



Lyve Cloud S3 Storage User manual



Hier klicken, um eine aktuelle Online-Version
dieses Dokuments aufzurufen. Auch finden Sie hier die aktuellsten Inhalte sowie erweiterbare
Illustrationen, eine übersichtlichere Navigation sowie Suchfunktionen.

Contents

1	Product Features	8
	Storage management	8
	Storage Analytics	9
	Access management and security	9
	• Access management	9
	• Security	9
	• Availability and durability	9
2	Quick Start Guide	11
	Understanding Account ID	11
	Signing in to Lyve Cloud	11
	• Login Sequence	12
	• Finding your Account ID in the Lyve Cloud console	13
	Using the Lyve Cloud console	14
	Understanding the home page dashboard	15
	Provisioning storage buckets	17
	• Create buckets	17
	• Create bucket permissions	18
	• Create service accounts	18
	Understanding Global Accounts	18
3	S3 API Endpoints	20
4	Using Account API	22
	What can I do with Account API?	22
	Generating Account API credentials	22
	Regenerating Account API credentials	23
	Deleting Account API credentials	24
5	Administrator's Guide - Introduction	25
	Console high-level workflow	25
	• Administrator workflow (admin role)	25
	• Storage management workflow (storage admin role)	25
	• Auditor workflow	26
	Console session management	26
	• Supported browsers	26
	Managing support tickets	27
	• Role-based access for the support page	27
	• Video: How to contact Lyve Cloud Support	28
	• Creating a support ticket	28
	• Editing a ticket	29
	• Viewing ticket details	29
	• Service availability	30
6	Administrator's Guide - Bucket Management	31

Role-based access to buckets	31
Creating buckets	31
Editing bucket properties	33
Listing buckets	34
• Video: Lyve Cloud - How to Create a Bucket	35
Using object immutability	35
• Video: Lyve Cloud - How to Prevent Objects From Being Deleted	35
• How does versioning work in object immutability?	35
• Setting duration	36
Managing bucket access permissions	37
• Role-based access to permission management	37
• Creating bucket access permissions	37
• Creating policy permissions	39
• How to get an IAM policy file from AWS	39
• Using policy permission files	40
• Example policy permission file	41
• Editing bucket permissions	43
• Deleting bucket permissions	44
• Viewing permissions	45

7 Administrator's Guide - Account Management 46

Role-based access to manage service accounts	46
Creating service accounts	46
Viewing service accounts	48
Editing service accounts	49
Changing the status of a service account	50
Deleting a service account	51
Cloning a service account	51
Service account settings	52
• Setting expiration duration	52

8 Administrator's Guide - Audit Log Management 54

Role-based access to permission	55
Video: Lyve Cloud - How to manage audit log settings in the Lyve Cloud console	55
S3 API audit logs	55
• Example S3 API audit log	55
• Enabling S3 API audit logs	57
• Disabling S3 API audit logs	59
Console audit log	59
• Example of Account API audit log structure	59
• Example of an account audit log	60
• Status code	63
• Example of an IAM audit log	64
• Enable account audit logs	72
• Disable account audit logs	72
• Edit audit log target bucket	72

9 Administrator's Guide - Identity and Access Management (IAM) 74

Role-based access to IAM	74
Managing users and roles	74

• About user roles	74
• About user and authentication types	75
• User types	75
• Authentication types	75
• Add a user	75
• Registering an IAM user in Lyve Cloud	76
• Registration workflow for password authentication type	77
• Registering after an email invitation link expires	77
• Viewing and editing a user	77
• Disabling or enabling a user	79
• Deleting a user	79
• Resetting password	79
• Video: Lyve Cloud - How to manage users and assign roles	79
Password policy	79
• Password change history	80
• Password expiration	80
• Restricting password	80
Using multi-factor authentication (MFA)	81
• About MFA	81
• How does MFA work?	82
• If MFA is enrolled	82
• If MFA is not enrolled	82
• Enrolling in MFA	83
• Resetting MFA for an individual IAM user	86
Configuring Federated Login	87
• Security Assertion Markup Language (SAML) protocol	87
• Configure Lyve Cloud as a SAML service provider	87
• Obtain metadata and certificate from your IdP administrator	87
• Configure Lyve Cloud as a service provider	87
• Add service provider metadata to the identity provider	89
• Configure the identity provider to send email attribute	89
• Auth0 and Azure	90
• G Suite	90
• Okta	90
• Troubleshooting SSO	91
• Updating the metadata file	91
• Deleting existing IdP configuration	92
• Generating XML metadata files for IdP	92
Generating XML metadata files for Okta	92
• Pre-requisites	92
• Generate an XML file for Okta	92
• How can an Okta user log in to Lyve Cloud?	95
Generating XML metadata files for Google GSuite	96
• How can a GSuite user log in to Lyve Cloud?	100
Generating XML metadata files for Microsoft Azure AD	101
• How can a Microsoft Azure user log in to Lyve Cloud?	103
Managing notification recipients	104
• Adding a recipient	105
• Editing a recipient	105
• Removing a recipient	105
• Video: Lyve Cloud - How to add a notification request	105

Master accounts	106
Using support	107
Managing Sub-accounts	108
• Role-based access to manage Sub-account	108
• Creating a Sub-account	108
• Listing Sub-accounts	112
• Exporting the Sub-account list	113
• Disabling/enabling a Sub-account	113
• Deleting a Sub-account	114
• Video: Lyve Cloud - How to use the Sub-Account feature of Lyve Cloud.....	115
Understanding the Sub-account usage dashboard	115
Understanding Sub-accounts billing	116
• Master account billing	116
• Invoicing master account	117
• Billing for a disabled Sub-account.....	119

11 .Connecting S3 Clients 120

Using Cyberduck	120
• Prerequisites	120
• Connect Cyberduck to Lyve Cloud	121
• Video: Lyve Cloud - How to connect Cyberduck to Lyve Cloud	123
• Managing data	123
• Uploading data to a bucket	123
• Downloading data to local storage	124
• Deleting data from the bucket	124
• Creating a new folder	124
• Deleting a folder	124
• Disconnecting Cyberduck from Lyve Cloud	125
Using S3 Browser	125
• Prerequisites	125
• Connect S3 Browser to Lyve Cloud	125
• Advanced S3 compatible storage settings.....	127
• Video: Lyve Cloud - How to use S3 Browser with Lyve Cloud	128
• Managing data	128
• Uploading data to a bucket	128
• Downloading data to local storage.....	128
• Deleting data from a bucket	129
• Creating a new folder	130
• Deleting a folder	130
• Disconnecting S3 Browser from Lyve Cloud.....	131
Using Mountain Duck	132
• Prerequisites	132
• Connect Mountain Duck to Lyve Cloud	133
• Video: Lyve Cloud - How to mount Lyve Cloud as a local drive on Windows with Mountain Duck.....	135
• Managing data	135
• Copying data	135
• Deleting bucket data	136
• Creating a folder in a bucket	136
• Disconnecting Mountain Duck from Lyve Cloud	136
Using Rclone	137
• Prerequisites	137

• Connecting to Lyve Cloud from Linux	137
• Installing Rclone	137
• Configuring rClone to connect to Lyve Cloud	137
• Video: Lyve Cloud - How to configure Rclone	142
• Mounting Lyve Cloud as a drive	142
• Managing data	142
• Viewing information about your buckets and directories	142
• Video: Lyve Cloud - How to use Rclone list commands	143
• Uploading data to a bucket	143
• Video: Lyve Cloud - How to use Rclone Copy-to and Copy-sync Commands	143
• Downloading data to local storage	143
• Delete bucket data	144
• Creating a new folder	144
• Deleting a folder	144
• Copying data	145
• Deleting a remote from Lyve Cloud	145
• Migrating Data	146
• Video: Lyve Cloud - Use Rclone Delete and Purge commands	146

12 .Frequently Asked Questions 147

General	147
Buckets	150
Service Account	154
User and Roles	155
Self Service	155
Sub-accounts	157
Billing	158
Security	159
Support	163
Partner	164
HIPAA	165

13 .Release Notes 167

Release 5.7	167
• Features	167
• Registration Process	167
• Trial Period	167
• Cancelling the Lyve Cloud account	168
Release 1.2.7 (Data Mover)	168
• Features	168
Release 5.6	168
• Features	168
Release 5.5	169
• Features	169
Release 5.4	169
• Features	169
• Enhancements	170
Release 5.3	170
• Features	170
• Enhancements	170
• Release features overview video	171

14 .Troubleshooting Guides 172

Login issue 172

• Solution 172

15 .Resources 175

16 .Lyve Cloud Compliance 176

HIPAA 176

• Lyve Cloud has a HIPAA Compliant report 176

ISO 27001 177

Type 2 SOC 2 178

• Lyve Cloud has a Type 2 SOC 2 Attestation report 178

Summary 178

17 .Lyve Cloud Limitations 180

Bucket and object limitations 180

Multipart upload limitation 180

Incompatible S3 API calls 181

Product Features

Lyve Cloud offers several features designed to support a variety of use cases. Customers can easily store, analyze, and manage data on secure, cost-efficient Seagate storage. Lyve Cloud provides an object storage solution that allows customers to move data to and from storage buckets through an HTTPS protocol. Admins can easily manage bucket access with user-specific access control lists. With Lyve Cloud's flexible application programming interface (API), customers can plug in their favorite S3-compatible applications to store data, run big data analytics, audit storage activity, and manage users across the platform.

Storage management

Lyve Cloud's single-tier design breaks away from traditional storage classes to provide uninterrupted data movement. Objects stored in Lyve Cloud can be uploaded, downloaded, updated, and erased anytime. Using S3 Select API calls, customers can easily connect to third-party clients to move and manage data. Applications are authenticated to Lyve Cloud using an access key and secret key provisioned at service account creation. Once authenticated, applications will access buckets and objects using the defined permissions set in the service account (read-only, write-only, or all operations).

All S3 API activity and actions within the Lyve Cloud console are tracked with [Audit logs](#). Audit logs record all S3-supported API calls and activities on the console to access audit functions and track suspicious activity.

Lyve Cloud also offers features to help prevent unintended data modifications and provide versioning. Using [Object Immutability](#) prevents objects from being deleted or overwritten by any user, including the account owner, for a specified retention duration. Object Immutability also supports Amazon S3 Object Lock to reinforce Write-Once-Read-Many (WORM) policies. Customers can toggle on Object Immutability at bucket creation to enable this feature, which also enables Versioning. Versioning allows customers to protect, recover, and restore every iteration of an object stored in a bucket in case of accidental deletions or failures. Versioning remains enabled even if Object Immutability is later disabled.

Lyve Cloud offers Global Account Management to allow customers to create buckets in different regions or create service accounts to access buckets in different regions. For more information, see [S3 API endpoints](#). This provides simplified management of multiple regions on the Lyve Cloud console and the ability to increase redundancy and availability. For more information, see [Understanding Global Accounts](#).

Customers also use Lyve Cloud Sub-Accounts to create, provision and manage additional sub-accounts to maintain a multi-level account structure. Each sub-account can function as its storage account with the ability to manage its users, create buckets and upload data.

Learn more about [Managing Lyve Cloud storage](#).

Storage Analytics

Customers can now analyze, process, move, and transform massive amounts of data on Lyve Cloud using Lyve Cloud Analytics. This platform uses big data frameworks such as Apache Spark, Trino, and ML to satisfy a variety of use cases, including scheduling, monitoring, machine learning, and more. For more information, visit [Getting Started with Analytics](#).

Lyve Cloud's flexible backend was designed to complement various computing applications. As a vendor-agnostic solution, Lyve Cloud can connect to public cloud environments such as AWS, Azure, and Google to utilize their analytics services on Lyve Cloud storage. This functionality allows businesses to consolidate, query, and analyse big data on cost-efficient Seagate storage. Customers can visit Lyve Cloud Marketplace to utilize validated Lyve Cloud partner solutions for computing, such as Zadara zCompute and Equinix Metal.

Access management and security

Access management

Account administrators have several tools to authorize access to Lyve Cloud users. [Identity and access management \(IAM\)](#) allows Lyve Cloud Administrators to manage users and their access to the console. Access is managed with user-defined roles that offer varying levels of accessibility. IAM users can use [Configuring multi-factor authentication \(MFA\)](#) for additional verification during login.

[Configuring Federated Login](#) requires Security Assertion Markup Language (SAML) protocol to provide a single sign-on authentication method through an organization's IDP (identity provider).

Security

Lyve Cloud offers security features to protect data in flight and at rest. To ensure data is protected in flight, Lyve Cloud aligns with Transport Layer Security (TLS) 1.2 protocol and leverages 256-bit Advanced Encryption Standard (AES) Galois/Counter Mode (GCM) encryption, establishing secure communications to the client. By default, all data is encrypted before it is stored.

Learn more by visiting the [Data Security Overview](#) and the [Lyve Cloud Data Security Whitepaper](#).

Availability and durability

Lyve Cloud data centers are located in multiple geographic locations, including Northern California, Virginia, and Singapore, with dedicated operations staff to ensure the Lyve Cloud services are available with a monthly uptime of 99.9%.

Data durability refers to long-term data protection against bit rot or other forms of corruption over long periods. Due to Lyve Cloud's industry-leading architecture, Lyve Cloud can achieve 11 9s of data durability making data loss statistically insignificant.

Quick Start Guide

Seagate on Vimeo: [Lyve Cloud - Getting Started with Lyve Cloud](#)

Understanding Account ID

Account ID is a unique identification that is associated with the Lyve Cloud account. An account ID is unique across all Lyve Cloud accounts globally and can include your company name, which is created during the onboarding process. The account ID helps to identify and distinguish resources in one account from the resources in another account.

While creating, the length of the account ID must be between 3 and 63 characters, where only lowercase characters, numbers, and "-" are allowed.

You cannot change the account ID once it is created.

The account ID is used to create the unique URL for the account's console URL with the following format: `https://<account_ID>.console.lyvecloud.seagate.com`

A single URL is used to access the Lyve Cloud console, which is authenticated by the account ID.

Signing in to Lyve Cloud

To use the Lyve Cloud console, you must sign in using your account credentials. To sign in to Lyve Cloud, you will need a login URL, which contains a unique account ID. The account ID can include your company name chosen during onboarding. You cannot change the account ID once it is created. The account ID is unique across Lyve Cloud accounts.

A single URL is used to access Lyve Cloud console which is authenticated by the account ID, and the URL has the following format: `https://<account_ID>.console.lyvecloud.seagate.com`.

If you know the Lyve Cloud account login URL

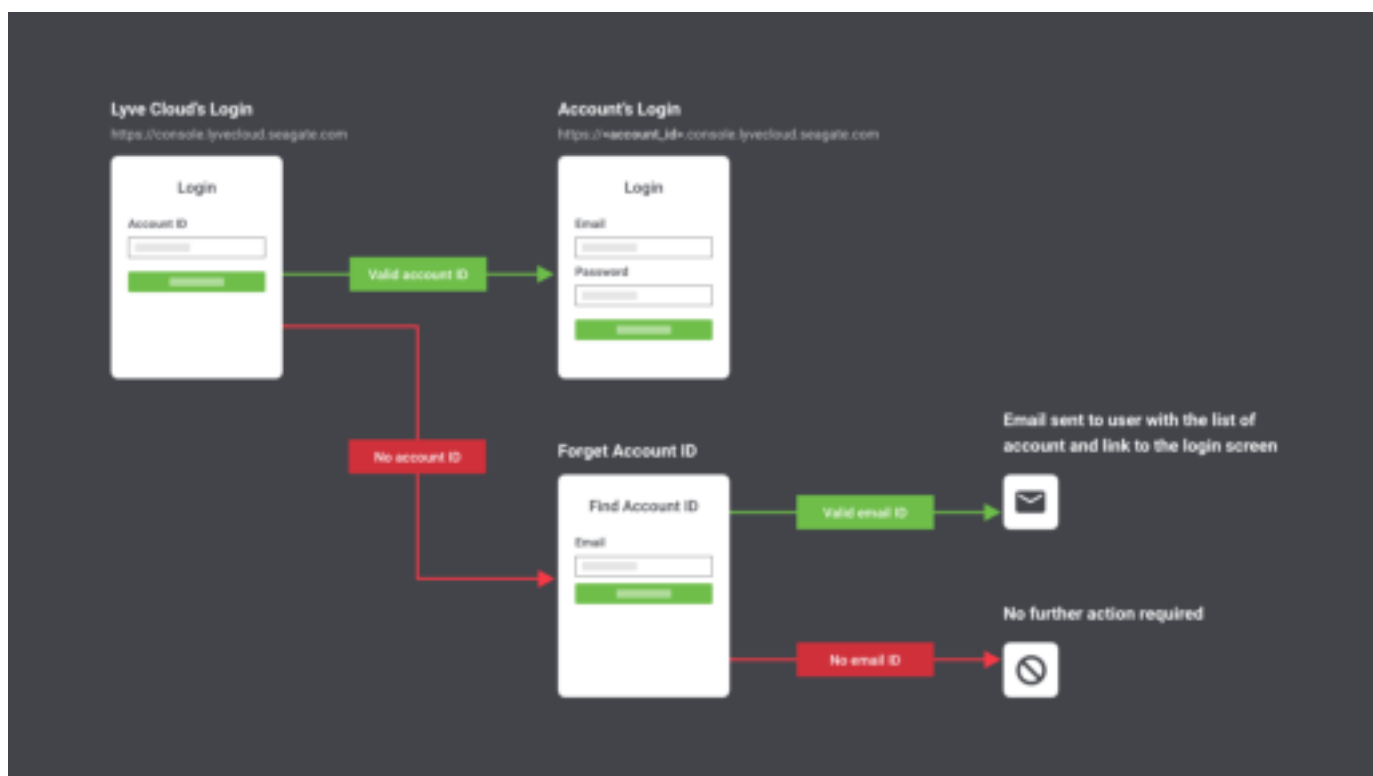
After successful onboarding, you will receive a welcome email. This email contains the Lyve Cloud URL. Using this URL, you can sign in to Lyve Cloud by creating a password.

This URL is in the following format: `https://<account_ID>.console.lyvecloud.seagate.com`. The user can bookmark or save this URL to log into the console.

If you do not know the Lyve Cloud account login URL

If you have not saved the URL: https://<account_ID>.console.lyvecloud.seagate.com you will not be able to login to Lyve Cloud console. But when you try to log in using the UR <https://console.lyvecloud.seagate.com> , you are directed to enter the Lyve Cloud Account ID to access the console.

Login Sequence



You have an Account ID—Enter the Account ID on the login page, followed by the registered email address. An email with all the details of the Lyve Cloud account is sent, which contains the Lyve Cloud URL to log in.

The screenshot shows the Lyve Cloud Login page. At the top is the 'LYVE Cloud' logo. Below it is the 'Login' heading. There is a text input field for 'Account ID' with a help icon. Below that is a checkbox labeled 'I'm not a robot' and a CAPTCHA image. At the bottom is a green 'Next' button and a link that says 'Forgot your account ID? Get help'.

You do not have an Account ID—You will receive your account ID by providing your registered email address. You must select **Get Help** on the Login page. You are directed to enter your email address.



The image shows the LYVE Cloud Login page. At the top is the LYVE Cloud logo. Below it is the title "Login". There is a text input field labeled "Account ID" with a help icon. Below the input field is a checkbox labeled "I'm not a robot" and a reCAPTCHA logo. A green "Next" button is below the checkbox. At the bottom, a red box highlights the link "Forgot your account ID? Get help".




The image shows the LYVE Cloud Find Account ID page. At the top is the LYVE Cloud logo. Below it is the title "Find Account ID". A paragraph of text reads: "If you have a valid account, you'll receive an email listing the details of all your Lyve Cloud Accounts associated with the specified email address." Below this is a text input field labeled "Email". Below the input field is a checkbox labeled "I'm not a robot" and a reCAPTCHA logo. At the bottom are two buttons: a green "Submit" button and a white "Cancel" button.

If the email address is not registered with Lyve Cloud, contact the support team at support.lyvecloud@seagate.com.

Finding your Account ID in the Lyve Cloud console

If you have already signed in to the Lyve Cloud account, you can view the Account ID from the Header pane. Select the username in the top right to view the Account ID.

