




# Flash Uses in High-Performance Platforms

Jim Cooke [JCooke@Micron.com](mailto:JCooke@Micron.com)

Sr. Manager, Enterprise Storage Segment

Micron Technologies

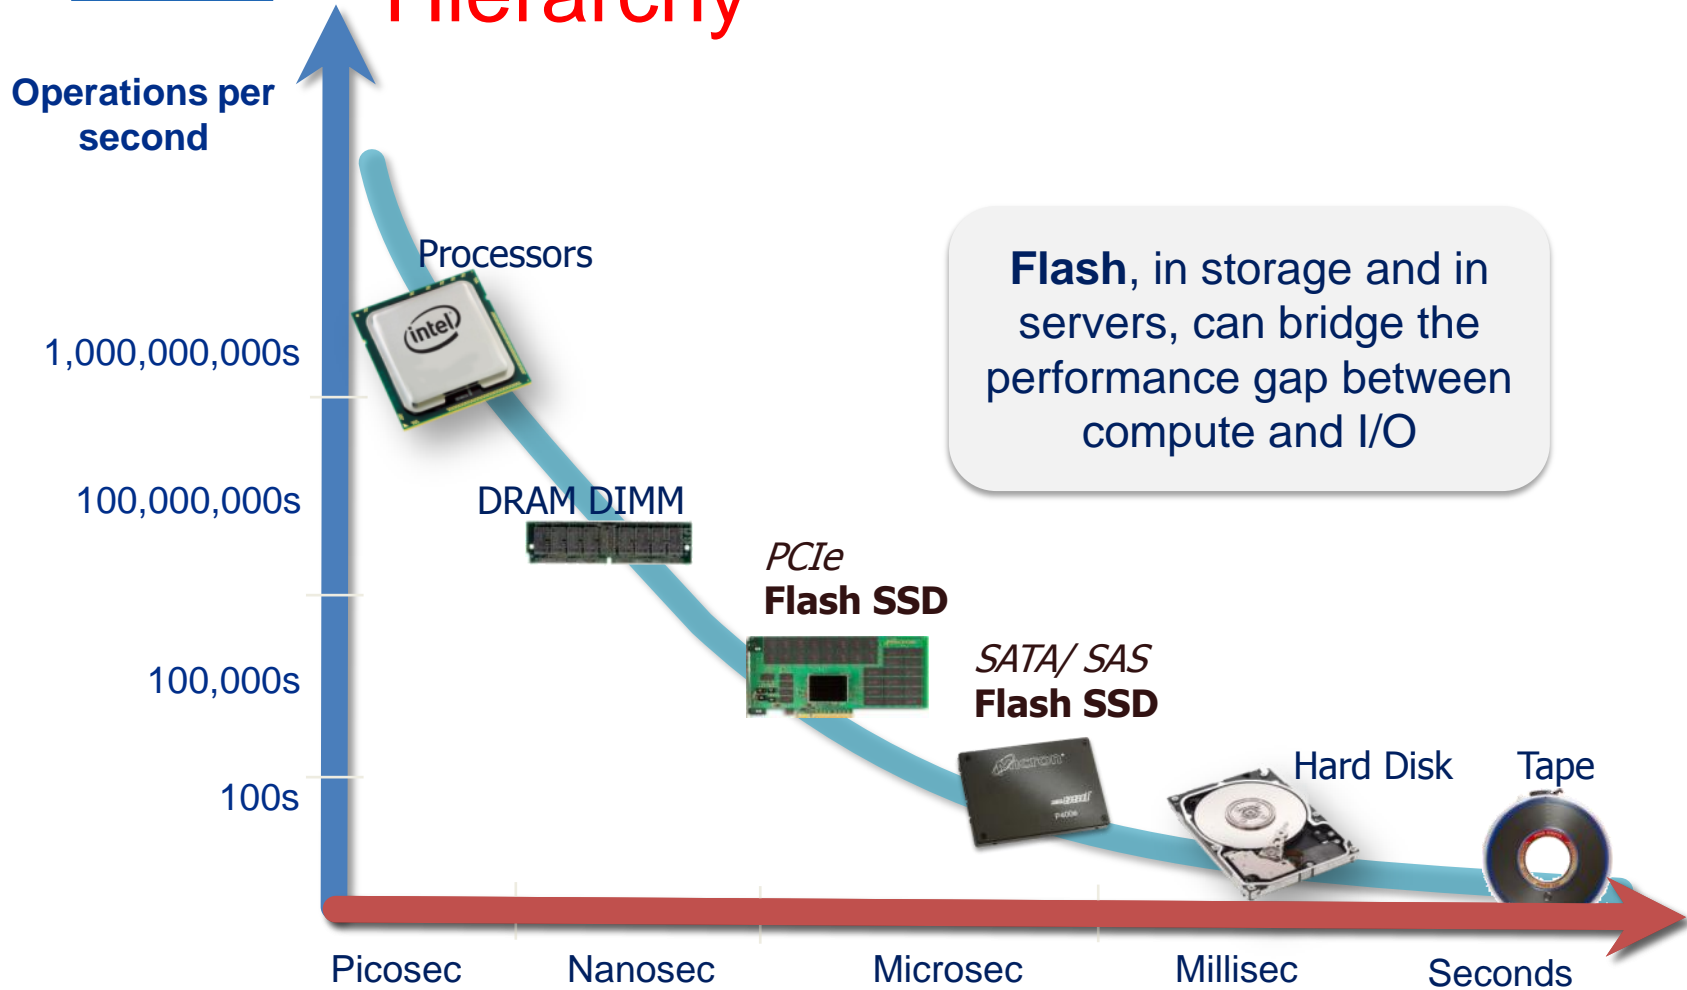
