
# I/O Virtualization: Enabling New Architectures in the Data Center for Server and I/O Connectivity

Marek Piekarski

marekp@micron.com


# Agenda

- I/O virtualization – ‘what’ and ‘why’
- I/O virtualization and PCIe Flash storage
- Impact on data center architectures

# What is I/O Virtualization?

Consolidate I/O into IOV appliance
- Adaptors are shared
- High utilization
- I/O capacity on demand

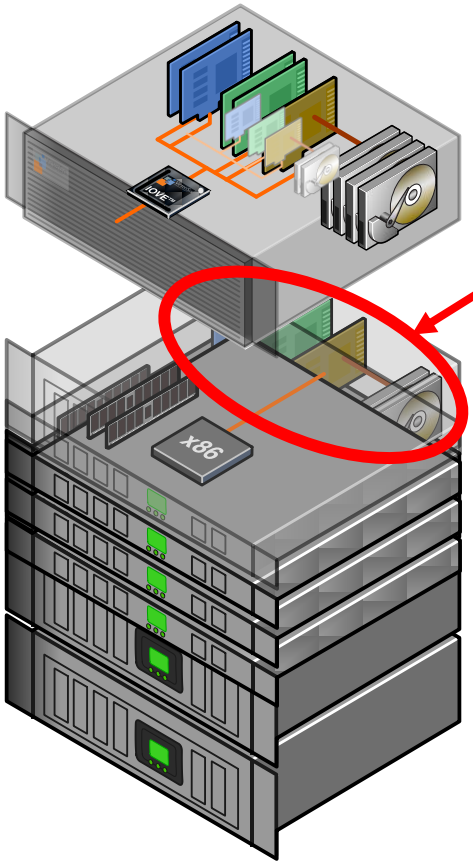


Separate I/O from servers

**Server I/O Subsystem**
- I/O adaptors (NICs, HBAs…)
- Cables & switch ports
- Direct Attached Storage (PCIe Flash)

Create virtual I/O in servers
- No disruption to server software

Servers become space/power efficient compute engines
- Just CPUs & memory

# I/O Virtualization Motivations

Virtualization (server, storage...)

- Applications decoupled from platform
- On demand server configuration
- On demand I/O bandwidth and connectivity
- On demand storage bandwidth and connectivity

# CPU & I/O-storage Evolution Rates

- New CPU
  - Every 12-18 months
  - Easy to upgrade (just the server, minimal disruption)
  - Significant cost/power/performance advantages

- New I/O-storage technology
  - Every 4-5 years?
  - Expensive to upgrade (large cost step function)
  - Hard to upgrade (server & infrastructure impact)
  - End user resistance
  - Disruption, cost & management complexity

# Why Virtualize PCIe Flash?

- Data center storage primarily networked
  - Capacity management
  - Data resilience and availability
  - Supports server virtualisation
  - Data sharing
- PCIe SSD
  - Very high throughput
  - Very low data latency
  - But very expensive