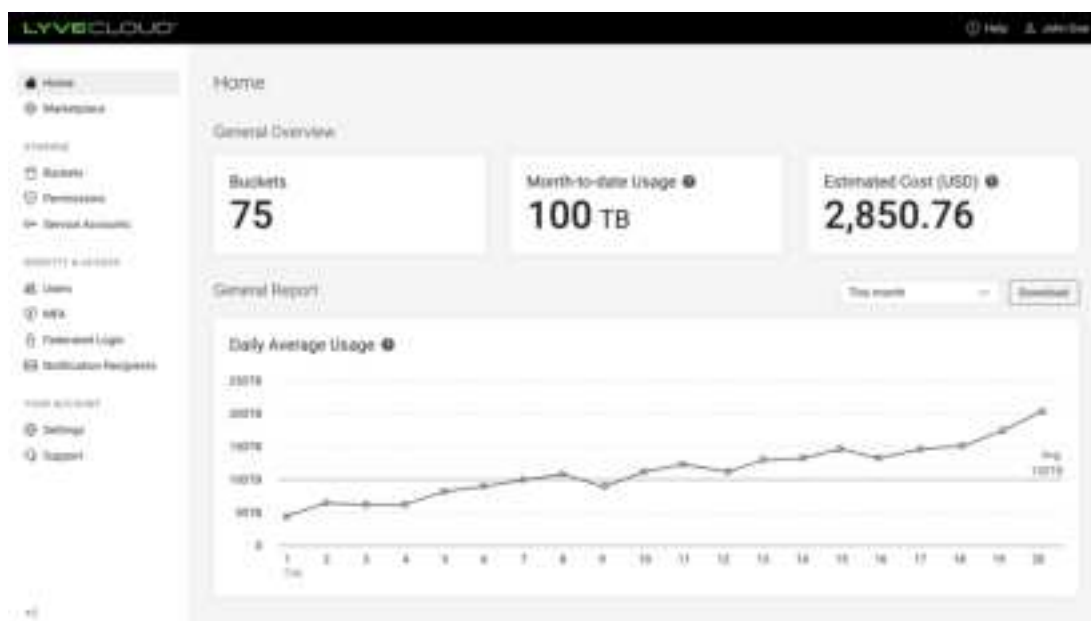
The following image highlights the Account ID in the console.



Using the Lyve Cloud console

The Lyve Cloud console includes three panes: the header pane, left menu, and main view.

The following image displays the three panes of the Lyve Cloud console.



- **Header pane:** Select the Lyve Cloud logo to return to the dashboard. The **Start Here** button, **Help** link, user name, and icon are displayed on the top right corner of the header pane. Open the **Start Here** window to find quickstart resources. View our documentation and training videos under **Help**. To exit Lyve Cloud, select the user name and then select **Logout**.
- **Left menu:** The left menu is organized as follows:
 - **Home page:** It is the landing page after you login to the console. It shows the number of buckets, reports, and usage and more. For more information, see [Understanding the home page dashboard](#) below.

- **Marketplace:** This section displays and provides more information on partner solutions like Backup and Recovery, Surveillance, Compute, etc. that are certified with Lyve Cloud.
- **STORAGE**
 - **Buckets:** Allows you to create and manage buckets.
 - **Permissions:** Allows you to set the permissions for buckets.
 - **Service Accounts:** Generates access credentials that enable S3 applications to perform S3 operations on the bucket.
- **IDENTITY & ACCESS**
 - **Users:** Allows you to create users and set user roles.
 - **MFA:** Allows you to add an additional factor to the login to prevent unauthorized access.
 - **Federated Login:** Allows you to enable federated single sign-on (SSO) from your organization's Identity Provider (IdP).
 - **Notification Recipients:** Allows you to add recipients to receive service and other important Lyve Cloud notifications via email.
- **SETTINGS**
 - **Settings:** Allows you to enable and disable audit logs. These logs are detailed records of activities in the Lyve Cloud console and S3 API operations.
 - **Billing:** Allows you to see each months' costs, and download and view previous monthly invoices.
 - **Support:** Allows you to open new support tickets for any issues related to Lyve Cloud services.



The non-administrator roles can only see a subset of the menu options.

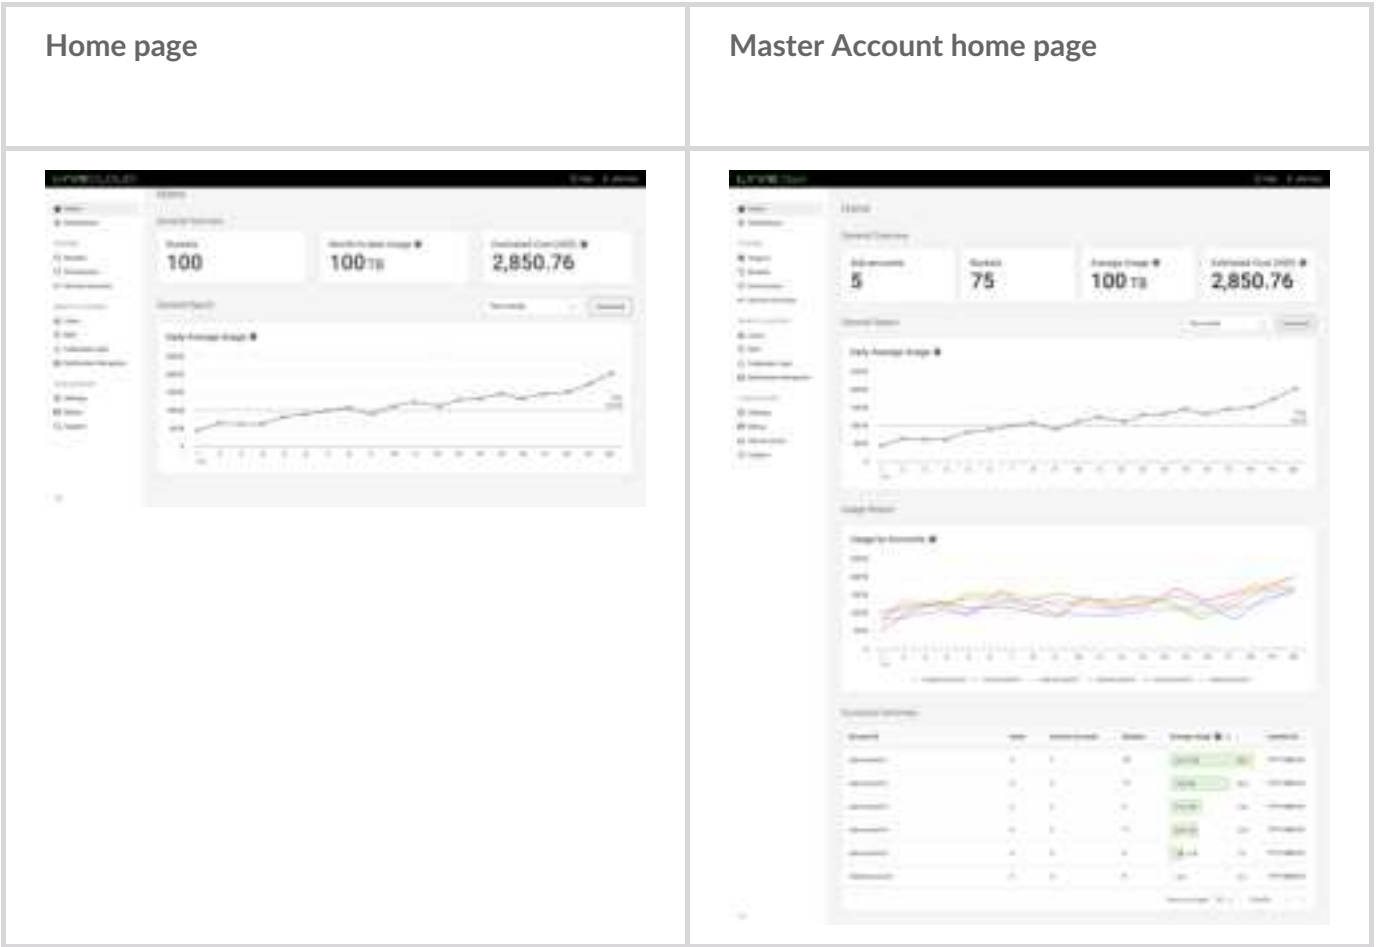
- **Main view:** Displays the information corresponding to the left menu item selected.

Understanding the home page dashboard

After you log in to the Lyve Cloud console, you are taken to the dashboard's home page, and the headings

on the page are displayed without data. However, if you have created buckets and are storing data in the buckets, the dashboard displays important details in the different sections.

The dashboard displays statistics of the storage system, usage, and estimated cost. A graphical view of usage trends, bucket count, and average usage are available.



- **Buckets:** Displays the total number of buckets.
- **Month-to-Date Usage:** Displays the average usage of the account from the beginning of the month until the current date.
- **Estimated Cost:** Displays the estimated monthly storage costs based on the current month’s usage trends. This cost is displayed in US dollars.
- **General Reports:**
 - **Daily Average Usage:** Displays the daily average from a series of four usage snapshots within a 24-hour period of data stored in all the buckets.
 - **Date range selection:** Select a current month, last six months, or custom time range to view usage trends.
 - **This month** is a default selection that displays the daily average usage trend for the current

month to date.

- Selecting the **Last 6 months** shows the usage trend of the last six months. Each data point displays the monthly average for that month.
- Selecting a **Custom time range** allows you to choose the monthly time range, and the data points display the monthly average usage.

Download the usage data in CSV format by selecting **Download**. Use the **Date range selection** to select the length of time of the report. This report shows the Date, Region Name, Bucket Name, Usage(byte), Usage (GB) in the excel sheet.

- **Usage Report:** Displays the usage of all the sub-accounts in the master account. The Sub-accounts Usage graph displays the usage of each sub-account on the same graph. The graph has different colour lines per account. Hovering over a particular day/month (depending on view scale) displays a tooltip with the time information for all selected accounts with the line colour, account name, and usage value per sub-account account.
- **Accounts Summary:** Displays the summary of each sub-account.
 - **Customers:** Lists the account ID of each sub-account in the master account.
 - **Users:** Lists the number of users for each sub-account.
 - **Service Accounts:** Lists the number of service accounts for each sub-account.
 - **Buckets:** Lists the number of buckets created by each sub-account.
 - **Average Usage:** Lists the average amount of data used per day for each sub-account, from the beginning of the month to the current date.
 - **Created On:** Displays the date when the sub-account is added to the Lyve Cloud master account.

Provisioning storage buckets

Seagate on Vimeo: [Lyve Cloud - How to Create a Bucket](#)

Create buckets

Begin by creating a bucket to add data.

1. On the left-hand menu, select **Buckets**, and then select **Create Bucket**.
2. Enter the **Bucket Name** and **Region**. Select **Create**. (Optionally, enable **Object Immutability**.)

After the bucket is created, it is listed on the Buckets page. For more information, see [Administrator's Guide - Bucket Management](#).

Create bucket permissions

Next, create and apply permissions to at least one bucket. Permissions define the type of operations that applications perform on the bucket: **Read**, **Write**, or **All Operations** (read, write, delete, and list).

1. On the left-hand menu, select **Permissions**, and then select **Create Bucket Permission**.
2. On the Create Bucket Permission dialog:
 - Name: Enter the name for permission. Permission names can contain any alphanumeric characters, dashes ("-"), underscores ("_"), or spaces.
 - Select one of the following from **Which buckets does this permission apply to?**
 - **One or more existing buckets**
 - **All buckets in this account with a prefix**
 - **All buckets in this account**
 - Select **Actions** to assign privileges as **All Operations**, **Read only**, **Write only**.
 - Select **Create** to save the permission for the bucket.



The Description of the permission assigned to the buckets is displayed.

3. Alternatively, you may import policy permission files to create new permissions. See [Using policy permission files](#).

For more information on buckets, see [Administrator's Guide - Bucket Management](#).

Create service accounts

Finally, after creating permissions for a bucket, create a service account to allow applications to authenticate and use these permissions. Applications use service account credentials in API calls to access buckets to add and delete data.

1. On the left-hand menu, select **Service Accounts**, and then select **Create Service Account**.
2. Enter the Service Account Name, and then select applicable Permissions from the available list.
3. Select **Create**. A confirmation displays the access key and secret key required to access the bucket.
4. Copy these account credentials or download them in CSV or JSON format before you close the dialog.



The access key and secret key cannot be retrieved later.

For more information on service accounts, see [Administrator's Guide - Account Management](#).

Understanding Global Accounts

Lyve Cloud Global Accounts let customers create buckets in different regions for increased provisioning

and data access. For more information, see [Creating buckets](#) .

Once you create buckets in different global accounts:

- The Lyve Cloud console lists all the buckets created for an account. For more information, see [Listing buckets](#).
- Listing buckets using the S3 API displays the buckets for the region that is specified in the API command.
- You can copy objects between different Lyve Cloud regions using S3 API commands.

To access data from buckets created in different global regions:

- Make direct requests to one of the Lyve Cloud S3 API endpoints. For more information on S3 access points, see [S3 API endpoints](#).



Lyve Cloud does not provide an S3 API global endpoint to access data across different global accounts. You must use the region specific endpoint to provision storage.

S3 API Endpoints

The following table shows Lyve Cloud regions where Lyve Cloud is currently available and the endpoints for these regions.

Region	Endpoint	Comment
US-East-1 (N. Virginia)	<code>https://s3.us-east-1.lyvecloud.seagate.com</code>	Standard Region
US-West-1 (N. California)	<code>https://s3.us-west-1.lyvecloud.seagate.com</code>	Standard Region
AP-Southeast-1 (Singapore)	<code>https://s3.ap-southeast-1.lyvecloud.seagate.com</code>	Standard Region
EU-West-1 (London)	<code>https://s3.eu-west-1.lyvecloud.seagate.com</code>	Standard Region
US-Central-2 (Texas)	<code>https://s3.us-central-2.lyvecloud.seagate.com</code>	Standard Region

Lyve Cloud supports path-style requests and virtual hosted-style requests available with AWS S3.

Use the URL format to access a bucket using a path-style endpoint or virtual hosted-style endpoint.



Lyve Cloud does not provide an S3 API global endpoint to access data across different regions.

Region	Path-style endpoint	Virtual hosted-style endpoint
US-East-1 (Virginia)	<code>https://s3.us-east-1.lyvecloud.seagate.com/[bucket_name]</code>	<code>https://[bucket_name].s3.us-east-1.lyvecloud.seagate.com</code>

Region	Path-style endpoint	Virtual hosted-style endpoint
US-West-1 (California)	<code>https://s3.us-west-1.lyvecloud.seagate.com/[bucket_name]</code>	<code>https://[bucket_name].s3.us-west-1.lyvecloud.seagate.com</code>
AP-Southeast-1 (Singapore)	<code>https://s3.ap-southeast-1.lyvecloud.seagate.com/[bucket_name]</code>	<code>https://[bucket_name].s3.ap-southeast-1.lyvecloud.seagate.com</code>
EU-West-1 (London)	<code>https://s3.eu-west-1.lyvecloud.seagate.com/[bucket_name]</code>	<code>https://[bucket_name].s3.eu-west-1.lyvecloud.seagate.com</code>
US-Central-2 (Texas)	<code>https://s3.us-central-2.lyvecloud.seagate.com/[bucket_name]</code>	<code>https://[bucket_name].s3.us-central-2.lyvecloud.seagate.com</code>

Using Account API

Account API allows you to access Lyve Cloud account information through an API endpoint. The account API can be generated only by the account administrators. You can perform all Lyve Cloud operations using the account API credentials.

What can I do with Account API?

Account API enables customers and sub-account administrators to leverage Lyve Cloud account's functionality programmatically.

You can perform the following actions using account API:

- Permissions: Create permissions, List permissions, Get permissions by ID, Delete permissions by ID, and Update permission
- Service Accounts: Create service account, List service account, Get service account data by ID, Update service account, Enable service account and Disable service account.
- Usage: Get current month storage usage and historical storage usage by month.

For more information, see [Lyve Cloud Account API version 2](#)

The API uses the secure HTTP/1.1 over TLS 1.2 protocol and operates mainly with JSON-formatted messages. All API responses are assigned specific numeric codes that help you quickly identify if a request to an endpoint is successful or unsuccessful. For more information on error codes, see [List of API error codes](#) in the Account API version 2.

Generating Account API credentials

The credentials never expire when generating API credentials unless you configure an expiration duration. You can change the default setting by setting an expiry duration for all newly created API credentials; see [Setting expiration duration](#). This limits the validity of the Account API credentials, which need to be changed again after the expiration. After the expiration date, the secret credentials cannot be used for authentication but will stay associated with the account until you delete or regenerate it.

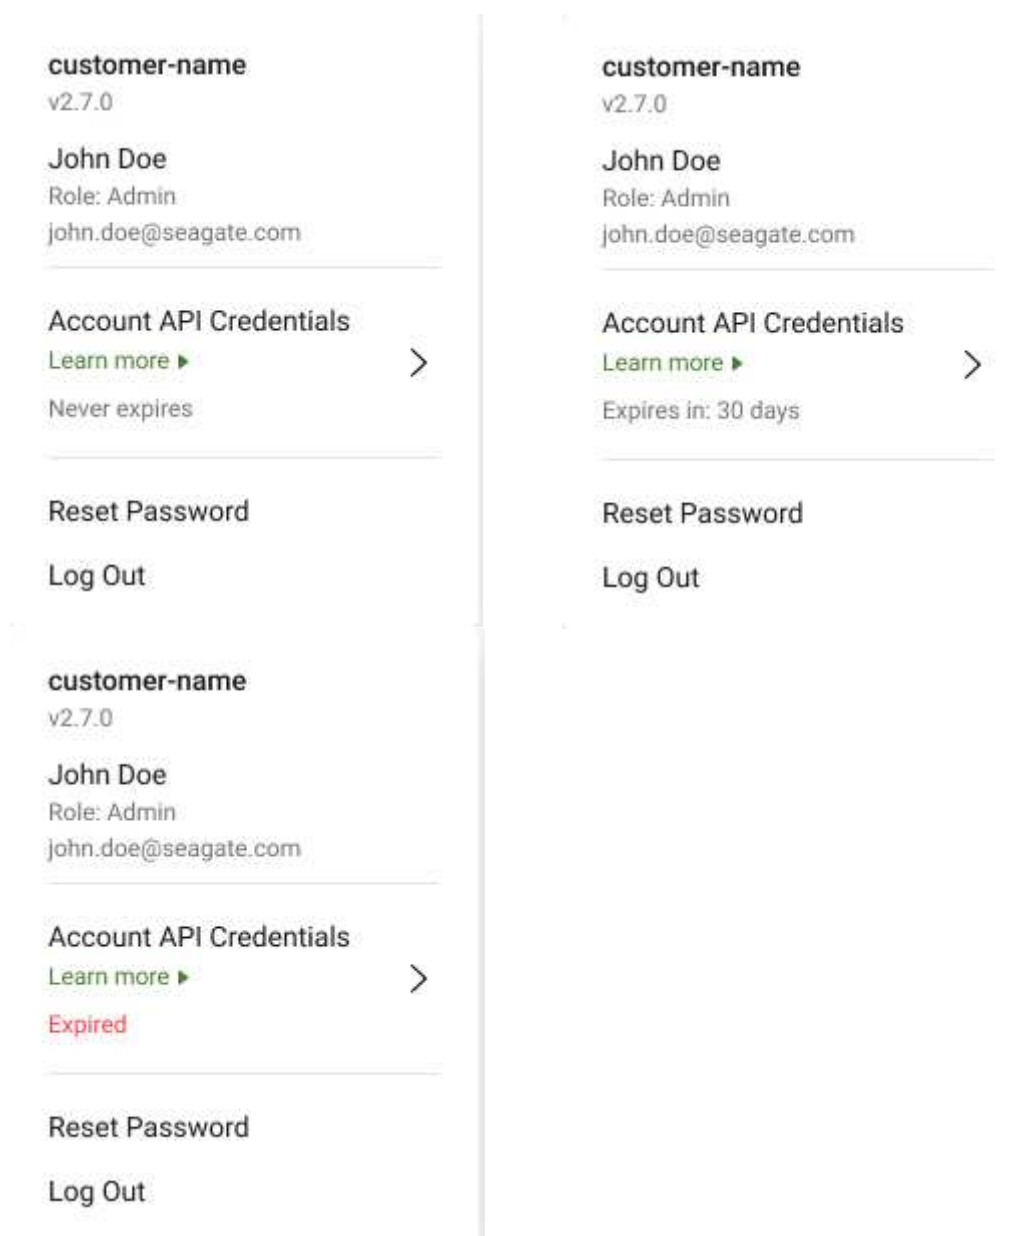
1. On the Header pane, select the username in the top right.
2. Select **Generate Account API Credentials**.
3. Copy or download the **Access Key** and **Secret Key** after you create the Account API credentials.



Download the key in CSV or JSON format, as you cannot retrieve the secret key details later.

The status of credentials will show as one of the following. The status is based on the expiration duration.

- **Expires in:** XX days: You must generate the API credentials to use the API after expiration.
- **Never Expires:** The security credentials are not set, and these credentials will never expire.
- **Expired:** The credentials have expired, and you cannot access the account API.