# The DRAM / HDD Speed Gap



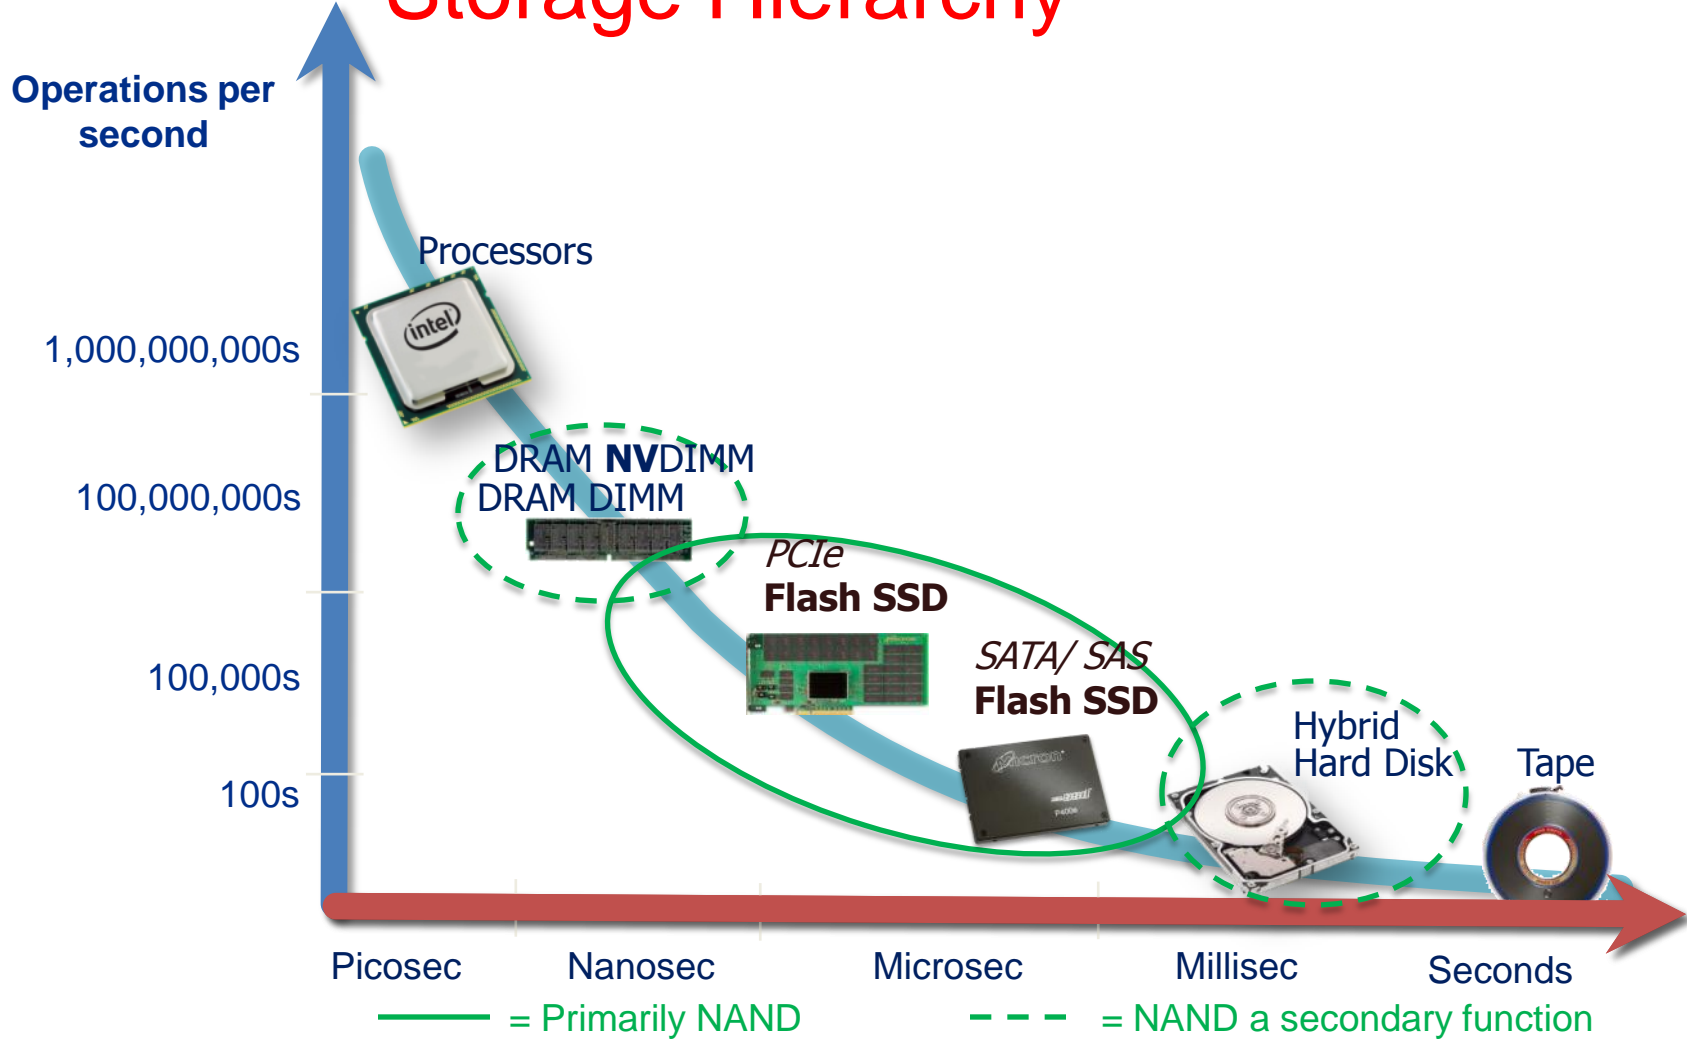
Source: Objective Analysis

Price per Gigabyte

# Evolving CPU / Memory / Storage Hierarchy



# NAND Uses in the Memory / Storage Hierarchy



# Some Examples



**Before:** 6 x 300GB of 10k SAS HDDs



| # Drives | Cost    | IOPS  | Power | IOPS/W | \$/IOPS |
|----------|---------|-------|-------|--------|---------|
| 6        | \$3,114 | 1,500 | 48W   | 31.25  | \$2.08  |

