

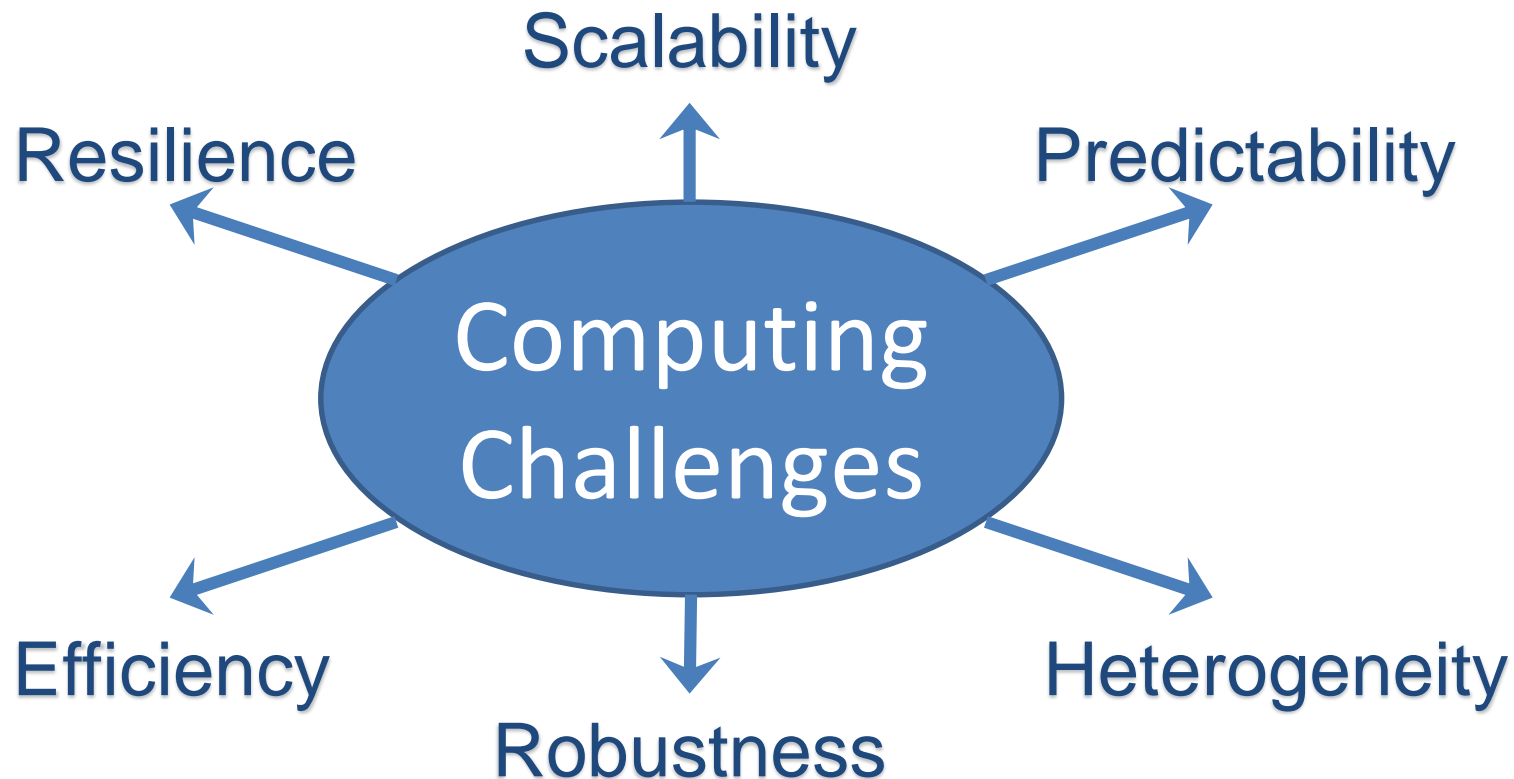
# Bridging The Gap Between Networking And Computing

A vision of future end-host computing

Noa Zilberman  
Andrew W. Moore

April 2015


# Can Networking Answer rack-scale Computing Challenges?



# Can Networking Answer rack-scale Computing Challenges?

...it looks like this.