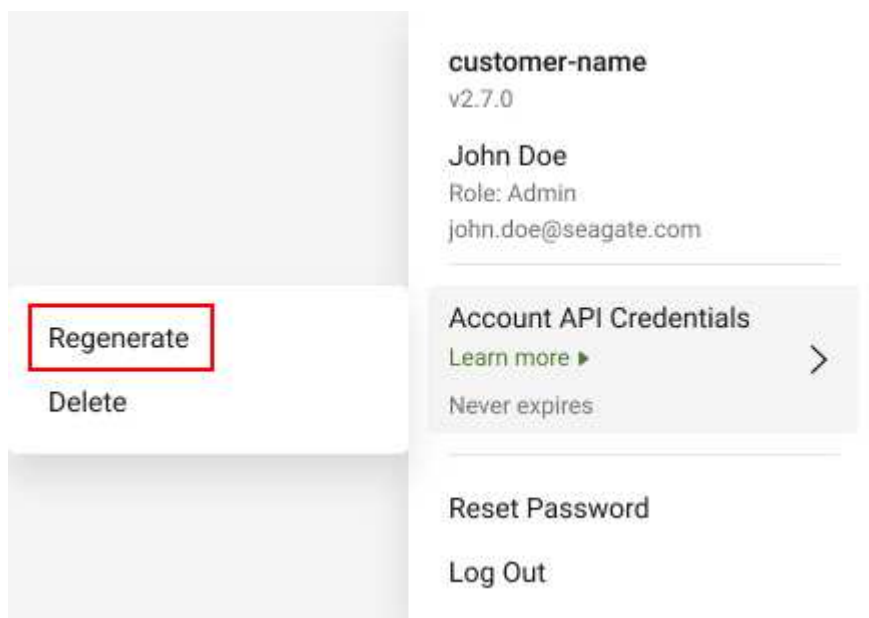
After the credentials are generated, use these credentials to generate a time-bound token. This token is used to authenticate the Lyve Cloud Account API and is passed as a Bearer header value.

You must attempt to generate credentials the second time if Account API generation fails.

Regenerating Account API credentials

You can re-generate the credentials regardless of their expiration status. If you already have active credentials and still regenerate new credentials, the old credentials become inactive.

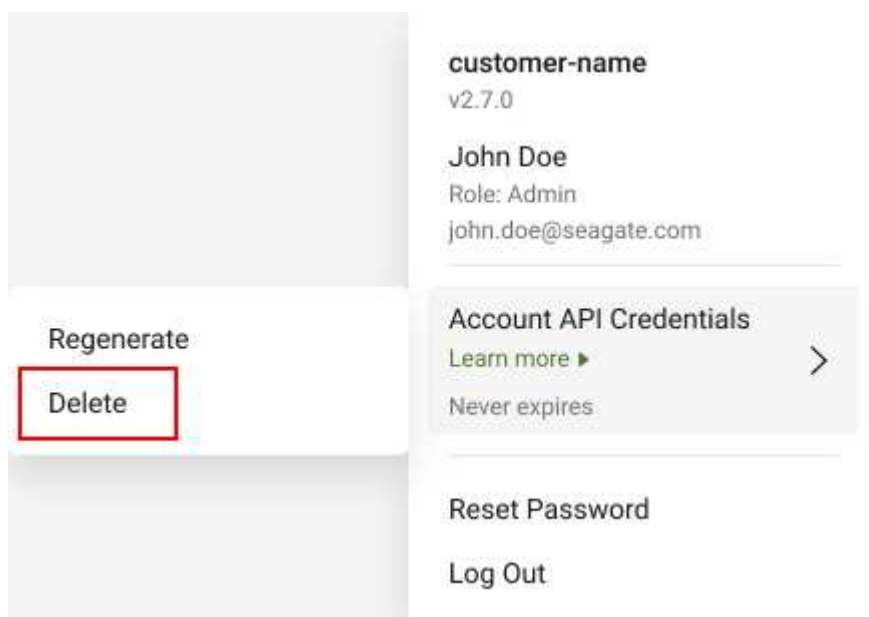
1. On the Header pane, select the username in the top right.
2. Select **Generate Account API Credentials**., and then select **Regenerate**.



Deleting Account API credentials

Once you delete the credentials, you can again generate new credentials. However, any workload that uses these credentials will immediately lose access to the resources.

1. On the Header pane, select the username in the top right.
2. Select **Generate Account API Credentials**., and then select **Delete**.



Administrator's Guide - Introduction

This guide provides instructions on creating buckets and managing bucket permissions and service accounts to authenticate and access data stored in the buckets. It describes identity and access management (IAM) to manage access to your Lyve Cloud resources. The Lyve Cloud console dashboard displays the storage system's overall statistics. See graphical views of usage trends, numerical values of buckets, and average usage.

Console high-level workflow

This section explains the console workflow as determined by user roles. For more information on your assigned role, see [Administrator's Guide - Identity and Access Management \(IAM\)](#).

There are three roles available in Lyve Cloud:

- Administrator
- Storage Administrator
- Auditor

Administrator workflow (admin role)

Administrators can perform all actions available in the Lyve Cloud console.

1. Once you sign in to the Lyve Cloud console, a dashboard is displayed. The dashboard shows details of buckets and usage-related information. For more information, see [Understanding the home page dashboard](#).
2. To manage Lyve Cloud storage:
 - A. Storage is managed and provisioned in buckets. For more information, see [Creating buckets](#).
 - B. Once you create a bucket, you must assign it Permissions, and define what operations are allowed for buckets. See [Creating bucket access permissions](#).
 - C. After you assign permissions to a bucket, create a service account. Service accounts are used by applications to authenticate API calls accessing the bucket. For more information, see [Creating service accounts](#).
3. Use Identity and Access Management (IAM) features to secure access to your Lyve Cloud account. For more information, see [Administrator's Guide - Identity and Access Management \(IAM\)](#).

Storage management workflow (storage admin role)

The storage admin user can perform all the storage operations as an [Administrator](#) user in Lyve Cloud, including managing buckets, managing permissions, and creating service accounts. The storage admin user is restricted from altering settings for Identity and Access Management (IAM) and Lyve Cloud

account billing.

Auditor workflow

Users with the Auditor role have **read only** access throughout the Lyve Cloud console, and are not permitted to perform any storage operations or alter settings.

Console session management

The user login session management increases the strength and security of the Lyve Cloud session. To provide more secure access, non-persistent sessions invalidate a Lyve Cloud console session cookie when the browser is closed.

By default, a user session timeout is 24 hours. Users are not required to log in with their credentials for up to 24 hours only if the Lyve Cloud session is active. The session is active after successful authentication by the user. The Lyve Cloud console automatically signs out the user after 24 hours.

When users close and re-open the browser, they get a prompt for re-authentication.

In summary, the Lyve Cloud console requires re-authentication in the following cases:

- When you sign out of the Lyve Cloud session
- The browser is closed without any active session, or the active Lyve Cloud tab is closed.
- The authenticated session is more than 24 hours.

The Lyve Cloud console session remains active in the following cases:

- At least one Lyve Cloud active session is open, and the authentication session is less than 24 hours.

Supported browsers

The Lyve Cloud Console supports the following browsers:

Browser	Version
Google Chrome	Last three versions
Mozilla Firefox	Last three versions
Microsoft Edge	Last three versions
Apple Safari	Last three versions

Managing support tickets

If you experience a problem with Lyve Cloud, use the Support page to create a support ticket. Please provide detailed information in the **Subject** and **Description** fields, and attach any relevant references for the support team. Detailed information helps us provide a more efficient and effective resolution, as the ticket response time is based on its severity level which is determined from the details provided.

Each support ticket is assigned a unique number. Use this ticket number to track the progress of the reported issue, and update the support ticket by adding a comment. Comments and resolutions are recorded in each ticket.

You can also send an email to support.lyvecloud@seagate.com to report an issue. A support ticket is opened based on the issue reported in the email.

The support team reviews ticket details and updates the ticket status.

- **New:** This status is assigned immediately when a ticket is created, and work is not yet started.
- **In Progress:** This status indicates that the ticket is under review, and a support engineer is investigating the issue.

After the ticket is updated, you will receive an email notification containing the ticket number, subject, and changes made. You will receive an email notification when a new ticket is opened, a ticket is updated, or an issue is resolved and the ticket is closed.

The Support page lists the number of new and in progress tickets.

i **Note**—Customers of a partner must report Lyve Cloud related issues to its partner. If you purchase Lyve Cloud through a reseller or partner, you will not have direct access to support. Please contact your reseller with all support queries.

Role-based access for the support page

The following table describes access to the Support page features based on the admin role.

Actions	Admin	Storage Admin	Auditor (Read only)
Create ticket	✓	✓	×
Edit ticket	✓	✓	×
View ticket	✓	✓	✓
Add Comments	✓	✓	×

Actions	Admin	Storage Admin	Auditor (Read only)
View Comments	✓	✓	✓

Video: How to contact Lyve Cloud Support

Seagate on Vimeo: [Lyve Cloud - How to Contact Lyve Cloud Support](#)

Creating a support ticket

To create a ticket:

1. On the left-hand menu, select **Support**.
2. On the Support page, select **Create New Ticket**.
3. In the Create New Ticket dialog, enter the following:
 - **Subject:** Enter a subject for the support ticket. This is a mandatory field.
 - **Description:** Enter the ticket details. This is an optional field that allows you to describe the problem summary.
 - **Attachments:** Add documents that provide more details about the issue. The file size must not exceed 4 MB. Select **Upload** and choose the file from the desired location to upload an attachment, then select **Open**.

After the file is uploaded, it is listed under Attachments. To remove the attachment, select the x to the right of the file name.

4. Select **Create**.

The new ticket displays in the ticket listing table.



Note—Once a ticket is saved, you cannot delete the attachments.

Editing a ticket

You can edit new and in progress tickets. Editing a ticket allows you to edit or add to the problem summary, description, customer name, and attachments.

To edit a ticket:

1. On the left-hand menu, select **Support**.
2. In the ticket listing table, select a ticket number to edit that ticket.
3. On the Details pane, select **Edit**.
4. Edit any of the following fields:
 - **Subject**
 - **Description**
 - **Attachments**: You can add new attachments, but you cannot delete attachments that were previously added.
 - **Add New Comment**
 - A. To add comments, select **Add New Comment**.
 - B. Enter a comment and select **Add**.

5. Select **Save**.

Viewing ticket details

To view a ticket:

1. On the left-hand menu, select **Support**.
2. On the Support page, select the ticket number to view its details.

Service availability

The Lyve Cloud availability in the following image shows the calculated service availability for the month.



Lyve Cloud service availability is calculated by subtracting the error rate from 100% within a five-minute interval. If a customer does not make any requests in a 5-minute interval, that interval is assumed to have an error rate of 0%.

The error rate is the total number of errors returned, divided by the total number of requests during that 5-minute interval.

Error rate = number of errors ÷ number of request

Availability = 100% - error rate

Administrator's Guide - Bucket Management

Lyve Cloud allows you to store objects (like files) in buckets (like folders). Before you add or store any object, you must create a bucket. When you create a bucket, you must specify the region where you want to create the bucket.

Role-based access to buckets

Bucket access levels are defined by the user roles. The following table describes console access to bucket features based on the user's role:

Actions	Admin	Storage Admin	Auditor (Read only)
Create bucket	✓	✓	×
Edit bucket	✓	✓	×
Delete	✓	✓	×
List and View	✓	✓	✓

Creating buckets

To create a bucket:

1. On the left-hand menu, select **Buckets**.
2. On the Buckets page, select **Create Bucket**.

3. Enter the bucket name:

Remember the following while creating a bucket name:

- A. The bucket name must be unique across all of Lyve Cloud.
- B. A bucket name containing a dot (.) is not allowed.
- C. After you create a bucket, you cannot change the bucket name.

4. Select the region (metro) from the drop-down, where you want the bucket to reside. For more information, see [Understanding Global Accounts](#).

- US - Virginia (us-east-1)
- US - California (us-west-1)
- AP - Singapore - (AP-Southeast-1)

i **Note**—You must create your first bucket in a region, using the console.

5. (Optional) Enable **Object Immutability** . For more information, see [Using object immutability](#).

If Object Immutability is not enabled when a bucket is created, you cannot turn it on later. However, if you switch it on while creating a bucket, you can later switch it off and on again as needed.

If you enable Object Immutability, you can also set a duration to retain the objects. For more information, see [Setting duration](#).

6. After you create a bucket, it is listed on the Buckets page.



Note—Sometimes there may be a delay in creating a bucket.

Editing bucket properties

The Buckets page displays the bucket list. It also displays the labels for each bucket, such as Immutable, Versioned, and Logged. For more information on the labels, see [Using object immutability](#) and [Administrator's Guide - Audit Log Management](#)

To edit a bucket:

1. On the left-hand menu, select **Buckets**.
2. On the Buckets page, choose and select the name to edit.
3. Perform any of the following actions in the bucket properties:
 - **S3 endpoint URL** allows copying the S3 endpoint URL to the clipboard. This URL is used to access the bucket. For more information on the S3 endpoint URL see [S3 API endpoints](#).
 - **Object Immutability** : You may choose to switch off Object Immutability if it is enabled. For more information, see [Using object immutability](#).
 - **Set Duration**: You can set duration only when Object Immutability is switched on. Select the pencil icon to edit the retention duration. For more information, see [Setting duration](#).
 - **S3 API Audit Logs**: Select the toggle switch to enable or disable the audit logs for this bucket. For more information on audit logs, see [Administrator's Guide - Audit Log Management](#) After you enable the audit logs for the selected bucket, the bucket is labeled as **Logged**, and once you disable the audit logs, the label is removed.
 - **Delete Bucket**: Select **Delete** to delete a bucket.

Before deleting a bucket, please make sure to:

- Delete **all data** from the bucket.
- Delete **all permissions** referencing this bucket. Deleting a bucket associated with bucket permissions is allowed only if you have applied permission to all buckets or all buckets with a prefix in the account.
- Verify that the bucket is **not** set as the **target bucket** for Audit Logs.



Listing buckets

To view the bucket list:

1. On the left-hand menu, select Buckets.



Note—This view displays the labels for each bucket, such as Immutable, and Logged. For more information on the labels, see [Using object immutability](#) and [Administrator's Guide - Audit Log Management](#)

By default, the Buckets page displays 10 buckets at a time. To increase or decrease the number of buckets per page, select the Rows per page arrow and select 10, 25, 50, or All.

2. Select the left or right arrow to move between the pages.

Buckets					+ Create Bucket
Create and manage storage buckets. Learn more					
Name	Region	Usage	Created On ↓		
testbucket22	US - California (US-West-1)	0B	2021-07-18	Immutable	Versioned
test-21	US - California (US-West-1)	0B	2021-07-18	Immutable	Versioned
test-test-4	US - California (US-West-1)	0B	2021-07-15	Logged	
test-test-2	US - California (US-West-1)	2.89MB	2021-07-14	Immutable	Versioned
test-test	US - California (US-West-1)	0B	2021-07-14	Immutable	Versioned
test-test	US - California (US-West-1)	241B	2021-07-11	Logged	

The following table displays the description to the column names of the bucket list.

Column Name	Description
Name	Displays the name of the bucket.

Column Name	Description
Region	Displays the region where the bucket is residing. You can select the region while creating a bucket. For more information, see Creating buckets .
Usage	Displays the total amount of data stored in the bucket in KiB, MiB, or GiB.
Created On	Displays when the bucket was created in YYYY-MM-DD format.
Immutable, Versioned, Logged	<p>Displays the bucket labels.</p> <ul style="list-style-type: none"> • Immutable: The label indicates that the bucket is in compliance mode. To disable the compliance mode, see Editing bucket properties. • Versioned: The label indicates that the bucket is versioned. The bucket version is not suspended even after you disable the Object Immutability • Logged: The label indicates that audit logs are enabled for the bucket. To disable the audit logs for buckets, see Editing bucket properties. <p>For more information on these labels, see Using object immutability and Administrator's Guide - Audit Log Management</p>

Video: Lyve Cloud - How to Create a Bucket

Seagate on Vimeo: [Lyve Cloud - How to Create a Bucket](#)

Using object immutability

Object immutability prevents objects from being deleted or overwritten by any user or application for a specified retention duration. This is especially useful when you want to meet regulatory data requirements or other scenarios where it is imperative that data cannot be changed or deleted. Object immutability must be used when you are certain that you do not want anyone, including the Administrator, to delete the objects during their retention duration. When you switch on object immutability, you must also [set the duration](#) and specify the default retention period.

Video: Lyve Cloud - How to Prevent Objects From Being Deleted

Seagate on Vimeo: [Lyve Cloud - How to Create a Bucket](#)

How does versioning work in object immutability?

Versioning allows saving multiple variants of an object in the same bucket. It allows you to preserve, retrieve, and restore every version of an object stored in the bucket. Versioning enables the recovery of objects from any unintended or accidental user actions and application failures.

After switching on object immutability for a bucket, versioning is automatically enabled, Lyve Cloud

automatically creates and stores an object version each time when:

- A new object is uploaded
- An existing object is overwritten
- An object is deleted



Note—Versioning may increase your storage capacity utilization.

For example, if you accidentally delete an object, instead of removing it permanently from Lyve Cloud, this deleted object becomes the current object version. You can then restore the previously available version.

When you create a bucket and switch on object immutability, you can switch off object immutability afterwards. However, versioning cannot be suspended for that bucket.

For example, if you accidentally delete an object, instead of removing it permanently from Lyve Cloud, this deleted object becomes the current object version. You can then restore the previously available version.

When you create a bucket and switch on object immutability, you can switch off object immutability afterwards. However, versioning cannot be suspended for that bucket.



Note—Switching on object immutability, the bucket is labelled as Immutable and Versioned. Switching off object immutability only removes the Immutable label.

Setting duration

The duration for immutability can be specified in days or years at the object level. When you set the duration, objects remain locked and cannot be overwritten or deleted. By default, the duration is set to 30 days. Setting the duration applies to individual object versions, and different versions of a single object can have different durations set.

For example, if you set duration to 10 days and then create an object A, object A will have its retention duration set to 10 days. If you later change the duration to 20 days and upload an object A again, in that case:

- The retention duration for the first version of object A remains to 10 days.
- The later version of the same object is set to 20 days.

When you place an object in the bucket, Lyve Cloud calculates the retention duration for an object version by adding the specified duration to the object version's creation timestamp. The calculated date is stored in an object's metadata and protects the object version until the retention duration ends. When retention duration ends for an object, you can retain or manually delete an object.

By default, object immutability is switched off, and you can switch it on only while creating a bucket. Once object immutability is switched on, Lyve Cloud automatically enables versioning for the bucket. For step-

by-step instructions see below.

To set object immutability:

- Enable object immutability when creating a new bucket, see [Creating buckets](#).
- Optionally, check the Delete objects after the retention duration ends check box.

Managing bucket access permissions

Permissions are used to control access to buckets and define which actions the service accounts are allowed for a bucket. Bucket permission and Policy permission are two options available for granting permission to your buckets.

- **Bucket permission:** Bucket permission is used to set **Read only**, **Write only**, or **All operations** permission for selected buckets. Using Bucket permission, you can grant access permissions to your bucket and the objects in the bucket. Only the admin and storage admin can associate permissions for the buckets. The permissions attached to the bucket apply to all of the objects in the bucket. For more information, see [Creating bucket access permissions](#).
- **Policy permission:** Policy permission is used for creating policy permission by uploading a JSON file. You can also import a file which is compatible with the AWS IAM policy file. Using the Policy permission, you can allow or deny requests at a granular level based on the elements in the policy, resources, and aspects or conditions of the request. For more information, see [Creating policy permissions](#).

Role-based access to permission management

The following table describes access to permission management features based on your role.

Actions	Admin	Storage Admin	Auditor (Read only)
Create permission	✓	✓	×
Edit	✓	✓	×
Delete	✓	✓	×
Status	✓	✓	×
List and view	✓	✓	✓

Creating bucket access permissions

You can create bucket permissions without any buckets in the account only if you apply permission to all buckets in the account or all buckets with a prefix.

To create bucket permissions:

1. On the left-hand menu, select **Permissions**.
2. On the Permissions page, select **Create bucket permission**.
3. In the Create bucket permission dialog, enter the following:

- **Name:** Enter a name for the permission.
- **Which buckets does this permission apply to?:** Select any one from the following:
 - **One or more existing buckets:** Choose one or more buckets from the Buckets list.
 - **Buckets:** The buckets field is displayed on when you select one or more existing buckets.
 - **All buckets in this account with a prefix:**

The bucket names must use the same few initial characters. For example, if four unique buckets for customer01 are created, such as customer01rawdata, customer01zipdata, customer01media and customer01, enter a prefix of the bucket names to assign and apply the permission. In this case, use the same beginning characters for each bucket for our prefix, customer01.



Note

- Only one prefix is allowed for a single permission.
- The prefix field allows a maximum of 64 characters.