**After:** SSD Caching with 3 x 750GB SATA HDDs & 1 x 50GB SSD



| # Drives | Cost    | IOPS   | Power | IOPS/W   | \$/IOPS |
|----------|---------|--------|-------|----------|---------|
| 4        | \$2,000 | 10,200 | 17.5W | 583      | \$0.21  |
| ↓ 33%    | ↓ 29%   | ↑ 580% | ↓ 64% | ↑ 1,765% | ↓ 90%   |

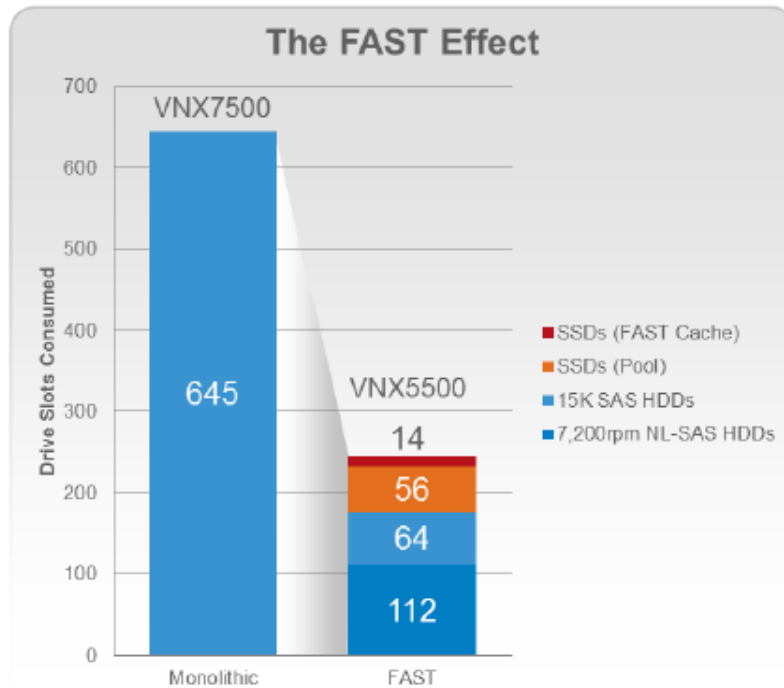
**5.8x increase in IOPS & 17x IOPS per Watt**  
**64% lower power & 33% less drives**  
**30% lower hardware cost & 90% lower \$/IOPS**

***Green IT & Data Center Efficiency to Drive Enterprise Flash Adoption***

Source: PMC Sierra

# Increased Storage Efficiency





100 TB Usable Capacity Example



- 645 X 600GB 15K
- Replaced by 246 drives in a 6/10/84 blend:
  - 70 X 200GB SSDs
  - 64 X 600 GB 15K HDDs
  - 112 X 2TB NL-HDDs
- Smaller frame needed
  - 62% smaller footprint
  - 75% less power
  - 27% lower acquisition cost

Source: Denis Vilfort, EMC, FMS2011

# What is most important to you?

|             | Hard Disk Drives<br>(HDD)   | Solid State Disks<br>(SSD)  |
|-------------|---|---|
| Density     |  |   |
| Reliability |   |  |
| Performance |   |  |
| Lower Power |   |  |

***SSDs far exceed HDDs in Reliability,  
Performance, and Power***

# New Value Metric – \$\$ per IOP

How much real performance do I get for my money?



HDD Enterprise

~ \$400

~150 IOPs

**\$2.66 / IOP**



SSD Enterprise

~ \$500

~9,500 IOPs

**\$0.05 / IOP**

**53X**



HDD Client

~ \$80

~120 IOPs

**\$.65 / IOP**



SSD Client

~ \$400

~2600 IOPs

**\$0.15 / IOP**

**4X**



# New Value Metric – IOPs per Watt

How much real performance do I get for electricity used?