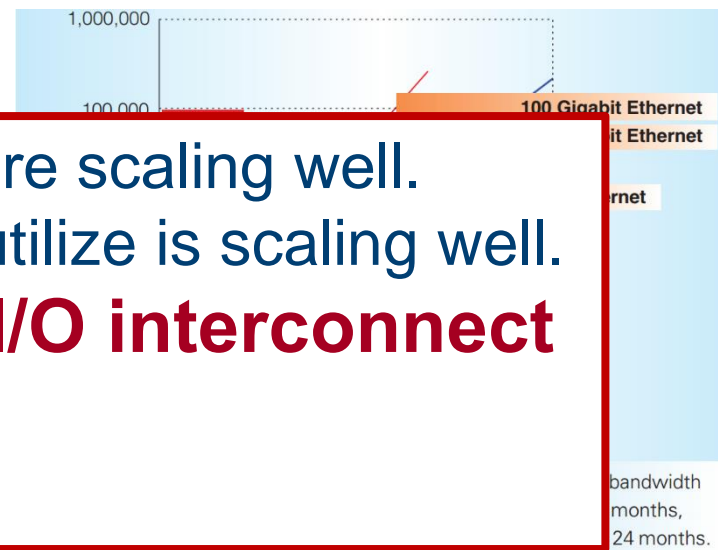
2-socket Xeon E5-2660, Linux 2.6.32, GCC 4.8.0



ORACLE

# The Performance Gap between Networking and Computing

- Core networking bandwidth doubles every 18 months
- Server I/O bandwidth doubles every 24 months
- There is already an order of magnitude gap!



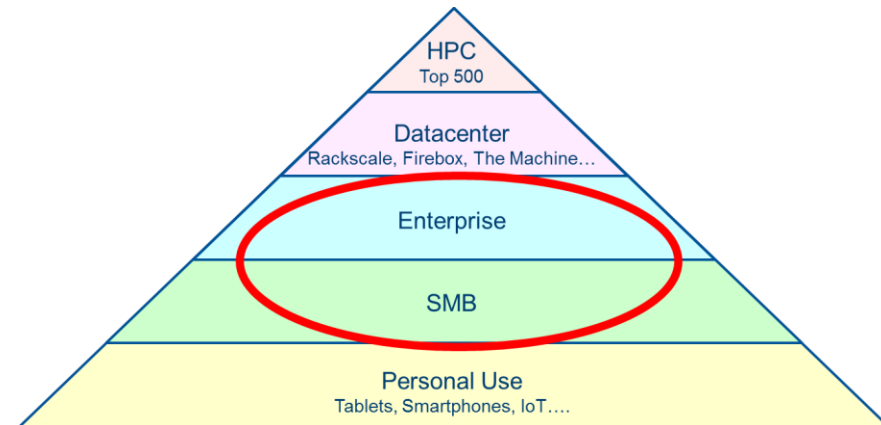
“The CPUs we’re bringing to market are scaling well. Memory bandwidth that those CPUs utilize is scaling well. **What’s not scaling well is the I/O interconnect—the I/O fabric.**”

Barry Davis, Intel, June 2014

(Copyright 2007, IEEE Higher Speed Study Group.)

