



# Keystone STaaS services

## Keystone

NetApp  
June 28, 2024

# Table of Contents

- Keystone STaaS services ..... 1
  - Metrics and definitions used in Keystone Services ..... 1
  - Quality of Service (QoS) in Keystone ..... 2
  - Supported storage in Keystone STaaS services ..... 5
  - Supported storage capacities in Keystone ..... 7
  - Service levels in Keystone ..... 9
  - Capacity requirements for service levels ..... 11
- Keystone subscription services | Version 1 ..... 14
- Add-on services ..... 15

# Keystone STaaS services

## Metrics and definitions used in Keystone Services

The following terms and definitions are used within the NetApp Keystone (Keystone) STaaS service to measure metrics:

- Capacity measurement units: GiB, TiB, and PiB
- IO density: IOPS/TiB: Number of input/output operations processed per second based on the total space that is being consumed by the workload, in tebibytes.
- Service availability
- Durability in accurate data access
- Latency and speed

### Metrics measurement

- **Capacity measurement in gibibyte (GiB), tebibyte (TiB), and pebibyte (PiB):** Measurements of data storage capacity using base of 1024 (1 GiB = 1024<sup>3</sup> bytes, 1 TiB = 1024<sup>4</sup> bytes, and 1 PiB = 1024<sup>5</sup> bytes).
- **Operations counter chart in IOPS/TiB:** The protocol operations per second, requested by the application, divided by the size of the volume used by workloads.
- **Availability:** Measured as a percentage of the number of I/O requests successfully responded to by the service, divided by total number of I/O requests made to the service. This is measured at the service demarcation in a month and does not include the scheduled service downtime or unavailability of the facilities, network, or other services provided by the customer.
- **Durability:** Percentage of data accessed without loss of fidelity, excluding customer-caused deletion or corruption.
- **Latency:** Time to service an I/O request received from a client, measured at the service demarcation (storage controller I/O port).

### Throughput performance metrics

Throughput performance metrics are applicable only for file and block services based on:

- 32 KB block sizes
- 70% read/30% write I/O mix

### Variations in IO density

IO density calculated in IOPS/TiB and/or MBps/TiB varies based on the following factors:

- Workload characteristics
- Latency, excluding the following:
  - Application latency
  - Host latency
  - Latency in the customer network while transferring data to and from the controller ports
  - Overhead latency associated with data transfer to the object store in the case of FabricPool

- The latency automatically applied by the QoS to keep IO within service level maximums
- The user and Snapshot copy data that is counted as part of the used capacity
- The allocated absolute minimum IOPS on each ONTAP volume, regardless of the amount of data in the volume:
  - Extreme: 1,000 IOPS
  - Premium: 500 IOPS
  - Performance, Standard, and Value: 75 IOPS
- While using the Advanced Data Protection add-on services, the target latency applies only to servicing IO requests from the local storage.

## Volume AQoS

Each ONTAP volume should have the applicable adaptive quality of service (AQoS) policy applied. Otherwise, the capacity within each volume that does not have an AQoS policy applied is billed at the rate of the highest Service Level.

## 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, such as those for booting of multiple systems, may consume most or all of the resources for a period, and affect other workloads. For information about QoS, see [Guarantee throughput with QoS overview](#).

## Adaptive QoS

Adaptive QoS (AQoS) is used by Keystone services to dynamically maintain the IOPS/TiB ratio based on the volume size. For information about AQoS policies, see [About adaptive QoS](#).

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

An ONTAP volume is non-compliant if it does not have an AQoS policy applied. A volume without a QoS policy is the last on the list of priority for the system to provide any available input-output operations. However, if any input-output operations are available, then the volume could consume all available IOs.



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

## Adaptive QoS settings

The Adaptive QoS (AQoS) settings vary with service levels.