- **All buckets in the account:** Apply permission to all current and future buckets in the account.
- **Actions:** Select actions to assign privileges as:
 - **All Operations:** Allows all operations in all buckets meeting the conditions defined under **Which buckets this permission applies to?**
 - **Read only:** This option allows you to perform a read only operation on one or more selected buckets and its objects.
 - **Write only:** This option allows you to write objects into the selected buckets without reading them back.

Create Bucket Permission [Learn more](#) ✕

Name

bucket1

Only alphanumeric names are allowed, including also `_, -, .`, or space

Which buckets does this permission apply to?

☒ One or more existing buckets

☐ All buckets in this account with a prefix

☐ All buckets in this account

Buckets

fmt-test-2 ✕ fmt-21 ✕

Actions

☒ All operations

☐ Read only

☐ Write only

Create

Description

Allow all operations on buckets: fmt-test-2,fmt-21

Once you select the desired options, the description of the permissions is displayed for that bucket permission.

4. Select **Create** to save the permission for a bucket.

The permissions list page displays all permissions. To manage permissions, see [Editing bucket permissions](#) and [Deleting bucket permissions](#).

Creating policy permissions

Lyve Cloud allows the migration of AWS IAM policy files to the Lyve Cloud policy permission, making it simple to start working with service accounts based on existing policies. A policy file uses a JSON file format that is compatible with an AWS IAM policy.

Working with policy files allows you to specify the Condition element. Query the exact request values to determine when a policy is in effect, or list specific actions such as `Action: ["s3:GetObject","s3:PutObject"]` and specify the Resource element for several buckets and objects. For more information, see [Example policy permission file](#).

How to get an IAM policy file from AWS

You must manually copy policy permission details from AWS IAM policy to use in Lyve Cloud:

1. Login to [AWS Management Console](#) using the credentials.
2. Select **Services** on the top left to view the list of services.
3. Select **IAM** in Security, Identity, & Compliance.
4. Under Access Management, select **Policies** and use the search field to find the relevant policy to copy the policy details.
5. Select the **JSON** tab, copy the policy details into a new file, and then save it as a JSON file.

Using policy permission files

The following table lists the mandatory, optional, and invalid elements in a policy permission file.



Note

- Invalid elements must be removed from the file before importing, as these elements are not used in the Lyve Cloud policy permission file.
- Remove tags from elements available in AWS IAM policy, as tags cannot be used in the policy permission file.

Elements	Mandatory/Optional/Invalid	Description
Statement	Mandatory	Contains a single statement or an array of individual statements.
Resource	Mandatory	Specifies object(s) or bucket(s) that is related to the statement.
Effect	Mandatory	Allows or denies access to the resource.
Action	Mandatory	Describes specific action(s) that will be allowed or denied.
Version	Mandatory	It defines the version of the policy language and specifies the language syntax rules that are to be used to process a policy file.
Condition	Optional	<p>Allows you to specify conditions when a policy is in effect.</p> <p>The Condition element includes expressions that match the condition keys and values in the policy file against keys and values in the request.</p> <p>Specifying invalid condition keys returns an error. For more information, see Known Issues.</p>

Elements	Mandatory/Optional/Invalid	Description
Sid	Optional	<p>A statement ID.</p> <p>The statement ID must be unique when assigned to statements in the statement array. This value is used as sub ID for policy document's ID.</p>
Id	Optional	A policy identifier, such as UUID (GUID).
Principal	Invalid	Specifies the service account that is allowed or denied to access a resource.
NotPrincipal	Invalid	The service accounts that are not specified, are allowed or denied access to the resource.
NotAction	Invalid	<p>Specifies that it matches everything except the specified list of actions.</p> <p>If this element is part of the permission file, you need to replace it with the Action element.</p>
NotResource	Invalid	<p>Specifies that it matches every resource except the available specified list.</p> <p>If this element is part of the permission file, you need to replace it with the resource element.</p>

Example policy permission file

In the following example, the policy permission has three statements:

- Statement1**: Allows object listing with a prefix **David** in the bucket **mybucket**. It is done using a Condition element.
- Statement2**: Allows read and write operations for objects with the prefix **David** in bucket **mybucket**.
- Statement3**: Denies delete object operation for two resources:
 - All the objects in **mybucket/David/***
 - All the objects in **mycorporatebucket/share/marketing/***

```
{ "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "statement1",
      "Action": ["s3:ListBucket"],
```

```

"Effect": "Allow",
"Resource": ["arn:aws:s3:::mybucket"],
"Condition": {"StringLike": {"s3:prefix": ["David/*"]}}
},
{
  "Sid": "statement2",
  "Action": [ "s3:GetObject", "s3:PutObject" ],
  "Effect": "Allow",
  "Resource": ["arn:aws:s3:::mybucket/David/*"]
},
{
  "Sid": "statement3",
  "Action": ["s3:DeleteObject"],
  "Effect": "Deny",
  "Resource": ["arn:aws:s3:::mybucket/David/*",
    "arn:aws:s3:::mycorporatebucket/share/marketing/*"]
}
]
}

```

Use the following policy to limit the bucket access to specific IP's:

```

{ "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Sid-1",
      "Action": ["s3:*"],
      "Effect": "Deny",
      "Resource": ["arn:aws:s3:::mybucket"],
      "Condition": {"NotIpAddress": {"aws:SourceIp": ["134.204.220.36/32"]}}
    },
    {
      "Sid": "Sid-2",
      "Action": [
        "s3:*"
      ],
      "Effect": "Allow",
      "Resource": ["arn:aws:s3:::mybucket", "arn:aws:s3:::mybucket/*"]
    }
  ]
}

```

To create policy permission:

1. On the left-hand menu, select **Permissions**.
2. On the Permissions page, select **Create Policy Permission**.
3. In the Create Policy Permission dialog:
 - Enter a name.
 - Edit the description if desired.
 - Drag and drop a policy permission file, or browse to upload a file.
 - Once the new policy permission file is available, download or replace the existing file.

[INSERT create-policy-permission-01.png]

4. Select **Create**.

You might encounter errors if the policy permission file (JSON) has any additional or missing elements. The following is the list of possible error messages. Read them carefully and update the policy permission file accordingly.

Error Message	Resolution
File Import Failed: Invalid JSON file.	Check the JSON file structure.
File Import Failed: Effect field is required.	Add this element to the policy permission file.
File Import Failed: Resource field is required.	
File Import Failed: Action field is required.	
File Import Failed: Statement is required.	
File Import Failed: Version field value is empty.	Add a value to this element.
File Import Failed: Action cannot be empty.	
File Import Failed: Resource cannot be empty.	
File Import Failed: Condition cannot be empty.	
File Import Failed: Effect value is invalid.	Add a valid value to this element.
File Import Failed: Action value < action> is not valid.	
File Import Failed: Resource value < resource> is not valid.	
File Import Failed: Condition name is not valid: <condition entered> .	Choose a valid condition name, such as <code>StringLike</code> .
File Import Failed: Condition key is not valid: <condition key entered> .	Choose a valid condition key, such as <code>s3:prefix</code> .

Editing bucket permissions

Edit existing permissions to change selected buckets and their associated actions.

To edit permissions:

1. On the left-hand menu, select **Permissions**.
2. On the Permissions page, select the ellipsis of the permission to modify, and select **Edit**.

To modify Policy Permission-type permissions:

In the Edit Policy Permission dialog, edit the following:

- Name
- Description
- Policy File: Download or replace the existing file.

Edit Policy Permission ✕

Upload a policy file compatible with an AWS IAM policy [Learn more](#)

Name

permission2

Only alphanumeric names are allowed, including also \\", 'Description

Allow access based on policy file

Policy File

policy.json Download Replace

Save

To modify Bucket Permission-type permissions:

In the Edit Policy Permission dialog, edit the following:

- Name
- Which buckets this permission applies to?
- Actions

4. Select **Save**.

These changes take effect as soon as the updated permission is saved, and any subsequent application API calls will be affected.

Deleting bucket permissions



Note—Permissions used by any service accounts cannot be deleted.

To delete permissions:

1. In the menu, select **Permissions**.

- 2. On the Permissions page, select the ellipsis (...).
- 3. Select **Delete**, and then select **OK** in the confirmation.

After you delete a permission, you cannot restore. However, you can create a new permission and reuse that permission name.

Viewing permissions

By default, the Permissions page displays 10 permissions at a time. You can sort the columns in the table.

To view all permissions:

- 1. In the left-hand navigation, select **Permissions**.



The screenshot shows a table with the following columns: Name, Description, Type, Service Accounts, and Created On. There are two rows of data. The first row has a name 'permissions', description 'Allow access based on policy file', type 'Policy Permission', service accounts '5', and created on '2021-07-20 09:48:19'. The second row has a name 'permissions', description 'Allow all conditions on bucket bucket1', type 'Bucket Permission', service accounts '10', and created on '2021-07-20 09:48:19'. At the bottom right, there is a 'Rows per page' dropdown set to '10' and a '1 of 2' indicator.

Name	Description	Type	Service Accounts	Created On
permissions	Allow access based on policy file	Policy Permission	5	2021-07-20 09:48:19
permissions	Allow all conditions on bucket bucket1	Bucket Permission	10	2021-07-20 09:48:19

The following table describes the columns used to list permissions.

Column Name	Description
Name	Displays name of the permission.
Description	Displays the permission description.
Type	Displays the type of permission created. The type can be Policy permission and Bucket permission .
Service Accounts	Displays the number of service accounts using that specific permission. You can hover the mouse on the number to view the names of the attached service account and the question mark icon to view the tooltip.
Creation On	Displays the date and time when the permission was created in the year, day, month YY:DD:MM AM/PM format.

- 2. Select the arrow next to **Rows per page** to change the number of permissions to list per page.

Administrator's Guide - Account Management

Service accounts allow applications to authenticate and access Lyve Cloud buckets and objects. The appropriate access and secret keys are generated when you create a service account. This information must be saved during the account creation, as you cannot recover key details afterwards. You must create buckets and assign permission to buckets before creating a service account. For more information, see [Creating buckets](#) and [Creating bucket access permissions](#).

Role-based access to manage service accounts

The following table describes access to service account features based on your role.

Actions	Admin	Storage Admin	Auditor (Read only)
Create service account	✓	✓	×
Edit	✓	✓	×
Clone	✓	✓	×
Delete	✓	✓	×
Status	✓	✓	×
List and view	✓	✓	✓
Service account expiration	✓	×	✓

Creating service accounts

You must have at least one associated permission before creating a service account. To set the duration of keys generated after service account creation, you must first configure the expiration period. If the expiration duration is not set, the service account will not have an expiration set, and the secret credentials will never expire. For more information, see [Setting expiration duration](#).

To create a service account:

1. On the left-hand menu, select Service Accounts.

2. Enter the Service Account Name.

- A. Select Permissions from the available list, and select Create.
- B. On the Service Accounts page, select Create Service Account.

i **Note**—Selecting permissions with different Actions (All operations, read only), the action with the least priority is applied to the account.

Create Service Account [Learn more](#)

Service Account Name

(*) Only alphanumeric names are allowed, including dash (-), _ and space.
(*) Maximum 255 characters

Secret Key Expiration Duration 60 days To edit the secret key expiration duration, please go to Settings

Select Permissions

Permission Name	Description
<input type="checkbox"/> owa-demo-permission-44	Allow create, delete and full access to all buckets with name starting with owa-
<input type="checkbox"/> adp2	Allow access based on policy file
<input type="checkbox"/> permission001	Allow all operations on bucket bucket001
<input type="checkbox"/> dudu	Allow access based on policy file
<input type="checkbox"/> mactest	Allow all operations on bucket mactest
<input type="checkbox"/> 3rd-test	Allow all operations on bucket 3rd-test
<input type="checkbox"/> adp2perm	Allow all operations on bucket adp2permtest
<input type="checkbox"/> march	Allow all operations on bucket marchtest
<input type="checkbox"/> perm-match0	Allow all operations on bucket 3rd-match
<input type="checkbox"/> fyecloud-demo-permission-45	Allow all operation on allnewworktest
<input type="checkbox"/> fyecloud-demo-permission-41	Allow create, delete and full access to all buckets with name starting with owa-

[Create](#)

i **Note**—When you configure the expiration duration, the Secret Key Expiration Duration displays the days when the secret key expires. Otherwise, the expiration duration is displayed as Never.

To change the expiration duration, see [Setting expiration duration](#).

If an administrator configures a new expiration duration during the same time frame as the storage administrator creates a service account, the storage administrator receives an information message about the new expiration duration.

Service Account Expiration Duration Change

Your service account expiration duration has recently been changed. Please review the new duration before creating the service account.


OK

3. A confirmation displays the **access key** and **secret keys** required to access the bucket.



Important—Before closing the dialogue, you must copy or download the service account credentials containing the access and secret keys. Download the key in CSV or JSON format, as the secret key details cannot be retrieved later.

The following image displays a generated access key and secret key.



Note—Once you create the service account, it may take a few minutes to replicate across other regions. If you cannot access your storage in a particular region, try after some time.



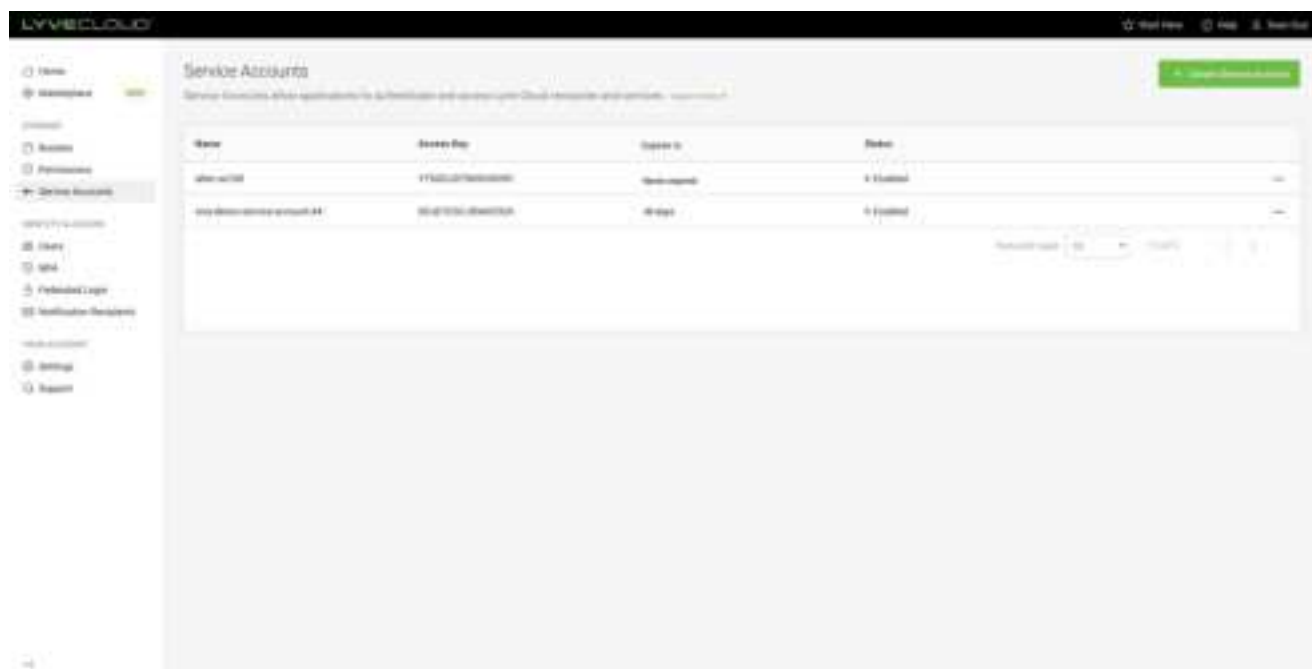
Note—Sometimes there may be a delay in creating a service account.

Viewing service accounts

The service account list displays the Access Key, expiration period, and the status of the service account.

The 'Expires in' column displays any of the following:

- **Expired:** If the service account is already expired.
- **Never Expires:** The expiration period for the service account is not configured.
- **Value:** Displays the remaining days for the service account to expire.



To view the service account list, select **Service Accounts** on the left-hand menu.

- You can view the list of service accounts.
- You can increase the number of service accounts per page.
- You can change the name from **Service_Account_1** to **Service_Account_01**.
- You can add permission3 (new permission) to permission0, permission1 and permission2 (existing). Or you can remove permission0 (existing) from the available list.

You can perform the following operations by selecting the ellipses for each service account:

- Edit service account
- Disable service account
- Clone service account
- Delete service account

Editing service accounts

Editing allows you to edit the service account name and permissions. Editing does not generate a new secret key (credentials) for a service account. To generate new credentials, you must create a new or clone an existing service account. While editing the service account, the access key and expiration period for the service account is displayed. However, you cannot edit them. The expiration period is set when you create a service account. For more information on the expiration period, see [Configuring expiration period](#).

i Note—You cannot edit a service account if the expiration period is over..

If you edit Service_Account_1:When you save this service account, the name and permission of the service account are changed. However, the secret credentials and expiration period remain the same as the original.

To edit a service account:

1. On the left menu, select **Service Accounts**.
2. On the Service Accounts page, select the service account to modify and then select **Edit**.
3. In the Edit Service Account dialog, you can edit the service account name and modify permissions.
4. Select or deselect the permissions to associate with the service account, and scroll to view all available permissions for the account.

Edit Service Account Learn more

Service Account Name

alice-sa100

Only alphanumeric names are allowed, including dash (-), underscore (_), or space

Maximum 255 characters

Access Key: YTSGDJD7W5GHD9C

Expires in: 30 days

Select Permissions

<input type="checkbox"/>	Permission Name	Description
<input type="checkbox"/>	alice-demo-permission-44	Allow create, delete and full access to all buckets with name starting with alice
<input type="checkbox"/>	adp2	Allow access based on policy file
<input type="checkbox"/>	permission001	Allow all operations on bucket bucket001
<input type="checkbox"/>	chub	Allow access based on policy file
<input type="checkbox"/>	bucket1	Allow all operations on bucket bucket1
<input type="checkbox"/>	third-test	Allow all operations on bucket third-test
<input type="checkbox"/>	ad3perm	Allow all operations on bucket adminarchtest
<input type="checkbox"/>	match	Allow all operations on bucket matchtest
<input type="checkbox"/>	rbucket	Allow all operations on bucket rbucket
<input type="checkbox"/>	third-test	Allow all operations on bucket third-test
<input type="checkbox"/>	ad3perm	Allow all operations on bucket adminarchtest
<input type="checkbox"/>	match	Allow all operations on bucket matchtest
<input type="checkbox"/>	per-match0	Allow all operations on bucket third-match
<input type="checkbox"/>	lyvecloud-demo-permission-45	Allow all operation on alianswqthent
<input type="checkbox"/>	lyvecloud-demo-permission-41	Allow create, delete and full access to all buckets with name starting with alice

Save

5. Select **Save** to save changes for the service account.

Changing the status of a service account

The service account is enabled by default. You can disable the service account anytime. Disabling a service account prevents you from using the secret key to authenticate.



Note—You cannot change the status of the service account if the expiration period is over..

To change the status of a service account:

1. On the left-hand menu, select **Service Accounts** to view the list of service accounts.
2. Set Status to **Enabled** or **Disabled** to change the account status.

Deleting a service account

Before you delete a service account, you can disable the key, and once you are sure that the service account is no longer needed, you can then delete the key. Deleting a service account permanently prevents you from using the secret key to authenticate.

To delete a service account:

1. On the left-hand menu, select **Service Accounts**.
2. On the Service Accounts page, select **Delete**.
3. Select **Yes** to delete the service account.

You cannot restore a deleted account. However, you can reuse the service account name to recreate a new service account.

Cloning a service account

Cloning a service account is a quick and easy way to create a duplicate service account. The values of the service account, like the service account name, associated permissions, etc., are the same as the original service account. However, it generates new access and secret keys. The name of the service account appears as a Copy of <service account name>, and you can change the name and associate different or same permission to this service account.

To clone a service account:

1. On the left-hand menu, select **Service Accounts** to view the list of service accounts.
2. Select the ellipses to clone the service account.
3. Select **Clone**, and edit the required fields of the service account.



New secret credentials are generated once you create a service account. For more information, see [Creating service accounts](#).

Service account settings

Adding the expiration duration to the service account enhances the security level of the service account. The existing Service Accounts are set as **Never Expires**. By default, the key never expires when creating a service account unless you configure an expiration duration. You can change the default setting by setting an expiry duration for all newly created service accounts; see [Setting expiration duration](#). This limits the validity of the service account, which needs to be changed again after the expiration duration. After the expiration date, the secret key cannot be used for authentication but will stay associated with the service account until you delete it. If you disable or delete a service account, any workload that uses the service account will immediately lose access to the resources.

As a best practice, change your secret keys regularly. You can create a new secret key by doing the following:

- Create a new service account or Clone the service account.
- Disable the old service account.
- Confirm that the old key is no longer in use.
- Delete the old service account.