# Introducing NES: Network Embedded at Scale

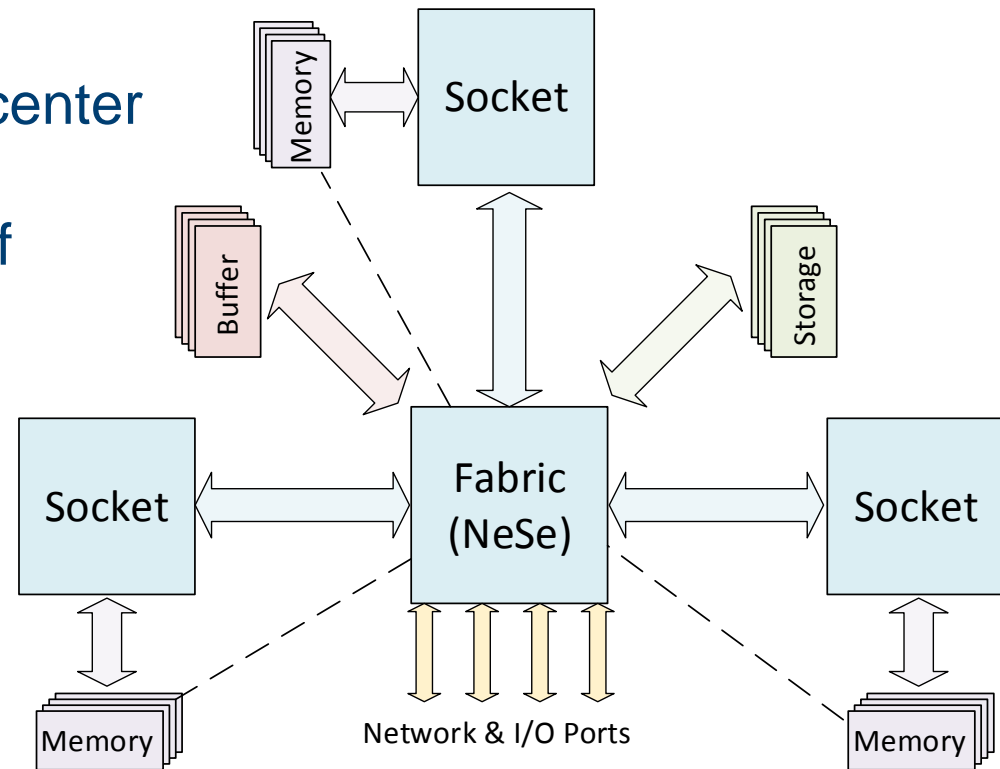
- A server level architecture
  - Applicable from processor level to rack-scale
- Scaling throughput with network-switching performance
- Offering performance guarantees in *hardware*
- Supporting 10K's to 100K's of processes
- For small to large enterprises and research institutes



# The Concept of NES

**Key: treat any transaction in the system as a networking transaction**

- Put a networking fabric at the center of the server
- The fabric connects all types of devices
- Any transaction is annotated with networking properties