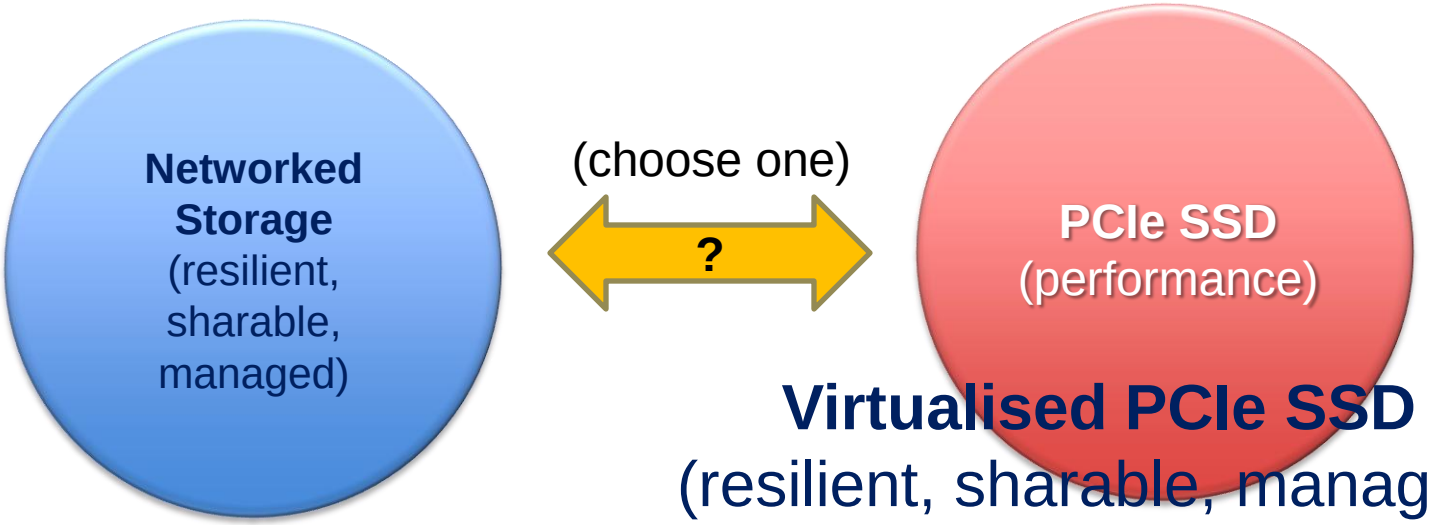


**Virtualised PCIe SSD**
(resilient, sharable, managed, high performance)

*... and a bit more affordable.*

# PCIe Flash Cache: Server-based Vs. Storage-based

- Write-through or write-back?
- Shared caches?
- Effect of failures?
- Effect on performance?
- Application tuning (size, policies...)?
- Cache warm-up times (hours rather than seconds)

# Too Much of a Good Thing?

- Most applications benefit from faster storage...
  - but very few need (can use) >100k sustained IOPS.
- Latency (performance) vs. IOPS (scalability)?
- Sharing...
  - Amortises costs
  - Provides useful performance at an affordable price
  - Performance can scale with demand
  - Provides entry point for new technologies into volume markets
  - Enables cost effective solutions for bursty usage models

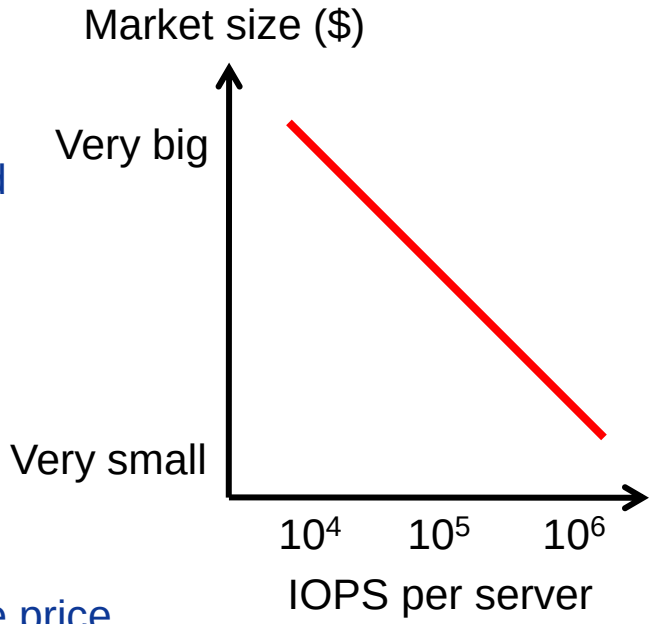

# Trends – PCIe Flash

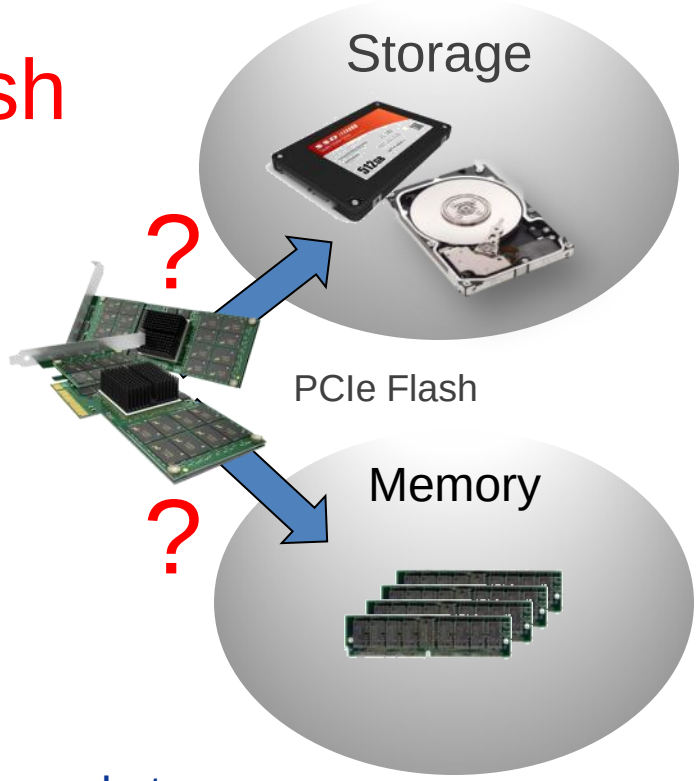


- Today: “Fast disk” - where next?
  - Not non-volatile DRAM, not fast disk.
  - Access models
    - Storage (async, block, persistent)
    - Memory (sync, byte, volatile)

- The physics of flash
  - Charge-based storage → resistance based storage
    - Density, power, endurance...
    - Note: Lower latency negates longer endurance! Need better solution.
  - 3D stacking → lower power, higher density & performance.

