

VMware vSphere APIs for I/O Filtering (VAIO)

First Published On: 10-09-2017
Last Updated On: 04-13-2018

Table of Contents

1. Introduction to vSphere APIs for I/O Filtering
 - 1.1. Introduction to vSphere APIs for I/O Filtering
2. VAIO Technical Overview
 - 2.1. VAIO Technical Overview
3. Storage Policy Based Management
 - 3.1. Storage Policy Based Management
4. VAIO Implementations
 - 4.1. VAIO Implementations
 - 4.2. Infinio
5. Summary
 - 5.1. Summary
6. Acknowledgements
 - 6.1. Acknowledgments
7. About the Author
 - 7.1. About the Author

1. Introduction to vSphere APIs for I/O Filtering

This section provides an introduction on how the VAIO “I/O Filters” enable VMware, and partners, to intercept and manipulate the I/O which can provide open-ended data services.

1.1 Introduction to vSphere APIs for I/O Filtering

The vSphere APIs for I/O Filtering (VAIO) were introduced in vSphere 6.0 Update 1. The VAIO framework and program were developed to provide VMware and partners the ability to insert filters for I/O into the data path of virtual machines. These “I/O Filters” enable VMware, and partners, to intercept and manipulate the I/O. This manipulation can provide open-ended data services, but thus far is limited to four use cases, two of which are currently exclusive to VMware and two which are open for partners. These use cases are: Replication; Caching; Quality of service (VMware only); Encryption (VMware only).

Prior to VAIO, partners had to resort to unsupported kernel level methods or inefficient virtual appliances to intercept and manipulate virtual machine I/O streams. This led to a variety of solutions, which were all implemented in different ways, and could potentially result in instability and operational complexity. The VAIO framework and program was specifically designed to address these challenges. Note that the program does entail a vetting and certification process for partners before a solution is officially VAIO certified. This is to ensure stability and provide a consistent user experience.

This paper provides a technical overview of VAIO. It highlights how VAIO enables the use of advanced data services without the need for expensive storage systems. The paper concludes with an overview of different filters developed by VMware and/or 3rd party software vendors.

2. VAIO Technical Overview

This section explains how I/O Filter is responsible for acknowledging the I/O to the Guest Operating System.

2.1 VAIO Technical Overview

The vSphere APIs for I/O Filtering enables the interception of I/O requests from Guest Operating Systems (GOS) to virtual disks (VMDK). This is enabled by inserting an I/O Filter in the vSphere I/O stack. This filter driver executes in user space (called User Worlds in ESXi) so that 3rd party filter code can run natively in ESXi without jeopardizing the stability of the ESXi kernel. At the same time, VAIO allows for filter drivers to intercept I/O operations as they go through the data path of the ESXi kernel without any perceived performance overhead. I/O operations are intercepted as soon as they go through the virtual SCSI emulation layer (vSCSI) in the kernel, as shown in Figure 1. As a result, data services can be applied before I/O is processed by the different storage virtualization modules of ESXi's kernel (VMFS, NFS, SCSI, vSAN, etc), which allows for generic data services that are agnostic to the storage backend. For the same reason, I/O is intercepted before it traverses the network providing security and integrity of data.

