

# Creating a Host Based Storage Acceleration Tier for VMware

## Table of Contents

<b>Executive Summary</b> .....	3
<b>Typical causes of poor performance in a VMware environment</b> .....	4
<b>Accelerating performance in a VMware environment</b> .....	5
<b>VirtuCache Differentiators Versus Other Caching Vendors</b> .....	8
<b>Summary</b> .....	11

## Executive Summary

Gain 3-20X performance improvement of your applications running in a VMware environment without any changes to the application or the SAN storage.

VirtuCache software can turbo charge your VM environment – now available in Australia with licensing starting at AUD\$166/month/Host on an annual plan – excluding local taxes.

Performance in VMware systems are typically hampered by the type of storage and the way in which the storage communicates with VMware and its applications.

In most virtualized data centers, storage I/O to shared storage is often a bottleneck. This is especially so if applications like relational databases, analytics, virtual desktops, search engines, and other throughput intensive or latency sensitive applications are deployed within virtual machines. Flash storage has been promoted as the solution to these bottlenecks.

It is advisable to understand the nature of the problem BEFORE opting for expensive FLASH storage arrays or undertaking a complete storage refresh.

80% to 95% of all data in most organizations is cold, stale or inactive and will probably never be used again – yet it is residing on your most expensive storage. Why?



## Typical causes of poor performance in a VMware environment

- a. **Disk Speeds:** Disk speeds and latencies have not kept up with improvements in CPU, Memory, Network and bus speeds.
- b. **Storage array controllers are a bottleneck as well.** Typically, a single pair of redundant controllers fronts a large number of disks. And there are only a limited number of disk drives and limited amount of I/O bandwidth that such a controller pair can support. Flash-based storage systems typically use SAS (serial-attached SCSI) links from their controller processors to communicate to flash. SAS relies on the legacy protocol SCSI, which was originally developed for use with slower hard-disk drives. SCSI can be used for communicating to flash, but it's not ideal.
- c. **Increased utilization of resources.** Increasing adoption of virtualization at enterprise datacenters has resulted in greater utilization of server resources, resulting in increased I/O to networked storage.
- d. **I/O blender effect** - In virtualised servers, by the time sequential I/O from each VM gets through the hypervisor, it gets interspersed with I/O from other VMs on the same server, and so the I/O pattern out of the virtualised server becomes mostly random. Since disk performance deteriorates rapidly with random I/O, this reduces storage performance further.
- e. **Latency** – this is the time it takes for a read or write to travel from the application, through the VMware hypervisor, over the network to the SAN storage controller, write the data to disk or read the data from disk, and return the acknowledgement or data back over the network, to the hypervisor and eventually back to the application. This is the #1 cause of poor performance in most VMware environments.

## Accelerating performance in a VMware environment

### 1. Using Flash arrays to solve the storage I/O bottleneck

Flash is the media that makes up the Solid-State Drive (SSD). Flash is ideal for solving storage throughput and latency issues for random workloads, as is the case with VMware workloads. However, since Flash is much more expensive than traditional hard drives, it needs to be used in smaller quantities to make the storage solution cost effective. Hence the need for caching software that backs slower disk or shared storage with smaller quantities of faster Flash based storage is more desirable.

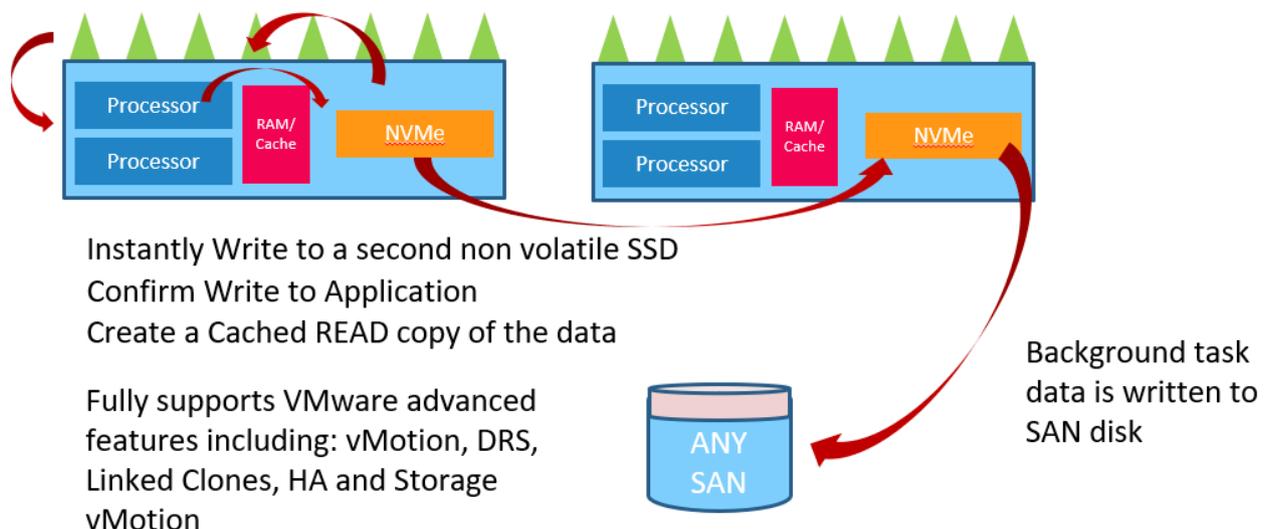
### 2. Use of VirtuCache to improve Storage Performance of VMware