Policy name	Extreme	Premium	Performance	Standard	Value
Expected IOPS	6,144	2,048	1,024	256	64

<b>Expected IOPS Allocation</b>	Allocated space				
<b>Peak IOPS</b>	12,288	4,096	2,048	512	128
<b>Peak IOPS Allocation</b>	Used space				
<b>Block Size</b>	32K				

## Configuration of adaptive QoS policy group

You can configure adaptive QoS (AQoS) policies to automatically scale a throughput ceiling or floor to volume size. Not all Keystone service levels are aligned with the default ONTAP QoS policies. You can create custom QoS policies for them. For configuring a policy, you should be aware of the following:

- **Policy group name:** The name of the AQoS policy group. For example, `Keystone_extreme`.
- **VServer:** The name of the VServer or storage VM (storage virtual machine).
- **Expected IOPS:** The minimum number of IOPS, per allocated TiB per volume, that the system attempts to provide when enough system IOPS are available.
- **Peak IOPS:** The maximum number of IOPS, per used TiB per volume, that the system allows the volume to reach before it throttles the IOPS through injection of latency.
- **Expected IOPS allocation:** This parameter 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 parameter 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 250 for `Performance`, and 75 for `Standard` and `Value` 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 (AQoS) 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 32K -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 32K -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 32K -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:

```
gos 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
32K -absolute-min-iops 75
```

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

### Settings for Value service level

Settings and commands for the Value service level:

- Sample command:

```
gos 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 32K -absolute
-min-iops 75
```

- 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 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:

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

## Supported storage in Keystone STaaS services

Keystone STaaS services support file and block storage of ONTAP, object storage of

## StorageGRID platform, and data management capabilities of Cloud Volumes ONTAP.

Keystone STaaS provides standard and optional services for your storage.

**Keystone STaaS standard services:** Standard services are included within the base subscription and are not charged separately.

**Keystone STaaS add-on services:** These are optional, chargeable services that provide additional utilities and benefits on top of standard Keystone STaaS subscription services.

Keystone STaaS services can coexist with each other. For example, a cloud storage subscription can co-term with file, block, and object storage subscriptions. A cloud service can be included at any point during the service term of an existing storage subscription. However, if you do not plan to renew an existing file, block, and object subscription, a cloud storage subscription cannot be added during the last 90 days of the subscription.

### Services for file, block, and object storage

Keystone STaaS services for ONTAP file and block storage, and StorageGRID object storage, support multiple features and protocols, and described in the following table:

Storage	Platform	Protocols	Supported features
File storage	ONTAP	NFS and CIFS	Supported ONTAP features: <ul style="list-style-type: none"><li>• FlexVol</li><li>• FlexGroup</li><li>• Snapshot copies</li><li>• SnapMirror (Asynchronous)</li><li>• SnapVault</li><li>• SnapLock Enterprise</li><li>• FabricPool/Cloud tiering</li><li>• SnapRestore</li><li>• FlexClone</li><li>• SnapCenter (license is included but is not a part of Keystone services, and management is not guaranteed)</li><li>• Autonomous ransomware protection<sup>1</sup></li></ul>



Storage	Platform	Protocols	Supported features
Block storage	ONTAP	FC and iSCSI	Supported ONTAP features: <ul style="list-style-type: none"> <li>• FlexVol</li> <li>• FlexGroup</li> <li>• Snapshot copies</li> <li>• SnapMirror (Asynchronous)</li> <li>• SnapVault</li> <li>• SnapLock Enterprise</li> <li>• FabricPool/Cloud tiering</li> <li>• SnapRestore</li> <li>• FlexClone</li> <li>• SnapCenter (license is included but is not a part of Keystone services, and management is not guaranteed)</li> </ul>
Object storage	StorageGRID	S3	Supports multiple information lifecycle management (ILM) policies across multiple sites <sup>2</sup>



<sup>1</sup> For information about ransomware protection in ONTAP, see [Autonomous Ransomware Protection](#).