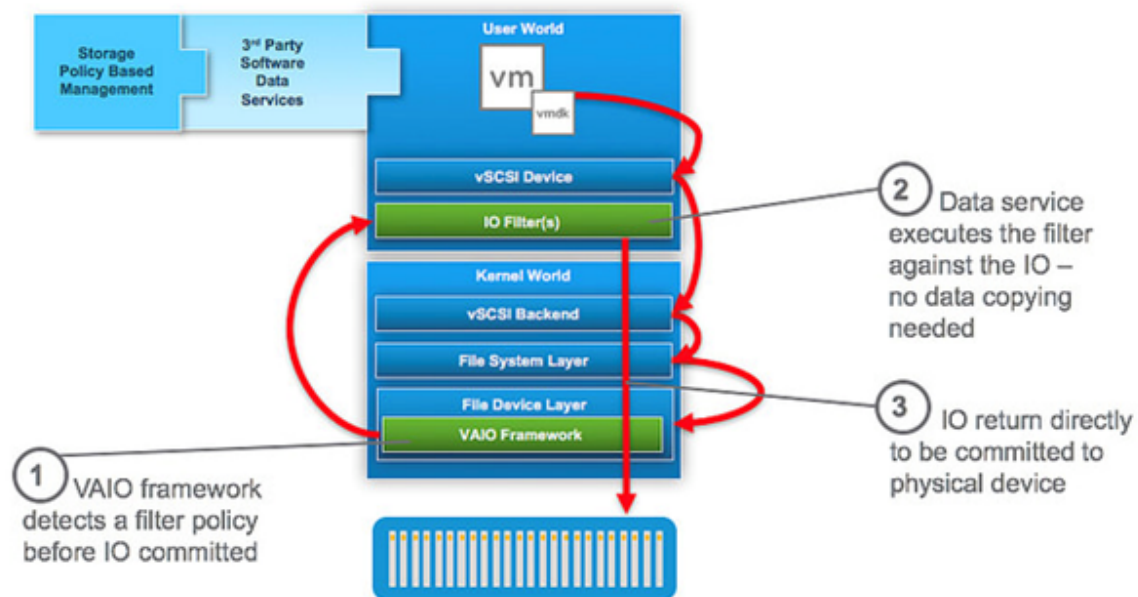


Figure 1 - VAIO Framework I/O Path

In the example above, a single I/O Filter is shown. The I/O Filter is implemented in the user space/world to avoid impacting the hypervisor should there be an issue with the I/O Filter. It is also possible to have multiple I/O Filters applied to the same VM or VMDK(s). If multiple I/O Filters are applied to the same VM or VMDK, the order in which it is applied is defined by the filter "class order" which is defined as part of the VAIO framework. This means, for instance, that a replication filter executes before a cache filter. Once the I/O has moved through all I/O filters, it moves to the next layer.

The I/O Filter not only intercepts and manipulates I/O depending on the implemented functionality, but it can also be responsible for acknowledging the I/O to the Guest Operating System. It is a bi-directional information and traffic flow. If an I/O Filter is written to do write caching, this can be implemented in a write-back or write-through modus. Depending on the implementation, the I/O Filter needs to be informed that the write data is committed to a storage device before placing it in cache.

When a replication I/O Filter is implemented, the replication can be done either synchronously or asynchronously. Synchronous replication includes the latency to reach the remote storage in the critical path of a write operation—both copies of the operation must be persisted before the operation is acknowledged to the Guest OS. Asynchronous replication is used when the additional latency is not acceptable. This means that the acknowledgement from the local storage system enables the VM to continue, while the I/O is (asynchronously) being replicated to a different location. Note that in this case, the successful asynchronous replication of the I/O is still acknowledged to the I/O Filter. The I/O filter will keep track of each I/O request on a per VMDK basis to ensure data consistency and integrity.

VMware vSphere APIs for I/O Filtering (VAIO)

Besides regular I/O, VAIO is also informed when certain control operations are performed. For instance, when a snapshot is requested by someone (or something), or a change to a VM or VMDK is requested (grow disk, stop VM, start VM, etc.) the I/O Filter is informed about this event so that this change can be taken in to account.

Requirements for implementing I/O Filters usually depend on the type of filter being used. There are, however, a couple of common requirements before an I/O Filter can be installed and configured:

- vSphere 6.0 Update 1 or higher
- vMotion with attached filters is supported, but requires the I/O Filter to be installed on the target host
- DRS is supported, but requires the I/O Filter to be installed on all hosts in the cluster
 - Note that with vSphere 6.0 U1 it was required to have DRS enabled, as of vSphere 6.0 U2 this is no longer the case.
- Permissions to install a VIB (vSphere Installation Bundle)

Later in this document we will look at how to install and configure an I/O Filter.

Now that we know more about the internals of I/O Filters, how are these applied?

3. Storage Policy Based Management

This section explains ways on how to apply an I/O Filter to a VMDK or a VM using a VM Storage Policy.

3.1 Storage Policy Based Management

I/O Filters are applied to VMs and/or VMDKs through policy. The framework which is used for this is the Storage Policy Based Management framework, also known as SPBM. In order to apply an I/O Filter to a VMDK or a VM, a VM Storage Policy will first need to be created. In order to create one, you must have the appropriate permissions.

