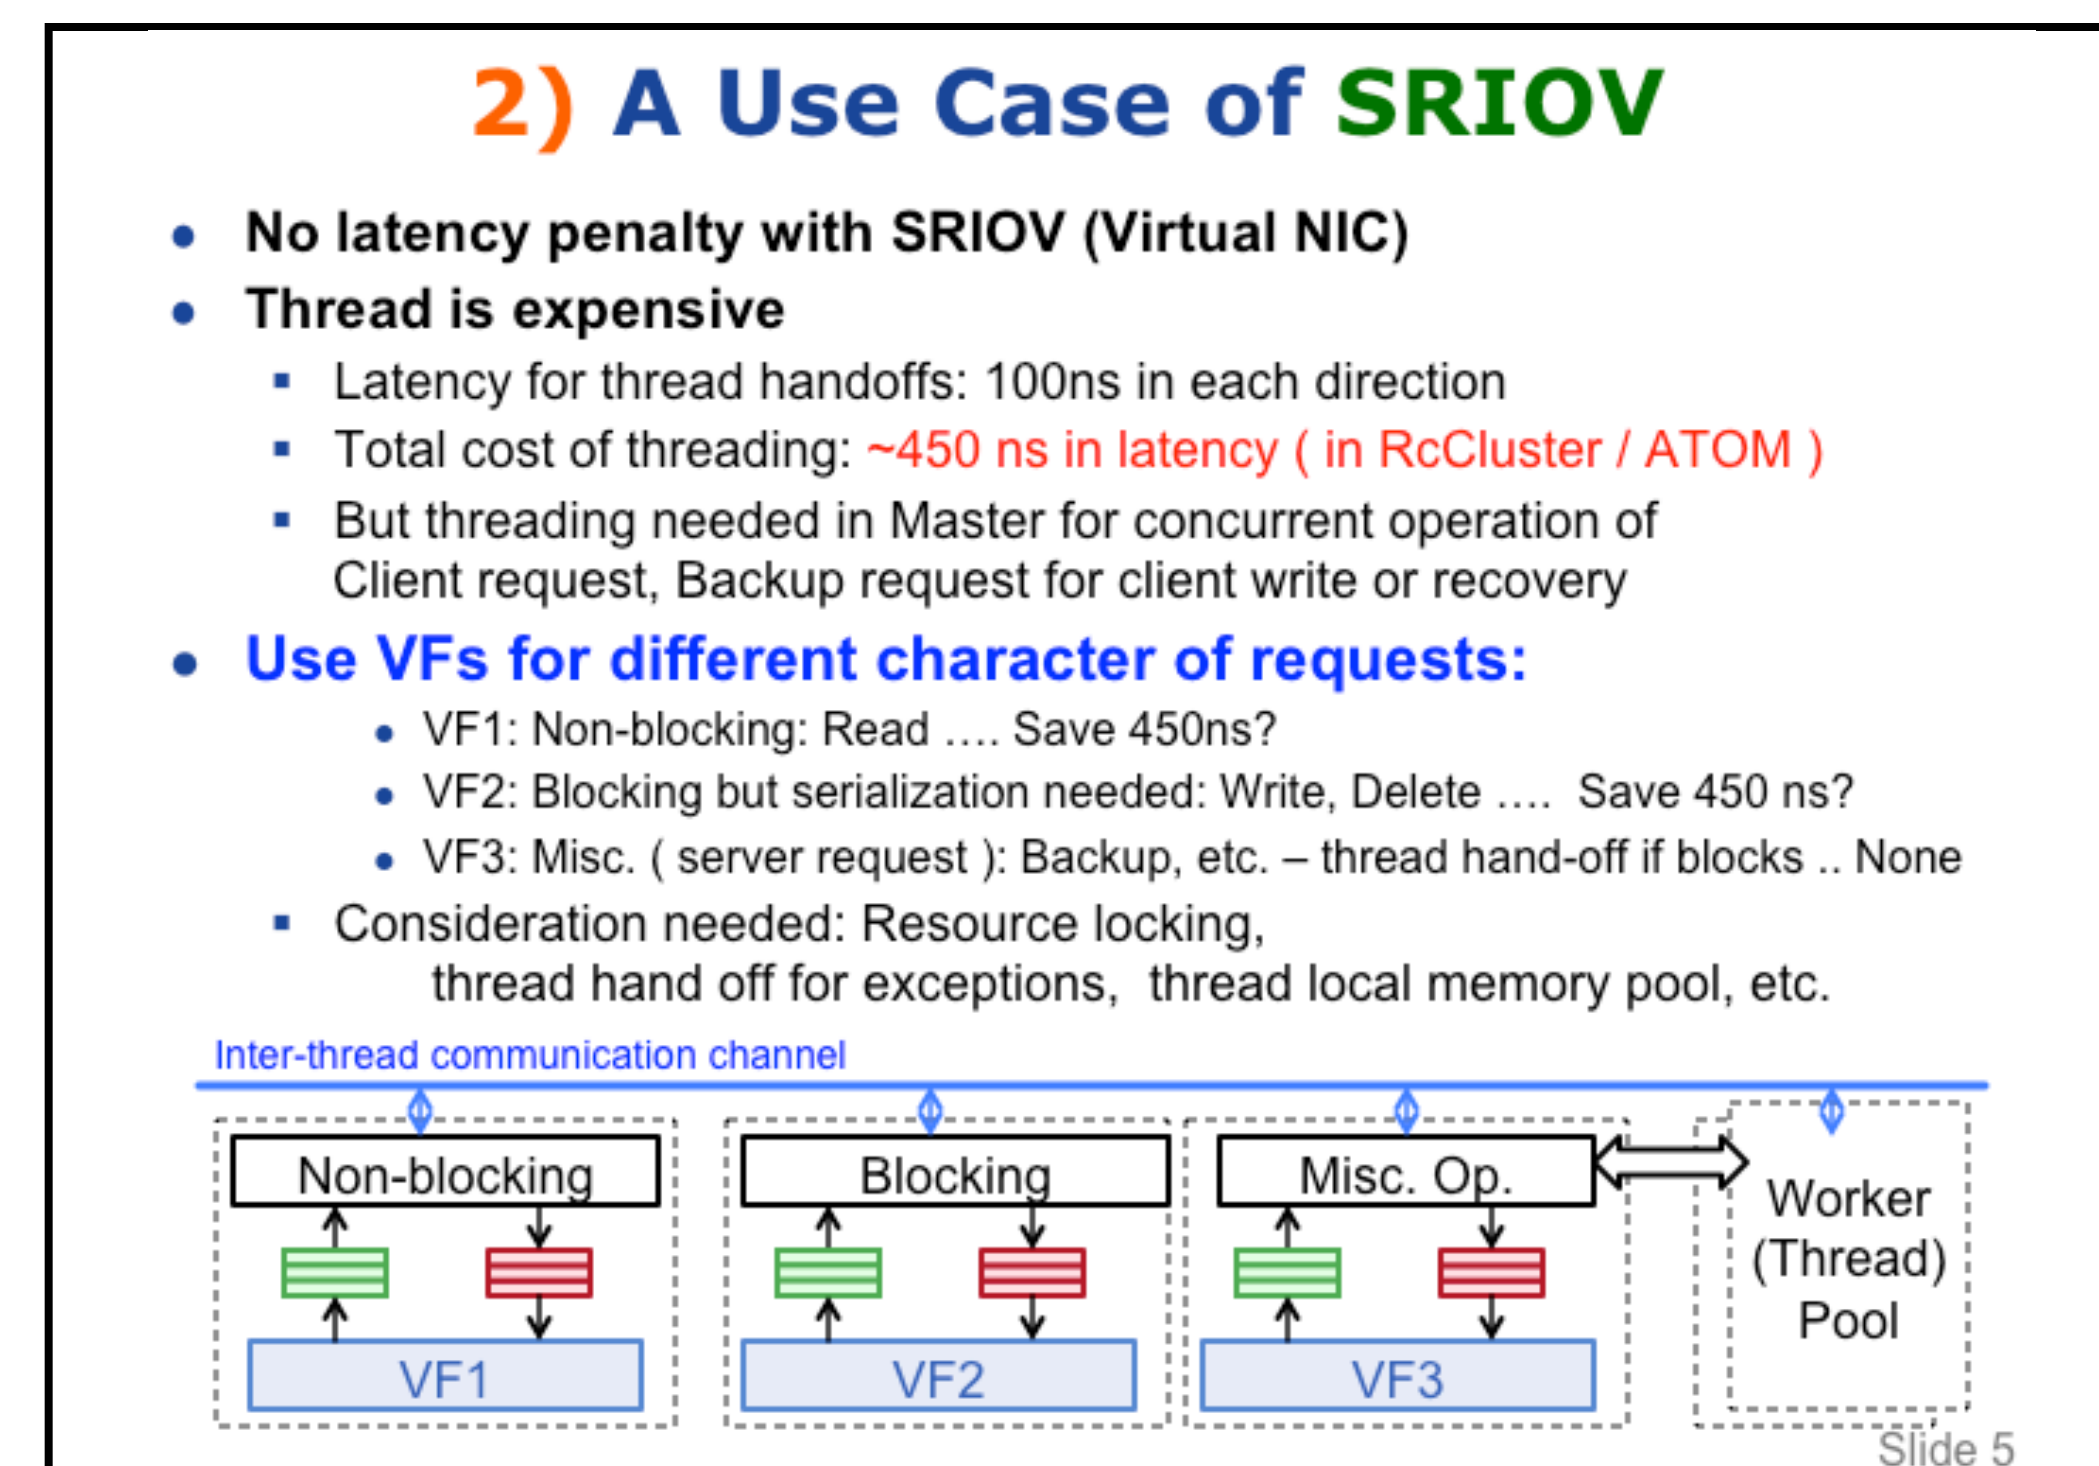
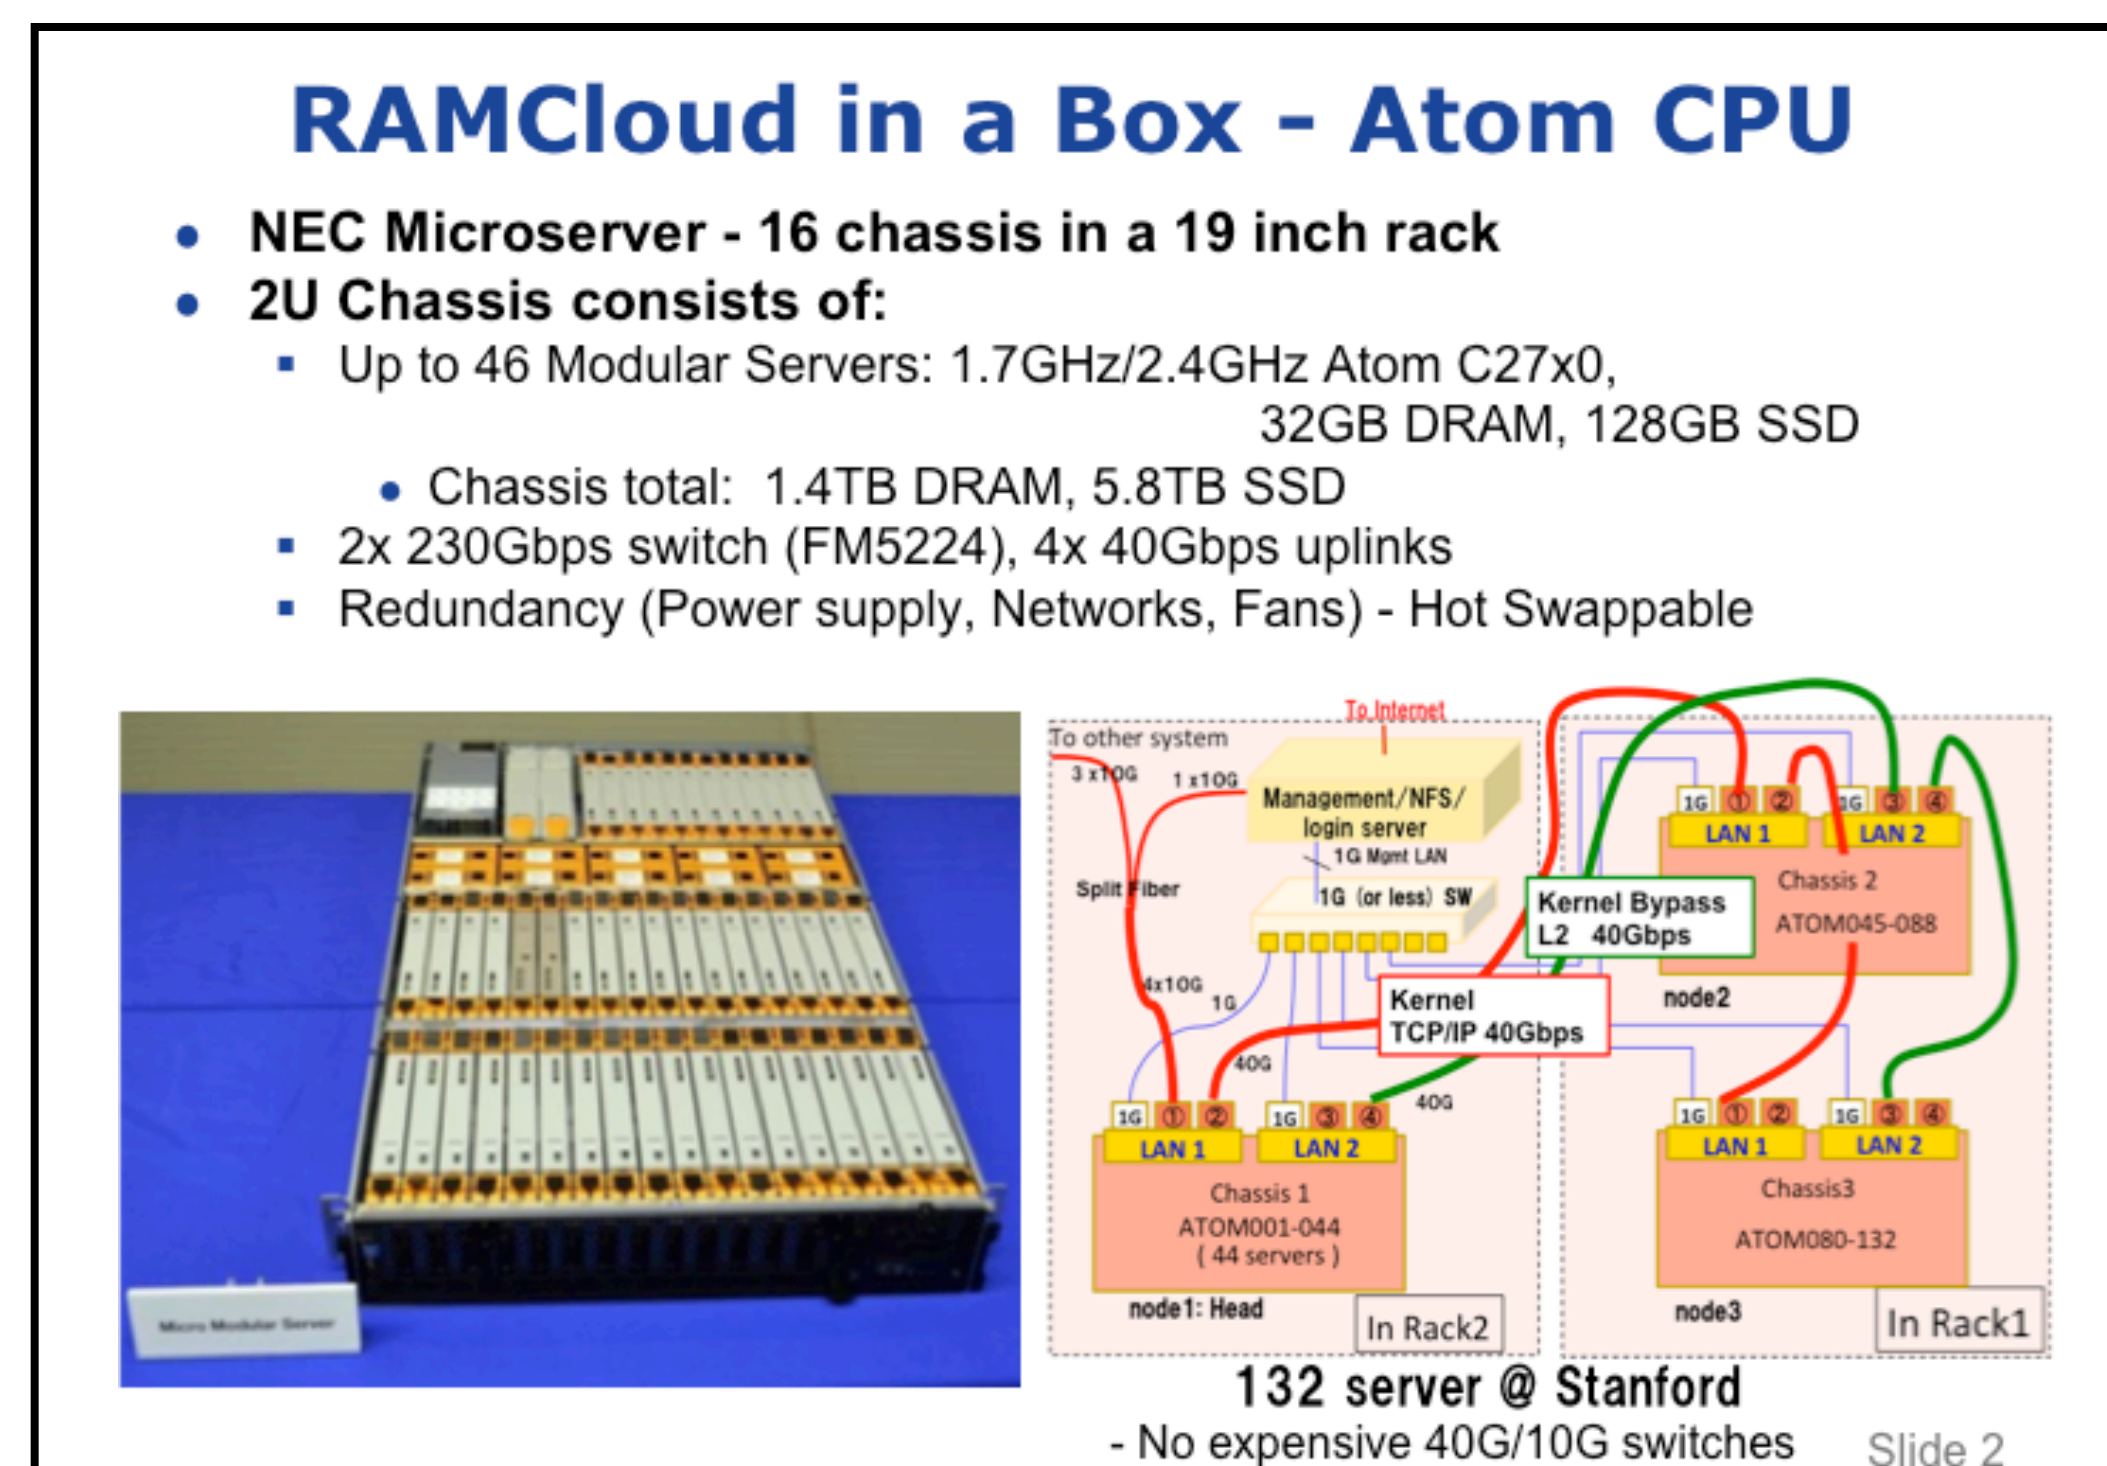
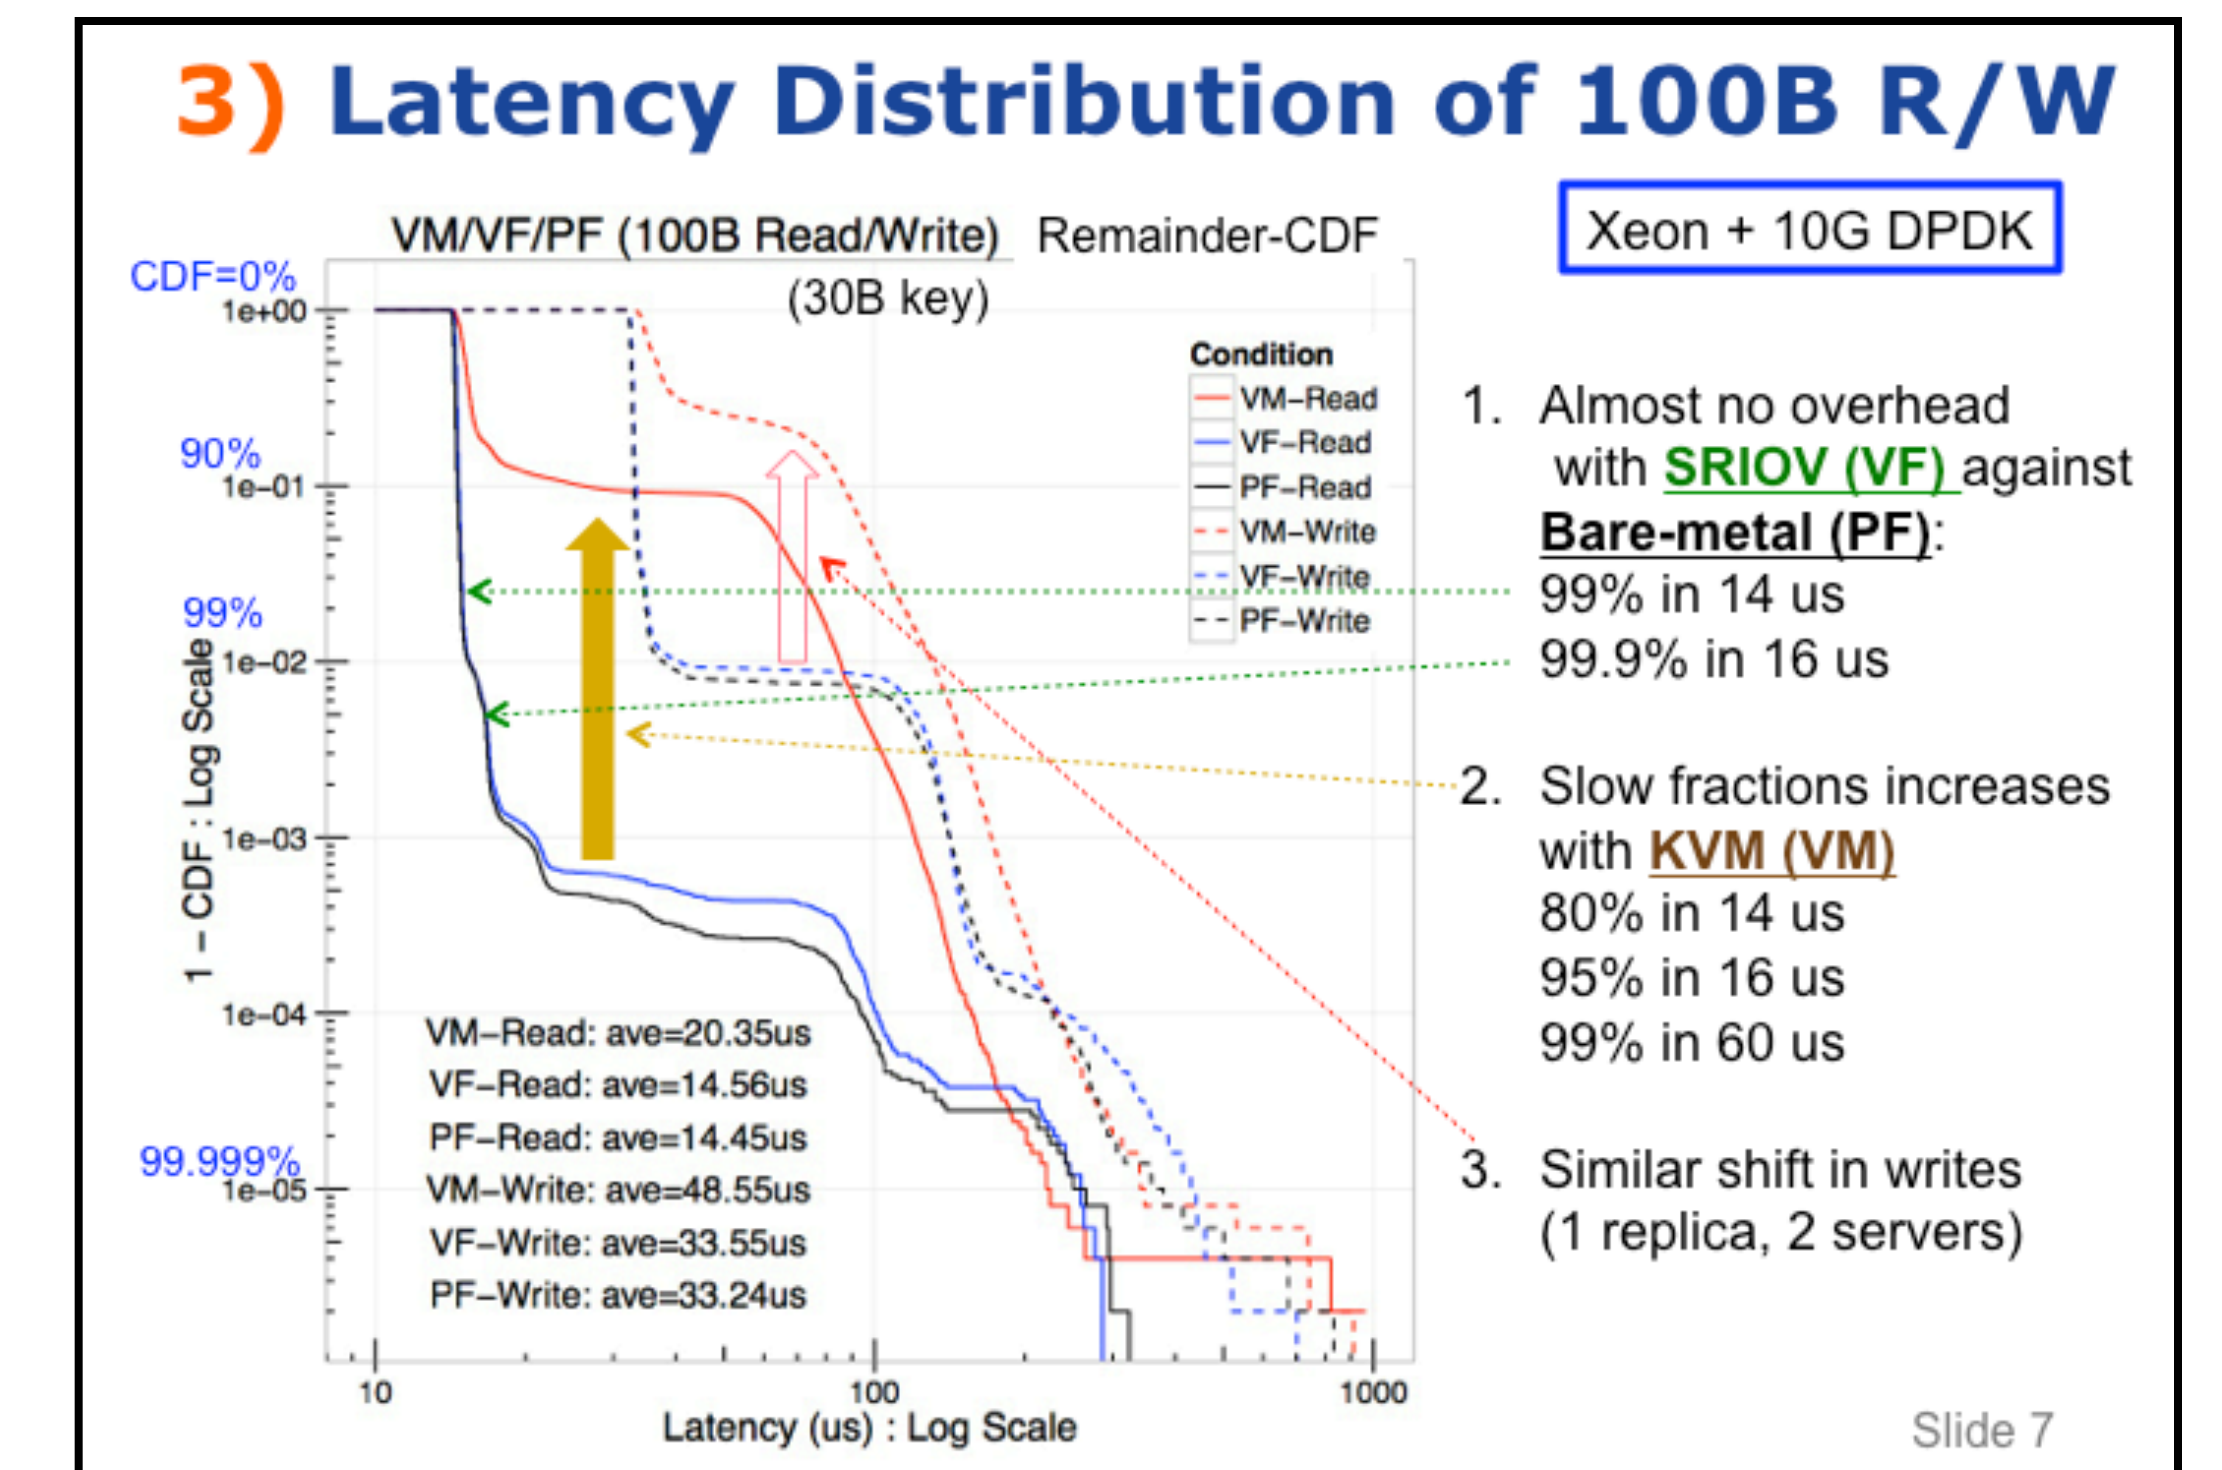
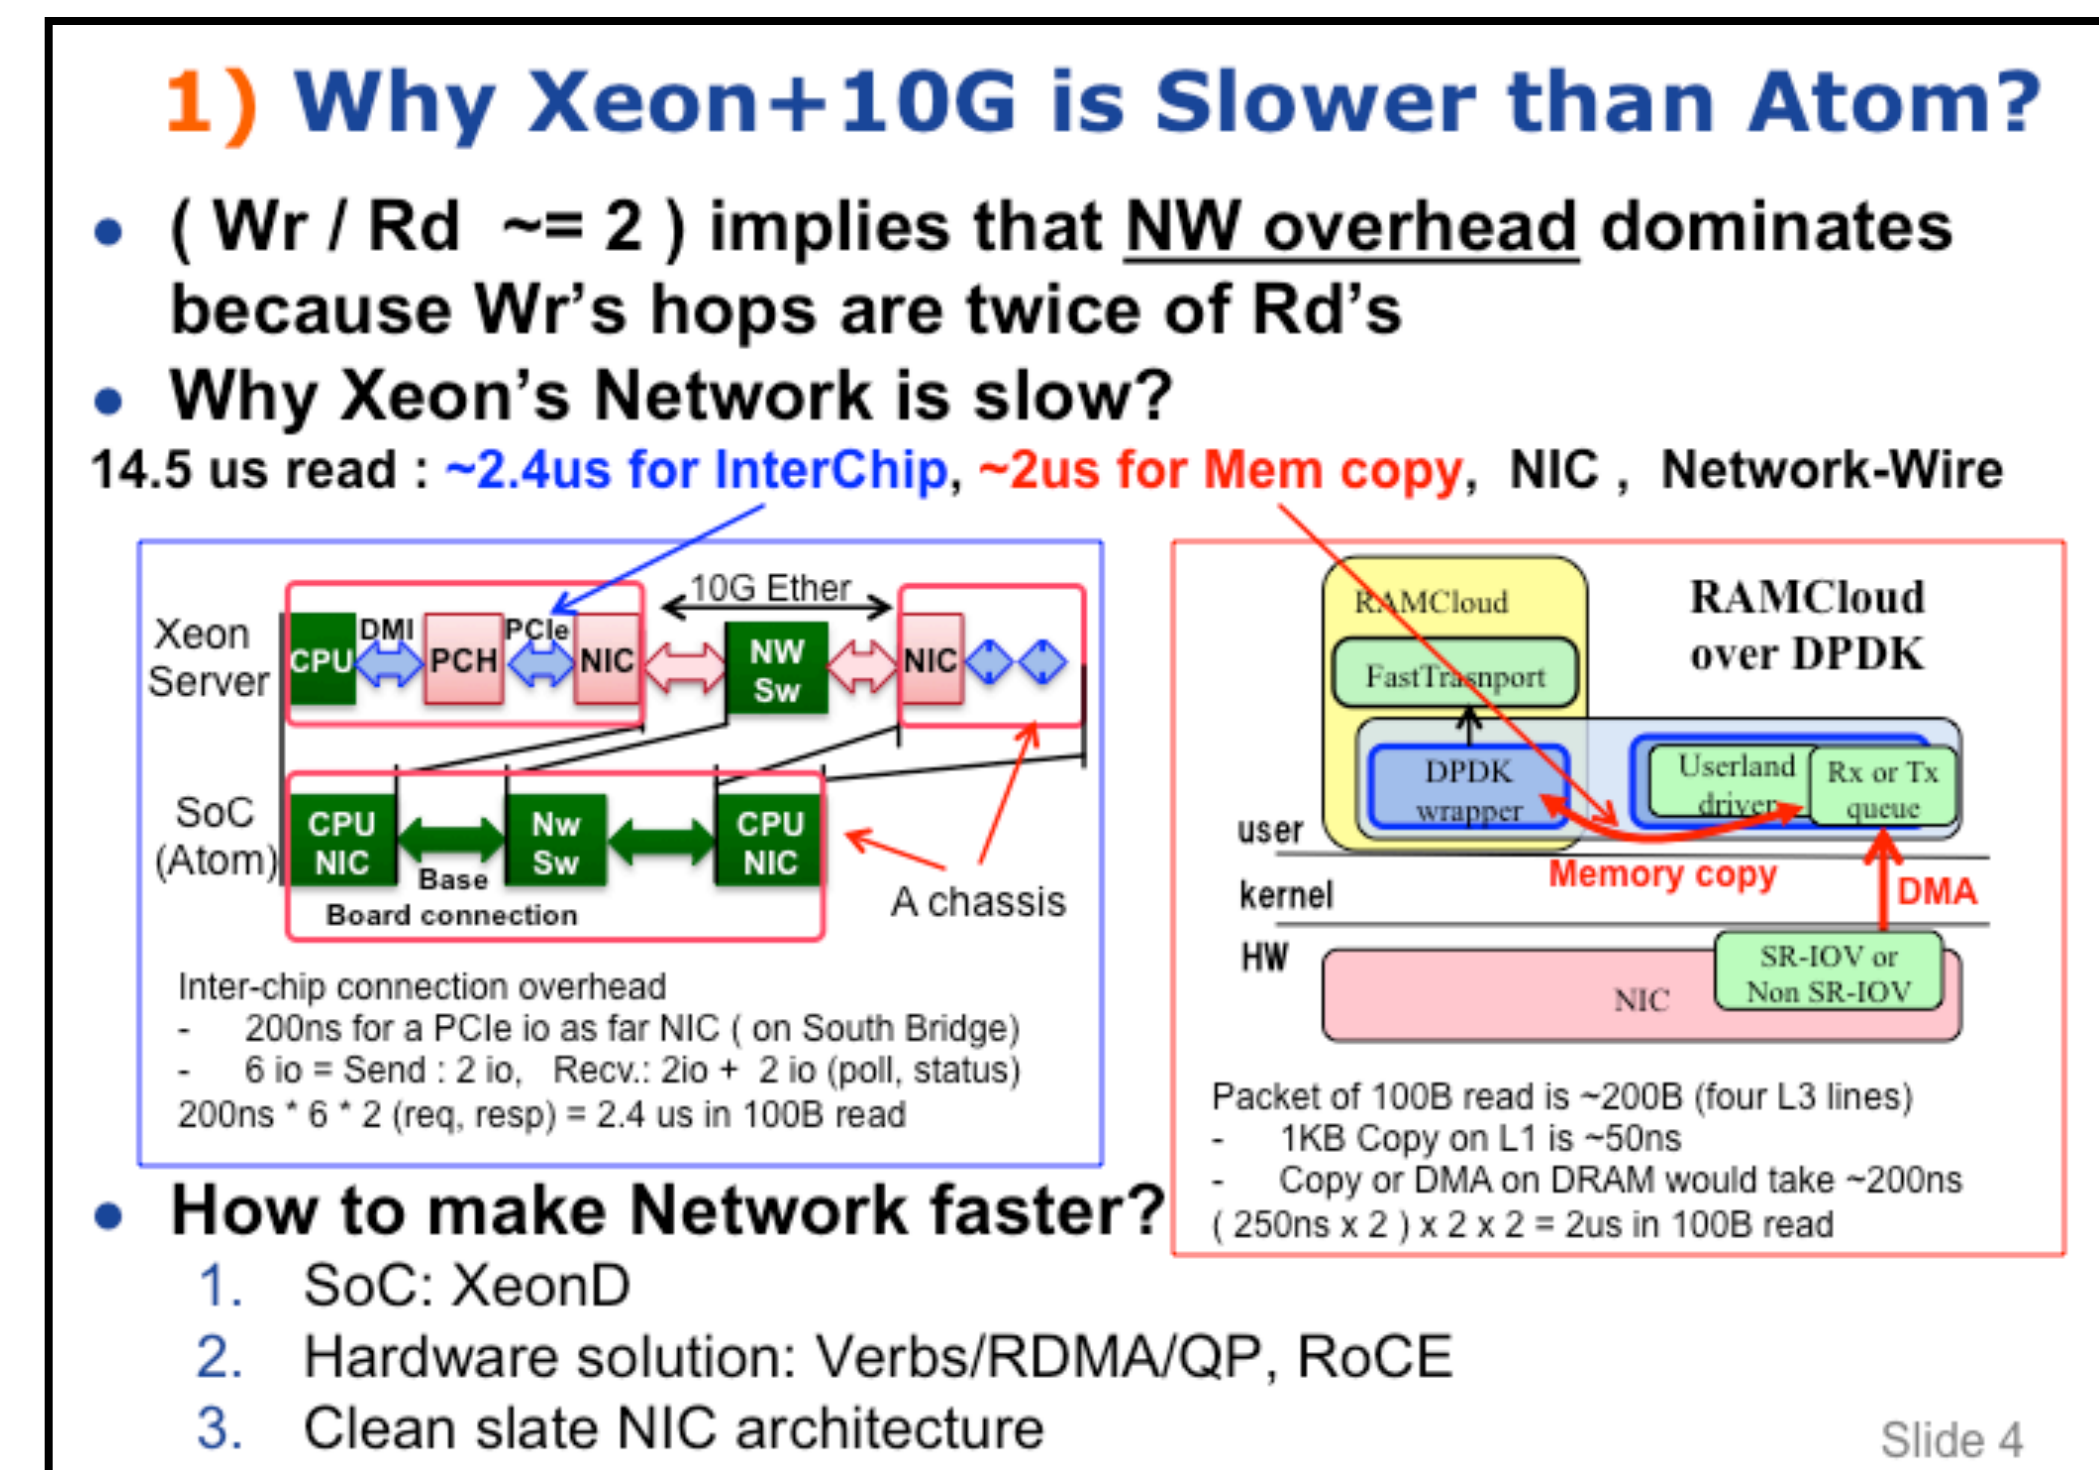
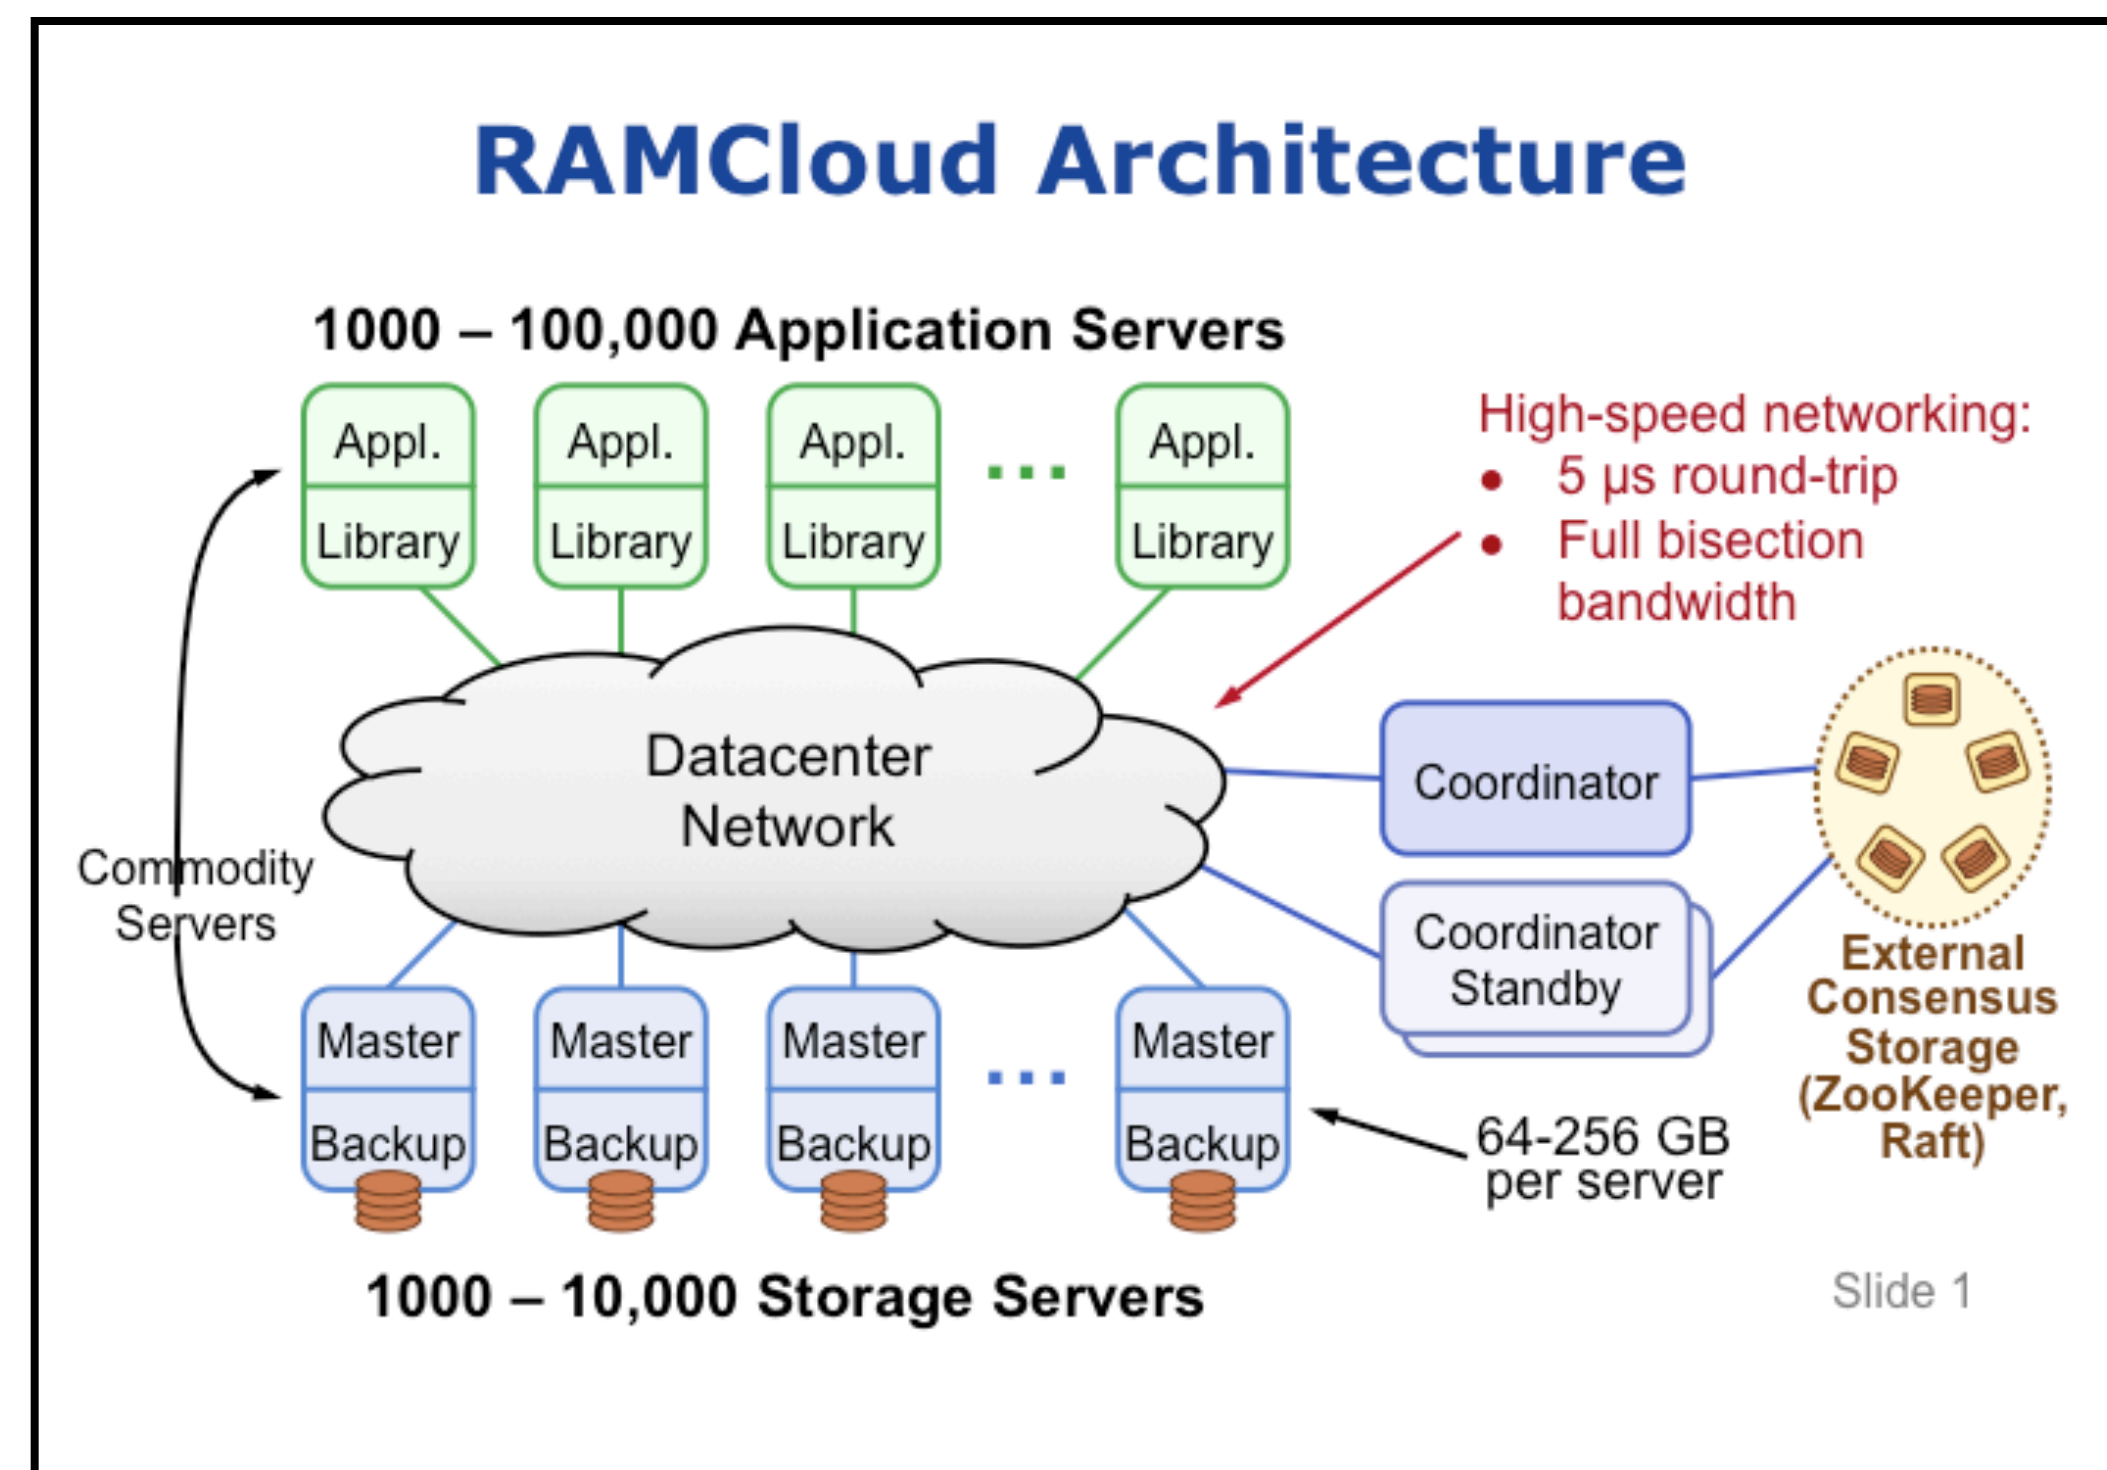