Setting expiration duration

Setting an expiration duration enables you to enforce additional security. The more often you change the service account keys, the less likely it is to be leaked. Hence, periodically invalidating your service account keys and creating new keys adds to security.

The Service Account Expiration defaults to **Off** (disabled). The service account key never expires when creating a service account without setting an expiration period. You can turn **On** (enable) the expiration and set the duration in days or years. All service accounts created after you turn on have an expiration period. For example, if you set the expiration duration as 365 days, any service account created after setting the duration has an expiration period of 365 days.

Based on the specified days, the service account expires at the end of the expiration date at 23:59:59 PM, regardless of the time the service account is created. For example, Setting the expiration duration to 30 days on the 1st of the month at 10:15:00 AM, the service account expires on day 30 at 23:59:59 PM.

To set expiration duration:

1. On the left-hand menu, select **Settings**.
2. Enable the **Service Account Expiration** toggle.
3. Enter the number of Day(s) or Years to set the expiration duration, and select **Save**.

Settings

Service Account Settings

Service Account Expiration

Set all service accounts with expiration duration to enforce secret rotation. [Learn more](#)

Expiration Duration

Day(s)

Only whole positive numbers allowed.

Save

Cancel

Audit Log Settings

Audit Log Target Bucket

All audit logs will be saved in this bucket. [Learn more](#)

commvault-webinar-backup-immutable (us-east-1)

Edit

S3 API Audit Logs

Create audit log files for all S3 API operations on a bucket. [Learn more](#)

☒ All buckets must be logged
☐ Individually set per bucket

Console Audit Logs

Create audit log files for console and IAM operations. [Learn more](#)

After the configuration is complete, all new service accounts created will have an expiration duration. Once it expires, you cannot perform any actions; however, you can only delete the service account.

You can configure the expiration duration as 90 days and create a service account. The Secret Key Expiration Duration in the create service account dialogue is set to 90 days. This value is displayed on pages where you create a service account, edit a service account, and list the service account. All service accounts created after configuring expiration duration will be, by default, set to 90 days.

After 90 days, the service account status will appear as **Expired**.

Administrator's Guide - Audit Log Management

Audit logs are detailed records of activities in the Lyve Cloud console and S3 API operations. Audit logs are used to access audit functions and track any suspicious activity.

When you enable audit logging, all audit logs are written to the selected target bucket. The target bucket must be immutable, which keeps audit logs immutable. For more information, see [Using object immutability](#). You cannot switch off object immutability for the target bucket. You can maintain three types of audit logs:

- [S3 API audit logs](#): This log records all supported S3 API calls. For more information, see [Supported S3 API calls](#).

S3 API audit logs are recorded in the `S3-<BUCKET-NAME>-<TIMESTAMP>.gz` format, where the `BUCKET-NAME` is the name of a bucket being logged. For more information, see [Example S3 API audit log](#).

- [IAM audit logs](#): This log includes all events corresponding to identity and access management actions.

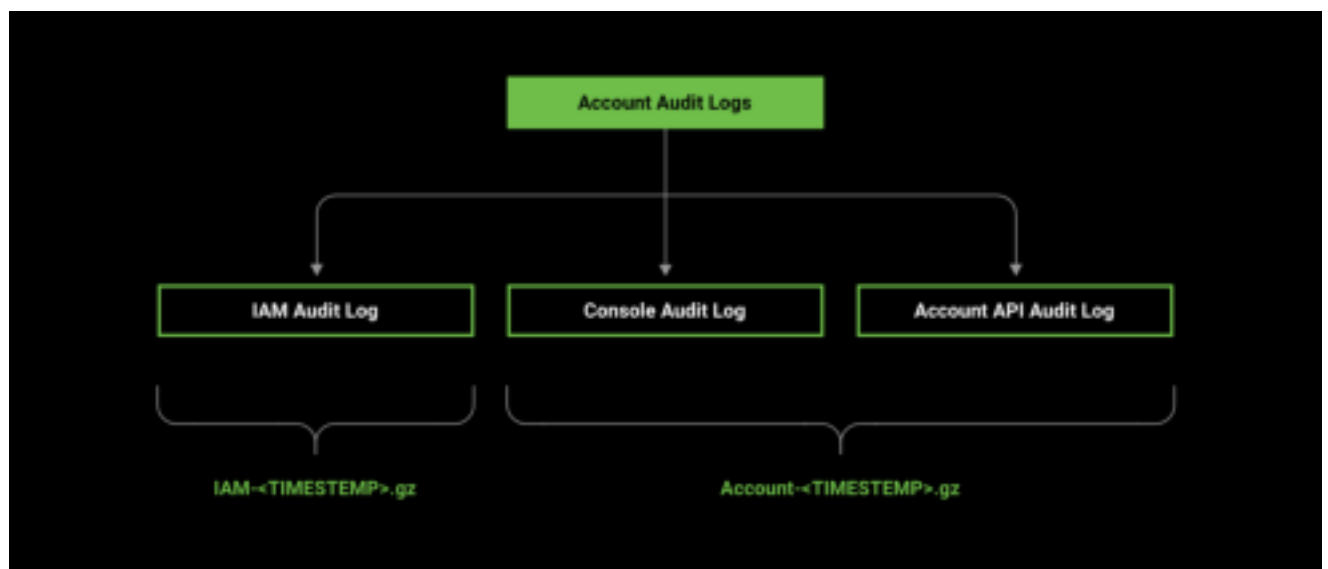
IAM audit logs are recorded in the `IAM-<TIMESTAMP>.gz` format. For more information, see [Example of IAM audit log](#).

- [Console audit logs](#): This log includes all the events that originated from the Lyve Cloud console's actions.

The console audit log is recorded in the `console-<TIMESTAMP>.gz` format. For more information, see [Example of the console audit log](#)



Note—Switching on Console Audit Logs enables both the Console audit logs and IAM audit logs that are written to the target bucket.



The audit log files have `TIMESTAMP` format: `yyyy-MM-dd-HH-mm-ss'` and are set to the UTC zone.

Audit log files keep sufficient information to establish which events occurred, when they occurred, and who caused them. Administrators can manually delete these audit log files after the specified retention duration ends. This helps you to manage the buckets cost-effectively. For more information, see [using object immutability](#).

Lyve Cloud periodically saves audit logs for specified buckets. The maximum size of a log file is 500 MB. If the file size reaches 500 MB, that log file is saved, and the logs continue to be written in a new file. Log files are saved to the target bucket as console audit log files, IAM audit logs, or S3 API logs.

Role-based access to permission

The following table describes access to enable and disable audit logs based on your role.

Actions	Admin	Storage Admin	Auditor (Read only)
Enable/disable S3 API audit logs	✓	×	×
Enable/disable account audit logs	✓	×	×
Edit audit log target bucket	✓	×	×
View audit log settings	✓	×	✓

Video: Lyve Cloud - How to manage audit log settings in the Lyve Cloud console

Seagate on Vimeo: [Lyve Cloud - How to manage audit log settings in the Lyve Cloud console](#)

S3 API audit logs

S3 API audit logs keep detailed records of activity in the Lyve Cloud console as well as S3 API operations. To enable S3 API audit logs, you must select buckets to be logged from the target buckets available in the account.

Example S3 API audit log

The following is an example of an S3 API audit log file.

```
{
  "serviceAccountCreatorId":
    "john.doe@email.com",
  "auditEntry":
    {
      "api":
        {
          "name": "PutObject",
          "bucket": "bucket-1",
          "object": "values-v2.yaml",
          "status": "OK",
          "statusCode": 200,
          "timeToResponse": "2246401314ns" },
          "time": "2021-01-22T10:49:30.699378337Z",
          "version": "1",
          "requestID": "165C883E70C2A5D0",
          "userAgent": "aws-sdk-java/1.12.25 Linux/4.15.0-135-generic OpenJDK_64-Bit_Server_VM/11.0.12+7 java/11.0.12 vendor/O
racle_Corporation cfg/retry- mode/legacy",
          "remotehost": "127.0.0.1",
          "deploymentid": "ef46b1cb-6be1-4aa2-9c14-e7ffbc11986b",
          "requestHeader":{
            "User-Agent": "aws-sdk-java/1.12.25 Linux/4.15.0-135-generic OpenJDK_64-Bit_Server_VM/11.0.12+7 java/11.0.12 vend
or/Oracle_Corporation cfg/retry-mode/legacy",
            "X-Amz-Date": "20210122T104928Z",
            "Content-Type": "text/yaml",
            "Authorization": "AWS4-HMAC-SHA256 Credential=AHPEVYIPHVQ3XNOY/20210122/us-east-1/s3/aws4_request, Signed
Headers=content-type;host;x-amz-content-sha256;x-amz-date, Signature=<redacted>",
            "Content-Length": "5637",
            "X-Amz-Content-Sha256": "UNSIGNED-PAYLOAD",
            "X-Amz-Server-Side-Encryption": "AES256" },
          "responseHeader":{
            "ETag": "219857b61eb0c3dc9a3916a0992fc803",
            "Vary": "Origin",
            "Server": "LyveCloud/DEVELOPMENT.2020-06-22T03-43-44Z",
            "Accept-Ranges": "bytes",
            "Content-Length": "0",
            "X-Amz-Request-Id": "165C883E70C2A5D0",
            "X-Xss-Protection": "1; mode=block",
            "Content-Security-Policy": "block-all-mixed-content",
            "X-Amz-Server-Side-Encryption": "AES256"
          }
        },
      "serviceAccountName": "serv-acc-01"
    }
}
```

The following table describes the parameters specified in the S3 API audit log file.

Parameter name	Description
serviceAccountCreatorId	A user who created the service account.

Parameter name	Description
<code>name</code>	Specifies the API name.
<code>bucket</code>	Specifies the bucket name.
<code>object</code>	Specifies the object name.
<code>status</code>	Specifies the HTTP status.
<code>statusCode</code>	Specifies the HTTP status code.
<code>timeToResponse</code>	Time for the entire request to complete.
<code>time</code>	The timestamp in UTC zone.
<code>version</code>	Represents the current version of Audit Log structure.
<code>requestID</code>	A unique request identifier.
<code>userAgent</code>	Specifies the User-Agent request header
<code>remotehost</code>	Displays IP address of the client who sent the request
<code>deploymentid</code>	A unique deployment identifier.
<code>requestHeader</code>	Specifies the request header content.
<code>responseHeader</code>	Specifies the response header content.
<code>serviceAccountName</code>	Displays the name of Service Account associated with buckets.

Enabling S3 API audit logs

To enable S3 API audit logs:

1. On the left-hand menu, select **Settings**.
2. On the Audit Logs Settings page, set S3 API Audit Logs to **ON** to begin saving audit logs.
3. In the Audit Log Target Bucket dialog, select the target bucket from the list to store the logs.

Set the target bucket only if you are setting the target bucket to write audit logs for the first time. However, if you have already set the target bucket while enabling console audit logs, you are not forced to select the target bucket.



Note—Only the buckets that are immutable are displayed in the list.

Audit Log Target Bucket



Audit logs can only be saved in a bucket that is in compliance mode

Select Bucket

bucket3



Save

4. Select **Save**.

After you enable the S3 API audit log:

To change the target bucket:

- On the Audit Log Settings page, a new 'Audit Log Target Bucket' section is displayed. This section displays the target bucket name and bucket region. To change the target bucket, see [Editing audit log target bucket](#).

To set the S3 API audit logs:

- Select which buckets will have audit log.
 - **All buckets must be logged:** Selecting this option allows you to set and enforce logging for all available buckets in the account. By default, this option is selected.
 - **Individually set per bucket:** Selecting this option allows you to edit each bucket to enable logging manually.



Note—The **All buckets must be logged** and **Individually set per bucket** options are available only after you enable S3 API audit logs, and the target bucket is set to store logs.

After selecting the **Individually set per bucket** option, you must choose each bucket individually and then enable the S3 API audit logs. To enable S3 API audit logs for an individual bucket, see [Editing bucket properties](#). Once S3 audit logs are enabled, the selected bucket in the account is labeled **logged**.

S3 API Audit Logs



Create audit log files for all S3 API operations on a bucket. [Learn more](#) ▶

- ☒ All buckets must be logged
- ☐ Individually set per bucket

Disabling S3 API audit logs

While enabling S3 API Audit Logs, if you select the **individually set per bucket** option and later disable audit logs, the **S3 API Audit Logs** option will be unavailable in that individual bucket.

To disable S3 API audit logs:

1. On the left-hand menu, select **Settings**.
2. On the Audit Log Settings page, set **S3 API Audit Logs** to **Off**.

After you switch off S3 API audit logs, the **Logged** label is removed from all buckets.

Console audit log

Enabling account audit logs enables **Console Audit Logs**, **IAM Audit logs**, and **Account API Audit logs**.

Account Audit Logs

Create audit log files for account and IAM operations. [Learn more](#)



Before you enable them, become familiar with the account audit logs by reviewing the examples below.

Example of Account API audit log structure

```
{
  "ApiEvent":
  {
    EventName: "<eventname>",
    Version: "2",
    Request: {
      AccountName: "service account name",
      AccessKey: "<accesskey>",
      RequestTime: "<hh:mm:ss>",
      RequestParams: {
        .... Parameters.. or Body
      },
      SourceIP: "XXYYXX"
    },
    Response: {
      ResponseTime: "<response time>",
      ResponseCode: "<response code>",
      ResponseError: "<error>",
      ResponseBody: {
        ....
        body.... (with secret redacted) ...
      }
    }
  }
}
```

Event name	API action
<code>authenticate-and-get-session-token</code>	Authenticate and get a session token
<code>Test-a-session-token-for-validity</code>	Test a session token for validity
<code>get-historical-storage-usage-by-month</code>	Get historical storage usage by month
<code>get-current-month-storage-usage</code>	Get current month storage usage
<code>list-permissions</code>	List permissions
<code>create-a-new-permission</code>	Create a new permission
<code>get-permission-by-id</code>	Get permission by id
<code>delete-a-permission-by-id</code>	Delete a permission by id
<code>update-a-permission</code>	Update a permission
<code>list-service-accounts</code>	List service accounts
<code>create-a-new-service-account</code>	Create a new service account
<code>get-service-account-data-by-id</code>	Get service account data by id
<code>delete-a-service-account-by-id</code>	Delete a service account by id
<code>update-a-service-account</code>	Update a service account
<code>enable-a-service-account</code>	Enable a service account
<code>disable-a-service-account</code>	Disable a service account
<code>enable-audit-logging-on-a-bucket</code>	Enable Audit logging on a bucket
<code>disable-audit-logging-on-a-bucket</code>	Disable Audit logging on a bucket

Example of an account audit log

The following is an example of the console audit log file.

```

{
  "ConsoleVersion": "DEVELOPMENT",
  "DeploymentID": "dell2",
  "LoginTime": "2021-01-25T09:19:11.622206Z",
  "UserIdentity": {
    "EventSource": "https://dell2.console.localhost:32428",
    "UserName": "john.doe@email.com",
    "Role": "admin",
    "IPAddress": "10.244.142.100:34310"
  },
  "ConsoleEvent": {
    "Eventname": "add-new-notification-recipient",
    "Status": "Error while inserting data to table: ",
    "StatusCode": 13,
    "EventResponse": "{\"Action\":\"Add NotificationRecipient\",\"FirstName\":\"Fname\",\"LastName\":\"Lname\",\"Email\":\"john.doe@email.com\",\"Partner\":\"dell2\",\"AddedBy\":\"john.doe@email.com\"}",
    "EventTime": "2021-01-25 09:37:01.505980988 +0000 UTC m=+1421.228517562"
  }
}

```

The following table includes console operations recorded in the console audit log. The **Event name** column displays the names inside the console audit log as an `eventname` parameter value.

Event name	Console operation
<code>create-bucket</code>	Create bucket
<code>delete-bucket</code>	Delete bucket
<code>create-permission</code>	Create permission
<code>set-object-immutability</code>	Set object immutability
<code>create-permission-from-imported-file</code>	Create permission from an imported file
<code>edit-permission</code>	Edit permission
<code>delete-permission</code>	Delete permission
<code>create-service-account</code>	Create service account
<code>edit-service-account</code>	Edit service account
<code>service-account-status-change</code>	Service account status change
<code>service-account-deletion</code>	Delete service account

Event name	Console operation
<code>add-user</code>	Add user
<code>user-password-reset</code>	User password reset
<code>edit-user</code>	Edit user
<code>user-enabled-disabled</code>	User enabled/disabled
<code>user-logout</code>	User log out
<code>create-support-ticket</code>	Create support ticket
<code>edit-support-ticket</code>	Edit support ticket
<code>new-comment</code>	New comment
<code>add-new-notification-recipient</code>	Add new notification recipient
<code>remove-notification-recipient</code>	Remove notification recipient
<code>edit-notification-recipient</code>	Edit notification recipient
<code>on-off-s3-api-audit-log</code>	On/off S3 API audit log
<code>on-off-s3-console-audit-log</code>	On/off Console audit log
<code>s3-api-audit-log-setting</code>	S3 API audit log setting
<code>s3-api-audit-log-bucket-setting</code>	3 API audit log bucket setting

The following table describes the parameters specified in the console audit log file.

Parameter name	Description
<code>consoleVersion</code>	Displays the console version
<code>deploymentid</code>	The unique deployment identifier
<code>loginTime</code>	The timestamp in UTC zone
<code>eventSource</code>	The console URL path
<code>userName</code>	The login ID of the user

Parameter name	Description
<code>role</code>	The Lyve Cloud user role
<code>ipAddress</code>	The user identity IP address
<code>eventName</code>	Specifies the console operation
<code>status</code>	Displays the human-readable message
<code>statusCode</code>	Displays the status numeric code. For more information, see Status Code table.
<code>eventResponse</code>	Displays the resulting action performed by the event name
<code>eventTime</code>	Displays the timestamp in UTC zone

Status code

The following tables provide descriptions for the `StatusCode` parameter.

Status code	Error	Error details
0	OK Code	OK is returned on success.
1	Cancelled Code	The operation is cancelled by the client.
2	Unknown Code	Specifies an unknown error. For example, errors raised by APIs that do not return enough error information.
3	InvalidArgument Code	The client specifies an invalid argument.
4	DeadlineExceeded Code	The operation has expired before completion. This error may be returned even if the operation has completed successfully, however, the response is delayed. For example, a successful response from a server could have been delayed long enough for the deadline to expire.
5	NotFound Code	Requested entity (file or directory) was not found.
6	AlreadyExists CodeA	An attempt to create an entity failed because one entity already exists.

Status code	Error	Error details
7	PermissionDenied Code	The caller does not have permission to execute the specified operation.
8	ResourceExhausted Code	Some resource has exhausted.
9	FailedPrecondition Code	The operation was rejected because the system is not in a state to execute the operation. For example, directory to be deleted may not be empty.
10	Aborted CodeT	The operation was aborted due to a concurrent issue like sequencer check failures, transaction aborts, etc.
11	OutOfRange Code	The operation was attempted past the valid range. For example, seeking or reading past end of a file.
12	Unimplemented Code	The operation is not implemented, supported, or enabled for this service.
13	Internal Code	Indicates internal errors, where some invariants have broken.
14	Unavailable Code	The service is currently unavailable.
15	DataLoss Code	Indicates unrecoverable data loss or corruption.
16	Unauthenticated Code	The request does not have valid authentication credentials for the operation.

Example of an IAM audit log

The following is an example of an IAM audit log file.

```
{
  "created_date": "2021-01-20T02:04:12.000Z",
  "organization": "random-org",
  "org_type": "TENANT",
  "source": "console",
  "created_by": "IAM",
  "content": {
    "date": "2016-02-23T19:57:29.532Z",
    "type": "sapi",
    "description": "",
    "connection": ""
  }
}
```



```

"connection_id": "",
"client_id": "AaiyAPdpYdesoKnqjj8HJqRn4T5titww",
"client_name": "My application Name",
"ip": "190.257.209.19",
"hostname": "190.257.209.19",
"user_id": "auth0|56c75c4e42b6359e98374bc2",
"user_name": "",
"audience": "",
"scope": "",
"strategy": "",
"strategy_type": "",
"log_id": "",
"isMobile": false,
"details": {},
"user_agent": "",
"location_info": {
  "country_code": "",
  "country_code3": "",
  "country_name": "",
  "city_name": "",
  "latitude": "",
  "longitude": "",
  "time_zone": "",
  "continent_code": ""
}
}
"bucket_name":""
}

```

The following is an example of an IAM audit log file.

Parameter name	Description
<code>created_date</code>	Date when the event occurred in ISO 8601 format.
<code>organization</code>	Name of the account.
<code>org_type</code>	Displays the type of organization Partner Tenant.
<code>source</code>	Displays source of the Log.
<code>created_by</code>	Displays which service created the log.
<code>content</code>	IAM log content.
<code>bucket_name</code>	Target bucket name where the log files are stored.Optional and can be left blank.