### HDD Enterprise

~ 150 IOPs

~ 16.5 Watts

**9 IOPs / Watt**



### SSD Enterprise

~ 9500 IOPs

~ 7 Watts

**1357 IOPs / Watt**

**150X**



### HDD Client

~ 120 IOPs

~ 9.3 Watts

**13 IOPs / Watt**



### SSD Client

~ 2600 IOPs

~ 3 Watts

**866 IOPs / Watt**

**66x**

~ = based on an average of several drives currently available


# 2.5" PCIe Form Factor



All the performance of PCIe with the serviceability standards of SATA/SAS

# SSD Small Form Factor Working Group: 2.5" PCIe Focus

- Hot plug
- External access
- Combo connector
- PCIe performance
- Common SFF 15mm form factor
- Combo connector ([SFF-8639](http://www.sff-ta.com/standards/SFF-8639))



[ssdformfactor.org](http://ssdformfactor.org)


# 2.5" PCIe Advantages



- PCIe performance
- Common Form Factor
- Compactness

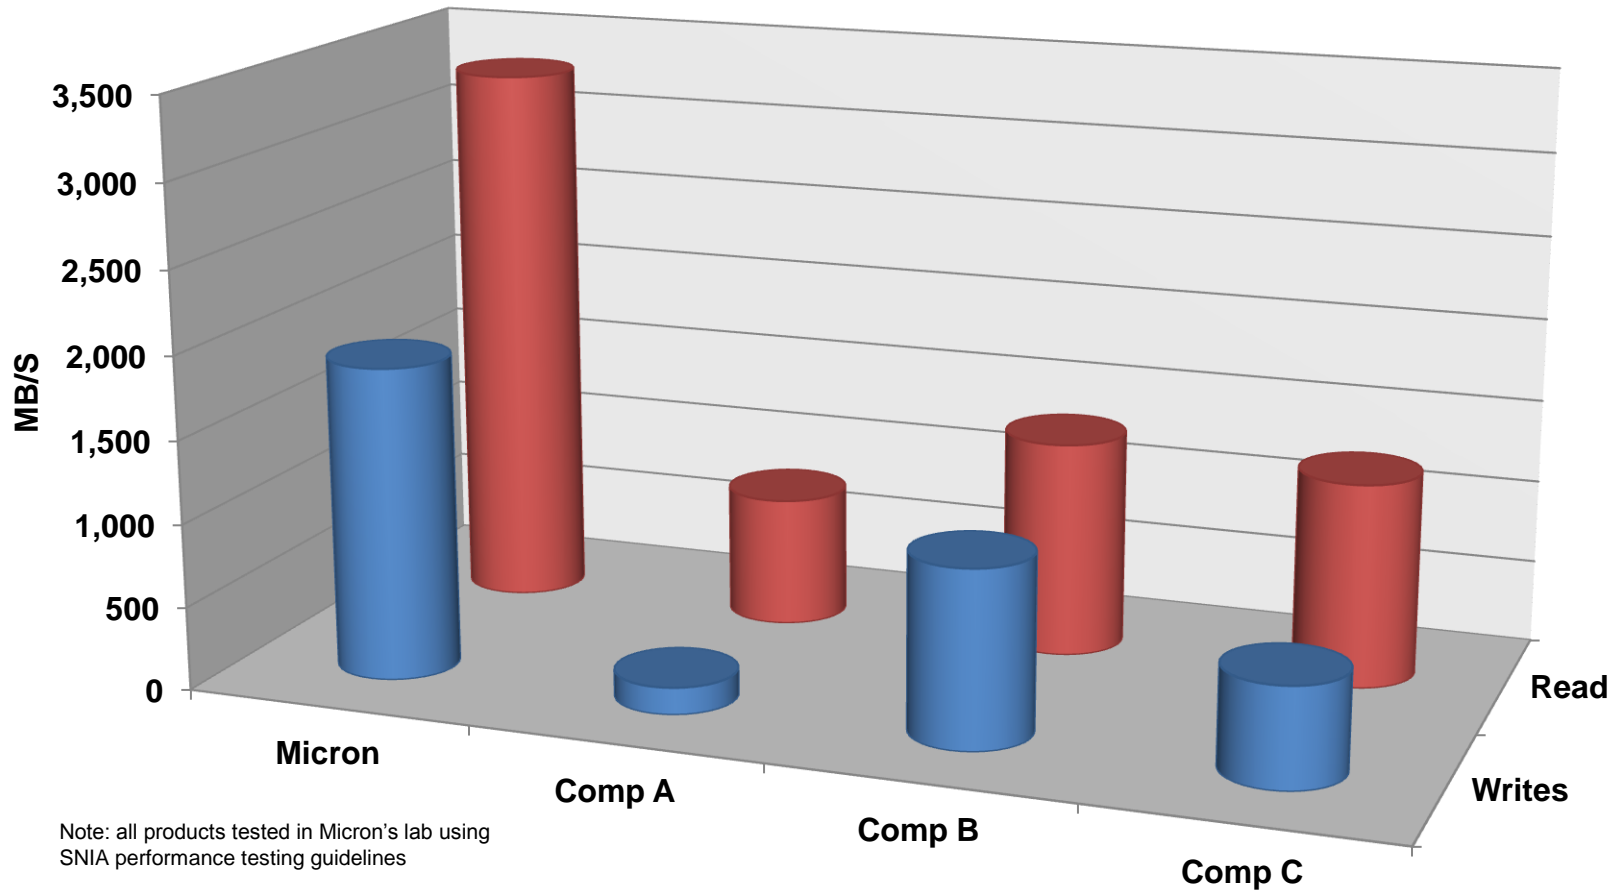
- Serviceability
- Lower TCO
- Supports RAID

