

VM Aware Fibre Channel

Virtual Machine Traffic Visibility for SANs



StorFusion™ VM-ID feature on QLogic Gen6 (32G) and Enhanced Gen5 (16G) Fibre Channel

KEY BENEFITS

Increases Visibility of VM traffic: VM-ID tags can be used by the Hypervisor and FC switches to understand the data flowing across its environment to better monitor outages or determining heavier traffic at different points along the path to help mitigate these situations.

Allows End-to-End Quality of Service (QoS): Ability to apply specific levels of QoS on a per application basis to direct Fibre Channel (FC) traffic from a specific VM through the fabric and onto the end storage device.

Improves Usage of Storage Devices: Target systems with the capability to utilize VM-ID can optimize the proper performance level of storage device (Flash-based or HDD-based) to be best utilized and match the need of the VM.

EXECUTIVE SUMMARY

Mysteries of the world, such as how the Monarch butterflies can find their way migration paths all the way back to their species' origination point even though they had never been there before, but these occurrences in nature should remain a complete mystery. With VM clusters generating an increasing amount of FC traffic that crisscrosses across SANs within enterprise/datacenter ecosystem, these paths and accesses should not remain a mystery any longer.

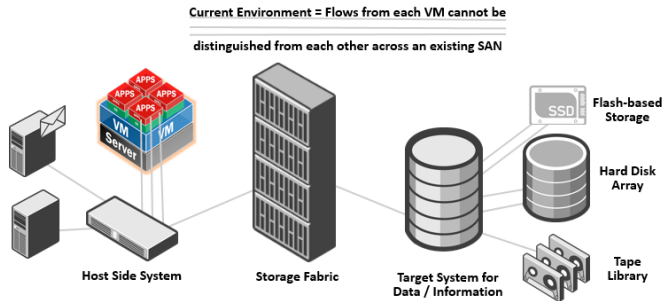
With QLogic StorFusion VM-ID technology the unknown can now be come the known and thereby help improve customers' experience by providing a more fully informed understanding of how each storage devices' usage levels is being accessed by every VM in the network. Increased visibility to better monitor the health aspects of each VM connection from end-to-end can assist service providers with better ways to fix problem areas and prevent degradations within their environments

QLogic VM-ID is supported with Gen 6 32 G and Enhanced Gen 5 16 G Fibre Channel HBAs. It works with Brocade's VM Insight product to provide our mutual customers, the ability to effectively monitor and manage their Fibre Channel storage networks, load balance VM clusters and storage and ensure efficient use of the storage resources. Efficient management of storage resources allows customers to avoid costly upgrades and still meet SLAs demanded by applications.

SERVER VIRTUALIZATION AND THE EXISTING SAN

Server virtualization has allowed for greater usage of a common link, such as with Fibre Channel, to be shared among a larger number of virtual machines (VMs). While this helps to better fully utilize the maximum available bandwidth along with CPU cores, memory, and other system resources.

The data being sent from any VM along with other physical systems become mixed together and all the traffic that travels across the Storage Area Network (SAN) appear the same and cannot be seen individually.



VIRTUAL MACHINE IDENTIFICATION FOR SANs

By utilizing VM-ID from QLogic, an end-to-end solution which uses frame tagging to associate the different VMs and their I/O flows across the SAN. QLogic has enabled this ability on our latest Gen6 and Enhanced Gen5 FC Host Bus Adapters (HBAs) and when used in parallel with Brocade's VM Insight [a part of its Fabric Vision feature set] this technology has a built-in Application Services monitor which works by gathering the globally unique ID from the hypervisor level, such as from VMware ESX, and can then interpret the different IDs from every VM to perform intelligent monitoring and give the potential to apply QoS policies to each VMs traffic from one end to the other.

With the VM-ID / VM Insight technologies working together brings deep level vision to how and where the I/O flows are originating at which VM then through the fabric and SAN managers will gain the ability to control and direct application level services to each virtual workload within a QLogic and Brocade Fibre Channel environment.

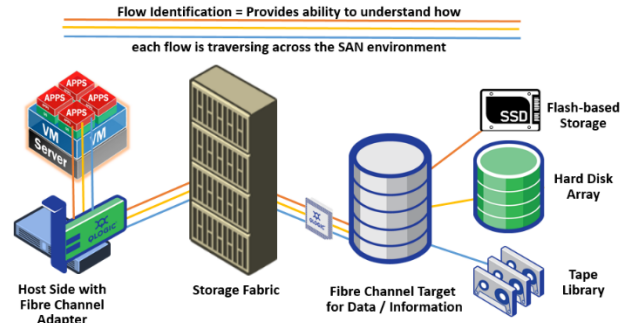
APPLICATIONS AND USES CASES

There are many different and varieties of applications and uses cases for VM-ID. As the technology becomes more fully ingrained into more products, systems and environments the potential benefits will be realized to their greatest extent. Further, we will discuss a few of the top uses.

