



# Documentation for NetApp Keystone Flex Subscription

Keystone

NetApp  
June 09, 2022

This PDF was generated from <https://docs.netapp.com/us-en/keystone/index.html> on June 09, 2022.  
Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

- Documentation for NetApp Keystone Flex Subscription . . . . . 1
  - Flex Subscription . . . . . 1
  - Flex Subscription service terms and descriptions . . . . . 3
  - Operational model, roles, and responsibilities . . . . . 12
  - Tenancy and multi-tenancy in Flex Subscription . . . . . 14
  - Flex Subscription infrastructure . . . . . 16
  - Site requirements . . . . . 18
  - Flex Subscription Services Operations . . . . . 19
  - What can Flex Subscription customers view in Active IQ? . . . . . 20
- Release Notes . . . . . 22
  - What’s new in this release of NetApp Keystone Flex Subscription services . . . . . 22
  - Fixed issues in NetApp Service Engine . . . . . 27
  - Known issues in NetApp Service Engine . . . . . 28
- NetApp Keystone frequently asked questions (FAQs) . . . . . 31
  - NetApp Keystone Flex Subscription FAQ . . . . . 32
  - Flex Subscription service offer details . . . . . 34
  - Operational models and responsibilities . . . . . 35
  - NetApp Service Engine/Self-service access portal . . . . . 37
- NetApp Service Engine web interface . . . . . 39
  - Billing accounts, subscriptions, services, and performance . . . . . 40
  - Get started . . . . . 46
  - View Flex Subscription Dashboard . . . . . 52
  - View billing . . . . . 54
  - Overview . . . . . 54
  - Overview . . . . . 64
  - Work with object storage . . . . . 72
  - Manage Cloud services . . . . . 78
  - View reports . . . . . 86
  - Back up file shares and disks . . . . . 89
  - Managing subscriptions . . . . . 91
  - Manage service requests . . . . . 93
  - Perform administrative tasks . . . . . 97
  - Define network configurations for tenants and subtenants . . . . . 101
- Overview of NetApp Service Engine APIs . . . . . 103
  - Target audience . . . . . 103
  - NetApp Service Engine API access and categories . . . . . 103
  - Key NetApp Service Engine concepts . . . . . 103
  - Authorization and authentication . . . . . 104
  - NetApp Service Engine REST APIs . . . . . 106
  - Consumer APIs . . . . . 106
  - Administrator APIs . . . . . 167

# Documentation for NetApp Keystone Flex Subscription


NetApp Keystone Flex Subscription is a part of the NetApp Keystone service offering. NetApp Keystone is a portfolio of on-premises capital expenditure (capex) alternatives to enable the on-ramp to cloud journey, which includes NetApp Keystone Flex Pay and NetApp Keystone Flex Subscription services. NetApp Keystone offers a seamless hybrid cloud experience with storage services that span on your premises and in the cloud.

For more information about NetApp Keystone, see [NetApp Keystone](#)

For more information about NetApp Keystone Flex Subscription, see [Keystone Flex Subscription](#)

Flex Subscription is a storage-as-a-service offering that delivers a cloud-like experience on your on-premises systems with connectivity to the cloud — available on your terms.

- **NetApp Keystone Flex Pay** - NetApp Keystone Flex Pay (Flex Pay) is a portfolio of flexible financial alternatives that include traditional financing, leasing, and fixed/variable options to meet your customer's cash flow needs.
- **NetApp Keystone Flex Subscription** - NetApp Keystone Flex Subscription (Flex Subscription) is pay-as-you-grow, outcome-based subscription service that brings an on-premises cloud-like experience.



## Flex Subscription

Flex Subscription is an pay-as-you-grow subscription-based service model that delivers a seamless hybrid cloud experience for those preferring OpEx consumption models to upfront CapEx or leasing. It enables customers to accelerate time to value by reducing the hurdles related to managing the unpredictable capacity growth and going through the complex procurement cycles. Flex Subscription allows customers to align economics and operations to their business priorities.



**Prefer 100% opex**  
(no asset ownership)



**Single orchestration**



**Monitor, manage, & optimize usage**



**Flexible terms**  
(choice of term & capacity)



**Cloud economics on premises**

Flex Subscription is a subscription-based service offering that provides storage capacity, at predefined performance service levels, for block, file, and object data types that can be deployed on-premises and can be operated by NetApp, a partner, or the customer. In addition to the base service that Flex Subscription provides, add-on services can be chosen at an additional cost. These services are described in the below figure.

- **Data Protection Basic** service provides NetApp SnapMirror and NetApp SnapVault capabilities with default settings.
- **Data Protection Advanced** provides a service with SLAs of RPO=0 by leveraging the NetApp MetroCluster capabilities.



Flex Subscription uses SnapVault technology for backup operations and SnapMirror for disaster recovery to replicate the data remotely; it does not automatically fail over and restore applications.



**Base package**

All Flex Subscription services include:

- FlexVol®, FlexGroup®, protocols, clone, and encryption capabilities by default
- Snapshot™ default: 4 hours, 7 days retention
- Free SnapMirror® for initial ingest, provided Premium Bundle on source system



**DP Basic\***  
(for additional fee)

Data protection Basic Services subscription

- Includes SnapVault® and SnapMirror
- Requires additional storage subscription to store replicated data at the target; target can be already purchased and owned NetApp storage
- Disaster recovery snapshot every hour; replicated every 4 hours; retained 7 days
- Backup with SnapVault (snapshot every 4 hours; replicated every 24 hours; retained 7 days)



**DP Adv.\***  
(for additional fee)

Data protection Advanced Services subscription

- MetroCluster™ / RPO 0 Service
- Pre-requisite – data protection Basic

\* Available only for file and block data services.

## Benefits of Flex Subscription

Flex Subscription provides the following benefits:

- Frees up IT staff from complicated storage-related tasks and allows them to focus on application management
- Reduces upfront capital investment
- Allows customers to meet their demands without overprovisioning

- Aligns data storage costs with business needs/activity
- Simplifies infrastructure provisioning by bypassing complex organizational procurement procedures
- Keeps data secure on their premises
- Enables proper control over compliance, performance, and security



**Align costs and usage**  
Reduce upfront cash and pay only for what you use, avoiding overbuying and overprovisioning



**Free IT staff**  
Free up IT to focus more on innovation and business priorities and less on typical storage tasks



**Unlock the best of both clouds**  
Scale across clouds easily and leverage the public cloud for bursting, data migrations, DR, backup, and tiering



**Meet regulatory requirements**  
Gain improved performance, data protection, compliance, and security with certified and field-proven deployments and services

## Flex Subscription service terms and descriptions

- NetApp Keystone Flex Subscription (Flex subscription) is available for a minimum of one year and up to three years. After the initial term, the service is renewable on an annual basis. Capacity can be increased in increments as small as 1 TiB.
- The minimum capacity is 100 TiB per site, and each site can have one or more clusters to meet the minimum capacity requirement. In a partner-operated model, subscriptions with flexible minimums are created for a customer, per site and across service levels.
- The 100 TiB capacity can be one single performance level or a combination of levels.
- Tenant subscriptions are limited to service levels that partners are subscribed to.
- 20% of burst capacity is available at each site; any burst usage is billed only for that billing period. If you need an additional burst requirements that is greater than 20%, contact support.
- Committed capacity or performance levels cannot be altered during a contract term.
- Increasing capacity or changing to higher performance level during term is allowed; however, moving from a higher performance level to a lower level is not permitted.
- Any change request in the last 90 days of the term requires the customer to renew the service for a minimum of one year.

## Flex Subscription service capacity definitions

The NetApp Keystone Flex Subscription (Flex Subscription) service capacities include:

### Logical capacity

This is the data placed into the Flex Subscription infrastructure by a customer. All Flex Subscription capacities refer to a logical capacity. For example, if a 1 TiB file is stored on the Flex Subscription infrastructure, then at least 1 TiB of capacity must be purchased.

## Committed capacity

The minimum logical capacity billed each month for the duration of the term:

- Capacity is committed to each performance level.
- Committed capacity cannot be decreased during the term.
- Committed capacity and additional performance 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 cannot be decreased
- 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 NetApp Service Engine interface or through Keystone Success Manager.  
For information about requesting changes, see [Raise a service request](#).

## Burst capacity

This is the logical capacity that has exceeded the committed capacity. Note the following points:

- Flex Subscription service provides 20% capacity in excess of the committed capacity.
- Burst capacity can be consumed on an elastic basis and is charged on a daily basis of the consumed average.
- Burst capacity up to 20% is charged at a same rate as the committed capacity.
- Burst capacity greater than 20% of committed is charged at a premium rate. Contact support for any additional burst requirements greater than 20%.

## Consumed/provisioned capacity

Consumed capacity refers to the capacity in TiB of storage currently being consumed on the service. Flex Subscription service considers the sum of the provisioned sizes (not the logical or physical capacity used) of all volumes on a particular Performance Service Level to be considered as consumed capacity for that Performance Service Level. This includes:

- The capacity that is provisioned through the creation, modification, deletion, or potential auto-growth of volumes.
- The Snapshot copies and clones.



The amount of data stored within provisioned capacity, or the amount of data actually written to disk is not considered.

## 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. For more information on billing, see [Flex Subscription billing](#).

## Performance Service Levels

NetApp Keystone Flex Subscription (Flex Subscription) offers capacity at predefined performance levels.

Each Performance Service Level is defined by its I/O density, which is the ratio of performance (input/output operations per second [IOPS]) and used storage (TiB of stored data) which is IOPS/TiB.

Each volume managed by the Flex Subscription services is associated with a performance service level. All I/O operations and all used storage used on the respective volume are factored into the volume's I/O density calculation.

The below table defines the performance service levels.

### Performance service levels for difference storage types

The service levels for file, block, and object storage are listed here.

I/O density calculations at the volume level are reported to show peak I/O density during the prior week. The peak performance is determined on an hourly time interval. I/O density reports by volume are generated monthly to gauge adherence to the respective service levels.

#### File service

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

Service level	Extreme	Premium	Standard	Value
<b>Workload type</b>	Analytics, databases	VDI, virtualization apps, Software dev	File shares, web servers	Backup
<b>Target IOPS/TiB</b>	6,144	2,048	128	N/A
<b>Max IOPS/TiB</b>	12,288	4,096	512	N/A
<b>Max throughput MBps (32KB/IOP)</b>	384	128	16	N/A
<b>Latency</b>	<1 ms	<2 ms	<17 ms	N/A
<b>Minimum capacity<sup>1,2</sup></b>	100 TiB <sup>1</sup>			
	15 TiB <sup>2</sup>	25 TiB <sup>2</sup>	50 TiB <sup>2</sup>	50 TiB <sup>2</sup>



<sup>1</sup> Minimum one-year term and 100TiB minimum capacity for a combination of any file share and block Performance Service Levels for NetApp direct opportunities.

<sup>2</sup> Minimum one-year term and minimum stated capacity per each service tier selected for channel led opportunities.

#### Block storage

**Supported protocols:** FC and iSCSI

Service level	Extreme	Premium	Standard
Workload type	HPC	Video surveillance	Backup
Target IOPS/TiB	N/A		
Max IOPS/TiB	5,500	4,000	N/A
Max throughput MBps (32KB/IOP)	43	31	N/A
Latency	<0.5 ms	<0.5 ms	N/A
Minimum capacity <sup>1,2</sup>	100 TiB <sup>3</sup>	100 TiB <sup>3</sup>	300 TiB <sup>3</sup>



<sup>1</sup> Minimum one-year term and 100TiB minimum capacity for a combination of any file share and block Performance Service Levels for NetApp direct opportunities.

<sup>2</sup> Minimum one-year term and minimum stated capacity per service tier selected for channel led opportunities.

<sup>3</sup> Minimum one-year term and minimum stated capacity per service tier selected for either opportunity type.

### Object storage

#### Supported protocol: S3

Service level	Object
Workload type	Media repository, archiving
Target IOPS/TiB	N/A
Max IOPS/TiB	N/A
Max throughput MBps (32KB/IOP)	N/A
Latency	N/A
Minimum capacity <sup>1,2</sup>	500 TiB <sup>3</sup>



<sup>1</sup> Minimum one-year term and 100TiB minimum capacity for a combination of any file share and block Performance Service Levels for NetApp direct opportunities.

<sup>2</sup> Minimum one-year term and minimum stated capacity per service tier selected for channel led opportunities.

<sup>3</sup> Minimum one-year term and minimum stated capacity per service tier selected for either opportunity type.

## Service Level metrics and definitions

The following terms and definitions are used within the NetApp Keystone Flex Subscription (Flex subscription) service:

- **GiB, TiB, and 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 1PiB = 1024<sup>5</sup> bytes).
- **IOPS/TiB.** The protocol operations per second requested by the application divided by the allocated logical size of the volume.



- **Availability** is measured as a percentage of number of I/O requests successfully responded to by the service, divided by total number of I/O requests made of the service, measured at the service demarcation, in a given month, not including scheduled service downtime or unavailability of required facilities, network or other services to be provided by customer.
- **Durability** is the percentage of data accessed without loss of fidelity, excluding customer-caused deletion or corruption.
- **Target IOPS per TiB.** The guaranteed IOPS for all I/O requests made to a volume before the target IOPS per TiB threshold is reached. Performance on the volume is capped at the selected IOPS per TiB.



The target IOPS per TiB performance metric is calculated based on the logical consumed capacity in TiB.

- **Latency.** Time to service an I/O request received from a client, measured at the service demarcation (storage controller I/O port).

## Flex Subscription billing

NetApp Keystone Flex Subscription (Flex Subscription) enables predictable and upfront pricing for your storage subscription.

If you prefer operational expenditures (OpEx) consumption model to capital expenditure (CapEx) or leasing, you can opt for the Flex Subscription pay-as-you-grow model for your flexible and scalable consumption needs.

Flex Subscription provides you with the following billing facilities:

- You can pay based on IOPS and latency committed capacity to meet various workload needs. The different performance service tiers - Extreme, Premium, Standard, and Value enable you to manage your storage based on your purchased service level for your Flex Subscription.
- It presents predictable billing for the committed capacity and pay-per-use for variable (burst) capacity usage.
- You can select a bundle price for hardware, core OS, and support for one \$/TiB price. You have a single invoice for each storage type, file, block, object, or cloud storage services.
- Select a flexible term for the services and payment: You can opt for 12 months, 100TiB, or more per site. Thereafter you can auto renew for 12 months or go month-to-month.

Flex Subscription billing is based on committed capacity and variable burst consumption.

For information about committed and burst capacity usage, see [Flex Subscription service capacity definitions](#).

For information about how to view billing details, see [View billing](#).

### Billing based on committed capacity

Committed capacity refers to the capacities for various services in a single subscription, agreed upon by the parties involved (NetApp/partner and customer). This capacity is stated on each Flex Subscription order and is billed, regardless of the actual consumption.

### Metering of consumed capacity

As a part of the Flex Subscription service deployment, NetApp continuously monitors and measures the consumption of the service. Every five minutes, a consumption record is generated by the system, detailing the

current consumed capacity for the subscription. These records are aggregated over the billing period to generate invoices and usage reports.

### **Billing based on burst consumption**

When the consumed capacity is greater than the committed capacity for a given Performance Service Level, burst consumption is recorded, and charges are applied accordingly. This process occurs for each consumption record generated. Burst consumption, therefore, is a reflection of both the amount and tenure of your over-consumed capacities on top of your committed capacities.

### **Billing schedules**

Flex Subscription services are billed monthly and yearly.

#### **Monthly billing**

Invoices are sent monthly. For the month in which the services are availed, an invoice is sent in the next month. For example, the invoice for the services you have used in January is delivered at the beginning of February. This invoice includes the charges for the committed capacity and if applicable, any burst usage.

#### **Annual billing**

An invoice is generated at the beginning of each subscription year for the minimum payment of the committed capacity. It is generated on the start date of the subscription.

Another invoice is sent at the end of a subscription quarter, summing up the applicable charges of any burst usage accrued in that quarter.

If the committed capacity is changed during a subscription, then an invoice is sent on the same day the change in the committed capacity is effective, for the prorated minimum payments for the rest of that subscription year.

### **Miscellaneous scenarios for Flex Subscription billing**

There are several scenarios for Flex Subscription billing and you should be familiar with those scenarios.

#### **Billing for cloned volumes**

If volumes are cloned in ONTAP and you use them for backing up and restoring your data, you can continue using the clones without any additional payments. However, cloned volumes used for any other purpose in your business for an extensive duration are charged.

Note the following:

- The consumption on the cloned volumes is not considered during the first 24 hours of the clone creation; and no charges are incurred during this time.
- Cloned volumes with the Standard and Value Performance Service Levels are not considered for consumption, if new data is not written to the volume. For example, in a backup or restore use case, backing up or restoring data from a cloned volume does not change the data in the cloned volume itself, and it is not considered as consumption.

#### **Billing for MetroCluster**

Advanced Data Protection uses NetApp MetroCluster to mirror data between two physically separated clusters. On MetroCluster mirrored aggregates, data is written twice, once on each cluster. Flex Subscription service charges for consumption on each side independently, resulting in two identical consumption records.

If you monitor your clusters through ONTAP System Manager (System Manager) or Active IQ Unified Manager (Unified Manager), you might see a discrepancy between the consumption reported on these tools and Flex Subscription. System Manager and Unified Manager do not report volumes on the mirrored (remote) cluster,

and in doing so, reports half the consumption metrics that the Flex Subscription service reports.

For Example:

Site A and Site B are set up in a MetroCluster configuration. When a user creates a volume of 10TB in site A, an identical volume of 10TB is created in site B. Flex Subscription distinguishes both the volumes and records an additional 10TB of consumption in each site, for a total increase of 20TB. System Manager and Unified Manager reports a 10TB volume created in site A.

### **Billing for temporary volumes**

Occasionally, temporary (TMP) volumes are created by ONTAP when moving volumes. These temporary volumes are short-lived, and the consumption on these volumes is not measured for billing.

### **Billing and adaptive QoS policies**

Flex Subscription measures consumption based on Performance Service Levels. Each Performance Service Level is associated with a specific adaptive quality of service (QoS) policy. During deployment, you will be informed of the details of each QoS policy for your subscribed Flex Subscription services. During storage management operations, ensure that your volumes have the appropriate QoS policies assigned as per your subscribed Performance Service Levels, to avoid unexpected billing.

For more information about QoS policies in ONTAP, see [Guarantee throughput with QoS overview](#).

### **Billing for SnapMirror destinations**

The pricing for the SnapMirror destination volume governed by the QoS policy for the Performance Service Level (service level) assigned on the source. However, if the source does not have an associated QoS policy, the destination is billed based on the lowest available service level.

### **Billing for FlexGroups**

FlexGroups are billed based on the adaptive QoS policy of the FlexGroup. The QoS policies of its constituents are not considered.

### **Billing for LUNs**

For LUNs, usually the same billing pattern is followed as for the volumes that are governed by QoS policies. If separate QoS policies are set on LUNs, then:

- The size of the LUN is counted for consumption according to the associated service level of that LUN.
- The remainder of the space in the volume, if any, is charged according to the QoS policy of the service level set on the volume.

### **System and root volumes**

System and root volumes are monitored as a part of the overall monitoring of the Flex Subscription service but are not counted or billed. The consumption on these volumes is exempted for billing.

## **Data Protection**

NetApp Keystone Flex Subscription (Flex Subscription) data protection service can back up your data and is able to recover it if required. The available data protection services are:

- Snapshots of disks and shares

- Backups of disks and shares (requires data protection service as part of the subscription)
- Disaster recovery for disks and shares (requires data protection service as a part of the subscription)



Backup and disaster recovery services are available as add-on services, while Snapshot is available as a part of the basic storage service.

	<b>Single Region Snapshots (Available as a part of the basic storage service)</b>	<b>Multi-region Backup (data protection add-on)</b>	<b>Multi-region Disaster Recovery (data protection add-on)</b>
Use case	Mitigate the risk of user or application data deletion or corruption, not against infrastructure loss or failure	Mitigate the risk of complete loss of data on the primary volume due to infrastructure loss or failure	Mitigate the risk of complete loss of data on the primary volume due to infrastructure loss or failure with a recovery time objective
Policy	Hourly, daily, weekly, and monthly	Number of backups to retain based on hourly, daily, weekly, and monthly Snapshots	1 hour, 4 hours, and daily
Topology	Source only	Backup	Async replication target
Target replication service level <sup>1</sup>	n/a	Standard	Same as primary

<sup>1</sup>Additional storage capacity to be subscribed



Subscription to a basic Flex Subscription service does not automatically back up your data. You should subscribe to add-on data protection services and configure your system for data backup and disaster recovery services. If your storage system is not managed by Flex Subscription services, NetApp can still support protecting the data on your storage system and help in connecting it with your Flex Subscription services. However, NetApp is not responsible for any backup failures.

## Tiering

NetApp Keystone Flex Subscription service includes a tiering capability that identifies less frequently used data and tiers it to a cold storage that is owned, deployed, and managed by NetApp.

The tiering capability leverages the NetApp FabricPool technology that enables automated tiering of data to low-cost object storage tiers either on or off premises. With this capability, infrequently accessed data is automatically tiered to a lower cost storage either on premises or in the cloud, based on the services agreed upon.

Partners and tenants can avail this capability easily by opting for the two preconfigured service levels, the Extreme-tiering and Premium-tiering service levels while provisioning their storage. The Extreme-tiering has the same QoS policies as the Standard, Extreme, and Premium service levels.

The add-on tiering capability is available only with Extreme and Premium service tiers. NetApp assumes 25%

of data is hot and 75% is less frequently used and can be moved to a cold storage. Billing is determined based on the duration per volume is in each service level.

The following features are enabled:

- You can create reports of the inactive data for your disks and file shares and decide upon whether to change the service level. On moving or changing the tiering policy, the latencies can be higher if data is accessed from cold tier.
- You can change the service level of the volumes from Extreme and Premium to Extreme-tiering and Premium tiering respectively, provided that the destination tiering is enabled on the cluster.
- Likewise, you can change the tiering service levels to non-tiering for your volumes.
- Enable and disable backups for a volume on a tiering service level.
- Enable and disable disaster recovery for a volume on a tiering service level.

## **Non-returnable disk offering**

As a part of NetApp Keystone Flex Subscription (Flex Subscription), NetApp extends the non-returnable disk (NRD) offering.

If you purchase the NRD offering for Flex Subscription, NetApp does not recover the physical storage media used during the entire service tenure because of support and maintenance activities, 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:

- Even on purchasing this service, you can opt for NetApp to recover the physical storage media.
- In case NetApp is not responsible for recovering the media, you are entitled to destroy the storage media or disks used in the delivery of the Flex Subscription service at the end of the service tenure.
- You can add, modify, or terminate the NRD offering during the renewal of the subscription and not in the middle of the tenure.
- The cost associated with the NRD offering changes based on the committed capacity of the subscription. That is, if you opt to increase your committed capacity in the middle of the subscription period, the cost of NRD is revised likewise. The increase will be proportional to the increase in the committed capacity.
- You can retain only the physical storage media used in your service. Controllers, shelves, cables, switches, network cards, and any other equipment owned by NetApp will be recovered by NetApp.

## **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 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>

## Flex Subscription powered by Equinix

NetApp has partnered with Equinix for hosting NetApp Keystone Flex Subscription (Flex Subscription) in an Equinix data center to ensure the delivery of a unified solution for you.

Flex Subscription powered by Equinix is unchanged from the standard Flex Subscription offering.

In addition to the standard Flex Subscription offering, you will need to select an Equinix datacenter to host your Flex Subscription equipment.

# Operational model, roles, and responsibilities

NetApp Keystone Flex Subscription (Flex Subscription) is based on the model of tenancy.

Flex Subscription offers three operational models for service delivery.

- **NetApp-operated model** allows the customer to subscribe to the offered services (according to the selected performance tiers and storage service types) and selects the NetApp-operated option at an extra cost. NetApp defines the architecture and products, installs at the customer premises, and manages the day-to-day infrastructure management operations by using NetApp storage and IT resources. Available storage service types are file, block, and object. Storage subscriptions based on Cloud Volumes Service for GCP and AWS can also be managed through your NetApp Keystone instance.
- **Partner-operated model** is similar to the NetApp-operated model, but with the partner operating the service for their end customers. In this model, the partner is the referenced contracted party. Tenants are customers of partners or service providers and have no billing relationship with NetApp. A partner-operated model usually has a multi-tenant environment where tenants and end customers/subtenants have their own subscriptions that are billed by the service provider/partner. The partner admin performs the administrative tasks for all the tenants. The functions that an admin can perform in a partner-operated model are different from that of an admin in a NetApp-operated model.
- **Customer-operated model** allows the customer to subscribe to an offered service, according to the selected performance tiers and storage service types. NetApp defines the architecture and products and installs at the customer premises and allows customers to manage the infrastructure using their storage and IT resources. A customer can be tied to NetApp or a partner/service provider, and based on that, the service requests can be raised and addressed to NetApp or the service provider. A customer admin can perform the administrative tasks in a customer-operated environment. These tasks are tied to the tenants

and subtenants for the specific customer.

The features and options in the offering vary based on the models. For information, see [Roles and operations of service providers and customers](#)

## Roles and responsibilities across the service lifecycle

- **NetApp-operated model:** The end to end management of installation, deployment, operations, monitoring, optimization and support is performed by NetApp.
- **Partner-operated model:** The share of roles and responsibilities depends on the SLA between you and the service provider or partner. Contact your service provider for information.
- **Customer-operated model:** The following table summarizes the overall service lifecycle model and the roles and responsibilities associated with them in a customer-operated environment.

Task	NetApp	Customer
Installation and related tasks <ul style="list-style-type: none"> <li>• Install</li> <li>• Configure</li> <li>• Deploy</li> <li>• Onboard</li> </ul>	✓	None
Administration and monitoring <ul style="list-style-type: none"> <li>• Monitor</li> <li>• Report</li> <li>• Perform administrative tasks</li> <li>• Alert</li> </ul>	None	✓
Operations and optimization <ul style="list-style-type: none"> <li>• Manage capacity</li> <li>• Manage performance</li> <li>• Manage SLA</li> </ul>	None	✓
Support <ul style="list-style-type: none"> <li>• Support customer</li> <li>• Hardware break fix</li> <li>• Software support</li> <li>• Upgrades and patches</li> </ul>	✓	None

### Roles and responsibilities summary

The following list summarizes NetApp's roles and responsibilities:

- NetApp delivers, installs, configures, and enables the applicable service (including applicable version of NetApp Service Engine at a customer-designated data center or a CoLo). NetApp is responsible for the uninstallation at the end of the contract term or if the customer chooses to terminate the contract sooner.
- While interfacing with other IT service providers, NetApp cooperates with the customer's IT service providers or their technical team.
- Performance of the NetApp-operated services are included, with the assumption that shared responsibilities are applicable to the supported environment.

## Tenancy and multi-tenancy in Flex Subscription

NetApp Keystone Flex Subscription (Flex Subscription) uses the concepts of **Tenant** and **Subtenant** as hierarchical entities that own logical storage resources. The concept of multi-tenancy is also supported, where multiple tenants are tied to a partner or service provider. The entities partner and service provider are used interchangeably here.



In the context of Flex Subscription, a single tenancy is a NetApp-operated model, while a multi-tenancy is a partner-operated model.

### Tenants


A tenant is the primary billing entity defined in Flex Subscription. Each customer that is onboarded into Flex Subscription and holds a subscription to the services exists as a Flex Subscription tenant. The customer provides NetApp with a customer name and/or identifier for the purposes of onboarding tenants and tracking subscription dates and commitment terms.

In a multi-tenancy, a partner is a tenant/customer of NetApp in a Flex Subscription environment. The partner, in turn, can bear multiple tenants or customers, who are linked to end customers/subtenants. In this model, tenants/subtenants are the customers of the service providers, and have no billing relationship with NetApp.

### Subtenants

Flex Subscription subtenants exist entirely within a parent tenant with a many-to-one relationship. Subtenants provide logical resource separation within a tenancy and are used as a basis for show-back reporting. All storage resources (that is file services, block storage, and object storage) belong to a subtenant.





## Subscription model in a NetApp-operated (single tenancy) environment

- A subscription is created by NetApp admin or GSSC for a tenant and zone.
- The subscription names are derived from the billing entity for which the tenant is subscribed.
- In the subscription, the start-date and contract term are configured.
- There can be only one active subscription for each tenant-zone
- One 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.
- Service levels can include:
  - Extreme
  - Extreme-Tiering
  - Performance
  - Performance-Tiering
  - Value
  - Data protection for each of the service levels
  - Advanced Data Protection for Extreme, Performance, Value
  - Storage objects

## Subscription model in a multi-tenant environment

- Service Providers are Flex Subscription customers and have subscriptions as tenants. The subscriptions are based on:
  - Commitment per service level and zone
  - Charged on allocated capacity with 100TiB minimum
  - Burst charges apply for 100-120% of committed capacity
- NetApp charges the providers monthly, as a part of their usual tenancy terms.

- For a service level to be available to tenants, the service provider or partner should first have a Flex Subscription in place for the service level.
- The service provider creates tenant subscriptions per service level, zone, and flexible minimums.
- Service providers can sell more capacity to their tenants than they have purchased from NetApp (oversubscription). Therefore, the capacity used by tenants is not limited by the capacity the service provider has subscribed to.
- Tenants can use storage capacity over their subscribed amount, that is listed as 'burst' on usage reports.
- Tenant usage reports are available to partners for viewing on a daily or monthly basis.
- Tenants can create subscriptions for longer duration as compared to the corresponding Flex Subscription, but a warning message is displayed to the end customer during that activity.
- Flex Subscription for a partner is configured by NetApp admins or GSSC. Management of Flex Subscription and tenant subscriptions is performed by a user with Partner admin role.
- Users with the tenant admin roles can only view the tenant subscription (not partner's Flex Subscription). They can update the given subscription to change capacity and service level. They can raise service requests for additional subscriptions.
- The partner admin can create another subscription either when the existing subscription is expired, or for a future date when the existing subscription is no longer valid. The start date for a new subscription must be greater than or equal to current end date.

## Flex Subscription infrastructure

This section describes the NetApp Keystone Flex Subscription (Flex Subscription) infrastructure architecture and management application for the NetApp and customer-operated environments.

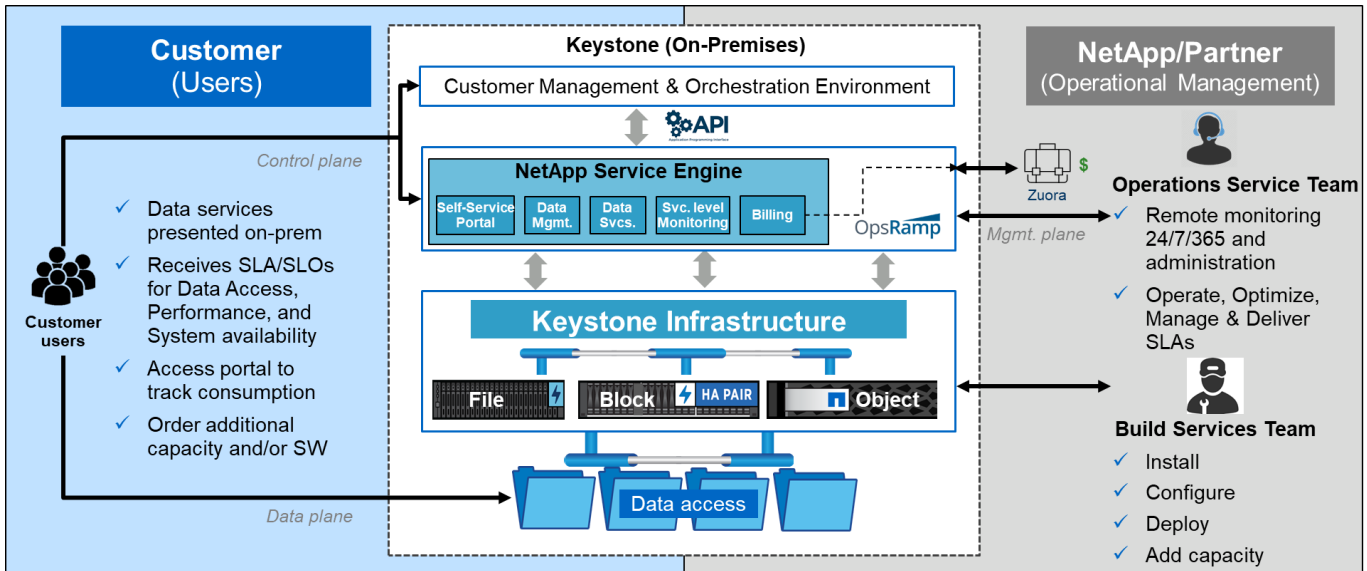
Flex Subscription infrastructure architecture, design, choice of technology, and component products reside solely with NetApp. NetApp reserves the rights to take the following actions:

- Select, substitute, or repurpose products.
- Refresh products with new technology when deemed appropriate.
- Increase or decrease capacity of the products to meet service requirements.
- Modify architecture, technology, and/or products to meet service requirements.

The Flex Subscription infrastructure includes multiple components:

- User interface (web portal) of NetApp Service Engine.
- NetApp Service Engine APIs for integration.
- The Flex Subscription infrastructure that includes storage controllers
- Tools to manage and operate the service such as OpsRamp, Active IQ, and Active IQ Unified Manager.

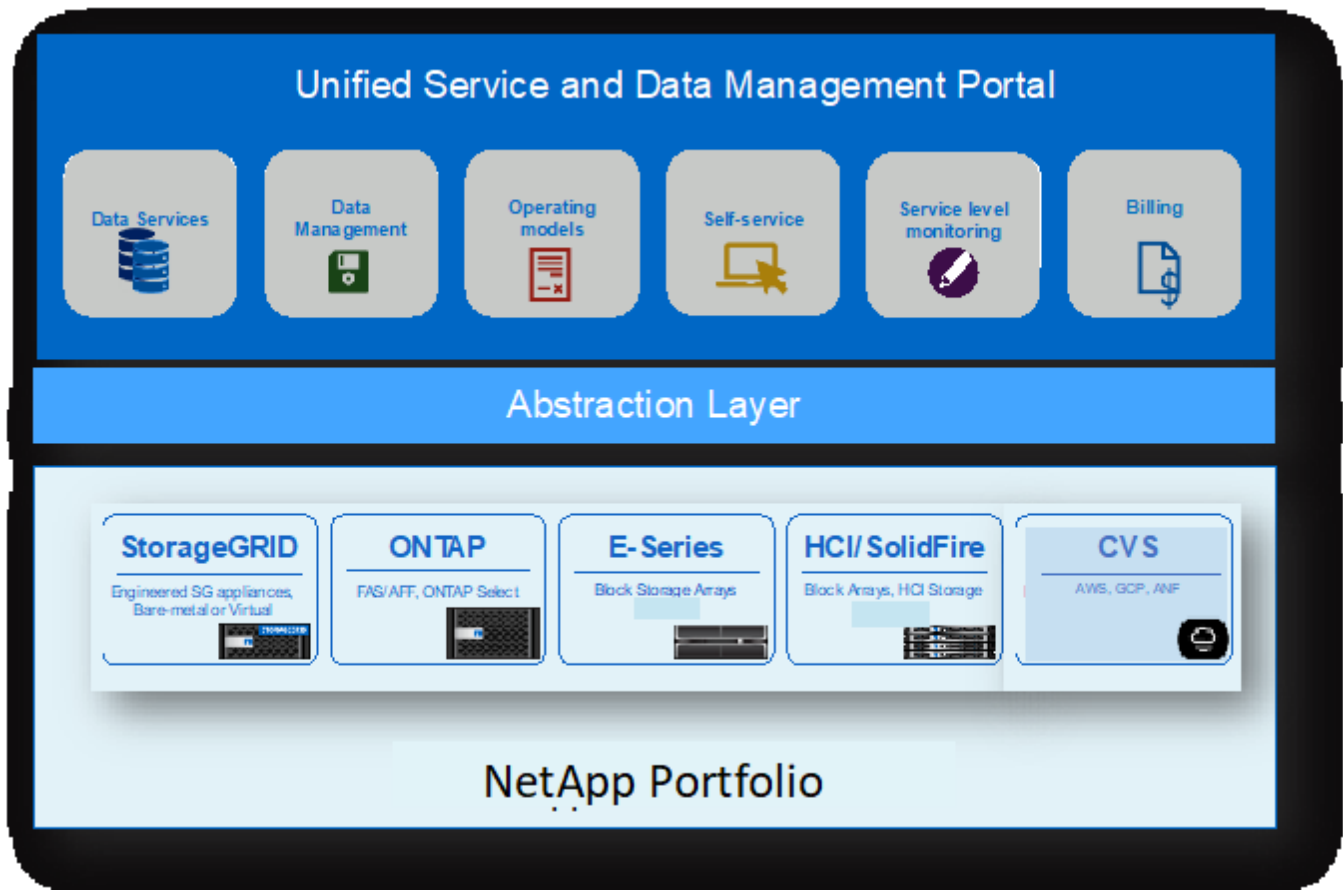
For more information about NetApp Keystone integration with Active IQ, see [Active IQ Digital Advisor Documentation](#)



## NetApp Service Engine

You can use the NetApp Service Engine web portal to manage and monitor your Flex Subscription service in a NetApp, partner, or customer-operated model. The portal consists of the following features:

- A graphical interface (NetApp Service Engine web interface) that supports monitoring and simple storage provisioning.
- A set of NetApp Service Engine REST APIs that allow more advanced setup and storage management actions.



The NetApp Service Engine portal is a single interface that allows customers to perform the following tasks:

- Subscription dashboard. View status of subscribed capacity and how much is being consumed.
- Storage provisioning. Workflows to create your NFS/CIFS file shares, FC/iSCSI disks, and S3 bucket.
- Data protection. Create snapshots and set backup policy or data replication on the provisioned file share or disk.
- Reports. View file shares and disks inventory and also trend on capacity usage against different service levels.
- Manage subscription. Order additional capacity to an existing subscription or a new service through the portal.
- Submit service requests or support issues and view their status

The complete functionality of NetApp Service Engine is available only in the NetApp-operated model. In a customer or partner-operated model, the NetApp Service Engine functionalities are limited to reporting and subscription management.

## Site requirements

There are multiple requirements for enabling NetApp Keystone Flex Subscription (Flex Subscription) services in your environment, such as space, racks, PDUs, power, and cooling, with additional network and security requirements discussed later.

## Space

Floor space to host the Flex Subscription infrastructure equipment (to be provided by customers). NetApp provides the weight specifications based on the final configuration.

## Racks

Four post racks in the customer-operated offering (to be provided by customers). In the NetApp-operated offering, either NetApp or the customer can provide the racks, depending on requirements. NetApp provides 42 deep racks.

## PDUs

You should be providing the power distribution units (PDUs), connected to two separate, protected circuits with sufficient C13 outlets. In the customer-operated offering, in some cases, C19 outlets are required. In the NetApp-operated offering, either NetApp or the customer can provide the PDUs, depending on requirements.

## Power

You should provide the required power. NetApp will provide the power requirement specifications based on 200V rating (Typical A, Max A, Typical W, Max W, Power cord type, and quantity), based on the final configuration. All components have redundant power supplies. NetApp will provide the in-cabinet power cords.

## Cooling

NetApp can provide the cooling requirement specifications (Typical BTU, Max BTU), based on the final configuration and requirement.

## Networking

Depending on customer requirements and the storage controllers used, NetApp can provide 10 Gb, 40 Gb, and 100 Gb connectivity at the customer site.

NetApp provides the required transceivers for NetApp-provided Flex Subscription infrastructure devices only. You should supply transceivers required for customer devices and cabling to the NetApp-provided Flex Subscription infrastructure devices.

# Flex Subscription Services Operations

NetApp Keystone Flex Subscription Services are run by leveraging the NetApp Global Services and Support Center (GSSC), NetApp Keystone Success Manager, and NetApp Keystone Service Delivery Manager.

## NetApp Global Services and Support Center (GSSC)

NetApp provides operational services remotely to NetApp Keystone customers. These services encompass a range of operational disciplines across storage management activities, including asset and configuration management, capacity and performance management, change management, event, incident and problem management, service request fulfillment, and reporting. NetApp will demonstrate a state of control and supporting evidence available as required.

## Additional Information and Support Contact

The NetApp Global Services and Support Center (GSSC) team primarily support the services to NetApp Keystone Flex Subscription customers.

Raise a ticket directly from the NetApp Service Engine portal (**Support > Service Requests**) with the appropriate details of the issue for assistance.

You can also use the following information to reach out to the support team.

- Global Service Contacts :  
<https://www.netapp.com/us/contact-us/support.aspx>
- If you have an open case/ticket that needs to be escalated, please send an email to one of the following addresses:  
[Keystone.services@netapp.com](mailto:Keystone.services@netapp.com)  
[Keystone.escalations@netapp.com](mailto:Keystone.escalations@netapp.com)
- NetApp uses OpsRamp, a cloud-based remote gateway solution to proactively monitor and connect to the NetApp Keystone environment for troubleshooting purposes. For information on OpsRamp, see <https://www.opsramp.com/#>.



In a partner-operated model, the tenant and subtenant's service requests are assigned to the partner's service desk. The partner's support tool might have integration with OpsRamp and GSSC applications. Only L3 issues are escalated to NetApp through GSSC.

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp Keystone  
<https://www.netapp.com/us/solutions/keystone/index.aspx>
- NetApp Product Documentation  
<https://docs.netapp.com>

## What can Flex Subscription customers view in Active IQ?

If you have subscribed to NetApp Keystone Flex Subscription (Flex Subscription) services, you can view your usage reports and graphs on the Active IQ Digital Advisor dashboard.

