SSRLabs

High-Performance Dense Hybrid Memory Cube (HMC)-based Flash Memory for Storage Appliances

August 2017

Axel Kloth, President & CEO SSRLabs

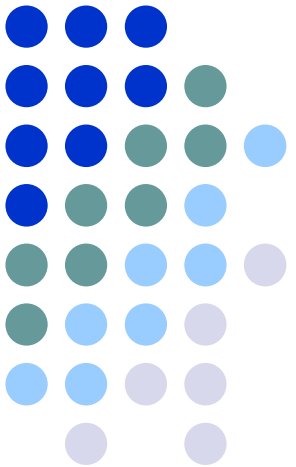

Overview and Motivation

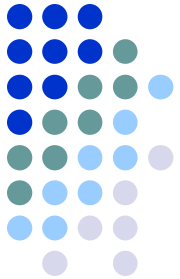

- Our appetite for data is ever-increasing.
- We store more data, retrieve more data, classify more data and pre-process more data.
- The Internet and Data Centers worldwide already use 7% to 14% of global electricity.
- We simply cannot multiply the number of data centers by 10 every 18 months.
- Processing, storing and retrieving data must improve in performance without increased power consumption.

Bandwidths in a Storage Appliance

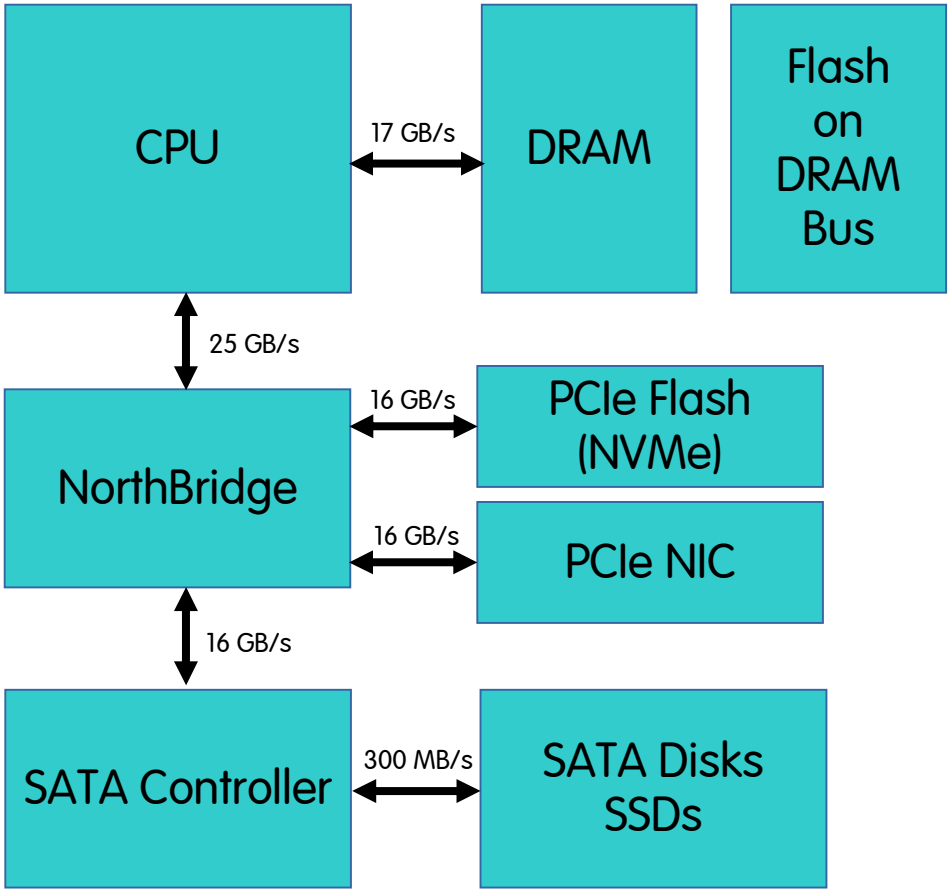

A new approach is needed

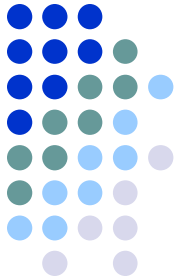

- Flash at Memory Speeds is not good enough.
  - Memory "Speed" = DDR3/4 DRAM "Speed"
  - DDR3/4 DRAM bandwidth: ~ 17 GB/s per channel
- Flash alone is not flexible enough.
  - While Flash is getting more cost effective, it is not yet at the same cost per bit as hard disks
  - What if Flash is not fast/dense/cheap enough?
- Scale-out using PCIe alone is not fast enough.
  - PCIe is limited to ~16 GB/s, 100 GbE is roughly 11 GB/s, SATA-12G is 1.5 GB/s.

