

# VirtuCache Case Study



## USE CASE

- Storage Performance

## LOCATION

- USA and France

## VIRTUALIZATION AND STORAGE ENVIRONMENT

- HP BL460C G10 blades running VMware 6.5
- 256 TB of 3PAR storage accessed over 10gbps iSCSI

## CHALLENGES

- Log Management and SIEM applications running Splunk and Elastic Search were not able to analyse log and network related security in real-time.

## SOLUTION

- VirtuCache installed on each of the HP BL460c G10 blades caching 3PAR Datastores to 1.9TB of Samsung SM863 enterprise grade SSD locally attached in each ESXi host. throughput.

## BENEFITS

Using VirtuCache ensured that application level latencies for Splunk and Elastic Search applications were under 10ms regardless of very high and random throughput.

## VirtuCache Improves Splunk and Elastic Search latencies for better performance

Business Intelligence, Log Management, Security Information & Event Management (SIEM), Search and Analytic software like Splunk, Elastic Search, Cognos, HP Vertica, HP Autonomy, need to provide real-time visibility into large volumes of fast changing data. When these applications are deployed in traditional VMware VMs connected to centralized storage, such large volume of write and read operations puts pressure on existing storage infrastructure resulting in much slower than real-time ingest and analysis speeds that are expected of such applications.

Especially at service provider scale, this is a difficult problem to solve at reasonable costs. To give you an idea of scale, we are talking about multiple 32 node VMware clusters generating few hundred MBps of storage throughput per host, with requirements of sub 10ms latencies at the application level. Many all-flash arrays have problems meeting these requirements, since storage IO from all hosts connected to it terminates on the same set of controllers and inbound HBA/NIC ports. The storage network and/or storage controllers soon become the bottleneck as throughput starts to increase.

VirtuCache does a better job of tackling this problem because VirtuCache uses SSDs in each VMware host as caching media, and this SSD is connected to the host CPU on a dedicated SATA or PCIe bus. As a result a large percentage of storage traffic from each host is serviced by the local SSD. Also since SSDs are deployed on each VMware host, storage performance scales linearly as SSDs are added hosts.

An enterprise grade SATA SSD like the Samsung SM863 or Intel S3710 is more than adequate to service 100MBps random write and 200MBps random read host level throughput, at under 10ms VM latencies.

**Tellabs' use of VirtuCache to improve the performance of their Service Provider customer's SIEM and log management applications.**

Tellabs was processing large amounts of log data from various networking and security equipment on their customer's network. They would feed this data into various log management applications that would then be analyzed for security threats and vulnerabilities. They wanted near real-time visibility into any security threat in their environment, and hence the requirement to ingest and analyze data rapidly.

They have multiple 32-node VMware 6.5 clusters connected to 3PAR Hybrid appliances.

# VirtuCache Case Study



## VirtuCache Improves Splunk and Elastic Search latencies for better performance

### Before and After VirtuCache.

Before VirtuCache, VMs on each host were processing 5MBps of read and write traffic and VM level latencies were on average 50ms with many peaks of greater than 200 milliseconds. The total storage traffic hitting the backend 3PAR appliance from each 32-node cluster was around 1TBps. Because of high VM level latencies, there was 10-15 minutes of lag time between ingesting data and analyzing it, which introduced the possibility of a security breach occurring during that time.

Once VirtuCache was installed in each ESXi host and it was caching to 1.9TB of Samsung SM863 SSD in each host, each VM was now able to process 20MBps of storage IO, with VM level latencies now under 10ms at all times. Also the amount of storage IO hitting 3PAR was much reduced since almost all the reads were offloaded to the in-VMware host SSD.

10gbps network and high queue depth RAID controller are key to low latencies in a high throughput environment.

If your requirements are to index/ingest (write) and analyze (read) greater than 100MBps of data from each host and at near real-time speeds, and you are using VirtuCache to improve the performance of your existing SAN storage appliance, then you need to ensure that you are not gated by the RAID controller on the host and not gated by your IP network. VirtuCache caching to enterprise grade SATA SSDs like Samsung SM863 or Intel S3710 can easily perform at these levels. However the most common bottleneck that we run into are RAID controller queue depths and the IP network bandwidth.

Regarding RAID controller Queue Depth: Queue depth for the Raid controller is the number of I/O requests that can be processed simultaneously by the controller. A higher queue depth results in proportionately higher throughput. To ensure 10ms VM level latencies at over 100MBps throughput, you need a Queue Depth of 256 or more for the SSD. Now enterprise grade SSDs like Samsung SM863 or Intel S3710 can easily process 256 requests simultaneously, and hence their Queue Depths are high. However since the SSD is behind the RAID controller, you need to ensure that the RAID controller queue depth is higher than 256 as well, else the RAID controller will become the bottleneck.

Regarding bandwidth of your IP network: VirtuCache replicates writes to another SSD in another host in the same VMware cluster. Such mirrored writes go over your IP network. We do this so that if the local host were to fail, we can immediately sync the backend storage appliance with backup copy of writes from the peer host, thus preventing data loss in case of host failure. To ensure low latencies for such write replication, especially in a high write throughput situation, we recommend a 10gbps IP network.