For more information about Active IQ Digital Advisor, see [Active IQ Digital Advisor documentation](#)

If your site has NetApp Service Engine installed by default, you can view the consumption charts through the NetApp Service Engine interface in addition to Active IQ Digital Advisor. However, in a customer-managed environment, you do not have access to the NetApp Service Engine interface, and you can view your capacity charts and usage reports only through the Active IQ Digital Advisor dashboard.

For information about customer-managed environments, see [Operational model, roles, and responsibilities](#)

The interval of metrics data collection from your site is 5 minutes.

As a customer, if you have subscribed to the Flex Subscription services, you can view the following information on Active IQ Digital Advisor:

- **Keystone Flex Subscription** widget: If you have logged in to Active IQ Digital Advisor, you can see the **Keystone Flex Subscription** widget that summarizes the capacity usage against your purchased Flex Subscription services.
- If you click **View Details** on the widget, you can see the **Keystone - Capacity Utilization** dashboard. The **Keystone - Capacity Utilization** dashboard provides further break down and analysis of the capacity and usage data as per the services subscribed, for example, the service levels, data protection, and advanced data protection services. The dashboard also displays the consumed, committed, and burst capacity usage graphs for the last 1, 7, and 30 days, with the mean accrual burst calculated and split over your current billing period.

For more information, see [View capacity utilization with NetApp Keystone Flex Subscription](#).

[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 of the geographical location at which the Keystone systems are deployed.

# Release Notes

## What's new in this release of NetApp Keystone Flex Subscription services

The release notes inform you of the new features and enhancements introduced in NetApp Keystone Flex Subscription (Flex Subscription) services, along with the known limitations and fixes.

### View your Cloud Volumes ONTAP usage and billing (05 December 2021)

If you are a NetApp Keystone Flex Subscription (Flex Subscription) customer, and also have Cloud Volumes ONTAP subscription, you can use Flex Subscription to charge your Cloud Volumes ONTAP usage and view your billing and consumption details. Cloud Manager can now discover your Flex Subscription accounts and display the details on the Cloud Manager interface. On clicking the **All Services > Digital Wallet > Keystone Flex Subscription** tab, you can perform the following activities:

- View a welcome screen from which you can request your Flex Subscription administrator to associate your Cloud Manager user ID with your Flex Subscription user ID; and thereby gain access to the **Keystone Flex Subscription** tab.
- View the consumption data of the committed and consumed capacities of the subscribed services on your on-premise systems and Cloud Volumes ONTAP.
- View the Flex Subscription details associated with your Cloud Manager account. Details, such as committed, consumed, and burst capacity are displayed.
- Link your Cloud Volumes ONTAP account with your Flex Subscription IDs to associate the account for billing and capacity usage reports.
- Request increasing or decreasing the committed capacity for the subscribed Flex Subscription services.
- View and use the linked Flex Subscription ID as a charging method for billing the Cloud Volumes ONTAP consumption (only linked IDs).

For more information, see the following links in the Cloud Manager documents:

[Keystone Flex Subscription](#)

[Manage Keystone Flex Subscriptions](#)

The following features have been introduced in different releases of NetApp Service Engine to support enhanced functionalities offered through NetApp Keystone Flex Subscription (Flex Subscription):

### Features introduced in NetApp Service Engine 2.2

The new features in this release include a revamped dashboard for new widgets on billing, capacity utilization, service requests, and alerts. This release also includes new screens for billing and alerts management, and renaming of the Subscribed Services menu to Cloud Services.



## Enhanced dashboard view

The NetApp Service Engine dashboard has been redesigned to include the following new components:

- **Capacity Utilization** For viewing the utilized capacity for your subscribed services.
- **Monthly Charges (Billing)** For viewing the aggregated monthly charges for all your subscriptions.
- **Alerts** For viewing the summary of the most recent alerts in your environment.
- **Service Requests** For viewing the list of the most recent service requests generated in your environment. For more information, see [View Flex Subscription Dashboard](#).

## New Billing screen

A new **Billing** screen has been added for easy accessibility and calculation of your historical billing data. The screen provides a holistic view of the monthly charges associated with all your subscribed services.

Navigate to this screen from the **Monthly Charges** widget on the dashboard or from **ADMINISTRATION > Billing** to view a monthly, subscription-level break up of the charges for all your subscribed services. The billing data is based on your committed and burst capacity usage, and is available for your usage in the previous months.

For more information, see [View billing](#).

## New Alerts screen

A new **Alerts** screen is introduced in this release that lists all system-generated and user-generated alerts. The screen also enables you to create custom alert messages for critical events concerning your environment and convey them to other users. They can view and dismiss the alerts, as required.

For more information, see [Create and manage alerts](#).

## Enhanced user interface

The following enhancements have been made on the NetApp Service Engine user interface:

- The **Dashboard** menu on the left navigation pane presents an intuitive navigation point to access the dashboard.
- The **Subscribed Services** menu on the left navigation pane has been renamed to **Cloud Services**.

## Features introduced in NetApp Service Engine 2.1

The new features in this release include supporting multi-tenancy in a Flex Subscription environment, and tiering capability that facilitates moving of inactive data to a lower cost local or cloud tier.

### Introducing Flex Subscription services for service providers

NetApp Service Engine now supports the management of a multi-tenant environment by a service provider. You can perform the functions of provisioning, reporting, billing, and managing customers having their own subscriptions. For supporting this feature, the following enhancements have been made:

- **Dashboard:** The Dashboard displays information on the storage subscriptions, such as service tiers, capacity usage for each service level, and add-on data protection services, for a specific subscription number. As a service provider, you can view the details of your NetApp Keystone Flex Subscription and

tenant subscriptions. As a tenant administrator, you can view the details of all the tenant subscriptions for your tenancy.

- **Reporting:** You can create capacity and performance reports with respect to your NetApp Keystone Flex Subscription usage and also for your tenant usage. As a partner administrator, you can view the capacity report for your Flex Subscription usage from **Reports > Keystone Usage**. As a partner admin, you can view the capacity usage reports for a specific tenant from **Reports > Tenant Usage/Capacity Usage**. As a tenant administrator, you can view the tenancy reports from **Reports > Tenant Usage**.
- **Subscription:** As a partner admin, you can view and update your Flex Subscription and tenant subscriptions from **SUBSCRIPTIONS > Keystone Subscriptions** and **SUBSCRIPTIONS > Tenant Subscriptions** respectively. As a tenant administrator, you can only view your tenant subscriptions.
- **Users:** Based on your role, you can assign privileges to a new or existing user within a tenancy as per the requirement. The role can be NetApp administrator, NetApp administrator with read only privileges, partner administrator, or tenant administrator. As a partner administrator, you can assign only partner administrator or tenant administrator roles to new users. A tenant administrator user can assign only the tenant administrator role to other users.
- **Networks menu:** As a partner administrator, you can view the networks defined for your tenancy. You can also create subnets for your subtenant and zone from **NETWORKS > Subnets**. This is required while provisioning storage by the end customers or subtenants.
- **API support:** The `/tenants/{tenant_id}/zones/{zone_name}/subnets` and `/tenants/{tenant_id}/zones/{zone_name}/subnets/{id}/tags` APIs are being offered as a part of this release to create and view subnets for subtenants.

For more information on this feature, see the following links:

- [Operational model, roles, and responsibilities](#)
- [Tenancy and multi-tenancy in Flex Subscription](#)
- [View Flex Subscription Dashboard](#)
- [View reports](#)
- [Managing subscriptions](#)
- [Managing tenants and subtenants](#)
- [Define networks for tenants and subtenants](#)

## Tiering

NetApp Keystone Flex Subscription service now includes a tiering capability that leverages the NetApp FabricPool technology. It identifies less frequently used data and tiers it to a cold storage that is owned, deployed, and managed on-premises by NetApp. You can opt for tiering by subscribing to the extreme-tiering or premium-tiering performance levels.

The following APIs have been modified to include new attribute values for the new tiering service levels:

- File services APIs
- Block store APIs

For more information, see the following links:

- [Tiering](#)
- [Performance Service Levels](#)

## Features introduced in NetApp Service Engine 2.0.1

The new features in this release include the following:

### Support extended to Cloud Volumes Services for Google Cloud Platform

NetApp Service Engine now has the ability to support Cloud Volumes Services for Google Cloud Platform (GCP) in addition to its existing support for Azure NetApp Files. You can now manage subscribed services, and provision and modify Google Cloud Volumes from NetApp Service Engine.



Subscriptions to Cloud Volumes Services are managed outside of NetApp Service Engine. The relevant credentials are provided to NetApp Service Engine to allow connection to the cloud services.

### Ability to manage objects provisioned outside of NetApp Service Engine

The volumes (disks and file shares) that already exist in the customer environment and belong to the storage VMs configured in NetApp Service Engine, can now be viewed and managed as a part of your NetApp Keystone Flex Subscription (Flex Subscription). The volumes provisioned outside of the NetApp Service Engine are now listed on the **Shares** and **Disks** pages with appropriate status codes. A background process runs at a periodic interval and imports the foreign workloads within your NetApp Service Engine instance.

The imported disks and file shares may not be in the same standard as the existing disks and file shares on NetApp Service Engine. After import, these disks and file shares are categorized with `Non-Standard` status. You can raise a service request from **Support > Service Request > New Service Request** for them to be standardized and managed through the NetApp Service Engine portal.

### SnapCenter integration with NetApp Service Engine

As a part of SnapCenter integration with NetApp Service Engine, you can now clone your disks and file shares from the Snapshots created in your SnapCenter environment, outside of your NetApp Service Engine instance. While cloning a file share or disk from an existing Snapshot on the NetApp Service Engine portal, these Snapshots are listed for your selection. An acquisition process runs in the background at a periodic interval to import the Snapshots within your NetApp Service Engine instance.

### New screen for maintaining backups

The new **Backup** screen enables you to view and manage the backups of the disks and file shares created in your environment. You can edit the backup policies, break the backup relationship with the source volume, and also delete the backup volume with all its recovery points. This feature allows the backups to be retained (as orphan backups) even when the source volumes are deleted, for later restoration. For restoring a file share or disk from a specific recovery point, you can raise a service request from **Support > Service Request > New Service Request**.

### Provision for restricting user access on CIFS shares

You can now specify the Access Control List (ACL) for restricting user access on a CIFS (SMB) or multi-protocol share. You can specify Windows users or groups based on the Active Directory (AD) settings to add to the ACL.

[Learn more.](#)

## Features introduced in NetApp Service Engine 2.0

The new features in this release include the following:

### MetroCluster support

NetApp Service Engine supports sites configured with MetroCluster configurations. MetroCluster is a data protection feature of ONTAP that provides recovery point objectives (RPO) 0 or recovery time objectives (RTO) 0 using synchronous mirror for continuously available storage.

MetroCluster support translates to a synchronous disaster recovery feature within NetApp Service Engine. Each side of an MetroCluster instance is registered as a separate zone, each with its own subscription that includes a Data Protection Advanced rate plan.

Shares or disks created in a MetroCluster-enabled zone synchronously replicate to the second zone. The consumption of the replicated zone follows the Data Protection Advanced rate plan applicable to the zone where storage is provisioned.

### Cloud Volumes Services support

NetApp Service Engine now has the ability to support Cloud Volumes Services. It can now support Azure NetApp Files.



Subscriptions to Cloud Volumes Services are managed outside of NetApp Service Engine. The relevant credentials are provided to NetApp Service Engine to allow connection to the cloud services.

NetApp Service Engine supports:

- Provisioning or modifying the Cloud Volumes Services volumes (including the ability to take snapshots)
- Backing up data to a Cloud Volumes Services zone
- Viewing Cloud Volumes Services volumes in NSE inventory
- Viewing Cloud Volumes Services usage.

### Host groups

NetApp Service Engine supports the use of host groups. A host group is a group of FC protocol host worldwide port names (WWPNs) or iSCSI host node names (IQNs). You can define host groups and map them to disks to control which initiators have access to the disks.

Host groups replace the need to specify individual initiators for every disk and allow for the following:

- An additional disk to be presented to the same set of initiators
- Updating the set of initiators across multiple disks

### Burst usage and notifications

Some NetApp Service Engine-supported storage subscriptions allow customers to use a burst capacity over their committed capacity, which is charged separately over and above the subscribed committed capacity. It is important for users to understand when they are about to use or have used burst capacity to control their usage and costs.

#### Notification when a proposed change results in using burst capacity

A notification to display a change in the proposed provisioning that will cause a subscription to go into burst.

The user can choose to continue, knowing that will put the subscription into burst or choose not to continue with the action.

[Learn more.](#)

#### **Notification when subscription is in burst**

A notification banner is displayed when a subscription is in burst.

[Learn more.](#)

#### **Capacity report shows burst usage**

Capacity report showing the number of days the subscription has been in burst and the quantity of burst capacity used.

[Learn more.](#)

#### **Performance Report**

A new Performance Report in the NetApp Service Engine web interface displays information about the performance of individual disks or shares on the following performance measures:

- IOPS/TiB (Input/Output operations per second per terabyte): The rate at which input and output operations per second (IOPS) occur on the storage device.
- Throughput in MBps: The data transfer rate to and from the storage media in megabytes per second.
- Latency (ms): The average time for reads and writes from the disk or share in milliseconds.

#### **Subscription management**

Subscription management has been enhanced. You can now:

- Request a data protection add-on, or request additional capacity for a data protection add-on for a subscription or service
- View data protection usage capacity

#### **Billing enhancement**

Billing now supports the ability to measure and bill for snapshot usage for ONTAP (file and block) storage.

#### **Hidden CIFS shares**

NetApp Service Engine supports creating hidden CIFS shares.

## **Fixed issues in NetApp Service Engine**

The following issues that were found in a previous release of NetApp Service Engine have been fixed for you to successfully use your NetApp Keystone Flex Subscription services.

Issue Description	After the fix	Fixed in version
Volume move was automatically triggered when non-FabricPool aggregate existed on the cluster.  Any modifications to volumes or disks triggered a volume move to another aggregate.	No volume move is triggered for volume operations.	NetApp Service Engine 2.2
Host groups deletion removed host groups from the NetApp Service Engine user interface (UI), but not from the cluster.	Resolved.	NetApp Service Engine 2.2
Host groups could be unmapped from the disks on the NetApp Service Engine UI, but not from the cluster.	Resolved.	NetApp Service Engine 2.2
Export policies could not be deleted from the NetApp Service Engine UI.	The changed policies can be saved from the UI.	NetApp Service Engine 2.2

## Known issues in NetApp Service Engine

The following known issues have been reported in NetApp Service Engine. You might encounter these issues when you provision or use your storage as a part of your Flex Subscription.

Known Issue	Description	Workaround
Limitations in synchronous data protection	There is an issue where the VLANs, IPspaces, and Broadcast domains are not defined on the secondary partner of an MetroCluster cluster. This issue can affect recovery of data from the replica zone.	Place a service request for GSSC to perform a manual network configuration on the partner cluster. The network configuration can be done in advance if the network components (VLAN, IPspace, and Broadcast domains) are known.
Limitations in disabling and deleting volumes that are disaster recovery enabled	If there are two or more volumes in a storage VM that are disaster recovery enabled, disaster recovery cannot be disabled for a file share or disk.	Raise a service request for GSSC to resolve the issue.
Limitations in deleting file servers and block stores that are disaster recovery enabled	Deleting a disaster recovery enabled block store or file share might fail because peering is not removed due to timeout.	Failed deletion is available in "Contact Support" and Support will be notified via OpsRamp Ticket.
Limitation when creating multiple snapshots	When multiple snapshots are triggered for creation, the web page initially displays only the last snapshot issued.	Auto job picks up the other snapshots within 15 minutes.

Known Issue	Description	Workaround
Disaster recovery enabled file servers on the secondary controllers are not deleted.	All primary volumes and the replication relationships are removed. However, the disaster recovery object is not deleted. This object does not consume any capacity or networking.	There is no workaround for this issue.
Limitation in disabling backup on volumes	Raise a service request for GSSC to resolve the issue.	There is no workaround for this issue.
Limitation in deleting network subnets	Customers are unable to delete network subnets.	Raise a service request for GSSC to resolve the issue.
Limitation in deleting VLAN using API	Deleting a VLAN that has a subnet associated with it leaves an empty subnet stub in the UI.	Raise a service request for GSSC to resolve the issue.
Limitation in importing storage VMs with non-standard subnet	Storage VMs (storage virtual machines, also known as SVMs), created outside of NetApp Service Engine cannot be operational due to a non-standard subnet. The storage VMs are imported with the status <code>Contacted Support</code> .	Raise a service request for GSSC to resolve the issue.
Capacity and billing reports might be inaccurate in NetApp-managed environments	In a NetApp-managed environment, the capacity and billing reports generated from the UI might be affected due to missing LUN size metrics and service levels on backup volumes.	Raise a service request for GSSC to resolve the issue.
A 15-minute lag in reporting the committed capacity on graphs after subscription creation.	When a tenant subscription is created, and the report graphs are checked immediately (within 15 minutes of the subscription creation), the committed capacity is not reported correctly.	Check the report graphs at least 15 minutes after the subscription is created.
Change in Performance Service Levels does not move volumes	If two different aggregates in the same cluster are used for two service levels, changing the service level on a volume to another does not move the volume to the other aggregate. Only the adaptive QoS policy is applied.	Raise a service request for GSSC to resolve the issue.
Creation of FCP SVM requires iSCSI activation	For creating a storage virtual machine (SVM) and enabling FCP support on it, an iSCSI interface activation is required.	Raise a service request for GSSC to resolve the issue.

<b>Known Issue</b>	<b>Description</b>	<b>Workaround</b>
Limitation with multiple host groups with same initiator	If multiple host groups are created with the same initiator name, and disks are mapped to all the host groups, the disks get mapped to only one host group on the cluster.	Raise a service request for GSSC to resolve the issue.



# NetApp Keystone frequently asked questions (FAQs)

The following questions give answers to your frequent queries about NetApp Keystone.

## What is NetApp Keystone?

NetApp Keystone is a portfolio of on-premises capital expenditure (capex) alternatives, consisting of NetApp Keystone Flex Pay (Flex Pay), and NetApp Keystone Flex Subscription (Flex Subscription).

- **Flex Pay:** A portfolio of payment solutions including traditional financing, leasing, and fixed/variable options to meet your customer's cash flow needs.
- **Flex Subscription:** Pay-as-you-grow subscription-based service that brings a cloud-like experience on premises, with an outcome-based service option (NetApp-operated).  
For more information, see [here](#).

## How does NetApp Keystone benefit my customers?

NetApp Keystone is the bridge that connects the pillars of our capex and hybrid cloud strategy— delivering agility, financial flexibility, and reduced financial risk that helps customers meet their cash flow and business needs.

The NetApp Keystone brand offers a portfolio of flexible payment solutions that include traditional financing, leasing, and fixed/variable options for cash-conscious customers along with Flex Subscription.

- **Flex Pay (financial flexibility, payment solutions):**
  - Prefer to own the asset title, but need payment/financial flexibility (rental model, lease, loan, installments)
  - Stringent security requirements with no external data connectivity options to enable subscription
  - Workloads have predictable capacity growth and aligns with capex budgets
  - Well-managed infrastructure with high asset utilization
  - Prefer to procure storage on a raw-capacity basis and retain efficiency benefits
- **Flex Subscription (OPEX, cloud-like experience):**
  - Prefer 100% OPEX, so assets will not end up on your customer's balance sheet
  - Internal/external Service Providers looking to align costs with usage/revenue
  - Workloads with unpredictable capacity growth
  - Reallocated IT resources from typical storage tasks (migrations, tech refresh, upgrades, and so on)
  - Short-term solution before migrating workloads to the cloud
  - Repatriating workloads back from the cloud to on- premises

## What is Flex Subscription?

Flex Subscription is a new, flexible, on-premises subscription-based procurement model. It enables customers to accelerate time to value by removing the hurdles around managing the resources and going through the lengthy procurement cycle. Flex Subscription allows customers to align economics to their business priorities. For more information, see [here](#).

## What does on-premises mean?

On-premises is defined as a customer-owned data center or customer-owned space in a colocation facility. The customer is responsible for the space, power, and cooling.

## What are the benefits of Flex Subscription services?

Some benefits of Flex Subscription services are:

- Frees up IT staff from complicated storage-related tasks and allows them to focus on application management
- Reduces upfront capital investment
- Allows customers to meet their demands without overprovisioning
- Aligns data storage costs with business needs/activity
- Simplifies infrastructure provisioning by bypassing complex organizational procurement procedures
- Keeps data secure on their premises
- Enables proper control over compliance, performance, and security  
For more information, see [here](#).

# NetApp Keystone Flex Subscription FAQ

The following questions give answers to your frequent queries about NetApp Keystone Flex Subscription services.

## What is offered as part of NetApp Keystone Flex Subscription?

NetApp Keystone Flex Subscription (Flex Subscription) is a subscription-based service offering for block, file, and object data services that can be deployed on-premises and can be operated by NetApp, a partner, or a customer.

## What storage service offers are provided as part of Flex Subscription?

For information, see [Performance Service Levels](#)

## What add-on services are supported?

Advanced data protection (backup and disaster recovery) and Hybrid Cloud Tiering with FabricPool are add-on services that can be chosen at an additional cost.

## What service levels does NetApp guarantee with the service?

In a NetApp operated scenario, NetApp Keystone guarantees IOPS/TiB for the storage that is provisioned and the latency for each service level.

## What does Flex Subscription map to?

Flex Subscription maps to a single site or a single data center and it can comprise of different performance service levels.

## What are the benefits of extreme-tiering and premium-tiering performance levels?

Tiering is enabled in the extreme-tiering and premium-tiering performance levels, which enables you to reduce your storage footprint and associated costs. NetApp assumes that 25% of your data is hot, while the remaining 75% is less frequently used or cold, and moves it to cold storage. Additionally, you can check usage reports to understand how frequently data is accessed and enable tiering service based on the information.

## Can partners sell more capacity than they have purchased from NetApp to customers?

Tenant subscriptions are not limited by the capacity that the partner has purchased. Partners can sell more capacity than they have purchased from NetApp to their customers. The capacity that is in excess of the purchased capacity is referred to as oversubscription.

### **Are APIs provided to integrate with customer tools?**

Yes, RESTful APIs are available to integrate into your own applications. Navigate to **SUPPORT > API Documentation** to see the API documentation for Flex Subscription services. For more information, see the [Keystone Flex Subscription API guide](#).

### **What is burst capacity?**

You can increase and decrease usage up to 20% above the committed capacity. The burst capacity usage is measured on a daily basis and billed only when used.

For example, if the committed capacity is 100 TiB, you can burst up to 120 TiB.

### **Is there a premium charge for using burst?**

The burst capacity usage up to 20% of committed capacity is billed at the same rate as committed capacity, any usage above 20% of committed capacity is billed at 50% premium.

### **What is the benefit of burst capacity?**

Burst capacity gives you the flexibility to consume storage on demand versus committing for the long term.

### **How is burst capacity allocated to tenants?**

Burst capacity is allocated to partners, who further allocate it to their customers based on requirements.

### **Where can I see the committed and burst capacity usage?**

NetApp Service Engine has built-in dashboards to report consumed capacity against committed capacity.

### **Will there be any notifications if I reach a certain percentage in committed capacity usage?**

Yes, the management tools provide notifications on capacity usage through the NetApp Service Engine dashboard.

### **How do I view the Flex Subscription usage?**

NetApp Service Engine provides a dashboard view, with information on all the services that are subscribed to and how much is consumed. For details about NetApp Service Engine, see [here](#).

### **How do I report any issues with the service?**

NetApp Keystone support can be reached through these various channels:

- Support email: [keystone.services@netapp.com](mailto:keystone.services@netapp.com)
- Escalations email: [keystone.escalations@netapp.com](mailto:keystone.escalations@netapp.com)

### **Can I order new storage service?**

Yes, new storage service or expansion to storage service can be requested from the NetApp Service Engine portal. The request is processed by the NetApp Keystone operations team before making it available for use.

### **Are increases to storage commitments available immediately?**

Depending the amount of capacity requested, a determination is made whether the capacity is already deployed, or it requires additional equipment to be shipped and installed.

### **Can workloads be moved between the tiers?**

Yes, workloads can be moved between tiers, provided the user has subscribed to the tier the workload is moving to. However, we do not recommend moving from a higher tier to a lower tier because it can cause a performance degradation. The process is achieved by simply editing the file share and changing the service level setting.

**What software version (for example, ONTAP) is installed as default?**

Depending on the service tier subscribed (for example, ONTAP with File and Block services, SANtricity for Block, and StorageGrid for Object) the support team installs the latest stable release with no security or feature issues.

## **Flex Subscription service offer details**

The following questions give answers to your frequent queries about NetApp Keystone Flex Subscription service offering.

**What is the minimum committed capacity?**

The minimum committed capacity for a NetApp Keystone Flex Subscription (Flex Subscription) is 100 TiB, per site across one or more service tiers.

**What is the typical length of term of a Flex Subscription agreement?**

Flex Subscription offers 12, 24, and 36-month term periods.

**How can I access the storage?**

In a NetApp-operated (standard) model, the storage controllers (ONTAP System Manager or administrative access to the systems) are owned and managed by NetApp. You can monitor and manage your storage only through NetApp Service Engine UI and APIs.

In a customer-managed (or Lite) deployment, where the NetApp Service Engine UI and APIs are used mainly for billing functions, you can access the storage controllers, such as Active IQ Unified Manager and ONTAP System Manager, and directly access ONTAP clusters.

**How do I manage the service?**

NetApp Service Engine is the orchestration and management tool that you can use to provision storage and get reports on the service usage.

**How can I increase the committed capacity in a subscription?**

You can submit a capacity addition request through the NetApp Service Engine management tool or through their NetApp Keystone success manager.

**Does increasing the capacity extend the term?**

All the additions are co-term to the existing term period, except if the request is made in the last 90 days of the subscription, in which case the term must extend for at least 12 months.

**Does a new subscription have flexibility to come with a new yearly term?**

Yes, new subscriptions can have new terms separate from any existing subscriptions.

**Can tenant subscription terms extend beyond the subscription term of the partner?**

Tenant subscriptions can extend beyond the current Flex Subscription term of the partner. A warning will be displayed when subscriptions are created and also in usage reports.

**Can I mix multiple subscriptions on the same ONTAP cluster?**

No, each cluster is assigned to a particular subscription.

### **Who does the monitoring and operations?**

For the NetApp-operated service, NetApp is responsible for monitoring the infrastructure remotely so that the service is delivered according to the agreed expectations.

For the partner-operated service, partner is responsible for monitoring the infrastructure remotely so that the service is delivered according to the agreed expectations.

For the customer-operated service, customer is responsible for monitoring the infrastructure and raise any issues to NetApp.

### **What happens if I terminate the service early?**

The minimum service commitment is 12 months. If you cancel the service early, the residual value needs to be paid upfront.

### **Is there an opportunity to convert into a purchase after the initial term is completed?**

No. The offer does not include an option to convert into a purchase. Additionally, previously purchased NetApp products are outside the scope of this program.

### **Can I request a particular software version?**

No. NetApp Keystone has standardized the software version across all its customer base.

### **Will I be informed about software upgrades?**

Yes. All maintenance/upgrade activities are communicated to and scheduled at a mutually agreeable time with the customer by the NetApp Keystone Success Manager.

## **Operational models and responsibilities**

There are three operational models in Flex Subscriptions. These FAQs are related to those operational models.

### **What are the different operating models and who is responsible for the major activities?**

The following chart is an overview of the three operating models that a customer can select: NetApp Operated, Partner Operated, and Customer Operated.

- **NetApp-operated model:** The end to end management of installation, deployment, operations, monitoring, optimization and support is performed by NetApp.
- **Partner-operated model:** The share of roles and responsibilities depends on the SLA between you and the service provider or partner. Contact your service provider for information.
- **Customer-operated model:** The following table summarizes the overall service lifecycle model and the roles and responsibilities associated with them in a customer-operated environment.

Task	NetApp	Customer
Installation and related tasks <ul style="list-style-type: none"> <li>• Install</li> <li>• Configure</li> <li>• Deploy</li> <li>• Onboard</li> </ul>	✓	None
Administration and monitoring <ul style="list-style-type: none"> <li>• Monitor</li> <li>• Report</li> <li>• Perform administrative tasks</li> <li>• Alert</li> </ul>	None	✓
Operations and optimization <ul style="list-style-type: none"> <li>• Manage capacity</li> <li>• Manage performance</li> <li>• Manage SLA</li> </ul>	None	✓
Support <ul style="list-style-type: none"> <li>• Support customer</li> <li>• Hardware break fix</li> <li>• Software support</li> <li>• Upgrades and patches</li> </ul>	✓	None

### What is a NetApp-operated model?

This operating model allows the customer to subscribe to the offered services, according to the selected performance tiers and storage service types and selects the NetApp operated option at an extra cost. NetApp defines the architecture and products, installs at the customer premises, and NetApp manages the day-to-day infrastructure management operations using our storage and IT resources. Available storage service types are file, block (iSCSI), and object. Cloud Volumes Service for GCP and AWS are also supported.

NetApp also creates and manages the partners, tenants, as applicable, and manages the subscriptions.

### What is a partner-operated model?

The operating model for the partner or service provider is similar to the NetApp-operated model, but with the partner operating the service for their end customer. In this model, the partner is the referenced contracted party. Tenants are customers of the service providers and have no billing relationship with NetApp. They manage their tenancy and customers. The tenants support requests are first triaged by the service provider before being escalated to NetApp.

### What is a customer-operated model?

This operating model allows the customer to subscribe to an offered service, according to the selected performance tiers and storage service types. NetApp defines the architecture and products and installs at the customer premises and allows customers to manage the infrastructure using their storage and IT resources. Available storage service types are file, block (iSCSI), and object. In this model, the customer referenced the contracted party, and this can be an end-user or partner.

### **Who owns the equipment?**

In all three operating models, NetApp owns the title to all the hardware and software installed at the customer premises.

## **NetApp Service Engine/Self-service access portal**

The following FAQ applies to NetApp Service Engine.

### **What is NetApp Service Engine?**

NetApp Service Engine is a self-service portal that is available in the NetApp-operated model for you to log into and provision storage based on your NetApp Keystone Flex Subscription service. The tool also provides reports on what their consumption levels are against their subscription and initiate any service requests or subscription changes.

### **NetApp Service Engine required in the customer-operated model?**

In a customer-operated model, NetApp Service Engine is required. It provides basic reports on the consumption details and is required to collect and report the billing information.

### **Where is NetApp Service Engine installed?**

In a NetApp-operated model, NetApp Service Engine is installed locally on the NetApp provided compute resources. In a customer-operated model, NetApp Service Engine is installed on the customer-provided compute resources.

### **Who can log into NetApp Service Engine?**

Users can be authenticated against NetApp SSO or local users configured in NetApp Service Engine.

### **How is access controlled?**

NetApp Service Engine provides role-based access control (RBAC), and each user can be associated to a role, which defines what actions they can perform. The RBAC assignment is done by the customer using NetApp Service Engine.

### **What access controls are available?**

The following access controls are available:

- **Partner Admin.** An administrator from the partner side who has rights to create and manage tenants, manage subscriptions for tenants, view usage reports, and manage technical aspects of storage.
- **Customer/tenant Admin.** An administrative person from the customer side who has rights to request changes to subscription, create new users and subtenants, and also create and view file shares, disks, and buckets.
- **NetApp Admin (read).** A NetApp administrator, who has access to all components of the web portal, and can view all details, such as all NetApp Keystone Flex Subscription, partners and tenants subscriptions,

and storage. However, this user has no permissions to create, edit, or delete.

- **NetApp Admin.** This user has full access and permissions to perform all functions in NetApp Service Engine web portal for managing NetApp Keystone Flex Subscription, and all activities for partner and tenant administration.

### **Where can I see the billing against my usage?**

You can view the aggregated monthly charges for your subscriptions for the last three months in the **Monthly Charges** widget on the Dashboard. You can view detailed billing information by clicking through the widget or navigating to **ADMINISTRATION > Billing**.



# NetApp Service Engine web interface

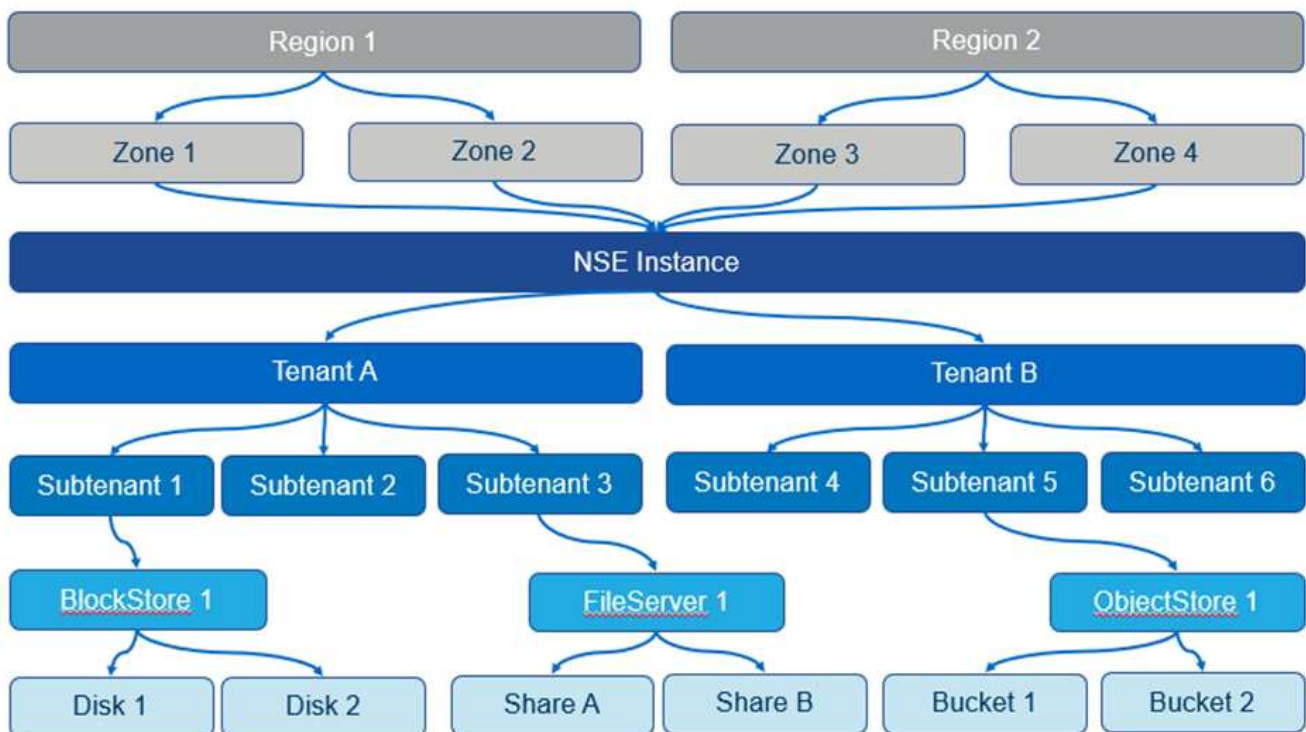
The NetApp Service Engine web portal allows you to manage and monitor a NetApp Keystone Flex Subscription (Flex Subscription) service. The portal consists of the following components:

- A graphical user interface (NetApp Service Engine web interface) that supports monitoring and simple storage provisioning. The menu options on the user interface (UI) are available based on the services and subscriptions of the tenant and subtenant.
- A set of REST APIs that allow more advanced setup and storage management actions. This guide describes how to use the NetApp Service Engine web interface. Functionality that is only available through the REST APIs is not covered in this document.

This section describes the NetApp Service Engine concepts and elements.

## Key NetApp Service Engine concepts

NetApp Service Engine supports the concepts of regions and zones. A region represents a data center or site, while a zone represents a storage subunit within the region (technically a cluster within a data center or site). Multiple zones support data availability and data protection (DP) features.



A single NetApp Service Engine instance can support one or more tenants. NetApp Service Engine uses the concepts of tenant and subtenant as hierarchical entities that own or manage the storage service.

A tenant can be a customer, partner, or a billing entity. A tenant holds the subscription (or multiple subscriptions) to the Flex Subscription service.

A subtenant is an entity wholly within the tenant. It can be used for show back, security separation, and so on.

Storage items are created per subtenant. Each subtenant can hold multiple storage items of the type and number suitable for that subtenant.

NetApp Service Engine supports the following types of storage:

- Block storage in block stores and disks
- File storage in file servers and file shares
- Object storage in buckets
- Cloud Volume Services for Google Cloud Platform and Azure NetApp Files

## Feature availability based on subscriptions

Based on your subscription, tenancy, and role in the NetApp Service Engine operations model, the features and options (screens and tabs) become available for you. For information, see [Service providers and customers and Operational model, roles, and responsibilities](#).

# Billing accounts, subscriptions, services, and performance

A subscription storage service is billed to a billing account. Each billing account is linked to a tenant. A billing account can be billed for one or more subscriptions.

A subscription refers to a group of storage services that are subscribed to and billed as one package. A subscription:

- Has one or more storage services
- Has a committed duration with a subscription end date
- Can have add-ons associated with the subscription



If storage is required in multiple data centers, a separate subscription is required for each data center, with separate commitments.

A storage service is a subscribed-to committed storage capacity with an associated performance level. NetApp Service Engine offers file and block storage at extreme, extreme-tiering, premium, premium-tiering, and standard performance levels, and object storage at an object performance service level.

The extreme-tiering and premium-tiering performance levels enable you to reduce your storage footprint and associated costs by monitoring and tiering your cold or inactive data to low-cost object storage tiers. The tiering policy is set to auto, where data left inactive for 31 days is tiered by default. You can modify this time period to be any number of days from 3 to 61 days.

When created, file and block storage items are associated with a performance level. Moving workloads between performance service levels is possible as requirements change. The extreme, extreme-tiering, premium, premium-tiering, and standard performance levels offer different levels of I/O performance (IOPS) and throughput (MBps) so storage can be tailored to requirements.

Burst usage is allowed on services up to a certain point; it is monitored and billed at separate charge rates (as defined in the subscription). For more information about capacity and usage, see [Committed, consumed, and burst capacity, and excess usage](#). DP services that support backups and disaster recovery are also offered.

## Committed, consumed, and burst capacity, and excess usage

Consumed capacity is how much capacity has been allocated (but not necessarily used). Committed capacity is the capacity that is committed to in a subscription; the subscriber is billed at a fixed rate for the committed capacity, regardless of how much is used.

Burst capacity is the capacity above the committed capacity that is allocated.

Burst capacity = consumed capacity – committed capacity

NetApp Service Engine monitors consumed capacity, checks the usage against the subscription, and bills for any consumed capacity that exceeds the committed capacity at the burst rate specified in the subscription. Usage is captured in five-minute increments and a daily summary is submitted to the billing engine for burst charge calculation. (Billing time is based on the local time for the underlying infrastructure for the NetApp Service Engine installation. )



In addition to primary storage, features such as snapshots, backup, and disaster recovery replicas contribute to and are included in usage calculations.

### Burst usage notifications

As burst usage incurs additional cost, the NetApp Service Engine GUI displays:

- A notification when a proposed change in provisioning will result in using burst capacity.
- A notification to a Customer Admin when a subscription has gone in to burst usage.
- How many days and in what quantity burst usage has been used for a service, in the Capacity Report. For more information see [Capacity usage](#).

### Notification when a proposed change will result in using burst capacity

This figure shows an example of a notification displayed when a proposed provisioning change will cause a subscription to go into burst. You can choose to continue knowing that it will put the subscription into burst or choose not to continue with the action.

#### Capacity



The following table lists when such burst notifications are displayed.

Action	Impact at the Source	Impact at the Destination
Create or resize a share/disk.	Exceeds commitment on the service level in the zone.	n/a
Move a share/disk to a new service level.	Exceeds commitment on the service level in the zone.	n/a

Action	Impact at the Source	Impact at the Destination
Create or resize a share/disk on a file server/block store with disaster recovery enabled.	<ul style="list-style-type: none"> <li>Exceeds commitment on the service level in the zone</li> <li>Exceeds commitment on DP in the zone</li> </ul>	Exceeds commitment on the service level in the destination zone by the automatically created destination share/disk.
Move a share/disk to a new service level on a file server/block store with disaster recovery enabled.	<ul style="list-style-type: none"> <li>Exceeds commitment on the service level in the zone.</li> <li>Exceeds commitment on DP in the zone.</li> </ul>	Exceeds commitment on the service level in the destination zone by the relocated destination share/disk.
Enable backups on a share/disk.	Exceeds commitment for DP.	Exceeds commitment on the service level in the destination zone by the automatically created destination share/disk.
Create a new object store tenancy.	The commitment for object capacity could be exceeded.	n/a
Increase the quota on an object store tenancy	The commitment for object capacity could be exceeded.	n/a

### Notification when subscription is in burst

The following notification banner is displayed when a subscription is in burst. The notification is displayed to the customer administrator for the tenancy and is displayed until the notification is acknowledged.



## Data protection

DP refers to methods that support back up of data and the ability to recover it if required.

NetApp Service Engine DP features include:

- Snapshots of disks and shares
- Backups of disks and shares (requires DP service as part of the subscription)
- Disaster recovery for disks and shares (requires DP or DP Advanced service as part of the subscription)

### Snapshots

Snapshots are point-in-time copies of data. Snapshots can be cloned to form a new disk or share with the same or similar features.

Snapshots can be created adhoc or automatically on a schedule as defined in a snapshot policy. The snapshot policy determines when snapshots are captured and how long they are retained.



