Empowered by Innovation NEC

RAMCloud in a Box

Satoshi Matsushita
NEC Corporation / Stanford CS visiting scholar

Hitoshi Takagi
NEC Corporation

Two Design Criteria

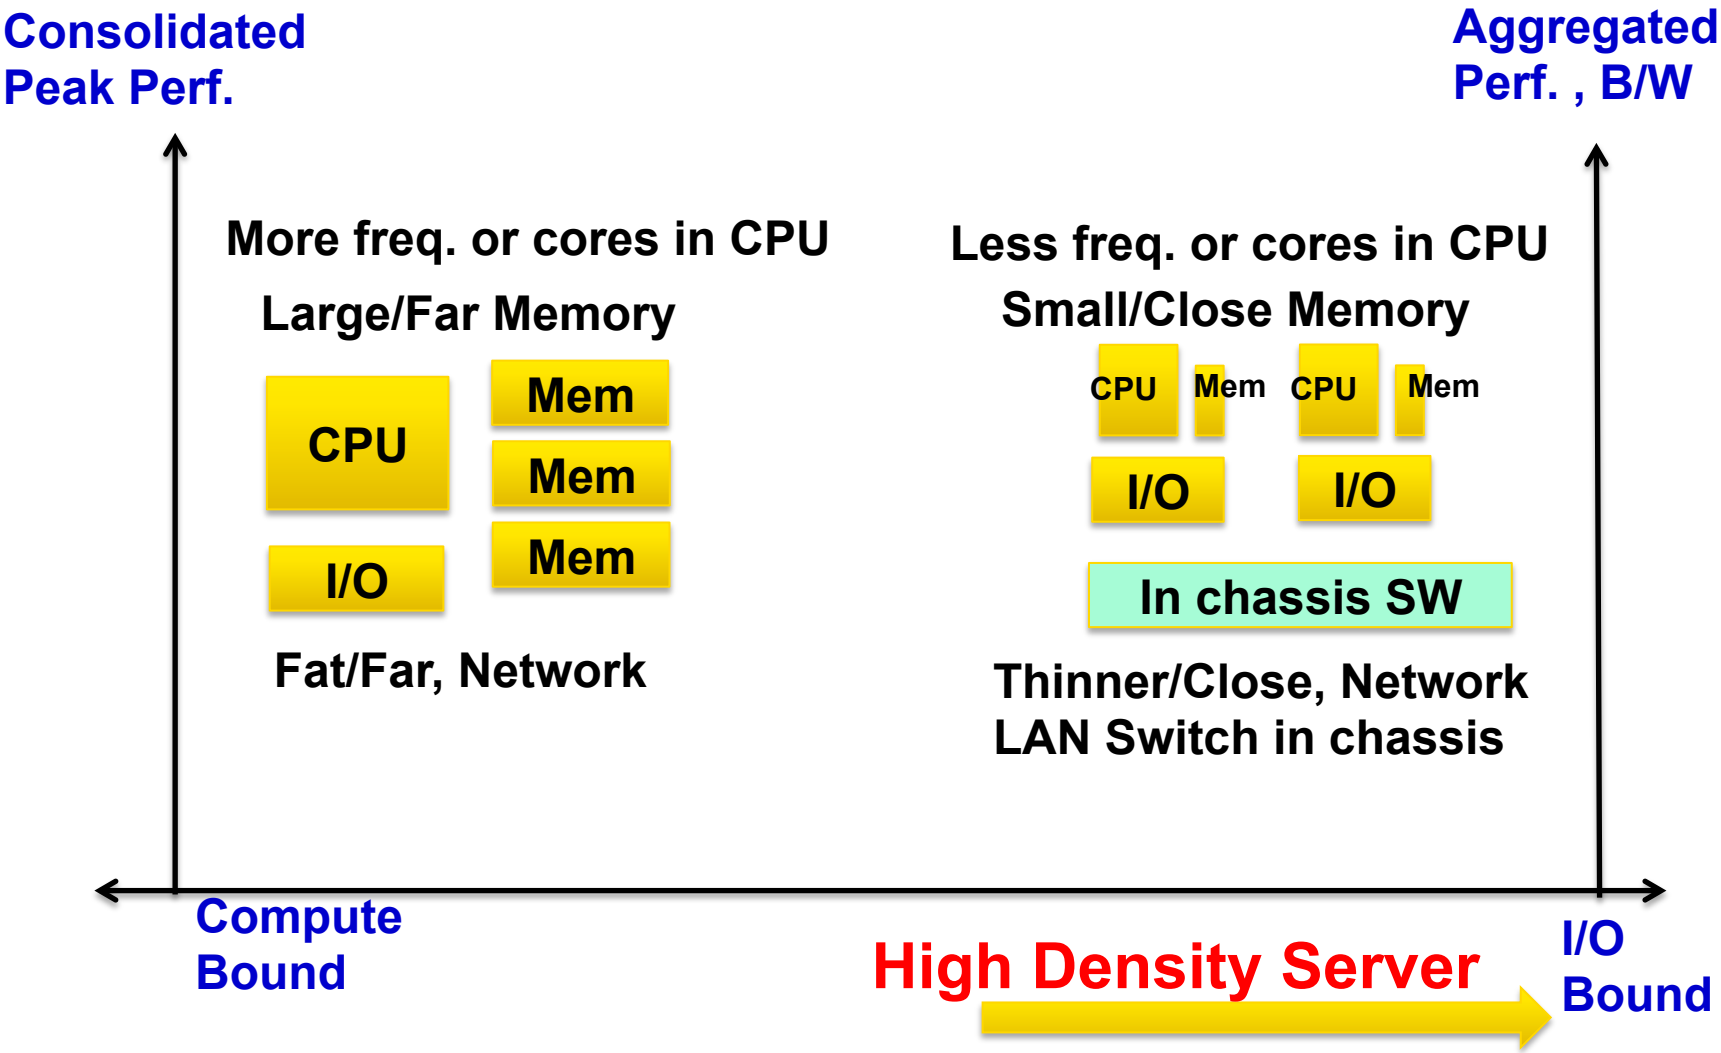

Quantitative Comparison of Two Design Criteria

Compute vs. I/O Bound : getting closer due to power density limitation and process miniaturization: e.g. Dark Silicon