For those who are familiar with Storage Policy Based Management and the creation of a VM Storage Policy, the creation of a policy with an I/O Filter is slightly different. When an I/O Filter is installed, typically the I/O Filter will include an example for a policy, very similar to how this is done for VMware vSAN and VVols. This policy typically will contain a pre-selected “common rule” which is where the data services are described. These “common rules” are defined through Storage Policy Components. By default, various Storage Policy Components are pre-created from vSphere’s built-in filters, and more are added as I/O Filters are installed. A vSphere 6.5 system without any external I/O Filters installed will come with Storage I/O Control and VM Encryption storage policy components as shown in Figure 2.

Name	Description	Category	VC
Default encryption properties	Storage policy component for V...	Encryption	so-rdops-vm15-dhcp-57-189.eng.vmware.com
Low IO shares allocation	Storage policy component for Lb...	Storage I/O Control	so-rdops-vm15-dhcp-57-189.eng.vmware.com
High IO shares allocation	Storage policy component for HI...	Storage I/O Control	so-rdops-vm15-dhcp-57-189.eng.vmware.com
Normal IO shares allocation	Storage policy component for M...	Storage I/O Control	so-rdops-vm15-dhcp-57-189.eng.vmware.com

Figure 2 - Storage Policy Component

These storage policy components can then be used in VM Storage Policies. Notice that vSphere will create only these storage policy components, instead of creating full VM Storage Policies with a rule-set, similar to what is done with VMware vSAN and VVols. In the current release of vSphere, a VM or VMDK can only have a single policy assigned and applied. By using these Storage Policy Components, I/O Filter capabilities can be included in any VM Storage Policies. In other words, by using Storage Policy Components we can enable, for example, filter-based capabilities such as VM Encryption on a VVol based VM, which otherwise would not be possible were the filters represented as full policies rather than policy components.

Of course, it is also possible to create a VM Storage Policy which only contains a Storage Policy Component; this would be useful for VMs stored on traditional storage where Storage Policy Based Management is tag-based and tends not to be as prevalent. Figure 3 shows a VM Storage Policy with multiple Storage Policy Components included.

VMware vSphere APIs for I/O Filtering (VAIO)