Snapshots contribute to the consumed capacity of a service.

## Backups

Backup refers to taking a copy of an item, replicating it, and storing the copy in a zone other than the original zone, which has the respective protocol enabled (in case of block storage only) and is non-MetroCluster enabled. NetApp Service Engine offers backups on file and block storage (requires a DP service on the subscription). Backups of shares/disks are stored in the backup zone on the lowest cost performance tier (that is Standard) on subscription.

Backups can be configured at the time of creation of a new share/disk or later added to an existing share/disk.

### Notes:

- Backups occur at a fixed time, around 0:00 UTC.
- Backups occur as defined by the backup policy set for the share/disk. The backup policy determines:
  - If backups are enabled
  - The zone to which the backups are replicated; a backup zone is any zone in NetApp Service Engine other than the zone in which the original share or disk resides, which has the respective protocol enabled (in case of block storage only) and is non-MetroCluster enabled. Once set, the backup zone cannot be changed.
  - The number of backups to keep (retention) of each interval (daily, weekly, or monthly).

Scheduled backups are taken regularly and cannot be deleted but will be aged out as determined by the retention policy.

- Backup replication occurs daily.
- Backups of disks or shares cannot be configured in an NetApp Service Engine instance that contains only one zone.
- Deleting a primary share or disk will delete all associated backups.
- Backups contribute to the total consumed capacity. In addition, backups incur cost at the DP subscription rate. See also [Data Protection, Consumed Capacity, and Charges](#).
- Restore from backup: raise a service request to restore a share or disk from backup.

## Disaster recovery

Disaster recovery refers to the ability to recover to normal operations in the event of a disaster.

NetApp Service Engine supports two forms of disaster recovery: Asynchronous and Synchronous.



Support for disaster recovery is dependent on the infrastructure supported by the NetApp Service Engine instance.

### Disaster recovery—asynchronous

NetApp Service Engine supports asynchronous disaster recovery by providing the ability to:

- Asynchronously replicate primary volumes to a disaster recovery zone
- Failover/failback (available by service request only)

Asynchronous disaster recovery is available on file and block storage and requires a DP service on the subscription.

The disaster recovery zone must be a zone within NetApp Service Engine that is different to the zone in which the primary volume is created, and should not be a MetroCluster partner if the source zone is MetroCluster enabled. Disaster recovery replicas of shares/disks are stored in the disaster recovery zone at the same performance tier as the original share/disk.

Enabling asynchronous disaster recovery replication for a primary volume requires:

- Configuring the file server or block store on which the volume resides to support disaster recovery.
- Enabling or disabling disaster recovery replication of the file share or disk. By default, shares and disks are enabled for disaster recovery replication if disaster recovery is configured.

### **Configure file server or block store to support asynchronous disaster recovery**

Enable asynchronous disaster recovery on a file server or block store at creation or at a later date. After it is enabled, disaster recovery cannot be disabled, and the disaster recovery zone cannot be changed. The disaster recovery schedule specifies how often the data is replicated to the disaster recovery location (hourly, four hourly, or daily).

### **Enable asynchronous disaster recovery on file share or disk**

A file share or disk can only be configured for asynchronous disaster recovery replication if the parent file server or block store is first configured for asynchronous disaster recovery. By default, if replication is enabled in the parent, replication is enabled in the file shares or disks that the parent hosts. You can exclude replication of a particular share or disk by disabling disaster recovery on that share/disk. It is possible to toggle between enabling and disabling replication on these shares/disks.

### **Notes:**

- Deleting a primary file server or block store will delete all disaster recovery replicated copies.
- Only one disaster recovery zone can be configured per file server or block store.
- Disaster recovery copies contribute to the total consumed capacity. In addition, disaster recovery incurs cost at the disaster recovery subscription rate. See also [Data Protection, Consumed Capacity, and Charges](#).

### **Disaster recovery—synchronous**

MetroCluster is a DP feature which synchronously replicates data and configuration between two distinct zones which reside in separate locations or failure domains. In the event of a disaster at one site, an administrator can enable data to be served from the surviving site.

NetApp Service Engine managed sites that are configured with MetroCluster can support synchronous disaster recovery for File and Block storage in the following way.

- Zones can be configured to support synchronous disaster recovery.
- Disks/shares created in these zones synchronously replicate to the disaster recovery zone.

### **Notes:**

- Synchronous disaster recovery incurs costs at synchronous disaster recovery subscription rate. See also [Data Protection, Consumed Capacity, and Charges](#).


## Data protection, consumed capacity, and charges

The figures in this section describe how DP charges are calculated.

### Asynchronous disaster recovery


In asynchronous disaster recovery, usage and cost is made up of the following charges:

- Original volume capacity charged at the performance tier on which it resides.
- Disaster recovery copy charged at the same performance tier at the destination or disaster recovery zone (disaster recovery copies are stored at the same tier).
- DP service charge (for the capacity of the original volume).



### Synchronous disaster recovery

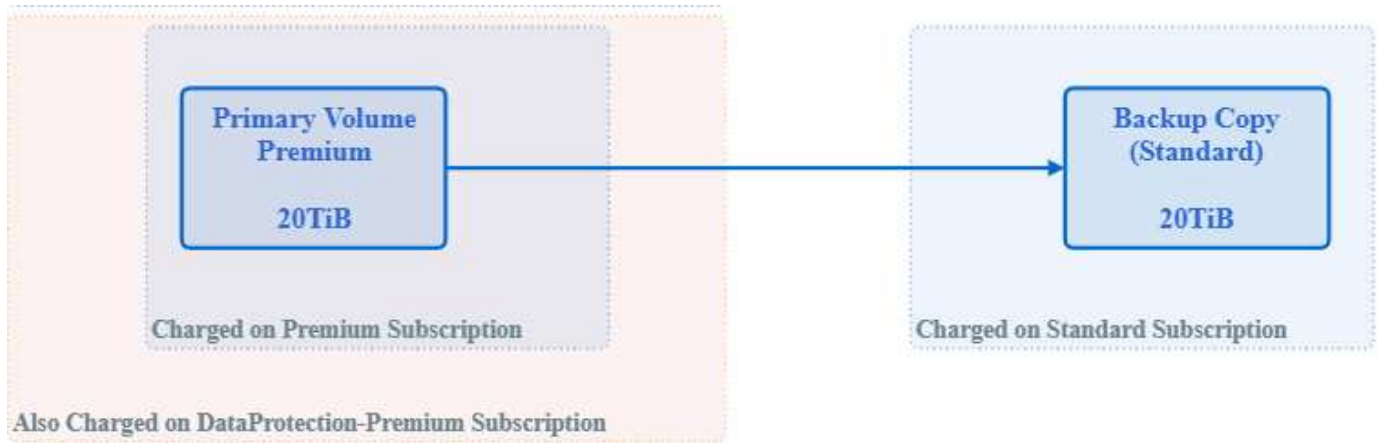
In synchronous disaster recovery, usage and cost is made up of the following charges:



### Backup

In backup, usage and cost are made up of the following charges:

- Original volume capacity charged at the performance tier on which it resides.
- Backup volumes charged at the lowest available performance tier (backup copies are stored on the lowest cost available tier).
- DP service charge (for the capacity of the original volume).



## Get started

### Overview

This section describes how to get started by using the NetApp Service Engine portal. It includes:

- Log in to the NetApp Service Engine web interface:
  - Log in with a user name and password
  - Log in with NetApp SSO
  - Log out of the NetApp Service Engine web interface
- Select tenant
- NetApp Service Engine web interface overview

### Log in to the NetApp Service Engine web interface

To use the NetApp Service Engine web interface, you must have an account. Your account is assigned one or more roles (by your NetApp Service Engine administrator) that determine your permissions and therefore which activities you can perform in the NetApp Service Engine portal.

NetApp Service Engine allows you to log in using the following credentials:

- User name and password (see [Log in with user name and password](#))
- NetApp SSO (see [Log in with NetApp SSO](#))

Confirm the sign-in option in use with your NetApp Service Engine instance with your NetApp Service Engine administrator.

### Role-based access

The following table lists the role-based access descriptions.



Role	Access
Partner admin/Account owner	Can perform all tasks of managing tenants and their subscriptions. For more information, see <a href="#">this link</a> .
Customer/tenant administrator	The customer or tenant administrator can perform all actions related with managing subtenants. For more information, see <a href="#">this link</a> .
NetApp admin (read)	Read-only access across all components.
NetApp admin	Full access and permissions to perform all functions in NetApp Service Engine web portal for managing NetApp Keystone Flex Subscription, and all activities for partner and tenant administration.

### Log in with user name and password

To log in with a user name and password, you need:

- Your NetApp Service Engine web interface user name and password
- The URL to the NetApp Service Engine portal
- A web browser

### Steps

1. In your web browser, go to the URL for your NetApp Service Engine portal. The log- in page is displayed.
2. Select **Local user sign in**.
3. On the Log in to NetApp Keystone page, enter your user name and password and click **Log In**.
4. After successful login, the NetApp Service Engine web interface loads, open at the dashboard. For an overview of the GUI, see [NetApp Service Engine Web Interface Overview](#).



If your login is successful but you cannot see the dashboard, check with your NetApp Service Engine portal administrator to ensure that your NetApp Service Engine portal account has been assigned the correct role.

### Log in with NetApp SSO

To log in with NetApp SSO, you need:

- A NetApp SSO account. You can request an account at the NetApp Support site; from the log-in screen by selecting **Create NetApp SSO account** and following the next steps.
- The URL to the NetApp Service Engine portal.
- A web browser.

### Steps

1. In your web browser, go to the URL for your NetApp Service Engine portal. The log- in page is displayed.
2. Select NetApp SSO.
3. On the SSO login page, enter your user name and password and click **Sign In**.

After the successful login the NetApp Service Engine web interface loads, the Flex Subscription dashboard is displayed. For an overview of the GUI, see [NetApp Service Engine Web Interface Overview](#).

## Log out of the NetApp Service Engine web interface

### Steps

1. To log out of the interface, click the user icon and click **Sign Out**.



## Roles and operations of service providers and customers

Based on your role of a service provider or partner (used interchangeably) or a customer or tenant (used interchangeably) administrator, the views and functions on the NetApp Service Engine web portal are determined.

### Activities that you can perform as a service provider administrator

As a service provider admin in a multi-tenant environment, you can perform the functions of provisioning, reporting, billing, and managing customers having their own subscriptions, through the NetApp Service Engine web UI. You can perform the following functions:

- View the storage subscribed and provisioned as a part of their Flex Subscription service.



As a service provide, you can sell more capacity to your customers than purchased from NetApp (oversubscription). Tenant subscriptions can extend beyond the current Flex Subscription period, however a warning is displayed at subscription creation time and in dashboards/reports.

- Create, view, modify, and delete zones, regions, tenants, subtenants, and users.
- Manage and host multiple tenants within the Flex Subscription service: Create and manage tenant subscriptions, and view your own Flex Subscription details.



The tenant subscriptions can be created based on the corresponding zones and Keystone Flex Subscriptions. They are created as per the service levels provided under the subscribed capacity of your Flex Subscription.

- Create file servers or block stores and initialize block, file, and object storage as you are aware of the network requirements for provisioning each storage type for your tenants.
- You can also create and manage subnets for your tenants,
- View file shares, disks, and buckets created as a part of your Flex Subscription. Also create and manage object storage groups and users.
- Create reports for consumption against your NetApp Keystone Flex Subscription, and tenants subscriptions, and usage for all tenants, for using it in the respective billing and rating systems.
- View the billing and reporting for all the tenants and customers and for only the storage subscribed to or provisioned.
- View and manage service requests raised by customers within your own (partner/service provider's)

service request tool (integrated with OpsRamp / GSSC).

### Activities that you can perform as a customer/tenant administrator

As a tenant administrator, you can be entitled to view the subscription details of your own tenancy and perform administrative tasks for your subtenants. You can perform the following actions:

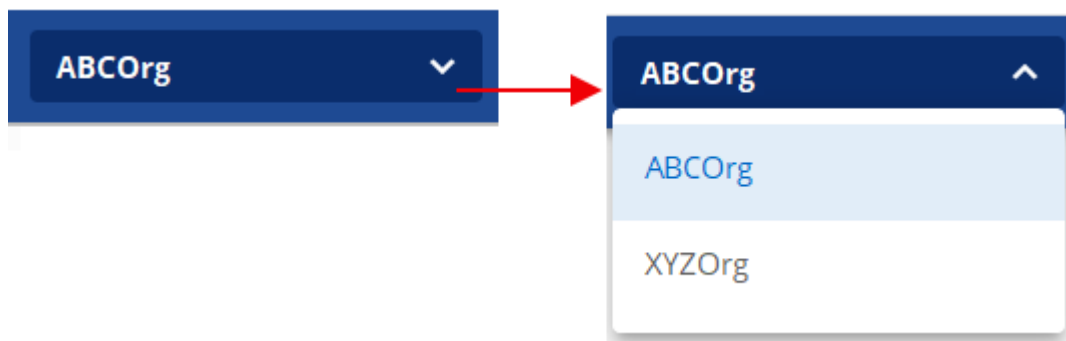
- View the storage subscribed and provisioned as a part of your Tenant Subscription Service.
- Manage and host multiple subtenants and users within the Tenant Subscription Service.
- View your zones, regions, tenancy, and subnets.
- View file servers and create file shares, disks, and buckets.
- Create reports for tenants and subtenant to view the consumption data for only the storage subscribed to or provisioned for your tenancy and subtenants.
- Raise and update service requests within the service provider's service request tool.



For additional subscriptions, you can make use of your service request tool.

### Select tenant

When working with the NetApp Service Engine web interface, any data that you see and any activities that you perform relate to the selected tenant. You can view the tenant at the top of the screen, as shown below.



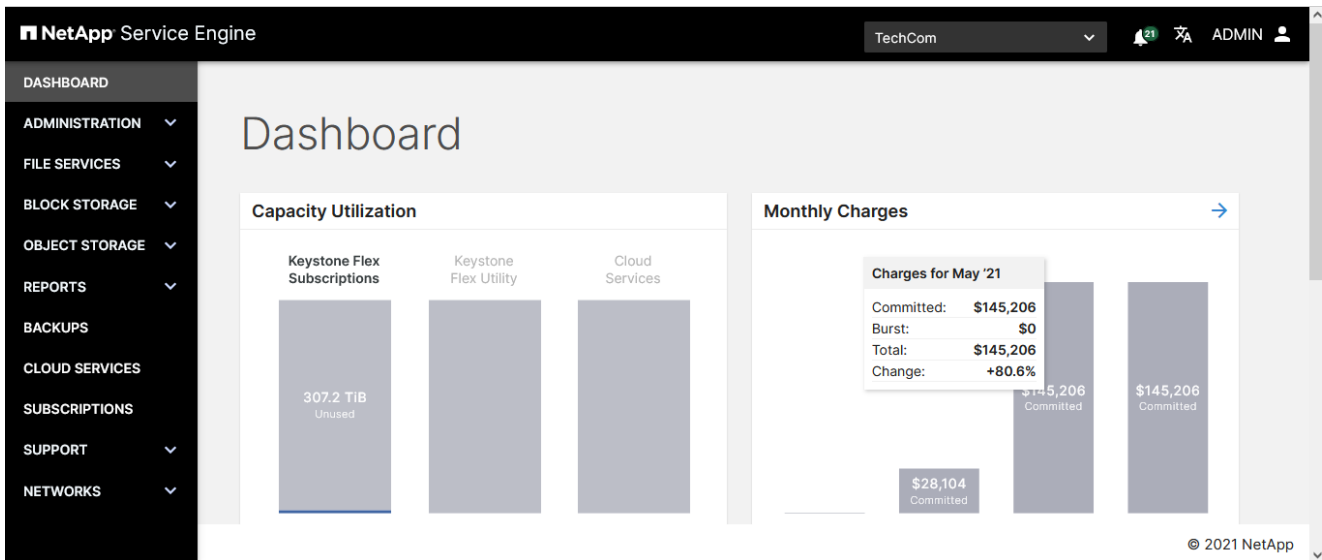
Most NetApp Service Engine web interface users have access to only one tenant. However, in instances where you might have access to multiple tenants, you can change the tenant by selecting another tenant in the Tenant field.

### NetApp Service Engine web interface overview

The figure below is an example of the NetApp Service Engine web Interface. It consists of the following components:

- **Display area.** This is the main working area of the screen; it has two views:
  - Dashboard view: Displays various tools to monitor the usage, billing, alerts, and service requests.
  - List view: Displays lists of items which can be navigated and sorted (more on that later).
- **Menu.** Use the menu to view and manage storage items, view reports, and access support.
- **NetApp logo.** Click this at any time to return the display to the dashboard view.

- **Tenant.** Displays the currently selected tenant. To change the tenant, see [Select the tenant](#).
- **Jobs.** Click to view the status of the most recent provisioning jobs. The icon changes color to display the configuration job status. For more information, see [Jobs and job status indicator](#).
- **User login.** Displays the name of the currently logged-in user. Click this icon to sign out. For more information, see [Log out of the NetApp Service Engine web interface](#).



The screenshot shows the Alerts and Service Requests section of the NetApp Service Engine Dashboard. The left sidebar is the same as in the previous screenshot. The main content area is divided into two panels: 'Alerts' and 'Service Requests'.

**Alerts:** This panel displays a table of alerts with columns for Status, Message, and Created. The alerts are as follows:

Status	Message	Created
Warning	premium-tiering service level of A-S00003875 has 1.71 TiB unused	about 1 hour ago
Warning	No subscription of dataprotect-premium-tiering is in burst by 300 Gib	about 1 hour ago
Warning	premium-tiering service level of A-S00003875 has 1.61 TiB unused	about 1 hour ago
Warning	premium-tiering service level of A-S00003875 has 1.9 TiB unused	about 15 hours ago
Warning	extreme service level of A-S00003875 has 9.66 TiB unused	about 15 hours ago
Warning	No subscription of premium-tiering is in burst by 150 Gib	3 days ago

**Service Requests:** This panel displays a table of service requests with columns for ID, Priority, Status, and Created. The requests are as follows:

ID	Priority	Status	Created
SRQ0045542704	Urgent	New	about 1 month ago
SRQ0045567639	Urgent	New	about 1 month ago
SRQ0045567849	Urgent	New	about 1 month ago
SRQ0045567964	Urgent	New	about 1 month ago
SRQ0045604969	Urgent	New	29 days ago
SRQ0045605054	Urgent	New	29 days ago
SRQ0045608534	Urgent	New	29 days ago
SRQ0045608974	Urgent	New	29 days ago
SRQ0045777564	Urgent	New	23 days ago
SRQ0045778884	Urgent	New	23 days ago

## Dashboard view

The Dashboard displays information about your NetApp Keystone storage subscriptions, such as the capacity used, billing against your usage, and recent alerts and service requests. It is the default view when you first log in to NetApp Service Engine. For more information on the Dashboard, see [View Flex Subscription Dashboard](#).

## List view

A List view is used to view a list of objects. For example, the list of servers that support the file shares in a subscription are displayed in a List view, as shown in the screenshot below.

From a List view, you can:

- Perform actions on the items in the list: see [List view actions](#).

### List view actions

A List view displays a list of items and provides a quick view of some of the item details (including state). In a List view, you can perform the actions listed in the following table.

Action	Description
Create an item	Use the Create button to create a new item.
Use Action icons	Use the Action icons to perform an action on the list item.
Sort the list	Use the arrow icons in the list column to sort the list by that column in ascending or descending order. The arrow icons are visible when you hover the cursor near the column name.
Change number of items displayed, navigate the list	Change the number of items displayed on the page and navigate the list using the Items per page field and the < and > icons at the bottom of the list.
Refresh the page	Refresh the page using the refresh icon:

### Object states

During provisioning and modification, the storage objects go through a series of states before they become operational. The state of the storage objects is displayed in the List view for those items. The objects may be in one of the following states:

- **Creating.** The storage resource is being created.
- **Updating.** The item is currently being modified.

Occurs when there is a change to storage resource (file server, files hare, blockstore, disk, and so on). It includes resizing shares, changing snapshot policy settings, changing export policy, taking a snapshot, renaming items, and so on.

- **Operational.** Denotes that the storage resource has been provisioned/modified correctly and is available, online, and functional.
- **Deleting.** Object is getting deleted and is being processed.
- **Queued.** Object is in queued state and is being processed.
- **Imported** Objects provisioned outside of the NetApp Service Engine are imported with this status when they do not fulfil the criteria of NetApp Keystone Flex Subscription (Flex Subscription).



This status typically implies that no QoS policy is applied on the object and therefore it cannot be managed by NetApp Service Engine as a part of your Flex Subscription. You can modify the object (through the edit pen) and assign an appropriate Service Level. This changes the status of that object to operational.

- **Non-Standard.** Objects provisioned outside of the NetApp Service Engine are imported with this status when they do not fulfil the criteria of NetApp Keystone Flex Subscription (Flex Subscription).



This status typically implies that the object misses one or more criteria to be managed by NetApp Service Engine, as a part of your Flex Subscription. You can [raise a service request](#) for them to be standardized and managed through the NetApp Service Engine portal and made operational by the support team.

- **Contact support.** This state occurs when the provisioning or modification task fails to fully complete. For items in this state, raise a Support Request to address the issue.
- **Operational (for the source object of a backup).** The following are the possible statuses:
  - With edit pen : Indicates that the backup object is operational and functioning.
  - Without edit pen: Indicates that the backup is orphaned, that is the SnapMirror relationship of the backup object with its source volume is broken, or the source volume has been removed.

### Jobs and job status indicator

Some provisioning tasks in NetApp Service Engine, such as create, modify or delete storage items, might take some time to complete. Rather than being executed immediately (as for synchronous tasks), these tasks are executed asynchronously. When initiating such tasks, NetApp Service Engine returns a job record. The status can be tracked through the bell icon on the top right corner that indicates whether the submitted task was completed successfully. The status of the job can be also be tracked through the APIs. For information, see [here](#)

Indicator color	Description
Black	A task is currently running.
Red	The last task failed to complete.
Green	The last task completed successfully.

Click on the status indicator to view the status of the 10 most recent tasks.

## View Flex Subscription Dashboard

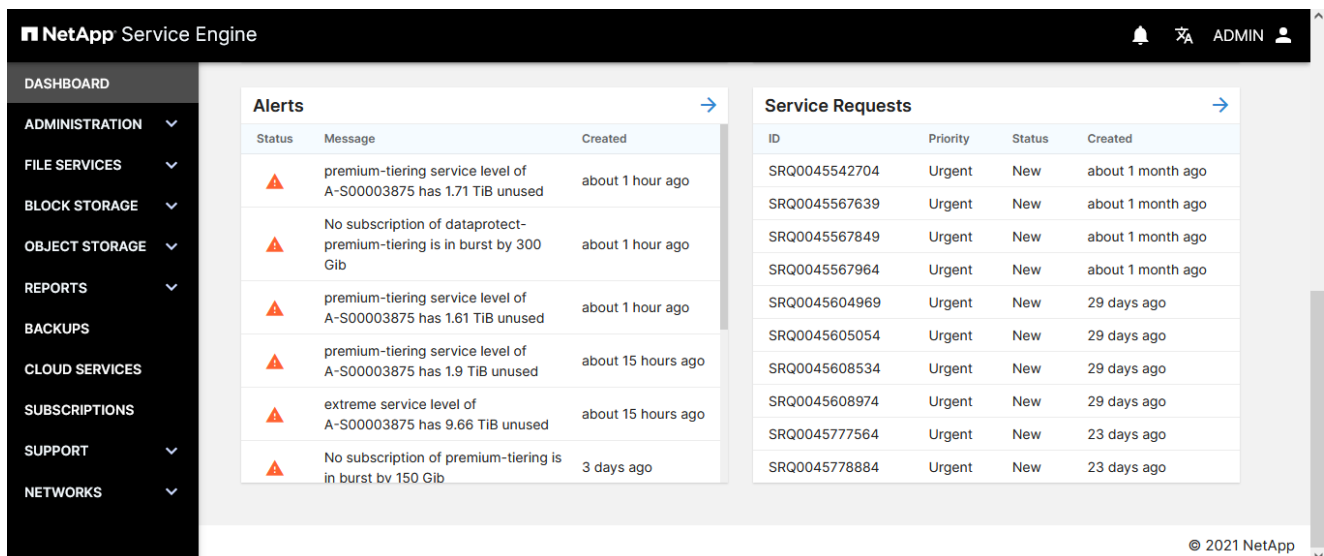
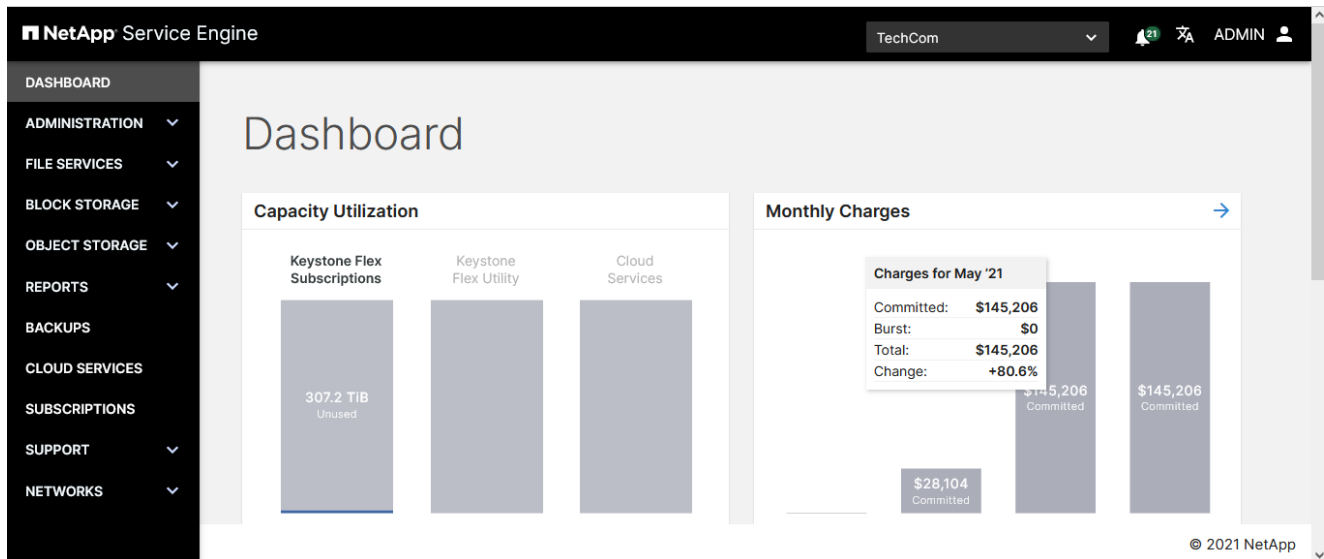
The Dashboard provides a quick overview of your NetApp Keystone storage subscriptions. It displays information about the capacity used, billing against your usage, and recent alerts and service requests.

The Dashboard displays important information across the following widgets, which can be clicked through for in-depth analysis:

- **Capacity Utilization:** Displays your capacity utilization for your subscribed services, for example, Keystone Flex Subscription, Keystone Flex Utility, and Cloud Services. Hovering your mouse over the bar charts displays the break up of the capacity utilized. You can view the committed and consumed capacities, as well as the capacity in burst and identify if you need to add more capacity to your subscriptions. Click the bar chart to view all the subscriptions for your subscribed service on a new pop-up. On clicking the subscription charts on this pop-up, you can see a break up of the capacity utilized per service level. You can also view the utilized capacity for your add-on services, such as data protection and advanced protection, if applicable.
- **Monthly Charges (Billing):** Displays the aggregated monthly charges for all your subscriptions. You can view the billing details for the last three months. On hovering your cursor over the chart, you can view the charges split up as per your committed and burst usage. If you have multiple subscriptions for a subscribed service, on clicking the bar chart for a month, the subscription list is displayed on the Billing screen. For

more information on billing, see [View billing](#).

- **Alerts:** Displays the summary of the most recent alerts, such as the status, message, and when it was created. For more information on alerts, see [Create and manage alerts](#).
- **Service Requests:** Displays the summary of the most recent service requests, sorted by priority and when they were created. For more information on service requests, see [Raise a service request](#).



## Dashboard view for a service provider in a multi-tenancy

If you are partner administrator in a multi-tenant environment, the dashboard provides a high-level overview of all your subscribed services for NetApp Keystone. The widget data of the capacity utilization and billing details, among others, represents your subscription and billing relationship with NetApp.

You can also view the capacity utilization by each tenant, by selecting a specific tenant from the top right corner. The data in the capacity utilization widget represents the data usage by the selected tenant, whereas the data in all the other widgets remain constant and represent your storage subscriptions.

## Dashboard view for a tenant in a multi-tenancy

If you are a tenant administrator in a multi-tenant environment, you can view the capacity utilization, recent alerts, and service requests for your tenancy, and click through the widgets for more details. Contact your service provider for your billing details.

## View billing

You can view the monthly charges for all the subscribed services on your NetApp Keystone Flex Subscription (Flex Subscription) on the NetApp Service Engine dashboard and Billing screen.

You can view the billing details for all your subscribed services for Flex Subscription for the previous months. The billing against the capacity in burst for each month is highlighted.



If any of your subscriptions are billed annually, the invoice reflects the prorated amount instead of actual monthly invoice amount.

You can sort the data in the table by month and view the list of all the storage subscriptions for any given month.

### Steps

1. Select **ADMINISTRATION > Billing** from the menu.



You can also click the arrow on the top right corner of the **Monthly Charges** widget on the dashboard or any bar that represents a month in the widget to view the Billing page. All your subscriptions for the selected month are displayed.

2. Optionally, hover over the bar representing a month to view additional details.
3. Select any month listed in the table.  
All your subscriptions for the month, along with the associated billing, is displayed in the table. The charges are split up as per the committed capacity and burst usage.

## Overview

This section describes how to manage your file servers and NFS/CIFS file shares. You can view information about your file servers and share, create, modify, and delete them.

### View servers

The Servers list displays the file servers belonging to the selected tenant. To view the list, select **File Services > Servers** from the menu.

The list displays simple information about each server such as:

- Server name
- IP address
- Subtenant



- Zone
- Operational state
- Protocols in use (NFS, CIFS)
- The CIFS server name (if relevant).

For more information on how to use the features of a list, see [List view](#).

## Create a file server

File servers belong to a subtenant and are created within a zone. When creating a server, you can optionally:

- Enable disaster recovery DP for the server. For more information about how disaster recovery works in NetApp Service Engine, see [Disaster recovery](#).
- Make it CIFS-enabled. For CIFS-enabled servers:
  - You must provide the Active Directory user name, Active Directory password, domain, DNS servers, server name and, optionally, the Active Directory Organizational Unit (OU).
  - The Active Directory credentials (Active Directory user name and Active Directory password) must be for a user that has the privilege to join a computer to the Active Directory domain.
  - When the Active Directory OU structure is hierarchical, as shown in the image below, specify the OUs from the lowest level to the top. In this example, to specify the Melbourne OU, set `cifs_ou` as `"cifs_ou": "ou=melbourne,ou=cifs"`.



### Before you begin

Make sure you have the following to create the server:

- The subtenant that will host the server.
- The region and zone in which the server belongs.
- Networking details such as the subnet and IP address (optional). If you are unfamiliar with your network, check with your IT department for the appropriate values.
- To enable asynchronous disaster recovery on the file server, the disaster recovery zone (the zone to which the file server will be replicated).

### Steps

1. View the [File Servers list](#).
2. Click **Create Server**.
3. On the Create Server page complete the following fields:

Field	Description
Subtenant	Select the Subtenant from the list.

Field	Description
Region	Select the region in which the server will reside.
Zone	Select the zone in which the server will reside.
Name	Enter the server name.
Subnet	Select any predefined subnet from the list.
IP address	(Optional) Specify an IP address. If not specified, the server will be given the next available IP address.

#### 4. Select the services:

NFS is enabled by default. The NFS protocol in use is displayed.

If creating a CIFS-enabled file server:

- a. Toggle the CIFS-enabled button to view the CIFS related fields.
- b. Complete the Active Directory Username, Active Directory Password, Domain, DNS Servers, Server Name and, optionally, the Active Directory Organizational Unit. The Active Directory credentials must be for a user that has the privilege to join a computer to the Active Directory domain.

#### 5. To enable asynchronous disaster recovery DP on this file server:

- a. Toggle the Asynchronous Disaster Recovery button to enable it.
- b. Select the disaster recovery region and zone.
- c. Select the disaster recovery replication schedule.

#### 6. If synchronous disaster recovery DP is enabled, the Synchronous Disaster Recovery toggle is enabled and cannot be disabled.

#### 7. Click **Create**. This creates a job to create the server.

### After you finish

Create server is run as an asynchronous job. You can:

- Check the status of the job in the jobs list.
- After the job is finished, check the status of the server in the Servers list.

### Modify file server

You can make the following changes to an existing server:

1. Change the server name
2. Make the server CIFS-enabled, and specify the Active Directory user name and password, Active Directory domain, DNS Server, Server name and optionally the Active Directory Organizational Unit. The Active Directory credentials must be for a user that has the privilege to join a computer to the Active Directory domain.
3. Enable asynchronous disaster recovery DP by specifying a region or zone to replicate the server to.



If asynchronous disaster recovery is already enabled, it cannot be disabled. For more information, see [Disaster recovery](#).

### Steps

1. View the [File Servers list](#).
2. Locate the server in the list and click the Edit icon for that server. (For details about working with items in lists, see [List view actions](#)).
3. Make any changes as required; refer to [Create a file server](#) for field descriptions.
4. Click **Done**. This creates a job to modify the server.

### After you finish

Modify server is run as an asynchronous job. You can:

- Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
- After the job is finished, check the status of the sever in the Servers list.

### Delete file server

**Attention:** Deleting a file server will also delete the following:

- All backups associated with the file server
  - All disaster recovery replicated copies associated with the file server
- You cannot undo deletion of a server.

### Before you begin

To delete a file server, you must first delete all shares that exist on the server.

### Steps

1. View the [File Servers list](#).
2. Locate the server in the list and click the Delete icon for that server. (For details about working with items in lists, see [List view actions](#)).
3. In the Confirm Delete dialog box, enter the file server name to confirm that you want to delete the file server.



4. Click **Confirm**. This creates a job to delete the server.

#### After you finish

- When deleting CIFS-enabled file servers, the Active Directory computer object remains. Ask your Active Directory administrator to manually remove the computer object for the deleted file server from Active Directory.
- Delete server is run as an asynchronous job. You can:
  - Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
  - After the job is finished, verify that the sever has been removed from the Servers list.

## View file shares

The **Shares** list displays the file shares belonging to the selected Tenant. To view the list, select **FILE SERVICES > Shares** from the menu.

The file shares that are already a part of your existing environment and belong to the storage VMs configured in your NetApp Service Engine, can also be viewed on this screen and be managed as a part of your NetApp Keystone Flex Subscription (Flex Subscription) services. The file shares provisioned outside of the NetApp Service Engine are periodically imported and listed on this page with appropriate status codes.

If the imported file shares are in acceptable standards of NetApp Service Engine, that is all the parameters that are required for making the shares operational are available, they are imported with the status as `Operational` and can be directly managed through NetApp Service Engine. However, some shares might not be in the same standard as the existing shares on NetApp Service Engine. After import, these file shares are categorized with `Imported` or `Non-Standard` status. For understanding volume statuses and the steps to be taken to make them operational, see [Object states](#)

The Shares list displays simple information about each share. For more information about how to use the features of a list, see [List view](#).

- Share name

- Server on which it resides
- Share path
- CIFS share path (used for mounting the CIFS share with DNS integration)
- Subtenant to which it belongs
- Zone in which it exists
- Service level
- Operational state (operational, updating, or contact support)
- Creation date

## Create a file share

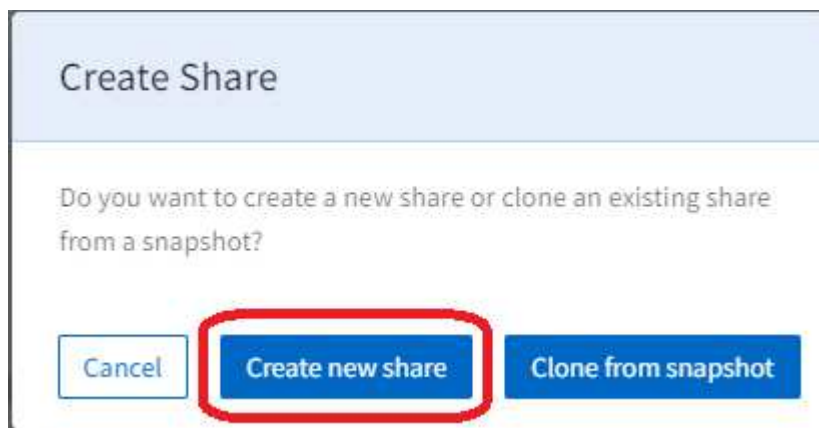
This section describes how to create a new share by directly specifying the share details. To create a new share based on a Snapshot of an existing share, see [Create Share from Snapshot](#).

### Before you begin

- A share is created on a file server. The file server must exist and be in an operational state before you can create a new file share.
- For creating a CIFS or NFS file share, the server must be enabled for the respective service. For a multi-protocol file shares, the server should support both CIFS and NFS services.
- To enable asynchronous disaster recovery options for the share, you must create the share on a server that has asynchronous disaster recovery enabled. For more information, see [Disaster recovery](#).
- To enable synchronous disaster recovery for a share, create the share in a zone that is MetroCluster-enabled.
- You can define a backup policy to capture backups of the file share on a scheduled basis. For more information, see [Backups](#).
- You can define a Snapshot policy to capture Snapshots of the file share on a scheduled basis. For more information, see [Snapshots](#).

### Steps

1. Go to **FILE SERVICES > Shares**.
2. Click **Create Share**.
3. In the Create Share dialog box, select Create New Share.



The Create Share page is displayed.

4. Select the share type: NFS, CIFS, or Multi-protocol. The options are enabled based on the services that your server supports.
5. Complete the following fields:

Field	Description
Name	Enter the share name.
Share Path	Enter the path for the file share. For CIFS shares, adding a \$ character to the end of the share path will make it a hidden share (for example, pathatomyhiddenshare\$).
Region	Select the region in which the share resides.
Zone	Select the zone of the share.
File Server	Select the file server to host the share. The list of the file server depends on the region, zone, and share type selected.
Security Style	Select the security style applicable to the file share. This list is automatically populated based on the share type selected.

6. Select a Performance Service Level. The IOPS and throughput limits are displayed based on the service level selected.



Select an option to view the performance details for that level (as peak/expected IOPS/throughput). Select the service level that best matches your needs.

7. Specify the capacity of the file share.



NetApp Service Engine displays a warning and the capacity bar changes color if the specified capacity puts the consumed capacity into burst (or even more into burst if it is already in burst). The capacity check is performed against the total capacity for all subscriptions in the tenancy.

8. If asynchronous disaster recovery is enabled on the underlying file server, asynchronous disaster recovery replication is automatically enabled for the new share. If you wish to exclude the share from asynchronous disaster recovery replication, toggle the Asynchronous Disaster Recovery button so that it is disabled.
9. If the share is being created in a zone that is MetroCluster-enabled, the Synchronous Disaster Recovery button is automatically enabled and cannot be disabled. The share will be replicated to the zone displayed in the panel below the Synchronous Disaster Recovery toggle.
10. If Snapshots are required for this file share:
  - a. Toggle to enable the Snapshot Policy and view the Snapshot Policy fields.
  - b. Specify when to create the Snapshots:
    - **Hourly.** Specify which minute (of the hour) to take Snapshot and the number of hourly Snapshots to retain.
    - **Daily.** Specify when (hour and minute) to take the Snapshot the number of daily Snapshots to

retain. If you want to specify multiple hours when the Snapshot has to be taken daily, you can add the values of the hours in a comma-separated list, for example, 5, 10, 15, and so forth.

- **Weekly.** Specify when (day of the week, hour, and minute) to take Snapshot and the number of weekly Snapshots to retain.
- **Monthly.** Specify when (day of the month, hour, and minute) to take Snapshot and the number of monthly Snapshots to retain.

11. To enable backups for this file share:

- a. Toggle to enable the Backup Policy and the Backup Policy fields.
- b. Specify the backup zone.
- c. Specify how many of each type of backup to keep: daily, weekly, and/or monthly.

12. For NFS or multi-protocol shares, specify the Export Policy. You can apply multiple export policies on a share. This section is available for only NFS and multi-protocol shares.

- a. Add the IPv4 address (with a subnet mask expressed as a number of bits) of the client to which the rule applies.
- b. Specify the read and write access, and whether the client has root access (superuser).

13. For a CIFS (SMB) or multi-protocol shares, specify the Access Control List (ACL) for restricting user access. This section is available for only CIFS and multi-protocol shares.

- a. Specify the Windows user or group based on the Active Directory (AD) settings to add to the ACL. If you specify the user name, include the user's domain in the <domain>\<username> format. The default value is `Everyone`.
- b. Specify the Windows permission. The default value is `Full control`. If a user is a part of two groups, the permissions of the group that has higher privileges get applied on the user's access.



The user or group name should follow the standard AD format. If the entered user or group does not match the user or user group configured on ONTAP, the ACL validation fails during a CIFS operation, even when the file share is operational.