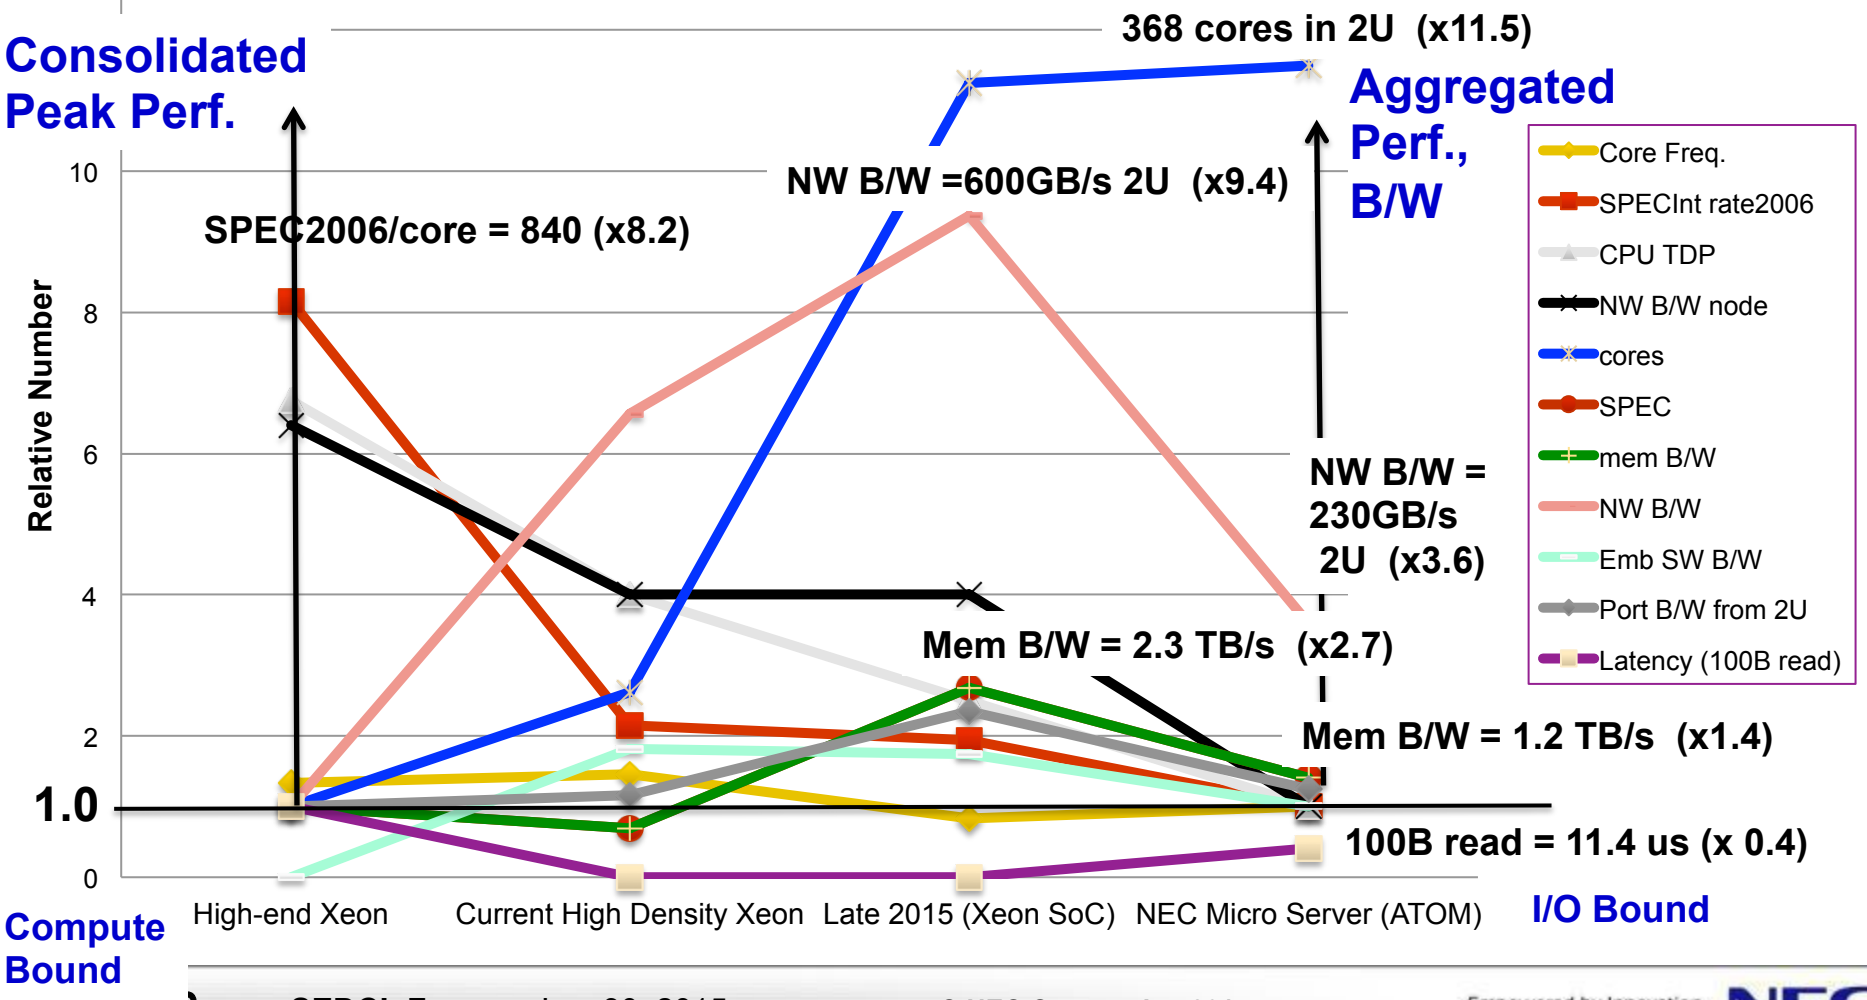

RAMCloud in a Box w. Micro Server (46 ATOMs in 2U)

- Difficulty in setup and management
  1. Both network and server knowledge
  2. RAMCloud specific
- Solution : RAMCloud in a Box
  1. Automated setup for RAMCloud
     1. Maintains continuous IPadrr (hostname)
        … Can isolate broken servers
     2. Applicable to remote management service