<sup>2</sup> Each site requires a separate subscription.

## Services for cloud storage

Keystone STaaS provides cloud storage services. Keystone STaaS supports Cloud Volumes ONTAP data management capabilities on Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform.



Hyperscaler-based compute, storage, and network services required by Cloud Volumes ONTAP are not provided by NetApp as a part of Keystone STaaS subscriptions; these subscriptions need to be procured directly from hyperscaler cloud service providers.

## Supported storage capacities in Keystone

The NetApp Keystone STaaS service capacities include:

## Logical capacity

This is the data placed on the Keystone infrastructure by a customer. All Keystone capacities refer to a logical capacity.

For example, if a 1 TiB file is stored on the Keystone infrastructure then a minimum of 1 TiB of capacity should be purchased.

## Committed capacity

The minimum logical capacity billed each month during the subscription:

- Capacity is committed to each service level.
- Committed capacity and additional service levels can be added during the term.

## Changes to committed capacity

During the tenure of a subscription, you can change the committed capacities. However, there are certain preconditions:

- The committed capacity can be decreased based on certain conditions. For information, see [Capacity reduction](#).
- The committed capacity cannot be increased 90 days prior to the expiry of your subscription, unless the subscription is to be renewed for an additional 12-month term.
- You can request changes to committed capacity through the BlueXP interface or through Keystone Success Manager (KSM).  
For information about requesting changes, see [NetApp Global Services Support Center](#).

## Consumed capacity

Consumed capacity refers to the capacity (in TiB of storage) currently being consumed on the service. Keystone service considers the sum of the logical used sizes (not the physical capacity used) of all volumes on a particular service level to calculate the consumed capacity for that service level.

## Burst capacity

Keystone service enables you to use additional capacity on top of the committed capacity for a service level. This is the burst capacity usage. Note these points:

- Burst capacity is agreed upon in the Keystone agreement. It is usually set up to 20% above the committed capacity, and is charged at the same rate as the committed capacity.
- Burst capacity can be consumed on an elastic basis and is charged on a daily basis for the consumed average.

## Billed capacity

Monthly bill = (committed capacity [TiB] \* committed rate [\$/TiB]) + (daily average provisioned burst capacity [TiB] \* burst rate [\$/TiB]). The monthly bill contains a minimum charge based on the committed capacity.

The monthly bill varies beyond the minimum charge based on daily average burst capacity consumption.

# Service levels in Keystone

Keystone STaaS offers data storage capacity at pre-defined performance service levels (service levels) on a subscription basis. Each volume managed by the Keystone services is associated with a service level.

A subscription can have multiple rate plans and each rate plan corresponds to a service level. Each rate plan has a committed capacity per service level.

Each service level is defined by its I/O density, that is IOPS/TiB/volume. This is the ratio of performance (input/output operations per second [IOPS]) and used storage capacity (TiB) which is IOPS/TiB at average latency per volume.

You select service levels based on your storage environment, and storage and consumption needs. The base service levels are available for you by default. Specific service levels are additionally available, if you have opted for add-on services. For example, for the advanced data protection add-on service, the *Advanced Data-Protect* service level is assigned to your subscription.



A detailed service description for NetApp Keystone STaaS service levels is available [here](#).

The base service levels for the supported storage types, file, block, object, and cloud services are described in the following sections:

## Service levels for file and block storage

**Supported protocols:** NFS, CIFS, iSCSI, and FC

Service level	Extreme	Premium	Performance	Standard	Value
<b>Sample workload types</b>	Analytics, databases, mission-critical apps	VDI, VSI, software development	OLTP, OLAP, containers, software development	File shares, web servers	Backup
<b>Maximum IOPS/logical TiBs stored per volume</b>	12,288	4,096	2,048	512	128
<b>Maximum IOPS/logical TiBs allocated per volume</b>	6,144	2,048	1,024	256	64
<b>Maximum MBps/logical TiBs stored per volume @ 32K B/S</b>	384	128	64	16	4
<b>Target 90<sup>th</sup> percentile latency</b>	<1 ms	<2 ms	<4 ms	<4 ms	<17 ms
<b>Block size</b>	32K				