The following table describes the data field for `contents` in the IAM audit log file.

Parameter name	Description
<code>date</code>	Date when the event occurred in ISO 8601 format.
<code>type</code>	Type of event. For more information, see Event code list associated with each log event.
<code>description</code>	Description of the event.
<code>connection</code>	Name of the connection for the event.
<code>connection_id</code>	ID of the connection for the event.
<code>client_id</code>	ID of the client (application).
<code>client_name</code>	Name of the client (application).
<code>ip</code>	The IP address of the log event source.
<code>hostname</code>	Hostname where the event is applied.
<code>user_id</code>	User ID involved in the event.
<code>user_name</code>	User name involved in the event.
<code>audience</code>	API audience for whom the event is applied.
<code>scope</code>	Scope permissions applied to the event.
<code>strategy</code>	Name of the strategy involved in the event.
<code>strategy_type</code>	Type of strategy involved in the event.
<code>log_id</code>	Unique identifier of the event.
<code>isMobile</code>	Specifies if the client is a mobile device (true), desktop, laptop, or server (false).
<code>details</code>	Additional details about the event (the structure is dependent upon event type).
<code>user_agent</code>	User agents details from the client device that caused the event.
<code>location_info</code>	Displays information about the location that triggered this event based on the IP.

The following table describes the data field for `location_info`.

Parameter name	Description
<code>country_code</code>	Displays the country code in two-letter <code>Alpha-2 ISO 3166-1</code> format.
<code>country_code3</code>	Displays the country code in a three-letter <code>Alpha-3 ISO 3166-1</code> format.
<code>country_name</code>	Full country name.
<code>city_name</code>	Full city name.
<code>latitude</code>	Global latitude (horizontal) position.
<code>longitude</code>	Global longitude (vertical) position.
<code>time_zone</code>	Time zone name.
<code>continent_cide</code>	Displays continent of the country. For example, AF (Africa), AN (Antarctica), AS (Asia), EU (Europe), NA (North America), OC (Oceania) or SA (South America).

The following table describes the event code associated with each log event.

Event Code	Description
<code>admin_update_launch</code>	Update launched.
<code>api_limit</code>	The maximum number of requests to the authentication API in given time has been reached.
<code>cls</code>	Passwordless login code/link has been sent.
<code>coff</code>	AD/LDAP connector is offline.
<code>con</code>	AD/LDAP connector is online and working.
<code>cs</code>	Passwordless login code has been sent.
<code>depnote</code>	Deprecation notice.
<code>du</code>	User has been deleted.

Event Code	Description
f	Failed login.
fc	Failed by connector.
fce	Failed to change user email.
fco	Origin is not in the allowed origins list for the specified application.
fcoa	Failed cross-origin authentication.
fcp	Failed change password.
fcph	Failed post change password hook.
fcpn	Failed change phone number.
fcpr	Failed change password request.
fcpro	Failed to provision a AD/LDAP connector.
fcu	Failed to change username.
fd	Failed to generate delegation token.
fdeac	Failed to activate device.
fdeaz	Device authorization request failed.
fdecc	User did not confirm device.
fd�	Failed user deletion.
feactf	Failed to exchange authorization code for access token.
fecctf	Failed exchange of access token for a client credentials grant.
fede	Failed to exchange device code for access token.
fens	Failed exchange for native social login.
feoobft	Failed exchange of password and OOB challenge for access token.
feotptf	Failed exchange of password and OTP challenge for access token.

Event Code	Description
<code>fepft</code>	Failed exchange of password for access token.
<code>fepotpft</code>	Failed exchange of passwordless OTP for access token.
<code>fercft</code>	Failed exchange of password and MFA recovery code for access token.
<code>fertft</code>	Failed exchange of refresh token for access token.
<code>ferrt</code>	Failed exchange of rotating refresh token.
<code>flo</code>	User logout failed.
<code>fn</code>	Failed to send email notification.
<code>fp</code>	Failed login (incorrect password).
<code>fs</code>	Failed signup.
<code>fsa</code>	Failed silent auth.
<code>fu</code>	Failed login (invalid email/username).
<code>fui</code>	Failed to import users.
<code>fv</code>	Failed to send verification email.
<code>fvr</code>	Failed to process verification email request.
<code>gd_auth_failed</code>	Multi-factor authentication failed. This could happen due to a wrong code entered for SMS/Voice/Email/TOTP factors, or a system failure.
<code>gd_auth_rejected</code>	A user rejected a Multi-factor authentication request via push-notification.
<code>gd_auth_succeed</code>	Multi-factor authentication success.
<code>gd_enrollment_complete</code>	A first time MFA user has successfully enrolled using one of the factors.
<code>gd_otp_rate_limit_exceed</code>	A user, during enrollment or authentication, enters an incorrect code more than the maximum allowed number of times. Ex: A user enrolling in SMS enters the 6-digit code wrong more than 10 times in a row.

Event Code	Description
<code>gd_recovery_failed</code>	A user enters a wrong recovery code when attempting to authenticate.
<code>gd_recovery_rate_limit_exceed</code>	A user enters a wrong recovery code too many times.
<code>gd_recovery_succeed</code>	A user successfully authenticates with a recovery code.
<code>gd_send_pn</code>	Push notification for MFA sent successfully sent.
<code>gd_send_sms</code>	SMS for MFA successfully sent.
<code>gd_send_sms_failure</code>	Attempt to send SMS for MFA failed.
<code>gd_send_voice</code>	Voice call for MFA successfully made.
<code>gd_send_voice_failure</code>	Attempt to make Voice call for MFA failed.
<code>gd_start_auth</code>	Second factor authentication event started for MFA.
<code>gd_start_enroll</code>	Multi-factor authentication enroll has started.
<code>gd_tenant_update</code>	Guardian tenant update.
<code>gd_unenroll</code>	Device used for second factor authentication has beenunenrolled.
<code>gd_update_device_account</code>	Device used for second factor authentication has beenupdated.
<code>limit_delegation</code>	Rate limit exceeded to <code>/delegation</code> endpoint.
<code>limit_mu</code>	An IP address is blocked with 100 failed loginattempts using different usernames, all with incorrect passwords in 24 hours, or 50 sign-up attempts per minute from the same IP address.
<code>limit_wc</code>	An IP address is blocked with 10 failed loginattempts into a single account from the same IP address.
<code>pwd_leak</code>	Someone behind the IP address: <code>ip</code> attempted to login with a leaked password.
<code>s</code>	Successful sign-on event.
<code>sapi</code>	Success API operation.
<code>sce</code>	Success change email.
<code>scoa</code>	Success cross-origin authentication.

Event Code	Description
scp	Success change password.
scph	Success post change password hook.
scpn	Success change phone number.
scpr	Success change password request.
scu	Success change username.
sd	Success delegation.
sdu	User successfully deleted.
seacft	Successful exchange of authorization code for access token.
seccft	Successful exchange of access token for a client credentials grant.
sede	Successful exchange of device code for access token.
sens	Native social login.
seoobft	Successful exchange of password and OOB challenge for access token.
seotpft	Successful exchange of password and OTP challenge for access token.
sepft	Successful exchange of password for access token.
sercft	Successful exchange of password and MFA recovery code for access token.
sertft	Successful exchange of refresh token for access token.
srrt	Successfully revoked a refresh token.
slo	User successfully signed out.
ss	Success signup.
ssa	Silent auth.
sui	Successfully imported users.
sv	Verification email.

Event Code	Description
svr	Verification email request.
sys_os_update_end	Update ended.
sys_os_update_start	Update started.
sys_update_end	Update ended.
sys_update_start	Update started.
ublkdu	User block setup by anomaly detection has been released.
w	Warnings during login.

Enable account audit logs

To enable the account audit log:

1. On the left-hand menu, select **Settings**.
2. On the Settings page, set **Account Audit Logs** to **ON**.
3. in the Audit Log Target Bucket dialog, select the target bucket from the list to store the logs.

You must set the target bucket only if you are setting the target bucket to write audit logs for the first time. However, if you have set the target bucket while enabling S3 API audit logs, you are not forced to select the target bucket again.

Only the buckets that are in immutable are displayed in the list.

4. Select **Save**.

Disable account audit logs

To disable account audit logs:

1. On the left-hand menu, select **Settings**.
2. On the Settings page, set **Account Audit Logs** to **OFF**.

After you disable audit logs:

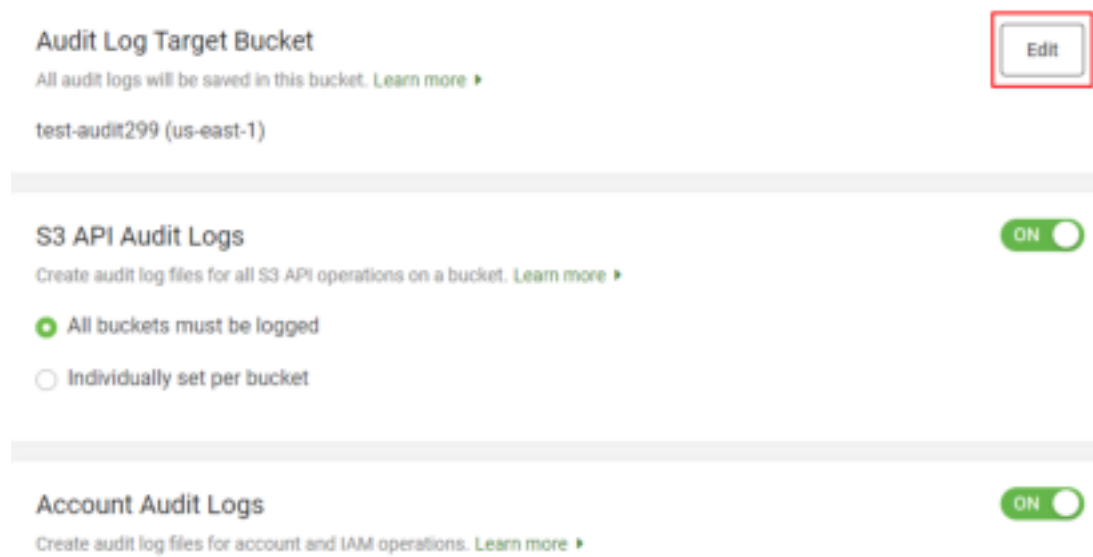
- Logs are no longer saved to the target bucket.
- The target bucket is still visible in the Audit Logs Target Bucket section.

Edit audit log target bucket

While editing the target bucket to save audit logs, only immutable buckets are displayed for selection.

To edit console audit logs:

1. On the left-hand menu, select **Settings**.
2. In the Audit Log Target Bucket section, select **Edit**.



3. On the Edit Audit Log Target Bucket dialog, select the target bucket from the **Select bucket** list and then select **Save**.

Administrator's Guide - Identity and Access Management (IAM)

Lyve Cloud's identity and access management (IAM) enables you to manage users and access to Lyve Cloud console resources.

Create users and assign roles, add multi-factor authentication for enhanced security, or set up federated login using your organization's IDP (identity provider) and SAML 2.0.

Role-based access to IAM

The following table describes access to IAM features based on your role.

Actions	Admin	Storage Admin	Auditor (Read only)
IAM	✓	×	✓
Users	✓	×	✓
MFA	✓	×	✓
SAML Federation	✓	×	✓
Notification Recipients	✓	×	✓

Managing users and roles

The Users page allows you to create users and set user roles. A user is an individual customer who can perform various actions in the Lyve Cloud console based on the assigned role. A role restricts the actions a user may perform, which prevents unauthorized access to Lyve Cloud features.

About user roles

You can set distinct roles for Lyve Cloud users. These users can perform actions based on assigned roles, see Role-based access sections in the respective topics.

- **Administrator** —An administrator can perform all the operations in the Lyve Cloud console.
- **Storage Administrator** —The storage administrator can manage all storage-related actions that includes managing buckets, permissions and service accounts in Lyve Cloud.

- **Auditor** —An auditor has read only access to the Lyve cloud console, and thus cannot perform any actions.

About user and authentication types

User types

In Lyve Cloud, there are two distinct user types:

- **Password user**—Users whose username and password are managed in Lyve Cloud
- **Federated user**—Users who are authenticated via their organization's identity provider (IdP).



Note—Federated users only exist when SAML Federation is configured on the account.

Authentication types

Set an authentication type while creating a user. The following are the available authentication types for each of the user types.

Multi-factor authentication (MFA) is enabled by default for IAM users. Multi-Factor Authentication required two authentication methods:

- **Password**—Set a password as the first factor of authentication. All IAM users must set a password. For more information, see [Registration workflow for password authentication type](#)

A password policy is applied to all user accounts created and managed directly in Lyve Cloud. For more information, see [Password policy](#).

- **OTP or SMS**—Set SMS or OTP as the second authentication factor. For more information, see [Using multi-factor authentication \(MFA\)](#).

Federated users have the following authentication type:

- **Federated**—This option is available only when configuring SAML Federation for the account. For more information, see [Configuring Federated Login](#).

Add a user

To add a user:

1. On the left-hand menu, select **Users**.
2. On the Users page, select **Add User**.
3. In the Add New User dialog box, enter the following:
 - **First Name**: Enter the first name of the user.
 - **Last Name**: Enter the last name of the user.
 - **Email**: Enter the email address of the user.



Note—You cannot modify the email address after adding a user.

- **Role:** Select from the options Admin, Storage Admin, Auditor.
- **Authentication Type:**

If SAML Federation is not configured, no selection is required, and the following are possible display options:

- **Password**

If SAML Federation is configured, no selection is required, and the following are possible display options:

- **Federated**
- **Password**

4. Select **Add User**.

An invitation email is sent to the IAM user to complete the registration process. For information, see [Registering an IAM user in Lyve Cloud](#).

The following image displays the Add New User dialog box.



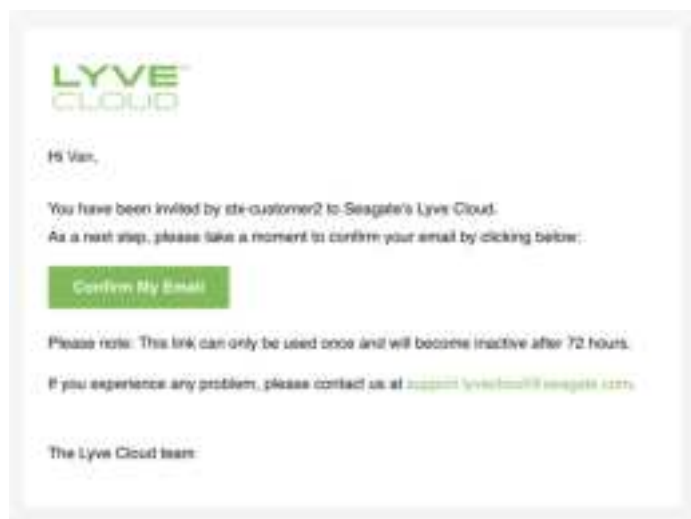
The image shows a screenshot of the 'Add New User' dialog box in Lyve Cloud. The dialog has a title bar with a close button. Below the title, there are three text input fields labeled 'First Name', 'Last Name', and 'Email'. Below these fields is a 'Role' section with three radio button options: 'Admin' (selected), 'Storage Admin', and 'Auditor'. Each option has a description: 'Full access' for Admin, 'Manage buckets, permissions and service accounts' for Storage Admin, and 'Read-only access' for Auditor. Below the role selection, it says 'Authentication type is: Password' and provides a link to 'enable Federated Authentication'. At the bottom is a green 'Add User' button.

Registering an IAM user in Lyve Cloud

When an IAM user is registered in Lyve Cloud, the user receives an email invitation. They must register in

Lyve Cloud by Multi-factor Authentication (MFA), a security method that is set by default for all accounts.

The following image displays a sample email invite. This email invitation link expires within 72 hours.



If the user doesn't select the link within 72 hours, they should select **Forgot Password** on the login page. For more information, see [Registering after an email invitation link expires](#). Check your spam folder if you believe you did not receive an email invitation or contact the support team at support.lyveloud@seagate.com to complete the registration process.

Registration workflow for password authentication type

After the user is registered in Lyve Cloud and receives an email invitation, they should complete the registration.

To complete the registration:

1. Select the link provided in the invitation email to get started.
2. in the Create Password dialog, enter and confirm your password and then select **Create**.



Note—Refer to the [Password Policy](#) while creating a new password.

3. Once a password is created, the user is taken to the Lyve Cloud Login page.

Registering after an email invitation link expires

To complete registration after an email invitation expires:

1. Select the **Forgot Password** link. This page appears after you select the **Click on this link to get started** link in the invitation mail.
2. Follow step 2 onwards from the [Registration workflow for Password Authentication Type](#).

Viewing and editing a user

An administrator can change the first name, last name, and the assigned role of an IAM user. Only administrators can edit or redefine roles for users, they cannot edit or change roles defined for themselves. If a change to an administrator role is desired, a different administrator must make the change.

i Note—Once a profile is edited, the respective user must log out of Lyve Cloud and log back in for role changes to take effect.


The following table describes the column names in the user's table:

Column Name	Description
First Name	Displays the user's first name.
Last Name	Displays the user's last name.
Email	Displays the user's email address.
Authentication Type	Displays the user's authentication type. For more information, see Using multi-factor authentication .
Role	Displays role of the selected user. See About user roles .
Status	Displays the user's status as either Enabled or Disabled .

To view or edit a user:

1. On the left-hand menu, select **Users**.
2. On the Users page, find the user you want to change.
3. Select the ellipses (...) in the right-most column of the user's role and then select **Edit**.

First Name ↑	Last Name	Email	Authentication Type	Role	Status
Ali	Ree	ali.ree@regala.co...	Password	Admin	Enabled
Michael	Wright	michael.wright@reg...	Password	Admin	Enabled
Bill	Hansen	bill.hansen@regala...	Password	Admin	Enabled
Francis	Tucker	francis.tucker@reg...	Password	Admin	Enabled
Samuel	Markus	samuel.markus@re...	Password	Admin	Enabled



4. In the Edit User dialog box, edit the following and select Save.

- First Name: Enter the first name of the user.
- Last Name: Enter the last name of the user.
- Select a Role to modify from the following options: Admin, Storage Admin, or Auditor.

You can also see the Authentication Type of the user, which is read-only and not editable.



Note—While editing a user, you cannot modify the email address.

Disabling or enabling a user

To enable or disable a user:

1. On the left-hand menu, select **Users**.
2. On the Users page, find the user you want to enable/disable.
3. Select the ellipsis and select **Disable** or **Enable**.
4. Select **Yes** in the confirmation to change the status.

Deleting a user

To delete a user:

1. On the left-hand menu, select **Users**.
2. On the Users page, find the user you want to delete.
3. Select the ellipsis and select **Delete**.
4. Select **Yes** in the confirmation to change the status.

Resetting password

To reset user password:

1. On the left-hand menu, select **Users**.
2. On the Users page, find the user whose password you want to reset.
3. Select the ellipsis and select **Reset Password**.
4. Select **Yes** in the confirmation to reset the password.

Video: Lyve Cloud - How to manage users and assign roles

Seagate on Vimeo: [Lyve Cloud - How to manage users and assign roles](#)

Password policy

A password policy is applied to all user accounts that are created and managed directly in Lyve Cloud.



Note—The password policy is not applicable for federated users.

The following password policy options are defined and must be fulfilled.

Property	Requirements
Characters allowed	A - Z, a - z, 0 - 9, (!, @, #, \$, %, ^, &;, *)
Number of characters	between 8 - 128 characters
Password Type	<p>The password must contain three of the following four character types:</p> <ul style="list-style-type: none"> • A lower-case letter • An upper-case letter • A number • A special character (!, @, #, \$, %, ^, &;, *)
Password restrictions	<ul style="list-style-type: none"> • Setting passwords to common options like password, 123456, 12345678, 1234, qwerty, etc. • Setting passwords that contain their personal data like name, username and nickname. The first part of the user's email will also be checked firstpart@example.com

Password change history

This prevents users from recycling old passwords; the last five passwords can't be used again when the user changes a password. The password change history determines the number of unique new passwords associated with a user account before an old password can be reused.

Password expiration

The password expiration policy determines the period of time (in days) that a password can be used before it requires the user to change it. The password will expire 180 days from the date when the password is updated. The password expiration date is updated to 180 days once the user changes the password.

Example: The password that is changed on 1 January 2022 will be set to expire on 30 June 2022.

Users will receive two email notifications, the first one before seven days, and another one before three days to reset the password. This email includes a link to change the password.

Restricting password

The password policy does not allow users to use the most commonly used passwords. The following restrictions include:

- **Commonly used password**

See the [restricted list](#) to view the list of passwords that are not allowed.

- **Personal data**

It prohibits users from setting passwords that contain any of their personal data.

