The Ultimate storage Virtualization

Felix Xavier
CloudByte Inc.

Abstract

- The concept of virtualization is not new for the storage. Logical Volume Manager is the earliest form of storage virtualization which allowed to create virtual volumes abstracted from the underlying hardware. Subsequently scale-out NAS virtualized the NFS mount point through virtual IP address where NFS mount point is floating above the group of hardware storage nodes. Virtual Server or commonly known as VM in the server world triggered the new form of virtualization in the storage world, it is storage virtual machines. This is in fact reverse of scale-out NAS problem. Due to the arrival of SSDs, each storage systems are becoming larger and larger in terms of storage performance, hence this hardware storage system needs to be shared among multiple applications. This brings out the need for storage virtual machine which abstract the physical characteristics of storage such as IOPS, throughput, latency and capacity into software and credibly share the larger storage system with multiple applications. This is the ultimate form of storage virtualization.

Logical volume manager

- Earliest form of storage virtualization
- Abstracted underlying drives
- Creates logical pool of storage
- Logical volumes defined on the storage pool
- Data on the volumes stripped across the underlying drives
- Provided the ability to use thin provision the storage
- Volumes are no longer constraint to small set of drives.
- This was done to aggregate both capacity and performance

virtual Mount point – scale-out NAS

- Larger Storage performance is achieved through spreading the workload across many drives.
- Performance from each nodes got limited by the number of drives it can handle.
- One large application needed storage performance more than single storage node can handle
- Scale-out NAS emerged to aggregate multiple storage nodes to provide single virtual mount point across nodes.

Why server virtualization is successful

- Applications grow along with hardware. Except a small segment of HPC application, they are big always.
- Applications delivering similar functionalities with few MBs RAM during 1990s consumes more RAM in 2000s
- However there were hardware capabilities multiplied very fast during early 2000s, out competed application requirements
- Applications were not catching up with those capabilities
- Gross under utilization of compute resources since there no credible ways to share those resources with multiple applications
- Result – grant success of server virtualizations

1M IOPS – What do I do with this?

- Those days are gone where we need to aggregate 100s of drives to generate more IOPS
- Those days are gone where we need to aggregate 10s of storage nodes to generate required number of IOPS
- We face the same problem where server world faced during early 2000s.
- With SSDs, each storage nodes packs with tons of IOPS but there is no credible ways to share with multiple application
- Result – gross under utilization of storage resources.

Storage is NOT just disks

- Storage box has other resources like CPU, cache, network as well
- Historically, disks were bottleneck , nobody worried about other resources.
- With SSDs bottleneck shifted to the other components.
- Now, new form of virtualization needs to consider sharing of all the resources with multiple application.
- Server virtualization all over again inside storage?

What makes it complex

- Server virtualization is straight forward since applications ask CPU, RAM, network, storage IO from the server.
- This is just abstracting the hardware resources in the software and provide vCPU, vNIC, vdisk, etc.
- However, applications ask IOPS, throughput, latency from the storage node.
- Each applications are different in terms of storage access, 1000 IOPS from MS exchange is quite different from 1000 IOPS from oracle.
- This make storage virtualization complex

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 and it is over kill, impacts the performance
- Server virtualization was designed to support heterogeneous OS and applications
- This will be handling homogeneous storage processing for multiple application.
- Needs a scalable architecture to support storage for 100s of applications

Storage Needs Intelligence

- Now I have scalable architecture with the thin storage virtualization, am I done?
- Real complexity starts here.
- 100 IOPS from Oracle needs difference amount of storage resources in terms of CPU, network, cache and disk IO than 100 IOPS from MS exchange.
- Key is to determine the resource requirement for each application workload and allocate required amount of resources.

Virtualization by resource allocation

- End of the day, resource allocation is more important for the successful virtualization.
- By allocating right amount of resources, each appliance feels the dedicated storage environment out of shard storage platform
- The diagram shows this architecture

Summary

- Storage is at cross roads due to the arrival of SSDs into the enterprise storage market
- Storage needs new form of virtualization
- The virtualization which abstracts the underlying storage hardware and provides the ability move the storage performance between sharing applications on the fly.
- This enables the sharing of power packed all flash storage array with 10s of application.