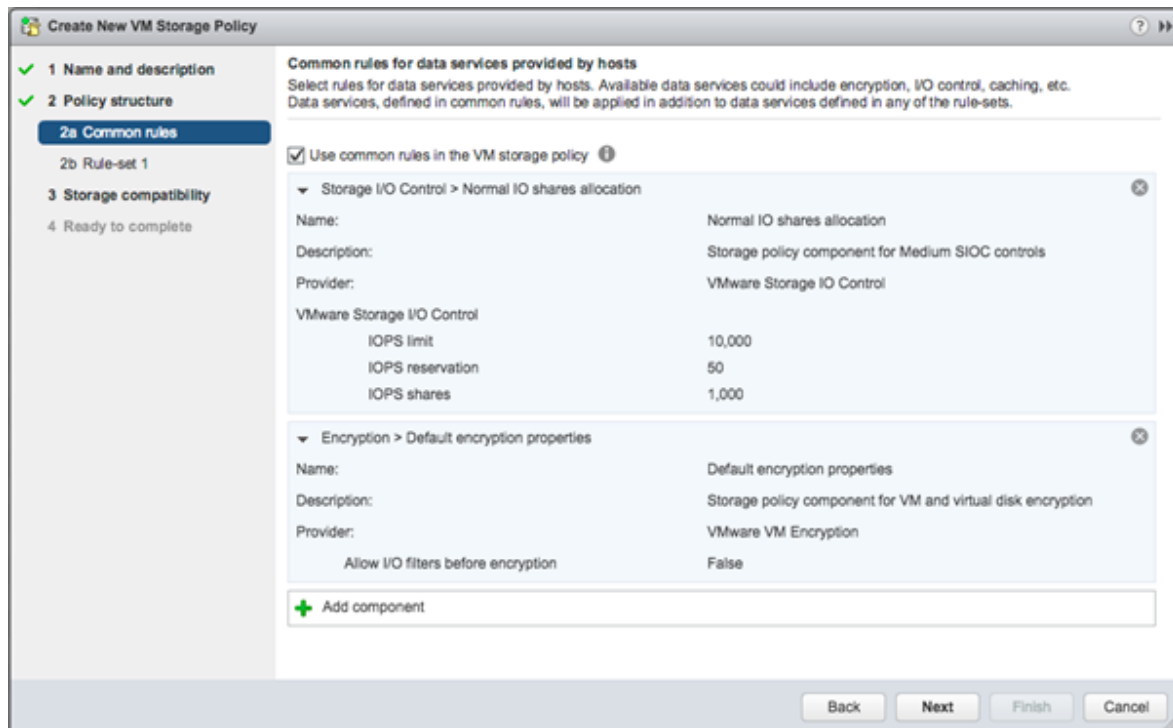


Figure 3 - VM Storage Policy with Components

After the creation of one (or more) VM Storage Policies, this policy can simply be applied to a VM or virtual disk. When the policy is applied, the I/O Filter will be active for that particular VM or VMDK.

4. VAIO Implementations

In this section, we show an example of an I/O Filter, how it is installed, configured and applied to VMs or VMDKs.

4.1 VAIO Implementations

Thus far we have described I/O Filters in a fairly abstract manner. In this section, we will show an example of an I/O Filter and how it is installed, configured and applied to VMs or VMDKs. Before buying or installing a VAIO I/O Filter VMware recommends to verify in the VMware Compatibility Guide if the I/O Filter is supported for the planned vSphere version and the storage platform. Note that in some cases I/O Filters are only certified for specific types of storage (NFS, VMFS, vSAN or VVol based) as shown in Figure 4. <https://www.vmware.com/resources/compatibility/search.php?deviceCategory=vaio>

Release	Model Version	Filter Class	VIB Version	Features
ESXi 6.5 U1	4.0.0	Cache	sngac4.0.0.1718	View
Feature Category		Features		
DATASTORE		NFS, VMFS5, vSAN		
ESXi 6.5 U1	4.0.0	Cache	sngac4.0.0.1618	View
ESXi 6.5	4.0.0	Cache	sngac4.0.0.1718	View
ESXi 6.5	4.0.0	Cache	sngac4.0.0.1618	View
ESXi 6.0 U3	4.0.0	Cache	sngac4.0.0.1718	View
ESXi 6.0 U3	4.0.0	Cache	sngac4.0.0.1618	View
ESXi 6.0 U3	4.0.0	Cache	sngac4.0.0.1397	View
ESXi 6.0 U2	4.0.0	Cache	sngac4.0.0.1718	View
ESXi 6.0 U2	4.0.0	Cache	sngac4.0.0.1618	View
ESXi 6.0 U2	4.0.0	Cache	sngac4.0.0.1397	View

Figure 4 - VMware Compatibility Guide and Storage Types

4.2 Infinio

The I/O Filter we will look at is a caching solution. This I/O Filter is by Infinio, who in their first versions mainly focused on acceleration of NFS based storage. Depending on the version of their accelerator implemented, different types of storage are supported. When we wrote the first version of this document version 3.2.0 only supported VMFS and vSAN. However, Infinio has recently completed the certification of 3.3.0 against ESX 6.5 U1 and it does support VMFS5, vSAN and VVols as shown in Figure 5.

Release	Model Version	Filter Class	VIB Version	Features
ESXi 6.5 U1	3.3.0	Cache	iniaccel3.3.0.1858	View
Feature Category		Features		
DATASTORE		VMFS5, vSAN, vVOL		

Figure 5 - VCG for Infinio Accelerator

VMware vSphere APIs for I/O Filtering (VAIO)

The Infinio installation uses a Microsoft Windows executable that provides an installation wizard. The installation wizard will ask for vCenter Server credentials, chosen host and datastore for deployment and network details. Once all information has been provided, a virtual appliance is installed and configured that provides the management console. The management console is then used to assign resources to the caching tier. Before you can create a caching layer, a flash resource will first need to be claimed for Virtual Flash usage which is shown in Figure 6.