Requirements

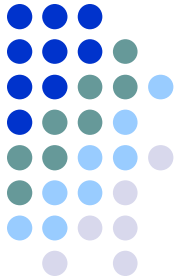

- New interfaces on processor and memory
  - Faster than QPI, DMI, PCIe and SSTL-2 DDR3/4
  - Lower power per bit transferred
  - Fewer pins on the package
- Flash and other bulk memory directly attached to the processor, similar or identical to DRAM but not using DDR3/4 SSTL-2
  - Reuse the serial link concept that was so successful for XAUI, PCIe, InfiniBand and others

New Interfaces

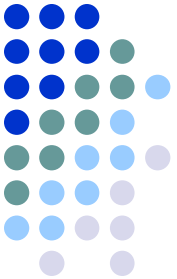

- QPI, DMI and PCIe do not provide enough bandwidth.
- Wider Interfaces? Too many pins!
- Trunked High Speed Serial Links provide a solution.
- More bandwidth, lower pin count.
- Use that Interface for memory and processor-to-processor communication.
- Enable legacy access by providing PCIe lanes.

Scaling Out

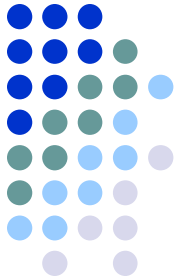

- A single Storage Appliance won't be able to support the demands of thousands of users.
- Scaling out is a necessity.
- 10 Gigabit Ethernet transfers ~1 GB/s – not fast enough to scale out.
- 100 Gigabit Ethernet provides roughly 10 GB/s. Close to maxing out 16 lane PCIe Gen3. Not fast enough…
- 1000 Gigabit Ethernet (yes, that is 1 Terabit Ethernet) provides ~100 GB/s for scaling out. Good enough.