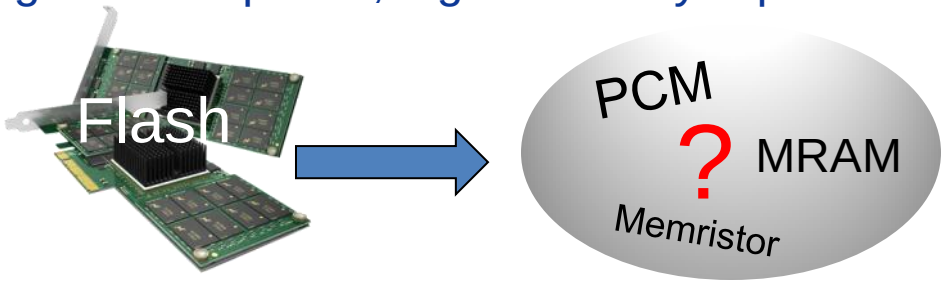

# Two Key Trends in the Data Center

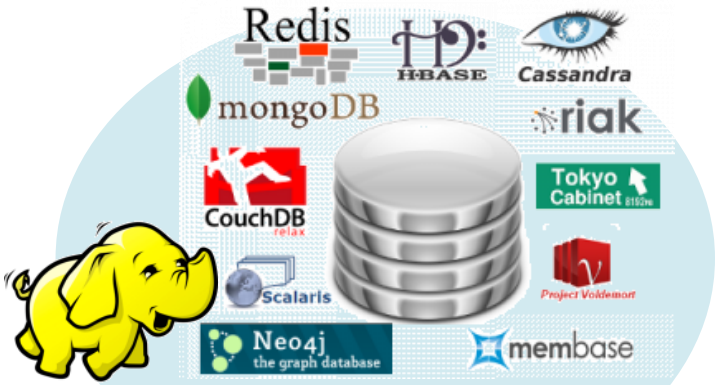


**Data-intensive workloads ("Big Data")**

Increasing volume, distributed/scalable, novel s/w arch.

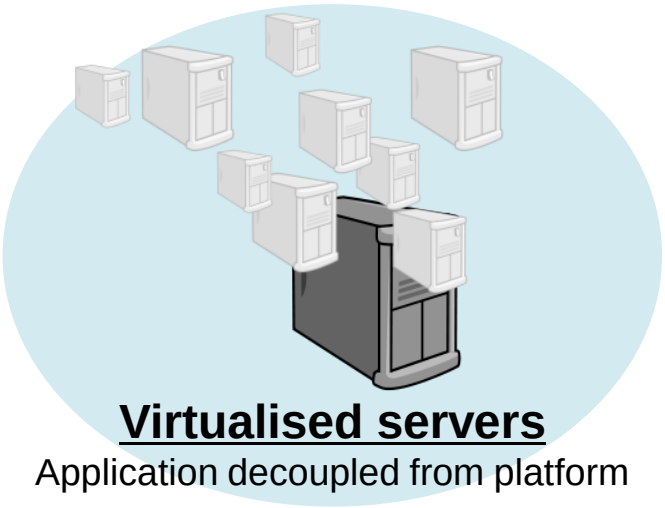

**Virtualised servers**

Application decoupled from platform

## Limitations and constraints

- Power/cooling
- Inertia (disruptive technology – a double edged sword)
- Compatibility (software and systems)
- Few have the luxury of 'clean sheets'
- Commodity hardware preferred

# Flash in the Data Center Trends: Big Data

- Solid-state Non-Volatile Storage (SSNVS) in server
  - But which?
  - How much investment?
  - How much integration?
- Point (single customer) solutions
- Short solution lifespan
- Little standardisation or cross-fertilisation
- Growth driven – but for how long?