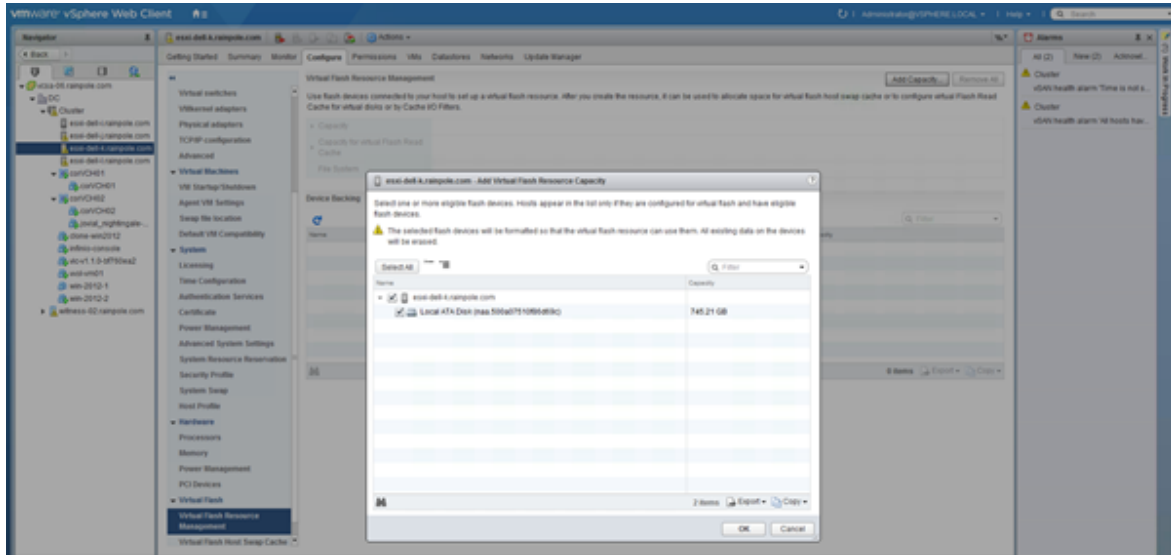
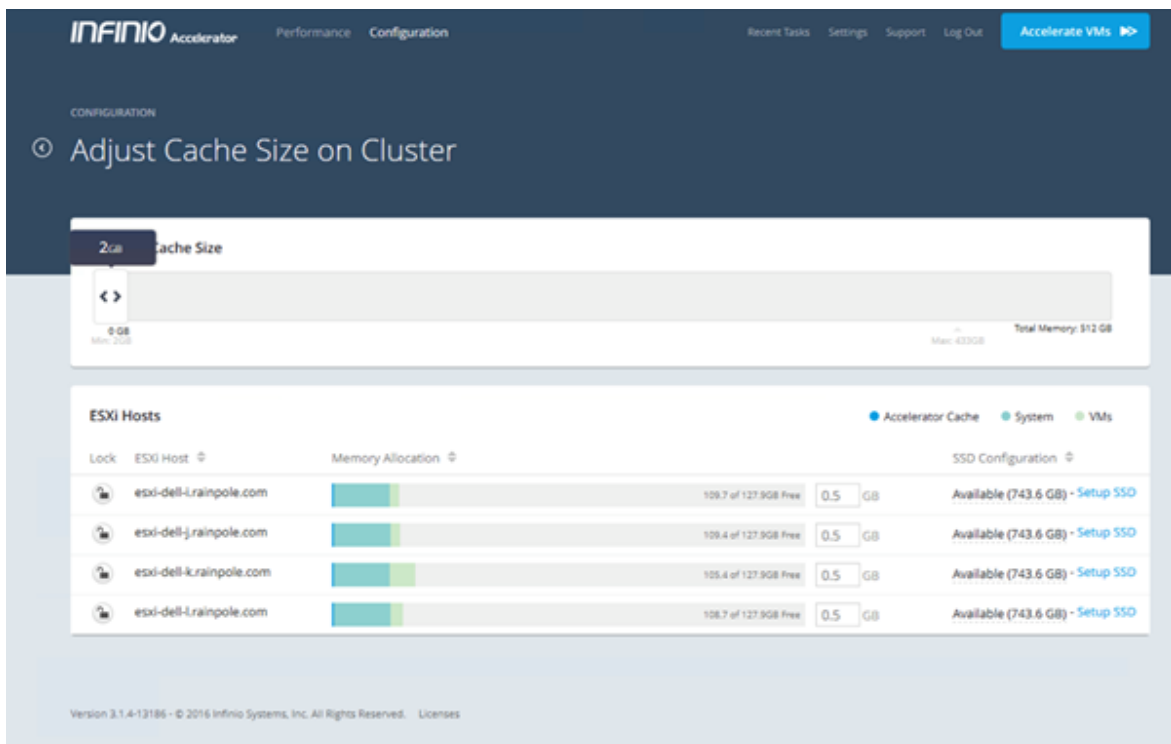


Figure 6 - Claim flash resources for Virtual Flash

During the configuration memory capacity and flash capacity can be assigned to the caching tier. Note that you could also assign only memory for caching. In our case we have selected both memory (0.5GB) and flash (743GB) for accelerating the workloads as shown in Figure 7.



