

Quality of Service (QOS) With Tintri VM-Aware Storage

Dominic Cheah

Technical Marketing Engineer (TME)

Oct 2015

Contents

Intended Audience..... 3

Consolidated List of Practices..... 3

Challenges With Traditional Storage Arrays 3

QoS Redefined..... 4

Per-VM Granularity With Tintri VMstore 5

Managing QoS through Tintri VMstore UI 6

Managing QoS At Scale With Tintri Global Center..... 8

Conclusion..... 16

References 17

Appendix A..... 17

Intended Audience

This document will discuss how to use Tintri's per-VM Quality of Service (QoS) to deliver performance Service Level Agreements (SLAs) and contain 'noisy neighbors' in a virtualized data center. Data center administrators and service providers will discover how to use Tintri's per-VM QoS to configure and manage different tiers of services base on their customers service tiers, requested performance requirements, and budget.

Data center administrators and service providers will see examples of how to apply QoS policies to manage mix workloads for different customers, with different service level requirements; while hosting multiple types of hypervisors on the same Tintri VMstore in a virtualized data center. Virtualization experts and service providers will also learn how Tintri's Global Center is a single pane of glass that can be used to manage multiple Tintri VMstores.

Consolidated List of Practices

The table below includes the recommended practices in this document. Click the text on any of the recommendations to jump to the section that corresponds to each recommendation for additional information.

[DO: Validate that the TGC query filter has correctly selected the available virtual machine resources for the Service Group before continuing with configuring QoS or configuring Service Group Protection. Fix the Service Group query filter as needed to ensure all current virtual machine resources are correctly selected.](#)

[DO: Test and verify all queries for new Service Groups created on TGC to ensure that new virtual machine resources will be dynamically added to the Service Groups. Service Group filters are case-sensitive, create advanced filters as needed to ensure virtual machines are dynamically added when more Tintri VMstores are added into the TGC environment.](#)

[DO: It is recommended to review the latest TGC release notes for the most up-to-date TGC and Service Group enhancements.](#)

Challenges With Traditional Storage Arrays

Traditional LUN base (Logical Unit Number) storage arrays are not designed with virtual machines and virtualized applications in mind. Each application or database on traditional storage arrays deployed on LUNs has to have RAID configurations customized based on the best practice guide of each storage array vendor for performance [and economic efficiencies](#). For example, on some LUN based storage array, an Oracle VM would use RAID 5 for datafiles and temp files, while requiring RAID 1 for redo log files. Another example is deploying SQL server VMs using RAID 10 with other traditional LUN storage.

As virtualized data center requirements for storage grows, finding storage and deploying the right type of LUN for each different application or database that is virtualized has become a bottleneck. The back and forth discussions for the raid type, the acceptable performance impact, and the low latency requirements take center stage. The result is deploying storage for virtual machines, just to host applications and databases, ends up taking valuable resources that could be better use to fulfill other customer service requests.

Tintri VM-Aware storage solution is designed for virtualization. From a virtualized data center administrator's point of view, this means that every [virtual disk in every](#) virtual machine resource that is deployed on a Tintri VMstore system has its own I/O lane (see figure 1). Applications and databases that

are hosted in virtual machines on Tintri VMstores [benefit from](#) optimized performance [and sub-millisecond latencies](#) [that dynamically adjust as the](#) virtual workload [changes](#).

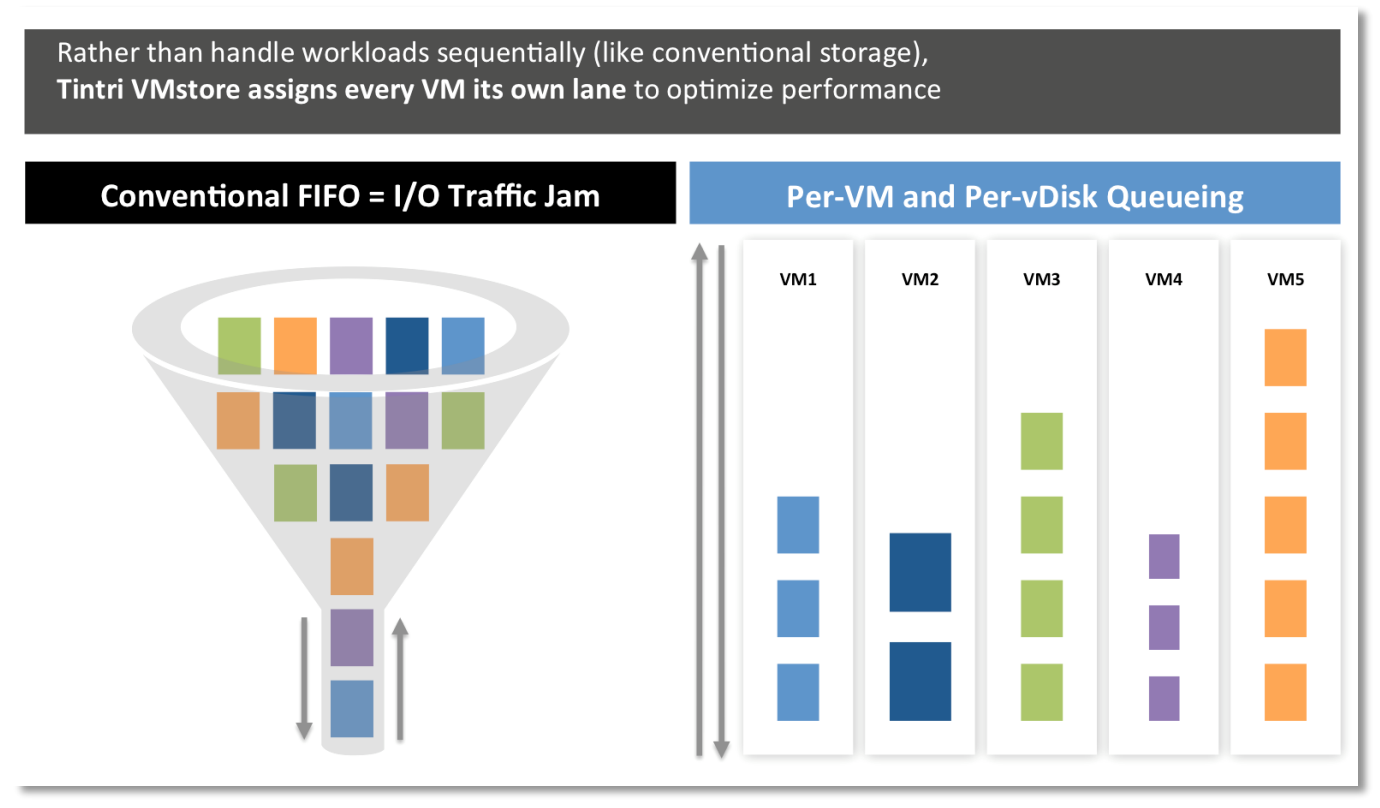


Figure 1 - Traditional storage vs. Tintri per-VM and per-virtual disk QoS

A storage system that is designed for virtualization, combined with the power of dynamic QoS policies, enables data center administrators to create and manage tiered services for different applications and databases, multi-tenancy, and multiple hypervisors on a scalable storage solution. For a list of supported hypervisor solutions, review [Tintri's Data Sheets](#).

QoS Redefined

Virtual machine QoS on traditional LUN results in every VM resource deployed on the same LUN receiving the same quality of service. When data center administrators are deploying a new virtualized data center and they have all the requirements in hand for each and every application and database, the data center administrators can easily carve out new LUNs on a new storage array and put the virtual machines into buckets base on the different service level agreements. However, managing LUNs and QoS becomes complex.

The process of carving out LUNs and customizing it base on different SLA agreements and placing applications and databases into buckets is not easy to manage. A data center administrator will eventually have to perform juggling acts and move VMs and virtual machines disks across LUNs just to guarantee the same level of QoS that was expected when the services were first deployed. As more and more virtual machine resources are dropped into the LUN buckets, the actual performance mileage for each application and database will eventually vary because QoS is at the LUN level and not at the per virtual machine level.