Flash-Based Storage Appliance

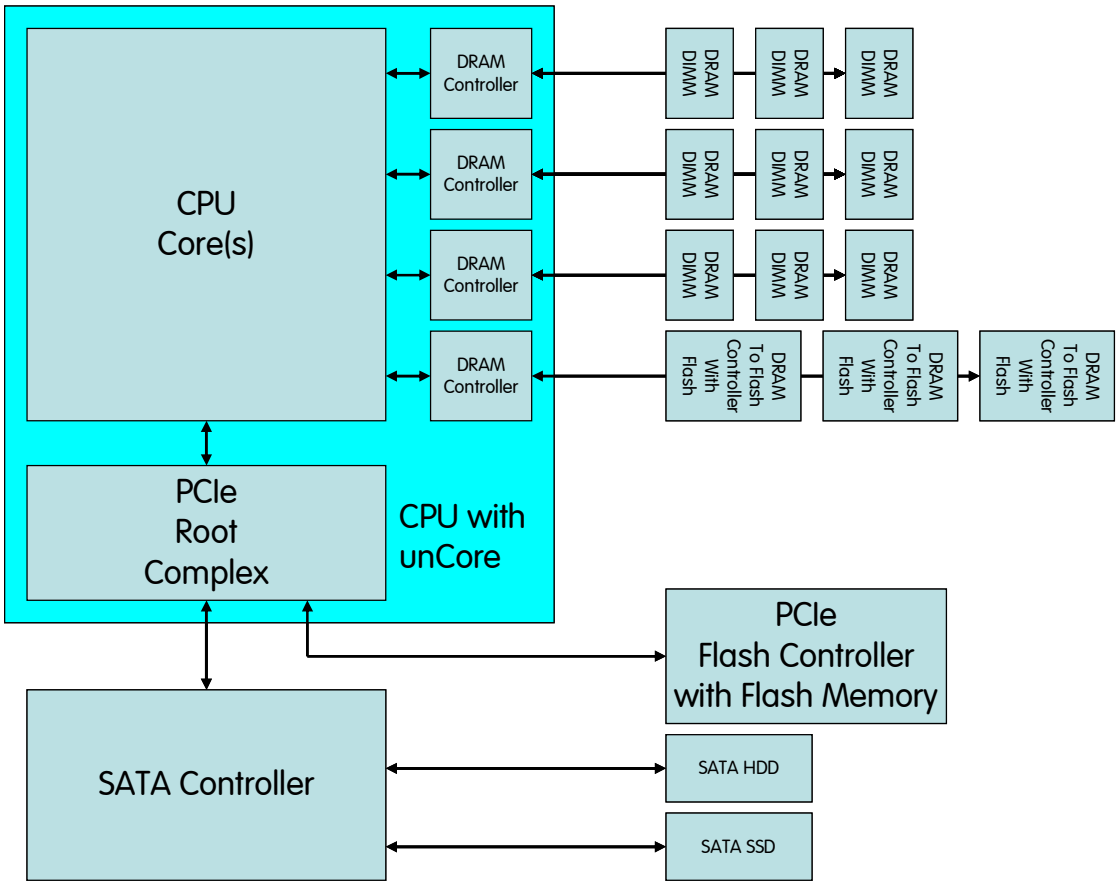

NG Storage Appliance

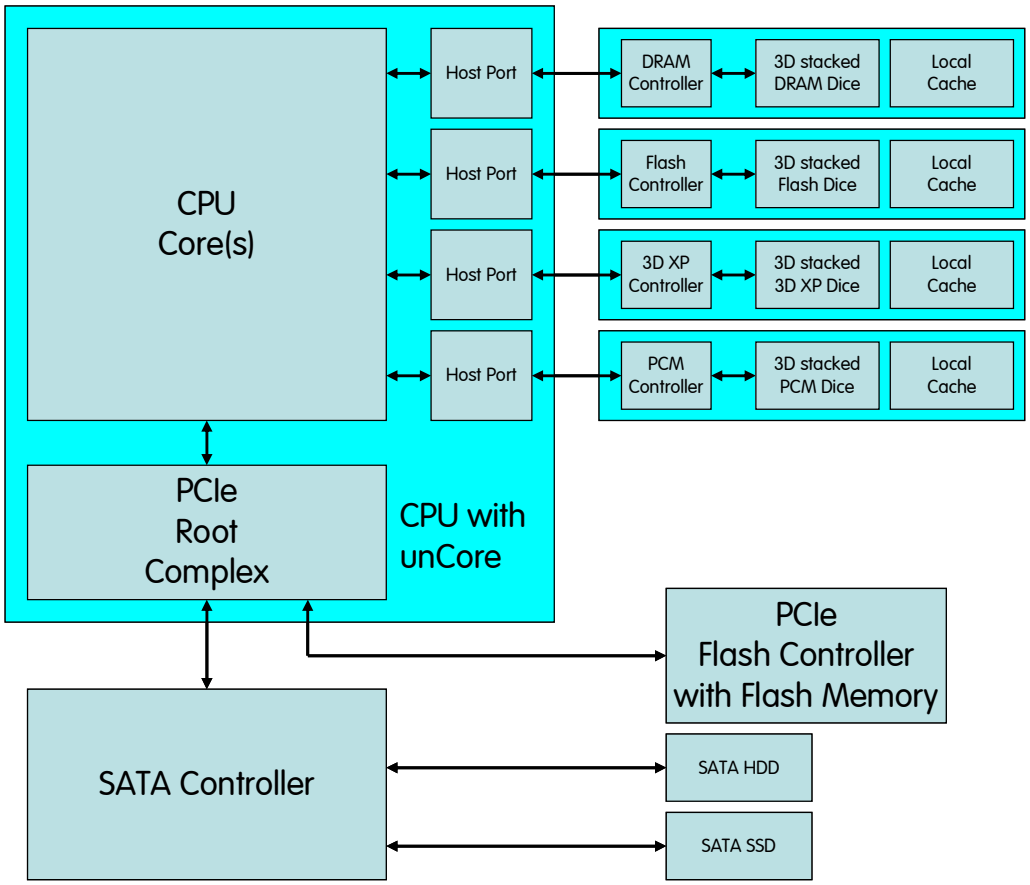

Storage Appliance Comparison

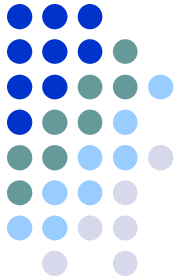

- The NG Storage Appliance looks pretty similar to the current Appliance – at least at first glance.
- That's intended as legacy disks, SSDs, PCIe Flash, NVMe and other I/O can be reused.
- However, performance is very different as the HMC-based UHP provide 60 GB/s of full-duplex I/O.
- We can implement a Storage Appliance with 480 GB/s of storage I/O bandwidth – storage at more than today's DDR3/4 DRAM memory speeds!
- Scaling out through two of those trunked ports to 1 Tbit Ethernet is possible.

NG Storage Appliance Overview

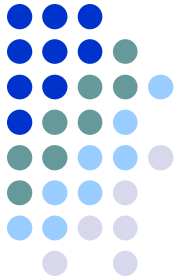

- Storage Processor
  - Any CPU that provides HMC superset ports can be used
  - Currently in talks with two processor companies
  - Ideally integer-only, multi-core
  - Memory port is independent of processor architecture
  - Could re-design Intel, ARM, others
- HMC superset memory (vlcRAM)
- 8 HMC superset ports, 3 or 4 vlcRAM per port
- Total memory/storage bandwidth: 480 GB/s FDX
- Total memory/storage size: 12 or 16 TB

HMC and HBM/HBM2

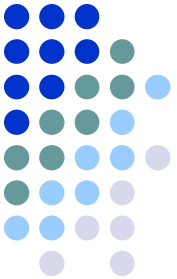

- We keep being asked why we don't use HBM or HBM2.
- The answer is simple: Both HBM and HBM2 are essentially L4 Caches as they must reside on the same multi chip module as the CPU or GPU. They require around 3000 connections.