## More on service levels for file and block storage

The base service level metrics depend on the following conditions:

- The service levels for file and block storage support ONTAP 9.7 and later.
- IOPS/TiB/volume, MBps/TiB/volume, and latency values for service levels are based on the amount of data stored in the volume, 32KB block size, and a random combination of 70% read and 30% write IO operations.
- Actual IOPS/TiB/volume and MBps/TiB/volume may vary based on the actual or assumed block size, system workload concurrency, or input-output operations.
- Latency does not include the following:
  - application or host latency
  - customer network latency to or from the controller ports
  - overheads associated with the data transfer to the object store in case of FabricPool
  - latency automatically applied by QoS to keep IO within service level maximums
- Latency values are not applicable to MetroCluster write operations. These write operations are dependent on the distance of remote systems.
- If one or more volumes on a storage system do not have an AQoS policy assigned, then these volumes are considered as non-compliant volumes, and no target service levels are applicable for those systems.
- *Expected IOPS* is targeted for FabricPool only if the tiering policy is set to "none" and no blocks are in the cloud. *Expected IOPS* is targeted for volumes that are not in a SnapMirror synchronous relationship.
- Workload IO operations need to be balanced across all deployed controllers, as determined by the Keystone order.

## Object storage

**Supported protocol:** S3

<b>Service level</b>	Object
<b>Workload type</b>	Media repository, archiving
<b>Maximum IOPS/logical TiB stored per volume</b>	N/A
<b>Maximum MBps/logical TiB stored per volume</b>	N/A
<b>Average Latency</b>	N/A



Latency does not include overheads associated with data transfer to the object store in case of FabricPool storage.

## Cloud storage

**Supported protocol:** NFS, CIFS, iSCSI, and S3 (AWS and Azure only)

<b>Service level</b>	Cloud Volumes ONTAP
<b>Workload type</b>	Disaster Recovery, software development/testing, business apps

<b>Maximum IOPS/logical TiB stored per volume</b>	N/A
<b>Maximum MBps/logical TiB stored per volume</b>	N/A
<b>Average Latency</b>	N/A



- Cloud native services, such as compute, storage, networking, are invoiced by cloud providers.
- These services are dependent on cloud storage and compute characteristics.

### Related information

- [Supported storage capacities](#)
- [Metrics and definitions used in Keystone Services](#)
- [Quality of Service \(QoS\) in Keystone](#)
- [Keystone pricing](#)