Tintri's patent pending technologies allows users to visualize contention, manage virtual machines and virtual disks [without LUNs](#). Tintri's dynamic QoS ensures quality of service at the virtual disk level, regardless of the supported hypervisors that are running on the Tintri VMstore system. Virtual resource contention and performance metrics of each virtual machine or virtual disks can be managed on each Tintri VMstore or at scale using [Tintri Global Center](#) (TGC).

NOTE: Tintri's QoS dynamically manages all virtual machines on a Tintri VMstore without any user interference. QoS enhancements made available in Tintri OS 3.x and later ensures that data center administrators and service providers can effectively manage their virtual machine resources for service tiers, chargeback, and also reigning in rogue virtual resources.

Per-VM Granularity With Tintri VMstore

In the Tintri VMstore UI, *Performance reserves* reflect the available headroom of all performance resources. For example, figure 2 shows a Tintri T820 VMstore managed by Derrick@ABC company in their Malaysian Data Center. On Oct 13th, Tintri VMstore-A shows **110%** performance reserve utilization (figure 2, point 1).

A visual inspection shows the performance reserve changers within the last week and *SQL VM 2* reflects a 44.4% (figure 2, point2) increase within the week. Because Tintri VMstores are able to visualize contention and resource utilization for each virtual machine workload, the data center administrators are able to drill into each virtual machine resource and review the performance changers.

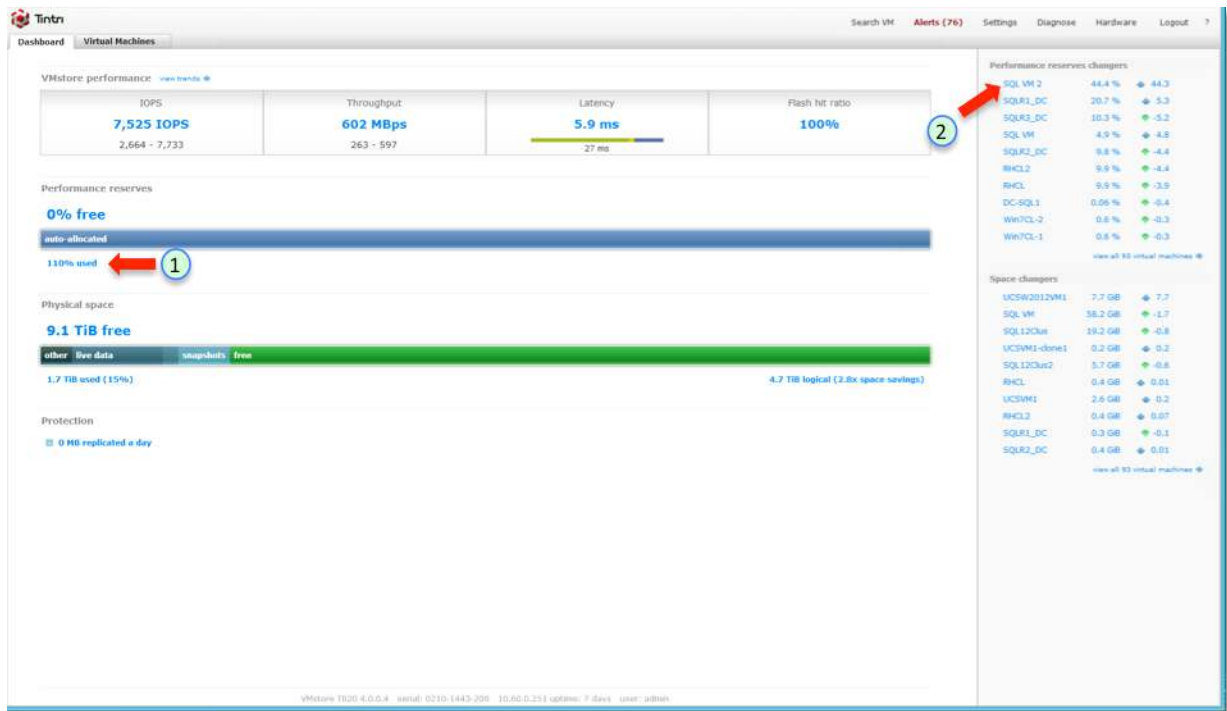


Figure 2 – Tintri T820 VMstore Dashboard view

Figure 3 shows *SQL VM 2* deployed in a Hyper-V environment (figure 3, point 1 and point 2). *SQL VM 2* is utilizing 40,525 [normalized IOPs](#) (figure 3, point 3) at 332MBps (figure 3, point 4) of reads and writes operations. Immediately, the administrators are able to visually confirm that this particular *SQL VM 2* has a read operation that is greater than 200 MBps at 5:20:36PM (figure 3, point4). The data center administrators need to guarantee that this specific virtual machine will receive the minimum QoS per the

SLA and throttle the max IOPs so that other virtual machines are guaranteed resources when performance reserves is at 0%.

From a resource utilization perspective, the activity on the *SQL VM 2* virtual machine can also look suspicious to the data center administrator. The administrators decides to review the historical trend of this particular virtual machine from the Tintri VMstore UI and they are able to confirm that this is not the typical performance workload that should be taking place on this particular virtual resource during this time period.

NOTE: To view normalized IOPS on a per virtual resource basis, the QoS checkbox must be selected for the virtual machine that is in focus in the Tintri VMstore UI or the Tintri Global Center UI. IOPS are normalized to 8KB per IOP. Normalizing IOPs to 8KB enables a simple standardized method on Tintri VMstores to measure virtual machine resource utilization. If IOPs were not normalized, virtual machine resources could have different block sizes for read and write operations on each virtual workload. It would make measuring performance and IOPs on these virtual machines near impossible when there are thousands of virtual machines deployed on a storage system. The use of normalized IOPs and the per-VM granular metrics allows for a more accurate and predictable method to enable chargeback base on each virtualized resource in a virtualized data center.

