



More about Service Levels

Keystone

NetApp
January 25, 2023

Table of Contents

- More about Service Levels 1
- Quality of Service (QoS) in Keystone 1
- Capacity requirements for service levels 4

More about Service Levels

Quality of Service (QoS) in Keystone

Storage quality of service (QoS) is a critical technology that ensures that applications obtain consistent and predictable performance. Without QoS, certain workloads like booting of multiple systems can take up entire resources for a period and affect other workloads. For information about QoS, see [Guarantee throughput with QoS overview](#).

Applying QoS ensures that applications get consistent performance every time. NetApp ONTAP defines two parameters for QoS: ceilings- [QoS Max](#), and throughput floor- [QoS Min](#). QoS Max represents the maximum number of IOPS or MBps or IOPS and MBps a workload can achieve. On the other hand, QoS Min guarantees that the throughput for a workload is not below a certain number of IOPS or Mbps, or IOPS and MBps.

Adaptive QoS

Adaptive QoS (AQoS) is used for Keystone services to maintain the ratio of IOPS to TBs/GBs as the size varies the QoS policy is automatically applied. For information about adaptive QoS, see [About adaptive QoS](#).

Keystone provides you with adaptive QoS policies that you can set up once your cluster is in production. You should ensure that all your volumes are associated with the correct adaptive QoS policies that are already created and available in your system.



If you have not applied adaptive QoS policies on your volumes correctly, those volumes will be measured and charged at the highest service level as per your subscription. This may result in unintended burst charges.

Configuration of adaptive QoS policy group

You can configure adaptive QoS policies to automatically scale a throughput ceiling or floor to volume size. For configuring a policy, you should be aware of the following:

- **Policy group name:** The name of the adaptive QoS policy group. For example, `Keystone_extreme`.
- **VServer:** The name of the VServer or storage VM (storage virtual machine).
- **Expected IOPS:** The number of IOPS/TiB that ONTAP provides as a minimum as long as there is performance guarantee.
- **Peak IOPS:** The maximum number of IOPS/TiB that will be available to a volume before performance is limited.
- **Expected IOPS allocation:** This field controls whether the expected IOPS available to the volume is based on the allocated or used size of the volume. In Keystone, this is based on the allocated space.
- **Peak IOPS allocation:** This field controls whether the peak IOPS available to the volume is based on the allocated or used size of the volume. In Keystone, this is based on the used space.
- **Absolute minimum IOPS:** The lowest number of expected IOPS that will be applied to a volume if the volume size is very small and would otherwise result in an unacceptable number of IOPS. This value defaults to 1,000 for `Extreme`, 500 for `Premium`, and 75 for other service levels.



This is not IOPS density (for example, 75 IOPS/TiB), but an absolute minimum number of IOPS.

For information about IO density, see [Metrics and definitions used in Keystone Services](#). For more information about AQoS policy groups, see [Use adaptive QoS policy groups](#).

Settings of adaptive QoS policies

The settings for adaptive QoS policies for each service level are described in the following sections. The minimum and maximum volume sizes for each service level provided here allow for optimal IOPs and latency values for a volume. Creating too many volumes outside of these guidelines may negatively impact performance in those volumes.

Settings for Extreme service level

Settings and commands for the Extreme service level:

- Sample command:

```
qos adaptive-policy-group create -policy-group <Keystone_extreme> -vserver <SVM_name> -expected-iops 6144 -peak-iops 12288 -expected-iops-allocation allocated-space -peak-iops-allocation used-space -block-size Any -absolute -min-iops 1000
```

- Minimum volume size: 100GiB, 0.1TiB
- Maximum volume size: 10TiB

Settings for Premium service level

Settings and commands for the Premium service level:

- Sample command:

```
qos adaptive-policy-group create -policy-group <Keystone_premium> -vserver <SVM_name> -expected-iops 2048 -peak-iops 4096 -expected-iops-allocation allocated-space -peak-iops-allocation used-space -block-size Any -absolute -min-iops 500
```

