

VirtuCache Case Study



STANFORD
UNIVERSITY

USE CASE

- Storage Performance

LOCATION

- Palo Alto, USA

VIRTUALIZATION AND STORAGE ENVIRONMENT

- Twenty six Dell R710 physical servers running VMware were connected to shared SAS JBODs, local HDDs, and Dell Equallogic appliances

CHALLENGES

- Hadoop workloads running in VMs were running slowly due to consistently higher than 30ms VM level write latencies

SOLUTION

- VirtuCache deployed on 26 Dell R series servers
- VirtuCache was caching reads and writes to in-host 800GB Intel S3500 SATA SSD from backend iSCSI Equallogic appliances, shared SAS JBOD appliances, and local Hard Drives.

BENEFITS

Using VirtuCache, Stanford saw 6-10x improvement in write latencies for VMs that ran their Hadoop workload.

VirtuCache Improves Hadoop Performance within VMs at Stanford

Typically Hadoop workloads are run on bare-metal servers. However since Stanford's School of Medicine was 100% virtualized, and because their security, monitoring & management tools were integrated within VMware, it was easier for them to deploy Hadoop within VMs, instead of provisioning new physical servers.

The biggest challenge Stanford faced with Hadoop VMs was low throughput and high latencies for writes to disk.

Deploying VirtuCache resulted in write latencies in Hadoop reducing to an eighth of what they were before.

VirtuCache was deployed on VMware hosts along with a 800 GB Intel S3500 SSD in each host. The 800GB S3500 costs \$800 and does about 60MBps Write IOPS and 180MBps Read IOPS at under 10ms latencies from within VMs. With VirtuCache, all recent writes were written to this SSD and all frequently used data was read from this same SSD. The writes were then asynchronously synced with SAS based shared disks or LUNs on Equallogic appliances. As is standard configuration with VirtuCache, all writes to local SSD were replicated to two other SSDs in two different VMware hosts to prevent against data loss in case the local SSD or local host failed.

By ensuring that most of the storage IO happened to the local in-host SSD, we ensured consistently low latencies even for bursty write-intensive workloads from within Hadoop VMs.

As of 2015, VirtuCache is the ONLY solution in the market that can accelerate writes to SAN based storage, shared SAS storage, and local disk based storage with a kernel only deployment and without requiring any storage reconfiguration. Support for these different types of local and shared storage technologies was a requirement at Stanford and differentiated us from the competition.

As a result, deploying VirtuCache resulted in a 6-10X improvement in write latencies for Hadoop workloads from within VMs.

The VirtuCache Difference

- VirtuCache is the only solution in the market that can improve latencies for writes from VMware VMs to SAN based storage, without requiring a storage upgrade and any storage reconfiguration.
- Since Hadoop is a write intensive workload, it was write latencies that needed to be improved