RAMCloud Latency over SRIOV/KVM on Xeon DPK

Satoshi Matsushita^{1,2}, Yoshihiro Hasebe², Kazuo Kuwata³

1) Stanford University, 2) NEC Corporation 3) NEC Solution Innovators, Ltd.



2015 SEDCL Retreat Poster

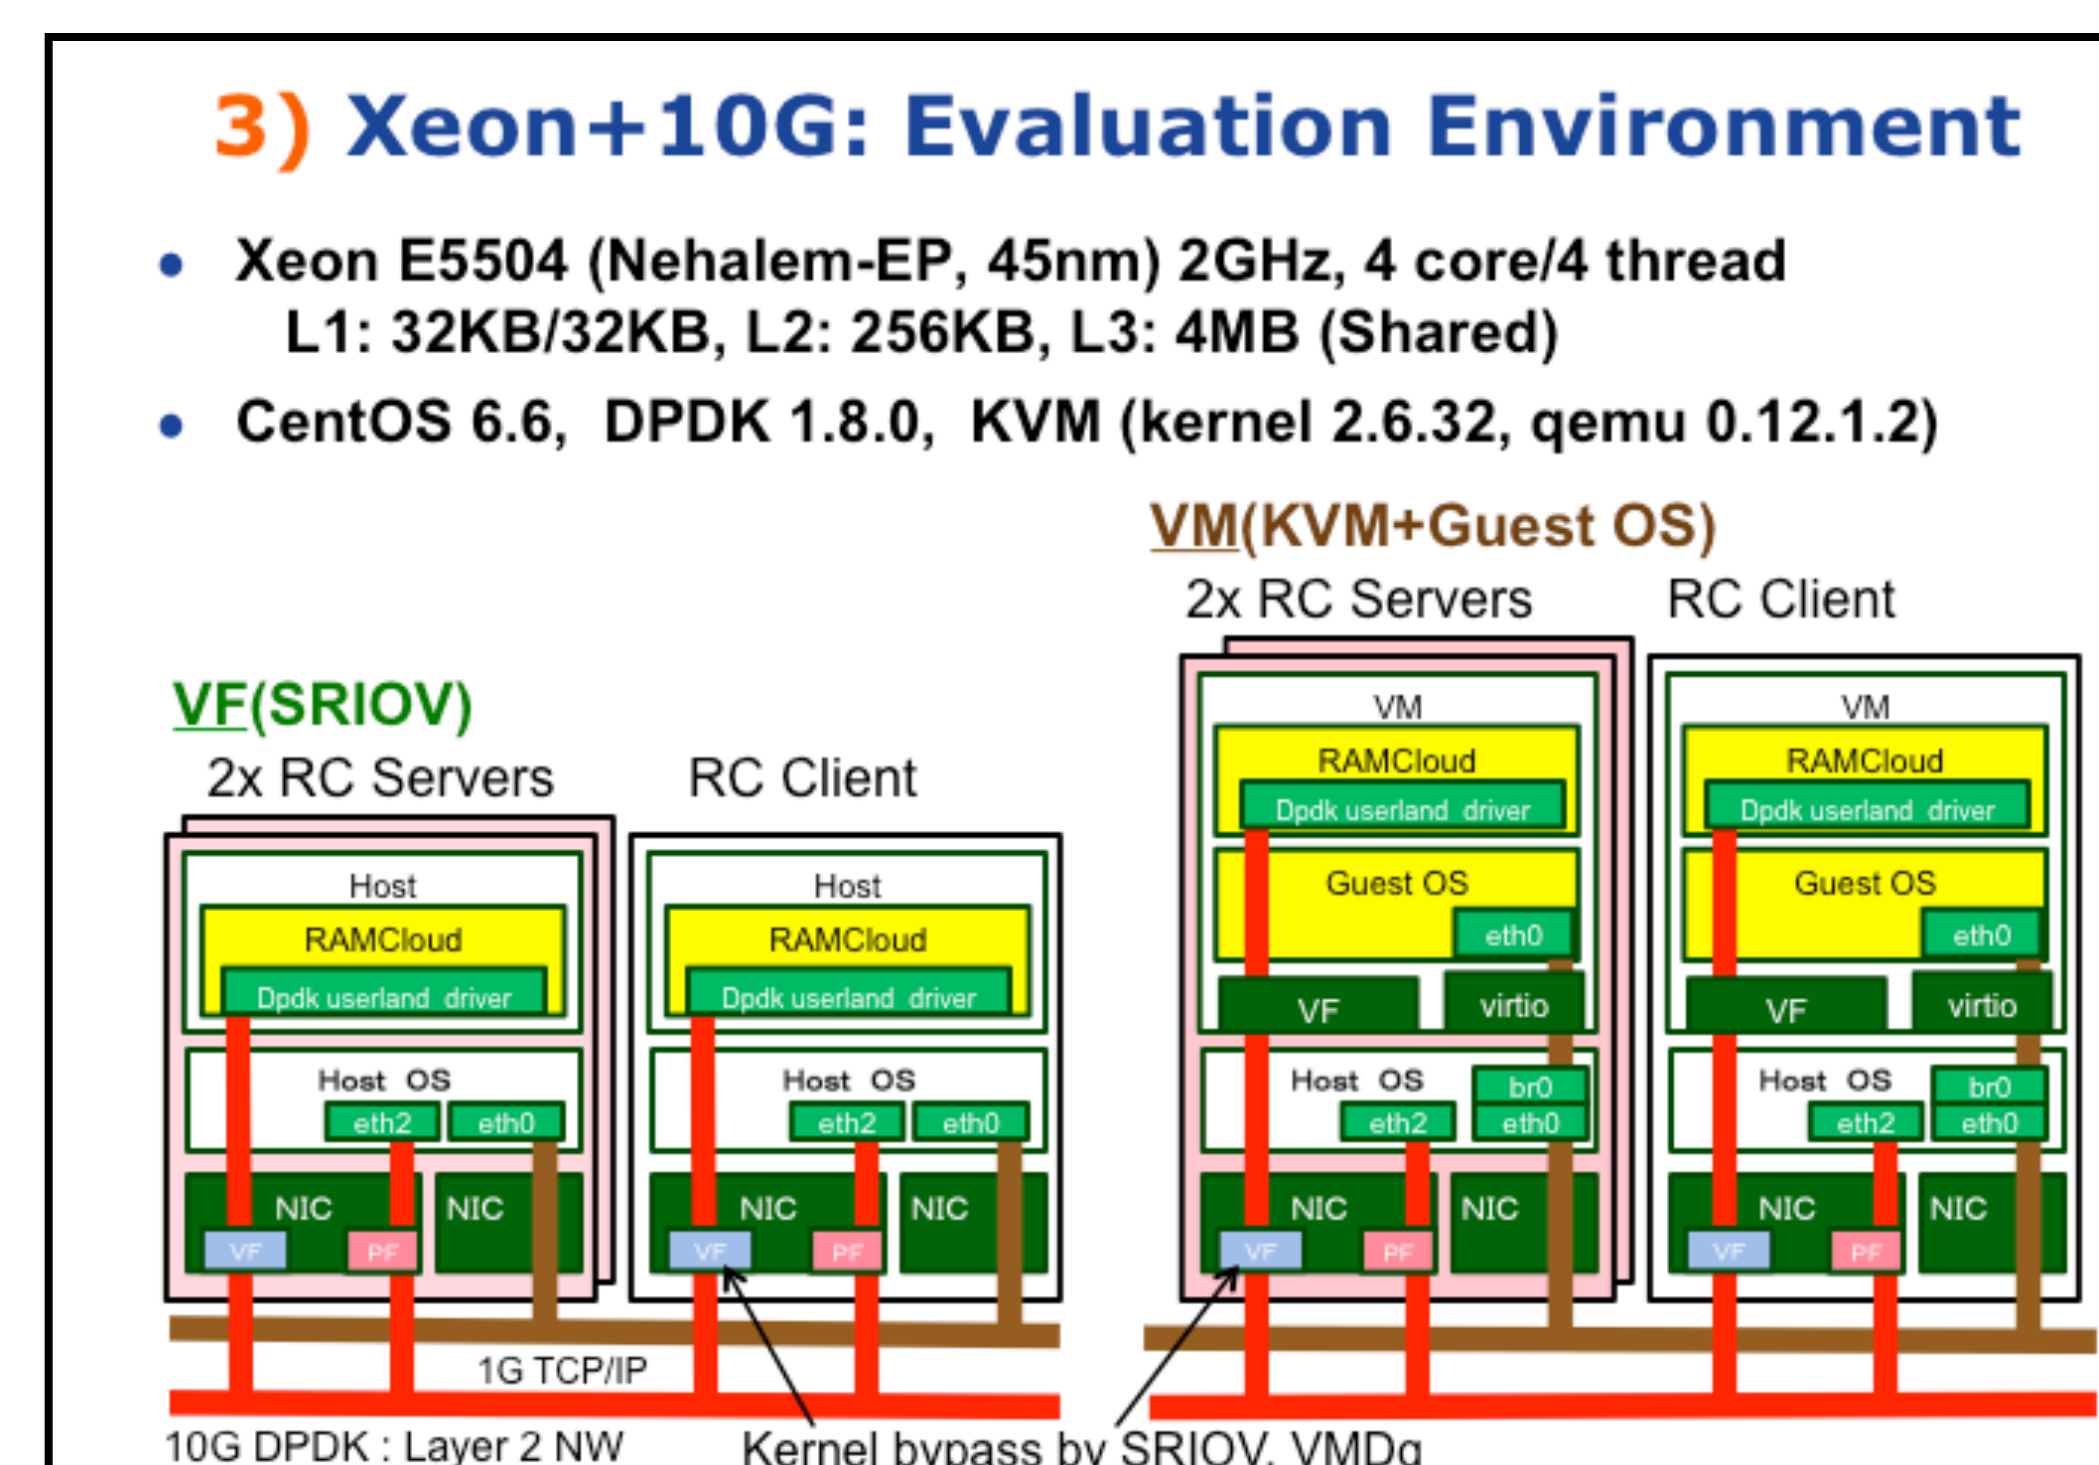


Atom or Xeon

- 1) Atom: Embedded NIC or Xeon: CPU speed
- 2) Features
 - Atom and Xeon: Intel VT-c "virtual NICs": SRIOV, VMDq