Caching software in operating systems has been in use for decades. Applying these same time-tested concepts to the relatively newer domain of virtualised servers that use networked storage, with in-server Flash as the caching media.

VirtuCache is a Kernel mode software for VMware that clusters together any in-Host SSDs installed across VMware Hosts in a VMware cluster and then caches frequently and recently used data from any SAN based primary storage appliance to this clustered pool of Host based SSDs.

Subsequently, by automatically serving more and more data from in-Host SSDs, VirtuCache substantially improves storage performance for VMware from our customer’s existing storage appliance, thus improving the performance of applications running within VMs and increasing the density of VMs running on each Host, without requiring an expensive upgrade to SSD based storage appliances.

## Data Path for a WRITE with VirtuCache



## How VirtuCache solves the issues:

### 1. Accelerating Reads

All read requests from the VMs on the Host are intercepted by VirtuCache software in the VMware Kernel. VirtuCache first looks at the metadata to determine if the data is on the local SSD. If the data is in the SSD, it is served to the VM from the SSD (called 'cache hit'). If the data is not in the SSD, the I/O path proceeds along its original course and VMware retrieves the data from the backend LUN/Volume. At that point VirtuCache copies the data to the local SSD as well. Subsequently if the same data is requested again by any VM on the Host, it is now served from the local SSD, instead of from the backend storage appliance. In this way VirtuCache accelerates read operations by serving up more and more data from in-Host SSDs.

### 2. Accelerating Writes

All writes from VMs on the Host are written to the local SSD without synchronously writing to the backend storage appliance. By writing to the in-server SSD, writes are substantially accelerated, however the fact that VirtuCache is not synchronously committing the writes to the backend storage appliance introduces the risk of data loss/corruption in case the local Host or SSD were to fail. To guard against this possibility, VirtuCache protects the local cache by replicating / mirroring the writes across SSD to other Hosts in a VMware cluster.

### 3. Syncing 'Dirty' Writes to backend storage

'Dirty' writes are writes on the local SSD cache that have not yet been synced with backend storage. VirtuCache has a background task that continuously syncs 'dirty' writes to the backend SAN based storage. VirtuCache adjusts the speed and frequency at which 'dirty' writes are synced based on the latencies exhibited by the SAN and appliance, so as not to choke the SAN by trying to sync to the backend appliance quickly. Also, at no point in time will more than a few minutes of 'dirty' writes be stored on the local SSD. This is to avoid large amounts of 'dirty' writes following the VM during a vMotion.

### 4. Cache Replication to protect against local Host or SSD failure

One of the main benefits of clustering SSDs across VMware Hosts is being able to mirror the cache across VMware Hosts in a distributed fashion. The administrator specifies the number (0, 1, or 2) of copies of cache that need to be kept for each local cache in the cluster.

The number of copies indicates the maximum number of node failures that can be sustained before there is data loss in the cluster. If a customer chooses to keep, say, 2 copies of cache for each local Host based cache, VirtuCache automatically replicates the 'dirty' writes across two SSDs on two VMware Hosts in the same cluster.

We default to using the vMotion network for such replication. However, a separate network can be configured as well. Reads are not replicated since the backend storage appliance is always consistent as far as reads go.

In the event of a Host or SSD failure, VirtuCache syncs the backup copy of the 'dirty' write cache from another Host to the backend storage appliance.

## 5. Flow control to prevent Write intensive VMs from taking over the SSD

Since the SSD capacity deployed within VMware Hosts is typically a small percentage of the total LUN capacity of the backend storage appliance, care needs to be taken to prevent write intensive VMs from taking over the entire SSD.