## 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

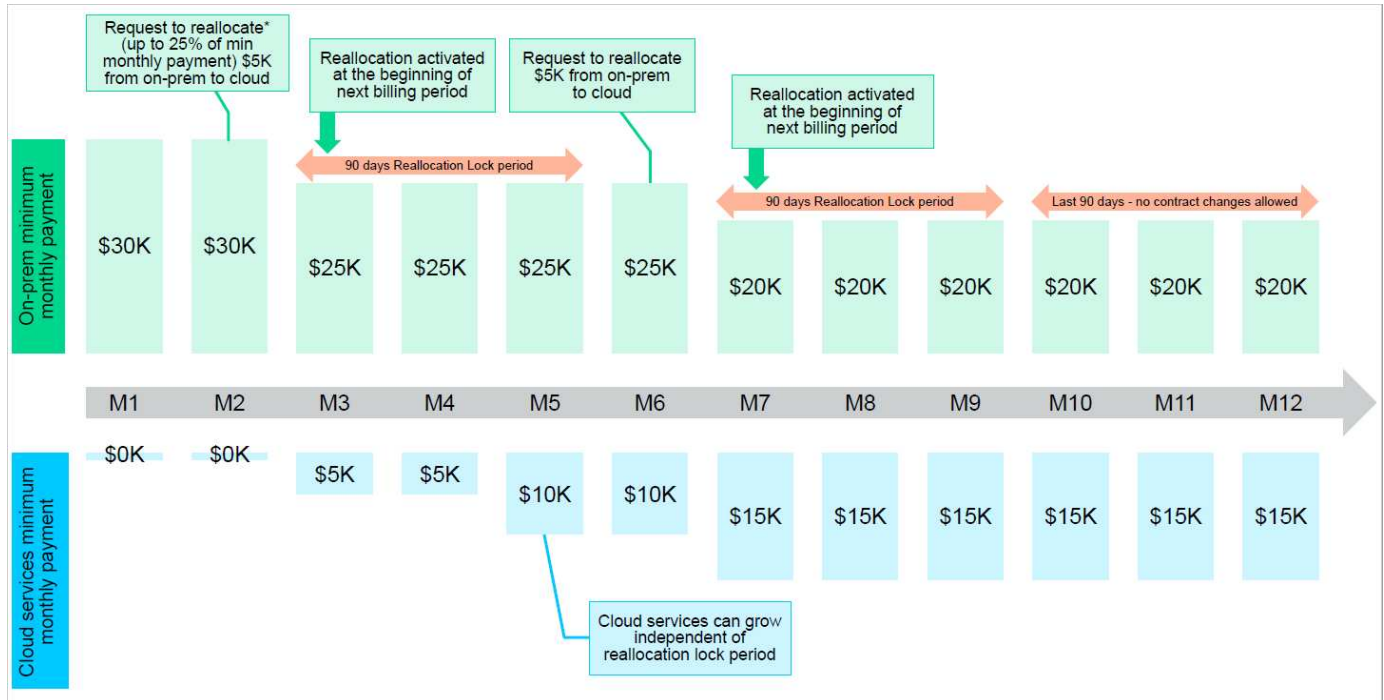
Note the following clauses for capacity adjustments:

- 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 have additional burst requirement greater than the capacity you 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.