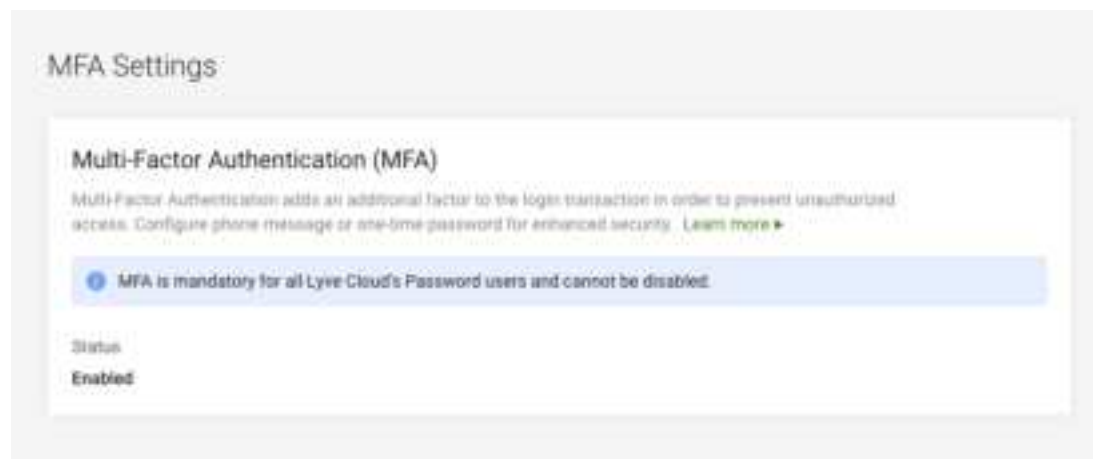
For example: name, username, nickname, user_metadata.name, user_metadata.first, user_metadata.last, first part of the user's email (firstpart@example.com)

If the user's name is John, the user would not be allowed to include John in their password. For example, John1234 will not be allowed.

Using multi-factor authentication (MFA)

Multi-factor authentication (MFA) is a security mechanism that adds a layer of protection to the sign-in process to access the Lyve Cloud console. Choose the one-time password (OTP) option (using third-party authenticator apps such as Google, Microsoft, or Oracle Mobile Authenticator) or the SMS text message as the second level of authentication.

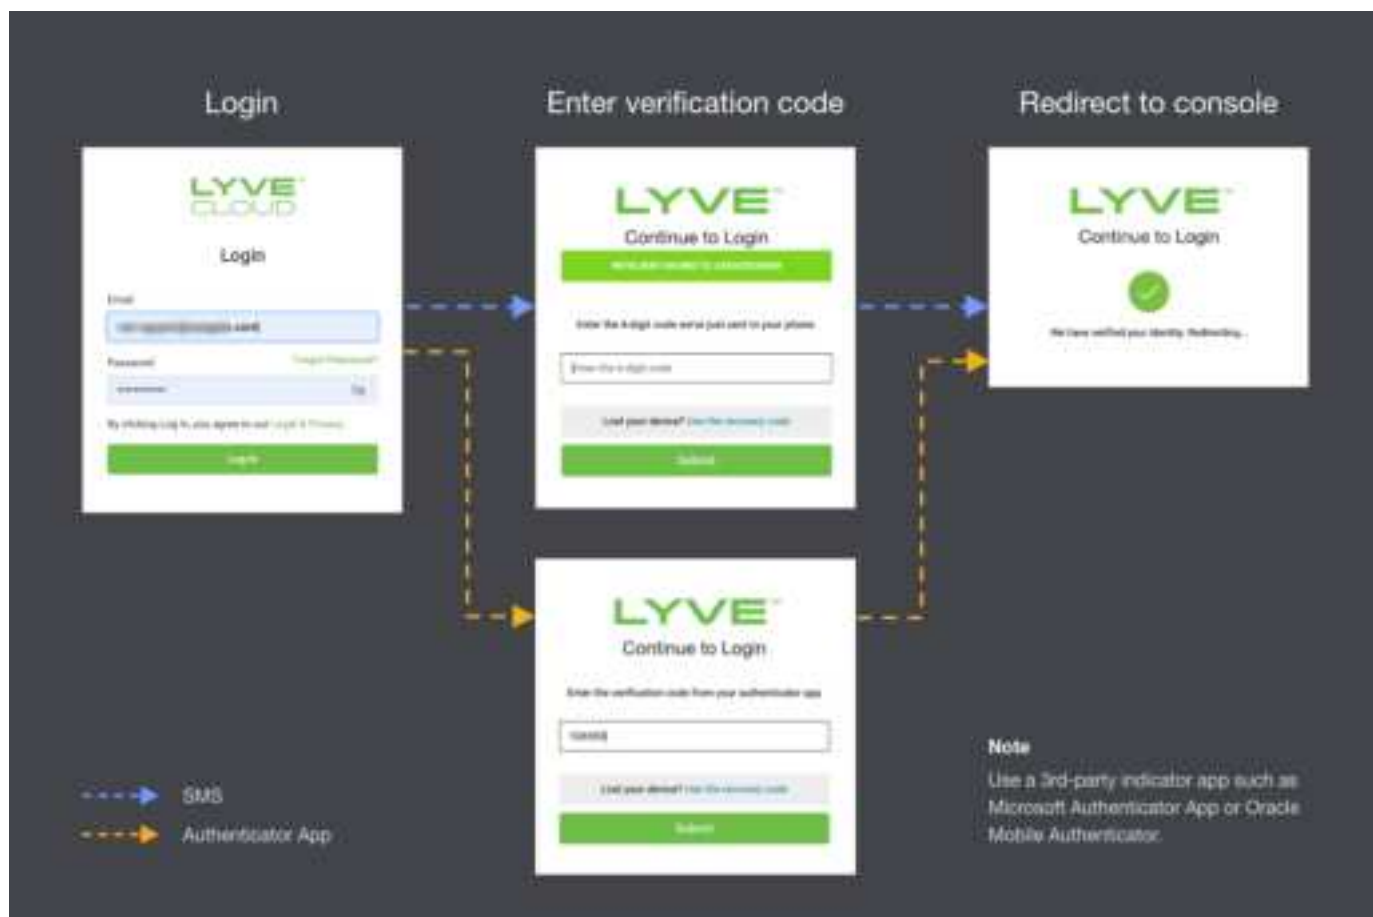
By default, MFA is configured for all Lyve Cloud users and users must go through an additional registration step to complete MFA setup. The registration occurs during their successful first sign-in attempt. Users are then prompted to use the second level of authentication for subsequent login attempts. Setting cannot be changed to disable MFA.



About MFA

MFA configuration provides these 2nd-factor authentication options:

- **Password + SMS text message**—Sets password and SMS as a two-factor authentication policy for users to sign in to the Lyve Cloud console.
- **Password + One-Time Password (OTP)**—Sets password and OTP as a two-factor authentication policy for users to sign in to the Lyve Cloud console. To use OTP, users must download a third-party authenticator to their phone.



How does MFA work?

If MFA is enrolled

All IAM users are required to set up their 2nd-factor authentication separate from their password. Users can choose one of these MFA methods:

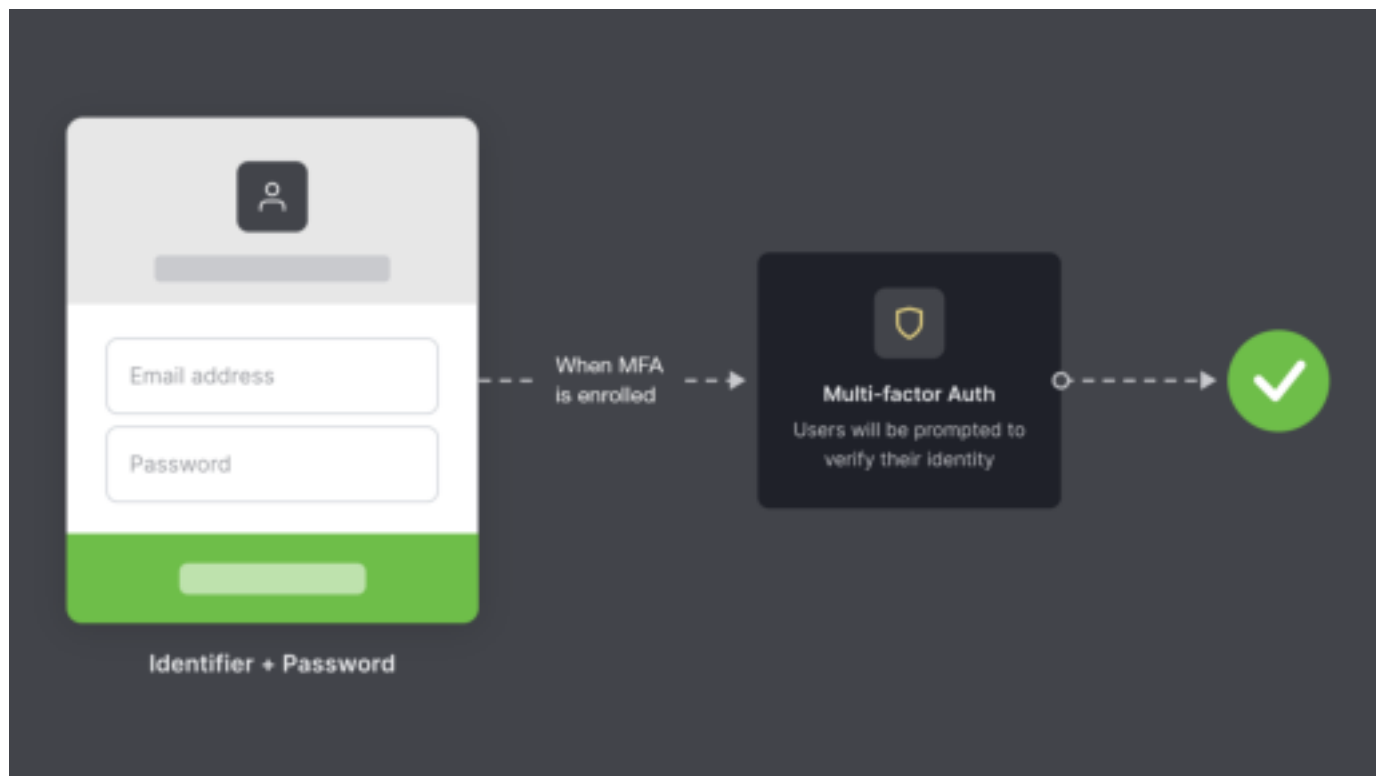
- Using the authenticator app—The authenticator app must be installed on your phone. Use the authenticator app to scan the QR code during the MFA enrollment. This enables the one-time password (auto-generated every 30 seconds) to be entered whenever you log in to Lyve Cloud.
- Using an SMS service—Use any mobile device with a phone number able to receive SMS text messages. When an MFA code is needed, Lyve Cloud sends a six-digit verification code to the phone number configured by the IAM user. Text messaging charges from the mobile carrier apply when choosing SMS-based MFA.

When a user attempts to login to Lyve Cloud, a code is sent via SMS, which the user has to enter to complete the transaction.

If MFA is not enrolled

If any user (new user or existing user) has set a password as the first method of authentication but does not enroll in MFA, in that case the user is prompted to enroll either using SMS service or the authenticator app after successful login using the password.

Users are not allowed to login into the Lyve Cloud console if they do not enroll in MFA.



Enrolling in MFA

MFA is configured for all Lyve Cloud users in the account to include an additional verification method. You can enroll after you have set the password. For more information, see [Registration workflow for password authentication type](#).

To enroll in MFA:

You must set up the required MFA enrollment to access Lyve Cloud. Every time you log in to Lyve Cloud, you must follow the two-step authentication process. Lyve Cloud requests users to enter the OTP generated from the authenticator application or SMS.

1. Log in to Lyve Cloud using your credentials.
2. Select either the authenticator app or SMS as your second authentication method:

If you choose the authenticator app option, scan the QR code from the authenticator app on your phone.



Continue to Login

Scan this QR code with any Authenticator app



[Trouble scanning?](#)

Enter your passcode here

I'd rather use [SMS](#)

Submit

Enter the the one-time passcode (OTP) displayed on the authenticator app and select Submit.



Note—Use any 3rd-party authenticator app such as Google, Microsoft or Oracle Mobile Authenticator. The authenticator app generates a random OTP and expires within a time limit.

If you choose to use SMS, select **I'd rather use SMS** located below the QR code field:



Continue to Login

Scan this QR code with any Authenticator app



[Trouble scanning?](#)

I'd rather use [SMS](#)

Submit

- Select the Country code.
- Enter the phone number to receive the SMS passcode and select **Continue**.
- Enter the code received on your phone as an SMS and then select **Submit**.

The image shows two screenshots of the LYVE login process. The left screenshot is the 'Continue to Login' screen, which includes a QR code, a passcode field, and a 'Submit' button. The right screenshot is the 'Continue to Login' screen, which includes a phone number field, a 'Submit' button, and a 'Trouble scanning?' link.

- Once the verification code is entered, save the recovery code.



Continue to Login

In the event that you need to login without your device you'll need a recovery code. Take a note and keep this somewhere safe:

A - Z - 7 - Z - A - T

☒ I have safely recorded this code

Submit



Note—Save a copy of the secret key in a secure place. If you lose the MFA device, you can use the recovery code to log in. The recovery code allows one-time login to the Lyve Cloud console.

4. Check the **I have safely recorded this code** checkbox and select **Submit** to complete MFA enrollment.

Resetting MFA for an individual IAM user

The Reset MFA feature allows admins to reset the IAM users' MFA enrollment. The reset action removes the old MFA entry. The user will then be unable to sign in to the Lyve Cloud console until they reset the MFA.

Make sure the users have an active phone number if you want to set Password + SMS Text Message, or an authenticator app installed on their phone if you want to set Password + One-Time Password (OTP) as your authentication type.

To reset MFA:

1. On the left-hand menu, select **Users**. A list of users is displayed on the Users page.
2. Select the ellipsis next to the user, then select **Reset MFA**.
3. To reset that user's MFA, select **Yes**.

After MFA is reset, users must re-enroll in MFA. See [Enrolling in MFA](#).

Configuring Federated Login

Federated Login provides authentication without revealing user login credentials to the Lyve Cloud service. Federated Login enables your users to use a single authentication method with the help of your organization's Identity Provider (hereafter referred to as IdP) for Lyve Cloud users. Once the Lyve Cloud user signs in and has access to your organization's domain, the user then has direct access to the Lyve Cloud console. Hence the user need not perform a separate login process. To use Federated Login feature, your organization must have an authentication system which uses the SAML 2.0 protocol.

To configure Federated Login, contact your organization's IdP administrator to obtain the metadata file in XML format. Upload this file to configure Federated Login.

Security Assertion Markup Language (SAML) protocol

The Security Assertion Markup Language (SAML) protocol is an open-standard, XML-based framework for authentication and authorization between two entities without a password:

- A Service Provider (SP) agrees to trust the identity provider to authenticate users.
- An Identity Provider (IdP) authenticates users and provides to service providers an authentication assertion that indicates a user has been authenticated.

In this scenario, Lyve Cloud is a Service Provider that will connect with your organization's Identity Provider to establish a Single Sign-On (SSO) access to your users.

Configure Lyve Cloud as a SAML service provider

To configure Lyve Cloud as a SAML service provider, you will need to:

1. Obtain metadata and certificate from your IdP administrator
2. Configure Lyve Cloud as a service provider
3. Add service provider metadata to the identity provider
4. Configure the identity provider to send email attribute

Obtain metadata and certificate from your IdP administrator

Contact your organizations IdP administrator and obtain the metadata file in XML format to upload and configure Federated Login.

For more information on generating a metadata file for Okta, Google Gsuite, and Microsoft Azure, see [Generating XML Metadata files for IdP](#).

Configure Lyve Cloud as a service provider

1. On the left-hand menu, click the SAML Federation menu.



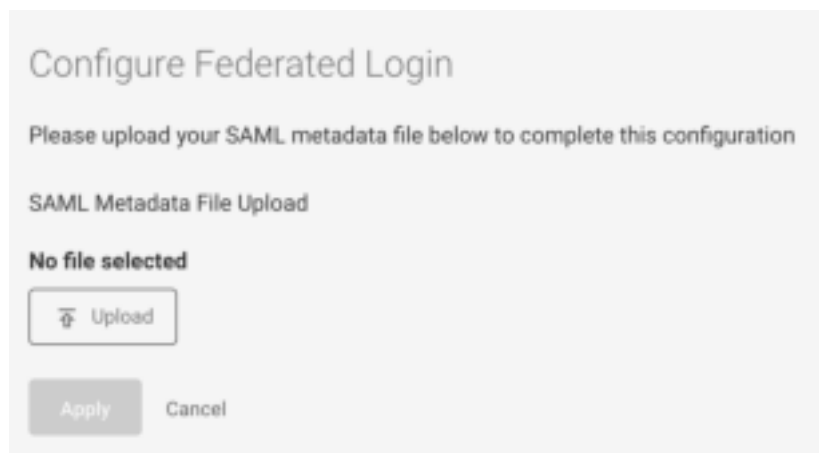
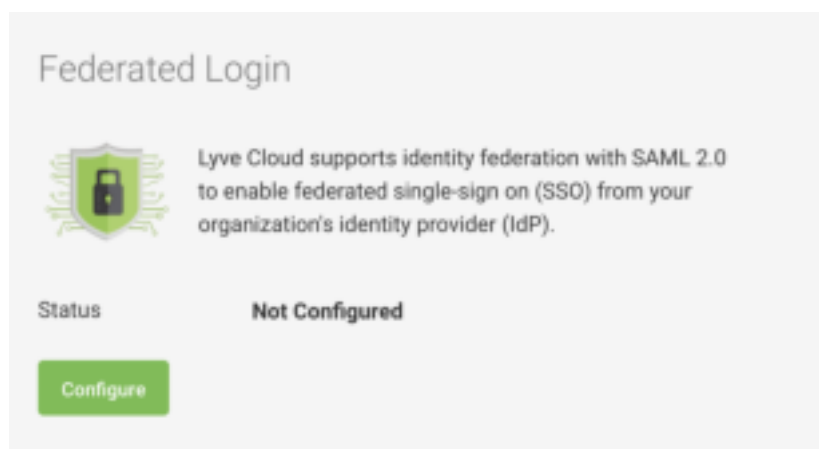
Note—If you have not configured Federated Login, the status is displayed as Not Configured.

1. On the Federated Login page, click Configure.
2. In the Configure Federated Login page, click Upload.
3. Select the XML file from the desired location, and select Open.
4. After SAML Metadata File is uploaded successfully, click Apply.



Note—You need to re-upload the file in case it is an invalid file.

The following image displays a Federated Login set up.



After configuration, the Federated Login page displays:

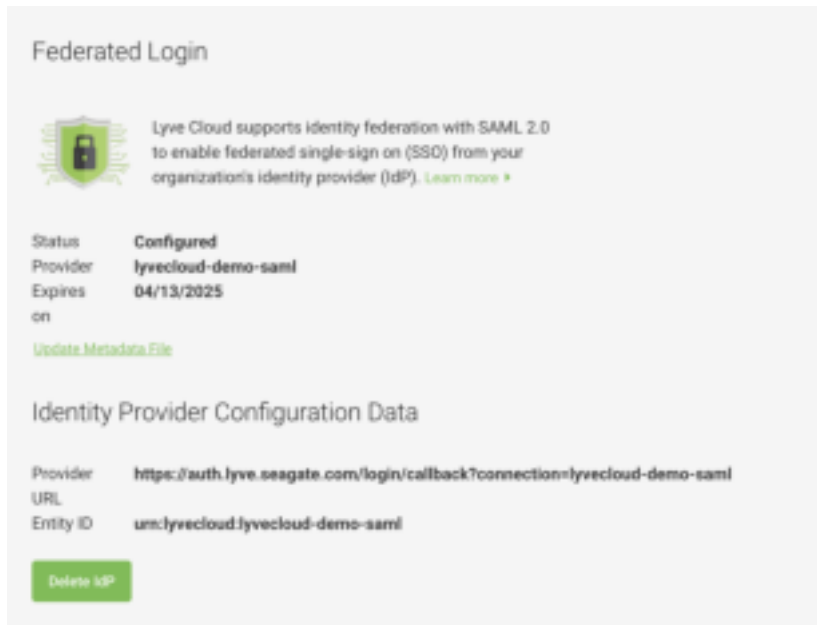
- Status as **Configured**
- Status name of the IdP Provider
- Metadata file's Expiry Date.

In addition, the Identity Provider configuration details are provided. The following attributes are used to configure the IdP:

- Provider URL

- Entity ID

The following image displays sample IdP configuration details.



Add service provider metadata to the identity provider

In this section, you will add some information to the IdP, so it knows how to receive and respond to SAML based authentication requests from the Lyve Cloud service provider. The instructions provided here are generic. You will need to find the appropriate screens and fields on the Identity Provider.

