



State of the Art Performance in Flash Storage Arrays

Steve Cooper – CEO
One Stop Systems

Santa Clara, CA
August 2016



State of the Art – 2016

- Highest performance
 - Significant improvements versus 2015
- Highest capacity and density
 - About the same as 2015



Two Types of Flash Arrays



- Direct Attached – JBOF
 - Communicates to host server(s)
 - PCIe over cable interface
 - Host server provides NAS functions and interfaces



- All-in-one – NAS
 - Communicates to other user computers
 - Ethernet and/or Infiniband interfaces
 - All-in-one package containing server motherboard and flash array



Direct Attached vs All-in-One

Flash Array


Server

Flash Array
Server Motherboard

- Direct Attached
 - Simpler – extension of server slots
 - More flexible, pick any server, any S/W
 - Server can run application S/W
 - Server and flash de-coupled
- All-in-one – NAS
 - More complete, ready to run
 - Pre-configured hardware and S/W
 - Factory optimized
 - Less flexible, all-in-one packaging



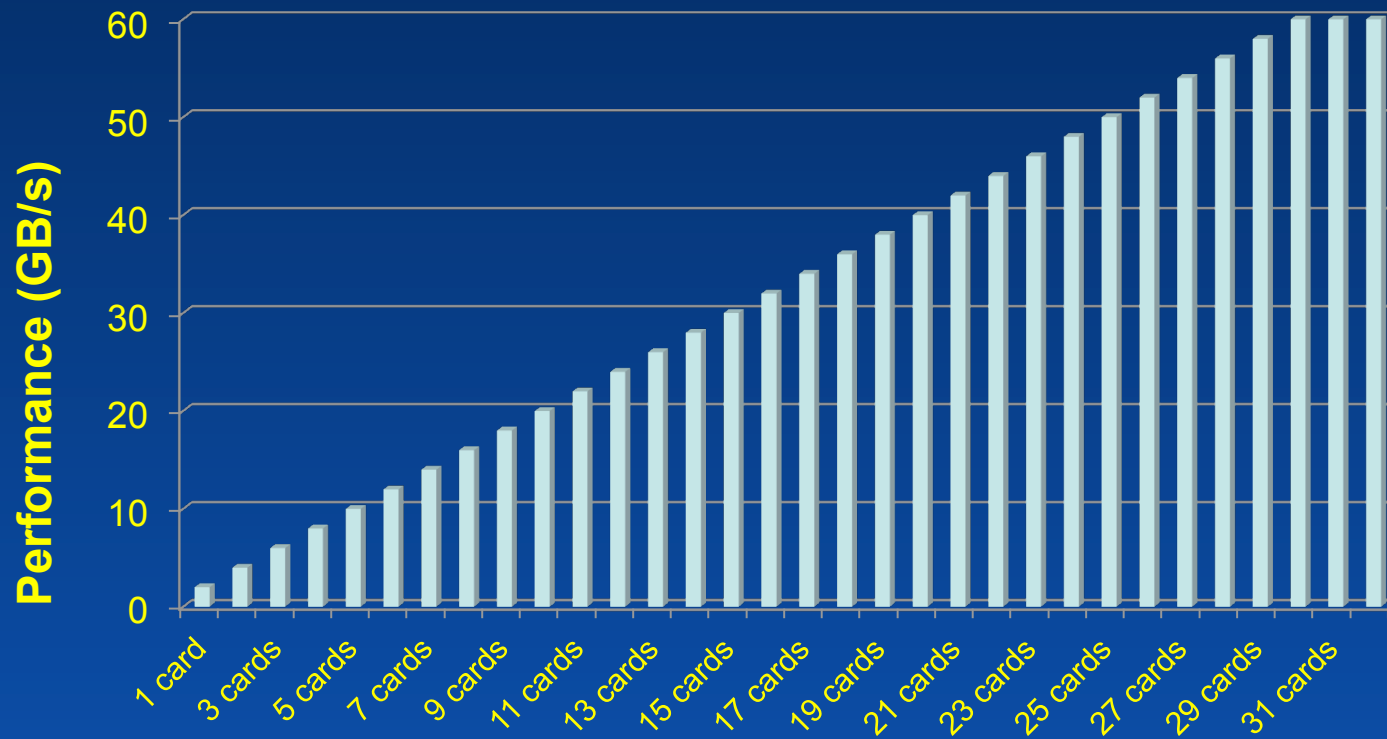
Keys to High Performance



- High performance cards
 - NVMe protocol
 - RAID flash cards together
- PCIe interfaces
 - PCIe Gen 3
 - X8 from each flash card
 - x16 to host
- Avoid conversions, compression or anything else that slows down transfers