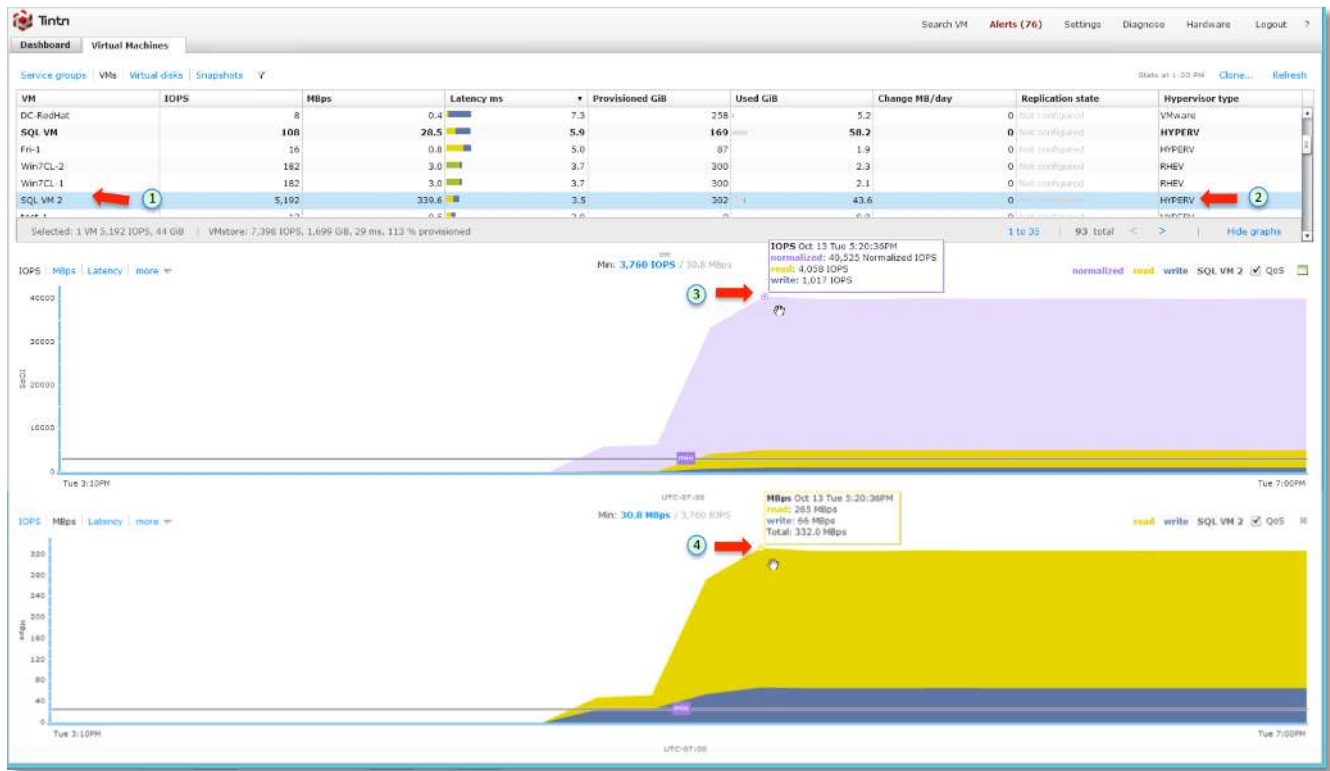


Figure 3 – Virtual Machines view of SQL VM 2 with IOPs and MBps.

Managing QoS through Tintri VMstore UI

The data center administrators need to guarantee the min IOPs on this *SQL VM 2* virtual machine and determine if there should be any additional security concern with this particular virtual resource. Using Tintri’s per virtual machine QoS, the data center administrators can manually tweak the min and max

IOPS of this single virtual resource without any negative QoS impacts to the other virtual machines on the Tintri VMstore as they investigate the issue.

Figure 4 shows the data center administrator setting the max IOPS on the virtual machine (figure 4, point 1). *SQL VM 2* is manually throttled to limit its resource utilization on the Tintri VMstore. The max IOPS throttling is immediate and is reflected in the live statistics of *SQL VM 2* (figure 4, point 2). A data center administrator can perform dynamic changes to the min and max QoS setting by dragging up or down the max and min performance line on a per-VM basis.

NOTE: For data center administrators who are open to scripting, Tintri Automation Toolkit can be downloaded from [Tintri Support](#). Scripting provides the added flexibility to manage multiple virtual machine resources across multiple Tintri VMstores without TGC. However, as a data center scales with additional Tintri VMstores, it is recommended to utilize TGC as a point of management for multiple VMstores across multiple data centers.

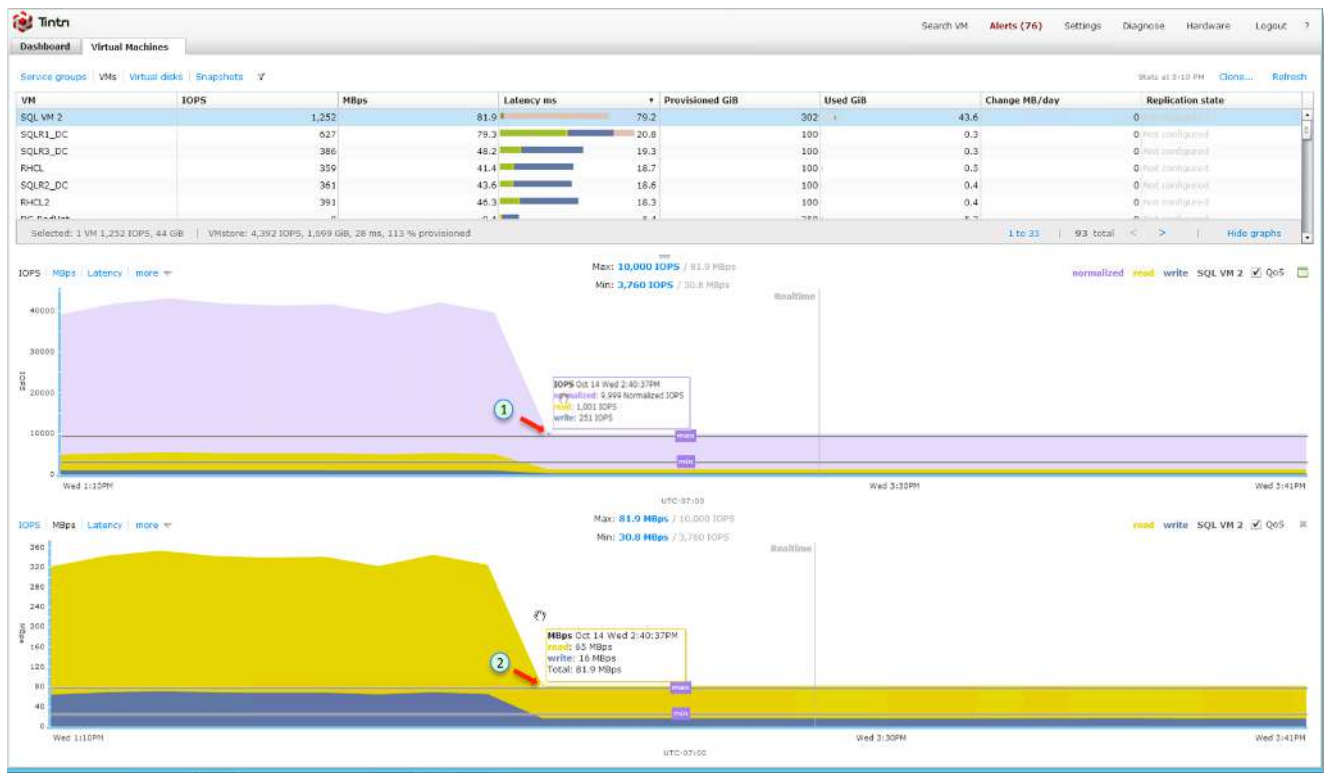


Figure 4 – Min and max IOPS thresholds set on SQL VM2.

SQL VM 2 is manually throttled to 10K max IOPs (figure 4, point 1). The read and write operations of this particular virtual machine has been limited to 81.9 MBps (figure 4, point 2) and is inline with other virtual resources that are hosted on the Tintri VMstore. A call to the SQL administrator eventually confirms their suspicion, a junior database administrator had accidentally triggered an SQL procedure that was producing lots of read operations every time it was executed and it was always running.

Performance management of runaway virtual machine resource issues used to take days to debug and resolve. Data center administrators @ABC company were able to troubleshoot within minutes on Tintri VMstores. Within an hour, the source of the rogue VM was determined with the help of the SQL database administrators.

To prevent future occurrences on this particular virtual server and to ensure that the virtual machine resource utilization will always be in-line with the rest of the other virtual resources, the data center administrators decided to enforce min and max IOPs on this virtual machine until the database team comes up with a plan to prevent rogue SQL procedures.

Figure 5 shows the data center administrator manually setting min normalized IOPs and max normalized IOPs (figure 5, point 1 and point 2) on this virtual machine just by clicking and dragging on the min and max IOPs toggles in the *Virtual Machines* view. Within seconds, the data center administrator was able to ensure the performance and IOPs requirements of a single virtual machine without negatively impacting any other virtual machine or other hypervisors that are dependent on the dynamic QoS that Tintri VMstores already delivers. Figure 5 (point 3 and point 4), shows the **min 3760** normalized IOPS and **max 15820** normalized IOPS that was dynamically tuned for the specific *SQL VM 2* virtual machine. All other virtual workloads are not affected by this specific min and max QoS setting because Tintri VMstores are VM-Aware and hypervisor agnostic.

