



# Optimizing SSDs with PCI Express

Efficient Use of PCIe Resources

Akber Kazmi

Director Marketing, PLX Technology



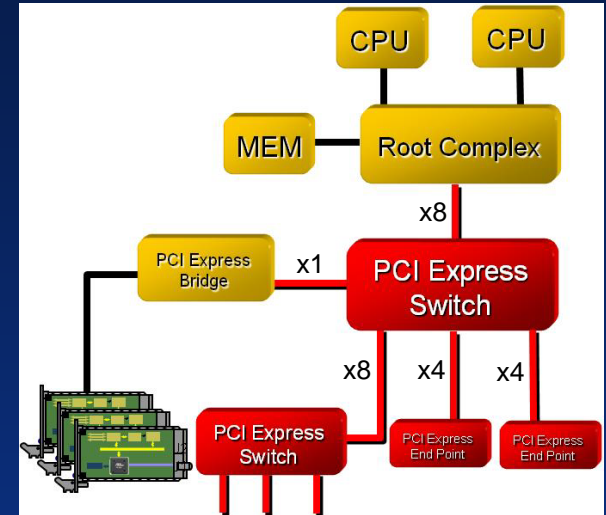
The logo for the Flash Memory Summit features a stylized yellow sunburst with multiple rays. Below the sunburst, the word "Flash" is written in red, "Memory" in blue, and "SUMMIT" in white on a blue rectangular background.

# Flash Memory Summit Agenda

- PCIe Technology Refresher
- PCIe Beyond a Fat-Pipe
- PCIe Usage Models in SSD Apps

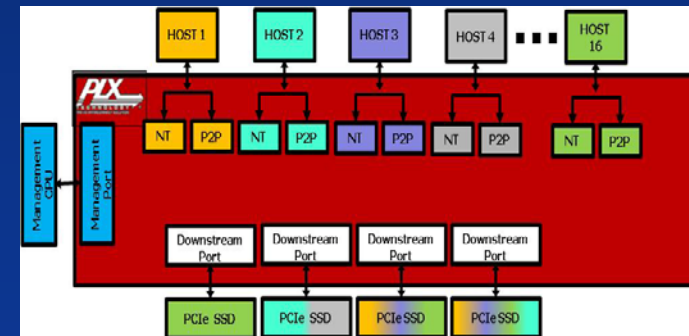
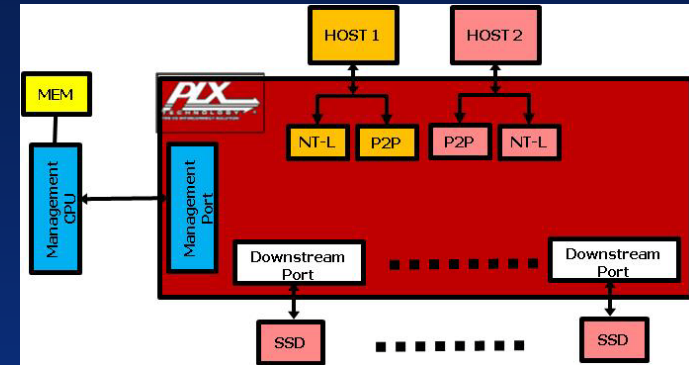
# PCIe Technology Refresher

- Replaces PCI as host bus (launched in 2004)
- Fully compatible with PCI software base
- Three generations of evolution – Gen 1, 2 & 3
  - Speeds - 2.5, 5.0 and 8.0 Gbps per lane
  - Gen 3 back-compatible to Gen 2 & 1
  - A system can run mix of Gen 1, 2 & 3 ports
  - Ports can scale to x1, x2, x4, x8, x16, x32
- Serial, point-to-point LVDS
- Embedded clock, no sideband signals
- Packet based, Credit based flow control
- Quality of Service
  - Virtual Channels
  - Traffic Classes
- Reliability, Availability and Serviceability



# PCIe Beyond a Fat-Pipe



- Host fail-over thru non-transparent bridging (NT)
  - Each host may be active – own some downstreams ports
- True peer-to-peer traffic with no CPU overhead
  - SSD controllers can talk to each other
- Host isolation thru a management CPU
  - SSD controller may pass interrupts to management CPU
- Multiple hosts sharing SR-IOV capable SSDs
  - No MR-IOV endpoints available
- Built-in DMA controllers in switches to move data



## PCIe Usage in SSD Apps

- Direct SSD controller w/PCIe host interface
- Multiple SSD controllers aggregated with a Switch
- Host Systems with more PCIe ports
- PCIe Fanout within the storage shelf
- IO Virtualization and Sharing