VMware vSphere APIs for I/O Filtering (VAIO)

Figure 7 - Assigning cache capacity

In order to accelerate the workloads, a policy will need to be defined and assigned. The first step will be validating the I/O Filter has been configured. This can be done in the “I/O Filters” section under “Configure / Storage” on a host as shown in Figure 8. Note that this was done by the installation wizard and no manual work was required from the user / administrator.


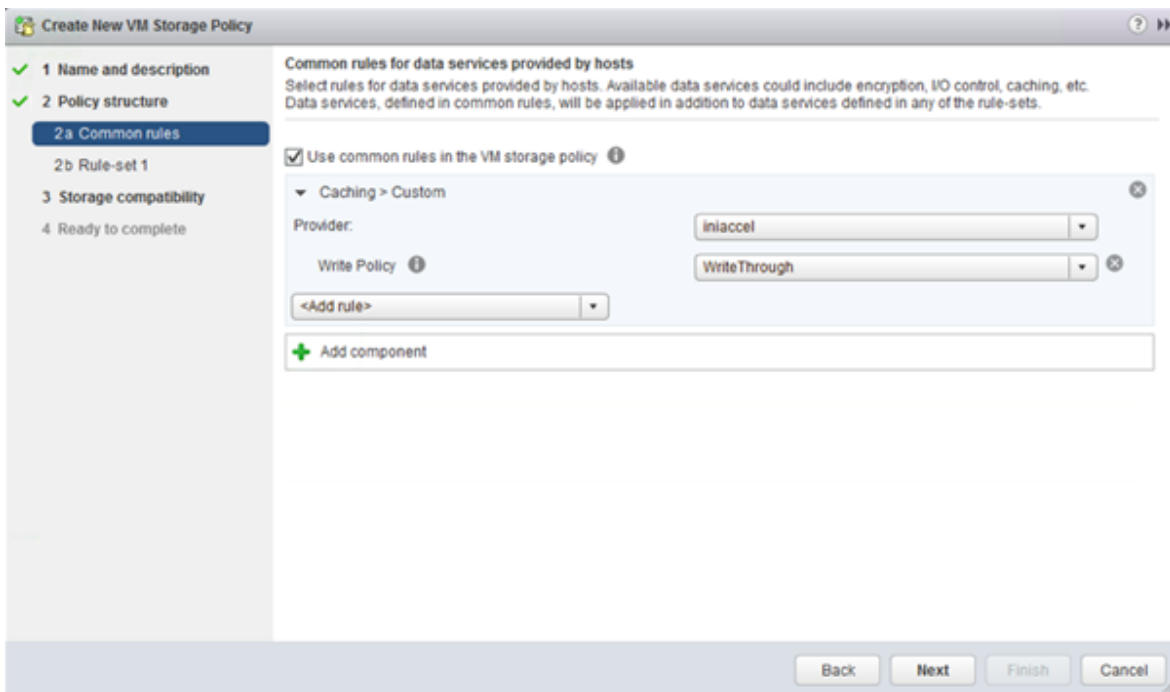


Figure 8 - I/O Filter configuration

Next a VM Storage Policy will be defined. During the creation of the VM Storage Policy a Common Rule will need to be selected. Select the Infinio Provider and next select the Write Policy.



Lastly, the policy will need to be applied to a VM, and then the workload acceleration will automatically occur. In the Infinio management console you are able to see for instance what the cache hit ratio is for your VMs, or the Top 10 VMs with the most cache hits as shown in Figure 9.