VirtuCache allows bursty writes from VMs to be written to the SSD at native SSD write speeds without synchronously syncing the data to the backend disk. However, for prolonged write intensive activity from VMs, VirtuCache's flow control feature throttles back the write speeds to the SSD. This helps ensure fair allocation of SSD capacity to other VMs on the Host and ensures orderly de-staging of writes from the SSD to the backend LUN.

## 6. Keeping the cache 'fresh'

We use a combination of Least Recently Used (LRU) and First-in-First-Out (FIFO) algorithms to replace less frequently used older data with newer data in cache, much like how traditional operating systems have been using these algorithms for Disk-to-Memory caching.

## VirtuCache Differentiators Versus Other Caching Vendors

VirtuCache is the highest performance Host side SAN acceleration solution on the market and has the lowest management overhead. In the sections below, we explain the reasons for both:

### 1. High Performance

We accelerate both read and write operations. Most of our competition accelerates only reads.

### 2. VirtuCache is installed in the Hypervisor layer of VMware.

There are no Userspace or VM level software components that are in the I/O path. We are unique in the VMware space because we intercept block I/O requests and make all caching decisions from within the VMware Kernel.

Installation is typically a 15 minute non-disruptive and non-intrusive process that does not require a reboot.

### 3. Block level caching software.

Block level caching is faster than file level caching, since it works at a lower level in the software stack than the file system.

### 4. We can deliver 3-20X IO acceleration for Reads and Writes

Testing with TPC-C, TPC-H, Sqlio, Iometer, and Fio benchmark tests, we have proven 3-20X performance improvement.

Additionally, we often see up to 3-7X improvement in VM densities allowing substantially better hardware utilization.

### 5. Minimal Administrative Overhead

Automatic caching of data with no human involvement minimises administration costs.

### 6. We automatically cache VM level storage I/O

We use time tested LRU and FIFO algorithms to keep the cache fresh. Unlike our competition, we do not need a dedicated virtual appliance VM per VMware server, or any software installed within guest VMs.

### 7. VirtuCache does not need storage for each VM reconfigured.

No need for the VMs or the Host to be restarted. VirtuCache operation is seamless to the organization's existing SAN based storage architecture, in the sense that end users and applications running within VMs do not realize that the data is being read from and written to the local SSD, instead of the backend storage appliance.

### 8. Important applications are typically isolated within VMs.

Despite being in the Kernel VirtuCache can correlate I/O to VMs and consequently can accelerate specific applications by enabling caching for only the VMs that the applications are installed on.

### **9. I/O de-duplication**

In Linked Clone deployments, VirtuCache caches parent VM blocks that are repeatedly being requested by multiple desktop VMs only once, thus conserving SSD usage. Since O/S and application installation files are duplicated across virtual machines in VMware Horizon View VDI deployments, this feature results in high cache hit ratios by using relatively small SSD capacities.

### **10. Advanced VMware Features**

Without any administrative overhead, VirtuCache automatically support VMware's advanced features – VAAI, Snapshots, DRS, Linked Clones, vMotion, High Availability, and Storage vMotion.

### **11. VMware Certified**

To satisfy VMware's certification requirement VirtuCache developed a kernel mode caching driver as a Path Selection Plug-in (PSP) for VMware Native Multipath Plug-in (NMP), using VMware's publicly available Multi-Pathing Plug-In framework. <https://kb.vmware.com/s/article/2116221>

## Technical FAQ's

### 1. Risk of data corruption

Most organisations start with evaluating VirtuCache in a READ ONLY mode. The benefits of doing this is that a small amount of RAM is used by the metadata to determine if the data requested is located in the cache. If it is located in the cache (RAM, NVMe, SSD) it is served up to the application directly from the cache. If the data is not located in the cache, the data path is unaffected by VirtuCache. A copy of the data is also written to the cache, so that if that same data is requested again, it is served up from the cache, not the slower spinning disks.

This eliminates the risk of corruption as only READS that are being cached.

With write caching, the data is synchronously written to a second non-volatile NVMe or SSD on a second host in the cluster for resilience and redundancy and once committed to the second NVMe/SSD, it only then tells the application that the data is safe and secure. In the background, the data is written from the NVMe/SSD cache to the spinning disk.

As a write occurs, and after it has been committed back to the application, a copy of the data is held in read cache in case it is required shortly afterwards.

Suggestion: Test and evaluate in READ ONLY mode before progressing to READ and WRITE mode.

