SDC
STORAGE DEVELOPER CONFERENCE
SNIA SANTA CLARA, 2014

The Ultimate Storage Virtualization

Felix Xavier, Founder
CloudByte Inc.

Agenda

- Trace through the different forms of storage virtualization
- Analyze server virtualization
- Analyze the impact of SSDs on storage virtualization
- Arrive at a solution for a complete storage virtualization

The Logical Volume Manager

- The earliest form of storage virtualization
- Abstracts the underlying drives
- Creates logical pool of storage
- Logical volumes defined on storage pool
- Data on volumes striped across the underlying drives
- Ability to thin provision storage
- Volumes not a constraint to small set of drives

The LVM virtualizes storage capacity

Virtual Mount Point – The Scale-out NAS

Larger storage performance by spreading the workload across many drives

Poor node performance limited by the number of drives each node can handle

One large application needs, for storage performance, more than what single storage node can handle

Node aggregation The Scale-out NAS emerged to aggregate multiple storage nodes to provide single virtual mount point across nodes

The Scale-out NAS virtualizes storage performance

Complete Storage Virtualization

- The LVM virtualizes storage capacity
- The Scale-out NAS virtualizes storage performance
- Combining the two technologies provides complete storage virtualization. *Is it true?*
- **NO.** Because SSDs entered the Enterprise storage market…
- Then what is complete storage virtualization?

SSDs changed the landscape of storage virtualization

Learning from the Server World

- **Hardware outcompetes app reqs** Hardware capabilities multiplied very fast during the early 2000s and out-competed application requirements
- **Applications didn't catch up** with hardware capabilities
- **Gross under utilization of compute resources** Since there is no credible way to share resources with multiple applications

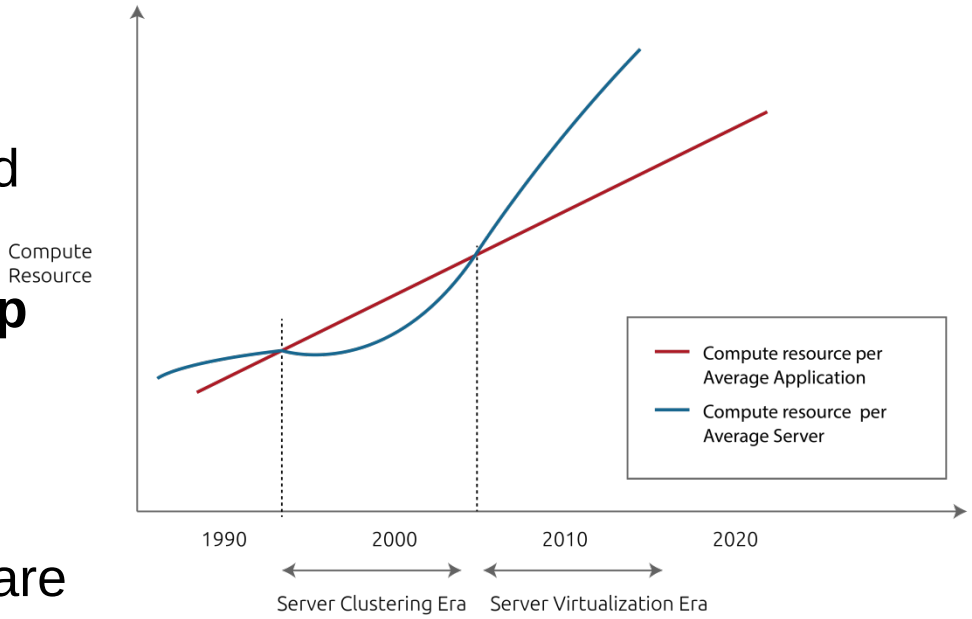

Result Grand success of server virtualizations

1M IOPS – What Do I Do with Them?