 - Xeon: Intel VT-d "Direct DMA" : reduce VM overhead

Result: RAMCloud 100B Read/Write Latency & Ratio \rightarrow Details in 3)

System	Transport	Features	Latency (us)		
			Rd	Wr	W/R
Stanford Xeon	40Gb Infiniband	Verb, RDMA	4.7	14.7	3.1
Solarflare Xeon	10Gb Ether+Fast	Kernel bypass, SRIOV	9.5	23.0	2.4
RAMCloud in a Box (Atom)	2.5Gb Ethernet+FastTransport	Intel DPK, Kernel bypass	11.6	34.7	3.0
Xeon + Intel 10G	10Gb Ethernet + FastTransport	Bare-metal, Kernel bypass	14.5	33.2	2.3
		SRIOV (VT-c)	14.6	33.6	2.3
		KVM Guest OS (VT-d)	20.3	48.6	2.4



- ### 3) Consideration on Latency Distribution
- Cliffs
 - {2.5us, 50-80us, 250us} slower than minimum in Ping/100B Rd. Similar to RcCluster Infiniband.
 - 2.5 us : L3 cache miss + CPU internal / Memory conflict?
 - 50-80 us : NIC or network?
 - 250 us : OS related?
 - 5 ms slower than minimum in VM ping – context switch, process in run queue with nanosleep()
 - SRIOV Overhead
 - Distribution shift derived from NIC/Network
 - KVM Overhead
 - Why all cliffs shift in the distribution?
- Slide 9