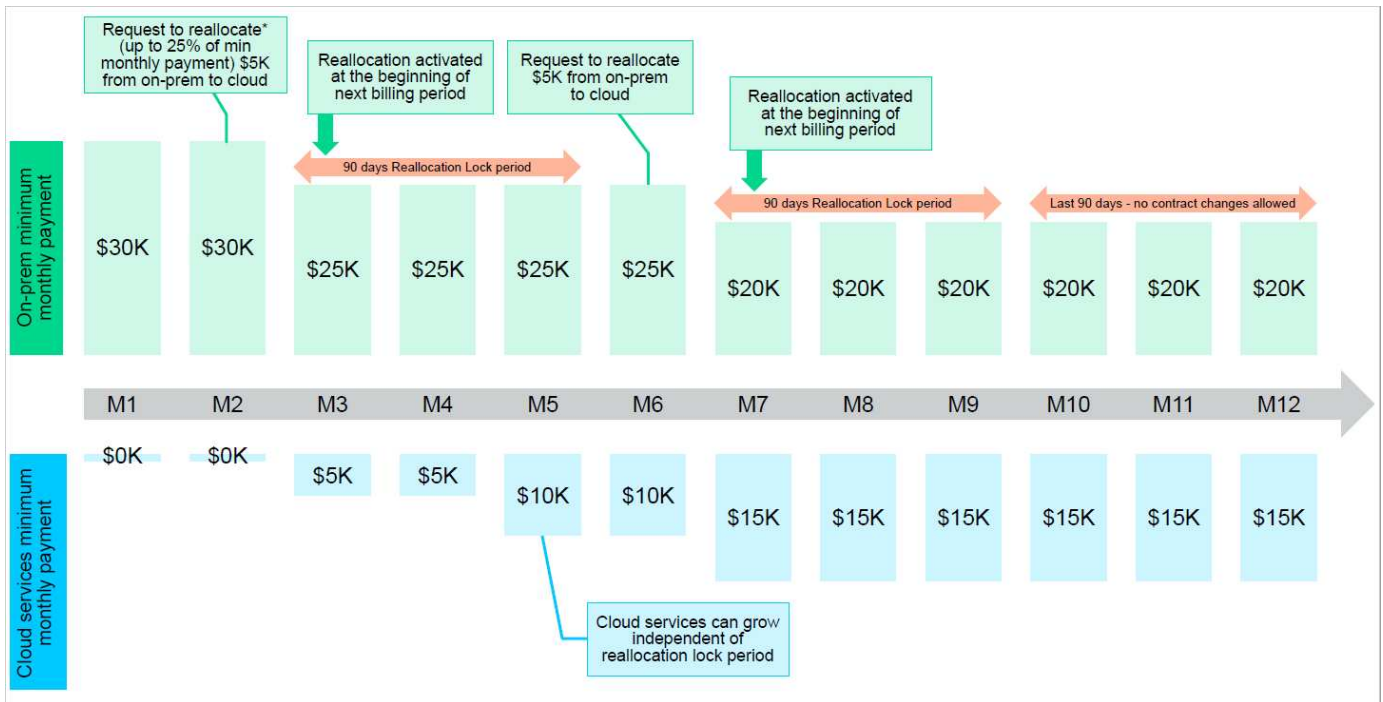
### Quarterly spend reallocation

Keystone STaaS offers you the option to reallocate on-premises service spend to Cloud Volumes ONTAP spend.

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.
- Capacity for Cloud Volumes ONTAP and Cloud Backup service 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.
- 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 services.
- 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 consecutive billing cycle.

You can allocate a portion of your expenses towards your subscribed file, block, or object storage service levels to hybrid cloud storage services. Up to 25% of the Annual Contract Value (ACV) can be reallocated on a quarterly basis to Cloud Volumes ONTAP Primary and Cloud Volumes ONTAP Secondary services:



This table provides a set of sample values to demonstrate how the reallocation of expenses works. In this example, \$5000 from the monthly spend is reallocated to hybrid cloud storage service.

Before allocation	Capacity (TiB)	Monthly designated expense
Extreme	125	37,376
After reallocation	Capacity (TiB)	Monthly designated expense
Extreme	108	37,376
Cloud Volumes ONTAP	47	5,000
		37,376

The reduction is of  $(125-108) = 17$  TiB of the capacity allocated for the Extreme service level. On spend reallocation, the allotted hybrid cloud storage is not of 17 TiB but an equivalent capacity that \$5000 can purchase. In this example, for \$5000, you can get 17 TiB on-prem storage capacity for the Extreme service level and 47 TiB hybrid cloud capacity of Cloud Volumes ONTAP service level. Therefore, the reallocation is with respect to the spend, not capacity.

Contact your Keystone Success Manager (KSM) if you want to reallocate expenses from your on-premises services to cloud services.

## Keystone subscription services | Version 1

Keystone STaaS was preceded by Keystone subscription services (previously known as Keystone Flex Subscription services).

While the navigation of the two offerings is similar in the [Keystone dashboard](#), Keystone subscription services differ from Keystone STaaS in the constituent service levels, service offering, and billing principles. As of April 2024, NetApp maintains and publishes documentation for only Keystone STaaS. If you are still using Keystone subscription services, contact your KSM for support in migrating to Keystone STaaS. If required, you can access a PDF version of the Keystone subscription services documentation [here](#):



- [English](#)
- [Japanese](#)
- [Korean](#)
- [Chinese \(Simplified\)](#)
- [Chinese \(Traditional\)](#)
- [German](#)
- [Spanish](#)
- [French](#)
- [Italian](#)

## Add-on services

### Advanced data protection

You can subscribe to advanced data protection add-on service as a part of your Keystone STaaS subscription. This add-on service leverages NetApp MetroCluster technology to ensure efficient data protection for your mission-critical workloads at 0 Recovery Point Objective (RPO).