Gone are the days where you had to aggregate

- Hundreds of drives to generate 100K IOPS
- Scores of storage nodes to generate 100K IOPS

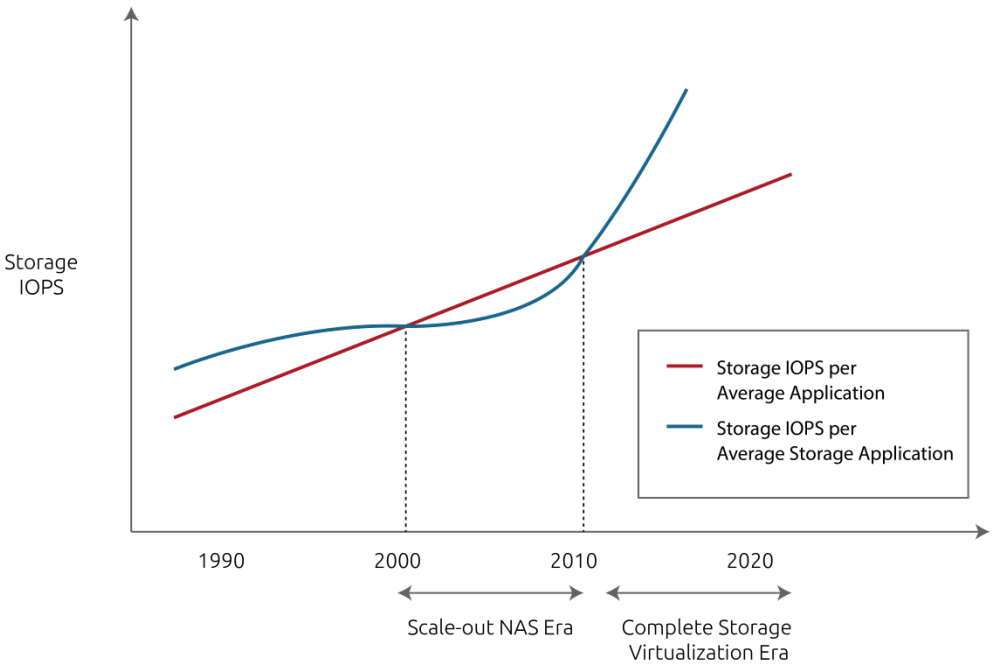

A problem that's analogous to what server world faced in the early 2000s

- With SSDs, each storage node packs tons of IOPS
- But no credible ways to share them with multiple applications

Server virtualization all over again inside storage?

Storage is NOT Just Disks

- Storage appliance has other resources such as CPU, cache, and network
- Historically, disks were the bottleneck; nobody worried about the other resources
- With SSDs, bottleneck shifted to the other resources

Now, the new form of virtualization needs to consider sharing of all the resources with multiple applications

All the storage resources need to be virtualized

Storage Needs New Virtualization

- With SSDs, each storage node is a performance power house
- Thick form of server virtualization can't be taken directly into storage. It is an overkill. It impacts performance.
- Server virtualization was designed to support heterogeneous OS and applications
- Handle homogeneous storage processing for multiple applications
- Needs a scalable architecture to support storage for 100s of applications

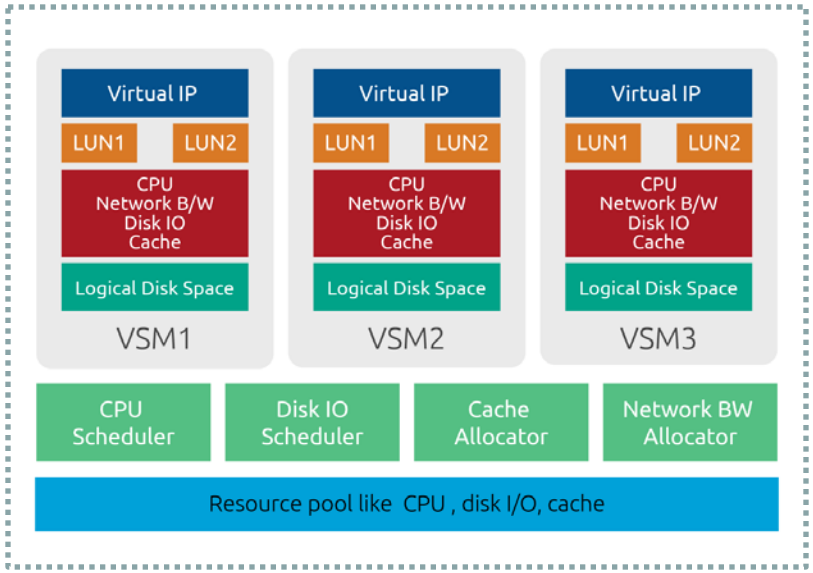

Storage Needs Intelligence

- Got scalable architecture with the thin storage virtualization.