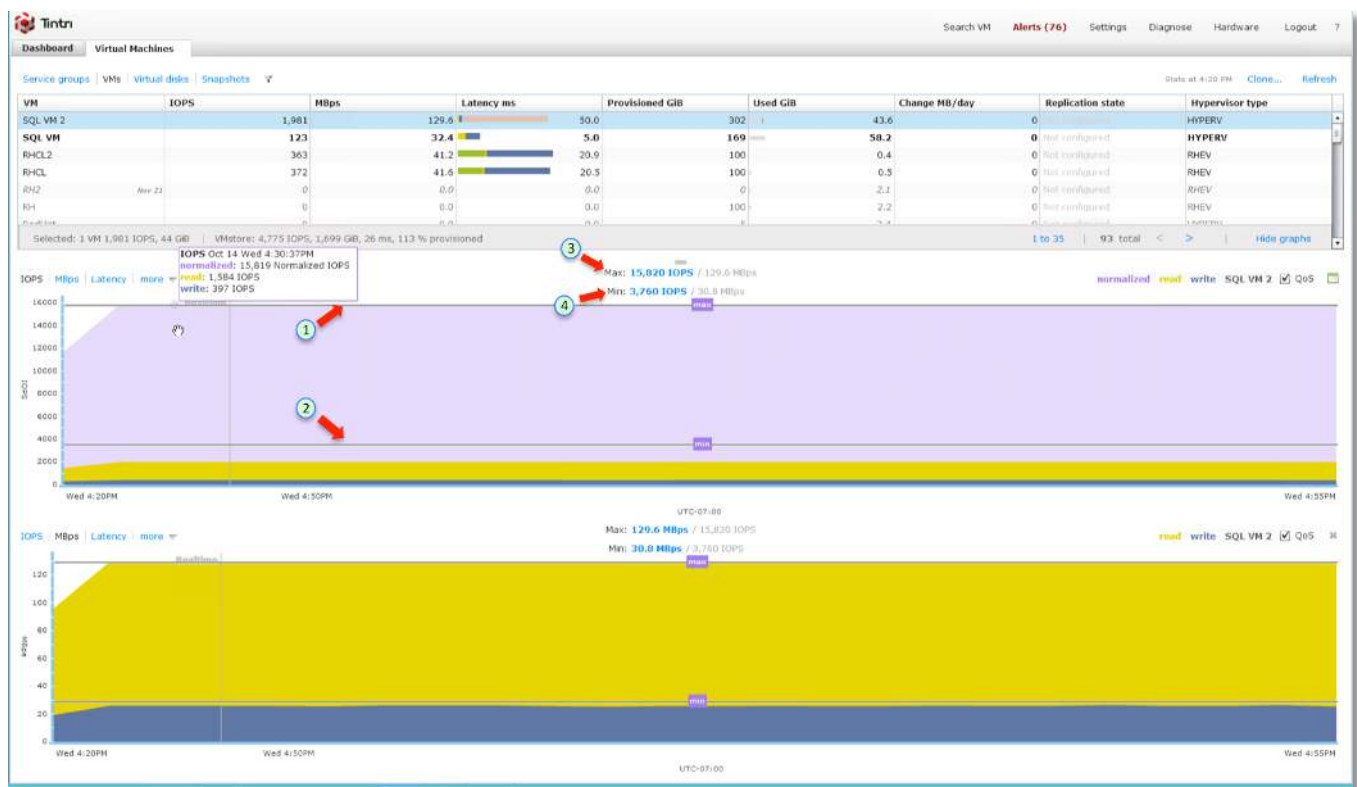


Figure 5 – SQL VM 2 with updated min and max IOPs.

A Tintri VMstore UI allows data center administrators to easily manage thousands of virtual machines, visualize contention, configure data protection, review performance metrics, and immediately remedy rogue virtual resource with QoS on each Tintri VMstore. These are some of the unique features and functionality that comes with a storage solution designed for virtualization.

Managing QoS At Scale With Tintri Global Center

With Tintri VMstores, data center administrators are able to perform granular storage operations such as:

1. Per-VM QoS
2. Per-VM snapshot
3. Per-VM clone

4. Per-VM replication
5. Review storage metrics at the VM level

With the power of a Tintri VMstore UI, REST API, and powershell, a data center administrator can fully leverage these features to enable application development teams to accelerate their development and test cycles, host applications and databases in virtualized environments, and enable cloud service providers to create different storage tiers base on their customers' performance requirements and budget.

Using Tintri's unique per-VM level QoS, managing thousands of virtual resources without the hassle of juggling VMs and virtual machine disks across storage is easy. Per-VM level QoS sets Tintri apart from the other traditional LUN storage arrays. As a virtualized data center scales with multiple Tintri VMstores, TGC can be deployed to protect virtual machines using Tintri's SnapVM, ReplicateVM, and manage QoS of hundreds of thousands of virtual machines across multiple Tintri VMstores in a single pane of glass.

Figure 6 shows *SQL VM 2* from the Tintri Global Center. Tintri's unique per-VM level granularity is reflected in the right hand column (figure 6, point 1) of the UI as each virtual machine across multiple Tintri VMstores can be selected for performance metric review from the left hand column (figure 6, point 2).

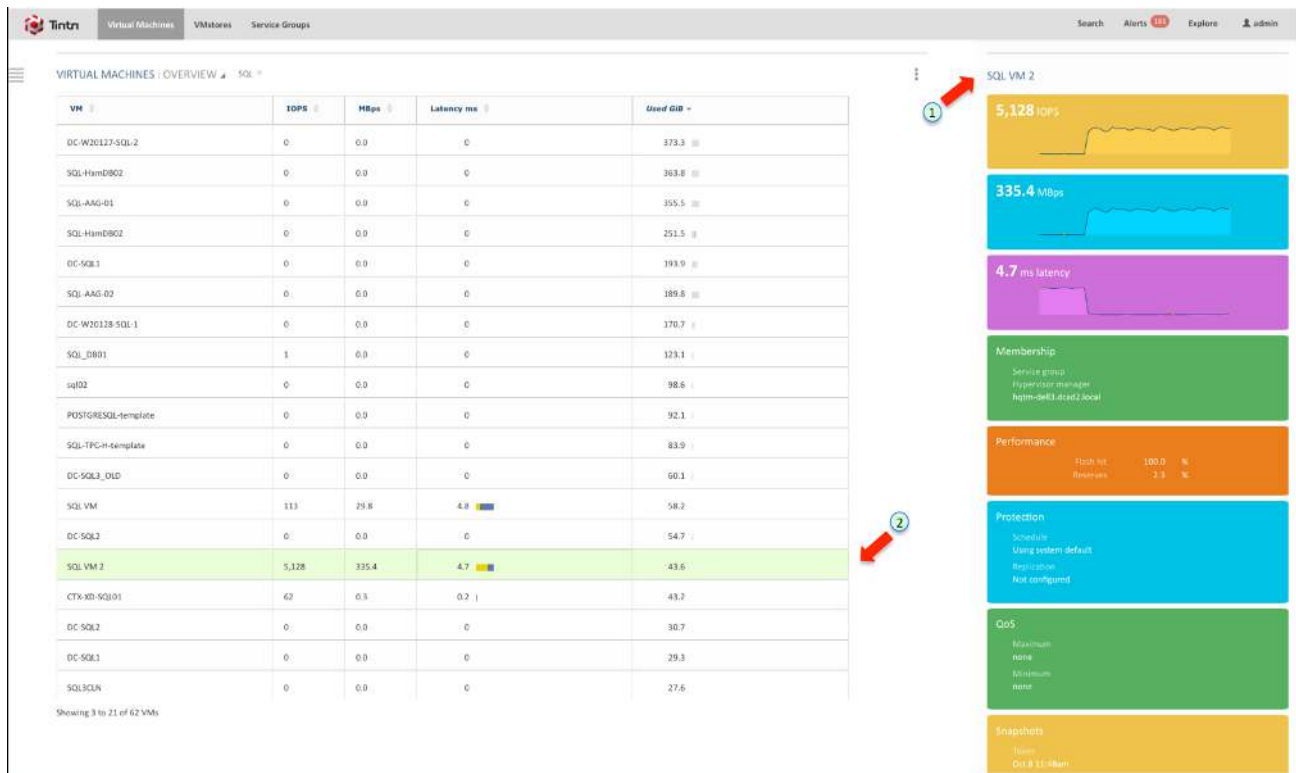


Figure 6 – Per-VM level granularity from the Tintri Global Center

A data center administrator or a cloud service provider can also easily use TGC to create tiered storage services across multiple Tintri VMstores and provide the per virtual workload visibility into performance and transparency to justify chargeback to their customers. Figure 7 shows how to create a SQL service group with the following information:

- Service Group Name: All Tier 1 SQL Servers
- A detailed description of the minimum IOPs requirements for the Tier 1 service group

- Membership rule with all VMs whose name starts with 'SQL'

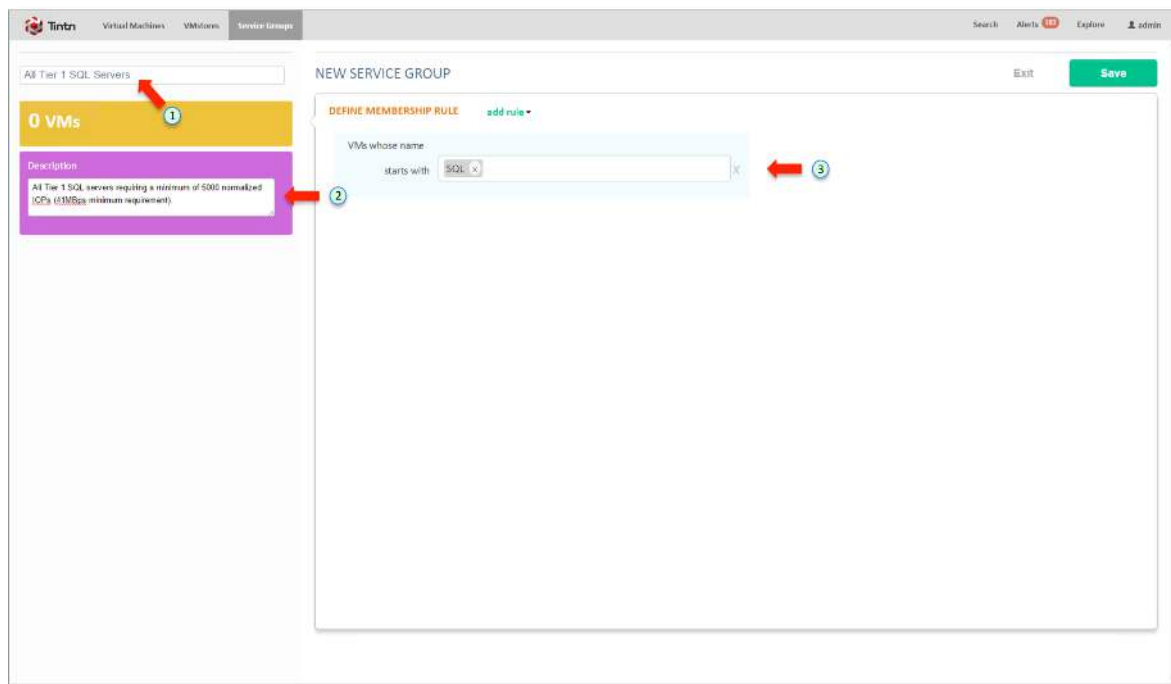


Figure 7 – Creating Service Groups on the Tintri Global Center

Click on **Save** and the new service group will be created. The TGC service group filter will select all virtual machine resources whose name starts with 'SQL' and automatically add those resources to this particular group. Figure 8 shows the updated 'All Tier 1 SQL Servers' service group with 20 virtual machines that meets the filter requirements across Tintri VMstores managed by TGC.

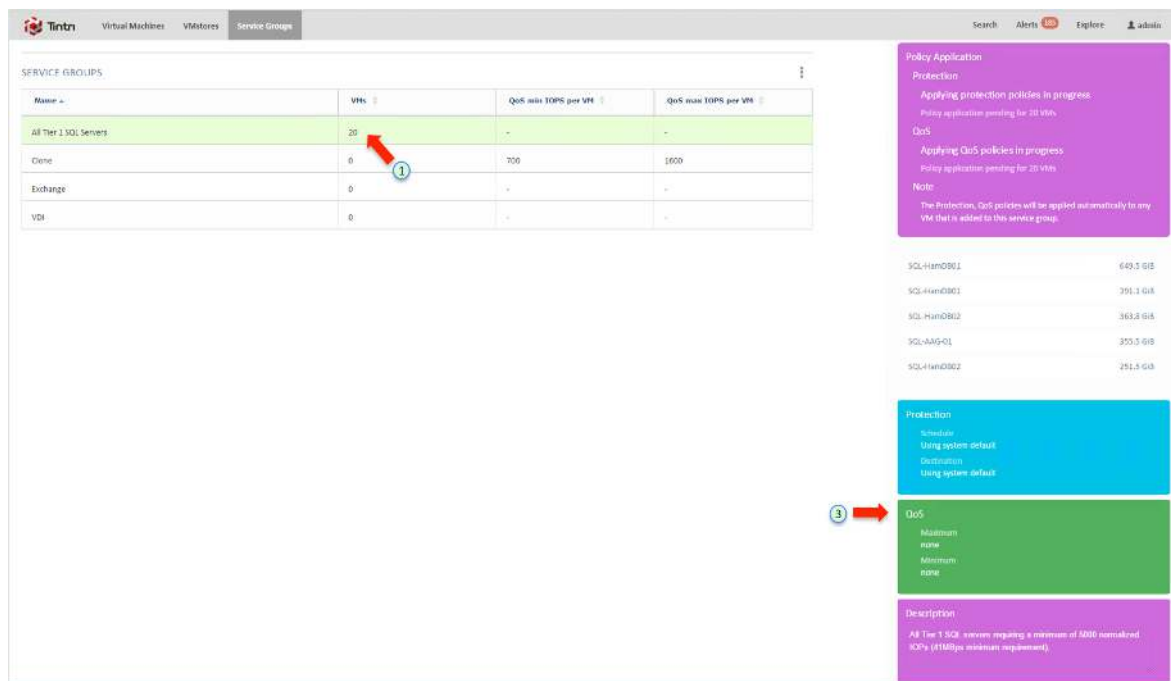


Figure 8 – Per-VM level granularity from the Tintri Global Center

Click on 'Show all **x** VMs' within the service group (see figure 9) to see all the virtual machine resources that meets the filter criteria.

NOTE: **x** denotes the number of virtual machine resources that have been selected and automatically added to the Service Group base on the query filter.