### 2. Difficulty in deployment

VirtuCache has been designed to be as simple as possible to deploy. No reboot is required. No special NVMe or SSD are required, as during testing and evaluation, you are able to use existing RAM to validate and prove the performance improvements.

No storage expertise is required, and a junior VMware administrator can normally install and configure VirtuCache in as little as 15 minutes.

### 3. Changes to applications

A key objective when developing VirtuCache is to ensure zero changes to any applications are required to test, evaluate and roll into full production.

### 4. Not having enough (NVMe/SSD) resources

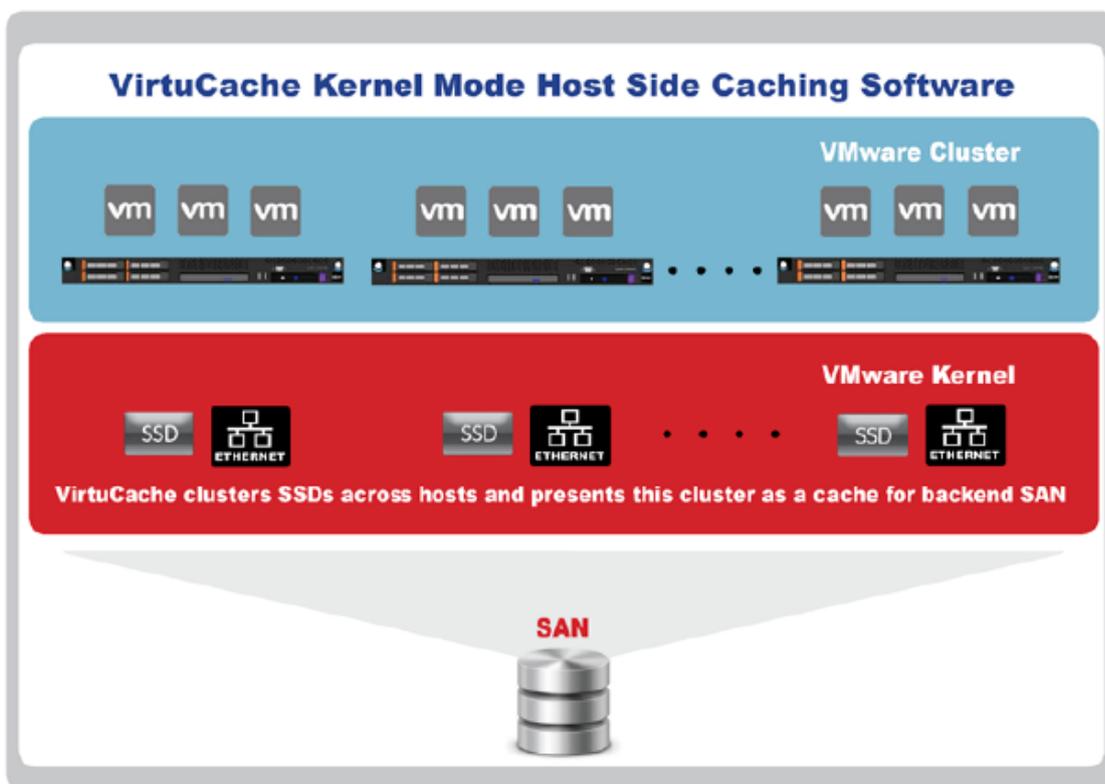
VirtuCache can be evaluated using existing RAM without the need for any NVMe or SSD's on the Host. Once you have determined the value of VirtuCache, it is a simple process to install the correct sized NVMe into the PCIe bus of your selected VMware Hosts.

## Summary

VirtuCache main design goals were to make VirtuCache high performance, near zero overhead, and VMware certified. VirtuCache will work on all current versions of ESX and once installed is transparent to all applications and supports VMware’s advanced features.

### Australia and New Zealand ONLY:

In Australia and New Zealand, VirtuCache software is available on an annual subscription of 1, 2, 3, and 5-year license. During that period the product updates and limited telephone support are available at no additional charge.



**Figure 2: VirtuCache deployment across a VMware cluster**

VirtuCache consists of two software components – Driver software installed in the Hypervisor layer of VMware vSphere Hosts that need to be accelerated, and one VirtuCache Management VM per vCenter to manage all the VirtuCache Driver software instances installed in Hosts managed by that vCenter instance.

## More information

Australia, New Zealand and SE Asia  
 vmSoftware: +61 2 8330 6717  
 e-mail: [info@vmSoftware.com.au](mailto:info@vmSoftware.com.au)  
 Web: [www.vmSoftware.com.au/VC](http://www.vmSoftware.com.au/VC)