Nope. The real complexity starts here.

100 IOPS from Oracle needs different amount of storage resources in terms of CPU, network, cache, and disk I/O than 100 IOPS from MS exchange.

The key is to determine the resource requirement for each application workload and allocate the required amount of resources.

What Makes Storage Virtualization Complex

- Server virtualization is straight forward since applications ask CPU, RAM, network, and storage I/O from the server.
- This is just abstracting the hardware resources in the software and providing vCPU, vNIC, vdisk, etc.
- However, applications ask IOPS, throughput, and latency from the storage node.
- Each application is different in terms of storage access. *1000 IOPS from MS Exchange is quite different from 1000 IOPS from Oracle.*

Performance Guarantee by Resource Allocation

- By allocating right amount of resources, each appliance feels dedicated storage environment out of the shared storage platform
- Needs an intelligent mechanism to determine the resource needed to meet the application workloads requirements

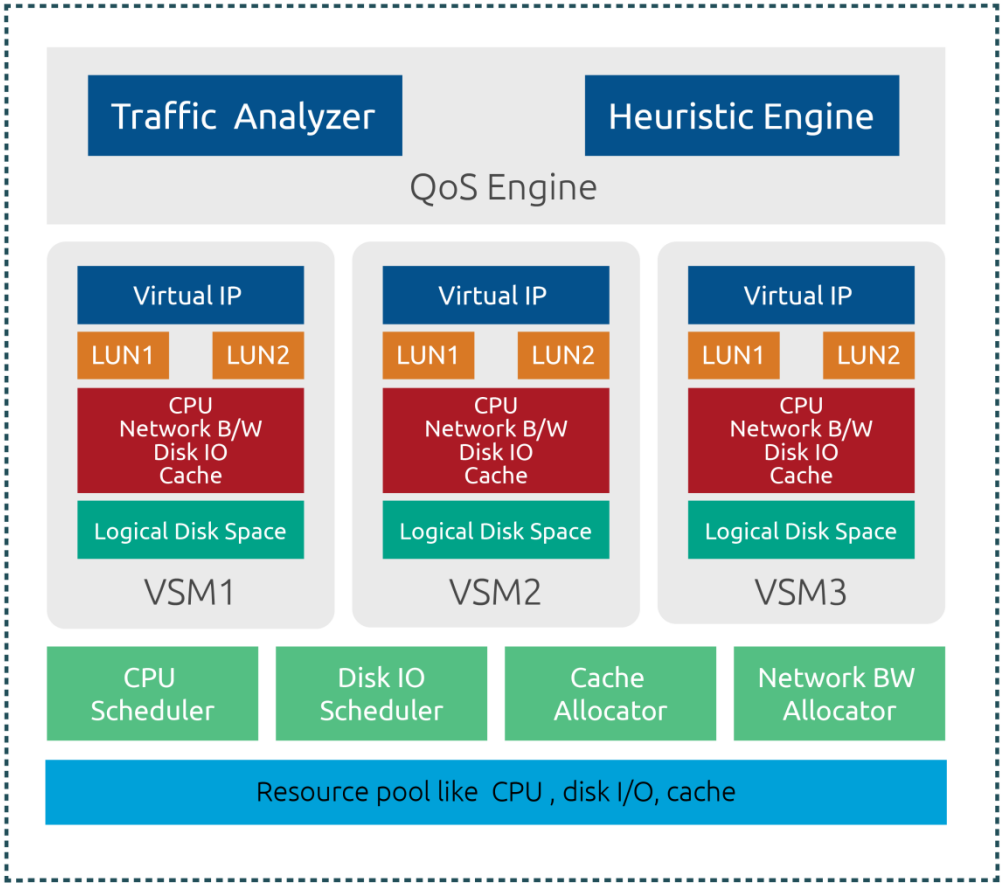

Summary

- **Storage at crossroads** Thanks to SSDs in the enterprise storage market
- **Storage needs new form of virtualization** A virtualization that abstracts the underlying storage hardware and provides the ability to move storage performance between the sharing applications on the fly

This enables the sharing of power packed all flash storage array with 100s of applications

Related Article from the Author

- Preserving Storage IO for critical applications in the virtualized environments