  2. RAMCloud source with kernel bypass driver on Intel DPDK, portable to future Intel NICs
  3. Pre-customized service by the host in the box: firewall / sshd / httpd / NFS / NIS / DHCP, etc

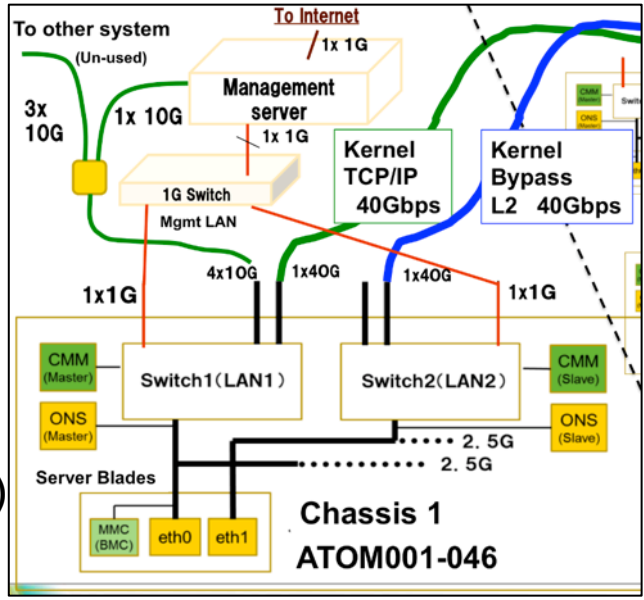

RAMCloud in a Box: Roadmap

- Prototype located in Stanford for evaluation
  - Ready for supplying other application evaluators and academics
- Compatibility in future line-ups:
  - source level compatibility
  - mixed operation