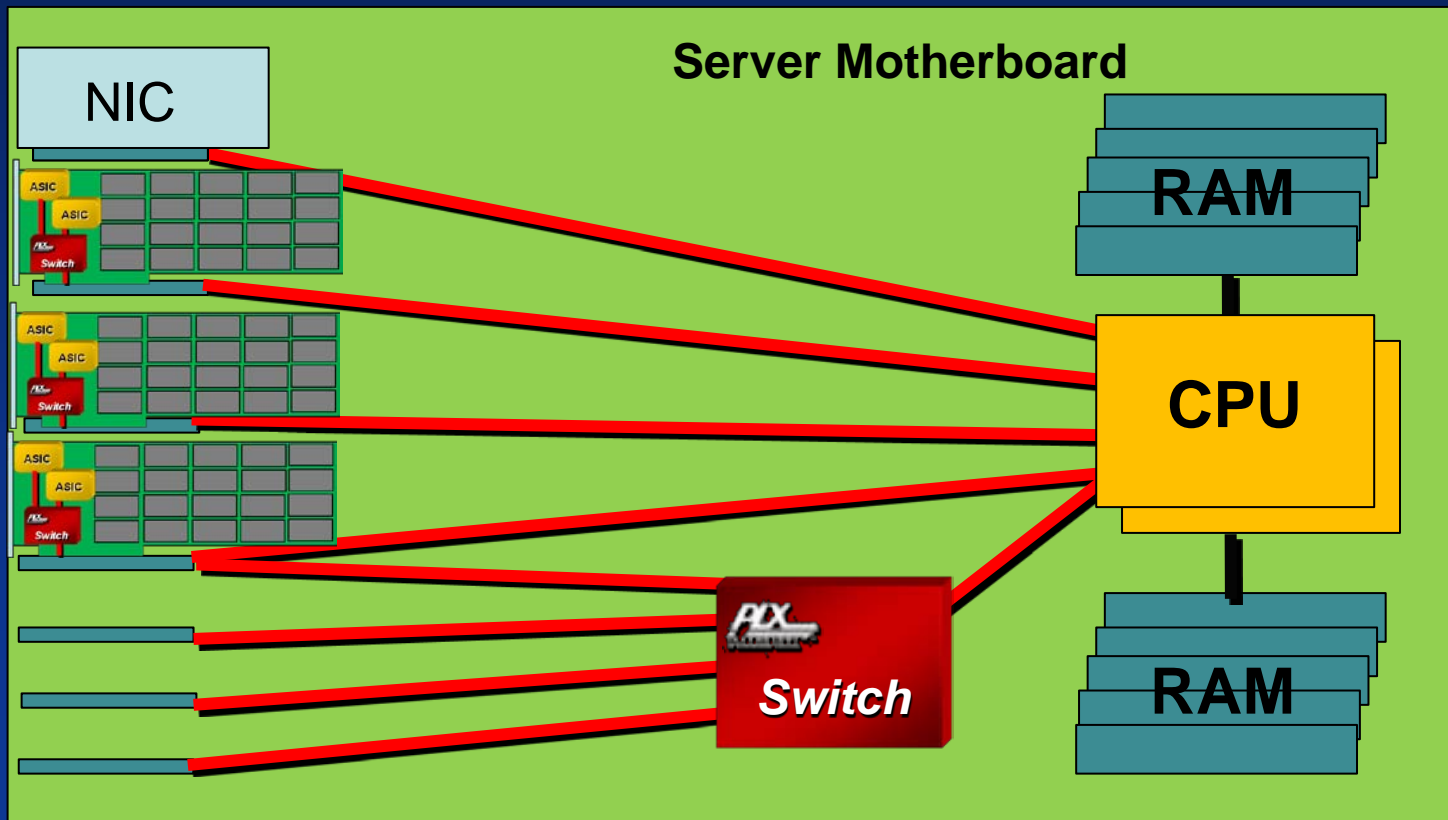
# Usage Models – 1



- SSD Controller directly linked to host
  - May not allow desired capacity
  - May not allow desired performance
  - May not provide desired longevity
- Two or more SSD controllers aggregated with a Switch
  - Double the capacity
  - Double the performance
  - Peer to peer ASIC communication for longevity
  - Reduce host/CPU overhead by allowing one ASIC to manage other ASIC