VMware vSphere APIs for I/O Filtering (VAIO)

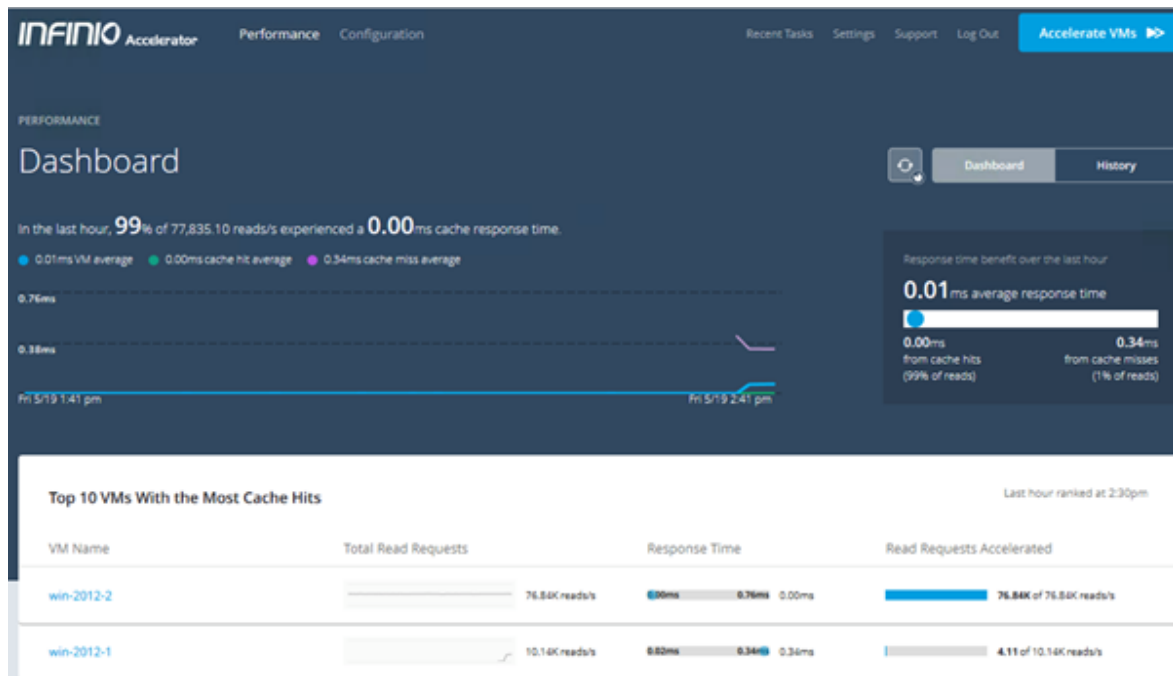


Figure 9 - Top 10 VMs

5. Summary

This section summarizes how storage vendors can speed up VMware® I/O operations that are more efficiently accomplished in storage hardware.

5.1 Summary

Using VMware's Storage Policy Based Management framework, virtual data services (I/O Filters) can be applied on a per-virtual disk or per-virtual machine level of precision. Ultimately this enables customers to provide a flexible Software Defined Data Center where availability, security and performance can be enhanced independently of the storage system.

Using VMware's Storage Policy Based Management framework, virtual data services (I/O Filters) can be applied on a per-virtual disk or per-virtual machine level of precision. Ultimately this enables customers to provide a flexible Software Defined Data Center where availability, security and performance can be enhanced independently of the storage system.

6. Acknowledgements

This section acknowledges the contributors who have helped the author.

6.1 Acknowledgments

Thank you to Cormac Hogan, Ken Werneburg and Christos Karamanolis for reviewing (and contributing to) the contents of this paper.

Thank you to Cormac Hogan, Ken Werneburg and Christos Karamanolis for reviewing (and contributing to) the contents of this paper.

7. About the Author

This section covers the details about the author of this guide.

7.1 About the Author

Duncan Epping is a chief technologist in the office of the CTO at VMware in the Storage and Availability business unit. He was among the first VMware certified design experts (VCDX 007). He is the co-author of several books, including the *VMware vSphere Clustering Deepdive* series and *Essential Virtual SAN*. He is the owner and main author of leading virtualization blog [yellow-bricks.com](http://www.yellow-bricks.com).

Follow Duncan's blogs at <http://www.yellow-bricks.com>.

Follow Duncan on Twitter: [@DuncanYB](https://twitter.com/DuncanYB).