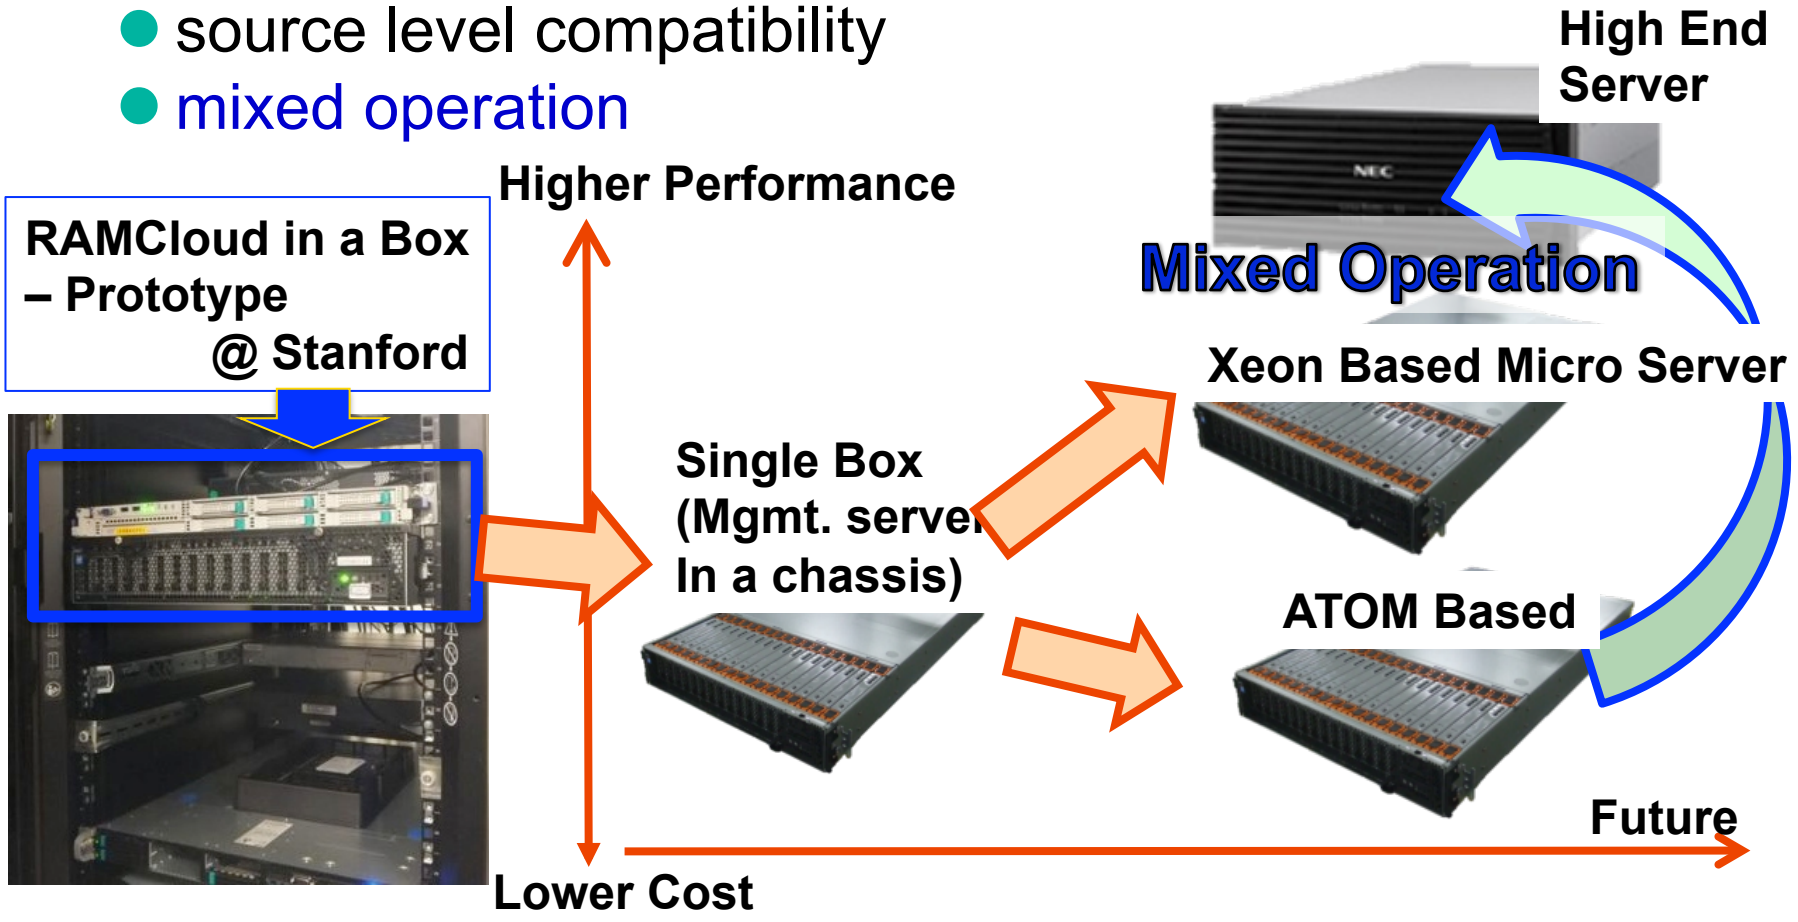

Estimation: Network Performance

- Evaluation on an high density equivalent low power Xeon
  - 2GHz Xeon, Intel 10G Ethernet