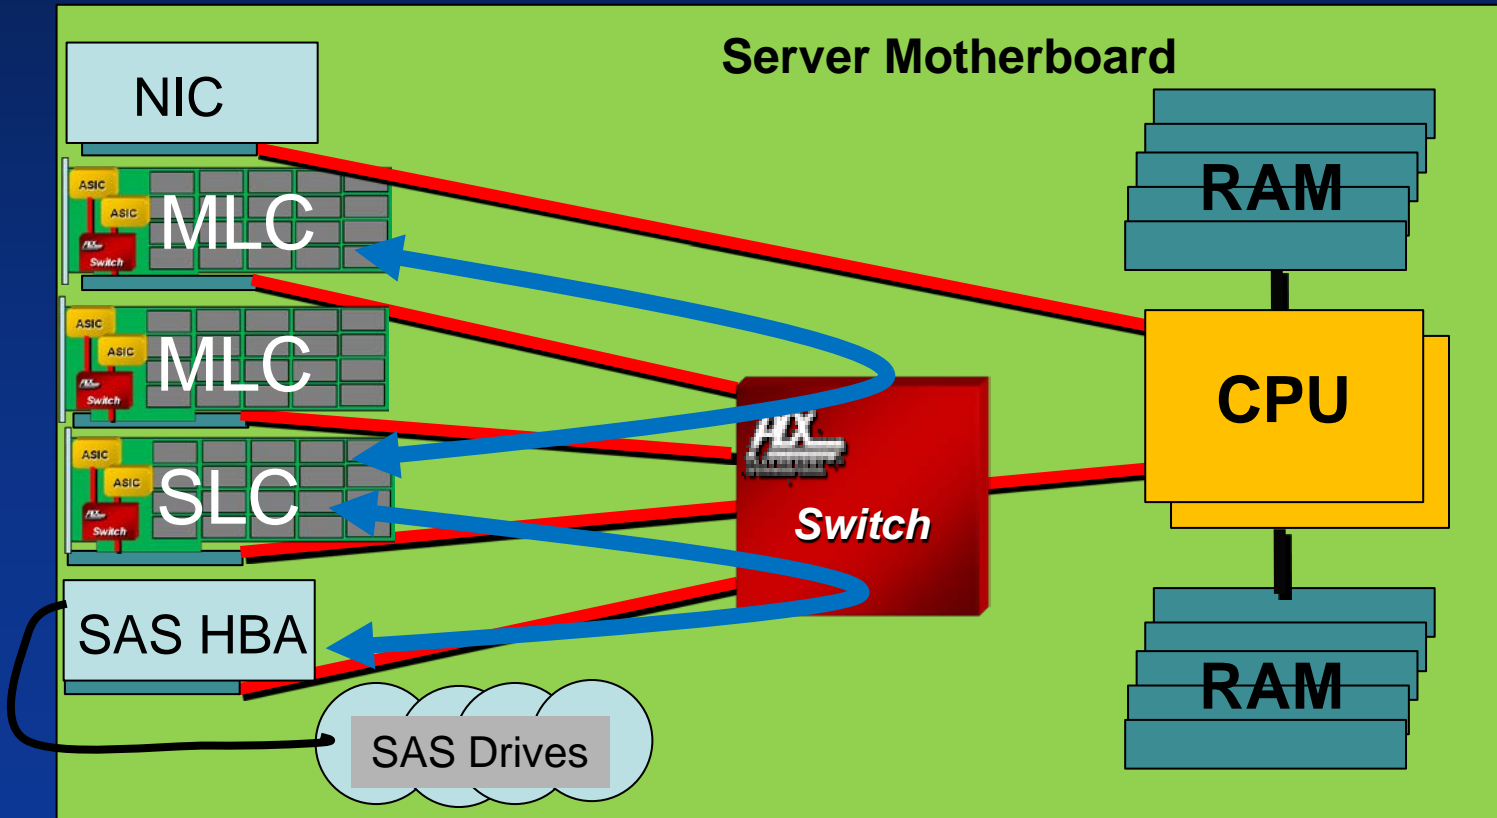
## Usage Models – 2

- Capacity expansion in server enclosures
  - Aggregation through PCIe Switch



# Usage Models – 3


- Mixing Storage Medium for cost/performance
  - Move data between these medium






## Usage Models – 4

- Creating Storage Expansion Box with PCIe
  - Connect to Server through PCIe CU/OP cable
  - No protocol translation overhead




# Usage Models – 5

- Share an array of SSD modules in a chassis
  - Associate a segment of SSDs to specific Server



