



Gridstore Storage Quality of Service (QoS) Precision I/O Control per VM

Gridstore delivers a Storage QoS that goes beyond the basic QoS features of WS2012/Hyper-V and greatly simplifies the management of Storage QoS across a set of VMs. This document outlines the basic features of Microsoft QoS and the differences of Gridstore QoS.

Key Benefits

- Priority Classes deliver QoS simply
- Guarantee performance to high priority VMs
- Manage on a per-VM basis

Summary

QoS Features	Microsoft QoS	Gridstore QoS
Priority Classes	Not available	VMs can be grouped into Priority Classes to simplify management of QoS. Up to six classes can be configured.
Per VM MAX-Normalized IOPS	VHDX level only. Each VHDX MAX must be set by administrator.	Set max IOPS per VM. Ensures VM does not exceed assigned max IOPS. This is aggregate per VM (includes multiple VHDXs).
Per VM MIN-Normalized IOPS	VHDX level only. Each VHDX MIN must be set by administrator. Notification only if threshold is not met. Can not enforce QoS MIN.	Available in V3.6 expected for Q4-2015. A VM can be assigned MIN IOPS. Ensures VM receives a minimum specified number of IOPS. This is aggregate per VM (includes multiple VHDXs).

Microsoft WS2012R2 Hyper-V Storage QoS

- **Priority Classes:** Microsoft does not have the concept of Priority Classes. Only MIN and MAX IOPS per VHDX. There is no automated provisioning of QoS.
- **MAX IOPS:** Microsoft QoS puts a MAX IOPS on a VHDX. This guarantees that IOPS for a specific VM's VHDX will not exceed the specified number of 8K normalized IOPS. This addresses the problem of a runaway VM that consumes storage resources at the expense of other VMs.

To configure MAX IOPS, the MAX needs to be set for each VHDX created. This is a tedious and time-consuming task. It is also error prone—if one is missed it can lead to problems if a VM consumes excessive resources.

- **MIN IOPS:** Microsoft does have the ability to set a MIN IOPS (a reserve number of IOPS for a VHDX). However, because Microsoft only operates in the hypervisor and does not control the storage, it is not in a position to enforce a MIN IOPS to a VHDX. Microsoft will alert when a MIN threshold is not reached.

Gridstore Storage QoS: Tiered QoS System

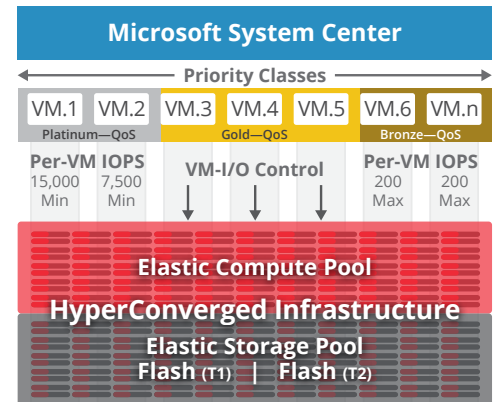
Gridstore implements QoS on two levels that can operate concurrently and are designed to simplify management of QoS. The first is priority classes. This can be an automated process that greatly simplifies the management of QoS by creating a prioritization hierarchy into which VMs are automatically placed when they attach to storage resources in the Grid. The second level of QoS is per-VM allocation. This tier provides for more granular control over the IOPS allocated to each VM.

1) Priority-Class Based

- This model works on a prioritization basis where higher-class VMs will receive more I/O than lower-class VMs when there is contention for a storage resource. There can be up to six priority classes (typically only three classes are used).
- If there is no contention for the storage resources, I/O will be served as they are received. In other words, if there is a backlog for a storage resource, I/O from higher classes will be served first. As the higher-class queues are worked down, priority will be served to the next highest class until the last class is reached.
- In this model, it is possible for the lowest class VMs to be starved of resources where there is contention. The VMs in the lower class will wait until higher classes are served.

AUTOMATED PROVISIONING

- When a VM is provisioned, it is allocated storage from a vStore. A property of a vStore is a priority class. When a VM is first attached to a vStore, it will automatically inherit this attribute and be placed in to the same priority class.
- This process eliminates the need to manually set QoS for each VM.
- The admin can at any time set the QoS for a VM to be higher or lower. The effect of this change is immediately reflected in the data path.
- This classification works on ALL VHDXs related to the VM.