- Minimum volume size: 500GiB, 0.5TiB
- Maximum volume size: 50TiB

Settings for Performance service level

Settings and commands for the Performance service level:

- Sample command:

```
qos adaptive-policy-group create -policy-group <Keystone_performance>
-vserver <SVM_name> -expected-iops 1024 -peak-iops 2048 -expected-iops
-allocation allocated-space -peak-iops-allocation used-space -block-size
Any -absolute-min-iops 250
```

- Minimum volume size: 500GiB, 0.5TiB
- Maximum volume size: 80TiB

Settings for Standard service level

Settings and commands for the Standard service level:

- Sample command:

```
qos adaptive-policy-group create -policy-group <Keystone_standard>
-vserver <SVM_name> -expected-iops 256 -peak-iops 512 -expected-iops
-allocation allocated-space -peak-iops-allocation used-space -block-size
Any -absolute-min-iops 77
```

- Minimum volume size: 1TiB
- Maximum volume size: 100TiB

Settings for Value service level

Settings and commands for the Value service level:

- Sample command:

```
qos adaptive-policy-group create -policy-group <Keystone_value> -vserver
<SVM_name> -expected-iops 64 -peak-iops 128 -expected-iops-allocation
allocated-space -peak-iops-allocation used-space -block-size Any -absolute
-min-iops 59
```

- Minimum volume size: 1TiB
- Maximum volume size: 100TiB

Block size calculation

Note these points before you calculate the block size by using these settings:

- $IOPS/TiB = MBps/TiB$ divided by $(block\ size * 1024)$
- Block size is in KB/IO
- TiB = 1024GiB; GiB = 1024MiB; MiB = 1024KiB; KiB = 1024Bytes; as per base 2
- TB = 1000GB; GB = 1000MB; MB = 1000KB; KB = 1000Bytes; as per base 10

Sample block size calculation

To calculate the throughput for the a service level, for example `Extreme` service level:

- Maximum IOPS: 12,288
- Block size per I/O: 32KB
- Maximum throughput = $(12288 * 32 * 1024) / (1024 * 1024) = 384\text{MBps/TiB}$

If a volume has 700GiB of logical used data, the available throughput will be:

$$\text{Maximum throughput} = 384 * 0.7 = 268.8\text{MBps}$$

Capacity requirements for service levels

The capacity requirements for Keystone STaaS service levels differ with the file, block, object, or cloud storage supported by the Keystone STaaS subscription.

Minimum capacity requirements for file and block services

The minimum capacity and incremental capacity allowed per subscription is described in the following table. The minimum capacity per service level is defined to be the same across Keystone sales motions. The capacity above the minimum capacity either at the beginning of the subscription, or as an add-on service to the subscription, or after reallocation during the subscription is also structured in the table.

Capacity	Extreme	Premium	Performance	Standard	Value
Minimum capacity [in TiB]	25			100	
Incremental capacity (and in multiples) allowed at start of subscription [in TiB]	25			25	
Incremental capacity (and in multiples) allowed as add-on during subscription [in TiB]	25			25	

Minimum capacity requirements for object storage

You can see the minimum capacity requirements for object storage in the following table:

Capacity	Data tiering	Object	Cloud Volumes ONTAP	Cloud Backup service
----------	--------------	--------	---------------------	----------------------

Minimum capacity [in TiB]	Not applicable	500	4	4
Incremental capacity (and in multiples) allowed at start of subscription [in TiB]	Not applicable	100	1	1
Incremental capacity (and in multiples) allowed as add-on during subscription [in TiB]	Not applicable	100	1	1