Keystone STaaS standard services for file and block storage offer default data protection services by leveraging NetApp technologies, such as SnapMirror, SnapVault, and Snapshot.

For information about the standard and cloud service, see [Keystone STaaS services](#).

Keystone advanced data protection service can synchronously mirror data to a secondary site. In case of a disaster at the primary site, the secondary site can take over, without any loss of data. This feature leverages the MetroCluster configuration between two sites to enable data protection. You can avail the advanced data protection add-on services for only your file and block storage services. As a part of this add-on service, the `Advanced Data-Protect` service level is assigned to your subscription.

For information about ONTAP MetroCluster, see [MetroCluster Documentation](#).


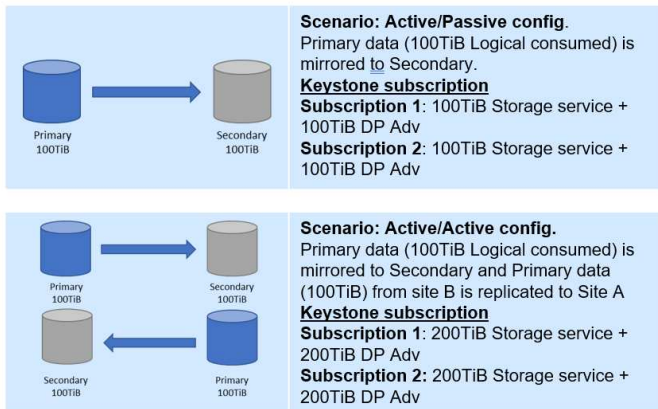
For information about how to view the consumption in a MetroCluster configuration, see [Reference charts for advanced data protection](#).

The add-on charges are applicable to all the capacities in the subscription: the source, mirrored copy, and unmirrored data.

Note the following:

- For this service, 100% of the committed capacity on an associated storage is configured as committed capacity.
- Storage is charged for both the source and the destination clusters. This add-on service is also charged for both the primary and secondary (mirrored) volumes.
- The charges are applicable only to your file and block storage.

The following MetroCluster scenarios are supported:



## Data tiering

Keystone STaaS standard services for file and block storage include tiering capabilities that identify less-frequently used data, and tiers it to Keystone STaaS-supported NetApp cold storage. You can use data tiering as an add-on service if you want to tier your cold data to any Keystone STaaS-supported, non-NetApp storage.

For information about standard data and add-on services, see [Keystone STaaS services](#).

For information about Service Levels, see [Service Levels in Keystone](#).

The tiering add-on service is required only when data is tiered to any non-NetApp storage such as Amazon Web Services (AWS) S3, Azure Blob, Google Cloud Platform (GCP), and other, Keystone STaaS-supported, S3-compatible, third party object storage.

The tiering capability leverages the NetApp FabricPool technology that enables automated tiering of infrequently accessed data to object storage tiers on and off premises.

The add-on data tiering service enables tiering from Extreme, Premium, Performance, Standard, and Value tier to an object storage target. The ratio of hot to cold data to be tiered is not fixed, and each tier is metered and invoiced separately.

For example, if the target for cold storage tier is:

- Keystone STaaS Value tier, Keystone STaaS StorageGRID Object Tier, or existing SGWS grid (customer owned) - There is no additional charge; it is part of the standard service.
- Public cloud (AWS, Azure, Google) or Keystone STaaS-supported, third party object storage - There is an additional charge for data capacity that is tiered to cold storage target.

The charges for add-on tiering services apply through the entire subscription term.



Hyperscaler-based compute, storage, and network services required by Cloud Volumes ONTAP are not provided by NetApp as a part of Keystone STaaS subscriptions; these services need to be procured directly from hyperscaler cloud service providers.