BENEFITS

This is the simplest approach to Storage QoS. Most organizations have an 80/20 split in workloads: 20% of workloads are very important, 80% is a mix of lower priority workloads. This model provides the simplest way to provision VMs with a prioritization attached to them that ensure the 20% of workloads that are important get the storage resources they demand.

2) Per VM I/O Control

- Gridstore also implements Storage QoS at a more granular per-VM level. This allows admins also to define either a MIN or MAX IOPS for a specific VM.
- MAX IOPS can be set to ensure a VM does not exceed a specific number of IOPS. This is an aggregate amount for the VM. If the VM has multiple VHDXs, then this MAX is for the total I/O to all VHDXs owned by the VM.
- MIN IOPS can be set to guarantee a minimum number of IOPS are reserved for a VM. This is an aggregate amount for the VM. If the VM has multiple VHDXs, then this MIN is for the total I/O to all VHDXs owned by the VM.
- If the MIN IOPS are not utilized by the VM, they are returned to the pool for other VMs.
- To allow these two QoS models to coexist, MIN IOPS can only be set on VMs in the highest priority class. MAX IOPS can only be set on VMs in the lowest priority class.
- MIN IOPS is a new feature that will be available in V3.6 scheduled for release in Q4-15.

PROVISIONING

Provisioning of MIN or MAX IOPS is done a per-VM basis. The MIN or MAX set is the aggregate IOPS for all VHDXs owned by that VM.

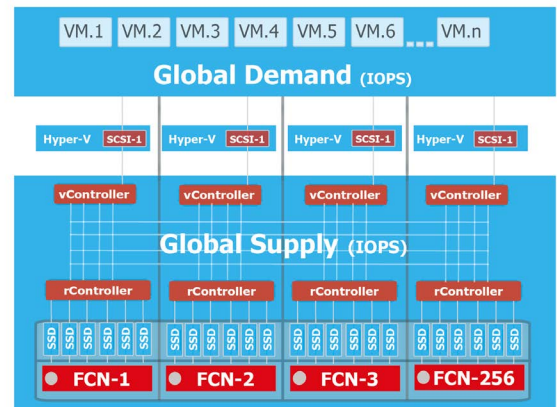
BENEFITS

Per-VM I/O control provides granular control to allocate IOPS to specific VMs. This model works well for high-performance workloads where a minimum number of IOPS need to be reserved to guarantee IOPS for a specific workload.

Gridstore Architecture

What makes the Gridstore storage QoS possible is a distributed resource management architecture that operates at both the host (demand) and storage (supply) sides of the storage network. Unlike a hypervisor that only controls demand from a single host, Gridstore's architecture provides a global view of demand across all hosts and a global view of IOPS supply across all storage resources.

Distributed Resource Management Architecture



Following are the major elements of the Gridstore QoS architecture:

- **I/O Tagging:** Each I/O is tagged with the VM's ID. Once tagged, I/O is put in the correct channel to access its storage resource.
- **I/O Channels:** Each VM's I/O is isolated into a unique channel across the network to its storage resources. A VM's I/O channel is a set of inbound/outbound socket pairs governed by vControllers™ on hosts and terminated by rControllers™ on the storage. I/O channels are grouped into priority classifications.
- **Demand Control (vControllers):** On the host side, vControllers schedule I/O requests for VMs. VM channels in higher classes are scheduled first followed by each subsequent lower class until all requests are satisfied. The scheduling algorithms allocate higher percentages of available resources at a given time to higher classes and less percentage to lower classes.
- **Supply Control (rControllers):** rControllers receive ordered I/O requests on VM channels within a Priority Group of channels. Outstanding I/O's are dealt with in order of Priority Class with special handling of VM channels with an IOPS Reserve. Subsequent lower Priority Classes are dealt with in order to the lowest class with special handling of VM IOPS Limit (MIN).



FOLLOW US  /gridstore  /company/gridstore  /user/GridstoreInc

t US 855.786.7065 US 650.316.5515 UK +44(0)20 3553 3662 e info@gridstore.com www.gridstore.com