14. If you want to add tags (key-value pairs) to the file share, specify them in the Tags section.

15. Click **Create**. This creates a job to create the share.

#### After you finish

- For CIFS type shares only: to make the shares available by host name, your domain administrator must update the DNS records with the CIFS server name and IP address. Otherwise, the share is only accessible through the IP address. For example:
  - With DNS records updated, use either the host name or IP to access the share: such as `\\hostname\share` or `\\IP\share`
  - With no DNS records updated, you must use the IP address to access the share i.e. `\\IP\share`
- Create share is run as an asynchronous job. You can:
  - Check the status of the job in the jobs list.
  - After the job is finished, check the status of the share in the Shares list.

## Create a file share from a Snapshot

You can create a new file share from an existing Snapshot. The new file share, cloned from the Snapshot, has the same properties as the file share from which the Snapshot is

created.

### Steps

1. Select **FILE SERVICES** from the left navigation pane and select **Shares**.
2. Click **Create Share** and select **Clone from snapshot**.  
The **Select Share** screen is displayed with all the file shares for the tenant. You can filter file shares by region, zone, and subtenant. You can select any file share that is in operational state.
3. Select the checkbox next to the file share that you want and click **Next**.  
The **Select Snapshot** screen is displayed with all the Snapshots for the file share.



For the selected file share, if you have some Snapshots created in your SnapCenter environment outside of NetApp Service Engine, you can find these Snapshots imported and listed for your selection. You can select these imported Snapshots and clone the new file shares from them.

You can search for a particular Snapshot or select the schedule type to filter the Snapshots.

4. Select the checkbox next to the Snapshot that you want to clone from and click **Next**.  
The new file share inherits the properties of the selected Snapshot.
5. Add **Name** and **Share Path**. Update the other settings, such as assigning a **Service Level**, and click **Create**.

### After you finish

- For CIFS type shares only: To make the shares available by host name, your domain administrator should update the DNS records with the CIFS server name and IP address. Otherwise, the share is only accessible through the IP address. For example:
  - With DNS records updated, use either the host name or IP to access the share: such as `\\hostname\share` or `\\IP\share`
  - With no DNS records updated, you should use the IP address to access the share i.e. `\\IP\share`
- **Create Share** is run as an asynchronous job. You can:
  - Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
  - After the job is finished, check the status of the share in the **Shares** list.

## Modify a file share

You can change the share name, the share type (CIFS, NFS, multi-protocol), service level, capacity, snapshot policy, export policy, Access Control List (ACL), and tags.



Using this method, you can move your shares to different performance levels if available. You can change the share type only if the server supports the respective services.

### Before you begin

The file share must be in the operational state. For understanding volume statuses and the steps to be taken to make them operational, see [View disks](#) and [Object states](#)

### Steps

1. View the [Shares list](#).



2. Locate the share in the list and click the Edit icon for that share. (For details about working with items in lists, see [List view actions](#)).
3. Make any changes as required; for field descriptions, see [Create a new file share](#).
4. Click **Done**. This creates a job to modify the share.

### After you finish

Modify share is run as an asynchronous job. You can:

- Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
- After the job is finished, check the status of the share in the Shares list.

## Delete a file share

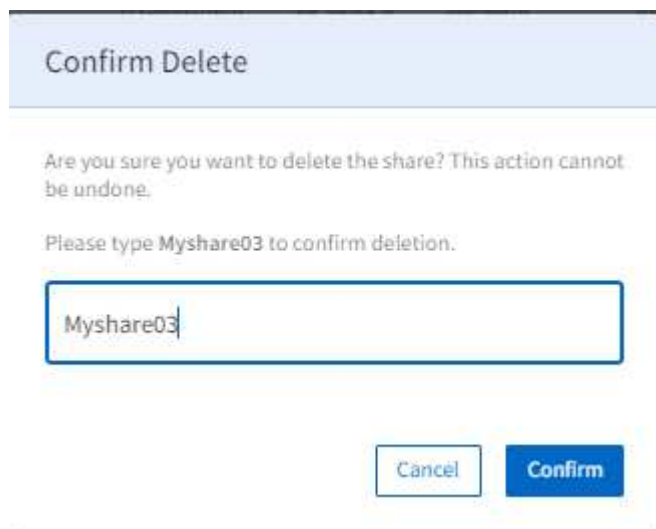
This section describes how to delete a file share.

### Attention:

- You cannot undo deletion of a share. After it is deleted, data cannot be recovered.
- Deleting a primary file share will delete all associated backups

### Steps

1. View the [Shares list](#).
2. Locate the share in the list and click the Delete icon for that share. (For details about working with items in lists, see [List view actions](#)).
3. In the Confirm Delete dialog box, enter the file share name to confirm that you want to delete the file share.



Confirm Delete

Are you sure you want to delete the share? This action cannot be undone.

Please type Myshare03 to confirm deletion.

Cancel Confirm

4. Click **Confirm**. This creates a job to delete the share.

### After you finish

Delete share is run as an asynchronous job. You can:

- Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
- After the job is finished, verify that the share has been removed from the Shares list.

## Create adhoc snapshot of a file share

This section describes how to create an adhoc snapshot of a file share.

### Steps

1. View the [Shares list](#).
2. Locate the share in the list and click the Snapshot icon for that share. (For details about working with items in lists, see [List view actions](#).)
3. In the Create Snapshot dialog, enter a name for your snapshot and click **Create**.

### After you finish

The snapshot can take a few minutes to become available.

## Overview

In block storage systems, the data storage is broken up into individual pieces each with a unique identifier. NetApp Service Engine refers to the block data storage as a block store, and the individual pieces as disks.

Block stores belong to a subtenant and are specified within a zone (one block store per zone per subtenant). A block store has networking attributes (for example, IP address and VLAN ID) which are used to access disks through the iSCSI or FC protocol. Disaster recovery DP can be enabled on a block store. For more information, see [Disaster recovery](#).

Block stores must be initialized before they can be used. Where block storage is available and it has not been initialized, it can be initialized prior to creating the first disk on the block store as part of the Create Disk process.

Disks are created on block stores. Disks have many configurable attributes including capacity and associated service level. DP options such as [Snapshots](#) and [Disaster recovery](#) can be enabled for a disk.

Access to disks is controlled through host groups. Host groups consist of initiator node names; by mapping one or more host groups to a disk, you can define which initiators have access to the disk.

Host groups:

- Are protocol specific. They can be either:
  - FC protocol host groups: these consist of initiators that are FC World Wide Port Names (WWPNs). For example, `20:56:00:a0:98:5c:0d:da`.
  - iSCSI protocol host groups: these consist of initiators that are iSCSI qualified names (IQNs). For example, `iqn.1998-01.com.vmware:esx2`.
- Consist of alias/initiator pairs. An alias allows a simple way to identify the initiator. For example, `esxserver1`.
- Can be created without any initiators. Empty host groups can be mapped to disks as placeholders but must be fully defined to allow access to the disk. Using host groups allows for:
  - Mapping multiple disks to the same set of initiators
  - Updating the set of initiators across multiple disks.

This section contains information on:

- Working with host groups:
  - View host groups
  - Create a host group
  - Modify host group initiators
  - Delete a host group
- Working with disks:
  - View disks
  - Create a new disk
  - Create a disk from snapshot
  - Modify a disk
  - Delete a disk
  - Create an adhoc snapshot of a disk

## Work with host groups

Host groups are defined to determine access to disks. Based on the initiator nodes assigned to a host groups, the access of that host group is determined.

### View host groups

To view the Host Group list, select **BLOCK STORAGE > Host Groups** from the menu.

The list displays the defined host groups.

From this page you can create a new host group, modify a host group, and delete a host group.

## Create a host group

There are two ways to create a host group:

- From the Host Groups page, described below.
- As part of creating a new disk. Use this method when you need to create a host group on a block store that has not yet been initialized. For more information see [Create a new disk](#).

It is possible to create an empty host group and map it to a disk as a placeholder. You must update the empty host group to add initiators before you will be able to access the storage.

### Before you begin

You need the following to create the host group:

- The subtenant, region and zone in which to create the host group.



If the block store for a subtenant/zone combination has not been previously initialized, you will not be able to create a host group using this method. An alternative is to follow the [Create a new disk](#) process, which allows you to initialize the block store and create a host group as part of the process.

- A name for the host group
- The host group protocol: iSCSI or FCP
- The list of initiators to add to the group: WWPNs for FC hosts nodes or IQNs for iSCSI host node names.
- An alias for each initiator; an alias is a simple name to identify the initiator server, or an individual port/interface on the server. For example, Server 4.

### Steps

1. [View the host groups](#) list.
2. Click **Create Host Group**.
3. On the Create Host Group page:
  - a. Select the protocol: iSCSI or FCP.
  - b. Select the subtenant, region, and zone and for the host group.
4. Specify a name for the host group.
5. Select the OS Type: the disk operating system.
6. Add the Initiators for the group. For each initiator, specify the alias and the initiator.
7. If required, add tags (key-value pairs) to the host group in the Tags section.
8. Click **Create**. This creates the host group.

### After you finish

After the host group is created, it is available for mapping to disks.

### Modify host Groups

You can modify a host group to add, remove or amend initiators.

Modifying a host group will modify access for each disk mapped to the host group.

You cannot modify the alias of an initiator. To change the alias, delete the initiator from the group and then re-create it.

### Steps

1. View the [Host Groups](#) list.
2. Locate the host group in the list and click the Edit icon for that host group.

To modify an existing initiator, locate the initiator in the list, edit the initiator value and click **Update**.

To add an initiator:

- a. click **Add Initiator**.
- b. Specify the Alias and Initiator.
- c. Click **Create**.

To remove an initiator from the host group, locate the initiator in the list and click the Delete icon.

3. Click **Done**.

## Delete a host groups

You can delete a host group if there are no disks mapped to the host group.

### Steps

1. View the [Host Groups](#) list.
2. Locate the host group in the list and click the Delete icon for that host group.
3. At the Confirm Delete dialog, enter the host group name to confirm that you want to delete the host group.
4. Click **Confirm**.

## View disks

The Disks list displays the disks belonging to the selected tenant. To view the list, select **BLOCK STORAGE > Disks** from the menu.

The disks that are already a part of your existing environment and belong to the storage VMs configured in your NetApp Service Engine, can also be viewed on this screen and be managed as a part of your NetApp Keystone Flex Subscription (Flex Subscription) services. The disks provisioned outside of the NetApp Service Engine are periodically imported and listed on this page with appropriate status codes.

If the imported disks are in acceptable standards of NetApp Service Engine, that is, if all the parameters that are required for making the disks operational are available, they are imported with the status as `Operational` and can be directly managed through NetApp Service Engine. However, some disks might not be in the same standard as the existing disks on NetApp Service Engine. After import, these disks are categorized with `Imported` or `Non-Standard` status. For understanding the disk statuses and the steps to be taken to make them operational, see [Object states](#)

In the Disks list, view simple information. For more information about how to use the features of a list, see [List view](#).

- Disk name
- Path to the disk
- Disk size
- Protocol
- Subtenant to which the disk belongs
- Zone in which disk exists
- Operational state

## Create a disk

This section describes how to create a new disk by directly specifying the disk details. For instructions on how to create a disk based on a Snapshot of an existing disk, see [Create a disk from a Snapshot](#).

### Before you begin

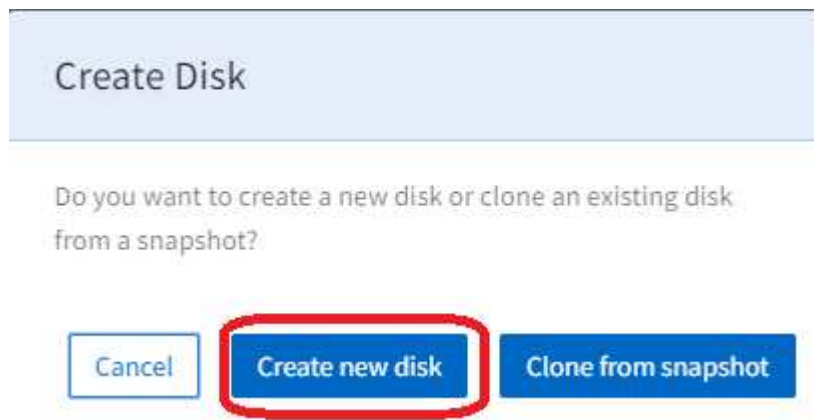
- If block storage is available but not initialized, you will be asked to initialize it before you can create the disk. To initialize the block store, you need:
  - The subnet for the block store. If you are not familiar with your network environment, please check with

you IT department for the relevant values.

- The protocol to be used. By default, block stores have the iSCSI service enabled. Ensure that the subnet has already been created to initialize block storage services (iSCSI) for the subtenant in the specified zone. You can optionally enable the FCP service if the infrastructure allows it.
  - The disaster recovery region, zone, and schedule if you want to enable asynchronous disaster recovery DP for the disk. For more information, see [Disaster recovery](#).
- Identify or define the host groups to be mapped to the disk. You can also create a host group as part of the disk creation.
  - To enable asynchronous disaster recovery DP options for the disk, you must create the disk on a block store that is asynchronous disaster recovery enabled. For more information, see [Disaster recovery](#).
  - To enable synchronous disaster recovery for a disk, create the disk in a zone that is MetroCluster-enabled.
  - You can define a backup policy to capture backups of the disk on a scheduled basis. For more information, see [Backups](#).


### Steps

1. View the [Disks list](#).
2. Click **Create Disk**.
3. In the Create Disk dialog box, select Create New Disk.



The Create Disk page is displayed.

4. Select the protocol for accessing the disk: iSCSI or FCP.
5. Select the subtenant, region, and zone, and subtenant for the new disk. The block store details for the selected region, zone, and subtenant display.
6. If block storage is not enabled in the selected zone for the selected subtenant, a message displays to initialize it.



7. Click **Initialize Block Storage**.

8. In the Initialize Block Storage dialog box:
  - a. In the iSCSI panel, mandatorily specify the subnet.  
The iSCSI service is enabled for all disks by default.
  - b. If required, enable the FCP protocol (only available if the underlying infrastructure supports it).
  - c. If required, enable the Asynchronous Disaster Recovery options (enable and select Region, Zone, and Schedule).
  - d. If the zone selected is MCC-enabled, the Synchronous Disaster Recovery toggle is enabled and cannot be disabled. Disks created in this block store synchronously replicate to the zone displayed in the Synchronous Disaster Recovery panel.
  - e. Click **Create** to initialize the block store. Wait until the block store initializes. The display returns to Create Disk page.
9. Complete the following fields:

Field	Description
Name	Enter the disk name.
Disk Path	Enter the path for the disk.
OS Type	Select an operating system for the disk.
Host Groups	This list displays existing host groups that match the protocol, OS type, subtenant and zone selected for the disk. Select one or more host groups. To define a new host group, see the next step.

10. If required, create a new host group:
  - a. Click **Create Host Group**. The Create Host Group dialog is displayed.
  - b. Specify the Name of the host group.
  - c. Add the Initiators for the group. For each initiator, specify the alias and the initiator.
  - d. If required, add tags (key-value pairs) to the host group in the Tags section.
  - e. Click **Create**. The system creates the host group and displays a message when it is successfully created.
  - f. To map the newly created host group to the disk, go to the Host Groups field and select it.

11. Select a performance service level.

Select an option to view the performance details for that level (as peak/expected IOPS/throughput). Select the service level that best matches your needs.

12. Specify the capacity of the disk.

13. If Snapshots are required for this disk:

- a. Toggle to enable the Snapshot Policy to view the Snapshot Policy fields.
- b. Specify when to create the Snapshots:
  - **Hourly**. Specify which minute (of the hour) to take Snapshot (check) and the number of hourly Snapshots to retain.
  - **Daily**. Specify when (hour and minute) to take the Snapshot (check) and the number of hourly

Snapshots to retain. If you want to specify multiple hours when the Snapshot has to be taken daily, you can add the values of the hours in a comma-separated list, for example, 5, 10, 15, and so forth.

- **Weekly.** Specify when (day of the week, hour and minute) to take Snapshot (check) and the number of weekly Snapshots to retain.
  - **Monthly.** Specify when (day of the month, hour, and minute) to take Snapshot and the number of monthly snapshots to retain.
14. If asynchronous disaster recovery is enabled on the underlying block store, asynchronous disaster recovery replication is automatically enabled for the new disk. If you wish to exclude the disk from asynchronous disaster recovery replications, toggle the Asynchronous Disaster Recovery toggle so that asynchronous disaster recovery is disabled.
  15. If the disk is being created in a zone that is MetroCluster-enabled, the Synchronous Disaster Recovery button is enabled and cannot be disabled. The disk will be replicated to the zone displayed in the Synchronous Disaster Recovery panel.
  16. To enable backups for this disk:
    - a. Toggle to enable the Backup Policy to view the Backup Policy fields.
    - b. Specify the backup zone.
    - c. Specify how many of each type of backup to keep: daily, weekly, and/or monthly.
  17. If you want to add tags (key-value pairs) to the disk, specify them in the Tags section.
  18. Click **Create**. This creates a job to create the disk.

#### After you finish

Create disk is run as an asynchronous job. You can:

- Check the status of the job in the jobs list.
- After the job is finished, check the status of the disk in the Disks list.

## Create a disk from a Snapshot

You can create a new disk from an existing Snapshot. The new disk, cloned from the Snapshot, has the same properties as the disk from which the Snapshot is created.

#### Steps

1. Select **BLOCK STORAGE** from the left navigation pane and select **Disks**.
2. Click **Create Disk** and select **Clone from snapshot**.  
The **Select Disk** screen is displayed with all the disks for the tenant. You can filter disks by region, zone, and subtenant. You can select any disk that is in operational state.
3. Select the checkbox next to the disk that you want and click **Next**.  
The **Select Snapshot** screen is displayed with all the Snapshots for the disk.



For the selected disk, if you have some Snapshots created in your SnapCenter environment outside of NetApp Service Engine, you can find these Snapshots imported and listed for your selection. You can select these imported Snapshots and clone the new disks from them.

You can search for a particular Snapshot or select the schedule type to filter the Snapshots.

4. Select the checkbox next to the Snapshot that you want to clone from and click **Next**.



The new disk inherits the properties of the selected Snapshot.

5. Add **Name** and **Disk Path**. Update the other settings, such as assigning a **Service Level**, and click **Create**.

### After you finish

**Create Disk** is run as an asynchronous job. You can:

- Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
- After the job is finished, check the status of the disk in the **Disks** list.

## Modify a disk

You can change the disk name, the host group mapping, performance service level, capacity, and snapshot policy. Using this method, you can move your disks to different service levels if available.

### Before you begin

The disk must be in an operational state. For understanding volume statuses and the steps to be taken to make them operational, see [View disks](#) and [Object states](#)

### Steps

1. View the [Disks list](#).
2. Locate the disk in the list and click the Edit icon for that disk. (For details about working with items in lists, see [List view actions](#)).
3. Make any changes as required; refer to [Create a new disk](#) for field descriptions.
4. Click **Done**. This creates a job to modify the disk.

### After you finish

Modify disk is run as an asynchronous job. You can:

- Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
- After the job is finished, check the status of the disk in the Disks list.

## Delete a disk

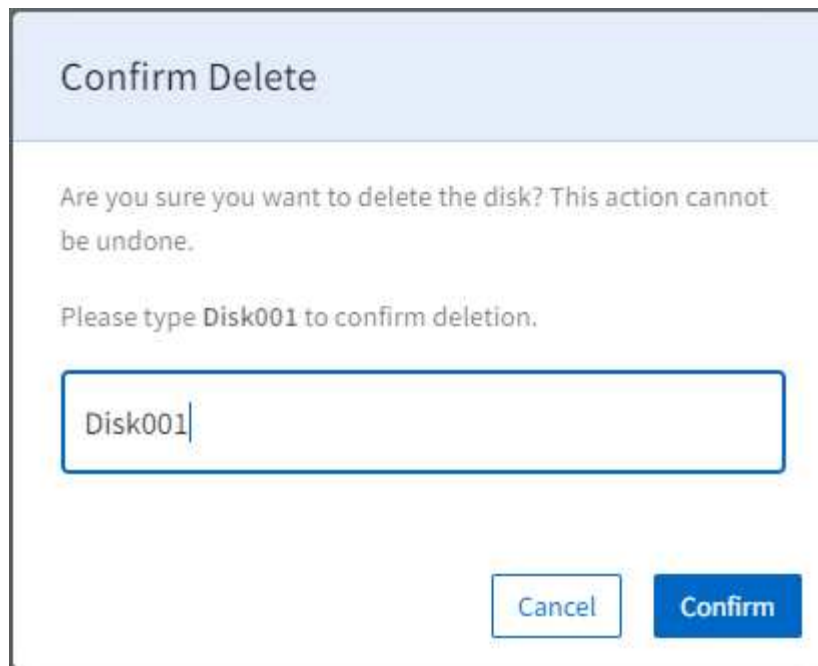
This section describes how to delete a disk.

### Attention:

- You cannot undo deletion of a disk.
- Deleting a primary disk will delete all associated backups

### Steps

1. View the [Disks list](#).
2. Locate the disk in the list and click the Delete icon for that disk. (For details about working with items in lists, see [List view actions](#)).
3. In the Confirm Delete dialog box, enter the disk name to confirm that you want to delete the disk.



4. Click **Confirm**. This creates a job to delete the disk.

#### After you finish

Delete disk is run as an asynchronous job. You can:

- Check the status of the job in the jobs list. For information about tracking jobs, see [here](#).
- After the job is finished, verify that the disk has been removed from the Disks list.

## Create an adhoc snapshot of a disk

This section describes how to create an adhoc snapshot of a disk.

#### Steps

1. View the [Disks list](#).
2. Locate the disk in the list and click the Snapshot icon for that disk. (For details about working with items in lists, see [List view actions](#)).
3. In the Create Snapshot dialog box, enter a name for your snapshot and click **Create**.

#### After you finish

The snapshot might take a few minutes to become available.

## Work with object storage

### Overview

Object storage requires an object storage service subscription. When an object storage service is available as part of the subscription, the storage must be initialized before it can be used.

With object storage, objects are stored in S3 buckets. Access to the S3 buckets is managed through

permissions set on object storage groups. Object storage users are granted membership to one or more object storage groups, inheriting permissions from the group membership.

Each object storage user has associated S3 keys that allow access to the object storage.



Access to the object storage is through an S3 compatible browser.

This section describes how to manage your object storage, including how to:

- Initialize object storage.
- Create S3 buckets.

It is not possible to delete S3 buckets through the NetApp Service Engine web portal; delete these buckets using an S3 compatible browser.

- Create and manage object storage groups.
- Create and manage object storage users, including creating S3 keys.

## View buckets

The Buckets list displays the existing buckets for a selected Subtenant.

1. To view the list, select Object Storage > Buckets from the menu.
2. Use the Subtenant drop-down menu to select another subtenant if required.

Simple information for each bucket is displayed, such as:

- Name
- Date created
- Usage

For more information about how to use the features of a list, see [List view](#).

## Initialize object storage

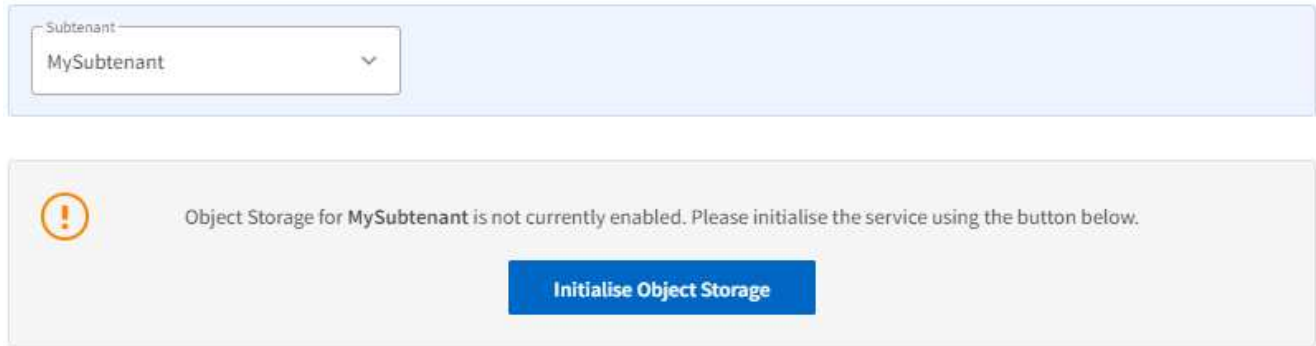
Object storage must be initialized before it can be used.

If you attempt to create a storage bucket on object storage that has not yet been initialized, you will see the prompt to initialize the storage. The following section describes the process.

### Steps


1. View the [Buckets list](#) and select the Subtenant under which to create the bucket.
2. If the object storage has not been initialized for this subtenant, the following message is displayed:

# Buckets



The screenshot shows a dropdown menu for 'Subtenant' with 'MySubtenant' selected. Below it is a notification box with an orange warning icon and the text: 'Object Storage for MySubtenant is not currently enabled. Please initialise the service using the button below.' A blue button labeled 'Initialise Object Storage' is centered in the notification box.

3. Click **Initialize Object Storage**.
4. In the Initialize Object Storage dialog box, select the quota for the object storage and click **Create**.



The 'Initialise Object Storage' dialog box features a 'Quota Usage' section with a progress bar showing 2/5 TiB. Below this is a 'Quota' section with a slider set to 100, radio buttons for 'GiB' (selected) and 'TiB', and 'Cancel' and 'Create' buttons at the bottom right.

5. Check the jobs list to make sure the object storage is successfully initialized.

## Create buckets

Buckets are created under a subtenant.

### Steps

1. View the [Buckets list](#) and select the subtenant under which to create the bucket.
2. Click **Create Bucket**.
3. In the Create Bucket dialog box, enter the name of the new bucket and click **Create**. This creates a job to create the bucket.

### After you finish

Create bucket is run as an asynchronous job. You can:

- Check the status of the job in the jobs list.
- After the job is finished, check the status of the server in the Buckets list.

## View the object storage group and users

The User Management list displays the object storage users and groups for a specified subtenant.

### Steps

1. From the menu, select Object Storage > Users.
2. Use the Subtenant drop-down menu to select another subtenant if required.
3. Use the Users and Groups links to toggle the display between users and groups.
  - The Users list displays the object storage users defined for the selected subtenant. For each user, it displays their name and the object storage groups to which they belong.
  - The Groups list displays the object storage groups defined for the selected subtenant. For each group, it displays the group name and S3 policy.

For more information about how to use the features of a list, see [List view](#).

## Create an object storage group

Use this method to create a new object storage group.

### Steps

1. View the [Users List](#).
2. Make sure the correct subtenant is selected and the display is showing groups (not users).
3. Click **Create Group**.
4. In the Create Group dialog box, enter the group name.
5. Select the S3 policy from the list.
6. Click **Create**. This creates a job to create the group with the specified settings.

### After you finish

Create object storage group is run as an asynchronous job. You can check the status of the job in the jobs list. For information about tracking jobs, see [here](#).

## Modify an object storage group

Use this method to modify details of an object storage group such as the group name and S3 policy.

### Steps

1. View the [Groups List](#).
2. Make sure the correct subtenant is selected and the display is showing groups (not users).
3. Locate the group in the list and click the Edit icon for that group.

For details about working with items in lists, see [List view actions](#).

4. In the Edit Group dialog box, enter the user name.
5. Select the S3 policy that applies to the modified group.
6. Click **Create**. This creates a job to modify the group with the specified settings.

## Delete an object storage group

Use this method to delete an Object Storage group.

### Steps

1. View [the Groups List](#).
2. Make sure the correct subtenant is selected and the display is showing groups (not users).
3. Locate the group in the list and click the Delete icon for that group.
4. To confirm the delete action, click **Confirm**. The object storage group is removed from the list.

## Create Object Storage User

### Before you begin

Because an object storage user belongs to an object storage group, the object storage group must exist before you can create the user.

### Steps

1. View the [Users List](#).
2. Make sure the correct subtenant is selected and the display is showing users (not groups).
3. Click **Create User**.
4. In the Create User dialog box:
  - a. Enter the user name.
  - b. Select the group to which the member belongs.
5. Click **Create**. This creates a job to create the user with the specified settings.

### After you finish

Create user is run as an asynchronous job. You can:

- Check the status of the job in the jobs list.
- After the job is finished, check the status of the share in the Users list.

## Modify an object storage user

Use this method to modify details of an object storage user such as the user name and the group to which they are assigned.

### Steps

1. View the [Users List](#).
2. Make sure the correct subtenant is selected and the display is showing Users (not Groups).
3. Locate the user in the list and click the Edit icon for that group.

For details about working with items in lists, see [List view actions](#).

4. In the Edit User dialog box:
  - a. Enter the user name.
  - b. Select the object storage group to which the user belongs.
5. Click **Update**. This creates a job to create the user with the specified settings.

## Create an S3 key for a user

Use this method to create an S3 key for a user.

Make sure that you capture the key immediately after it is created. There is no way to retrieve key details after it has been created.

### Steps

1. View the [Users List](#).
2. Make sure the correct subtenant is selected and the display is showing users (not groups).
3. Locate the user in the list and click the Key icon for that group.

For details about working with items in lists, see [List view actions](#).
4. The Manage S3 Keys dialog box displays showing the list of keys for the user.
5. If there are no existing keys for the user, click **Create S3 Key** to view the fields to create a key.
6. If you want the key to expire, specify the expiry date and expiry time in UTC. Otherwise, leave the default values for these fields.
7. To create the key, click the checkmark. The S3 Key Details dialog box is displayed showing the access key and secret key.
8. Copy both the access and secret key in preparation of advising the user of the details.

Make sure that you capture the key details before closing the S3 Key Details dialog box. After closed, the key details will be obscured and you will not be able to retrieve them.

9. Click **Close**.
10. Click **Close** again to close the Manage S3 Keys dialog.

## Delete an S3 key for a user

Use this method to delete an S3 key for a user.

### Steps

1. View the [Users List](#).
2. Make sure the correct subtenant is selected and the display is showing Users (not Groups).
3. Locate the user in the list and click the Key icon for that group.

For details about working with items in lists, see [List view actions](#).
4. The Manage S3 Keys dialog displays showing the list of keys for the user.
5. In the list of keys, locate the key to delete and click the Delete icon.
6. Click **Confirm** to confirm the delete action. The key is removed from the list of user keys.

7. Click **Close** to close the Manage S3 Keys dialog box.

## Manage Cloud services

### Managing Cloud services

#### Overview

When registered with the NetApp Service Engine, you can manage your cloud volumes services (CVS) through the NetApp Service Engine web portal. This includes:

- Azure NetApp Files (ANF)
- (Future) Amazon Web Services (AWS)
- (Future) Google Cloud Platform (GCP)

You can only register existing CVS subscriptions with NetApp Service Engine; to create a new subscription, create the CVS subscription in the normal way and then register it with the NetApp Service Engine using your subscription credentials.

This section describes how to manage cloud services



Managing cloud services requires Customer Admin access.

You can:

- View a list of cloud services
- Add a cloud service
- Modify cloud service details
- Delete a cloud service

#### View a list of cloud services

The Cloud Services list displays a list of all the cloud services in the selected tenancy.

To view the list, select **Cloud Services** > **Cloud Services** from the menu.

The list displays simple information about each cloud service.

The action icons next to each cloud service allow you to modify or delete the cloud service.

#### Add a cloud service

This section describes how to add a cloud service.

NetApp Service Engine supports Azure NetApp Files.

#### Before you begin

You need the information about the cloud services, including credentials to connect to the service. The information varies with the cloud service type.



## Steps

1. Select **Cloud Services** from the menu.
2. Click **Create**.
3. On the Create Cloud Service page, select the cloud service type and then complete the relevant information for that service type:

### Azure:

- a. Select the Subtenant.
  - b. Enter a name for the OCCM Working Environment that will be created for the service.
  - c. Enter the Azure Directory (tenant) ID, the Azure Client ID, Azure Client Secret and the Name of the cloud service.
4. Specify tags for the service if required.
  5. Click **Create**.

## Modify the cloud service details

Use this method to modify the details of a cloud service. You can:

- Modify the details for each service, as listed in the table below.
- Add or remove tags for the service.

Cloud Platform	You Can Modify:
Azure	Azure Client ID, Azure Client Secret, and name
Azure US Government	TBA
GCP	TBA
AWS	TBA

## Steps

1. Select **Cloud Services > Cloud Services** from the menu.
2. Locate the cloud service in the list and click the Edit icon.
3. On the Update Cloud Service dialog box, make any changes as required.
4. Click **Update**.

## Delete a cloud service

Use this method to delete a CVS subscription from the NetApp Service Engine web portal.

## Steps

1. Select **Cloud Services > Cloud Services** from the menu.
2. Locate the cloud service in the list and click the Delete icon.
3. In the Confirm Delete dialog box, enter the name of the cloud service to be deleted.
4. Click **Confirm**.

# Managing Azure NetApp Files

## Managing Azure NetApp accounts

This section describes how to manage Azure NetApp accounts:



Managing Azure NetApp accounts requires Customer Admin access.

In this section:

- View a list of azure NetApp accounts
- Add an Azure NetApp account
- Modify Azure NetApp account details
- Delete an Azure NetApp account

### View a list of Azure NetApp accounts

The Azure NetApp Accounts list displays a list of all the Azure NetApp accounts in the tenancies that the logged-in user has access to. To view the list, select **ADMINISTRATION > Azure NetApp Accounts** from the menu.

The list displays simple information about each Azure NetApp account. The action icons next to each Azure NetApp account allow you to modify or delete the Azure NetApp account.

### Add an Azure NetApp account

This section describes how to add an Azure NetApp account.

#### Before you begin

You will need:

- The subtenant to which the Azure NetApp Account belongs
- The ANF Instance
- The Azure location (region)
- The Azure Resource Group
- A name for the account

#### Steps

1. Select **ADMINISTRATION > Azure NetApp Accounts** from the menu.
2. Click **Create**.
3. In the Create Azure NetApp Account dialog box, specify the subtenant, the Azure location (region), the Azure resource group, and a name for the Azure NetApp account.
4. If needed, add tags to the account.
5. Click **Create**. When created successfully, the account state will be Succeeded (can be viewed in the Azure NetApp Accounts list).

## Modify Azure NetApp account tags

Use this method to modify tags associated with an Azure NetApp account.

### Steps

1. Select **ADMINISTRATION > Azure NetApp Accounts** from the menu.
2. Locate the Azure NetApp Account in the list and click the edit icon.
3. On the Manage Azure NetApp Account dialog box, add, remove, or edit tags as required.
4. Click **Update**.

## Delete an Azure NetApp account

Use this method to delete an Azure NetApp Account.

### Steps

1. Select **ADMINISTRATION > Azure NetApp Accounts** from the menu.
2. Locate the Azure NetApp account in the list and click the delete icon.
3. In the Confirm Delete dialog box, enter the name of the Azure NetApp account to be deleted.
4. Click **Confirm**.

## Managing Azure capacity pools

This section describes how to manage Azure capacity pools:



Managing Azure capacity pools requires Customer Admin access.

This section describes how to perform the following tasks:

- View a list of Azure capacity pools
- Add an Azure capacity pool
- Modify Azure capacity pool details
- Delete an Azure capacity pool

### View a List of Azure capacity pools

The Azure capacity pools list displays a list of capacity pools for a specified subtenant and location. To view the list, select **FILE SERVICES > Azure Capacity Pools** from the menu.

The list displays simple information about each Azure capacity pool. The action icons next to each Azure capacity pool allow you to modify or delete the Azure capacity pool.

### Add an Azure capacity pool

This section describes how to add an Azure capacity pool.

### Before you begin

You will need:

- The subtenant to which the capacity belongs

- The ANF account
- The service level
- A name for the capacity pool
- The size of the pool

### Steps

1. Select **FILE SERVICES > Azure Capacity Pools** from the menu.
2. Click **Create**.
3. In the Create Azure Capacity Pool dialog box, specify the subtenant, the ANF account, the service level, the name and the size.
4. Add tags to the capacity pool if required.
5. Click **Create**. When the capacity pool is created the state changes to *Succeeded* (can be viewed in the Azure Capacity Pools list).

### Modify Azure capacity pool details

Use this method to modify the details of an Azure capacity pool. You can modify the service level the size of the pool, and you can add or remove tags.

### Steps

1. Select **FILE SERVICES > Azure Capacity Pools** from the menu.
2. From the list, locate the Azure capacity pool and click the edit icon.
3. In the Manage Azure Capacity Pool dialog box, make any changes as required.
4. Click **Update**.

### Delete an Azure capacity pool

Use this method to delete an Azure capacity pool.

### Steps

1. Select **FILE SERVICES > Azure Capacity Pools** from the menu.
2. From the list, locate the Azure capacity pool and click the delete icon.
3. In the Confirm Delete dialog box, enter the name of the Azure capacity pool to be deleted.
4. Click **Confirm**.

### Managing Azure volumes

This section describes how to manage Azure volumes.



Managing Azure Volumes requires Customer Admin access.

This section describes how to perform the following tasks:

- View a list of Azure volumes.
- Add an Azure volume.
- Modify Azure volume details.

- Delete an Azure volume.

### View a list of Azure volumes

The Azure Volumes list displays a list of Azure volumes pools for a specified subtenant and location. To view the list, select **FILE SERVICES > Azure Volumes** from the menu.

The list displays simple information about each Azure volume. The action icons next to each Azure volume allow you to modify or delete the Azure volume.

### Add an Azure volume

This section describes how to add an Azure volume.

#### Before you begin

You will need:

- The subtenant to which the Azure volume belongs
- The ANF Account
- Azure Virtual Network (VNet) and Subnet
- The capacity pool to house the volume
- A name for the volume
- The file path
- The quota (size) of the volume

#### Steps

1. Select **FILE SERVICES > Azure Volumes** from the menu.
2. Click **Create**.
3. In the Create Azure Volume dialog box, specify the subtenant, the ANF account, the VNet, the subnet, the capacity pool, the name of the volume, the file path and the quota. Ensure that you enter a unique file path for a successful volume creation.
4. If required, add tags to the volume.
5. Click **Create**. When an Azure volume is created, the state of the volume changes to `Succeeded` (can be viewed in the Azure Volumes list).

#### Modify Azure volume details

Use this method to modify an Azure Volume's details. You can modify the file path and quota, and add or remove tags.

#### Steps

1. Select **FILE SERVICES > Azure Volumes** from the menu.
2. From the list, locate the Azure volume and click the edit icon.
3. In the Manage Azure Volume dialog box, make the required changes. If you are modifying the file path or quota, ensure that the file path is unique and that the quota does not exceed the available quota in the pool.
4. Click **Update**.

## Delete an Azure volume

Use this method to delete an Azure volume.

### Steps

1. Select **FILE SERVICES > Azure Volumes** from the menu.
2. Locate the Azure volume in the list and click the delete icon.
3. In the Confirm Delete dialog box, enter the name of the Azure volume to be deleted.
4. Click **Confirm**.

## Managing Cloud Volumes Service for GCP

Subscriptions to Cloud Volumes Service for Google Cloud Platform (GCP) are associated with specific subtenants under a tenant. A subtenant is subscribed to a single Google Cloud service. You can create and manage cloud services (Google Cloud instances) for your tenants and subtenants and provision storage on Google Cloud based on the tenancies. The Cloud Volumes Service accounts of the tenants and subtenants should be set up before you can manage Google Cloud instances through NetApp Service Engine. For information about Cloud Volumes Service on GCP, see [Learn about Cloud Volumes Service for Google Cloud](#).



The Cloud Volumes Service accounts for the tenants and subtenants should already be configured for GCP. Contact support for the service account details. You need customer administrator access to manage Google Cloud Volumes.

## Managing your cloud services on GCP

You can manage your cloud services on Google Cloud, from the **CLOUD SERVICES** menu on the left navigation pane. The Cloud Services page displays all the cloud services for all subtenants of a selected tenant. You can view, modify, and delete the cloud services from this page. You can also add a new service by following these steps.

### Before you begin

You need to ensure that:

- The subtenant is subscribed to the Google Cloud service.
- You have the JSON file for the subscriber's service account and you are aware of the GCP project number that is associated with the Google Cloud instance.
- To ensure that Cloud Volumes Service for GCP is already configured on Cloud Manager by your support team. For more information, see [Set up Cloud Volumes Service for Google Cloud](#).

### Steps

1. Select **CLOUD SERVICES** from the menu.
2. Click **Create**, and then click **Google Cloud**.
3. On the Create Cloud Service page for Google Cloud, select the subtenant for which you want to create the service, and add a name and the GCP project number.
4. In the Google Cloud Credentials section, upload the JSON file with the credentials for the service account.

5. If required, add tags to the service.
6. Click **Create**. When created, the service state changes to `Succeeded` on the Cloud Services page.

## Managing volumes on GCP

Before managing cloud volumes on GCP, the Cloud Volumes Service account for the tenant should be created and the account should be subscribed to cloud volumes. Contact support for ensuring that these steps have been complete.

You can manage the volumes for a subtenant on GCP from **FILE SERVICES > GCP Volumes** on the left navigation pane. The GCP Volumes page lists the existing cloud volumes created on GCP for the selected subtenant and corresponding region. You can view, modify, and delete the existing volumes from this page. You can also provision a new cloud volume by following these steps.


### Steps

1. Select **FILE SERVICES > GCP Volumes** from the menu.
2. Click **Create**. The Create GCP Volume page is displayed.
3. Select the subtenant for which you want to create the volume.
4. Select the region and network. The list is filtered based on the Cloud Volumes Service account for the subtenant on GCP.
5. Add a name for the volume.
6. Select the appropriate Service Level.
7. Enter the file path of the volume as the creation token.
8. Assign the volume quota in bytes.
9. If required, add tags for the volume.
10. Click **Create**. When created, the volume state changes to `Succeeded` on the GCP Volumes page.