# PCIe SSD: 4K Random R/W Performance






# PCIe SSD: 128K Sequential R/W Performance






# Is there a way to share High Performance SSD's?

# What's Next To Share - PCIe SSD



- Dedicated PCIe SSD Card
- Increased management
  - Poor resource utilization
  - Expensive
  - Fixed performance

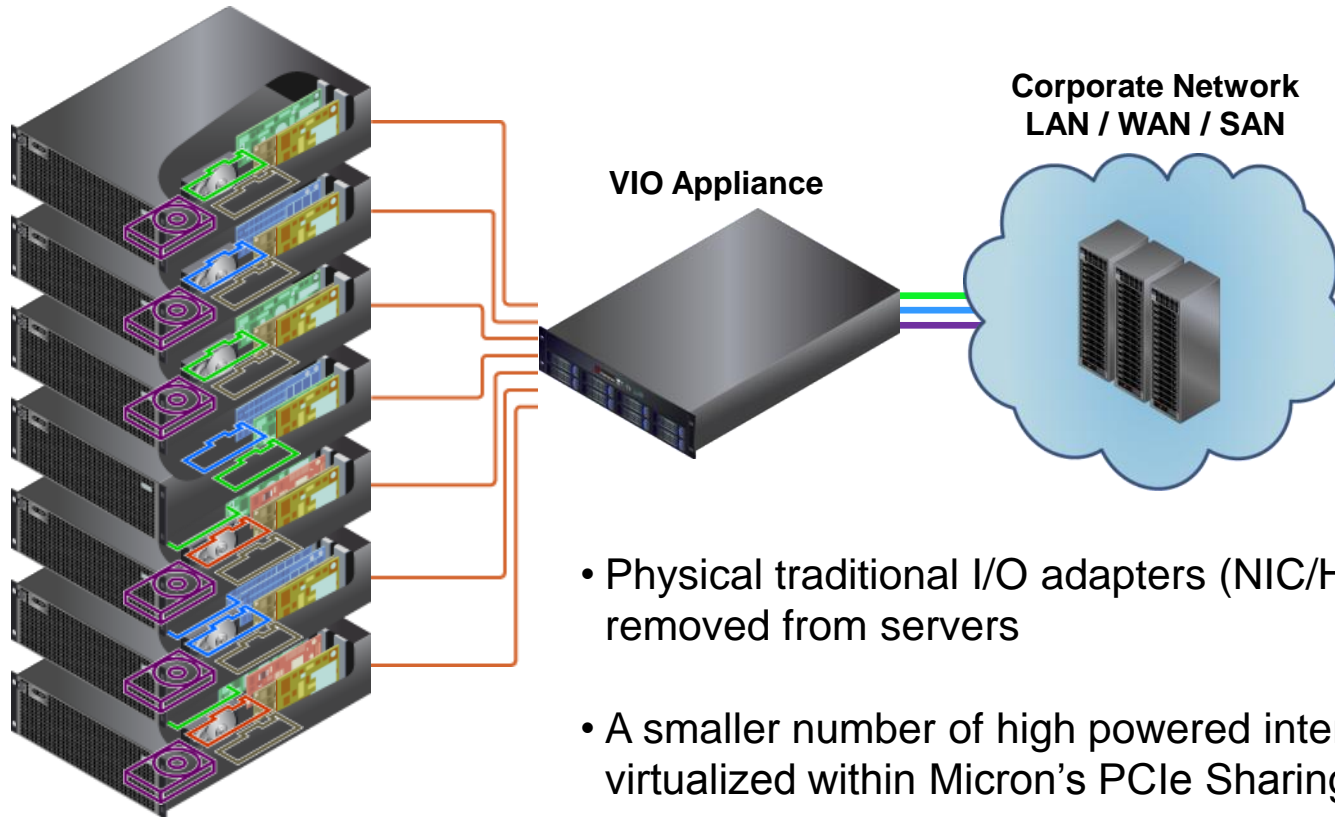
Sharing provides precisely the capacity and performance a server/application needs when it needs it