INCREASED VISIBILITY ACROSS THE ENVIRONMENT

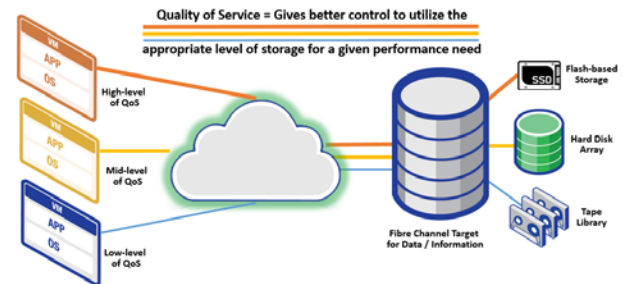
Working in conjunction our switch partner, Brocade, the VM-ID capability can be used with their product feature called VM Insight, which is a part of their

Fabric Vision technology. By having the ability to view how different flows from even a single VM can be sent to the associated tiered storage for proper handling. Additionally, this allows for improved troubleshooting and reassurance to know when and where an issue has occurred.



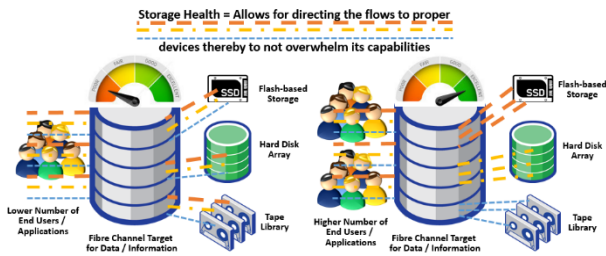
QUALITY-OF-SERVICE FOR ANY APPLICATION

VM-ID allows the end application which requires more storage performance to access the end device that can support its needs and other applications that needs lower performance to be directed to its required storage type. This helps to ensure how each storage device can be utilized to the best of its capabilities.



IMPROVED STORAGE USAGE AND PERFORMANCE

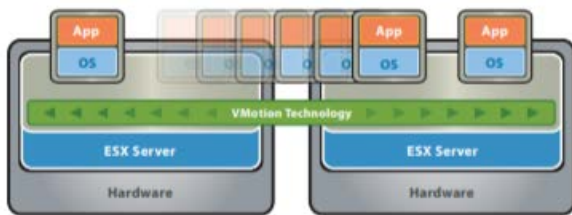
Managing the different needs of end users or related applications to the respective storage was not possible before VM-ID. The overall health of the storage target could be poor because of misusing the right storage for the actual performance needs. With VM-ID, this can be solved by granting the needs for higher IOPs or bandwidth to the devices that have the ones with the best ability, as well as, for slower or higher capacity backups should be sent to tape storage.



Load Balancing Clusters using VMware VMotion and VM-ID

VMware ESXi host clusters run workloads inside the virtual machines. ESXi clusters share networked storage and are connected to same network for management and VM traffic. The virtualization layer and shared storage and networking, makes the VMs mobile within the cluster. If the ESXi servers fail, the VMs can be quickly restarted on another server in the cluster within minutes and mostly automatically.

VMware VMotion technology allows planned movement of VMs from one server to another without any downtime whatsoever. The entire VM is copied from one server to another without any loss of data or any break in processing. Even the memory state is copied so the VM can execute on the new server exactly where it left off.



VMware vCenter monitors CPU and memory utilization of each VM on each server within a cluster. The Dynamic Resource Scheduler (VMware DRS) uses this data to automatically move VMs from a highly utilized server to a relatively less utilized server within the cluster. This keeps the hosts in the cluster load balanced.

Unfortunately, VMware DRS has a major limitation. While DRS monitors CPU and memory utilization, it does not monitor storage use or Fibre Channel storage network utilization. So DRS ensures that the clusters are load balanced from the CPU and memory perspective, but not from the perspective of the storage network utilization. VMs that are performing IO at a high rate can be concentrated on a single ESXi server. This situation ends up creating storage bottlenecks and reduce overall

performance of an application running inside a VM.

Without VM-ID implementation, IO from all the VMs on the host is combined into a single stream and sent to the switch in the fabric. There is no way for the switch to know which VM is sending which packet. In fact, the switch may not be aware that there are more than one resource using the port. Without this information, there is no way to monitor traffic to gain visibility into the storage traffic patterns.