Watch the following video to learn how to create a volume on GCP through Cloud Volumes Service and NetApp Service Engine:

# Creating a Volume on the GCP Through CVS and NetApp Keystone User Interface

NetApp Keystone



© 2021 NetApp, Inc. All rights reserved.

## View reports

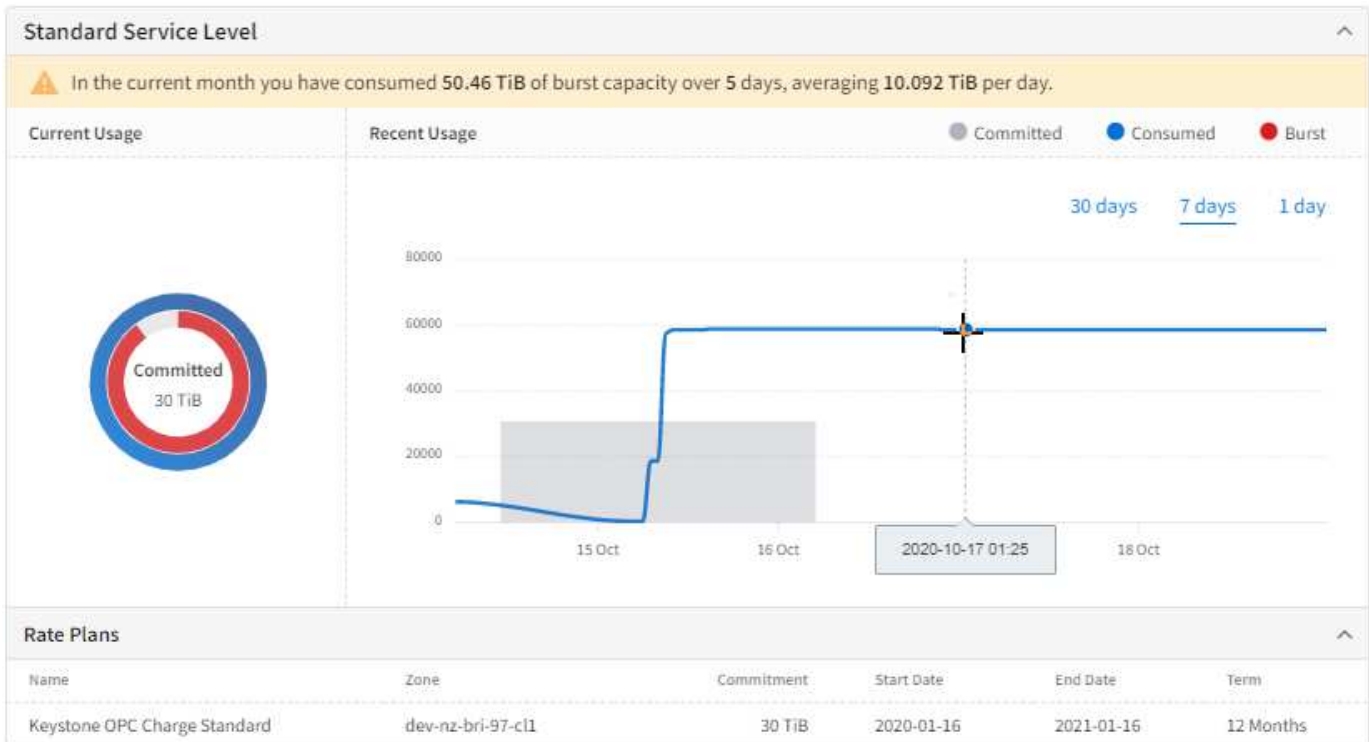
You can create capacity and performance reports with respect to your NetApp Keystone Flex Subscription usage and also for your tenant usage.

### Keystone subscription capacity usage report

The Capacity Usage for Keystone Subscriptions page displays capacity usage for each storage service in a subscription over time. As a NetApp administrator, You can view the capacity usage report for all tenants and partners in a subscription. As a partner administrator, you can view the capacity report for your Flex Subscription usage.

Use the graphical reports on this page to view the trends usage for all the storage services, as well as add-on data protection services, in separate tabs. When a service is in burst, a banner displays the burst capacity used in the current billing cycle.





To view the Capacity Usage for Keystone Subscriptions page, select **Reports > Keystone Usage** from the menu.

To view the capacity usage for a service, perform the following steps:

### Steps

1. Select the subscription number containing the service from the **Subscription** drop-down list.
2. You can choose to view the capacity usage for base service levels or data protection services by selecting the different tabs. The page displays the service level view by default.
3. You can scroll the page to view the services, and use the period filters to limit the display to a selected period.

## Tenant subscription capacity usage report

The Capacity Usage for Tenant Subscriptions page displays capacity usage over time for each tenant for the subscribed storage services in Flex Subscription. This page is available to NetApp, partner, and tenant or customer admins to view the capacity usage reports for a specific tenant.



For the tenant administrator, the page appears as "Capacity Usage".

Use the graphical reports on this page to view the trends usage for both all the storage services, as well as add-on services, such as data protection, in separate tabs. When a service is in burst, a banner displays the burst capacity used in the current billing cycle.

To view the Capacity Usage for Tenant Subscriptions page, select **Reports > Tenant Usage/Capacity Usage** from the menu.

To view the capacity usage for a tenant, perform the following steps:

### Steps

1. Select the subscription number containing the service from the Subscription drop-down list.
2. You can choose to view the capacity usage for base service levels or data protection services by selecting the respective tabs. The page displays the service level view by default.
3. You can scroll the page to view the services, and use the period filters to limit the display to a selected period.

## Performance report

The Performance page (shown in the image below) displays information about the performance of individual disks or shares. It displays information on three performance measures:

- Input/output operations per second per terabyte (IOPS/TiB).

The rate at which IOPS is occurring on the storage device.

- Throughput in MBps.

Throughput measures the data transfer rate to and from the storage media in megabytes per second.

- Latency (ms).

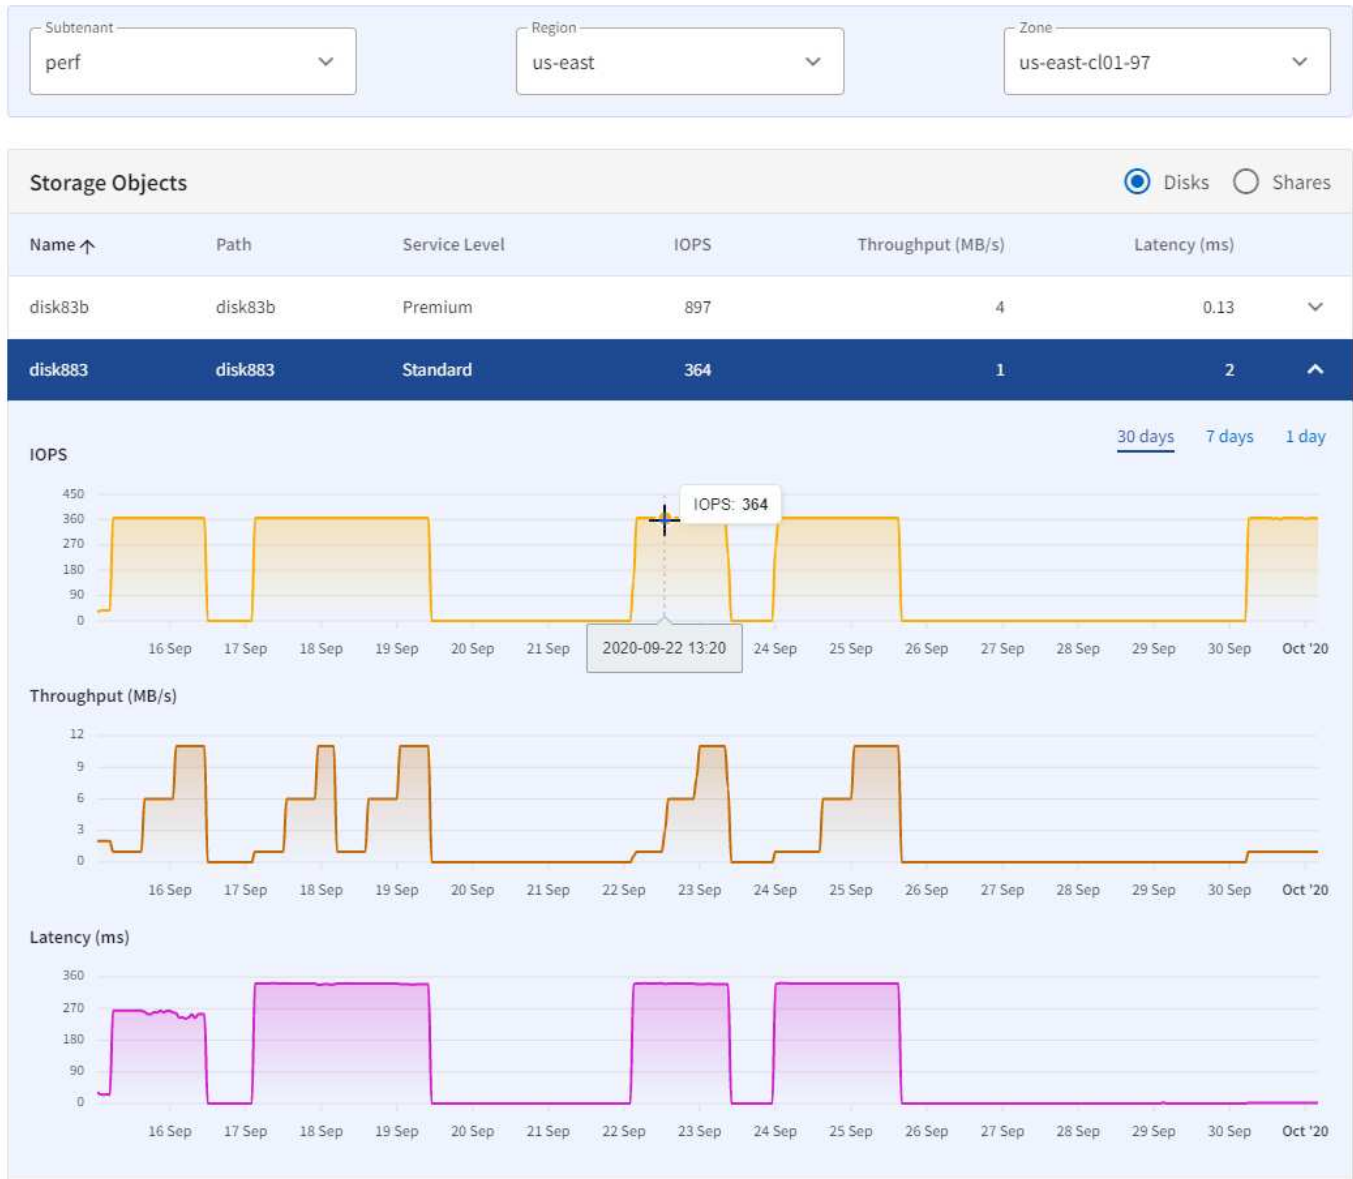
Average time for reads/writes from the disk/share in milliseconds.

To view the Performance page, select Reports > Performance from the menu.

To view the performance details for a disk/share, complete the following steps.

1. Select the **Subtenant**, **Region**, and **Zone**, and then the storage object type to view either **Disks** or **Shares** for the disk or share. The page displays the list of storage objects meeting the selected criteria, showing the latest performance data for those objects.
2. To view the trend in performance data for a selected share or disk, locate the storage object in the list and click to expand the panel. The performance graphs for the selected object is displayed.

# Performance



The graphs show the performance of the storage over time. You can:

- Choose the period to display by selecting the period filters or click and drag on the graph.
- Hover the cursor over a point in the graph to see more information for that point.

## Back up file shares and disks

You can view and manage the Snapshots (backups or recovery points) of the volumes (file shares and disks) created in your environment to maintain multiple copies of primary file shares and disks remotely.

The replicated file shares and disks from different zones are stored on a dedicated backup storage VM (storage virtual machine, also known as SVM) on each cluster. Backups can be performed between two separate ONTAP clusters in different zones. The standard backup policies supported by NetApp Service

Engine are:

- Adhoc
- Daily
- Weekly
- Monthly

If your volume does not follow the standard backup policies supported by NetApp Service Engine, you can [raise a service request](#) of the **Non Standard Volume** category to fix the issue.



An adhoc (non-scheduled) backup can be performed at any time, manually through the GUI or by using APIs.

You can view the list of Snapshots created for your file shares and disks on the **Backups** page. Backups are displayed based on the selected tenant.

## Create a backup object

You can create Snapshots for your volumes by defining the backup policy when you [create a file share](#) or [create a disk](#).

### Steps

1. Enable **Backup Policy**.
2. Specify the backup zone and the number of backups to be taken on an adhoc, daily, weekly, and/or monthly basis.
3. Click **Create**.



You can also create or schedule backups by editing the backup policy for an existing file share or disk.

## Modify a backup object

For a selected backup object of a file share or disk, you can disable the backup or detach the backup object from its source and make it an orphan backup. This operation breaks the SnapMirror relationship with the source. You can also modify the number of backups to be taken on an adhoc, daily, weekly, and/or monthly basis.

### Steps

1. Locate the backup in the list and click the edit icon, at the end of the row, for that object.
2. Make changes to the backup source and the number of each type of backup as required.  
You can break the SnapMirror relationship with the source volume by changing the backup source to **None (orphaned)**. The backups are retained as orphan backups for restoration even when the source volumes are deleted. The state of the backup object changes, and the status of the source volume (orphan or deleted) is appended to the source volume name on the same list.  
For information about object states, see [here](#)
3. Click **Done**.  
The Snapshot is modified. If you have broken the SnapMirror relationship with the source, that is, orphaned the backup object, the name of the Source in the **Backups** tab appears as `deleted`.



You can also break the SnapMirror relationship with the source volume by disabling **Backup Policy** for the volume.

## Delete a backup object

Delete an unused backup volume to increase your storage space.

### Steps

1. Locate the backup in the list and click the delete icon, at the end of the row, for that object.
2. In the **Confirm Delete** dialog box, enter the backup name (source volume name) to confirm that you want to delete the backup object.
3. Click **Confirm**.



When a backup is deleted, all the recovery points for the backup are deleted. Also, even if you delete a primary file share or disk, the backup is retained as an orphan backup with all the recovery points for restoration. To delete the backup volume and recovery points of the deleted source, you should explicitly delete the backup object from the **Backups** tab.

## Backup Restore

You can [raise a service request](#) to restore a file share or disk from a specific recovery point.

### Steps

1. Select **SUPPORT** from the left navigation pane and select **Service Requests**.
2. Click **New Service Request**.
3. Select **Backup Restore** as the category for the service request.
4. Specify other details and proceed.

## Managing subscriptions

Managing subscriptions includes assigning service levels to your NetApp Keystone Flex Subscription (Flex Subscription) and tenant subscription services or allocating and changing capacities to the service levels by using the NetApp Service Engine interface.

The Flex Subscriptions represented by the relationship between the service providers/partners and NetApp and the tenant subscriptions represent the relationship between the service provider and their tenants and subtenants (end customers). When NetApp admins or GSSC creates Flex Subscription services for partners and tenants, it implies that the storage allocation has been accomplished. However, unless the service levels are created and the corresponding capacity allocation is complete, storage cannot be consumed by subtenants or end customers. Partner admins can manage storage for their Flex Subscription and tenant subscriptions. Tenant admins can only view the tenant subscription details and services for their own tenancy.

### Manage Flex Subscription (Keystone Subscriptions)

If you are a partner admin, you can view and manage your Flex Subscription services from **SUBSCRIPTIONS > Keystone Subscriptions**. **Keystone subscriptions** represent the Flex Subscription services created by NetApp admins for partners. Partner admins can view their Flex Subscription services and create service requests to:

- Increase the capacity of a subscription
- Add a service to a subscription (assign service level and capacity)

Service requests are received and addressed by the support team.

You can also view and manage the data protection and advanced protection services that you are subscribed to, in separate tabs.

### **Increase the capacity of a subscription**

You can increase the committed capacity for any of the subscribed service levels to meet growth requirements.

#### **Steps**

1. Select the subscription that you to want to increase capacity for from the drop-down menu.
2. Click **Increase Capacity** for the service level that you want.
3. Specify the capacity (in TiB) required.
4. Click **Confirm**.

### **Add a service to a subscription**

You can add a service to a subscription. Adding a service includes assigning a service level to your subscription and specifying the capacity for the same service level.

A subscription might only have one of each performance level, for example, if your subscription already contains an Extreme and a Standard service, you can add a Premium or Block service.

#### **Steps**

1. Select the subscription that you to want to increase capacity for from the drop-down menu.
2. Click **Add service**.
3. Select the subscription service that you want to add from the drop-down menu and add the capacity (in TiB) that you want to commit to the service level.
4. Click **Confirm**.



If you are subscribed to extreme-tiering or premium-tiering, you can edit the performance level of any existing subscription and opt for a tiering-enabled performance level. You can also opt to move from a tiering-enabled performance level to a non-tiering performance level.

## **Manage tenant subscriptions**

If you are a partner admin, you can create, view, and update your tenant subscriptions from **SUBSCRIPTIONS > Tenant Subscriptions**. If you are a tenant admin, you can only view the subscriptions.

### **Create a tenant subscription**

You can create new subscriptions for your tenants.

#### **Steps**

1. Click **New Subscription**.
2. Specify the region, zone, subscription term, rate plans, and other details.

The tenant subscription is assigned to a Keystone Subscription based on the zone selected. This is because each Keystone Subscription is tied to a zone.



You cannot create overlapping subscriptions. The start date should be at least the next day as subscriptions start a day ahead of the specified start date. Additionally, it cannot be more than 30 days from the current date.

3. Click **Create**.

You cannot delete any subscriptions. You can only edit existing tenant subscriptions and add new service levels, increase the committed capacity for a service level, and add tags to subscriptions.

## Manage service requests

### Overview

Place a service request for:

- Backup restoration
  - Disaster recovery failover
  - A technical issue with NetApp Service Engine
  - Any other issue that is not covered in this list
- This section describes:
- How to raise a service request
  - How to track a service request
  - The service request process

### Raise a service request

The Service Requests screen enables you to raise new service requests to be taken up by your own support team or NetApp GSSC.

You can view the summary of the most recent service requests, sorted by priority and creation time, on the **Service Requests** widget on the dashboard. You can also navigate to the Service Requests screen to search and view the details of all the service requests raised in your environment.

For creating a new service request, follow this procedure.

#### Steps

1. From the menu, select **SUPPORT > Service Requests** and click **New Service Request**.
2. On the New Service Request page, select the category and priority for the request, and then click **Next**.
3. The fields on the Details screen change based on the category that you selected. Enter the required details. Click **Next**.
4. Review the request details. If you want to correct anything, use the Back button to return to the previous page and correct the details. If all the details are correct, click **Submit**.
5. The service request is created and listed on the Service Requests screen. A newly-opened service request has a **New** status.

An email is sent to the requestor indicating the request has been created. The email is sent to the email

addressed registered for the requesting NetApp Service Engine user.

## Track a service request

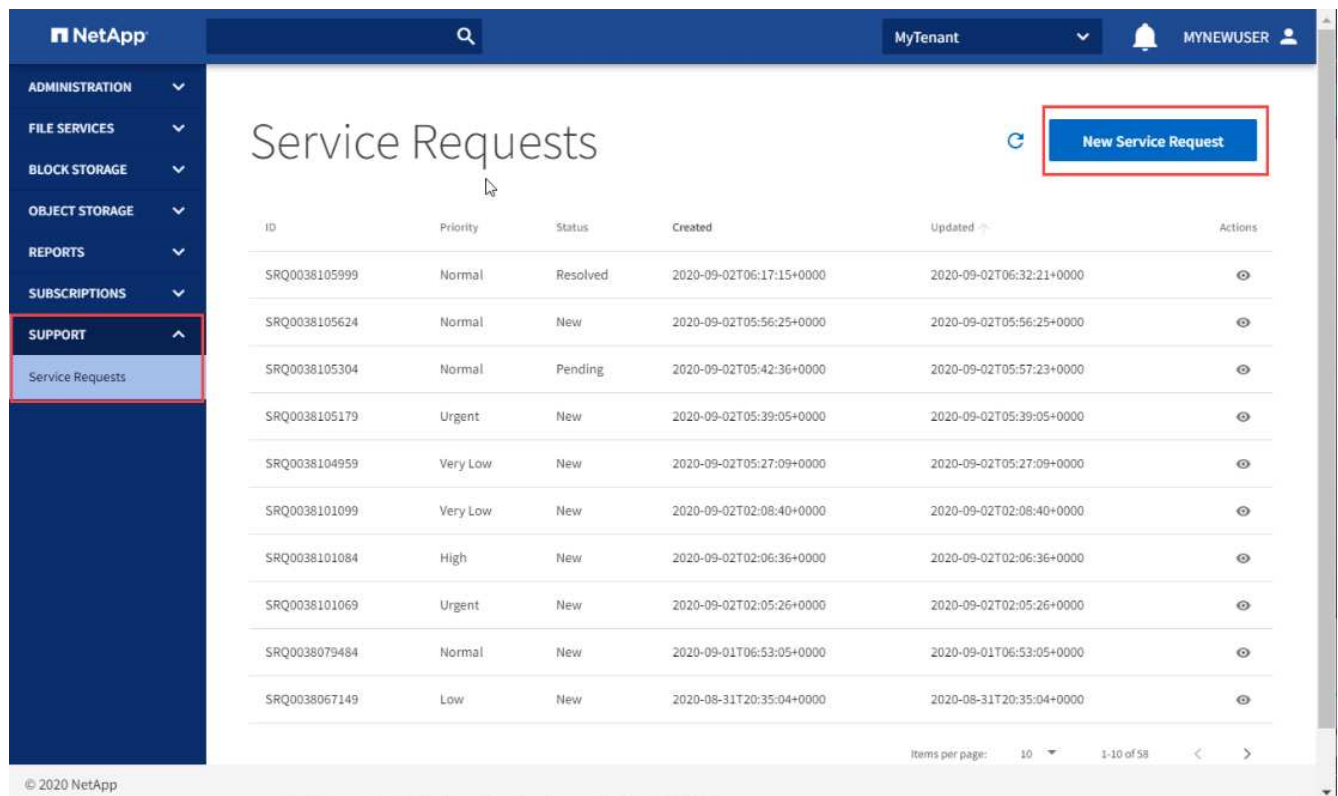
A service request follows the process described in this section.

The following table provides a list of the service request statuses.

Status	Description
New	This is a newly created request.
Open	The request is under review/fulfilment.
Pending	The request is waiting on action from a third party (i.e. waiting for additional information).
Resolved	The request has been resolved.
On hold	The request is waiting for an activity that is outside the control of support and can take some time before the request can be addressed.
Closed	The request is closed.

You can follow the status of a service request by reviewing the status in NetApp Service Engine, as described below, or you can follow the status in the automated emails sent as part of the process.

1. From the menu, select Support > Service requests. The Service Requests list is displayed.



The screenshot displays the NetApp Service Engine interface. The left-hand navigation menu is expanded to show the 'SUPPORT' section, with 'Service Requests' highlighted. The main content area is titled 'Service Requests' and features a 'New Service Request' button in the top right corner. Below the title is a table listing service requests with columns for ID, Priority, Status, Created, Updated, and Actions. The table contains 10 rows of data, including requests with statuses like Resolved, New, Pending, and Urgent.

ID	Priority	Status	Created	Updated	Actions
SRQ0038105999	Normal	Resolved	2020-09-02T06:17:15+0000	2020-09-02T06:32:21+0000	👁
SRQ0038105624	Normal	New	2020-09-02T05:56:25+0000	2020-09-02T05:56:25+0000	👁
SRQ0038105304	Normal	Pending	2020-09-02T05:42:36+0000	2020-09-02T05:57:23+0000	👁
SRQ0038105179	Urgent	New	2020-09-02T05:39:05+0000	2020-09-02T05:39:05+0000	👁
SRQ0038104959	Very Low	New	2020-09-02T05:27:09+0000	2020-09-02T05:27:09+0000	👁
SRQ0038101099	Very Low	New	2020-09-02T02:08:40+0000	2020-09-02T02:08:40+0000	👁
SRQ0038101084	High	New	2020-09-02T02:06:36+0000	2020-09-02T02:06:36+0000	👁
SRQ0038101069	Urgent	New	2020-09-02T02:05:26+0000	2020-09-02T02:05:26+0000	👁
SRQ0038079484	Normal	New	2020-09-01T06:53:05+0000	2020-09-01T06:53:05+0000	👁
SRQ0038067149	Low	New	2020-08-31T20:35:04+0000	2020-08-31T20:35:04+0000	👁

2. Locate the service request in the list and view the status in the Status column.



NetApp MyTenant MYNEWUSER

## Service Requests

[New Service Request](#)

ID	Priority	Status	Created	Updated	Actions
SRQ0038107664	Normal	New	2020-09-02T07:46:54+0000	2020-09-02T07:46:54+0000	👁
SRQ0038105999	Normal	Resolved	2020-09-02T06:17:15+0000	2020-09-02T06:32:21+0000	👁
SRQ0038105624	Normal	New	2020-09-02T05:56:25+0000	2020-09-02T05:56:25+0000	👁
SRQ0038105304	Normal	Pending	2020-09-02T05:42:36+0000	2020-09-02T05:57:23+0000	👁
SRQ0038105179	Urgent	New	2020-09-02T05:39:05+0000	2020-09-02T05:39:05+0000	👁
SRQ0038104959	Very Low	New	2020-09-02T05:27:09+0000	2020-09-02T05:27:09+0000	👁
SRQ0038101099	Very Low	New	2020-09-02T02:08:40+0000	2020-09-02T02:08:40+0000	👁
SRQ0038101084	High	New	2020-09-02T02:06:36+0000	2020-09-02T02:06:36+0000	👁
SRQ0038101069	Urgent	New	2020-09-02T02:05:26+0000	2020-09-02T02:05:26+0000	👁
SRQ0038079484	Normal	New	2020-09-01T06:53:05+0000	2020-09-01T06:53:05+0000	👁

Items per page: 10 1-10 of 59


Created service request with ID SRQ0038107664

## Service request process

This section describes a simplified service request process flow.

The figure below shows the flow as a diagram, and the table below it describes the actors, actions, and status of the service request as a result of the action.

## High Level Service Request Flow



	Who	Action	Status of Request in NetApp Service Engine after action	Notification
1	NetApp Service Engine	Generate a new service request.	New	Automated email confirming request has been created
2	NetApp GSSC	Acknowledge the request.	Open	<ul style="list-style-type: none"> <li>Acknowledgement email: GSSC</li> <li>Automated email indicating status in Open</li> </ul>
3	NetApp GSSC	Triage and assign to the Fulfilment Engineer.	Open	–
4	NetApp GSSC	Monitor and track the request throughout the request fulfilment process.	–	–
5	NetApp Fulfilment Engineer	Review the request.	Open	–

	Who	Action	Status of Request in NetApp Service Engine after action	Notification
6	NetApp Fulfilment Engineer	If additional information is required, email the requestor for additional information.	Pending	<ul style="list-style-type: none"> <li>Automated email indicating status is Pending.</li> <li>Communicate via email to request additional information</li> </ul>
7	Customer	Provide more information.	Pending	Communicate via email
8		Repeat steps 6 and 7 until there is sufficient information to progress the request.	Pending	Communicate via email to request additional information
9	NetApp Fulfilment Engineer	Fulfil request	Open	Automated email indicating status is Open.
10	NetApp Fulfilment Engineer	Resolve request and advise customer request is resolved.	Resolved	<ul style="list-style-type: none"> <li>Email to request confirmation Service Request is resolved</li> <li>Automated email indicating status is Resolved.</li> </ul>
11	Customer	Advise Service Request is resolved.	Resolved	Email to confirm Service Request is resolved.
12	Automatic	If there is no further action required after three day , the Service Request is automatically closed.	Closed	Automated email indicating status is Closed.

## Perform administrative tasks

### Manage users

#### View a list of users

The Users list displays a list of all the users in the tenancies that the logged-in user has access to. To view the list, select Administration > Users from the menu.

The list displays simple information about each user such as the user name, first and last name, email address, primary tenancy, and date created.

The action icons next to each user allow you to modify or delete the user.

## Add a user

This section describes how to add a user.

### Before you begin

You need:

- The tenants to which the user should be given access.
- The role that the user is to be granted within each tenancy. A user can have only one role within a tenancy.

### Steps

1. Select **ADMINISTRATION > Users** from the menu.
2. Click **Create User**.
3. In the Create User dialog box, specify the user name, email address, first name, last name, display name, and password for the user.
4. Select the primary tenant and the role of the user in that tenancy.



You can assign privileges to a user based on their role, such as a NetApp administrator, NetApp administrator with read only privileges, partner administrator, or tenant administrator. Based on your role, you can view the roles of different users within your environment and assign roles to newly created or existing users.

5. To add additional tenancies, click **Add Tenancy** to display new entry fields and select the tenant and the role of the user in that tenancy.
6. Click **Create**.

## Modify user details

Use this method to modify a user's details. You can:

- Modify the user's first name, last name, and display name.
- Add or remove tenancies to which the user has access.
- Modify the role of the user in each tenancy to which the user has access.

### Steps

1. Select **ADMINISTRATION > Users**.
2. Locate the user in the list and click the Edit icon.
3. On the Edit User dialog box, make the changes as required.
4. To change the role of the user in a tenancy, locate the tenancy in the Tenancies list and select the new role.
5. To grant the user access to another tenancy, click **Add Tenancy** and select the Tenant and the Role of the user in that tenancy in the new entry fields.
6. To remove access for the user to a tenancy, locate the tenancy in the Tenancies list and select the Delete icon.

7. Click **Update**.

## Delete a user

Use this method to delete a user.



It is not possible to recover a deleted user.

### Steps

1. Select Administration > Users from the menu.
2. Locate the user in the list and click the Delete icon.
3. In the Confirm Delete dialog box, enter the name of the user to be deleted.
4. Click **Confirm**.

## Managing tenants and subtenants

You can view and manage your tenants and subtenants by using NetApp Service Engine. The initial tenancy is set up by NetApp support for service providers/partners and tenants. The region, zone, and subscription are already configured by support, and are available when you create tenants and subtenants under your tenancy.

### Tenants

If you are a partner admin, you can create, modify, and delete a tenant from **ADMINISTRATION > Tenants**. Tenant admins can only view their respective tenant subscriptions and subtenants. For the tenants under your partner tenancy, you can also define the network.

#### Create a tenant

You can create a tenant if you are a partner admin. You can also create and manage networks for each tenant. For more information on creating networks, see [Define network configurations with subnets](#).

### Steps

1. Click **Create Tenant**.
2. Add a name and the code.  
The code that represents the tenant should be unique within a NetApp Service Engine instance and consist of lower-case characters.
3. Click **Create**.



You can only modify the name and description of an existing tenant. Additionally, you cannot delete a tenant that has subtenants or storage components associated with it.

### Subtenants

If you are a Partner admin or a Tenant admin, you can view and manage subtenants.

#### Create a subtenant

If you are a Partner admin or a Tenant admin, you can create, modify, and delete a subtenant from

**ADMINISTRATION > Subtenants.** You can also create and manage networks for each subtenant.

For more information on creating networks, see [Define network configurations with subnets](#).

### Steps

1. Click **Create Subtenant**.
2. Add a name and the code.  
The code that represents the subtenant should be unique within a NetApp Service Engine instance and consist of lower-case characters.
3. Click **Create**.



You can only modify the name of an existing subtenant. Additionally, you cannot delete a subtenant that has storage components associated with it.

## Create and manage alerts

Alerts are messages that are either triggered automatically or created manually to share information about critical events within your environment.

When you provision storage, automatic alerts are triggered in the following scenarios:

- The new disk or file share pushes the subscription into burst
- The new disk or file share is provisioned on a new subscription
- The consumed capacity within the subscription has crossed the threshold capacity, or is close to the committed capacity

The **Alerts** screen enables you to view the system-generated and user-generated alerts in your environment.

You can also create custom alerts and display them to other users. They can view and dismiss the alerts for themselves and other users within their tenancy, as required.

### Create an alert in your multitenant environment

If you are a partner admin, you can manually create and send alerts to a single tenant or all the tenants in your environment.



You can dismiss an alert that you created. However, it is dismissed only for tenants who have not viewed it yet.

### Steps

1. Select **ADMINISTRATION > Alerts** from the menu.
2. Click **Create alert**.
3. On **Create alert** screen, select the status and the tenant that you want to send the alert to, and add the message.



If you want to send the alert to all the tenants in your environment, select "All".

4. Click **Create**.  
The new alert is created.

## Create an alert in your tenancy

If you are a tenant admin, you can manually create and send alerts to all the users in your environment.



You can dismiss an alert that you created. However, it is dismissed only for users who have not viewed it yet.

### Steps

1. Select **ADMINISTRATION > Alerts** from the menu.
2. Click **Create alert**.
3. On **Create alert** screen, select the status and add the message.
4. Click **Create**.  
The new alert is created.

## Define network configurations for tenants and subtenants

A service provider or partner admin can define network configurations for tenants based on their subscription zones and assign VLANs to each network configuration. They can also create a subnet with each VLAN belonging to the same zone and tenancy, and assign it to a subtenant under that tenancy.

A preconfigured tenant network is required creating the subnet. A subnet is required when a subtenant or end customer provisions file servers or block stores. File and block storage VMs are owned by subtenants in NetApp Service Engine, whereas they are owned by IPspaces in ONTAP. When you define the networking boundary at a subtenant level, the subtenant in NetApp Service Engine maps to a single ONTAP IPspace for that zone. This eliminates risk and ensures that two distinct subtenants do not select the same VLAN ID and end up in the same IPspace.

- There can be multiple subnets created in each zone. For each zone under a tenancy, there is an allocated list of VLANs. For all those available VLANs, corresponding subnets can be created. There is a 1:1 mapping between a subnet and a VLAN. Whereas, there is a one to many mapping between a zone and subnets.
- One tenant can have multiple subnets, one for each available VLAN under a zone. They are displayed under a single network configuration for each zone.
- There can be only one VLAN for each network subnet. Each network configuration shows a list of available VLANs.
- Only one subnet can be created for a zone and VLAN combination for a subtenant.



A tenant admin can only view the network configurations for their tenancy and the subnets for their subtenants.

## Define a network configuration for a tenant

As a partner admin, you can create network configurations for each tenant. A configurable and unique VLAN pool can be allocated to each tenant based on their requirement. When a tenant is deleted, the VLANs associated with it are released back to the tenant VLAN pool.

### Before you begin

- Ensure the tenants have been created
- You should know the zone, VLANs, and port override configuration

### Steps

1. From **ADMINISTRATION > Tenants**, select the required tenant.
2. Click the network icon. The Networks screen is displayed with the list of network configurations for the tenant.
3. Click **Create Network**. Specify the zone, VLAN, and port override, if applicable.



The port override is configured by the support team during zone configuration. If a port override is not specified, the network uses the default port and creates a new VLAN in the IPspace for that tenant.

4. Click **Create**. The new network configuration is listed for the tenant on the Networks screen.

If the network configuration for a zone is created, you can only assign new VLANs to the network configuration. Click the **+** icon and add the VLAN and port override details. Click **Update**.

## Create a subnet for a subtenant

As a partner admin, you can create subnets for your subtenants from **NETWORKS > Subnets**. The subnet definition includes the VLAN ID that is implemented in ONTAP, the CIDR address range, and a route specification. A subnet is aligned to a VLAN within a tenant and zone. You cannot create two subnets with the same VLAN for a zone.

### Before you begin

- The VLAN and zone configuration should already have been complete for the tenant
- You should be aware of the requirements for the subnet to be created

### Steps

1. Select the required tenant. The subtenants for whom you create the subnets should be under this tenant.
2. Click **NETWORKS > Subnets**. You can view the list of subnets created for your subtenants on the **Subnets** page.
3. Click **Create Subnet**.
4. Specify the subtenant, zone, name of the subnet, VLAN, and Subnet CIDR. You can also add route specification and tags.



In route specification, the **Destination** should be unique and part of the mentioned VLAN, and multiple routes can have the same value for **Metric**.

5. Click **Create**.



You cannot edit or delete a subnet. You can only edit the tags associated with it.



# Overview of NetApp Service Engine APIs

NetApp Service Engine provides a set of APIs to manage storage resources on the supported storage systems through a RESTful web service interface for any third-party integration. The APIs cover all the functionalities that are supported through the GUI. NetApp provides the API documentation, but it is the customer's responsibility to integrate into their applications.

This NetApp Service Engine API Reference Guide provides you with information about APIs and sample codes. The information provided in the guide enables you to create RESTful clients of NetApp Service Engine. The APIs are based on the Representational State Transfer (REST) architectural style.

## Target audience

This guide is intended for developers creating applications that interface with the NetApp Service Engine software through REST APIs.

You should use this guide if you want to use the storage provider, NetApp ONTAP cluster, and management administration APIs for managing your storage.

## NetApp Service Engine API access and categories

### Constructing a URL to directly access REST APIs

You can access the REST APIs directly through a programming language, such as Python, C#, C++, JavaScript, and so on. To access the REST APIs in the `https:// <hostname> /api` format, enter the host name or IP address and the URL.

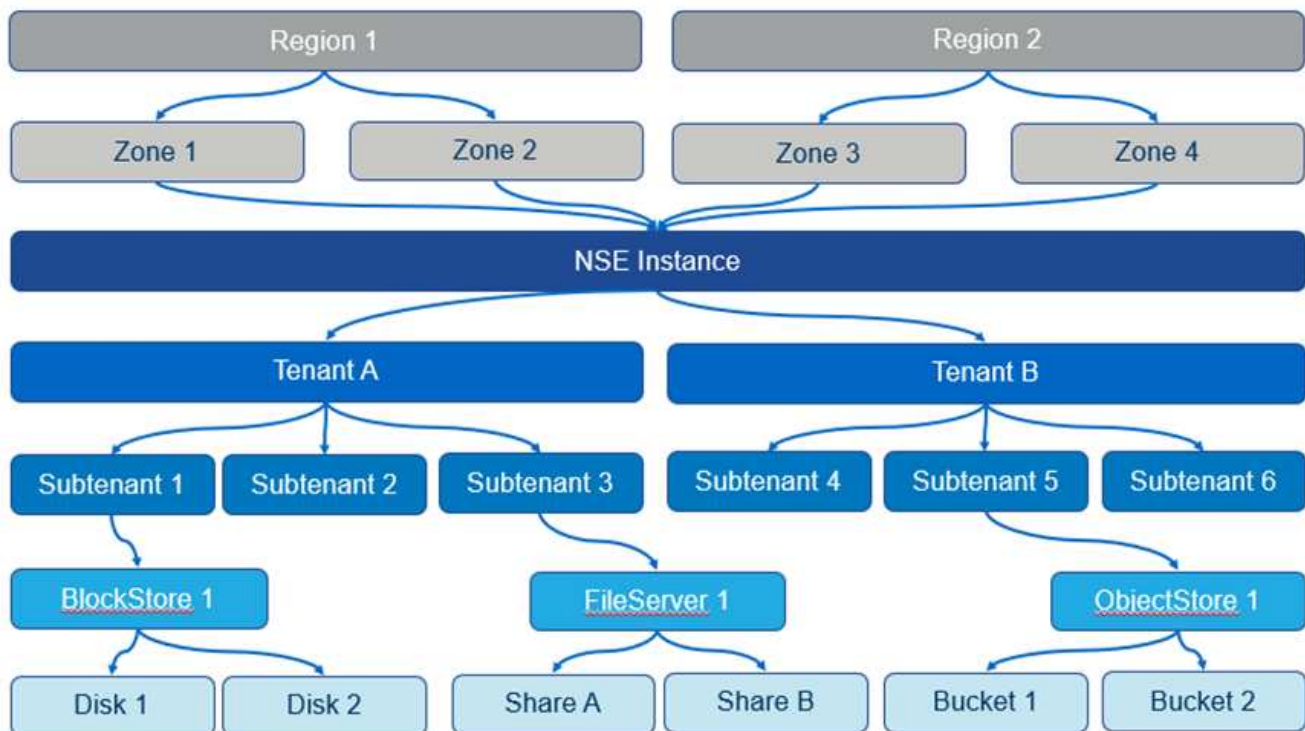
### Consumer and administrator APIs

This document splits the API descriptions into two sections consisting of:

- Consumer APIs
- Administrator APIs

## Key NetApp Service Engine concepts

NetApp Service Engine supports the concepts of regions and zones, as shown in figure below. A region represents a data center or site, while a zone represents a storage subunit within the region (technically, a cluster within a data center/site). Multiple zones support data availability and data protection (DP) features.



A single NetApp Service Engine instance can support one or more tenants. NetApp Service Engine uses the concepts of tenant and subtenant as hierarchical entities that own or manage the storage service.

A tenant can be a customer, partner, or a billing entity. A tenant holds the subscription (or multiple subscriptions) to the Flex Subscription service.

A subtenant is an entity wholly within the tenant. It can be used for showback, security separation, and so on.

Storage items are created per subtenant. Each subtenant can hold multiple storage items of the type and number suitable for that subtenant.

NetApp Service Engine supports the following types of storage:

- Block storage in block stores and disks
- File storage in file servers and file shares
- Object storage in buckets

## Authorization and authentication

The NetApp Service Engine REST API is accessible by using any web browser or programming platform that can issue HTTP requests. NetApp Service Engine supports basic HTTP authentication mechanism as well as Java Web Token (JWT) Bearer authentication. Before you call an API, you must authenticate a user.