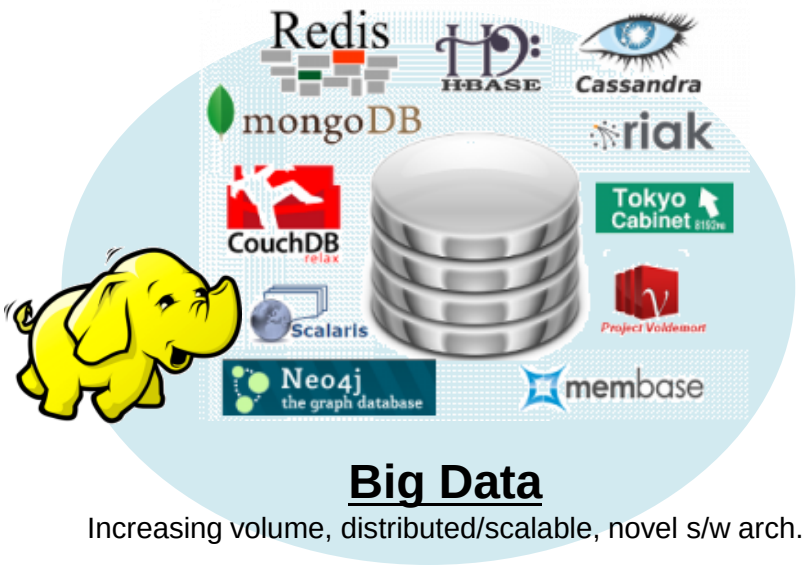


**Big Data**

Increasing volume, distributed/scalable, novel s/w arch.

# Flash in the Data Center Trends: Virtualized Data Center (short term)

- Flash = storage (async, block, consolidated)
  - Standard servers with DRAM
- Quirks of different SSNVS technologies hidden from system software
  - Hybrid solutions (e.g., PCM+ Flash)
- Increasing penetration of SSNVS-aware system software (removing HDD ‘optimisations’)

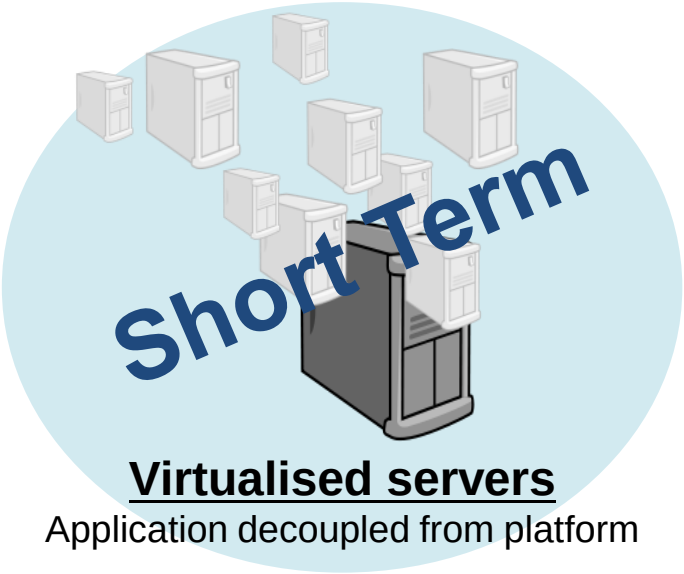


**Virtualised servers**
Application decoupled from platform

# Flash in the Data Center Trends: Virtualized Data Center (longer term)

- Opportunities of 3D stacking and improving NVS performance too good to ignore
- Intelligent storage - push some functionality into storage device.
  - Search, analyse, compress, filter, protect...
  - Performance come for free.
  - Standard functionality platform (small, power optimised ‘server’) - transparency.

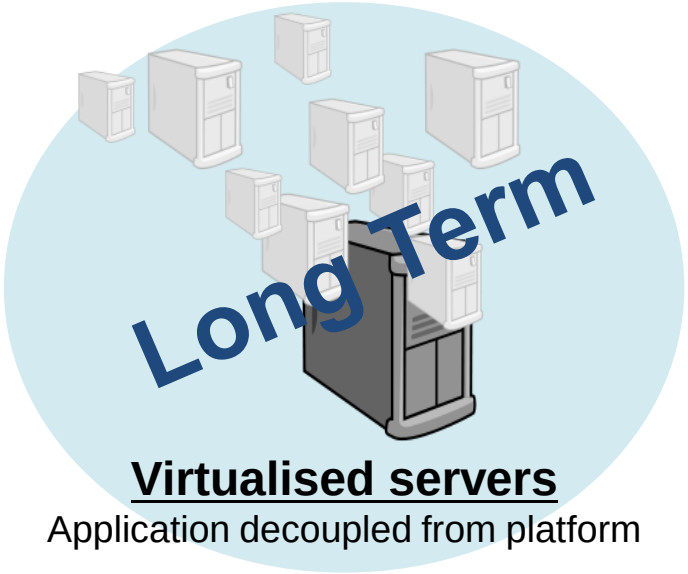


**Virtualised servers**
Application decoupled from platform

# Flash in the Data Center Trends

## Putting ***Data*** at Center of the Data Center