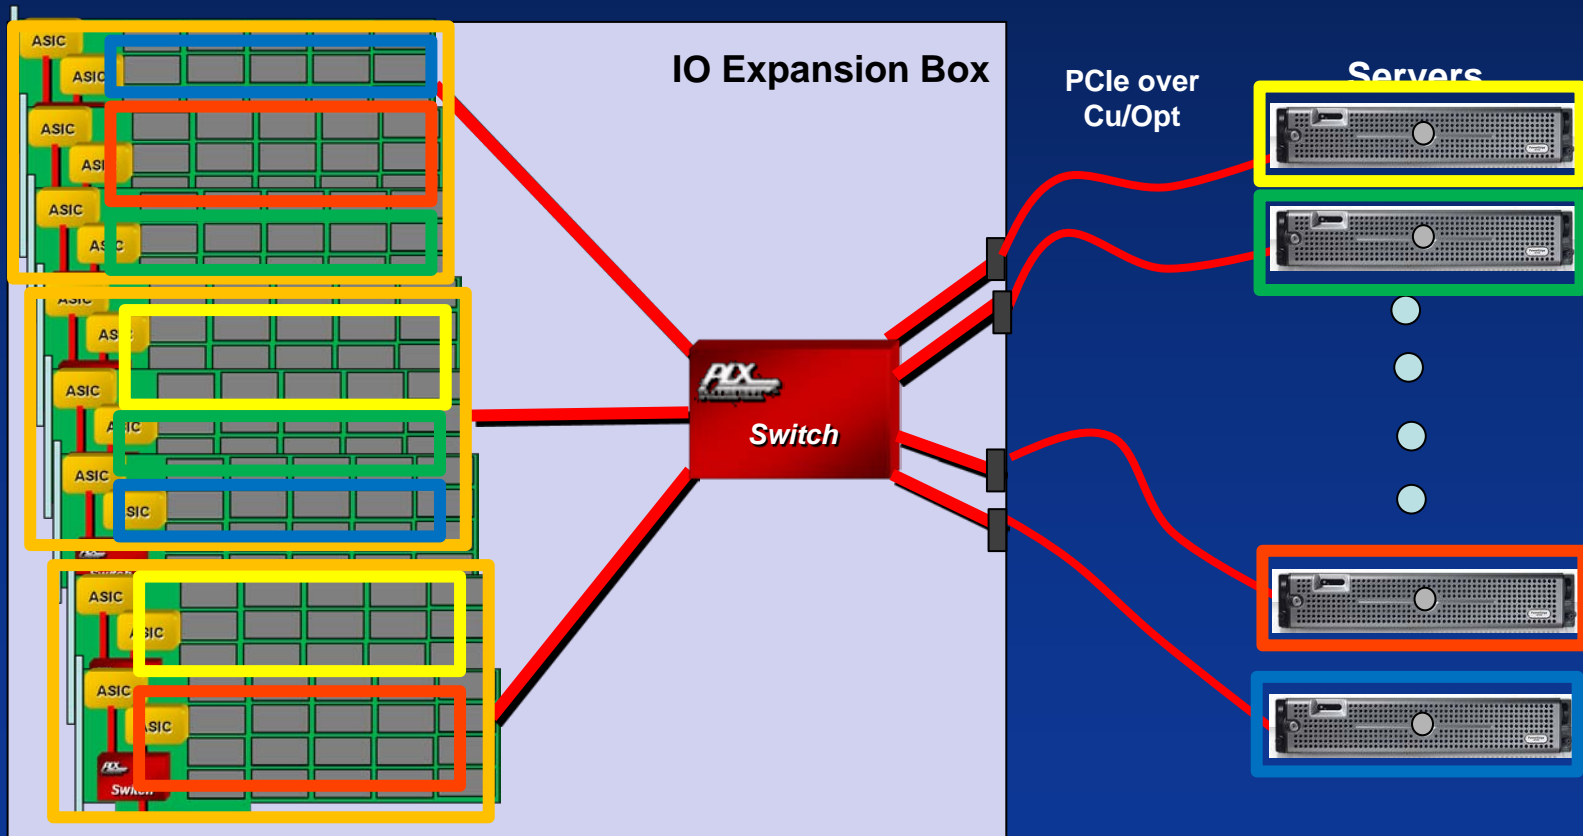
# Usage Models – 6

- Array of SSD modules in a chassis
  - Build Server Redundancy



# Usage Models – 7

- Share SR-IOV SSD modules in expansion chassis
  - Shared by multiple Servers/Hosts or Blade Servers



- PCIe offers ubiquity & bandwidth for SSD apps
  - Mature technology available everywhere
  - Gen 3 at 8Gb/s (can scale to 8Gx32)
  - Expected to go to 16G/s per lane with Gen 4
- PCIe switch vendor value-add features offer
  - Optimization of cost/performance
  - Scaling and expansion of capacity
  - Cost & power savings thru IO-sharing
  - Redundancy and fail-over



Thanks You!

Questions?

[akazmi@plxtech.com](mailto:akazmi@plxtech.com)