### HTTP status codes

While running the APIs or troubleshooting issues, you should be aware of the various HTTP status codes and error codes that are used by NetApp Service Engine APIs.

The following table lists the error codes related to authentication.

HTTP Status Code	Status Code Title	Description
200	OK	Returned on successful execution of synchronous API calls.
201	Created	Creation of new resources by synchronous calls, such as configuration of Active Directory.
202	Accepted	Returned on successful execution of asynchronous calls for provisioning functions, such as creating LUNs and files shares.
400	Invalid Request	Indicates input validation failure. User must correct the inputs, for example, valid keys in a request body.
401	Unauthorized Request	You are not authorized to view the resource/unauthorized.
403	Forbidden Request	Accessing the resource you were trying to reach is forbidden.
404	Resource Not Found	The resource you were trying to reach is not found.
405	Method Not Allowed	Method not allowed.
412	Precondition Failed	One or more preconditions were not met.
500	Internal Server Error	Internal server error. Failed to get a response from the server. This internal server error might or might not be permanent. For example, if you run a GET or GET ALL operation and receive this error, it is recommended that you repeat this operation for a minimum of five retries. If it is a permanent error, then the status code returned continues to be 500. If the operation succeeds, the status code returned is 200.

## Recommendations for using the APIs

When using the APIs, you should follow certain recommended practices:

- For a valid execution, make sure all response content type are in the following format:

```
application/json
```

- While updating array values, you must update the entire string of values. You cannot append values to an array—you can only replace an existing array.
- Avoid querying objects by using a combination of the filter operators wildcard (\*) and pipe (|). It might retrieve an incorrect number of objects.

## NetApp Service Engine REST APIs

Many API calls return a large number of records. It is possible to limit the number of records by specifying the `offset` and `limit` values:

- **Limit.** The number of records to return.
- **Offset.** The number of records to skip before returning the first record. For example, an offset of one will skip the first record and return records starting at two.

You can also sort and order the responses:

- **Sort\_by.** Sort the records using the specified attribute. For example, to sort by the create date, specify `created`.
- **Order\_by.** Orders the records by ascending (`asc`) or descending (`desc`) order.

## Consumer APIs

### Overview

This section describes the following consumer APIs:

- Cloud orchestration APIs
- Block storage APIs
- File services APIs
- Object storage APIs
- Reporting APIs
- (Consumer) administration APIs

### Cloud Orchestration APIs

The Cloud Orchestration APIs NetApp Service Engine support connecting Cloud Volumes Services to NetApp Service Engine and to work with them from within the NetApp Service Engine. Azure NetApp Files (ANF) and Google Cloud Platform (GCP) are supported for Cloud Volumes Services.

Subscriptions to Cloud Volumes Services are managed outside of NetApp Service Engine. The relevant credentials are provided to NetApp Service Engine to allow connection to the cloud services.

This section describes the APIs for:

- Managing Cloud Volumes Service instances
- Managing Azure NetApp Files accounts

- Managing Azure NetApp Files capacity pools
- Managing Azure NetApp Files Volumes
- Managing Google Cloud Volumes

### Cloud Volumes Service instances

Use the methods listed in the following table to retrieve and create Cloud Volumes Service instances.

HTTP Method	Path	Description
GET	/v2.1/cvs/instances	Retrieve all Cloud Volumes Service instances.
POST	/v2.1/cvs/instances	Create new Cloud Volumes Service instances.
GET	/v2.1/cvs/instances/{id}	Retrieve a Cloud Volumes Service instance by ID.
GET	/v2.1/anf/accounts	Retrieve Azure NetApp Files accounts.
POST	/v2.1/anf/accounts	Create an Azure NetApp Files accounts
GET	/v2.1/anf/accounts/{id}	Retrieve an Azure NetApp Files account by ID.
GET	/v2.1/anf/pools	Retrieve all capacity pools.
POST	/v2.1/anf/pools`	Create a capacity pool.
GET	/v2.1/anf/pools/{id}	Retrieve an Azure NetApp Files capacity pool by ID.

### Cloud Volumes Service instance attributes

The following table lists the Cloud Volumes Service instance attributes.

Attribute	Type	Description
id	String	Unique identifier for the Cloud Volumes Service instance.
name	String	Name of the Cloud Volumes Service instance.
cc_working_env	String	Working environment name within Cloud Manager.
cc_client_id	String	Auth0 service ID.
cc_refresh_token	String	Auth0 refresh token.
cc_account_id	String	Cloud Central account ID.
cc_account_name	String	Cloud Central account name.

Attribute	Type	Description
cc_credentials_id	String	ID of the Cloud Centra working environment.
azure_client_id	String	Application (Client) ID. The ID of an Active Directory service principal that Cloud Manager can use to authenticate with Azure Active Directory.
azure_client_secret	String	The value of client secret for service principal application.
azure_tenant_id	String	The Azure Active Directory ID. This is also known as tenant ID.
azure_subscription_id	String	The Azure Active Directory subscription ID.
subtenant_id	String	The subtenant ID.
tags		The key-value pairs.

#### Retrieve Cloud Volumes Service instance

HTTP Method	Path	Description	Path Parameters
GET	/v2.1/cvs/instances	Retrieve Cloud Volumes Service instances.	tenant_id: (Optional) Return the Cloud Volumes Service instances belonging to the specified tenant. See also <a href="#">Common Pagination</a> parameters.

Required request body parameters: none

#### Retrieve Cloud Volumes Service instance by ID

Use the method listed in the following table to retrieve a Cloud Volumes Service instance by its identifier.

HTTP Method	Path	Description	Parameters
GET	/v2.1/cvs/instances/{id}	Retrieve a Cloud Volumes Service instance by ID.	id (string): The unique identifier of the Cloud Volumes Service instance.

Required request body attributes: none

#### Create Cloud Volumes Service instances

Use the method listed in the following table to create a new Cloud Volumes Service instance.

HTTP Method	Path	Description	Parameters
POST	/v2.1/cvs/instances	Create a Cloud Volumes Service instance.	None

Required request body attributes: name, cc\_working\_env, cc\_client\_id, cc\_refresh\_token, cc\_account\_id, cc\_account\_name, azure\_client\_id, azure\_client\_secret, azure\_tenant\_id, azure\_subscription\_id, subtenant\_id

#### Request body example:

```
{
  "name": "instance1",
  "cc_working_env": "my-working-env",
  "cc_client_id": "Mu0V1ywgYteI6w1MbD15fKfVIUrNXGWC",
  "cc_refresh_token": "y1tMw3lNzE8JL9jtiE29oSRxOAZyu0cdnwS_2XhjQBr9G",
  "cc_account_id": "account-335jdf32",
  "cc_account_name": "my-account-name",
  "cc_credentials_id": "d336c449-aeb8-4bb3-af28-5b886c40dd00",
  "azure_client_id": "53ba6f2b-6d52-4f5c-8ae0-7adc20808854",
  "azure_client_secret": "NMubGVcDqkwwGnCs6fa01tqlkTisfUd4pBBYgcxxx=",
  "azure_tenant_id": "53ba6f2b-6d52-4f5c-8ae0-7adc20808854",
  "azure_subscription_id": "1933a261-d141-4c68-9d6c-13b607790910",
  "subtenant_id": "5d2fb0fb4f47df00015274e3",
  "tags": {
    "key1": "Value 1",
    "key2": "Value 2",
    "key3": "Value 3",
    "keyN": "Value N"
  }
}
```

#### Manage tags for Cloud Volumes Service instances

Use the method listed in the following table to specify tags for the named Cloud Volumes Service instance.

HTTP Method	Path	Description	Parameters
POST	/v2.1/cvs/instances/{id}/tags	Manage tags for a Cloud Volumes Service instance.	id (string): The unique identifier of the Cloud Volumes Service instance.

Required request body attributes: key-value pairs

#### Request body example:

```
{
  "env": "test"
}
```

## Azure NetApp Files accounts

### Azure NetApp Files accounts attributes

The following table lists the Azure NetApp Files account attributes.

Attribute	Type	Description
id	String	The unique identifier for the Azure NetApp Files account.
name	String	The name of the Azure NetApp Files account.
resource_group	String	The Azure resource group.
location	String	The Azure location (region/zone).
cvs_instance_id	String	The Cloud Volumes Service instance identifier.
tags	–	The key-value pairs.

### Retrieve Azure NetApp Files accounts

HTTP Method	Path	Description	Path Parameters
GET	/v2.1/anf/accounts	Retrieve Azure NetApp Files accounts.	subtenant_id: (Mandatory) The subtenant ID to which the Azure NetApp Files account belongs. tenant_id: (Optional) Returns the Azure NetApp Files accounts belonging to the specified tenant. See also <a href="#">Common Pagination</a> parameters.

Required request body parameters: none

### Retrieve Azure NetApp Files account by name

Use the method listed in the following table to retrieve an Azure NetApp Files account by name.



HTTP Method	Path	Description	Parameters
GET	/v2.1/anf/accounts/{name}	Retrieve an Azure NetApp Files account by name.	name (string): (Mandatory) The name of the Azure NetApp Files account. subtenant_id (string): (Mandatory) The subtenant ID to which the Azure NetApp Files account belongs.

Required request body attributes: none

### Create Azure NetApp Files accounts

Use the method listed in the following table to create a new Azure NetApp Files account.

HTTP Method	Path	Description	Parameters
POST	/v2.1/anf/accounts	Create a new Azure NetApp Files account.	None

Required request body attributes: name, resource\_group, location, cvs\_instance\_id

### Request body example:

```
{
  "name": "string",
  "resource_group": "string",
  "location": "string",
  "cvs_instance_id": "5d2fb0fb4f47df00015274e3",
  "tags": {
    "key1": "Value 1",
    "key2": "Value 2",
    "key3": "Value 3",
    "keyN": "Value N"
  }
}
```

### Azure NetApp Files capacity pools

#### Capacity pools attributes

The following table lists the capacity pool attributes.

Attribute	Type	Description
id	String	The unique identifier for the capacity pool.

Attribute	Type	Description
name	String	The name of the capacity pool.
resource_group	String	The Azure resource group.
location	String	The Azure location (region/zone).
size	Integer	The size of the capacity pool in TB.
service_level	String	The service level name applicable: Ultra, Premium, or Standard.
anf_account_name	String	The Azure NetApp Files account instance identifier.
subtenant_id	String	The subtenant ID.
tags	–	The key-value pairs.

### Retrieve capacity pools

HTTP Method	Path	Description	Path Parameters
GET	/v2.1/anf/pools	Retrieve capacity pools.	<p>subtenant_id: (Mandatory) The subtenant ID to which the ANF account belongs.</p> <p>tenant_id: (Optional) Return the capacity pools belonging to the specified tenant.</p> <p>See also <a href="#">Common Pagination</a> parameters.</p>

Required request body parameters: none

### Request body example:

```
none
```

### Retrieve capacity pool by name

Use the method listed in the following table to retrieve a capacity pool by name.

HTTP Method	Path	Description	Parameters
GET	/v2.1/anf/pools/{name}	Retrieve a capacity pool by name.	name (string): (Mandatory) The unique name of the capacity pool.  subtenant_id (string): (Mandatory) The subtenant ID to which the capacity pool belongs.

Required request body attributes: none

### Create capacity pools

Use the method listed in the following table to create a new capacity pool.

HTTP Method	Path	Description	Parameters
POST	/v2.1/anf/pools	Create a capacity pool.	None

Required request body attributes: name, resource\_group, location, size, service\_level, anf\_account\_name, subtenant\_id

### Request body example:

```
{
  "name": "string",
  "resource_group": "string",
  "location": "string",
  "size": 10,
  "service_level": "Standard",
  "anf_account_name": "myaccount",
  "subtenant_id": "5d2fb0fb4f47df00015274e3",
  "tags": {
    "key1": "Value 1",
    "key2": "Value 2",
    "key3": "Value 3",
    "keyN": "Value N"
  }
}
```

### Modify size of the capacity pool

Use the method listed in the following table to modify the size of the capacity pool.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/anf/pools/{name}	Modify the size of the capacity pool.	name (string): Mandatory: the unique name of the capacity pool.

Required request body attributes: name, resource\_group, location, anf\_account\_name, size, service\_level, subtenant\_id

#### Request body example:

```
{
  "name": "myaccount",
  "resource_group": "string",
  "location": "string",
  "anf_account_name": "myaccount",
  "size": 4,
  "service_level": "Standard",
  "subtenant_id": "5d2fb0fb4f47df00015274e3",
  "tags": {
    "key1": "Value 1",
    "key2": "Value 2",
    "key3": "Value 3",
    "keyN": "Value N"
  }
}
```

## Azure NetApp Files volumes

### Azure NetApp Files volume attributes

The following table lists the Azure NetApp Files volume attributes.

Attribute	Type	Description
id	String	The unique identifier for the Azure NetApp Files volume.
name	String	The name of the Azure NetApp Files volume.
resource_group	String	The Azure resource group.
subtenant_id	String	The subtenant ID.
anf_account_name	String	The Azure NetApp Files account name.
anf_pool_name	String	The Azure NetApp Files Pool name.
location	String	The Azure location (region/zone).

Attribute	Type	Description
file_path	String	Creation Token or File Path. A unique file path for accessing volume.
quota_size	Integer	Maximum storage quota allowed in GiB.
subNetID	String	The Azure Resource URL for a delegated subnet. Must have the delegation Microsoft NetApp/volumes.
tags	—	The key-value pairs.

### Retrieve Azure NetApp Files volumes

Use the method listed in the following table to retrieve Azure NetApp Files volumes. Specifying a `tenant_id` returns only the accounts belonging to that tenant.

HTTP Method	Path	Description	Path Parameters
GET	<code>/v2.1/anf/volumes</code>	Retrieve Azure NetApp Files volumes.	<p><code>subtenant_id</code>: (Mandatory) The subtenant ID to which the ANF volume belongs.</p> <p><code>tenant_id</code>: (Optional) Return the ANF Volumes belonging to the specified tenant.</p> <p>See also <a href="#">Common Pagination</a> parameters.</p>

Required request body parameters: none.

### Retrieve Azure NetApp Files volume by name

Use the method listed in the following table to retrieve an Azure NetApp Files volume by name.

HTTP Method	Path	Description	Parameters
GET	<code>/v2.1/anf/volumes/{name}</code>	Retrieve an Azure NetApp Files volume by name.	<p><code>name</code> (string): Mandatory: the unique name of the Azure NetApp Files volume.</p> <p><code>subtenant_id</code>: (string) Mandatory. The subtenant ID to which the Azure NetApp Files volume belongs.</p>

Required request body attributes: none

### Request body example:

```
none
```

### Create Azure NetApp Files volumes

Use the method listed in the following table to create a new Azure NetApp Files volume.

HTTP Method	Path	Description	Parameters
POST	/v2.1/anf/volumes	Create an Azure NetApp Files volume.	None

Required request body attributes: *name*, *resource\_group*, *subtenant\_id*, *anf\_account\_name*, *anf\_pool\_name*, *virtual\_network*, *location*, *file\_path*, *quota\_size*, *subNetID*

### Request body example:

```
{
  "name": "myVolume",
  "resource_group": "string",
  "subtenant_id": "5d2fb0fb4f47df00015274e3",
  "anf_account_name": "myaccount",
  "anf_pool_name": "myaccount",
  "virtual_network": "anf-vnet",
  "location": "string",
  "file_path": "myVolume",
  "quota_size": 100,
  "subNetId": "string",
  "protocol_types": [
    "string"
  ],
  "tags": {
    "key1": "Value 1",
    "key2": "Value 2",
    "key3": "Value 3",
    "keyN": "Value N"
  }
}
```

### Managing Cloud Volumes Service for Google Cloud

The `/v2.1/gcp/volumes` API under the Cloud Orchestration category enables you to manage cloud volumes for your Google Cloud instance. Before running this API, ensure that the Cloud Volumes Service account for Google Cloud Platform (GCP) subscription has been enabled for the subtenant.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/gcp/volumes	You can use the GET method to retrieve the details of all the Google Cloud Volumes created for your subtenant's Cloud Volumes Service subscription.	<p>offset: The number of items to skip before starting to collect the result set.</p> <p>limit: The numbers of items to return.</p> <p>subtenant_id: The ID of the subtenant subscribed to Google Cloud.</p> <p>region: The region of the subscribed service.</p>
GET	/v2.1/gcp/volumes/{id}	You can use this method to retrieve the details of a specific Google Cloud Volume created for your subtenant's Cloud Volumes Service subscription.	<p>id: The ID of the GCP volume.</p> <p>subtenant_id: The ID of the subtenant subscribed to Google Cloud.</p> <p>region: The region of the subscribed service.</p>

HTTP Verb	Path	Description	Mandatory parameters/Request body
POST	/v2.1/gcp/volumes	Create a GCP volume for a subtenant. Add the values in the request body to create a volume with the specified parameters.	<pre> ` {   "subtenant_id":   "&lt;ID&gt;",   "name":   "&lt;Volume_name&gt;",   "region":   "&lt;region&gt;",   "zone": "&lt;zone&gt;",   "creation_token":   "&lt;token&gt;",   "allowed_clients":   "&lt;IP address of the   clients allowed to   access GCP&gt;",   "network":   "&lt;network details   as entered for the   GCP subscribed   service&gt;",   "protocol_types": [   "&lt;Protocol for the   connection, such as   NFSv3&gt;"   ],   "quota_gib":   &lt;volume quota in   bytes&gt;,   "service_level":   "&lt;the type of   Performance Service   Level, such as   standard&gt;",   "labels":   ["&lt;tag_value&gt;"] } ` </pre>



HTTP Verb	Path	Description	Mandatory parameters/Request body
PUT	/v2.1/gcp/volumes/{id}	Modify a GCP volume already created for a subtenant. Add the volume ID of the volume that you want to modify and the value for the parameters that you want to modify, in the request body.	<pre> {   "subtenant_id":   "&lt;ID&gt;",   "name":   "&lt;Volume_name&gt;",   "region":   "&lt;region&gt;",   "zone": "&lt;zone&gt;",   "allowed_clients":   "&lt;IP address of the clients allowed to access GCP&gt;",   "quota_gib":   &lt;volume quota in bytes&gt;,   "service_level":   "&lt;the type of Performance Service Level, such as standard&gt;",   "protocol_types":   [&lt;Protocol for the connection, such as NFSv3&gt;"],   "labels":   [&lt;tag_value&gt;"] } </pre>
DELETE	/v2.1/gcp/volumes/{id}	You can use this method to delete a specific Google Cloud Volume created for your subtenant's Cloud Volumes Service subscription.	<p>id: The ID of the GCP volume.</p> <p>subtenant_id: The ID of the subtenant subscribed to Cloud Volumes Service for Google Cloud.</p> <p>region: The region of the subscribed service.</p>

## Block storage APIs

You can use Block Storage APIs to view and manage your block storage.

Before you create disks, you have to create block stores. When you create a disk, you have to create a host group or select an existing host group to access the disk. For more information, see [Work with block storage](#).

## Block stores

You can use Block Store APIs to retrieve and manage your block stores.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/blockstores	You can retrieve the details of all your block stores. Retrieves details of the block stores, such as block store ID, IP address, region, zone, subnet ID, and tags.	<code>offset</code> : The number of items to skip before starting to collect the result set. <code>limit</code> : The numbers of items to return.
GET	/v2.1/blockstores/{id}	You can retrieve the details of a specific block store. Retrieves details of the block store, such as IP address, region, zone, subnet ID, and tags based on the entered ID.	<code>id</code> : The ID of the block store.
POST	/v2.1/blockstores	You can create a block store. Add the values in the request body to create a block store with the specified parameters, such as service protocol, subtenant ID, zone, subnet ID, and tags.	<code>subtenant_id</code> : The ID of the subtenant. <code>zone</code> : The name of the zone. <code>subnet_id</code> : The ID of the subnet.
POST	/v2.1/blockstores/{id}/tags	You can create or replace tags for a block store. Add the ID of the block store, and the values for the tags in the "key:value pair" format in the request body.	<code>id</code> : The ID of the block store.
PUT	/v2.1/blockstores/{id}	You can modify any block store based on its ID. Add the ID of the block store, and the values that you want to modify in the request body, such as zone, service protocol, and tags.	<code>id</code> : The ID of the block store.
DELETE	/v2.1/blockstores/{id}	You can delete any block store by its ID.	<code>id</code> : The ID of the block store.



Before you delete a block store, you should delete all the disks mapped to it.

## Disks

You can use Disks APIs to retrieve and manage your disks.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/disks	You can retrieve the details of all your disks. Retrieves details of the disks, such as block store ID, name, disk path, protocol, snapshot policy, and tags.	offset: The number of items to skip before starting to collect the result set. limit: The numbers of items to return.
GET	/v2.1/disks/{id}	You can retrieve the details of a specific disk. Retrieves details of the disk, such as block store ID, name, disk path, protocol, snapshot policy, and tags based on the entered ID.	id: The ID of the disk.
POST	/v2.1/disks	You can create a disk. Add the values in the request body to create a disk with the specified parameters, such as subtenant ID, zone, name, disk path, snapshot policy, backup policy, and tags.	subtenant_id: The ID of the subtenant. zone: The name of the zone. name: The name of the disk. disk_path: The path of the disk. protocol: The storage protocol used for block device access. os_type: The type of the host operating system. hostgroup_mappings: The host group mappings. service_level: The service level name applicable: Standard, Premium, Premium-Tiering, Extreme, or Extreme-Tiering. size_gb: The size of the disk in GBs.
POST	/v2.1/disks/{id}/snapshot/{name}	You can create a snapshot of a disk.	id: The ID of the disk. name: The name of the snapshot.
POST	/v2.1/disks/{id}/tags	You can create or replace tags for a disk. Add the ID of the disk and the values for the tags in the "key:value pair" format in the request body.	id: The ID of the disk.
PUT	/v2.1/disks/{id}	You can modify any disk based on its ID. Add the ID of the disk, and the values that you want to modify in the request body, such as name, service level, snapshot policy, backup policy, and tags.	id: The ID of the disk. hostgroup_mappings: The host group mappings.

HTTP Verb	Path	Description	Mandatory parameters/Request body
DELETE	/v2.1/disks/{id}	You can delete any disk by its ID.	id: The ID of the disk.
DELETE	/v2.1/disks/{id}/snapshot/{name}	You can delete any snapshot of a disk by the ID of the disk and the name of the snapshot.	id: The ID of the disk. name: The name of the snapshot.

## Host groups

Access control to disks is managed with host groups. You can retrieve and manage host groups using Host Groups APIs.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/hostgroups	You can retrieve the details of all your host groups. Retrieves details of the host groups, such as name, subtenant details, tenant details, zone, protocol, initiators, disks using the host group, and tags.	offset: The number of items to skip before starting to collect the result set. limit: The numbers of items to return.
GET	/v2.1/hostgroups/{id}	You can retrieve the details of a specific host group. Retrieves details of the host group, such as name, subtenant details, tenant details, zone, protocol, initiators, disks using the host groups, and tags based on the entered ID.	id: The ID of the host group.
POST	/v2.1/hostgroups	You can create a host group. Add the values in the request body to create a host group with the specified parameters, such as name, subtenant ID, zone, protocol, initiators, and tags.	name: The name of the host group. subtenant_id: The ID of the subtenant. zone: The name of the zone. protocol: The storage protocol used for block device access. os_type: The type of the host operating system.
POST	/v2.1/hostgroups/{id}/tags	You can create or replace tags for a host group. Add the ID of the host group and the values for the tags in the "key:value pair" format in the request body.	id: The ID of the host group.

HTTP Verb	Path	Description	Mandatory parameters/Request body
DELETE	/v2.1/host groups/{id }	You can delete any host group by its ID.	id: The ID of the host group.

### Initiators in a host group

You can use Host Groups APIs to retrieve and manage the initiators mapped to your host groups.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/host groups/{id }/ initiators	You can retrieve the details of all your initiators. Retrieves initiators and their aliases.	id: The ID of the host group.
GET	/v2.1/host groups/{id }/ initiators /{alias}	You can retrieve the details of a specific initiator. Retrieves the initiator based on the entered ID and alias.	id: The ID of the host group. alias: The alias name of the initiator.
POST	/v2.1/host groups/{id }/ initiators	You can create an initiator for a host group. Add the values for the initiator and its alias in the request body to create an initiator for the host group.	id: The ID of the host group. alias: The alias name of the initiator. initiator: The initiator (iSCSI Qualified Names or FC WWPNS).
PATCH	/v2.1/host groups/{id }/ initiators /{alias}	You can modify an initiator. Add the new initiator in the request body.	id: The ID of the host group. alias: The alias name of the initiator. initiator: The initiator (iSCSI Qualified Names or FC WWPNS).
DELETE	/v2.1/host groups/{id }/ initiators /{alias}	You can delete an initiator by the ID of the host group and the alias of the initiator.	id: The ID of the host group. alias: The alias name of the initiator.



When adding initiators to a host group, the initiator should match the host group protocol. You should use IQNs for host groups with iSCSI protocol, and WWPNs for host groups with FC protocol.

Deleting an initiator from a host group affects all the disks to which the host group is mapped to.

## File services APIs

You can use File Services APIs to view and manage your file servers and file shares.

Before you create file shares, you have to create file servers to host them.

### File Servers APIs

You can use File Servers APIs to view and manage your file servers. For more information on file servers, see [Create a file server](#).

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/file servers	You can retrieve the details of all your file servers. Retrieves details of the file servers, such as file server ID, name, region, zone, tenant, and tags.	<code>offset</code> : The number of items to skip before starting to collect the result set. <code>limit</code> : The numbers of items to return.
GET	/v2.1/file servers/{id}	You can retrieve the details of a specific file server. Retrieves details of the file server, such as name, region, zone, tenant, and tags based on the entered ID.	<code>id</code> : The ID of the file server.
POST	/v2.1/file servers	You can create a file server. Add the values in the request body to create a file server with the specified parameters, such as name, subtenant ID, zone, subnet ID.	<code>name</code> : The name of the file server. <code>subtenant_id</code> : The ID of the subtenant. <code>zone</code> : The name of the zone. <code>subnet_id</code> : The ID of the subnet.
POST	/v2.1/file servers/{id}/tags	You can create or replace tags for a file server. Add the ID of the file server and the values for the tags in the "key:value pair" format in the request body.	<code>id</code> : The ID of the file server.
PUT	/v2.1/file servers/{id}	You can modify any file server based on its ID. Add the ID of the file server, and the values that you want to modify in the request body, such as such as name, protocol, zone, region, and tags.	<code>id</code> : The ID of the file server.

HTTP Verb	Path	Description	Mandatory parameters/Request body
DELETE	/v2.1/file servers/{id}	You can delete any file server by its ID.	id: The ID of the file server.

Before you delete a file server, you should delete all the file shares mapped to it.



After deleting CIFS-enabled file servers, the Active Directory computer object remains. Ask your Active Directory admin to manually remove the computer object for the deleted file server from Active Directory.

## File Shares APIs

You can use File Shares APIs to view and manage your file shares. For more information on file shares, see [Create a file share](#).

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/file shares	You can retrieve the details of all your file shares. Retrieves details of the file shares, such as ID, name, snapshot policy, protocols, file server IP, and tags.	offset: The number of items to skip before starting to collect the result set. limit: The numbers of items to return.
GET	/v2.1/file shares/{id}	You can retrieve the details of a specific file share. Retrieves details of the file share, such as ID, name, snapshot policy, protocols, file server IP, and tags based on the entered ID.	id: The ID of the file share.

HTTP Verb	Path	Description	Mandatory parameters/Request body
POST	/v2.1/file shares	You can create a file share. Add the values in the request body to create a file share with the specified parameters, such as ID, name, snapshot policy, protocols, file server IP, and tags.	<p>name: The name of the file share.</p> <p>share_path: The path of the file share.</p> <p>fileserver_id: The ID of the file server.</p> <p>size_gb: The size of the file share in GBs.</p> <p>service_level: The service level name applicable: Standard, Premium, Premium-Tiering, Extreme, or Extreme-Tiering.</p> <p>protocol: The protocol used to access the file share (NFS, CIFS, or multi-protocol).</p> <p>security_style: The security style (Unix or NTFS).</p> <p>export_policy: The export policy of the file share.</p>
POST	/v2.1/file shares/{id }/ snapshot/{ name}	You can create a snapshot of a file share.	<p>id: The ID of the file share.</p> <p>name: The name of the snapshot.</p>
POST	/v2.1/file shares/{id }/ tags	You can create or replace tags for a file share. Add the ID of the file share and the values for the tags in the "key:value pair" format in the request body.	<p>id: The ID of the file share.</p>
PUT	/v2.1/file shares/{id }	You can modify any file share based on its ID. Add the ID of the file server and the values that you want to modify in the request body, such as such as name, protocol, snapshot policy, backup policy, and tags.	<p>id: The ID of the file share.</p> <p>name: The name of the file share.</p> <p>size_gb: The size of the file share in GBs.</p> <p>service_level: The service level name applicable: Standard, Premium, Premium-Tiering, Extreme, or Extreme-Tiering.</p> <p>protocol: The protocol used to access the file share (NFS, CIFS, or multi-protocol).</p> <p>export_policy: The export policy of the file share.</p>



HTTP Verb	Path	Description	Mandatory parameters/Request body
DELETE	/v2.1/file shares/{id }	You can delete any file share by its ID.	id: The ID of the file share.
DELETE	/v2.1/file shares/{id }/ snapshot/{ name }	You can delete any snapshot of a file share by the ID of the file share and the name of the snapshot.	id: The ID of the file share. name: The name of the snapshot.



For CIFS shares, adding a \$ character to the end of the share path will make it a hidden share, for example, path\to\my\hidden\share\$.

## Object storage APIs

This section provides the APIs to manage your object storage, object storage users, and object storage group.

The object storage workflow includes these tasks:

- Create an object storage account.
- Create an object storage group.
- Create an object storage user.
- Create an S3 key for user.

## Object storage groups

Use the methods listed in the following table to retrieve, create, or modify object storage groups.

HTTP Method	Path	Description
GET	/v2.1/objectiam/groups	Retrieve object storage groups.
GET	/v2.1/objectiam/groups/{id }	Retrieve an object storage group by ID.
POST	/v2.1/objectiam/groups	Create an object storage identity access management group.
PUT	/v2.1/objectiam/groups/{id }	Modify an object storage group identified by ID.
DELETE	/v2.1/objectiam/groups/{id }	Delete an object storage group identified by ID.

## Object storage group attributes

The following table lists the object storage attributes.

Attribute	Type	Description
id	String	The unique identifier for the object storage group.
name	String	The object storage group name.
subtenant	String	The name of the subtenant to which the group belongs.
subtenant_id	String	The identifier for the subtenant to which the group belongs.
tenant	String	The name of the tenant to which the group belongs.
tenant_id	String	The identifier for the tenant to which the group belongs.
s3_policy		<p>S3 policy For example:</p> <pre> "s3_policy": {   "Statement": [     {       "Effect": "Allow",       "Action": "s3:*",       "Resource": "arn:aws:s3:::*"     }   ] } </pre>

### Retrieve object storage groups

Use the method listed in the following table to retrieve all object storage groups or a subset of object storage groups. Specifying a `subtenant_id` will return only the object storage groups belonging to that subtenant.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectiam/groups	Retrieve all object storage groups. Optionally, specify a subtenant ID to retrieve only the object storage groups associated with the subtenant.	<p><code>subtenant_id</code> (string): The subtenant ID associated with the Identity and Access Management (IAM) users/groups.</p> <p><code>offset</code> and <code>limit</code>— see <a href="#">Common Pagination</a></p>

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "total_records": 1,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 0,
    "limit": 20,
    "records": [
      {
        "id": "5e1bleff8bc5c0300011c989c",
        "name": "MyGroup",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant",
        "subtenant_id": "5e7c3af8aab46c00014ce878",
        "subtenant": "MySubtenant",
        "s3_policy": {
          "Statement": [
            {
              "Action": [
                "s3:*"
              ],
              "Effect": "Allow",
              "Resource": "arn:aws:s3:::*"
            }
          ]
        }
      }
    ]
  }
}
```

### Retrieve an object storage group by ID

Use the method listed in the following table to retrieve an object storage group by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectiam/groups/{id}	Retrieve an object storage group by ID.	id (string): The unique identifier of the object storage group.

Required request body attributes: none

#### Request body example:

```
none
```

#### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5e1eff8bc5c0300011c989c",
        "name": "MyGroup",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant",
        "subtenant_id": "5e7c3af8aab46c00014ce878",
        "subtenant": "MySubtenant",
        "s3_policy": {
          "Statement": [
            {
              "Action": [
                "s3:*"
              ],
              "Effect": "Allow",
              "Resource": "arn:aws:s3:::*"
            }
          ]
        }
      }
    ]
  }
}

```

### Create an object storage group

Use the method listed in the following to create an object storage group.

HTTP Method	Path	Description	Parameters
POST	/v2.1/objectiam/groups/	Create a new object storage group service to host object storage users.	None

Required request body attributes: name, subtenant\_id, s3Policy

### Request body example:

```

{
  "name": "MyNewGroup",
  "subtenant_id": "5e7c3af8aab46c00014ce878",
  "s3_policy": {
    "Statement": [
      {
        "Effect": "Allow",
        "Action": "s3:*",
        "Resource": "arn:aws:s3:::*"
      }
    ]
  }
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed5fa312c356a0001a73841",
        "action": "create",
        "job_summary": "Create request is successfully submitted",
        "created": "2020-06-02T07:05:21.130260774Z",
        "updated": "2020-06-02T07:05:21.130260774Z",
        "object_id": "5ed5fa312c356a0001a73840",
        "object_type": "sg_groups",
        "object_name": "MyNewGroup",
        "status": "pending",
        "status_detail": "",
        "last_error": "",
        "user_id": "5ec626c0f038943eb46b0af1",
        "job_tasks": null
      }
    ]
  }
}

```

## Modify an object storage group

Use the method listed in the following table to modify an object storage group.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/objectiam/groups/{id}	Modify an object storage group.	id (string): The unique identifier of the object storage group.

Required request body attributes: name, subtenant\_id, s3Policy

### Request body example:

```
{
  "s3_policy": {
    "Statement": [
      {
        "Action": [
          "s3:ListAllMyBuckets",
          "s3:ListBucket",
          "s3:ListBucketVersions",
          "s3:GetObject",
          "s3:GetObjectTagging",
          "s3:GetObjectVersion",
          "s3:GetObjectVersionTagging"
        ],
        "Effect": "Allow",
        "Resource": "arn:aws:s3:::*"
      }
    ]
  }
}
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed5fe822c356a0001a73859",
        "action": "update",
        "job_summary": "Update request is successfully submitted",
        "created": "2020-06-02T07:23:46.43550235Z",
        "updated": "2020-06-02T07:23:46.43550235Z",
        "object_id": "5ed5fa312c356a0001a73840",
        "object_type": "sg_groups",
        "object_name": "MyNewGroup",
        "status": "pending",
        "status_detail": "",
        "last_error": "",
        "user_id": "5ec626c0f038943eb46b0af1",
        "job_tasks": null
      }
    ]
  }
}

```

### Delete an object storage group by ID

Use the method listed in the following table to delete an object storage group by ID.

HTTP Method	Path	Description	Parameters
Delete	/v2.1/objectiam/groups/{id}	Delete an object storage group by ID.	id (string): The unique identifier of the object storage group.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:



```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5e1e1eff8bc5c0300011c989c",
        "name": "MyGroup",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant",
        "subtenant_id": "5e7c3af8aab46c00014ce878",
        "subtenant": "MySubtenant",
        "s3_policy": {
          "Statement": [
            {
              "Action": [
                "s3:*"
              ],
              "Effect": "Allow",
              "Resource": "arn:aws:s3:::*"
            }
          ]
        }
      }
    ]
  }
}

```

## Object storage users

Use the methods listed in the following table to perform the following tasks:

- Retrieve, create, or modify object storage users.
- Create S3 keys, retrieve S3 keys for a user, or retrieve keys by key ID.

HTTP Method	Path	Description
GET	/v2.1/objectiam/users	Retrieve object storage users.
GET	/v2.1/objectiam/users/{id}	Retrieve an object storage user by ID.
POST	/v2.1/objectiam/users	Create an object storage user.

HTTP Method	Path	Description
PUT	/v2.1/objectiam/users/{id}	Modify an object storage user identified by ID.
DELETE	/v2.1/objectiam/users/{id}	Delete an object storage user by ID.
GET	/v2.1/objectiam/users/{user_id}/s3keys	Get all S3 keys mapped to a user.
POST	/v2.1/objectiam/users/{user_id}/s3keys	Create S3 keys.
GET	/v2.1/objectiam/users/{user_id}/s3keys/{key_id}	Get S3 keys by key ID.
DELETE	/v2.1/objectiam/users/{user_id}/s3keys/{key_id}	Delete S3 keys by key ID.

### Object storage user attributes

The following table lists the object storage user attributes.

Attribute	Type	Description
id	String	The unique identifier for the object storage user.
display_name	String	The display name of the user.
subtenant	String	The name of the subtenant to which the user belongs.
subtenant_id	String	The identifier for the subtenant to which the user belongs.
tenant	String	The name of the tenant to which the user belongs.
tenant_id	String	The identifier for the tenant to which the user belongs.
objectiam_user_urn	String	The URN.
sg_group_membership	String	NetApp StorageGRID group memberships. For example: "sg_group_membership": [ "5d2fb0fb4f47df00015274e3" ]

### Retrieve object storage users

Use the method listed in the following table to retrieve all object storage users or a subset of object storage users. Specifying a `subtenant_id` will return only the object storage groups belonging to that subtenant.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectiam/users	Retrieve all object storage users.	subtenant_id (string): The subtenant ID associated with the IAM users/groups.  offset and limit – see <a href="#">Common Pagination</a>

Required request body attributes: none

**Request body example:**

```
none
```

**Response body example:**

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "total_records": 1,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 0,
    "limit": 20,
    "records": [
      {
        "id": "5eb2212d1cbe3b000134762e",
        "display_name": "MyUser",
        "subtenant": "MySubtenant",
        "subtenant_id": "5e7c3af8aab46c00014ce878",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant",
        "objectiam_user_urn":
"urn:sgws:identity::96465636379595351967:user/myuser",
        "sg_group_membership": [
          "5eb1eff8bc5c0300011c989c"
        ]
      }
    ]
  }
}

```

### Retrieve an object storage user by ID

Use the method listed in the following table to retrieve an object storage use by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectiam/users{id}	Retrieve an object storage user by ID.	id: The object storage account ID.

Required request body attributes: none

### Request body example:

```
none
```

## Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5eb2212d1cbe3b000134762e",
        "display_name": "MyUser",
        "subtenant": "MySubtenant",
        "subtenant_id": "5e7c3af8aab46c00014ce878",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant",
        "objectiam_user_urn":
"urn:sgws:identity::96465636379595351967:user/myuser",
        "sg_group_membership": [
          "5eb1eff8bc5c0300011c989c"
        ]
      }
    ]
  }
}
```

## Create an object storage user

Use the method listed in the following table to create an object storage user.