Ideal Flash Performance Scaling

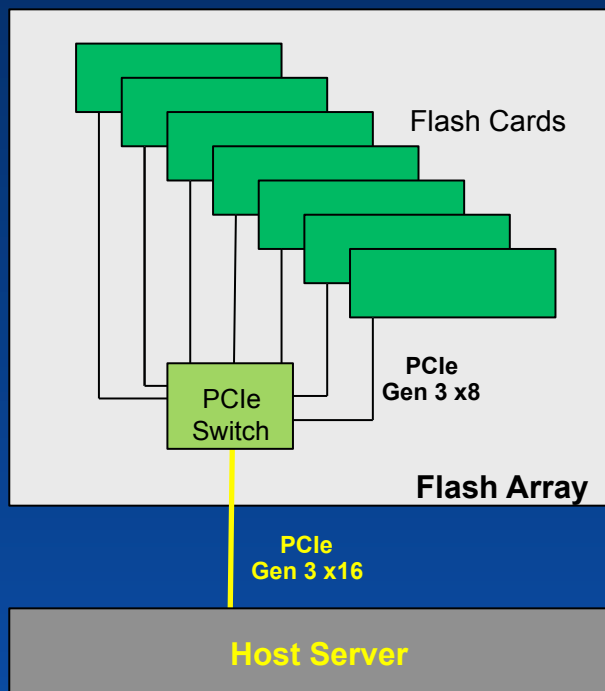




Performance Saturation Issues

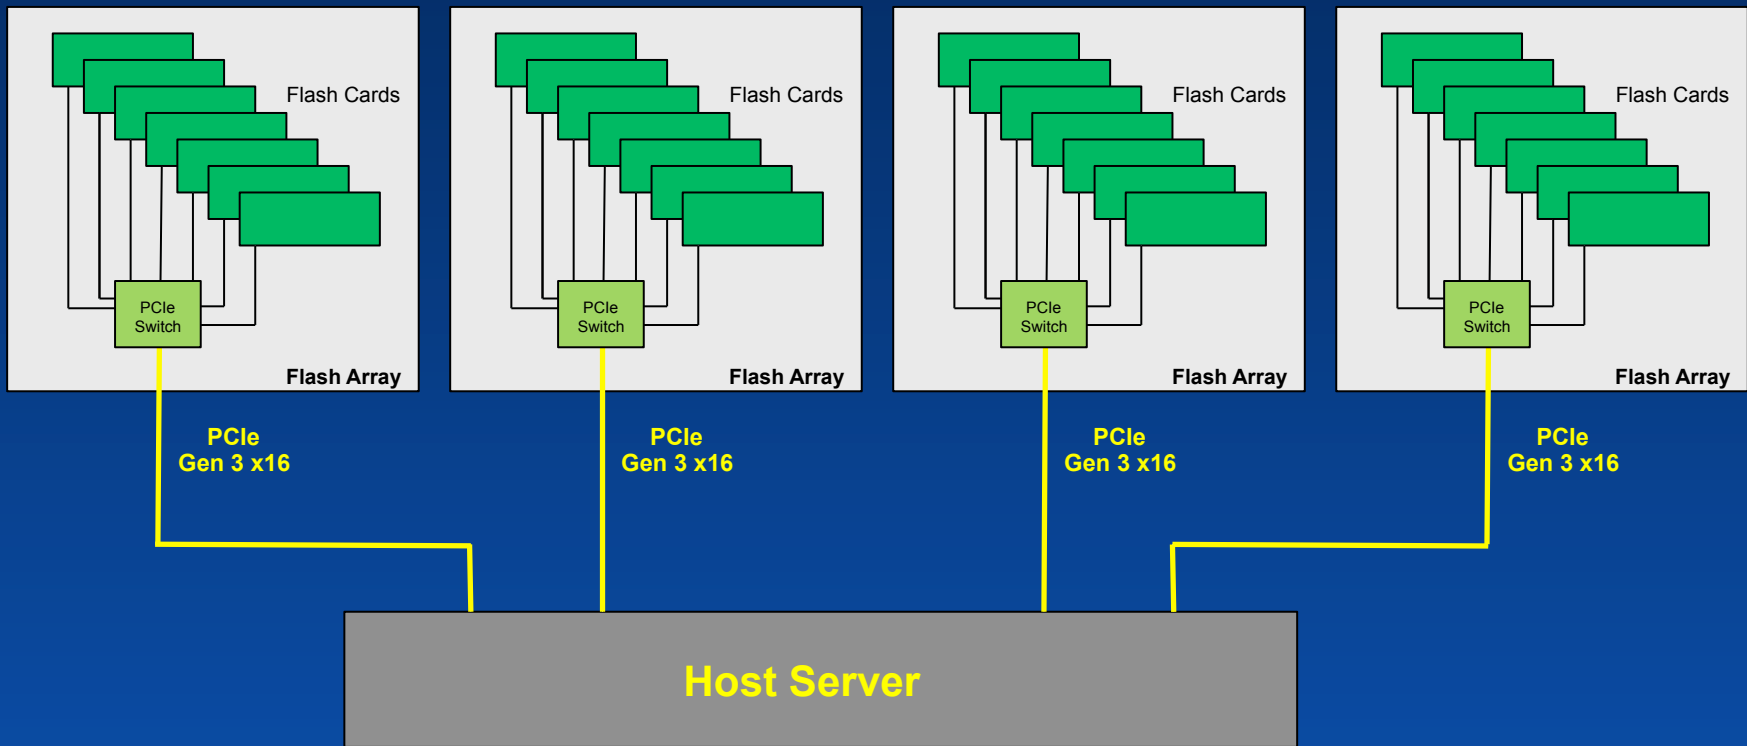
- PCIe interfaces
- Conversions to other interfaces
- Host server architecture

PCIe Performance