## Keystone co-location services powered by Equinix

NetApp has partnered with Equinix for hosting NetApp Keystone STaaS services in an Equinix data center to ensure the delivery of a unified solution for you.

Keystone co-location (Co-Lo) services powered by Equinix remain unchanged from the standard Keystone service offering.

In this service:

- Equinix provides space, power, cooling, networking, invoice, and storage, on a monthly basis.
- Support is provided for certain Keystone sales motions.
- The services are supported in the data centers across 11 countries.

Equinix has the data centers at these locations:

<b>Data center</b>	<b>Country</b>
Amsterdam	Netherlands
Atlanta	U.S.
Chicago	U.S.
Dallas	U.S.
Denver	U.S.
Frankfurt	Germany
London	U.K.
Los Angeles	U.S.
Madrid	Spain
Melbourne	Australia
Miami	U.S.
Milan	Italy
Osaka	Japan
Paris	France
Seattle	U.S.
Silicon Valley	U.S.
Sydney	Australia

Data center	Country
Tokyo	Japan
Toronto	Canada
Washington DC	U.S.
Zurich	Switzerland

## Non-returnable, non-volatile components and SnapLock compliance

As a part of NetApp Keystone subscription, NetApp extends the non-returnable, non-volatile components (NRNVC) offering for your file, block, and object services.

You can subscribe to this add-on service as a part of your Keystone subscription. For information about the standard and cloud services, see [Keystone STaaS services](#).

For information about Service Levels, see [Service Levels in Keystone](#).

NetApp does not recover the physical storage media used during the entire service tenure or at service termination when NetApp otherwise recovers all of its physical assets used in the delivery of the service.

If you have purchased this service, note the following:

- You do not need to return any drives and nonvolatile memory at end of the service term or if they failed or were found defective during the service term.
- However, you need to produce a certificate of destruction for the drives and/or nonvolatile memory and cannot be used for any other purpose.
- The additional cost associated with the NRNVC is charged as a percentage of the total subscription services (includes standard service, Advanced Data Protection, and data tiering) monthly bill.
- This service is applicable only to file, block, and object services.

### SnapLock compliance

The SnapLock technology enables the NRNVC feature by making the drive unusable after the expiry date set in the volume. For using the SnapLock technology on your volumes, you need to subscribe to NRNVC. This is applicable only to file and block services.

For information about SnapLock technology, see [What SnapLock is](#).

### U.S. Citizen Support (USCS)

United States Citizen Support (USCS) is an add-on offering for NetApp Keystone Subscriptions. It entitles you to receive delivery and support of ongoing Keystone services from U.S. citizens on U.S. soil.

Read the following sections to understand which elements of your subscriptions are bound by this add-on service and are provided under the terms of NetApp Keystone Agreement. <sup>[1]</sup>

## **NetApp Global Services Support Center monitoring**

NetApp Global Services and Support Center (GSSC) monitors the health of your products and subscribed services, provides remote support, and collaborates with your Keystone Success Manager. All personnel monitoring the products associated with the relevant Keystone subscription orders are U.S citizens operating on U.S. soil.

## **Keystone Success Manager**

The Keystone Success Manager (KSM) is a U.S. citizen operating on U.S. soil. Their responsibilities are specified in your NetApp Keystone Agreement.

## **Deployment activities**

Where available, onsite and remote deployment and installation activities are conducted by U.S. citizens on U.S. soil. <sup>[2]</sup>

## **Support**

Where available, the necessary onsite troubleshooting and support activities are conducted by U.S. citizens on U.S. soil. <sup>[2]</sup>

## Copyright information

Copyright © 2024 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.

[1] The services and offerings described here are subject to, and limited and governed by a fully-executed Keystone Agreement.

[2] Availability of appropriate personnel for onsite activities is dependent on the geographical location at which the Keystone systems are deployed.