HTTP Method	Path	Description	Parameters
POST	/v2.1/objectiam/users	Create a new object storage user.	None

Required request body attributes: display\_name, subtenant\_id, sg\_group\_membership

## Request body example:

```
{
  "display_name": "MyUserName",
  "subtenant_id": "5e7c3af8aab46c00014ce878",
  "sg_group_membership": [
    "5ed5fa312c356a0001a73840"
  ]
}
```

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed603712c356a0001a7386c",
        "action": "create",
        "job_summary": "Activate request is successfully submitted",
        "created": "2020-06-02T07:44:49.647815816Z",
        "updated": "2020-06-02T07:44:49.647815816Z",
        "object_id": "5ed603712c356a0001a7386d",
        "object_type": "sg_users",
        "object_name": "MyUserName",
        "status": "pending",
        "status_detail": "",
        "last_error": "",
        "user_id": "5ec626c0f038943eb46b0af1",
        "job_tasks": null
      }
    ]
  }
}
```

### Modify an object storage user

Use the method listed in the following table to modify an object storage user.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/objectiam/users/{id}	Modify an object storage user identified by ID.	id: The object storage user ID.

Required request body attributes: display\_name, subtenant\_id, sg\_group\_membership

**Request body example:**

```
{
  "display_name": "MyModifiedObjectStorageUser",
  "subtenant_id": "5e57a465896bd80001dd4961",
  "sg_group_membership": [
    "5e60754f9b64790001fe937b"
  ]
}
```

**Response body example:**

```

{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed604002c356a0001a73880",
        "action": "update",
        "job_summary": "Update request is successfully submitted",
        "created": "2020-06-02T07:47:12.205889873Z",
        "updated": "2020-06-02T07:47:12.205889873Z",
        "object_id": "5ed603712c356a0001a7386d",
        "object_type": "sg_users",
        "object_name": "MyUserName",
        "status": "pending",
        "status_detail": "",
        "last_error": "",
        "user_id": "5ec626c0f038943eb46b0af1",
        "job_tasks": null
      }
    ]
  }
}

```

### Map all S3 keys to an object storage user

Use the method listed in the following table to map all S3 keys to an object storage user.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectiam/users/{user_id}/s3keys	Create an S3 key for an object storage user.	user_id (string): The object storage user identifier.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:



```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5e66de2509a74c0001b895e7",
        "display_name": "*****HNDE",
        "subtenant_id": "5e57a465896bd80001dd4961",
        "subtenant": "BProject",
        "objectiam_user_id": "5e66c77809a74c0001b89598",
        "objectiam_user": "MyNewObjectStorageUser",
        "objectiam_user_urn":
"urn:sgws:identity::09936502886898621050:user/mynewobjectstorageuser",
        "expires": "2020-04-07T10:40:52Z"
      }
    ]
  }
}

```

### Create an S3 key for an object storage user

Use the method listed in the following to create an S3 key for an object storage user.

HTTP Method	Path	Description	Parameters
POST	/v2.1/objectiam/users/{user_id}/s3keys	Create an S3 key for an object storage user.	user_id (string): The object storage user identifier.

Required request body attributes: expires (string)



The key expiry date/time is set in UTC—it must be set in the future.

### Request body example:

```

{
  "expires": "2020-04-07T10:40:52Z"
}

```

### Response body example:

```

"status": {
  "user_message": "Okay. Returned 1 record.",
  "verbose_message": "",
  "code": 200
},
"result": {
  "total_records": 1,
  "records": [
    {
      "id": "5e66de2509a74c0001b895e7",
      "display_name": "*****HNDE",
      "subtenant_id": "5e57a465896bd80001dd4961",
      "subtenant": "BProject",
      "objectiam_user_id": "5e66c77809a74c0001b89598",
      "objectiam_user": "MyNewObjectStorageUser",
      "objectiam_user_urn":
"urn:sgws:identity::09936502886898621050:user/mynewobjectstorageuser",
      "expires": "2020-04-07T10:40:52Z",
      "access_key": "PL86KPEBN6XT4T7UHNDE",
      "secret_key": "F1D/YWAM7JMr9gG8pumU8dzvcTLMzLYtUe21NzcA"
    }
  ]
}
}

```

### Get S3 keys for an object storage user by key ID

Use the method listed in the following table to get S3 keys for an object storage user by key ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectiam/users/{user_id}/s3keys/{key_id}	Get S3 keys by key ID.	<ul style="list-style-type: none"> <li><b>user_id (string):</b> The object storage user ID. For example: 5e66c77809a74c0001b89598</li> <li><b>key_id (string):</b> S3 key For example: 5e66de2509a74c0001b895e7</li> </ul>

Required request body attributes: none

### Request body example:

none

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ecc7bb9b5d2730001f798fb",
        "display_name": "*****XCXD",
        "subtenant_id": "5e7c3af8aab46c00014ce878",
        "subtenant": "MySubtenant",
        "objectiam_user_id": "5eb2212d1cbe3b000134762e",
        "objectiam_user": "MyUser",
        "objectiam_user_urn":
"urn:sgws:identity::96465636379595351967:user/myuser",
        "expires": "2020-05-27T00:00:00Z"
      }
    ]
  }
}
```

### Delete an S3 key by key ID

Use the method listed in the following table to delete an S3 key by key ID.

HTTP Method	Path	Description	Parameters
Delete	/v2.1/objectiam/users/{user_id}/s3keys/{key_id}	Delete S3 key by key ID.	<ul style="list-style-type: none"><li>• <code>user_id</code> (string): The object storage user ID. For example: 5e66c77809a74c0001b89598</li><li>• <code>key_id</code> (string): S3 key For example: 5e66de2509a74c0001b895e7</li></ul>

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
No content to return for succesful execution
```

## Object storage accounts

Use the methods listed in the following table to perform the following tasks:

- Retrieve, activate, or modify object storage accounts.
- Create S3 buckets.

HTTP Method	Path	Description
GET	/v2.1/objectstorage/accounts	Retrieve object storage accounts.
GET	/v2.1/objectstorage/accounts/{id}	Retrieve an object storage account by ID.
POST	/v2.1/objectstorage/accounts	Create an object storage account.
PUT	/v2.1/objectstorage/accounts/{id}	Modify an object storage account identified by ID.
DELETE	/v2.1/objectstorage/accounts/{id}	Modify an object storage account identified by ID.
GET	/v2.1/objectstorage/buckets	Get S3 buckets.
POST	/v2.1/objectstorage/buckets	Create S3 buckets.

### Object storage account attributes

The following table lists the object storage account attributes.

Attribute	Type	Description
id	String	The unique identifier of the object storage user.
subtenant_id	String	The identifier of the instance of a subtenant object.
quota_gb	Integer	The size of the share or disk.

### Retrieve all object storage accounts

Use the method listed in the following table to retrieve all object storage accounts or a subset of object storage accounts.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectstorage/accounts	Retrieve all object storage users.	offset and limit- . see <a href="#">Common Pagination</a>

Required request body attributes: none

#### Request body example:

```
none
```

#### Response body example

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "total_records": 19,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 3,
    "limit": 1,
    "records": [
      {
        "id": "5ec6119e6344d000014cdc41",
        "name": "MyTenant - MySubtenant",
        "subtenant": " MySubtenant",
        "subtenant_id": "5ea8c5e083a9f80001b9d705",
        "tenant": "E- MyTenant",
        "tenant_id": "5d914499869caefed0f39eee",
        "sg_account_id": "29420999312809208626",
        "quota_gb": 100,
        "sg_instance_name": "NSE StorageGRID Dev1",
        "sg_instance_id": "5e3ba2840271823644cb8ab6"
      }
    ]
  }
}

```

### Retrieve an object storage account by ID

Use the method listed in the following table to retrieve an object storage account by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectstorage/accounts/{id}	Retrieve an object storage account by ID.	id: The object storage account ID.

Required request body attributes: none

### Request body example:

none

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ec6119e6344d000014cdc41",
        "name": "MyTenant - MySubtennant",
        "subtenant": " MySubtennant",
        "subtenant_id": "5ea8c5e083a9f80001b9d705",
        "tenant": " MyTenant",
        "tenant_id": "5d914499869caefed0f39eee",
        "sg_account_id": "29420999312809208626",
        "quota_gb": 100,
        "sg_instance_name": "NSE StorageGRID Dev1",
        "sg_instance_id": "5e3ba2840271823644cb8ab6"
      }
    ]
  }
}

```

### Activate an object storage account

Use the method listed in the following table to activate an object storage account.

HTTP Method	Path	Description	Parameters
POST	/v2.1/objectstorage/accounts	Activate an object storage service.	None

Required request body attributes: subtenant\_id, quota\_gb

### Request body example:

```

{
  "subtenant_id": "5ecefbbef418b40001f20bd6",
  "quota_gb": 20
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed608542c356a0001a73893",
        "action": "create",
        "job_summary": "Activate request for Sub Tenant MyNewSubtenant is
successfully submitted",
        "created": "2020-06-02T08:05:40.017362022Z",
        "updated": "2020-06-02T08:05:40.017362022Z",
        "object_id": "5ed608542c356a0001a73894",
        "object_type": "sg_accounts",
        "object_name": "MyTenant - MyNewSubtenant",
        "status": "pending",
        "status_detail": "",
        "last_error": "",
        "user_id": "5ec626c0f038943eb46b0af1",
        "job_tasks": null
      }
    ]
  }
}

```

### Modify an object storage account

Use the method listed in the following table to modify an object storage account.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/objectstorage /accounts/{id}	Modify an object storage service (such as, change the quota).	id (string): The object storage account ID.

Required request body attributes: name, subtenant\_id, quota\_gb

### Request body example:



```
{
  "name": "MyTenant - MyNewSubtenant",
  "subtenant_id": "5ecefbbef418b40001f20bd6",
  "quota_gb": 30
}
```

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed609162c356a0001a73899",
        "action": "update",
        "job_summary": "Update request is successfully submitted",
        "created": "2020-06-02T08:08:54.841652098Z",
        "updated": "2020-06-02T08:08:54.841652098Z",
        "object_id": "5ed608542c356a0001a73894",
        "object_type": "sg_accounts",
        "object_name": "MyTenant - MyNewSubtenant",
        "status": "pending",
        "status_detail": "",
        "last_error": "",
        "user_id": "5ec626c0f038943eb46b0af1",
        "job_tasks": null
      }
    ]
  }
}
```

### Delete an object storage account

Before you can delete an object storage account, you must first delete all associated groups, users, and buckets. Use the method listed in the following table to delete an object storage account.



Use your S3 compatible utility to delete buckets. It is not possible to delete buckets from NetApp Service Engine.

HTTP Method	Path	Description	Parameters
Delete	/v2.1/objectstorage/accounts/{id}	Delete an object storage account.	id (string): The object storage account ID.

Required request body attributes: none

#### Request body example:

```
{
  "name": "MyTenant - MyNewSubtenant",
  "subtenant_id": "5ecefbbef418b40001f20bd6",
  "quota_gb": 30
}
```

#### Response body example:

```
{
  "status": {
    "user_message": "string",
    "verbose_message": "string",
    "code": "string"
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5d2fb0fb4f47df00015274e3",
        "action": "string",
        "object_id": "5d2fb0fb4f47df00015274e3",
        "object_type": "string",
        "status": "string",
        "status_detail": "string",
        "last_error": "string",
        "user_id": "5d2fb0fb4f47df00015274e3",
        "link": "string"
      }
    ]
  }
}
```

### Object storage buckets

Use the APIs in the following table to create and retrieve object storage buckets.

HTTP Method	Path	Description
GET	/v2.1/objectstorage/buckets	Retrieve object storage buckets.
POST	/v2.1/objectstorage/buckets	Create an object storage bucket.

### Object storage bucket attributes

The following table lists the object storage bucket attributes.

Attribute	Type	Description
id	String	The unique identifier for the object storage user.
Name	String	The bucket name.
subtenant_id	String	The identifier of the subtenant to which the bucket belongs.

### Retrieve S3 buckets

Use the method listed in the following table to retrieve S3 buckets.

HTTP Method	Path	Description	Parameters
GET	/v2.1/objectstorage/buckets	Retrieve S3 buckets.	Subtenant_id: The subtenant that owns the bucket.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "creationTime": "2020-06-02T08:13:25.695Z",
        "name": "mybucket"
      }
    ]
  }
}

```

### Create S3 buckets

Use the method listed in the following table to create an S3 bucket.



Before you can create a bucket, an object storage account for the subtenant must exist.

HTTP Method	Path	Description	Parameters
POST	/v2.1/objectstorage /buckets	Create an S3 bucket.	None

Required request body attributes:

- name (string): S3 bucket name (lowercase or numeric characters only)
- subtenant\_id (string): ID of the subtenant to which the S3 bucket belongs

### Request body example:

```

{
  "name": "mybucket",
  "subtenant_id": "5ecefbbef418b40001f20bd6"
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed60a232c356a0001a7389e",
        "action": "create",
        "job_summary": "Create request is successfully submitted",
        "created": "2020-06-02T08:13:23.105015108Z",
        "updated": "2020-06-02T08:13:23.105015108Z",
        "object_id": "5ed60a232c356a0001a7389f",
        "object_type": "sg_buckets",
        "object_name": "mybucket",
        "status": "pending",
        "status_detail": "",
        "last_error": "",
        "user_id": "5ec626c0f038943eb46b0af1",
        "job_tasks": null
      }
    ]
  }
}


```

## Backups APIs

You can use Backup APIs to view and manage the Snapshots (backups or recovery points) of the volumes (file shares and disks) in your environment.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/backups	You can retrieve the details of all the backup objects for all the volumes created for the subtenant under tenant. Retrieves details of the source volume and the backup object, such as backup policy, zone, and tags.	<b>offset:</b> The number of items to skip before starting to collect the result set. <b>limit:</b> The numbers of items to return.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/backups/{id}	You can retrieve the details of a specific backup object created for a volume for the subtenant under tenant. Retrieves details of the source volume and the backup object, such as backup policy, zone, and tags, based on the entered ID.	id: The ID of the backup object.
GET	/v2.1/backups/{id}/recovery_points	You can retrieve details of all the recovery points of a specific backup object. Details, such as time stamp and name are retrieved.	id: The ID of the backup object.
GET	/v2.1/backups/{id}/recovery_points/{name}	You can retrieve details of any recovery point of a specific backup object. Details, such as time stamp and name are retrieved.	id: The ID of the backup object. name: The name of the recovery point.
POST	/v2.1/backups	You can create a backup object for a particular source volume. Add the values in the request body to create a backup object with the specified parameters, such as backup policy, number of backup copies to retain, and source resource ID.	<pre>{   "primary_resource_type": "&lt;Resource Type&gt;",   "source_resource_id": "&lt;ID&gt;",   "backup_zone": "&lt;Backup Zone&gt;",   "backup_policy": {     "daily_backups_to_keep": &lt;Number of daily backups to keep&gt;,     "weekly_backups_to_keep": &lt;Number of weekly backups to keep&gt;,     "monthly_backups_to_keep": &lt;Number of monthly backups to keep&gt;,     "adhoc_backups_to_keep": &lt;Number of adhoc backups to keep&gt;   },   "tags": {     "key1": "&lt;Tag 1&gt;",     "key2": "&lt;Tag 2&gt;",     "keyN": "&lt;Tag N&gt;"   } }</pre>

HTTP Verb	Path	Description	Mandatory parameters/Request body
POST	/v2.1/backups/{id}/tags	You can create or replace tags for a backup object for your subtenant. Add the ID of the backup object and the values for the tags in the "key:value pair" format in the request body.	<p>id: The ID of the backup object.</p> <pre>{   "key1": "&lt;Tag 1&gt;",   "key2": "&lt;Tag 2&gt;",   "keyN": "&lt;Tag N&gt;" }</pre>
PATCH	/v2.1/backups/{id}	<p>You can modify any backup object for a volume based on its ID. Add the ID of the backup object and the values that you want to modify in the request body, such as the backup policy details, number of backups to retain, and source resource ID.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  <p>You can make a backup object orphan by setting the value of the <code>source_resource_id</code> parameter as null in your PATCH request. This removes the SnapMirror relationship, but retains the backup volume and the recovery points under it.</p> </div>	<p>id: The ID of the backup object.</p> <pre>{   "source_resource_id":   "&lt;Resource ID&gt;",   "backup_policy": {     "daily_backups_to_keep":     &lt;Number of daily     backups to keep&gt;,     "weekly_backups_to_kee     p": &lt;Number of weekly     backups to keep&gt;,     "monthly_backups_to_ke     ep": &lt;Number of     monthly backups to     keep&gt;,     "adhoc_backups_to_keep     ": &lt;Number of adhoc     backups to keep&gt;   },   "tags": {     "key1": "&lt;Tag 1&gt;",     "key2": "&lt;Tag 2&gt;",     "keyN": "&lt;Tag 3&gt;"   } }</pre>
DELETE	/v2.1/backups/{id}	You can delete any backup object, along with all the recovery points, for a volume.	id: The ID of the backup object.
DELETE	/v2.1/backups/{id}/recovery_points/{name}	You can delete any recovery point in a specific backup object for a volume.	<p>id: The ID of the backup object.</p> <p>name: The name of the recovery point.</p>

## Reporting APIs

Use the methods in the following table to retrieve and generate reports.

HTTP Method	Path	Description
GET	/v2.1/reports	Retrieve all list of reports.
POST	/v2.1/reports/{filename}	Retrieve a report.

### Retrieve a list of reports

Use the method listed in the following table to retrieve a list of reports.

HTTP Method	Path	Description	Parameters
GET	/v2.1/reports	Retrieve all available reports for a tenant. Reports can be filtered by date range.	<ul style="list-style-type: none"><li>• <code>tenant_id</code> (string)</li><li>• <code>start</code> (string): Retrieve reports more recent than the date specified in <code>start</code>.</li><li>• <code>end</code> (string): Retrieve reports older than the date specified in <code>end</code>.</li></ul>

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:



```

{
  "status": {
    "user_message": "string",
    "verbose_message": "string",
    "code": "string"
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "filename": "string"
      }
    ]
  }
}

```

### Retrieve a report identified by file name

Use the method in the following table to retrieve a report identified by the file name.

HTTP Method	Path	Description	Parameters
GET	/v2.1/reports/{file name}	Retrieve a report identified by the file name.	filename (string): The report file name.

Required request body attributes: none

### Request body example:

none

### Response body example:

```

{
  "status": {
    "user_message": "string",
    "verbose_message": "string",
    "code": "string"
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "filename": "string"
      }
    ]
  }
}

```

## Define network configurations with Subnets APIs

You can use Subnets APIs to view and create subnets for your subtenant and zone. You can use them to define your network configuration.

HTTP Verb	Path	Description	Mandatory parameters/Request body
GET	/v2.1/tenants/{tenant_id}/zones/{zone_name}/subnets	You can retrieve all the subnet objects for a tenant and zone. Retrieves details of the subnet, such as VLAN, subtenant, zone, routes, and tags.	tenant_id: The ID of the tenant. zone_name: The name of the zone. offset: The number of items to skip before starting to collect the result set. limit: The numbers of items to return.
GET	/v2.1/tenants/{tenant_id}/zones/{zone_name}/subnets/{id}	You can retrieve the details of a specific subnet object created for a tenant and zone. Retrieves details of the subnet, such as VLAN, subtenant, zone, routes, and tags based on the entered ID.	tenant_id: The ID of the tenant. zone_name: The name of the zone. id: The ID of the subnet.

HTTP Verb	Path	Description	Mandatory parameters/Request body
POST	/v2.1/tenants/{tenant_id}/zones/{zone_name}/subnets	You can create a subnet object for a particular tenant and zone. Add the values in the request body to create a subnet object with the specified parameters, such as name, VLAN, CIDR, subnet ID, routes, and tags.	<p>tenant_id: The ID of the tenant.</p> <p>zone_name: The name of the zone.</p> <pre>{   "name": "string",   "vlan": "1000",   "cidr": "10.0.0.0/24",   "subnet_id":   "5d2fb0fb4f47df00015274e3",   "routes": [     {       "destination":       "10.0.0.0/24",       "gateway": "10.0.0.1",       "metric": "20"     }   ],   "tags": {     "key1": "Value 1",     "key2": "Value 2",     "keyN": "Value N"   } }</pre>
POST	/v2.1/tenants/{tenant_id}/zones/{zone_name}/subnets/{id}/tags	You can create or replace tags for a subnet object for your tenant. Add the ID of the subnet object and the values for the tags in the "key:value pair" format in the request body.	<p>tenant_id: The ID of the tenant.</p> <p>zone_name: The name of the zone.</p> <p>id: The ID of the backup object.</p> <pre>{   "key1": "&lt;Tag 1&gt;",   "key2": "&lt;Tag 2&gt;",   "keyN": "&lt;Tag N&gt;" }</pre>

## (Consumer) Administration APIs

The (consumer) administration APIs consist of methods that allow you to perform the following tasks:

- Sign in, set a password, and refresh an authentication token.

- Retrieve jobs and view job details.

See [Retrieve jobs](#) and [Retrieve a job-by-job ID](#).

- Retrieve regions.

See [Retrieve regions](#) and [Retrieve a region by name](#).

- Retrieve service levels.

See [Retrieve Service Levels](#) and [Retrieve Service Levels by Name](#).

- Work with subtenants.

See:

- [Retrieve all subtenants](#)
- [Retrieve a subtenant by ID](#)
- [Modify a subtenant](#)
- [Delete a subtenant by ID](#)

- Work with tenants.

See:

- [Retrieve all tenants](#)
- [Retrieve a tenant by ID](#)
- [Create a tenant](#)
- [Modify the tenant](#)
- [Delete the tenant](#)

- Retrieve users.

[Retrieve all users](#), [Retrieve a user by ID](#), and [Retrieve a user by user name](#).

- Retrieve zones.

See [Retrieve all zones](#) and [Retrieve a zone by name](#)

The following table lists the consumer APIs documented in this section.

HTTP Method	Path	Description
POST	/v2.1/auth/password	Set the password for a user.
POST	/v2.1/auth/password	Refresh authentication JWT.
POST	/v2.1/auth/signin	Sign in.
GET	/v2.1/auth/regions	Retrieve regions.
GET	/v2.1/auth/regions/{name}	Retrieve regions by name.
GET	/v2.1/auth/zones	Retrieve zones.

HTTP Method	Path	Description
GET	/v2.1/auth/zones/{name}	Retrieve zones by name.
GET	/v2.1/jobs/	Retrieve jobs.
GET	/v2.1/jobs/{id}	Retrieve jobs by ID.

### Reset user password

Use the method listed in the following table reset the user password.

HTTP Method	Path	Description	Parameters
POST	/v2.1/auth/password	Reset the password.	None

Required request body attributes: `username` (string), `new_password` (string)

### Request body example:

```
{
  "username": "MyName",
  "old_password": "oldPassword",
  "new_password": "newPassword"
}
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "user": {
          "id": "5e61aa814559c20001df1a5f",
          "username": "MyName",
          "firstName": "MyFirstName",
          "lastName": "MySurname",
          "displayName": "CallMeMYF",
          "email": "user@example.com",
          "tenancies": [
            {
              "id": "5e5f1c4f253c820001877839",
              "name": "MyTenant",
              "code": "testtenantmh",
              "role": "user"
            }
          ]
        },
        "token":
"eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJ1c2VybmFtZSI6ImV4cCI6MTU4Mzg3Njg3MX0.ZuRXjDPVtc2pH-e9wqgmszVKCBYS2PLqux2YwQ5uoAM"
      }
    ]
  }
}

```

### Refresh authentication token

Use the method listed in the following table to refresh the authentication token.

HTTP Method	Path	Description	Parameters
POST	/v2.1/auth/refresh	Refresh the authentication token.	None

Required request body attributes: none

### Request body example:

none

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "user": {
          "id": "5d914547869caefed0f3a00c",
          "username": "myusername",
          "firstName": "myfirstname",
          "lastName": "",
          "displayName": "Myfirstname Mysurname",
          "email": "",
          "tenancies": [
            {
              "id": "5d914499869caefed0f39eee",
              "name": "MyOrg",
              "code": "myorg",
              "role": "admin"
            },
            {
              "id": "5d9417aa869caefed0f7b4f9",
              "name": "ABCsafe",
              "code": "abcsafe",
              "role": "admin"
            }
          ]
        },
        "token":
        "eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJ1c2VybmFtZSI6ImVsbGlvdCI6ImV4cCI6MTU4MzgxNzA2N30.FdKD3QhPoNdWdbMfZ0bzCMTHluIt6HNP311F482K9AY"
      }
    ]
  }
}
```

## Sign in

Use the method listed in the following table to sign in.

HTTP Method	Path	Description	Parameters
POST	/v2.1/auth/signin	Log in as a user.	None

Required request body attributes: `username` (string), `new_password` (string)

### Request body example:

```
{
  "username": "MyName",
  "password": "newPassword"
}
```

### Response body example:



```

{
  "status": {
    "user_message": "Authentication succeeded.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "user": {
          "id": "5e61aa814559c20001df1a5f",
          "username": "MyName",
          "firstName": "MyFirstName",
          "lastName": "MySurname",
          "displayName": "CallMeMYF",
          "email": "user@example.com",
          "tenancies": [
            {
              "id": "5e5f1c4f253c820001877839",
              "name": "MyTenant",
              "code": "testtenantmh",
              "role": "user"
            }
          ]
        },
        "token":
"eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJ1c2VybmFtZSI6ImV4cCI6
MTU4MzgxNzQwMH0._u_UyYrZg_RewF-9ClIGoKQhfZYWrixZYBrsj1kh1hI"
      }
    ]
  }
}

```

## Administrator APIs

### Overview

This section describes the following administrator APIs:

- Tenants
- Subtenants
- Users

Activities such as resetting user passwords, refreshing tokens, or logging in as a user are available as part

of the Consumer API suite. See (Consumer) Administration APIs.

- Zones
- Regions
- ONTAP clusters
- StorageGRID instances
- Service levels
- Service requests
- Jobs

## Tenants

Use the methods listed in the following table to retrieve, create, modify, and delete tenants.

HTTP Method	Path	Description
GET	/v2.1/tenants	Retrieve a list of all tenants.
GET	/v2.1/tenants/{id}	Retrieve a tenant by the tenant ID.
POST	/v2.1/tenants	Create a new tenant.
PUT	/v2.1/tenants/{id}	Modify the details of a tenant.
DELETE	/v2.1/tenants/{id}	Delete a tenant.

### Tenant attributes

The following table lists the tenant attributes.

Attribute	Type	Description
id	String	The unique identifier of the tenant.
code	String	A customer-specified (or default) code that represents the tenant. This attribute can contain lowercase letters, numbers, and underscores.
name	String	The tenant name.
zuora_account_name	String	The billing account name: the name of the subscription in Zuora.
zuora_account_number	String	The billing account number: the subscription number in Zuora.
description	String	The description of the tenant.

Attribute	Type	Description
usage	–	The services and service details applicable to the tenant. For each service level, this attribute displays the following: <b>name:</b> Service level name <b>used_size_gb:</b> Service level name <b>role_name:</b> User role (user, admin, read, partner, or root)

### Retrieve all tenants

Use the method listed in the following table to retrieve all tenants or a subset of all tenants.

HTTP Method	Path	Description	Parameters
GET	/v2.1/tenants	Retrieve all tenants.	offset and limit – see <a href="#">Common Pagination</a>

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 2 records.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 2,
    "total_records": 23,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 0,
    "limit": 2,
    "records": [
      {
        "id": "5e7c3af7aab46c00014ce877",
        "name": "MyTenant",
        "zuora_account_name": "MyAccount",
        "zuora_account_number": "A00000415",
        "description": ""
      }
    ]
  }
}
```

```

"code": "mytenantcode",
"usage": {
  "A-S00003875": [
    {
      "service_level": "extreme",
      "consumed": 0,
      "committed": 10,
      "burst": 0
    },
    {
      "service_level": "standard",
      "consumed": 1.94,
      "committed": 30,
      "burst": 0
    }
  ],
  "A-S00004566": [
    {
      "service_level": "object",
      "consumed": 3.31,
      "committed": 300,
      "burst": 0
    }
  ]
}
},
{
  "id": "5d914499869caefed0f39eee",
  "name": "MyOrg",
  "zuora_account_name": "MyOrg Inc",
  "zuora_account_number": "A00000415",
  "description": "",
  "code": "myorg",
  "usage": {
    "A-S00003875": [
      {
        "service_level": "standard",
        "consumed": 12.33,
        "committed": 30,
        "burst": 0
      },
      {
        "service_level": "object",
        "consumed": 0,
        "committed": 40,
        "burst": 0
      }
    ]
  }
}

```

```

    }
  ],
  "A-S00003969": [
    {
      "service_level": "extreme",
      "consumed": 0,
      "committed": 5,
      "burst": 0
    }
  ]
}
]
}
}
}
}
}
}

```

### Retrieve a tenant by ID

Use the method listed in the following table to retrieve a tenant by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/tenants/{id}	Retrieve the tenant specified by the ID.	id (string): The unique identifier of the tenant.

Required request body attributes: none

Request body example:

```
none
```

Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5e7c3af7aab46c00014ce877",
        "name": "MyTenant",

```

```
"zuora_account_name": "MyAccount",
"zuora_account_number": "A00000415",
"description": "",
"code": "mytenantcode",
"usage": {
  "A-S00003875": [
    {
      "service_level": "extreme",
      "consumed": 0,
      "committed": 10,
      "burst": 0
    },
    {
      "service_level": "premium",
      "consumed": 2.4,
      "committed": 20,
      "burst": 0
    },
    {
      "service_level": "standard",
      "consumed": 1.94,
      "committed": 30,
      "burst": 0
    },
    {
      "service_level": "object",
      "consumed": 0,
      "committed": 40,
      "burst": 0
    }
  ],
  "A-S00003969": [
    {
      "service_level": "extreme",
      "consumed": 0,
      "committed": 5,
      "burst": 0
    },
    {
      "service_level": "standard",
      "consumed": 0,
      "committed": 30,
      "burst": 0
    }
  ],
  "A-S00004566": [
```

```

    {
      "service_level": "object",
      "consumed": 3.31,
      "committed": 300,
      "burst": 0
    }
  ]
}
]
}
}

```

### Create a tenant

Use the method listed in the following table to create a tenant.

HTTP Method	Path	Description	Parameters
POST	/v2.1/tenants	Create a new tenant.	None

Required request body attributes: code, name, zuora\_account\_name, zuora\_account\_number

### Request body example:

```

{
  "name": "MyNewTenant",
  "code": "mytenant",
  "zuora_account_name": "string",
  "zuora_account_number": "A00000415",
  "description": "DescriptionOfMyTenant"
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. New resource created.",
    "verbose_message": "",
    "code": 201
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed5ac802c356a0001a735af",
        "name": "MyNewTenant",
        "zuora_account_name": "string",
        "zuora_account_number": "A00000415",
        "description": "DescriptionOfMyTenant",
        "code": "mytenant",
        "usage": null
      }
    ]
  }
}

```

## Modify the tenant

Use the method listed in the following table to modify the tenant.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/tenants/{id}	Modify the tenant specified by the ID. You can change the name, the Zuora subscription details (account name or subscription number), and the description of the tenant.	id (string): The unique identifier of the tenant.

Required request body attributes: code

### Request body example:



```

{
  "name": "MyNewTenant",
  "code": "mytenant",
  "zuora_account_name": "string",
  "zuora_account_number": "A00000415",
  "description": "New description of my tenant"
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed5ac802c356a0001a735af",
        "name": "MyNewTenant",
        "zuora_account_name": "string",
        "zuora_account_number": "A00000415",
        "description": "New description of my tenant",
        "code": "mytenant",
        "usage": null
      }
    ]
  }
}

```

### Delete the tenant

Use the method listed in the following table to delete the tenant.

HTTP Method	Path	Description	Parameters
DELETE	/v2.1/tenants/{id}	Delete the tenant specified by the ID.	id (string): The unique identifier of the tenant.

Required request body attributes: none

### Request body example:

none

### Response body example:

No content for successful delete

## Subtenants

Use the methods listed in the following table to retrieve, create, modify, and delete subtenants.

HTTP Method	Path	Description
GET	/v2.1/subtenants	Retrieve subtenants.
GET	/v2.1/subtenants/{id}	Retrieve a subtenant by the subtenant ID.
POST	/v2.1/subtenants`	Create a new subtenant.
PUT	/v2.1/subtenants/{id}	Modify the details of a subtenant. You can modify the name of the subtenant.
DELETE	/v2.1/subtenants/{id}	Delete a subtenant.

### Subtenant attributes

The following table lists the subtenant attributes.

Attribute	Type	Description
id	String	The unique identifier of the subtenant.
name	String	The subtenant name.
code	String	A customer-specified (or default) code that represents the subtenant.
tenant_id	String	The identifier of the tenant to which the subtenant belongs.

### Retrieve all subtenants

Use the method listed in the following table to retrieve all subtenants or a subset of all subtenants. Specifying a `tenant_id` will return only the subtenants belonging to that tenant.

HTTP Method	Path	Description	Parameters
GET	/v2.1/subtenants	Retrieve subtenants.	tenant_id: (Optional) Return the subtenants belonging to the specified tenant. offset and limit– see

Required request body attributes: none

**Request body example:**

```
none
```

**Response body example:**

```

{
  "status": {
    "user_message": "Okay. Returned 2 records.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 2,
    "total_records": 202,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 0,
    "limit": 2,
    "records": [
      {
        "id": "5e7c3af8aab46c00014ce878",
        "description": "",
        "name": "MySubtenant",
        "code": "mysubtenant",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant"
      },
      {
        "id": "5d9144f3869caefed0f39f82",
        "description": "",
        "name": "MySubtenant2",
        "code": "myothersubtenant",
        "tenant_id": "5d914499869caefed0f39eee",
        "tenant": "MyTenant"
      }
    ]
  }
}

```

### Retrieve a subtenant by ID

Use the method listed in the following to retrieve a subtenant by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/subtenants/{id}	Retrieve the subtenant specified by the ID.	id (string): The unique identifier of the subtenant.

Required request body attributes: none

### Request body example:

none

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5e7c3af8aab46c00014ce878",
        "description": "",
        "name": "MySubtenant",
        "code": "subtenantcode",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant"
      }
    ]
  }
}
```

### Create a subtenant

Use the method listed in the following table to create a subtenant.

HTTP Method	Path	Description	Parameters
POST	/v2.1/subtenants	Create a new subtenant.	None

Required request body attributes: name, code, tenant\_id

### Request body example:

```
{
  "name": "MySubtenant",
  "code": "mynewsubtenant",
  "tenant_id": "5ed5ac802c356a0001a735af"
}
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. New resource created.",
    "verbose_message": "",
    "code": 201
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ecefbbef418b40001f20bd6",
        "description": "",
        "name": "MyNewSubtenant",
        "code": "mynewsubtenant",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant"
      }
    ]
  }
}

```

### Modify a subtenant by ID

Use the method listed in the following table to modify a subtenant by ID.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/subtenants/{id}	Modify the subtenant specified by the ID. You can change the subtenant name.	id (string): The unique identifier of the subtenant.

Required request body attributes: name

#### Request body example:

```

{
  "name": "MyModifiedSubtenant"
}

```

#### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ecefbbef418b40001f20bd6",
        "description": "",
        "name": "MyNewSubtenant",
        "code": "mynewsubtenant",
        "tenant_id": "5e7c3af7aab46c00014ce877",
        "tenant": "MyTenant"
      }
    ]
  }
}

```

### Delete a subtenant by ID

Use the method listed in the following table to delete a subtenant by ID.

HTTP Method	Path	Description	Parameters
DELETE	/v2.1/subtenants/{id}	Delete the subtenant specified by the ID.	id (string): The unique identifier of the subtenant.

Required request body attributes: none

#### Request body example:

```
none
```

#### Response body example:

```
No content for succesful delete
```

## Users

Use the methods listed in the following table to retrieve, create, modify, and delete subtenants.

HTTP Method	Path	Description
GET	/v2.1/users	Retrieve a list of all users.
GET	/v2.1/users/{id}	Retrieve a user by the user ID.
POST	/v2.1/users	Create a new user.
PUT	/v2.1/users/{id}	Modify the details of a user.
DELETE	/v2.1/users/{id}	Delete a user.
GET	/v2.1/users/{username}	Retrieve a user by the user name.

## User attributes

The following table lists the user attributes.

Attribute	Type	Description
id	String	The unique identifier of the user.
username	String	The user name.
password	String	The user's password.
firstName	String	The user's first name.
lastName	String	The user's last name.
displayName	String	The user's display name.
email	String	The user's email address.
phone	String	The user's phone number.
profileImageURL	String	The URL of the users' profile image.
tenant_id	String	The primary tenant identifier for this user.
tenancies	–	The tenancy to which the user belongs; an array consisting of: <ul style="list-style-type: none"> <li>• tenant_id, and</li> <li>• role_name This is the user role; one of user, admin, read, partner, or root.</li> </ul>
provider	String	Authentication provider: local or activeDirectory



Attribute	Type	Description
provider data	–	Authentication provider details, consisting of: <ul style="list-style-type: none"> <li>• email_address, and</li> <li>• member_of (group membership)</li> </ul>

### Retrieve all users

Use the method listed in the following table to retrieve all users.

HTTP Method	Path	Description	Parameters
GET	/v2.1/users	Retrieve all users.	None

Required request body attributes: none

### Request body example:

none

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5dddbe0ef071fe0001b889fd",
        "username": "TestUser3",
        "firstName": "Test",
        "lastName": "User",
        "displayName": "",
        "email": "testuser@netapp.com",
        "tenancies": [
          {
            "id": "5d914499869caefed0f39eee",
            "name": "MyOrg",
            "code": "myorg",
            "role": "admin"
          }
        ]
      }
    ]
  }
}

```

## Retrieve a user by ID

Use the method listed in the following table to retrieve a user by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/users	Retrieve a user by ID.	id (string): The unique identifier of the user.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5e585df6896bd80001dd4b44",
        "username": "testuser01",
        "firstName": "",
        "lastName": "",
        "displayName": "",
        "email": "",
        "tenancies": [
          {
            "id": "5d914499869caefed0f39eee",
            "name": "MyOrg",
            "code": "myorg",
            "role": "user"
          }
        ]
      }
    ]
  }
}

```

### Retrieve a user by user name

Use the method listed in the following table to retrieve a user by the user name.

HTTP Method	Path	Description	Parameters
GET	/v2.1/users	Retrieve a user by user name.	username (string): The user name of the user.

Required request body attributes: none

#### Request body example:

none

#### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5e61aa814559c20001df1a5f",
        "username": "MyName",
        "firstName": "MyFirstName",
        "lastName": "MySurname",
        "displayName": "CallMeMYF",
        "email": "user@example.com",
        "tenancies": [
          {
            "id": "5e5f1c4f253c820001877839",
            "name": "MyTenant",
            "code": "testtenantmh",
            "role": "user"
          }
        ]
      }
    ]
  }
}

```

### Create a user

Use the method listed in the following table to create a user.

HTTP Method	Path	Description	Parameters
POST	/v2.1/users	Create a new user.	None

Required request body attributes: username, tenant\_id, tenancies, provider

### Request body example:

```
{
  "username": "MyUser",
  "password": "mypassword",
  "firstName": "My",
  "lastName": "User",
  "displayName": "CallMeMyUser",
  "email": "user@example.com",
  "phone": "string",
  "profileImageURL": "string",
  "tenant_id": "5e7c3af7aab46c00014ce877",
  "tenancies": [
    {
      "tenant_id": "5e7c3af7aab46c00014ce877",
      "role_name": "admin"
    }
  ],
  "provider": "local",
  "provider_data": {
    "email": "user@example.com",
    "member_of": "string"
  }
}
```

**Response body example:**

```

{
  "status": {
    "user_message": "Okay. New resource created.",
    "verbose_message": "",
    "code": 201
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ed6f463129e5d000102f7e1",
        "username": "MyUser",
        "firstName": "My",
        "lastName": "User",
        "displayName": "CallMeMyUser",
        "email": "user@example.com",
        "tenancies": [
          {
            "id": "5e7c3af7aab46c00014ce877",
            "name": "MyTenant",
            "code": "mytenantcode",
            "role_name": "admin"
          }
        ]
      }
    ]
  }
}

```

## Modify a user by ID

Use the method listed in the following table to modify a user by user ID.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/users/{id}	Modify a user identified by the user ID. You can modify the user name, display name, password, email address, phone number, profile image URL, and tenancy details.	id (string): The unique identifier of the user.

Required request body attributes: none

### Request body example:

```
{
  "password": "MyNewPassword",
  "firstName": "MyFirstName",
  "lastName": "MySurname",
  "displayName": "CallMeMYF",
  "email": "user@example.com",
  "phone": "string",
  "profileImageURL": "string",
  "tenant_id": "5e5f1c4f253c820001877839",
  "tenancies": [
    {
      "tenant_id": "5e5f1c4f253c820001877839",
      "role_name": "user"
    }
  ]
}
```

**Response body example:**

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5e61aa814559c20001df1a5f",
        "username": "MyName",
        "firstName": "MyFirstName",
        "lastName": "MySurname",
        "displayName": "CallMeMYF",
        "email": "user@example.com",
        "tenancies": [
          {
            "id": "5e5f1c4f253c820001877839",
            "name": "MyTenant",
            "code": "testtenantmh",
            "role": "user"
          }
        ]
      }
    ]
  }
}

```

### Delete a user by ID

Use the method listed in the following table to delete a user by ID.

HTTP Method	Path	Description	Parameters
DELETE	/v2.1/users/{name}	Delete the user identified by the ID.	id (string): The unique identifier of the user.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:



No content for succesful delete

## Zones

Use the methods listed in the following table to create, modify, and delete zones. For APIs that allow you to retrieve zones, see the (Consumer) Administration APIs.

HTTP Method	Path	Description
POST	/v2.1/zones	Create a new zone.
PUT	/v2.1/zones/{name}	Modify the details of a zone.
DELETE	/v2.1/zones/{name}	Delete a zone.

### Zone attributes

The following table lists the zone attributes.

Attribute	Type	Description
id	String	the unique identifier of the zone.
name	String	The zone name.
description	String	The description of the zone.
region_name	String	The name of the region in which the zone resides.

### Retrieve all zones

Use the method listed in the following table to retrieve all zones or a subset of zones. Specifying a region will return only the block stores belonging to that tenant.

HTTP Method	Path	Description	Parameters
GET	/v2.1/zones	Retrieve zones.	(Optional) Region name (string) offset and limit – see <a href="#">Common Pagination</a>

Required request body attributes: none

### Request body example:

none

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "total_records": 5,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 2,
    "limit": 1,
    "records": [
      {
        "id": "5ce5d919b68d3b82dc34bef2",
        "name": "au-west1-a",
        "description": "au-west1-a",
        "region": "au-west1"
      }
    ]
  }
}

```

### Retrieve a zone by name

Use the method listed in the following table to retrieve a zone by the zone name.

HTTP Method	Path	Description	Parameters
GET	/v2.1/zones/{name}	Retrieve a zone by name.	Name (string): Zone name

Required request body attributes: none

#### Request body example:

```
none
```

#### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5ce5d919b68d3b82dc34bef2",
        "name": "au-west1-a",
        "description": "au-west1-a",
        "region": "au-west1"
      }
    ]
  }
}

```

### Create a zone

Use the method listed in the following table to create a zone.

HTTP Method	Path	Description	Parameters
POST	/v2.1/zones	Create a new zone within a region.	None

Required request body attributes: name, description, region\_name

### Request body example:

```

{
  "name": "MyZoneName",
  "description": "DescriptionOfMyZone",
  "region_name": "MyRegionName"
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. New resource created.",
    "verbose_message": "",
    "code": 201
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5e61741c9b64790001fe9663",
        "name": "MyZoneName",
        "description": "DescriptionOfMyZone",
        "region": "MyRegionName"
      }
    ]
  }
}

```

### Modify a zone

Use the method listed in the following table to modify a zone.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/zones{name}	Modify a zone identified by name.	name (string): Name of the zone.