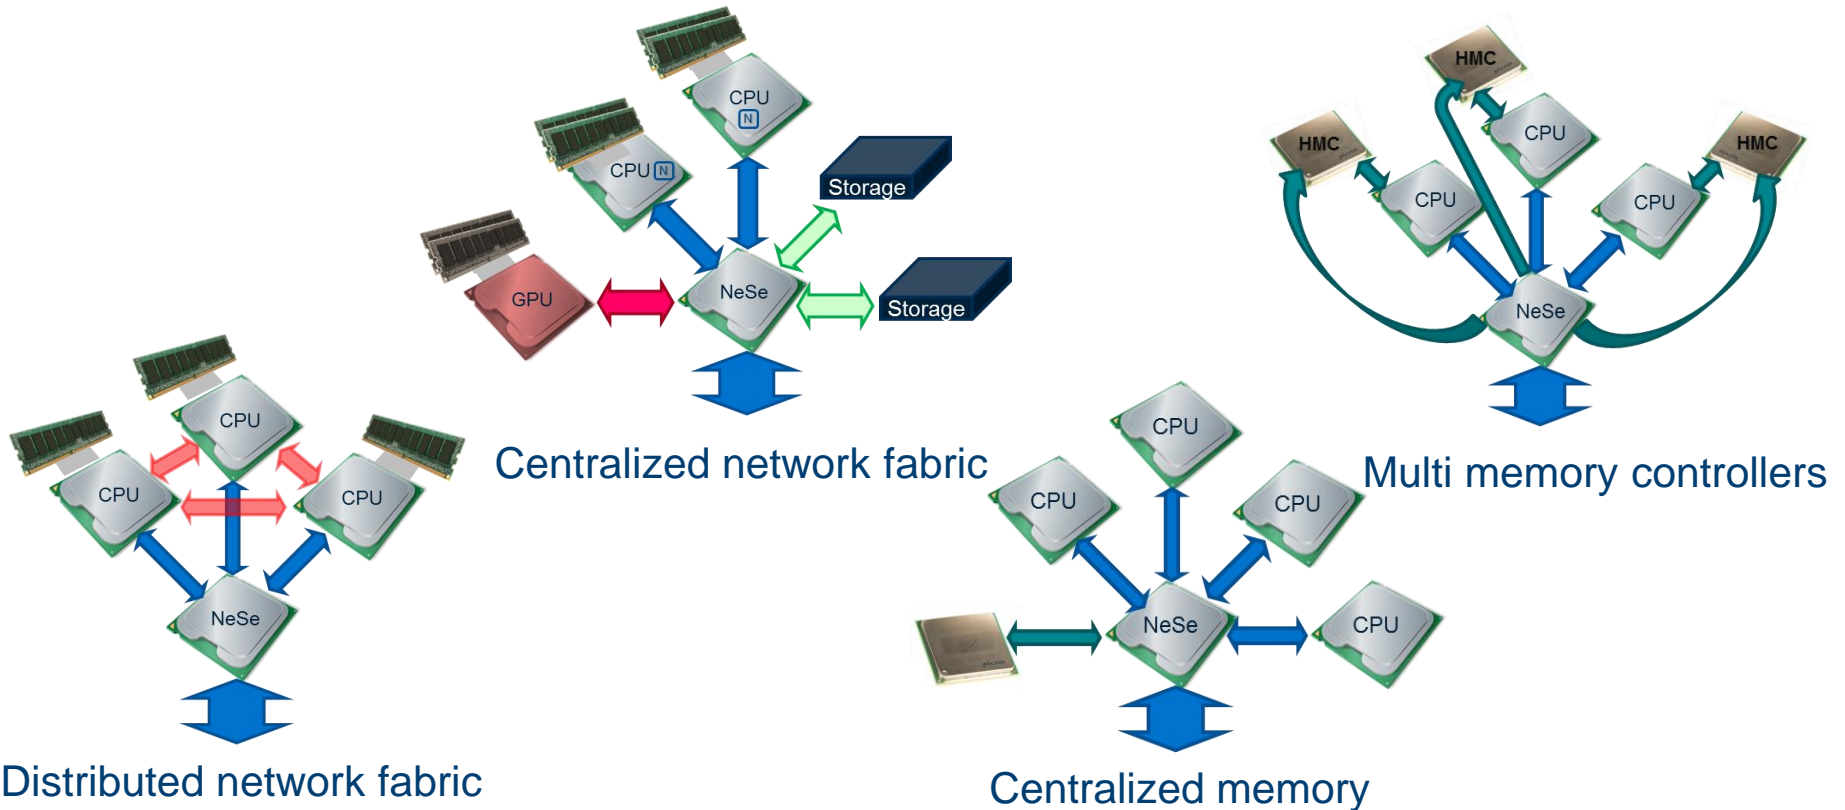
- An integrated HW/SW solution
- A conceptual drawing of a NES-enabled Server

# Properties of NES

- Provides per process: priority enforcement, guaranteed throughput over shared infrastructure, bounded latency, ...
  - Robust
  - Predictable
- Inter-server throughput scales with network-switching performance
- Intra-server throughput scales with computing performance
  - Avoids traffic explosion
- Resilient
- Affordable
- Power-efficient

# Flexible implementation

- NES enables different types of implementations
- An interface-agnostic fabric reduces deployment constraints





# Realizing NES

- A collection of efforts:  
NeSe, OS & Hypervisor support, Interconnect, ...
- Step 3: A fully-customized 1 Tbps NeS server
  - Optical switching, optimized processors, full-blown SW support,...
- Step 2: A NES fabric connecting a collection of commodity servers
- Step 1: Proof-of-concept using NetFPGA-SUME platform, CHERI (soft core) CPU and CHERI BSD