Item	Transport	Condition	2GHz Xeon High Density	ATOM Micro Modular server
Ping Latency	Kernel TCP	1 hop SW	24.0 us -25%	32.0 us
	L2 DPDK (kernel bypass)		9.0 us (*) +50%	6.0 us
RAMCloud 100B read		without SRIOV	TBD	11.6 us (**)
		with SRIOV	TBD	TBD
		On VM w. SRIOV, VT-D	TBD	VT-D unsupported

Note) * : Now investigating.
****: Current RAMCloud cluster) 4.7 us with 32Gbps Infiniband**
8 us (68%) out of 11.6 us spent in Intel driver…..

132 server RAMCloud in a Box @ Stanford: Chaining with 80Gbps without external SW

- No need for expensive 40G/10G switches

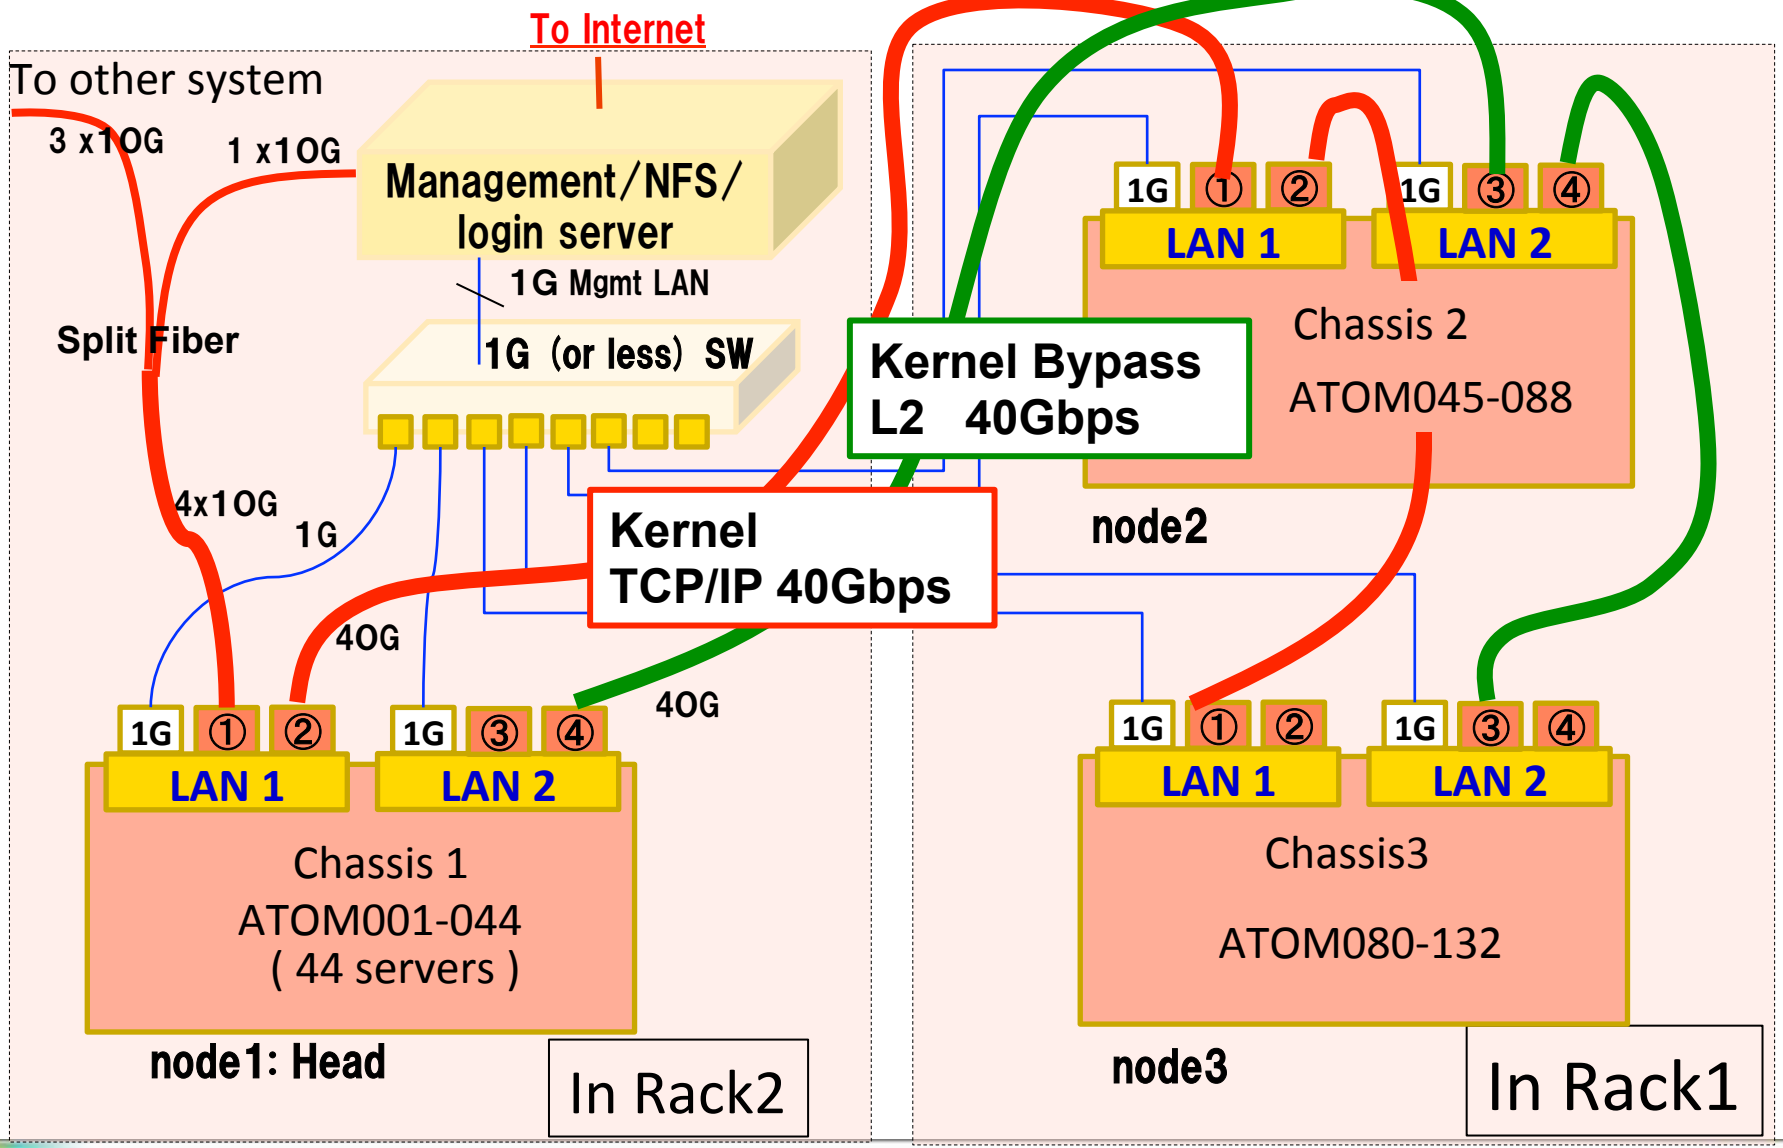

Questions?