

Advancements in Virtualization Technology in Cloud Computing SSDs

Session Name: **SOFT-301-1: Virtualization and Orchestration**

Session Date and Time: **8/10/2023 8:30 AM**

Douglas Arens, Technical Staff, Applications
Microchip Technology Inc.



Agenda

1. What is virtualization?
2. Why do we need virtualization?
3. How is virtualization used in cloud infrastructure?
4. How does virtualization reduce TCO?



Definitions of Virtualization

Storage virtualization creates an alternative way to utilize physical storage resources to enable more effective utilization from host computer systems.

Virtualization Types and Architecture

Types of Virtualization

1. Network virtualization
2. Desktop virtualization
3. Application virtualization
4. **Storage virtualization provides a similar access method to utilize storage with many different types and characteristics to reduce TCO**

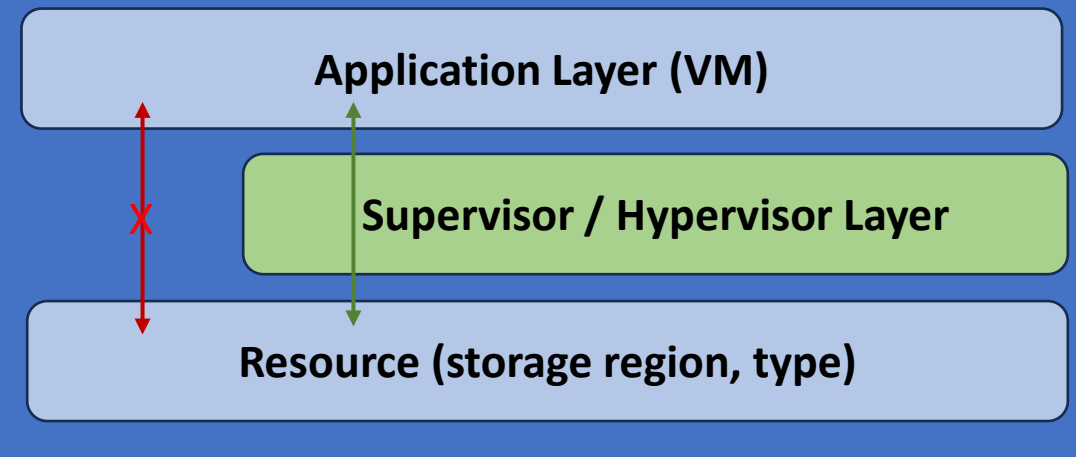
Architecture Layers

- Separate direct access from an Application (VM) to a resource
- Add a minimal supervisor layer to control access from Application to the Virtual resource.
- Add a configurable Virtual interface to map access to the actual HW resource

Benefits of Virtualization

- ✓ Reduced capital and operating costs
- ✓ Minimized or eliminated downtime.
- ✓ Increased IT productivity, efficiency, agility and responsiveness
- ✓ Faster provisioning of applications and resources

Storage Virtualization Architecture Layers



Industry Standards and Virtualization Features

PCIe® (Scalability, Extensibility)

- **SR-IOV**
- **Intel® Scalable I/O Virtualization**
- **DualHost**
- **NVMe-OF**

DevOps (CaaS)

- Google Kubernetes, Cloud Run
- Azure Container Instances
- AWS Fargate

NVMe® 2.0 (Scalability, Extensibility)

- **Namespaces**
- **Zoned Namespaces**
- **NVMe-KV**
- **MFND**
- **Endurance Groups**
- Technical Proposals

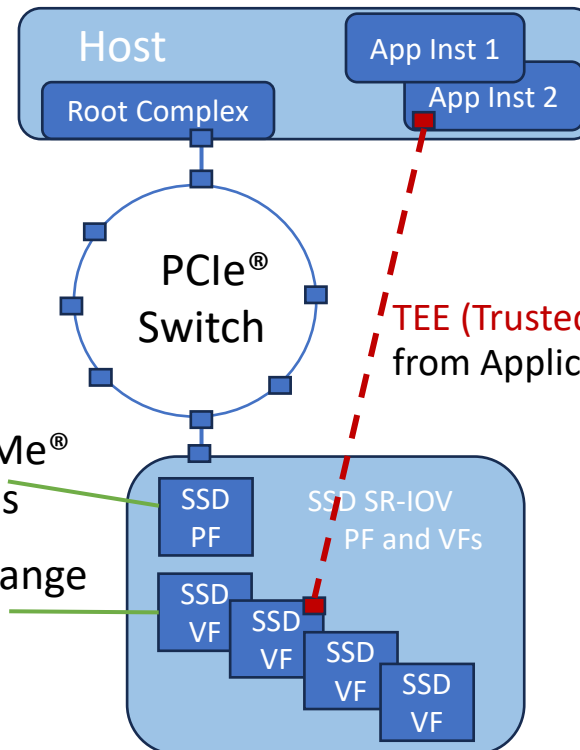
Other

- **RAID**
- SDS – Software Defined Storage
- STaaS – Storage as a Service

PCIe® SR-IOV (Single-root I/O Virtualization)

USE CASE: Multifunction device, better cost/features

SR-IOV allows a single Peripheral Component Interconnect Express (PCIe) physical device under a single root port to appear as multiple separate devices to the hypervisor or the guest operating system.