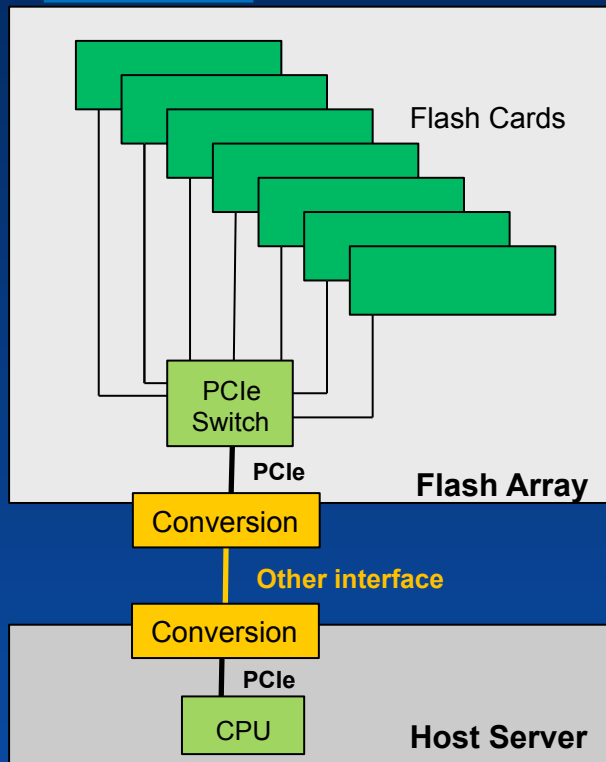


- PCIe
 - PCIe varieties
 - Flash cards – Gen 3 x8
 - Cards utilize 15-50% of bandwidth
 - Link to host – Gen 3 x16
 - Saturates with 4-8 cards
 - Multiple PCIe interfaces extend performance