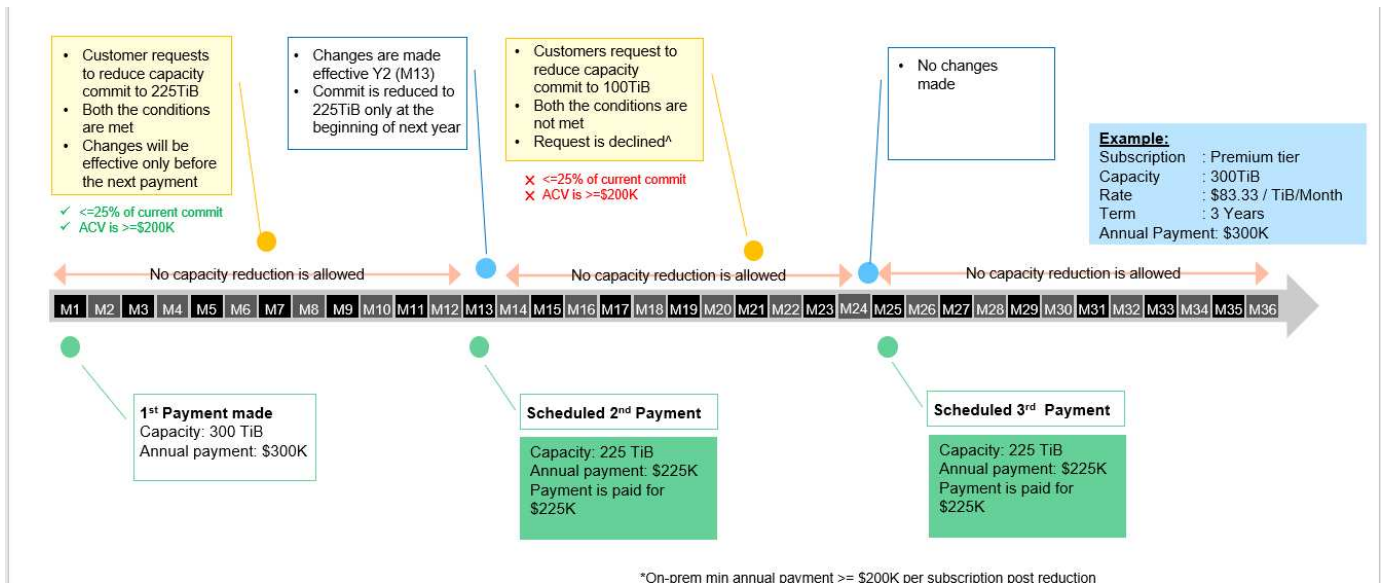
Capacity adjustments

Note the following clauses for capacity adjustments:

- Capacity can be added anytime during the term, except for the last 90 days of the contract term, in the increments per service level as described in the tables in the previous section. Addition of capacity and/or services is allowed within the last 90 days of the contract term as long as there is a consent of service renewal. Any addition in capacity, new service on-prem or cloud can co-term with the existing term. The invoice sent to you following the activation of the new services reflects the revised billing. Committed capacity of cloud services cannot be reduced at any point during the subscription term. Meanwhile, committed capacity and committed spend on the on-premises services during the term of the contract can be reduced based on certain criteria as defined in the following section *Capacity reduction*.
- A burst capacity is available at each site, based on the Keystone agreement. Usually it is 20% above the committed capacity for a service level. Any burst usage is billed only for that billing period. If you need an additional burst requirements that is greater than the capacity agreed upon, contact support.
- Committed capacity can be altered during a contract term, only under certain conditions, as described in the following section *Capacity reduction*.
- Increasing capacity or changing to higher service level during a subscription term is allowed. However, moving from a higher service level to a lower service level is not permitted.
- Any change request in the last 90 days of the service term requires a renewal of the service for a minimum of one year.

Capacity reduction

Capacity reduction (annual) is applicable to the *Annual in Advance* payment model and on-premises only deployments. It is not available for cloud services or hybrid cloud services. It provides provision for on-premises capacity, which can be reduced by up to 25% per service level per subscription. This reduction is allowed once every year to be made effective at the beginning of the next annual billing period. On-premises service-based annual payments should be $\geq \$200K$ anytime during the term in order to take advantage of capacity reduction. Because it is supported only for on-premises deployments, this billing model does not provide reallocation in spending from on-premises to cloud services. An example of annual capacity reduction is illustrated in the following image.