Physical Function- Real NVMe® device, configuration, access

Virtual Function- Only exchange of data

Virtualization Benefits:

- ✓ It makes it possible to run a large number of virtual machines per server, which reduces the need for separate hardware and the resultant costs of space and power required by hardware devices
- ✓ It adds the ability to run different operating systems on the same host machine, again reducing the need for discreet hardware


TEE (Trusted Execution Environments) adds back in trusted security relationship from Application instance to an SR-IOV Virtual Function.



Intel® Scalable I/O Virtualization

USE CASE: Highly Scalable Multifunction device

Intel Scalable IOV improves upon SR-IOV to allow a single SSD under a single root port to virtualize ~1000 separate storage devices for applications in hyper-scale datacenters



Physical Function- Real NVMe device, configuration, access

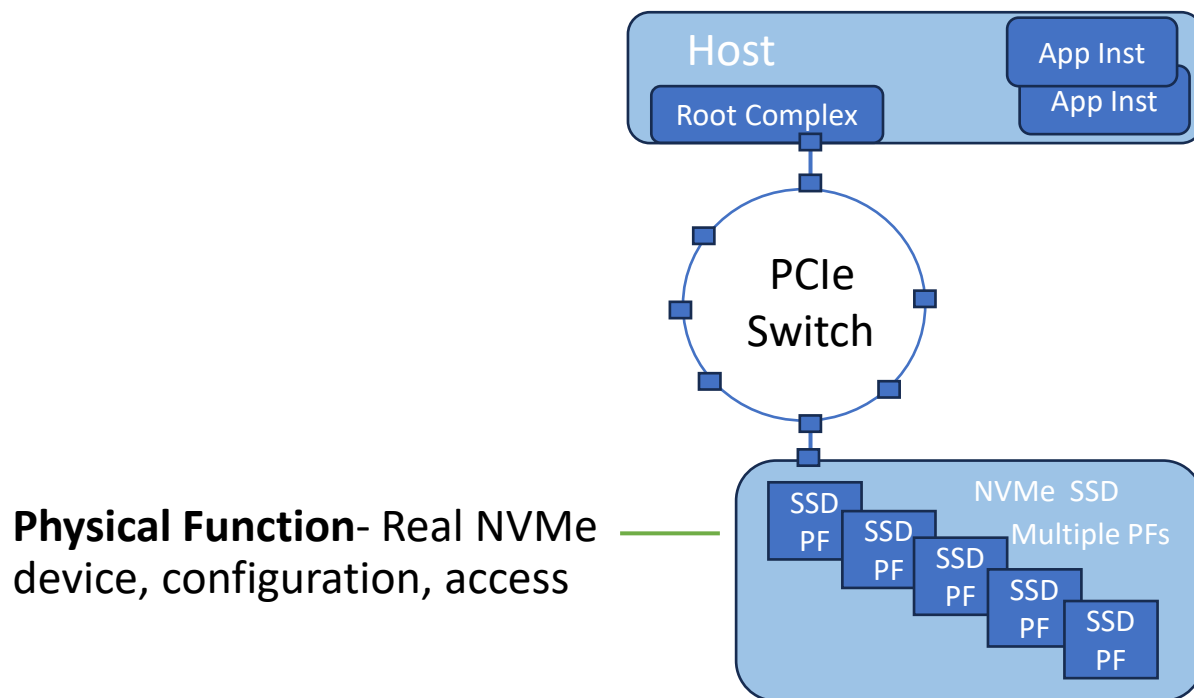
Device Resources- Exchange data (direct-addressing to queue level)

Intel Scalable I/O Benefits:

- ✓ It is theoretically possible to virtualize ~1000 storage devices per server, reducing the need for separate hardware, and costs for space and power required by hardware-centric solutions
- ✓ Storage is 16X more scalable, hardware-assisted IO paired with software flexibility. Scalable 20-bit ID (over 16-bit ID) between VM and DR.
- ✓ DMA and interrupt remapping
- ✓ Directed IO to NVME® queue level at Device
- ✓ Separation of fast path I/O from slow path (configuration, reset) for efficiency and lower latency
- ✓ VMM (Hypervisor) Reduces overhead, and provides more direct access to resources with interrupt remapping

NVMe[®] MFND (Multiple Physical Function Nonvolatile NVMe Device)

USE CASE: Multiple Functions at NVMe layer





Virtualization Benefits:

- ✓ MFND allows one physical device to contain multiple functions
- ✓ Reduce number of devices needed to support multiple users
- ✓ Provide different ways to access storage, access parameters, or QoS

NVMe® Namespaces, NVMe® Zoned Namespaces

USE CASE: Diverse use cases to access Storage Media



Virtualization Benefits:

- ✓ Namespaces provide the organization and security to allow access of separate regions (LBA ranges) of SSD storage as appropriate from different applications (VMs)
- ✓ Sub-divide a large drive into smaller drives, then paired with each application.