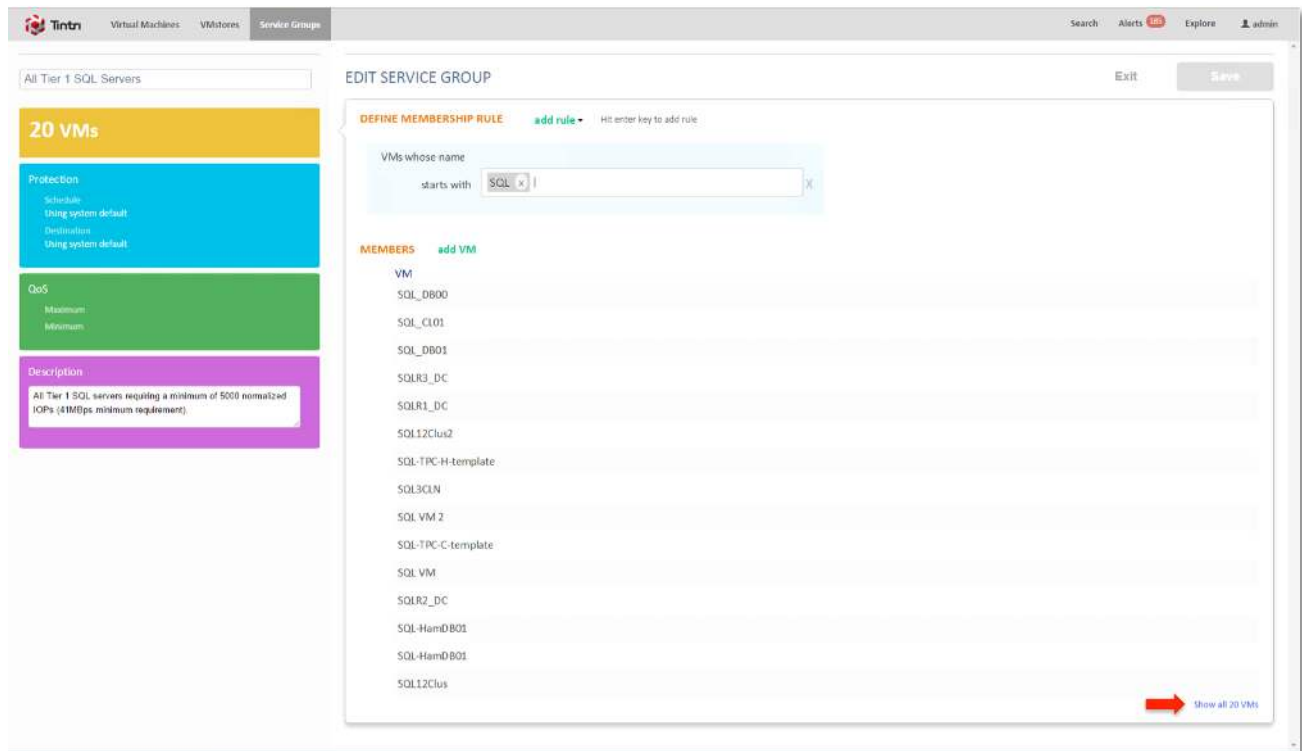


Figure 9 – Show all 20 VMs of 'All Tier 1 SQL Servers' Service Group.

The TGC administrator should validate that the query has correctly selected all VMs that meets the filter criteria and added the resources to the Service Group before configuring QoS or Data Protection at the Service Group level (see figure 10). If the Service Group does not correctly filter all the virtual resource requirements, update the query filter to ensure that it will automatically select virtual machines to be added into the Service Group base on the specified filter criteria(s).

Before configuring Service Group QoS or Service Group Protection, validate that the Service Group filter has correctly selected all virtual machine resources into the Service Group. If the query filter is not correct, edit the Service Group and update the query filter to ensure that the virtual machines that needs to be in the Service Group will be dynamically added.

New virtual machine resources that meet the query filter criteria will be dynamically added into the Service Group for QoS or data protection. This feature ensures that new virtual machines that meet the query criteria in TGC will be added to the existing Service Group without any user intervention.

Tintri's Automation Toolkit can also be used with TGC to manage Service Groups and QoS. Review Appendix A for an example of how to use Tintri Automation to manage Service Groups and QoS with TGC. The sample script is provided **as-is**. It is not recommended to use the script without testing and modifying it to fit your production environment.

NOTE: The query filter is case-sensitive; edit your existing query to ensure that VMs that must be added to the Service Group are automatically detected. It is Tintri's recommendation to test new Service group filters to ensure that virtual machine resources are properly protected.

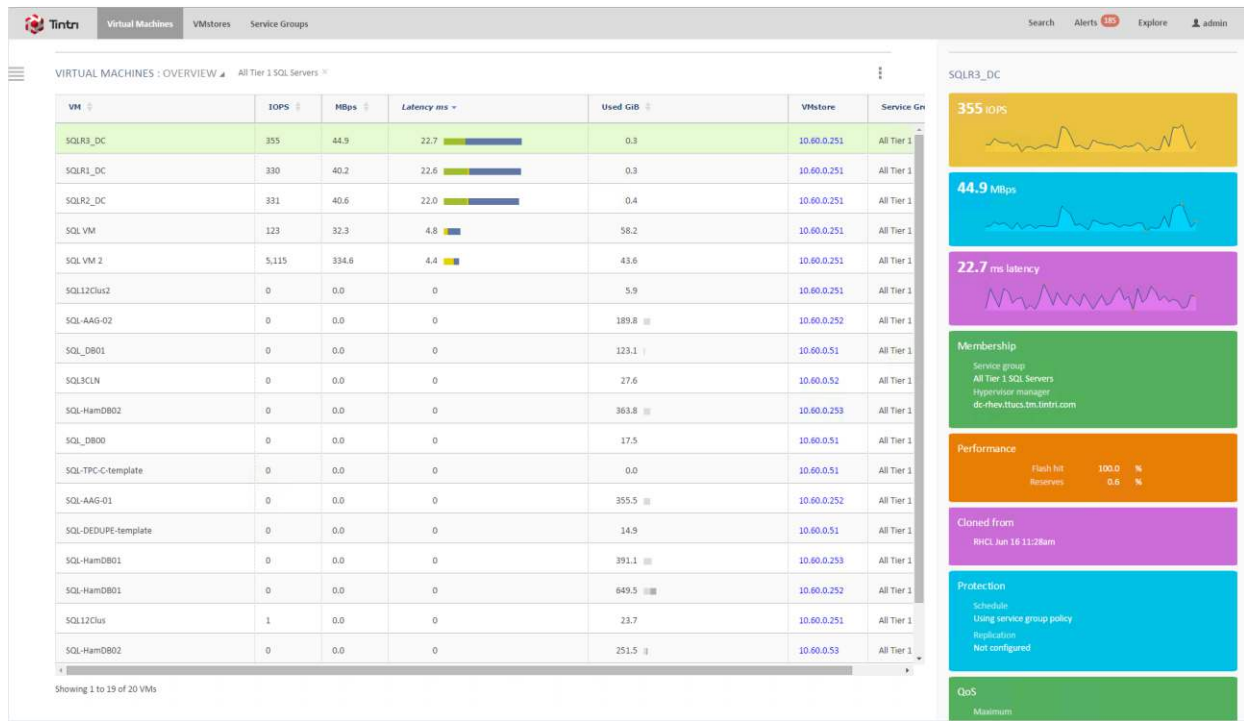


Figure 10 – Validate TGC query by reviewing all virtual machine resources in the Service Group.

DO: Validate that the TGC query filter has correctly selected the available virtual machine resources for the Service Group before continuing with configuring QoS or configuring Service Group Protection. Fix the Service Group query filter as needed to ensure all current virtual machine resources are correctly selected.

DO: Test and verify all queries for new Service Groups created on TGC to ensure that new virtual machine resources will be dynamically added to the Service Groups. Service Group filters are case-sensitive, create advanced filters as needed to ensure virtual machines are dynamically added when more Tintri VMstores are added into the TGC environment.

Figure 11 shows how to select and configure QoS by clicking on the menu tile and selecting **Configure QoS**.



Figure 11 – Select Configure QoS for ‘All Tier 1 SQL Servers’

Figure 12 shows min normalized IOPs is set to 5000 for the ‘All Tier 1 SQL Servers’ Service Group. Click on Save to configure the QoS for the Service Group.