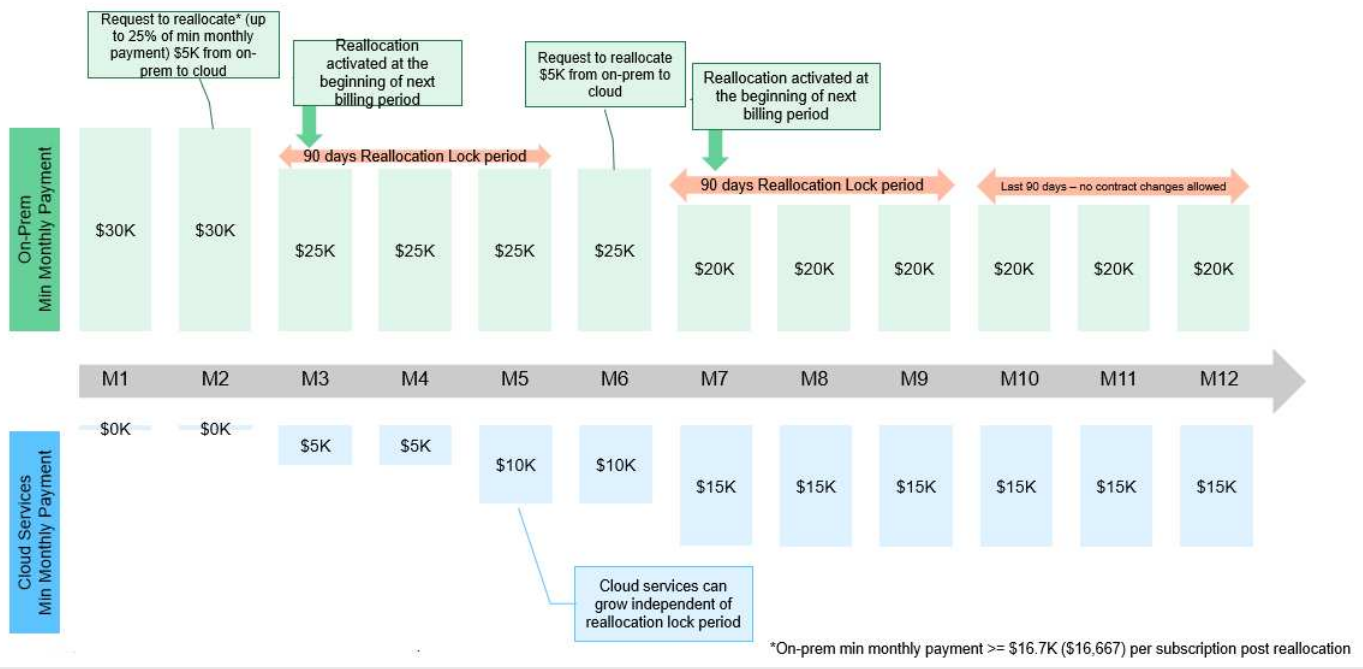


Quarterly spend reallocation

Requirements and conditions at a subscription level:

- Applies only to monthly billing in arrear model.
- Applies only to subscriptions with 1, 2, or 3-year term commitments.
- Cloud Volumes ONTAP and Cloud Backup service capacity should be purchased through Keystone.
- Up to 25% of the existing on-premises, service-based monthly payments can be used for reallocation to cloud services.
- On-premises, service-based monthly payments should be >=\$16.7K anytime during the term in order to take advantage of reallocation.
- Reallocation requests are made effective only after 90 days from the previous activation date of the reallocation.
- Reallocation cannot be done from cloud services back to on-premises.
- A request to reallocate should be formally submitted by the customer or partner to Keystone Success Manager (KSM) at least one week before the next billing cycle.
- New requests go into effect only from the next billing cycle.

An example of a quarterly spend reallocation is illustrated in the following image:



Copyright information

Copyright © 2023 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.