In NVMe 2.0, Zoned Namespaces provide equal-sized zones of non-overlapping LBA “scratchpad” regions:

- ✓ ZoneReset moves the Write Pointer to the beginning LBA of the zone
- ✓ Subsequent writes advance the Write Pointer

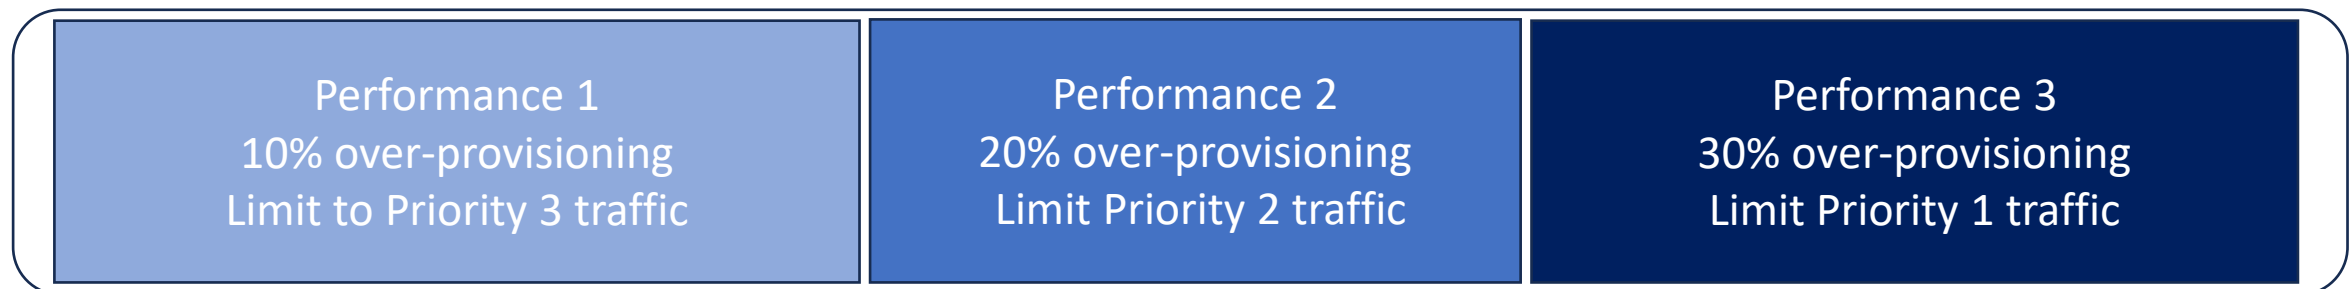
NVMe® Endurance Groups

USE CASE: Diverse use cases for storage characteristics as if they are separate mix of products/types

Virtualization Benefits:

- ✓ In NVME 2.0, Endurance Groups allow each application to select I/O Performance and wear-leveling effects
- ✓ Sub-divides a larger drive into smaller drives, each with different performance and endurance characteristics.


NVMe Endurance Groups: LBA Range



PCIe® Link Bifurcation for Dual Host

USE CASE: Hot fail-over from failed HOST/equipment

Leverages PCIe Link Bifurcation to split **one** PCIe link into **two** PCIe links (half as many lanes) for different hosts. Requires special consideration for clocking.




Virtualization Benefits:

- ✓ Fail-over mechanism
- ✓ Simultaneous access
- ✓ Upgrading server
- ✓ Upgrading storage

NVMe[®]-OF (NVMe over Fiber) NVMe-OE (NVMe over Ethernet)

USE CASE: Allow storage to be pooled, at a distance across the datacenter




Virtualization Benefits:

- ✓ Connectivity to storage over a longer distance inside the data center.
- ✓ Distances not practical for PCIe[®] repeaters/retimers
- ✓ Virtualization Protocols: RDMA, RoCE, FC-NVME, Infiniband, iWARP

RAID (Redundant Array of Inexpensive Disks)

USE CASEs: Create Virtual Drive, larger, redundant, or both


RAID 0: Create Larger Drive




RAID 1: 100% Redundant




RAID 5: Redundant Single Parity



RAID 6: Redundant Dual Parity



RAID 10: Layers of Raid 1 and Raid 0



Virtualization Benefits:

- ✓ Aggregate one large storage device from smaller devices
- ✓ Avoid backing up data each night.
- ✓ Live backup, ready for fail-over.
- ✓ Both scalability and redundancy
- ✓ Gains benefits without adding complexity on host system

Booth #419



Flexibility for Cloud

Highly-programmable
architecture with dynamic
configurability

Security for Cloud

Virtualization for Cloud

Computational Storage

Choose to work with a technology partner that can help de-risk your product delivery pipeline