Required request body attributes: none

### Request body example:

```

{
  "name": "MyZoneName",
  "description": "NewDescriptionOfMyZone"
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5e61741c9b64790001fe9663",
        "name": "MyZoneName",
        "description": "NewDescriptionOfMyZone",
        "region": "MyRegionName"
      }
    ]
  }
}

```

### Delete a zone

Use the method listed in the following table to delete a zone.

HTTP Method	Path	Description	Parameters
DELETE	/v2.1/zones{name}	Delete a single zone identified by name. All storage resources within a zone must be deleted first.	name (string): Name of the zone.

Required request body attributes: none

#### Request body example:

```
none
```

#### Response body example:

No content to return on a successful deletion.

## Regions

Use the methods listed in the following table to create, modify, and delete regions. For APIs that allow you to retrieve regions, see the (Consumer) Administration APIs.

HTTP Method	Path	Description
GET	/v2.1/regions	Get regions.
GET	/v2.1/regions/{name}	Get regions by name.
POST	/v2.1/regions	Create a new region.
PUT	/v2.1/regions/{name}	Modify the details of a region.
DELETE	/v2.1/regions/{name}	Delete a region.

### Region attributes

The following table lists the region attributes.

Attribute	Type	Description
id	String	The unique identifier of the region.
name	String	The region name.
description	String	The description of the region.

### Retrieve regions

Use the method listed in the following table to retrieve all regions or a subset of regions.

HTTP Method	Path	Description	Parameters
GET	/v2.1/regions	Retrieve regions.	offset and limit– see <a href="#">Common Pagination</a>

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "total_records": 4,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 0,
    "limit": 1,
    "records": [
      {
        "id": "5e7bf44daab46c00014ce77f",
        "name": "au-east8",
        "description": "This is the new region description",
        "zones": []
      }
    ]
  }
}

```

### Retrieve a region by name

Use the method listed in the following table to retrieve a region by name.

HTTP Method	Path	Description	Parameters
GET	/v2.1/regions/{name}	Retrieve a region by name.	name (string): The region name.

Required request body attributes: none

#### Request body example:

```
none
```

#### Response body example:

```

{
  "status": {
    "user_message": "string",
    "verbose_message": "string",
    "code": "string"
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5d2fb0fb4f47df00015274e3",
        "name": "au-east1",
        "description": "string",
        "zones": [
          "au-east1-a"
        ]
      }
    ]
  }
}

```

## Create a region

Use the method listed in the following table to create a region.

HTTP Method	Path	Description	Parameters
POST	/v2.1/regions	Create a new region.	None

Required request body attributes: name

### Request body example:

```

{
  "name": "MyRegionName",
  "description": "DescriptionOfMyRegion"
}

```

### Response body example:



```

{
  "status": {
    "user_message": "Okay. New resource created.",
    "verbose_message": "",
    "code": 201
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5e616f849b64790001fe9658",
        "name": "MyRegionName",
        "Description": "DescriptionOfMyRegion",
        "user_id": "5bbee380a2df7a04d43acae",
        "created": "0001-01-01T00:00:00Z",
        "tags": null
      }
    ]
  }
}

```

## Modify a region

Use the method listed in the following table to modify a region.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/regions/{name}	Modify a region identified by name. You can change the name and description of the region.	name (string): The name of the region.

Required request body attributes: none

### Request body example:

```

{
  "name": "MyRegionName",
  "description": "NewDescriptionOfMyRegion"
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "id": "5e616f849b64790001fe9658",
        "name": "MyRegionName",
        "description": "NewDescriptionOfMyRegion",
        "zones": []
      }
    ]
  }
}

```

## Delete a region

Use the method listed in the following table to delete a region.

HTTP Method	Path	Description	Parameters
DELETE	/v2.1/regions{name}	Delete a single region identified by name. All zones within a region must be deleted first.	Name (string): The name of the region.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
No content for succesful delete
```

## ONTAP clusters

Use the methods listed in the following table to retrieve, create, modify, and delete ONTAP clusters.

HTTP Method	Path	Description
GET	/v2.1/ontapclusters	Retrieve all ONTAP clusters.
GET	/v2.1/ontapclusters/{id}	Retrieve an ONTAP cluster by ID.
POST	/v2.1/ontapclusters	Create a new ONTAP cluster.
PUT	/v2.1/ontapclusters/{id}	Update ONTAP cluster inventory by ID.
DELETE	/v2.1/ontapclusters/{id}	Delete an ONTAP cluster.

### ONTAP clusters attributes

The following table lists the ONTAP cluster attributes.

Attribute	Type	Description
id	String	The unique identifier of the ONTAP cluster.
name	String	The ONTAP cluster name.
uuid	String	The ONTAP cluster universal unique identifier (UUID).
management_ip	String	The ONTAP cluster management IPv4 address.
username	String	The ONTAP cluster name
password	String	The ONTAP cluster password
provisioning_state	String	Identifies whether a cluster is available for provisioning operations. The options include: <ul style="list-style-type: none"> <li>• Open</li> <li>• Closed</li> </ul>
data_network_ip_cidr	String	The CIDR notation of a subnet.
data_network_default_gateway	String	The IPV4 address.

Attribute	Type	Description
data_network_ports	-	<p>A list of the ONTAP cluster data network ports.</p> <pre data-bbox="1047 262 1485 766"> For example: [   {     "node_name": "dev-ots-per01-01",     "port_name": "e0c-120",     "parent_port": "e0c"   } ] </pre>
intercluster_lifs	-	<p>The ONTAP cluster intercluster LIFs.</p> <pre data-bbox="1047 913 1485 1816"> For example: [   {     "name": "peer1",     "node": "aff-01",     "port": "a0a-103",     "address": "10.128.113.232",     "netmask": "255.255.255.0"   },   {     "name": "peer2",     "node": "aff-02",     "port": "a0a-103",     "address": "10.128.113.233",     "netmask": "255.255.255.0"   } ] </pre>

Attribute	Type	Description
svm_root_service_level	String	The ONTAP cluster storage virtual machine (SVM) root service level name. Applicable values are Standard, Extreme, or Premium. This service level is assigned by default to all SVMs created under the cluster. The cluster should have an associated aggregate for the service level mentioned.
zone	String	The zone name.
Subscription_number	String	The Zuora subscription.
services_available	–	List of services available and their state. For example: <div data-bbox="1047 730 1487 991" style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <pre>{   "fcg": false,   "iscsi": true,   "nas": true }</pre> </div>
data_fcp_ports	–	List of nodes and ports for FCP-enabled ONTAP clusters. For example: <div data-bbox="1047 1171 1487 1516" style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <pre>[ { "node_name":   "aff-01",   "port_name": "0g" },   { "node_name":   "aff-01",   "port_name": "0h"   } ]</pre> </div>
is_mcc	Boolean	Indicates whether the cluster is MetroCluster enabled or not. The default is False.
mcc_partner_cluster	String	The identifier of the partner cluster of the current cluster in a MetroCluster pair. Required if the cluster is enabled for MetroCluster.

## Retrieve all ONTAP clusters

Use the method listed in the following to retrieve all ONTAP clusters or a subset of ONTAP clusters.

HTTP Method	Path	Description	Parameters
GET	/v2.1/ontapclusters	Retrieve all ONTAP clusters.	offset and limit – see <a href="#">Common Pagination</a>

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 2 records.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 2,
    "total_records": 5,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 3,
    "limit": 2,
    "records": [
      {
        "id": "5c5bb9f16680a7002a5f7450",
        "name": "dev-ots-per01",
        "region": "au-west1",
        "zone": "au-west1-a",
        "uuid": "63053baa-ada4-11ea-b197-005056a4c0ef",
        "management_ip": "10.128.115.173",
        "username": "admin",
        "services_available": {
          "fcp": false,
          "iscsi": true,
          "nas": true
        },
        "provisioning_state": "open",
        "data_network_ports": [
          {
```

```

        "node_name": "dev-ots-per01-01",
        "port_name": "e0c-120",
        "parent_port": "e0c"
    }
],
"data_network_ip_cidr": "10.96.120.0/24",
"data_network_default_gateway": "10.96.120.1",
"svm_root_service_level": "performance",
"intercluster_lifs": [
    {
        "name": "dev-ots-per01-01-icl01",
        "node": "dev-ots-per01-01",
        "port": "e0b",
        "address": "10.128.115.144",
        "netmask": "255.255.255.0"
    }
],
"subscription_number": "A-S00003875",
"created": "2019-02-22T03:38:38.867Z",
"data_fcp_ports": []
},
{
    "id": "5eaf5249f038943eb46b6608",
    "name": "aff",
    "region": "au-east1",
    "zone": "au-east1-b",
    "uuid": "62d649d2-07a1-11e6-9549-00a0985c0dcb",
    "management_ip": "10.128.113.69",
    "username": "admin",
    "services_available": {
        "fcp": true,
        "iscsi": true,
        "nas": true
    },
    "provisioning_state": "open",
    "data_network_ports": [
        {
            "node_name": "aff-01",
            "port_name": "a0a-2000",
            "parent_port": "a0a"
        },
        {
            "node_name": "aff-02",
            "port_name": "a0a-2000",
            "parent_port": "a0a"
        }
    ]
}

```

```

],
"data_network_ip_cidr": "10.50.50.0/24",
"data_network_default_gateway": "10.50.50.1",
"svm_root_service_level": "premium",
"intercluster_lifs": [
  {
    "name": "peer1",
    "node": "aff-01",
    "port": "a0a-103",
    "address": "10.128.113.232",
    "netmask": "255.255.255.0"
  },
  {
    "name": "peer2",
    "node": "aff-02",
    "port": "a0a-103",
    "address": "10.128.113.233",
    "netmask": "255.255.255.0"
  }
],
"subscription_number": "A-S00004635",
"created": "2019-02-22T03:38:38.867Z",
"data_fcp_ports": [
  {
    "node_name": "aff-01",
    "port_name": "0g"
  },
  {
    "node_name": "aff-01",
    "port_name": "0h"
  },
  {
    "node_name": "aff-02",
    "port_name": "0g"
  },
  {
    "node_name": "aff-02",
    "port_name": "0h"
  }
],
"is_mcc": false,
"created": "1995-09-07T10:40:52Z"
}
]
}
}

```



## Retrieve ONTAP clusters by ID

Use the method listed in the following to retrieve ONTAP clusters by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/ontapclusters /{id}	Retrieve the ONTAP clusters identified by ID.	id (string): The unique identifier of ONTAP cluster.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5c5bb9f16680a7002a5f7450",
        "name": "dev-ots-per01",
        "region": "au-west1",
        "zone": "au-west1-a",
        "uuid": "63053baa-ada4-11ea-b197-005056a4c0ef",
        "management_ip": "10.128.115.173",
        "username": "admin",
        "services_available": {
          "fcp": false,
          "iscsi": true,
          "nas": true
        },
        "provisioning_state": "open",
        "data_network_ports": [
          {
            "node_name": "dev-ots-per01-01",
            "port_name": "e0c-120",
            "parent_port": "e0c"
          }
        ]
      }
    ]
  }
}
```

```

    ],
    "data_network_ip_cidr": "10.96.120.0/24",
    "data_network_default_gateway": "10.96.120.1",
    "svm_root_service_level": "performance",
    "intercluster_lifs": [
      {
        "name": "dev-ots-per01-01-icl01",
        "node": "dev-ots-per01-01",
        "port": "e0b",
        "address": "10.128.115.144",
        "netmask": "255.255.255.0"
      }
    ],
    "subscription_number": "A-S00003875",
    "created": "2019-02-22T03:38:38.867Z",
    "data_fcp_ports": [],
    "is_mcc": false,
    "created": "1995-09-07T10:40:52Z"
  }
]
}
}
}

```

### Create ONTAP cluster

Use the API in the following table to create an ONTAP cluster.

ONTAP clusters are always created with the iSCSI service enabled. Optionally, the FCP service can be enabled if the infrastructure supports it.

HTTP Method	Path	Description	Parameters
POST	/v2.1/ontapclusters	Create an ONTAP cluster.	None

Required request body attributes: name, uuid, management\_ip, username, password, data\_network\_ip\_cidr, data\_network\_default\_gateway, intercluster\_lifs, zone

If FCP is enabled (using the services\_available FCP attribute), the `data\_fcp\_ports` are required.

If is\_mcc is true, the mcc\_partner\_cluster is required.

### Request body example:

```

{
  "name": "clustername",
  "uuid": "49b6e08e-513a-11ea-b197-005056a4c0ef",
  "management_ip": "10.128.112.165",
  "username": "admin",
  "password": "ClusterPassword",
  "provisioning_state": "open",
  "data_network_ip_cidr": "10.96.112.0/24",
  "data_network_default_gateway": "10.96.112.1",
  "data_network_ports": [
    {
      "node_name": "clustername-01",
      "port_name": "e0c-112",
      "parent_port": "e0c"
    }
  ],
  "intercluster_lifs": [
    {
      "name": "clustername-01-icl01",
      "node": "clustername-01",
      "port": "e0b",
      "address": "10.128.112.222",
      "netmask": "255.255.255.0"
    }
  ],
  "svm_root_service_level": "extreme",
  "zone": "MyZone",
  "subscription_number": "",
  "services_available": {
    "fcp": false,
    "iscsi": true,
    "nas": true
  },
  "data_fcp_ports": [
  ],
  "is_mcc": false,
}

```

#### Response body example:

```

{
  "status": {
    "user_message": "Okay. New resource created.",
    "verbose_message": "",
    "code": 201
  }
}

```

```

},
"result": {
  "returned_records": 1,
  "records": [
    {
      "id": "5ef155b8f5591100010a75c5",
      "name": "clustername",
      "region": "MyRegion",
      "zone": "MyZone",
      "uuid": "49b6e08e-513a-11ea-b197-005056a4c0ef",
      "management_ip": "10.128.112.165",
      "username": "admin",
      "services_available": {
        "fcp": false,
        "iscsi": true,
        "nas": true
      },
      "provisioning_state": "open",
      "data_network_ports": [
        {
          "node_name": "clustername-01",
          "port_name": "e0c-112",
          "parent_port": "e0c"
        }
      ],
      "data_network_ip_cidr": "10.96.112.0/24",
      "data_network_default_gateway": "10.96.112.1",
      "svm_root_service_level": "extreme",
      "intercluster_lifs": [
        {
          "name": "clustername-01-icl01",
          "node": "clustername-01",
          "port": "e0b",
          "address": "10.128.112.222",
          "netmask": "255.255.255.0"
        }
      ],
      "subscription_number": "",
      "created": "2020-06-23T01:07:04.563Z",
      "data_fcp_ports": [],
      "is_mcc": false,
      "mcc_partner_cluster": "5d2fb0fb4f47df00015274e3",
      "created": "1995-09-07T10:40:52Z"
    }
  ]
}

```

```
}
```

## Modify ONTAP cluster

Use the method listed in the following to modify the ONTAP cluster.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/ontapclusters /{id}	Modify the details of the ONTAP cluster identified by ID.	id (string): The unique identifier of ONTAP cluster.

Required request body attributes: none

**Request body example:**

```

{
  "name": "clustername",
  "uuid": "49b6e08e-513a-11ea-b197-005056a4c0ef",
  "management_ip": "10.128.112.165",
  "username": "admin",
  "password": "ClusterPassword",
  "provisioning_state": "open",
  "data_network_ip_cidr": "10.96.112.0/24",
  "data_network_default_gateway": "10.96.112.1",
  "data_network_ports": [
    {
      "node_name": "dev-ots-syd01-01",
      "port_name": "e0c-112",
      "parent_port": "e0c"
    }
  ],
  "intercluster_lifs": [
    {
      "name": "dev-ots-syd01-01-icl01",
      "node": "dev-ots-syd01-01",
      "port": "e0b",
      "address": "10.128.112.222",
      "netmask": "255.255.255.0"
    }
  ],
  "svm_root_service_level": "standard",
  "zone": "MyZone",
  "subscription_number": "",
  "services_available": {
    "fcp": false,
    "iscsi": true,
    "nas": false
  },
  "data_fcp_ports": [
  ]
}

```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Accepted for processing.",
    "verbose_message": "",
    "code": 202
  },

```

```

"result": {
  "returned_records": 1,
  "records": [
    {
      "id": "5ef155b8f5591100010a75c5",
      "name": "clustername",
      "region": "MyRegion",
      "zone": "MyZone",
      "uuid": "49b6e08e-513a-11ea-b197-005056a4c0ef",
      "management_ip": "10.128.112.165",
      "username": "admin",
      "services_available": {
        "fcg": false,
        "iscsi": true,
        "nas": true
      },
      "provisioning_state": "open",
      "data_network_ports": [
        {
          "node_name": "dev-ots-syd01-01",
          "port_name": "e0c-112",
          "parent_port": "e0c"
        }
      ],
      "data_network_ip_cidr": "10.96.112.0/24",
      "data_network_default_gateway": "10.96.112.1",
      "svm_root_service_level": "standard",
      "intercluster_lifs": [
        {
          "name": "dev-ots-syd01-01-icl01",
          "node": "dev-ots-syd01-01",
          "port": "e0b",
          "address": "10.128.112.222",
          "netmask": "255.255.255.0"
        }
      ],
      "subscription_number": "",
      "created": "2020-06-23T01:07:04.563Z",
      "data_fcg_ports": [],
      "is_mcc": false,
      "mcc_partner_cluster": "5d2fb0fb4f47df00015274e3",
      "created": "1995-09-07T10:40:52Z"
    }
  ]
}

```

## Delete an ONTAP cluster

Use the method listed in the following table to delete an ONTAP cluster.

HTTP Method	Path	Description	Parameters
DELETE	/v2.1/ontapclusters /{id}	Delete the ONTAP cluster identified by ID.	id (string): The unique identifier of the ONTAP cluster.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
No content for succesful delete
```

## StorageGRID instances

Use the methods listed in the following table to set up and manage StorageGRID instances to support object storage.

HTTP Method	Path	Description
GET	/v2.1/storagegridinstances	Retrieve StorageGRID instances.
POST	/v2.1/storagegridinstances	Create a new StorageGRID instance to host object storage.
GET	/v2.1/storagegridinstances /{id}	Retrieve a StorageGRID instance by ID.
PUT	/v2.1/storagegridinstances /{id}	Update a StorageGRID instance.

### StorageGRID instance attributes

The following table lists the StorageGRID instance attributes.

Attribute	Type	Description
id	String	The unique identifier for the StorageGRID instance.
name	String	The StorageGRID instance name.
admin_rest_uri	String(\$uri)	The StorageGRID admin node endpoint.



Attribute	Type	Description
s3_endpoints	String(\$uri)	The StorageGRID endpoints. For example: [ "https://s3.examplegrid.com", "https://s3.location.company.com" ]
grid_username	String	The StorageGRID user name.
grid_password	String	The StorageGRID password.
tenant_username	String	The StorageGRID tenant user name.
tenant_password	String	The StorageGRID tenant password.
subscription_number	String	The Zuora subscription number.

### Retrieve StorageGRID instances

Use the method listed in the following table to retrieve the StorageGRID instances.

HTTP Method	Path	Description	Parameters
GET	/v2.1/storagegridinstances	Retrieve StorageGRID instances.	None

Required request body attributes: none

#### Request body example:

none

#### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5e3ba2840271823644cb8ab6",
        "name": "NSE StorageGRID Dev1",
        "user_id": "5bbee380a2df7a04d43acae",
        "admin_rest_uri": "https://sggmi-dev.dev.ausngs.netapp.au",
        "s3_endpoints": [
          "https://sgs3.dev.ausngs.netapp.au"
        ],
        "subscription_number": "A-S00004566"
      }
    ]
  }
}

```

### Retrieve StorageGRID instances by ID

Use the method listed in the following table to retrieve StorageGRID instances by ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/storagegridinstances/{id}	Retrieve a StorageGRID instance by ID.	id (string): The unique identifier of the StorageGRID instance.

Required request body attributes: none

#### Request body example:

```
none
```

#### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5e3ba2840271823644cb8ab6",
        "name": "NSE StorageGRID Dev1",
        "user_id": "5bbee380a2df7a04d43acae",
        "admin_rest_uri": "https://sggmi-dev.dev.ausngs.netapp.au",
        "s3_endpoints": [
          "https://sgs3.dev.ausngs.netapp.au"
        ],
        "subscription_number": "A-S00004566"
      }
    ]
  }
}

```

### Create a StorageGRID instance by ID

Use the method listed the following table to create a StorageGRID instance by ID.

HTTP Method	Path	Description	Parameters
POST`	/v2.1/storagegridinstances/{id}	Retrieve a StorageGRID instance by ID.	id (string): The unique identifier of the StorageGRID instance.

Required request body attributes: none

#### Request body example:

```

{
  "name": "Grid1",
  "admin_rest_uri": "https://examplegrid.com",
  "s3_endpoints": [
    "https://s3.examplegrid.com",
    "https://s3.location.company.com"
  ],
  "grid_username": "root",
  "grid_password": "string",
  "tenant_username": "root",
  "tenant_password": "string",
  "subscription_number": "A-S00003969"
}

```

### Response body example:

```

{
  "status": {
    "user_message": "string",
    "verbose_message": "string",
    "code": "string"
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5d2fb0fb4f47df00015274e3",
        "name": "Grid1",
        "admin_rest_uri": "https://examplegrid.com",
        "user_id": "5d2fb0fb4f47df00015274e3",
        "s3_endpoints": [
          "https://s3.examplegrid.com",
          "https://s3.location.company.com"
        ],
        "subscription_number": "A-S00003969"
      }
    ]
  }
}

```

### Modify a StorageGRID instance by ID

Use the method listed in the following table to modify a StorageGRID instance by ID.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/storagegridinstances/{id}	Modify a StorageGRID instance by ID.	id (string): The unique identifier of the StorageGRID instance.

Required Request Body attributes: none

**Request body example:**

```
{
  "name": "Grid1",
  "admin_rest_uri": "https://examplegrid.com",
  "s3_endpoints": [
    "https://s3.examplegrid.com",
    "https://s3.location.company.com"
  ],
  "grid_username": "root",
  "grid_password": "string",
  "tenant_username": "root",
  "tenant_password": "string",
  "subscription_number": "A-S00003969"
```

**Response body example:**

```

{
  "status": {
    "user_message": "string",
    "verbose_message": "string",
    "code": "string"
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "5d2fb0fb4f47df00015274e3",
        "name": "Grid1",
        "admin_rest_uri": "https://examplegrid.com",
        "user_id": "5d2fb0fb4f47df00015274e3",
        "s3_endpoints": [
          "https://s3.examplegrid.com",
          "https://s3.location.company.com"
        ],
        "subscription_number": "A-S00003969"
      }
    ]
  }
}

```

## Service levels

Use the methods listed in the following table to retrieve, create, modify, and delete service levels.

HTTP Method	Path	Description
GET	/v2.1/servicelevels	Retrieve all service levels.
GET	/v2.1/servicelevels/{id}	Retrieve a service level by ID.
POST	/v2.1/servicelevels	Create a new service level.
PUT	/v2.1/servicelevels/{id}	Modify service level details by ID.
DELETE	/v2.1/servicelevels/{id}	Delete a service level.

## Service level attributes

The following table lists the service level attributes.

Attribute	Type	Description
id	String	The unique identifier for the service level.

<b>Attribute</b>	<b>Type</b>	<b>Description</b>
name	String	The service level name.
description	String	The description of the service level.
policy_name	String	Service level quality of service (QoS) name. Allowed values: nse_value, nse_standard, nse_performance, and nse_extreme.
available	Boolean	Indicates whether the service level is available for use.
grandfathered	Boolean	Indicates whether the service level has been retired.
peak_iops_per_tb	Integer	The maximum possible IOPS per TiB.
expected_iops_per_tb	Integer	The minimum expected IOPS per TiB.
absolute_min_iops	Integer	The absolute minimum IOPS which is used as an override when the expected IOPS is less than this value.
peak_iops_allocation	String	The peak IOPS allocation. Allowed values: allocated_space and used_space.
io_block_size_kb	Integer	The I/O block size (KiB).
min_size_gb	Integer	The minimum size (GiB).
max_size_gb	Integer	The maximum size (GiB).
max_peak_iops	Integer	The maximum peak IOPS for the service level.
max_expected_iops	Integer	The maximum expected IOPS.
autogrow_max_percent	Integer	The autogrow maximum percent.

Attribute	Type	Description
ontap_aggregates	-	<p>The list of ONTAP aggregates. An ONTAP aggregate consists of:</p> <ul style="list-style-type: none"> <li>• cluster_UIID: ONTAP cluster UUID</li> <li>• aggr_name: ONTAP aggregate name</li> <li>• aggr_uuid: ONTAP aggregate UUID</li> <li>• node_name: ONTAP cluster node name</li> </ul> <p>For example:</p> <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;">"ontap_aggregates": [   {     "cluster_uuid":     "3fa85f64-5717-4562-     b3fc-2c963f66afa6",     "aggr_name":     "string",     "aggr_uuid":     "3fa85f64-5717-4562-     b3fc-2c963f66afa6",     "node_name":     "node01"   } ]</pre>



Attribute	Type	Description
primary_volume_defaults	-	<ul style="list-style-type: none"> <li>snapshot_auto_delete_target_free_space (integer): This option specifies the free space percentage at which the automatic deletion of Snapshot copies must stop.</li> <li>auto_size_mode (string): The autosize mode for the volume. Allowed values: off, grow, grow_shrink For example:</li> </ul> <pre>"primary_volume_defaults": {   "snapshot_auto_delete_target_free_space":   3,   "auto_size_mode":   "grow_shrink" }</pre>

### Retrieve service levels

Use the method listed in the following table to retrieve all service levels.

HTTP Method	Path	Description	Parameters
GET	/v2.1/servicelevels	Retrieve all service levels.	None

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 3 records.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
```

```

"total_records": 3,
"records": [
  {
    "name": "standard",
    "description": "Best suited for general purpose workloads",
    "slo": "1000IOPS/TB",
    "min_size": 137438953472,
    "io_block_size_kb": 32,
    "min_size_gb": 10,
    "max_size_gb": 40960,
    "min_iops": 100,
    "peak_iops_per_tb": 1000,
    "expected_iops_per_tb": 700,
    "max_peak_iops": 40000,
    "max_expected_iops": 28000,
    "max_peak_throughput": 1250,
    "max_expected_throughput": 875
  },
  {
    "name": "extreme",
    "description": "Best suited for performance-critical workloads",
    "slo": "12000IOPS/TB",
    "min_size": 91625968981,
    "io_block_size_kb": 32,
    "min_size_gb": 10,
    "max_size_gb": 10240,
    "min_iops": 500,
    "peak_iops_per_tb": 12000,
    "expected_iops_per_tb": 8000,
    "max_peak_iops": 120000,
    "max_expected_iops": 60000,
    "max_peak_throughput": 3750,
    "max_expected_throughput": 1875
  },
  {
    "name": "premium",
    "description": "Best suited for databases and high performance
workloads",
    "slo": "4000IOPS/TB",
    "min_size": 137438953472,
    "io_block_size_kb": 32,
    "min_size_gb": 10,
    "max_size_gb": 10240,
    "min_iops": 300,
    "peak_iops_per_tb": 4000,
    "expected_iops_per_tb": 3000,

```

```
    "max_peak_iops": 40000,  
    "max_expected_iops": 30000,  
    "max_peak_throughput": 1250,  
    "max_expected_throughput": 937  
  }  
]  
}  
}
```

### Retrieve service levels by name

Use the method listed in the following table to retrieve service levels by name.

HTTP Method	Path	Description	Parameters
GET	/v2.1/servicelevels /{name}	Retrieve a service level by name.	name (string): The name of the service level.

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "name": "premium",
        "description": "Best suited for databases and high performance
workloads",
        "slo": "4096IOPS/TB",
        "min_size": 137438953472,
        "io_block_size_kb": 32,
        "min_size_gb": 10,
        "max_size_gb": 10240,
        "min_iops": 300,
        "peak_iops_per_tb": 4096,
        "expected_iops_per_tb": 3000,
        "max_peak_iops": 40000,
        "max_expected_iops": 30000,
        "max_peak_throughput": 1250,
        "max_expected_throughput": 937
      }
    ]
  }
}

```

### Create a service level

Use the method listed in the following table to create a service level.

HTTP Method	Path	Description	Parameters
POST	/v2.1/servicelevels	Create a service level.	None

Required request body attributes: name, policy\_name

### Request body example:

```
{
  "name": "MyServiceLevelName",
  "description": "My new service level description",
  "policy_name": "nse_value",
  "available": true,
  "grandfathered": false,
  "peak_iops_per_tb": 1000,
  "expected_iops_per_tb": 700,
  "absolute_min_iops": 100,
  "peak_iops_allocation": "allocated_space",
  "io_block_size_kb": 32,
  "min_size_gb": 10,
  "max_size_gb": 40960,
  "max_peak_iops": 20000,
  "max_expected_iops": 5000,
  "autogrow_max_percent": 3,
  "ontap_aggregates": [
    {
      "cluster_uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
      "aggr_name": "string",
      "aggr_uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
      "node_name": "node01"
    }
  ],
  "primary_volume_defaults": {
    "snapshot_auto_delete_target_free_space": 3,
    "auto_size_mode": "grow_shrink"
  }
}
```

**Response body example:**

```

{
  "status": {
    "user_message": "Okay. New resource created.",
    "verbose_message": "",
    "code": 201
  },
  "result": {
    "total_records": 1,
    "records": [
      {
        "name": "MyServiceLevelName",
        "description": "My new service level description",
        "slo": "1000IOPS/TB",
        "min_size": 0,
        "io_block_size_kb": 32,
        "min_size_gb": 10,
        "max_size_gb": 40960,
        "min_iops": 100,
        "peak_iops_per_tb": 1000,
        "expected_iops_per_tb": 700,
        "max_peak_iops": 20000,
        "max_expected_iops": 5000,
        "max_peak_throughput": 625,
        "max_expected_throughput": 156
      }
    ]
  }
}

```

## Modify a service level

Use the method listed in the following table to modify a service level.

HTTP Method	Path	Description	Parameters
PUT	/v2.1/servicelevels /{name}	Modify the details of a service level.	name (string): The name of the service level.

Required request body attributes: none

### Request body example:

```

{
  "name": "MyNewServiceLevelName",
  "description": "Service level description",
  "policy_name": "nse_value",
  "available": false,
  "grandfathered": false,
  "peak_iops_per_tb": 1000,
  "expected_iops_per_tb": 700,
  "absolute_min_iops": 100,
  "peak_iops_allocation": "allocated_space",
  "io_block_size_kb": 32,
  "min_size_gb": 10,
  "max_size_gb": 40960,
  "max_peak_iops": 20000,
  "max_expected_iops": 5000,
  "autogrow_max_percent": 3,
  "ontap_aggregates": [
    {
      "cluster_uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
      "aggr_name": "string",
      "aggr_uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6",
      "node_name": "node01"
    }
  ],
  "primary_volume_defaults": {
    "snapshot_auto_delete_target_free_space": 3,
    "auto_size_mode": "grow_shrink"
  }
}

```

### Response body example:

TBA

### Delete service level by ID

Use the method listed in the following table to delete a service level by ID.

HTTP Method	Path	Description	Parameters
DELETE	/v2.1/servicelevels/{name}	Delete the service level identified by ID.	name (string): The name of the service level.

### Request body example:

none

### Response body example:

No content for succesful delete

## Service requests

Use the methods listed in the following table to create and retrieve service requests.

HTTP Method	Path	Description
GET	/v2.1/tenants/{tenant_id}/servicerequests	Retrieve service requests.
GET	/v2.1/tenants/{tenant_id}/servicerequests/{id}	Retrieve a service request by ID.
POST	/v2.1/tenants/{tenant_id}/servicerequests/	Create a service request.
GET	/v2.1/tenants/{tenant_id}/servicerequests/categories	Retrieve service request categories.

### Service request attributes

The following table lists the service request attributes.

Attribute	Type	Description
Id	String	An identifier for the service request. For example: SRQ0035952014.
subject	String	The subject of the service request.
comment	String	A comment on the service request
category	String	The category of the request: Backup, disaster recovery, technical, other, or subscription.
priority	String	The priority of the service request: very low, low, normal, high, or urgent.
subscription	String	The Zuora subscription number.



Attribute	Type	Description
commitment	–	Subscription commitment details, including: <ul style="list-style-type: none"> <li>• Service level (string): Service level: Standard, Premium, or Extreme.</li> <li>• New_commitment (integer): New total capacity required.</li> </ul>
zone	String	The zone name.
subtenant_id	String	The subtenant identifier.
storage_object_type	String	Storage object type: File server or block store.
storage_object_id	String	The identifier of the storage object. For example: 5d2fb0fb4f47df00015274e3
service_type	String	The service type: File services, block storage, or object storage.

### Retrieve service requests

Use the method listed in the following table to retrieve service request categories for the specified tenant.

HTTP Method	Path	Description	Parameters
GET`	/v2.1/tenants/{tenant_id}/servicerequests	Retrieve service requests.	tenant_id: (Optional) Return the service requests for the specified tenant.  offset and limit– see <a href="#">Common Pagination</a>

Required request body attributes: none

### Request body example:

```
none
```

### Response body example:

```

{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "total_records": 34,
    "sort_by": "created",
    "order_by": "desc",
    "offset": 6,
    "limit": 1,
    "records": [
      {
        "id": "SRQ0035952014",
        "subject": "DR Failover - fileserver",
        "description": "catgory:Disaster Recovery Failover \n subtenant:
DefaultSubtenants2 \n region: au-east2 \n zone: au-east2-a \n
fileserver: Demotsysserv1 \n tenant:MyOrg \n comments:comments",
        "priority": "Urgent",
        "status": "New",
        "createdDate": "2020-05-22T04:23:12+0000",
        "updatedDate": "2020-05-22T04:23:12+0000"
      }
    ]
  }
}

```

### Retrieve a service request by ID

Use the method listed in the following table to retrieve a service request by service request ID.

HTTP Method	Path	Description	Parameters
GET	/v2.1/tenants/{tenant_id}/servicerequests/{id}	Retrieve a service request by ID.	<ul style="list-style-type: none"> <li>tenant_id: Tenant ID</li> <li>id: Service request ID For example: SRQ0035952014</li> </ul>

Required request body attributes: none

### Request body example:

none

### Response body example:

```
{
  "status": {
    "user_message": "Okay. Returned 1 record.",
    "verbose_message": "",
    "code": 200
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "SRQ0035952014",
        "subject": "DR Failover - fileserver",
        "description": "catgory:Disaster Recovery Failover \n subtenant:
DefaultSubtenants2 \n region: au-east2 \n zone: au-east2-a \n
fileserver: Demotsysserv1 \n tenant:MyOrg \n comments:comments",
        "priority": "Urgent",
        "status": "New",
        "createdDate": "2020-05-22T04:23:12+0000",
        "updatedAt": "2020-05-22T04:23:12+0000"
      }
    ]
  }
}
```

### Create a service request

Use the method listed in the following table to create a service request.

==

HTTP Method	Path	Description	Parameters
-------------	------	-------------	------------

POST			
------	--	--	--

	/v2.1/tenants/{tenant_id}/servicerequests/categories		
--	------------------------------------------------------	--	--

		Create a service request.	
--	--	---------------------------	--

	tenant_id	The tenant identifier.	
--	-----------	------------------------	--

Required request body attributes: the required attributes are dependent on the category of service request. The following table lists the request body attributes.

Category	Required
----------	----------

Subscription	
--------------	--

|subscription and commitment

|Disaster recovery

|storage\_object\_type, subtenant\_id, and storage\_object\_id

|Technical

|subtenant\_id and service\_type

If service\_type is file services or block storage, zone is required.

|Other

|Zone

**Request body example:**

```
....
{
  "subject": "string",
  "comment": "string",
  "category": "subscription",
  "priority": "Normal",
  "subscription": "A-S00003969",
  "commitment": {
    "service_level": "standard",
    "new_commitment": 10
  },
  "zone": "au-east1-a",
  "subtenant_id": "5d2fb0fb4f47df00015274e3",
  "storage_object_type": "fileserver",
  "storage_object_id": "5d2fb0fb4f47df00015274e3",
  "service_type": "File Services"
}
....
```

**Response body example:**

```
....
{
  "status": {
    "user_message": "string",
    "verbose_message": "string",
    "code": "string"
  },
  "result": {
    "returned_records": 1,
    "records": [
      {
        "id": "string",
        "subject": "string",
        "description": "string",
        "status": "New",
        "priority": "Normal",
        "createdDate": "2020-05-12T03:18:25+0000",
        "updatedAt": "2020-05-12T03:18:25+0000"
      }
    ]
  }
}
....
```

== Retrieve service request categories

The following table lists the retrieve service request categories for a specified tenant.

HTTP Method	Path	Description	Parameters
GET	/v2.1/tenants/{tenant_id}/servicerequests/categories	Retrieve service requests categories.	

|tenant\_id: (Optional) Return the service requests for a specified tenant.

Required request body attributes: none

**Request body example:**

....  
none  
....

**Response body example:**

```
....  
{  
  "status": {  
    "user_message": "Okay. Returned 5 records.",  
    "verbose_message": "",  
    "code": 200  
  },  
  "result": {  
    "returned_records": 5,  
    "records": [  
      {  
        "key": "dr",  
        "value": "Disaster Recovery Failover"  
      },  
      {  
        "key": "technical",  
        "value": "Technical Issue"  
      },  
      {  
        "key": "other",  
        "value": "Other"  
      },  
      {  
        "key": "subscription",  
        "value": "Subscription Management"  
      },  
      {  
        "key": "backup",  
        "value": "Backup Restore"  
      }  
    ]  
  }  
}
```

```
....  
  
= Jobs  
:hardbreaks:  
:icons: font  
:linkattrs:  
:relative_path: ./  
:imagesdir: /tmp/d20220609-5437-17mplbb/source/././media/
```

|HTTP Method |Path |Description |Parameters

|GET

|/v2.1/jobs

|Retrieve jobs.

|tenant\_id: (Optional) Return the block stores belonging to the specified tenant.

offset and limit– see [Common Pagination](#)



Required request body attributes: none

### Request body example:

```
....  
none  
....
```

### Response body example:

```
....  
{  
  "status": {  
    "user_message": "Okay. Returned 1 record.",  
    "verbose_message": "",  
    "code": 200  
  },  
  "result": {  
    "returned_records": 1,  
    "total_records": 2625,  
    "sort_by": "created",  
    "order_by": "desc",  
    "offset": 0,  
    "limit": 1,  
    "records": [  
      {  
        "id": "5ed72c8c6342e90001439d54",  
        "action": "create",  
        "job_summary": "Create request is successfully submitted",  
        "created": "2020-06-03T04:52:28.478Z",  
        "updated": "2020-06-03T04:52:32.636Z",  
        "object_id": "5ed72c8c6342e90001439d55",  
        "object_type": "sg_buckets",  
        "object_name": "test1234",  
        "status": "successful",  
        "status_detail": "Creation of s3 bucket 'test1234' completed successfully.",  
        "last_error": "",  
        "user_id": "5e85225af038943eb4b74684",  
        "job_tasks": [  
          {  
            "id": "5ed72c8c6342e90001439d57",  
            "job_id": "5ed72c8c6342e90001439d54",  
            "action": "create",  
            "driver": "storagegrid_ansible",  
            "object_id": "5ed72c8c6342e90001439d55",  
            "object_type": "sg_buckets",  
            "resource_type": "sg_bucket",  
            "status": "successful",  
            "status_detail": "Worker completed task successfully.",  
            "last_error": "",  
            "user_id": "5e85225af038943eb4b74684",  
            "request_payload": {  
              "grid_account_id": "05336917559886003354",  
              "grid_admin_base_url": "https://sggmi-dev.dev.ausngs.netapp.au/api/v3",
```

|HTTP Method |Path |Description |Parameters

|GET

|/v2.1/jobs/{id}

|Retrieve a job by ID.

|id (string): The unique identifier for the job.

Required request body attributes: job identifier

### Request body example:

```
....  
none  
....
```

### Response body example:

```
....  
{  
  "status": {  
    "user_message": "Okay. Returned 1 record.",  
    "verbose_message": "",  
    "code": 200  
  },  
  "result": {  
    "total_records": 1,  
    "records": [  
      {  
        "id": "5e66f18e09a74c0001b89640",  
        "action": "create",  
        "job_summary": "Create S3 bucket for Sub Tenant",  
        "created": "2020-03-10T01:46:54.097Z",  
        "updated": "2020-03-10T01:46:57.664Z",  
        "object_id": "5e66f18e09a74c0001b89641",  
        "object_type": "sg_buckets",  
        "status": "successful",  
        "status_detail": "Creation of s3 bucket 'mys3bucket' completed successfully.",  
        "last_error": "",  
        "user_id": "5bbee380a2df7a04d43acaee",  
        "job_tasks": [  
          {  
            "id": "5e66f18e09a74c0001b89642",  
            "job_id": "5e66f18e09a74c0001b89640",  
            "action": "create",  
            "driver": "storagegrid_ansible",  
            "object_id": "5e66f18e09a74c0001b89641",  
            "object_type": "sg_buckets",  
            "resource_type": "sg_bucket",  
            "status": "successful",  
            "status_detail": "Worker completed task successfully.",  
            "last_error": "",  
            "user_id": "5bbee380a2df7a04d43acaee",  
            "request_payload": {  
              "grid_account_id": "47490102387197219062",  
              "grid_admin_base_url": "https://sggmi-dev.dev.ausngs.netapp.au/api/v3",  
              "org_password": "netapp01",  
              "org_username": "root",  
              "s3_bucket_name": "mys3bucket"  
            }  
          }  
        ]  
      }  
    ]  
  }  
}
```