- Shared PCIe SSD
- Centralized management, resource, data
  - High resource utilization
  - Cost effective
  - Higher performance




# Micron PCIe Sharing Technology




- Physical traditional I/O adapters (NIC/HBA) are removed from servers
- A smaller number of high powered interfaces are virtualized within Micron's PCIe Sharing Appliance
- A simple, low power, low profile PCIe extender card is inserted into the servers connected to the VIO through a PCI extension cable

# Sharing Benefits

- Dynamically provision bandwidth, capacity, and connectivity
- Move I/O resources from one machine to another without touching the systems
- Size resources based on average need, but get much higher peak performance for “bursty” application traffic



**Equipment Cost Savings**



**Power Cost Savings**

- SSDs are now a critical part of the storage hierarchy
  - Bridge the performance gap between HDD and DRAM
  - Increase storage efficiency
- SSD advantages change the way we measure value
  - \$ / IOP
  - IOPs / watt
- High performance platforms adapting to take advantage of PCIe
  - 2.5" PCIe form factor provides all the performance of PCIe with the ease of use of SATA/SAS
  - PCIe SSD sharing provides precise capacity and performance when you need it



# Other Micron Presentations

- NAND uses in high performance platforms
  - Tutorial A-11 – Tuesday, August 21<sup>st</sup> @ 8:30 am
- NAND flash architecture and specification trends
  - Tutorial B-11 – Tuesday, August 21<sup>st</sup> @ 8:30 am
- MLC media discussion
  - Tutorial C-11 – Tuesday, August 21<sup>st</sup> @ 8:30 am
- Next-generation storage and the mobile computing ecosystem
  - Session 101-B – Tuesday, August 21<sup>st</sup> @ 8:30 am
- Why ECC-free NAND is the best solution for high-performance applications
  - Session 102-A – Tuesday, August 21<sup>st</sup> @ 10:10 am
- How ONFI standards are fueling high-performance SSDs
  - Session 102-C – Tuesday, August 21<sup>st</sup> @ 10:10 am
- The need for differentiated MLC solutions
  - Tutorial F-21 – Wednesday, August 22<sup>nd</sup> @ 8:30 am
- Virtualized SSD storage for enterprise systems
  - Tutorial H-22 – Wednesday, August 22<sup>nd</sup> @ 4:30 pm
- Performance trade-offs of flash-based client storage solutions
  - Tutorial A-31 – Thursday, August 23<sup>rd</sup> @ 8:30 am
- Phase Change Memory – Panel Discussion
  - Session 302-D – Thursday, August 23<sup>rd</sup> @ 9:50 am
- 2.5-inch PCIe interface for enterprise flash cache – Panel Discussion
  - Session 303-B – Thursday, August 23<sup>rd</sup> @ 3:10 pm