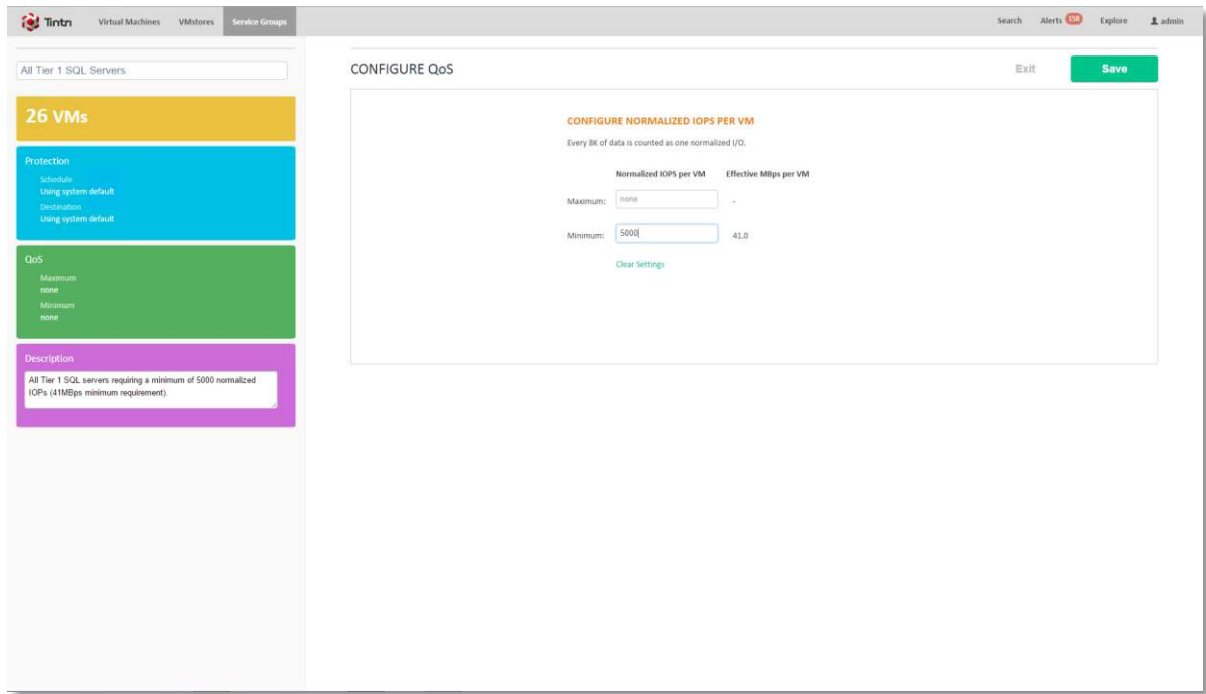


Figure 12 – Configure min QoS for ‘All Tier 1 SQL Servers’ Service Group.

A review of the Tintri VMstores that are connected to TGC will show the virtual machines that are in the ‘All Tier 1 SQL Servers’ Service Group adopting the min QoS (see figure 13, point 1 and point 2). New virtual machines across Tintri VMstores that are managed by TGC will dynamically be added into the Service Groups as along as the virtual machine resources meets the query filter criteria.

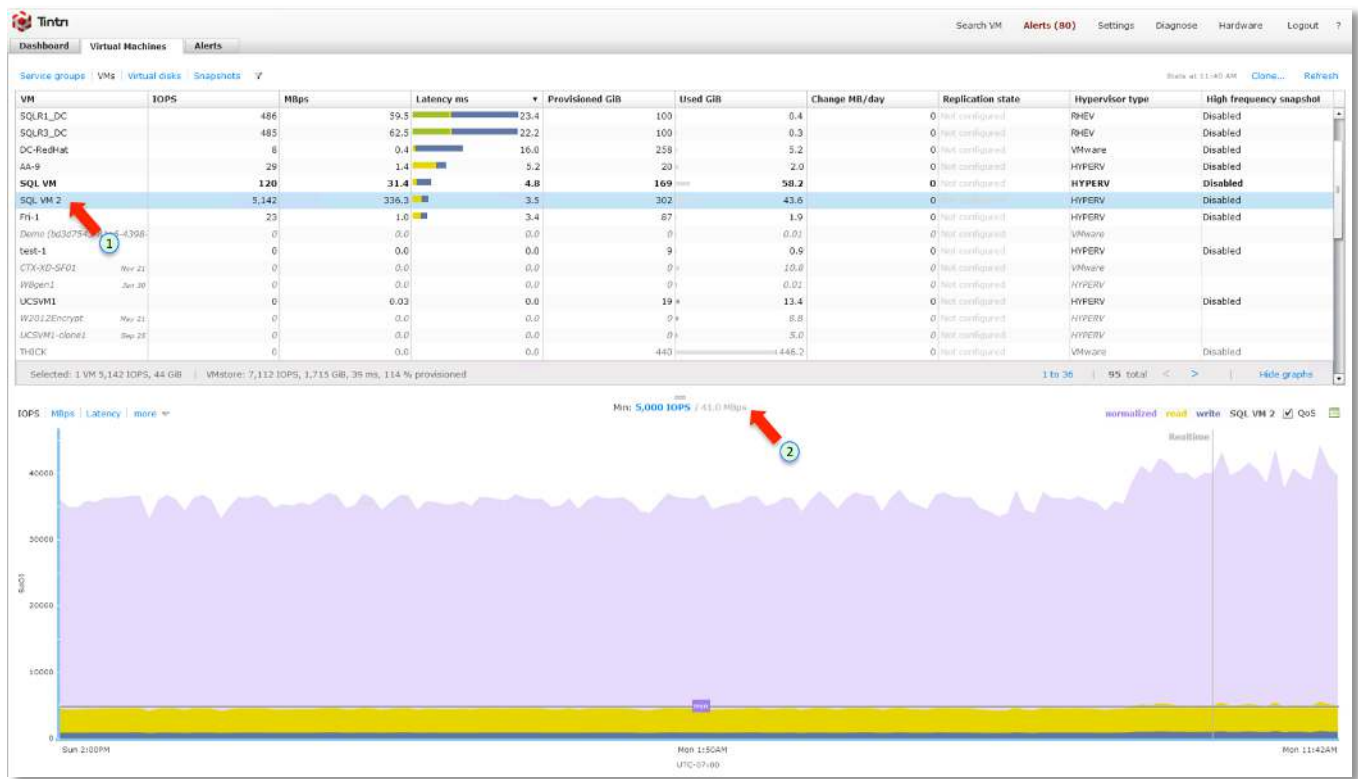


Figure 13 – SQL VM 2 and other SQL VM within the 'All Tier 1 SQL Servers' have been set with 5000 min normalized IOPs.

NOTE: QoS configuration changes for a virtual machine at the Tintri VMstore level will override QoS settings from the TGC Service Group. When tuning QoS for virtual machines in a group setting, use the TGC Service Group. If there are only a couple of virtual machine resources that require manual QoS setting on a Tintri VMstore, the administrator should manually modify QoS on a per virtual machine resource level.

Figure 14 shows how to override the manual QoS setting from TGC for each virtual machine that is added into the Service Group. The data center administrator must perform this on each virtual machine that has been dynamically added to the Service Group to ensure that the virtual resources will utilize the Service Group QoS setting. It is recommended to review the latest TGC release notes for Service Group enhancements.



Figure 14 – Override SQL VM 2 QoS setting to use TGC QoS Service Group.

DO: It is recommended to review the latest TGC release notes for the most up-to-date TGC and Service Group enhancements.

Using TGC Service Groups, cloud service providers using Tintri VMstores are able to create multiple types of service levels for their customers with QoS across multiple Tintri VMstores. For example:

- Platinum service policies for virtual machine resources that requires a minimum of 10K IOPs. This will guarantee that virtual machines across Tintri VMstores that belong to the service group will be guaranteed a minimum of 10K IOPs regardless of the performance reserves on the Tintri VMstores. In a scenario that requires more resources for these platinum customers, a cloud service provider would be able to create additional Service Groups to limit the resource consumption of other Silver or Bronze service level customers.
- Gold service policies for virtual machines resources that requires a minimum of 10K IOPs and a maximum of 15K IOPs will ensure that the virtual machines belonging to the service group will be guaranteed a minimum of 10K IOPs but will be capped at 15K IOPs.

The results of QoS on each virtual machine resource in a Service Group is instantaneous. Figure 15 is an example of how cloud service providers can utilize TGC QoS at the Service Group level to manage and control virtual machine resource usage to downgrade or upgrade customers virtual machine performance base on their SLA agreements.