PCIe to Host Performance

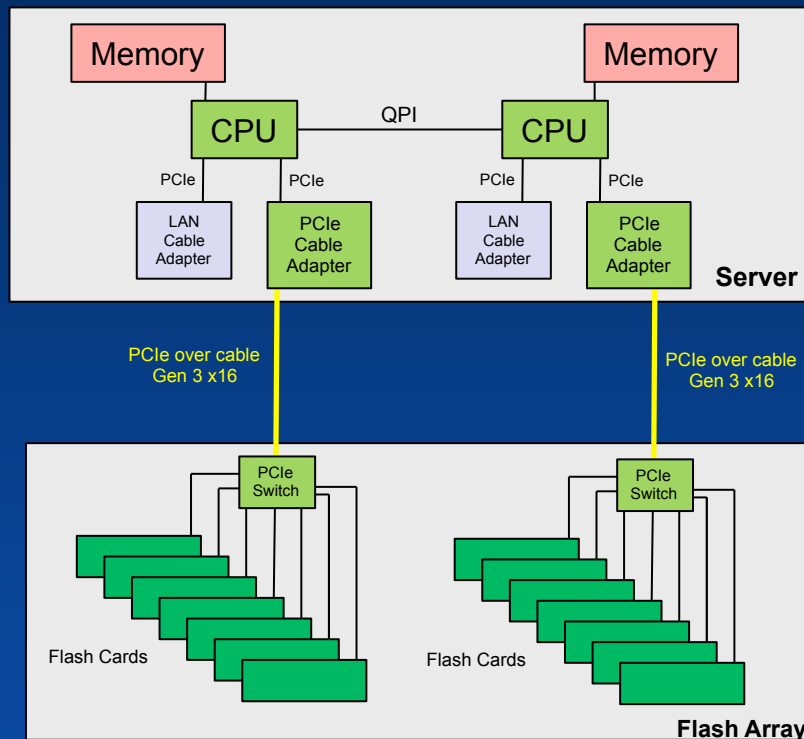


Conversions



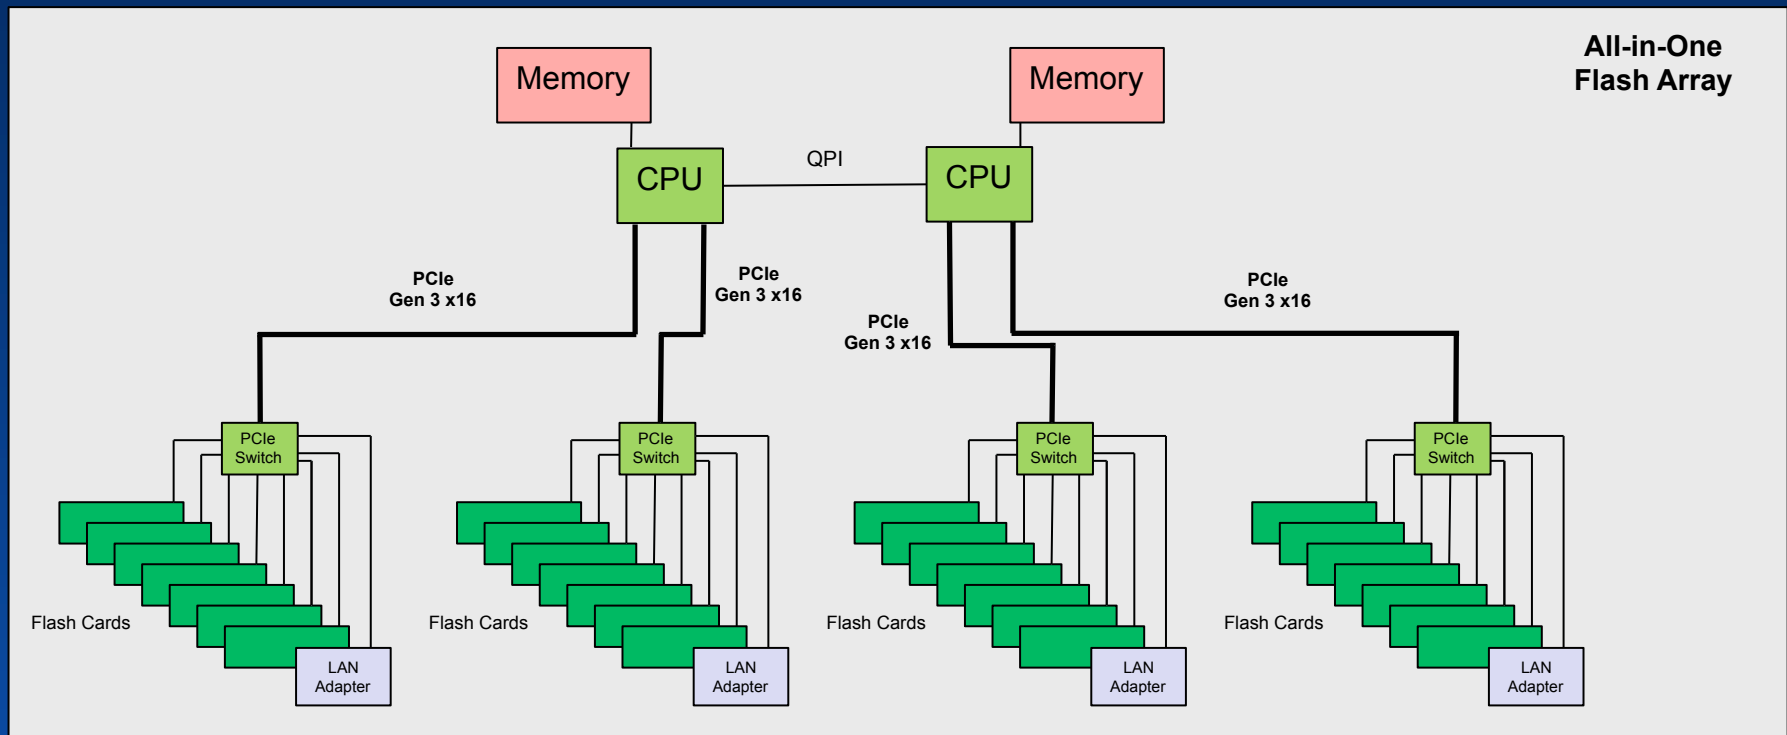
- PCIe interfaces are the standard
 - Flash controllers have PCIe interfaces
 - CPU components have PCIe interfaces
- Conversions to other interfaces
 - SAS, Infiniband, etc. conversions add latency
 - And these interfaces are slower

Host Server




- Host server speed and configuration
 - Speed of CPUs and memory
 - Left side, right side of dual CPU systems
 - Don't cross over, avoid QPI bus traffic
 - Physically understand slot configuration
 - Map CPU cores appropriately
 - Map memory appropriately
 - Network adapter slot location

All-in-One Array Performance






Examples High Performance Flash Array



- Direct Attach Array
- Holds 32 PCIe flash cards
 - Up to 200TB per system
 - Gen 3 x8 cards
 - 4 PCIe Gen 3 x16 interfaces to host servers
- 8 cards per PCIe link in four canisters
 - Cards in a canister RAIDed together
 - Canister performance of 14GB/s, 3M IOPS
- **Total performance**
 - **56GB/s transfer rate**
 - **12M IOPS**



Examples All-in-One Flash Array



- All-in-One Flash Array
- Holds 32 PCIe flash cards
 - Up to 200TB per system
 - Gen 3 x8 cards
 - PCIe Gen 3 x16 internal interfaces
- Four 100Gb Infiniband EDR interfaces
- **Total performance over Infiniband**
 - **40GB/s transfer rate**
 - **10M IOPS**



Summary

Flash Array State of the Art



- Direct Attached
 - 56GB/s transfers
 - 13M IOPS
- All-in-One
 - 40GB/s transfers
 - 10M IOPS