- In other words, HBM and HBM2 are not for externally attached memory, and densities are not enough for mass storage or bulk DRAM replacement.
- HBM and HBM2 complement our HMC superset, but they do not replace each other.
- HBM and HBM2 very likely won't exceed 16 GB in size.
- Our HMC-based superset provides 512 GB.

Contact

SSRLabs

Scalable Systems Research Labs Inc.

Energy- and Instruction-efficient
High Performance Computing

Ashok Babbar

Chief Executive Officer

458 San Mateo Avenue
San Bruno, CA 94066
email: ashok.babbar@ssrlabs.com
Office: (650) 322-6507
Cell: (408) 464-0015
www.ssrlabs.com

SSRLabs

Scalable Systems Research Labs Inc.

Energy- and Instruction-efficient
High Performance Computing

Axel Kloth
Founder
Chairman & CTO

458 San Mateo Avenue
San Bruno, CA 94066
email: axel.kloth@ssrlabs.com
Office: (650) 322-6507
Cell: (650) 283-3554
www.ssrlabs.com

Disclaimer

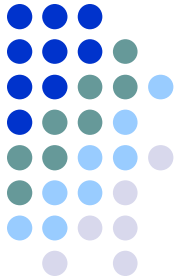

THIS PRESENTATION IS PROVIDED FOR INFORMATION PURPOSES ONLY

AND ITS CONTENTS WILL NOT BE UPDATED.

IT MUST NOT BE USED AS AN OFFICIAL REFERENCE.

FOR TECHNICAL DATA OR OPERATIONAL PROCEDURES, PLEASE REFER TO THE RELEVANT SCALABLE SYSTEMS RESEARCH LABS INC. ('SSRLABS') DOCUMENTATION

The statements made herein do not constitute an offer or a technical specification.

They are based on the assumptions shown and are expressed in good faith. Where the supporting grounds for these statements are not shown, the Company will be pleased to explain the basis thereof. This document is the property of Scalable Systems Research Labs Inc.