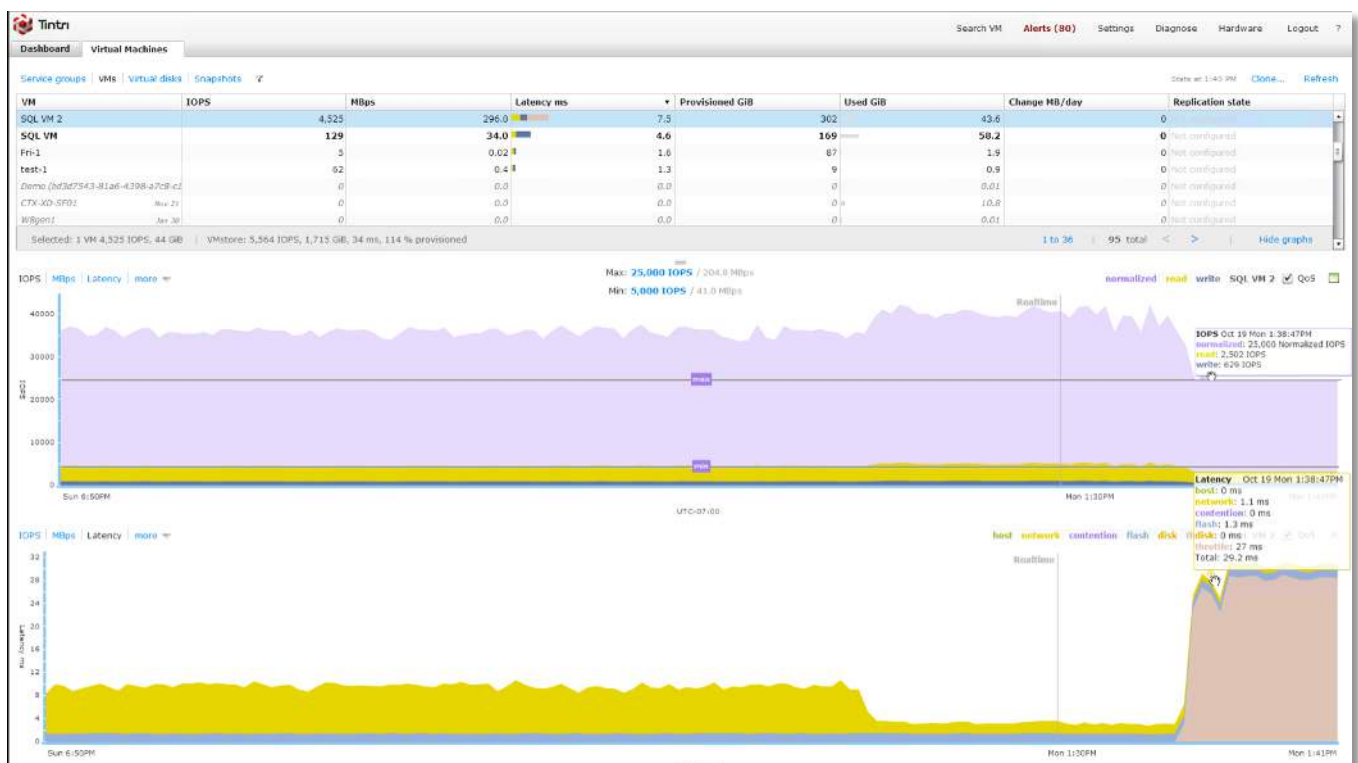


Figure 15 – Instantaneous QoS override base on TGC Service Group setting.

SQL VM 2, in this example, has a manual throttle of 27 milliseconds to ensure that the actual performance of the virtual resource does not exceed 25K normalized IOPs. Tintri VMstore's QoS is a simple and manageable way to ensure that a cloud service provider's customers can be guaranteed the level of performance that they can expect base their SLAs.

Tintri QoS feature can also be used to control performance of runaway VMs using normalized IOPs. If an SLA agreement is violated or a performance upgrade is required for a set of virtual machines, a cloud service provider can utilize QoS at the Service Group level, within TGC, to limit rogue VM resources or upgrade a set of virtual resources regardless of the supported hypervisors that are running on the multiple Tintri VMstores.

Conclusion

Figure 16 is an example of how a Tintri VMstore's UI or TGC can help take the guesswork out of troubleshooting resource utilization. For example, a mouse over to the *Latency* column lays bare, across your virtualized infrastructure hosted on Tintri VMstores, the sources of latencies. Troubleshooting and debugging latency issues that used to take hours or days can be traced and granularly tracked right down to each virtual machine resource within minutes.

Within Tintri VMstore's UI or TGC, the sources of contention is easily identified by the following:

- Host
- Network
- Storage
- Throttle (if QoS is enabled)



Figure 16 – Performance reserves and latency visualization

With Tintri VMstores and TGC, data center administrators and cloud service providers will spend 0% of their resources to manage storage LUNs. This is because Tintri VMstores are designed for virtualized data centers for hosting test and dev environments, applications and databases, and cloud service providers virtual resources.

Administrators will not need to spend any time juggling virtual machines around between storage because virtual machine resources that are hosted on Tintri VMstores are dynamically managed at the virtual machine resource level. There is no more guesswork required to determine which virtual resource is having issues or which virtual machine is a rogue resource. Data center administrators and cloud service providers can easily utilize QoS and Service Groups via TGC to reign in rogue VMs or promote performance of a set of VMs to meet customer SLAs.

References

- [Without LUNs](#)
- [Tintri Global Center](#)

Appendix A

The following script is provided **as-is**. Tintri does not recommend using the sample script without testing and modifying the script to fit your requirements. Data center administrators are encouraged to review, modify, and test the script to ensure that it works and meets your Service Group QoS requirement before deploying it in a production environment. Tintri also recommends the data center administrator to review the latest Tintri Automation Toolkit release notes to ensure that commands are not deprecated.

<#

The MIT License (MIT)

Copyright (c) 2015 Tintri, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

#>

<#

A Service Group is a logical collection of VMs in a Tintri Global Center (TGC).

Administrators can use service groups to manage a group of VMs

(apply protection policies and other settings) as they would manage a single VM.

The following code snippet shows how to apply a QoS policy to all VMs in a TGC service group using the Tintri Automation Toolkit 2.0.0.1.

More specifically, we'll apply a config that sets the minimum normalized IOPS.

* This script assumes that all the VMs that are part of the service group belong to exactly one VMstore (not multiple VMstores).

#>

Param(

[string] \$tgcName,

[string] \$vmStoreName,

[string] \$serviceName,

[int] \$minNormalizedIops

```

)

# Connect to the TGC, will prompt for credentials
Write-Host "Connecting to the Tintri Global Center $tgcName"
$tgc = Connect-TintriServer -Server $tgcName

if ($tgc -eq $null)
{
    Write-Host "Could not connect to $tgcName"
    return
}

# Get the service group on the TGC
Write-Host "Getting the service group $serviceGroupName on $tgcName"
$serviceGroup = Get-TintriServiceGroup -Name $serviceGroupName

# Fetch all the VMs of the service group
Write-Host "Getting the VMs that are members of the service group $serviceGroupName"
$serviceGroupVmsOnTgc = $serviceGroup | Get-TintriVM

# Resolve the corresponding VM objects on the VMstore.
# Connect to the VMstore, will prompt for credentials
Write-Host "Connecting to the VMstore $vmStoreName"
$ts = Connect-TintriServer $vmStoreName

if ($ts -eq $null)
{
    Write-Host "Could not connect to $vmStoreName"
    return
}

Write-Host "Resolving the corresponding VMs on $vmStoreName"
$serviceGroupVmsOnVmstore = $serviceGroupVmsOnTgc | Get-TintriVM -Uuid { $_.Uuid.Uuid } -TintriServer $ts

# We can apply QoS policies only on live VMs. Filter them.
$liveVms = $serviceGroupVmsOnVmstore | Where { $_.IsLive }

# Update the QoS setting for these VMs.
Write-Host "Updating the QoS setting (min normalized IOPS) for these VMs"
$liveVms | Set-TintriVMQos -MinNormalizedIops $minNormalizedIops

# Disconnecting from the TGC and VMstore.
Disconnect-TintriServer -All

```