QLogic StorFusion™ VM-ID technology working with Brocade's VM Insight tags each packet with a VM identifier, so that storage and VM administrators gain visibility into the traffic patterns for each VM. With each IO frame tagged with VM identify, administrators gain critical visibility that is used to:

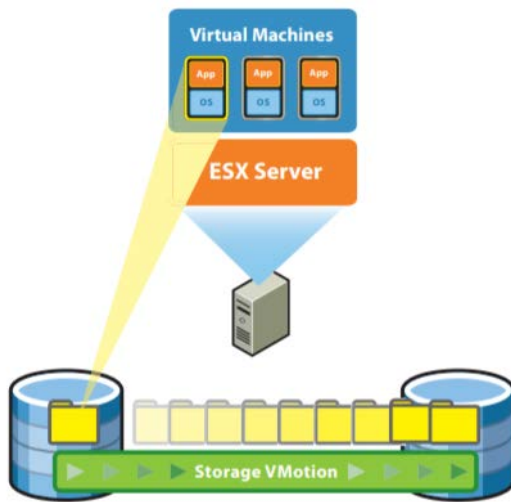
- Determine IO imbalance within the cluster at VM level: VM-ID enables counting of all IO in the cluster and attribute it to the VMs performing the IO. If the most chatty VMs are grouped on a few servers, VMotion can be used to move some of them to less used servers. This should alleviate bottlenecks from forming and increase IOPs for VMs running in the cluster.
- Determine actual IO latency from VM to the switch: Because of the tagging of each packet with a VM identifier, the switch can calculate actual observed time it took the package to originate from the VM to reach the switch. Increased latency is an early indicator of Fibre Channel becoming a constraint, which could lead to packet loss, degraded performance etc.

Load Balancing Storage LUNs

LUNs are shared with multiple VMs. Mostly, storage administrators assign LUNs of large size to the cluster and these LUNs are then shared by multiple Virtual Machines. When the VMs are reading and writing data, they are not aware of any other VMs trying to do the same. This can cause spike in LUN usage that cannot be controlled.

¹ https://www.vmware.com/pdf/vmotion_datasheet.pdf

Using VM-ID and VM Insight, administrators can monitor all the traffic going to a particular LUN as well as the VMs where the traffic is



originating. This visibility allows storage and VM administrators to monitor for hot spots on the LUNs. When two or more VMs that are demanding higher IO than other VMs, they can be moved using VMware's Storage VMotion technology. Storage VMotion, as the name suggests, moves running VMs from one LUN (VMFS) to another without any downtime. VMs and applications keep reading and writing data without interruption or loss.

Once the move is done, VM-ID continues to tag new packets with the VM identification for continued monitoring and analysis.

²http://www.dell.com/downloads/global/solutions/vslc/vmware_storage_vmotion_datasheet.pdf

SUMMARY

Cavium QLogic VM-ID technology strengthens the datacenter capabilities for assisting virtualized infrastructures to monitor, manage and control the separate flows from end-to-end. Business applications that require extremely fast IOPs can be augmented by ensuring it is only accessing flash-based storage. While other processes that do not require as fast but still perform critical

operations are met with the lower tier storage devices. The overall health of the storage end is maintained at more optimal levels by having this ability to use the correct storage for the right performance level. Quality-of-Service can further help datacenter operators by giving them the capability to service end customers with a better path to utilize a mix of storage that meets their exact needs.



Follow us:      

Corporate Headquarters: Cavium, Inc. | 2315 N. First Street San Jose, CA 95131 | 408-943-7100

2018 Cavium, Inc. All Rights reserved. ThunderX2 is registered trademarks of Cavium, Inc. All rights reserved worldwide. QLogic Corporation is a wholly owned subsidiary of Cavium, Inc. Cavium, QLogic are registered trademarks or trademarks of Cavium Incorporated, registered in the United States and other countries. All other brand and product names are registered trademarks or trademarks of their respective owners.

"Products are warranted only to meet Micron's production data sheet specifications. Products, programs and specifications are subject to change without notice. Dates are estimates only. ©2018 Micron Technology, Inc. All rights reserved. All information here is provided on an "AS IS" basis without warranties of any kind. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. NVMe is a trademark of NVM Express, Inc. PCIe is a registered trademark of PCI-SIG. All other trademarks are the property of their respective owners."

This document is provided for informational purposes only and may contain errors. Cavium reserves the right, without notice, to make changes to this document or in product design or specifications. Cavium disclaims any warranty of any kind, expressed or implied, and does not guarantee that any results or performance described in the document will be achieved by you. All statements regarding Cavium's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.