1. Locate the screens from the Identity Provider that allow you to configure SAML.
2. The IdP must know where to send the SAML assertions after it has authenticated a user. This is the **Provider URL** in Lyve Cloud. The IdP might call this **Assertion Consumer Service URL** or **Application Callback URL**. `https://YOUR_TENANT_URL/login/callback?connection=YOUR_CONNECTION_NAME`
3. The connection URL parameter is required for identity provider-initiated flow.



Note—If you have custom domains set up, use the custom domain-based URL rather than your Lyve Cloud domain in the following format:

`https://auth.lyve.seagate.com/login/callback?connection=YOUR_ACCOUNT_ID`

4. Enter **Entity ID** in the **Audience** or **Entity ID** field from Lyve Cloud:
`audience:urn:lyvecloud:YOUR_TENANT:YOUR_CONNECTION_NAME`
5. If IdP provides a choice for bindings, you should select **HTTP-Redirect** for Authentication Requests.
6. The **Single Logout Service URL**, where SAML logout requests and/or responses from the Identity Provider must be sent and should be configured as `https://YOUR_DOMAIN/logout`
7. **Signing Logout Requests:** When configuring the IdP, make sure that SAML Logout Requests sent to the service provider are signed.

Configure the identity provider to send email attribute

Lyvecloud reads “email” attribute from your identity profile. Some identity providers send “email” by default, while some require you to configure it to send “email”.

Auth0 and Azure

Sends email by default. No additional configuration is required.

G Suite

Requires a configuration to send an email attribute.

To configure an email attribute for G Suite:

1. Log in to the GSuite app. Click **Apps** in the left menu, and then click **Web and mobile apps**.
2. Select the SAML app to edit and update attribute mapping.
3. In the SAML attribute mapping section, click the arrow to edit.

SAML attribute mapping

Map Google Directory user profile fields to SAML service provider attributes.

email

Basic Information > Primary email



4. In the Attributes mapping section, add the following attributes and click **Save**.

Okta

Requires a configuration to send an email attribute.

To configure an email attribute for Okta:

1. On the left-hand menu, select **Applications** and then click **Applications**.
2. On the Applications page, click the application to edit, and then select **General**.
3. Click **Edit** in SAML settings.
4. Click **Next** in General Settings without making any change.
5. In the Attribute Statements (optional) section, click **Add Another** to add an attribute statement, and update the following attributes:

- Name = email
- Value = user.email

Attribute Statements (optional) [LEARN MORE](#)

Name	Name format (optional)	Value
email	Unspecified ▼	user.email ▼

[Add Another](#)

Troubleshooting SSO

If your application doesn't work the first time, you should clear your browser history and cookies before you test again. The browser may otherwise not pick up the latest version of your HTML page or it may have outdated cookies that impact execution.

While troubleshooting SSO:

- Capture an HTTP trace of the interaction: Use any of the available tools to capture the HTTP traffic from your browser for analysis.
 - Search for HTTP Trace
 - Capture the login sequence from start to finish and analyze the sequence of GETs to determine how far in the expected sequence is achieved.
 - See a redirect from your original site to the Service Provider and then to the Identity Provider.
 - A post of credentials if you had to log in.
 - Then a redirect back to the callback URL or the Service Provider.
 - Finally a redirect to the callback URL specified in your application.
- Ensure the cookies and javascript are enabled for your browser.
- Check to make sure that the callback URL specified by your application in its authentication request is listed in the Allowed Callback URLs field.
- The <http://samltool.io> tool can decode a SAML assertion and is a useful debugging tool.

Updating the metadata file

You need to update the metadata file before the certificate expires. Contact your IdP administrator to get the updated XML file. If you make any updates and regenerate metadata.xml, you must delete the old metadata file and reupload the updated file. If you just upload the updated file, it may not make changes to the old file.

To update the metadata file:

1. Log in to the AWS IAM console.

1. On the left-hand menu, click **SAML Federation**.
2. On the SAML Federation page, click **Update Metadata File**.
3. Select the XML file from the desired location and click **Open**.
4. After the SAML Metadata File is uploaded successfully, click **Apply**.

Deleting existing IdP configuration

To delete the IdP:

1. On the left-hand menu, click **SAML Federation**.
2. On the SAML Federation page, click **Delete IdP**.
3. In the Delete IdP configuration confirmation box, click **Yes**.

Generating XML metadata files for IdP

Different types of IdP products have their own way of generating XML metadata files. See the following sections for details on generating XML metadata files for:

- Okta
- Google GSuite
- Microsoft Azure Active Directory (AD)

Generating XML metadata files for Okta

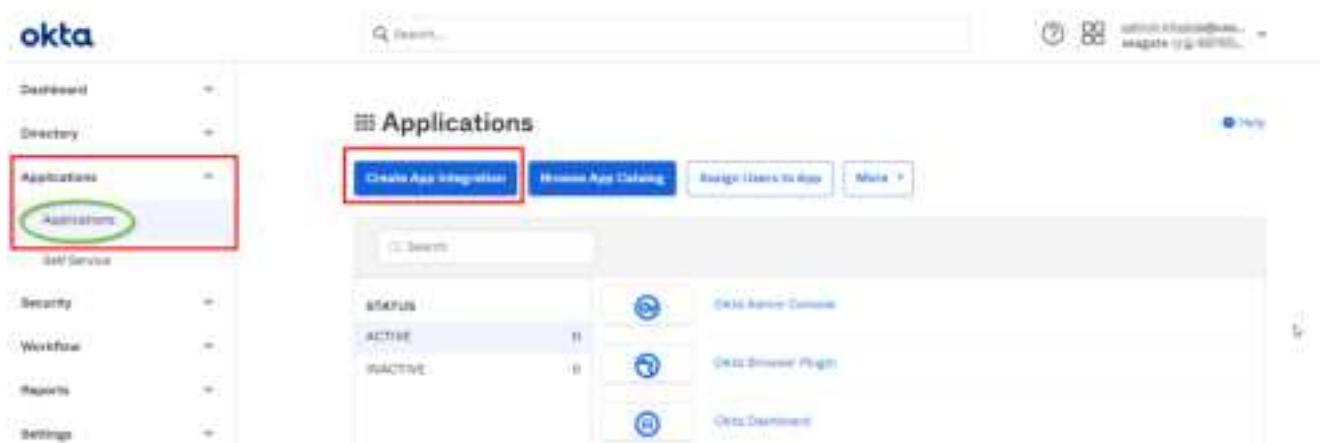
Pre-requisites

- Create an Okta account and add a user as an administrator for configuration.
- Lyve Cloud account name (Tenant name) and administrators account in the console.

Generate an XML file for Okta

To generate an XML file for Okta:

1. Create an application in Okta for Lyve Cloud and log in as administrator.
2. On the left-hand menu, select **Applications** and then click **Applications**.
3. Click **Create App Integration**.



4. In the Create a new app integration dialog, select **SAML 2.0**, and then click **Next**.



5. In the Create SAML Integration section, enter the App name in the General Settings.
6. In the Configure SAML section, select the following SAML Settings:

- A. **Audience URI (SP Entity ID):** Enter the SP entity ID in the following format: `urn:lyvecloud:<TENANT>-saml`
- B. **Single sign on URL:** Enter the URL in the following format: `https://auth.lyve.seagate.com/login/callback?connection=<TENANT>-saml`.

For example, consider your Lyve Cloud account (tenant) is mylctenant1 in this case:

- The Single sign on URL is: `https://auth.lyve.seagate.com/login/callback?connection=mylctenant1-saml`
- The SP Entity ID uri: `lyvecloud: mylctenant1 -saml`



Note—The Single sign on URL and SP Entity ID is generated in configuring Lyve Cloud federated login 2.2. If the values are different as mentioned in the above step, you must update the attributes in the application.

A SAML Settings

General

Single sign on URL ⓘ

https://auth.lyve.seagate.com/login/callback?connection=

☒ Use this for Recipient URL and Destination URL

☐ Allow this app to request other SSO URLs

Audience URI (SP Entity ID) ⓘ

:lyvecloud: mytenant1 -saml

Default RelayState ⓘ

If no value is set, a blank RelayState is sent

Name ID format ⓘ

Unspecified ▼

Application username ⓘ

Okta username ▼

Update application username on

Create and update ▼

[Show Advanced Settings](#)

C. In the Attribute Statements section set the following values and click **Next**.

- Name: email
- Value: user.email

Attribute Statements (optional)

[LEARN MORE](#)

Name	Name format (optional)	Value
email	Unspecified ▼	user.email ▼

[Add Another](#)

D. In the Feedback section, provide feedback to help them understand why the Okta application was configured.

E. Select the appropriate option if you are an OKTA customer or a partner, and click **Finish**.

3 Help Okta Support understand how you configured this application

Are you a customer or partner?

- ☒ I'm an Okta customer adding an internal app
- ☐ I'm a software vendor. I'd like to integrate my app with Okta

[Previous](#)

[Finish](#)

7. After the application is generated, you must retrieve the XML metadata file.

To retrieve XML metadata file:

- Click the **Sign On** tab.
- In the Settings section, click **Save** and add .xml extension to the file.
- This is the XML file that is used to configure Lyve Cloud federation.

General **Sign On** Import Assignments

Settings [Edit](#)

Sign on methods

The sign-on method determines how a user signs into and manages their credentials for an application. Some sign-on methods require additional configuration in the 3rd party application.

Application username is determined by the user profile mapping. [Configure profile mapping](#)

☒ SAML 2.0

Default Relay State

SAML 2.0 is not configured until you complete the setup instructions.

[View Setup Instructions](#)

Identity Provider metadata is available if this application supports dynamic configuration.

How can an Okta user log in to Lyve Cloud?

To login to Lyve Cloud user as an Okta user:

- Add a user to your Okta account and assign apps to the users.
- First assign the Okta application to the users. Create a user and assign the authentication type as Federated, ensure that the account is configured as Federated Login. For more information, see [Add a user](#).

https://seagatelyveps-admin.okta.com/admin/app/bulk-assign

Change your passw... All | Mobile Device... My files - OneDrive Sharepoint

40 days left in your trial. [Contact Sales](#)

See Okta's Business Continuity plan for COVID-19 here: [Our Commitment to Customer Success: People, Business, and Service Preparedness](#)

Search...

Assign Applications

1 Assign Apps to People 2 Confirm Assignments

[Cancel](#) [Next](#)

Applications 1

Application & Label	Sign-on
<input checked="" type="checkbox"/> okta-app-ic	SAML 2.0

First Previous 1 Next Last

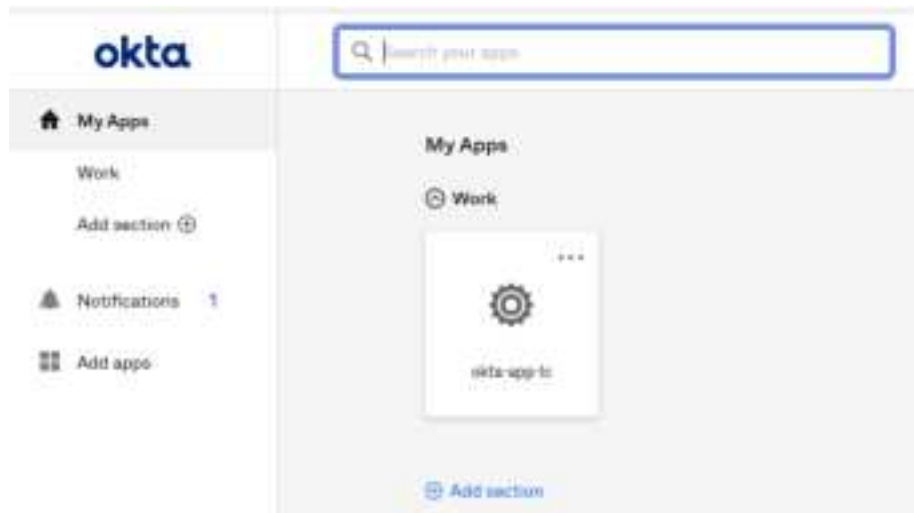
People 1

Person & Username	Status
<input type="checkbox"/> allen liu allen.liu@seagate.com	Active
<input checked="" type="checkbox"/> Allen Gmail Liu allenliu2000@gmail.com	Active

First Previous 1 Next Last

An Okta user can log in to Lyve Cloud account by two ways:

1. Select the tile on the Okta Home page to connect to Lyve Cloud. You are redirected to the Lyve Cloud console and are logged in automatically.



2. Copy the embed link from the General tab of the Okta application and paste into a browser.

App Embed Link

[Edit](#)

Embed Link

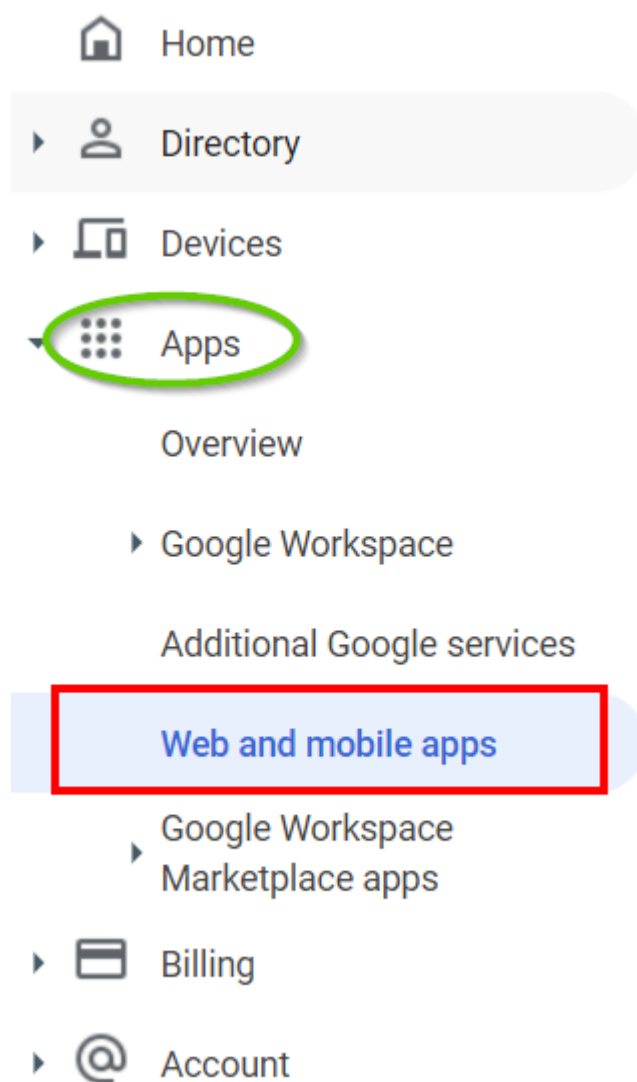
You can use the URL below to sign into okta-app-ic from a portal or other location outside of Okta.

https://seagate.lyveps.okta.com/home/seagate.lyveps_oktaappic_1/00a52sr2nAvEQ3cR696/aIn52vu5mJoKCTaJ4696

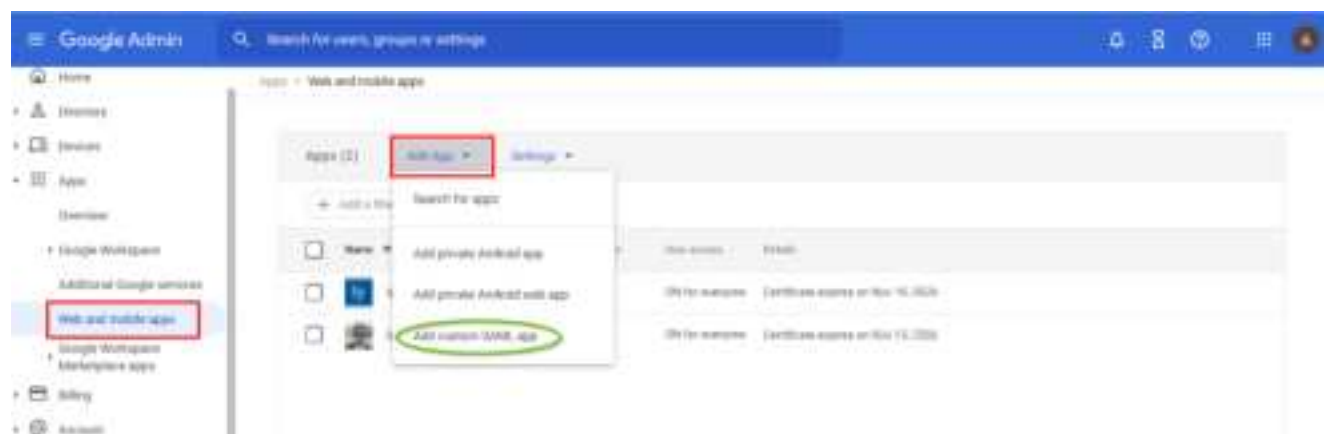
Generating XML metadata files for Google GSuite

To generate an XML file for Google GSuite:

1. Create an application in GSuite for Lyve Cloud and log in as administrator.
2. Click **Apps** in the left menu, and then click **Web and mobile apps**.




3. In the Web and mobile apps section, click Add App and then select Add custom SAML app.



4. In the App details section, provide the App name and optionally upload the App icon. Select **Continue**.

App details
Enter details for your custom SAML app. This information is shared with app users. [Learn more](#)

App name

App icon
Attach an app icon. Maximum upload file size: 4 MB


CANCEL

CONTINUE

5. Click **DOWNLOAD METADATA** and save the file as an .xml extension. Select **Next**.

[illegible]

6. In the Service Provider details section, enter the following info:

- **ACS URL:** Enter the URL in the following format: `https://auth.lyve.seagate.com/login/callback?connection=<TENANT>-saml`
- **Entity ID:** Enter the SP entity ID in the following format: `urn:lyvecloud:<TENANT>-saml`
- For example, Consider your Lyve Cloud account (tenant) is mylctenant1 in this case:
 - The Single sign on URL is `https://auth.lyve.seagate.com/login/callback?connection=mylctenant1 -saml`

- The SP Entity ID is urn:lyvecloud: mylctenant1 –saml

7. Enter the following info in the NAME ID section:

- In **Name ID Format**, select **EMAIL** from the drop-down list.
- Select **Basic Information > Primary email** from the drop-down list.

8. Select **Continue**.

Service provider details
To configure single sign-on, add service provider details, such as ACS URL and entity ID. [Learn more](#)

ACS URL
https://auth.lyve.seagate.com/login/callback?connection=mylctenant1-saml

Entity ID
urn:lyvecloud: mylctenant1 –saml

SSO URL (optional)
☐ Required username

Name ID
Defines the naming format supported by the identity provider. [Learn more](#)

Name ID Format
EMAIL

Name ID
Basic information > Primary email

BACK CANCEL CONTINUE

9. In the Attributes mapping section, add the following attributes and click **Finish**.

- **Google Directory attributes**: Select **Primary email**.
- **App attributes**: Select **email**.

Attributes
Add and select user fields to Google Directory, then map them to service provider attributes. Attributes marked with * are mandatory. [Learn more](#)

Google Directory attributes App attributes

Back (optional) *

Primary email → email

ADD MAPPING

BACK CANCEL FINISH

/p>

10. Once the app is successfully created, the configuration details are displayed.



11. In the User access section, enable the user access. Select **ON for everyone** and then select **Save**.



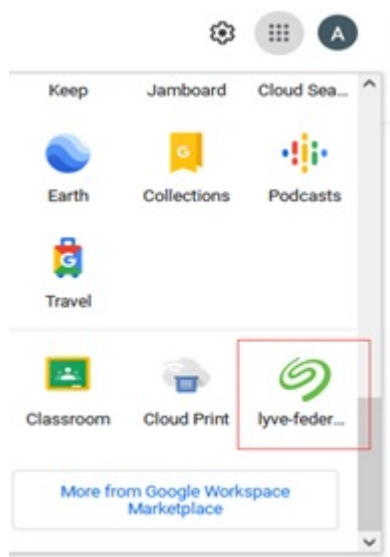
12. Download the Metadata file from the account.



How can a GSuite user log in to Lyve Cloud?

To login to Lyve Cloud user as an GSuite user:

1. Log in to your Google account at www.google.com
2. In the right-up corner of the page, select the Lyve Cloud app icon:



Generating XML metadata files for Microsoft Azure AD

To generate an XML file for Microsoft Azure AD:

1. Create an application in Azure portal for Lyve Cloud and log in as administrator.
2. In the navigation pane, click **Azure Active Directory** and then select **Enterprise applications**.
3. In the Enterprise application section, click **New application**.



4. In the **Browse Azure AD Gallery** section, click **Create your own application**.



5. In **Create your own application**, enter a name for the application and select the purpose of using the application as **Integrate any other application you don't find in the gallery (Non-gallery)**, and then select **Create**.

Create your own application



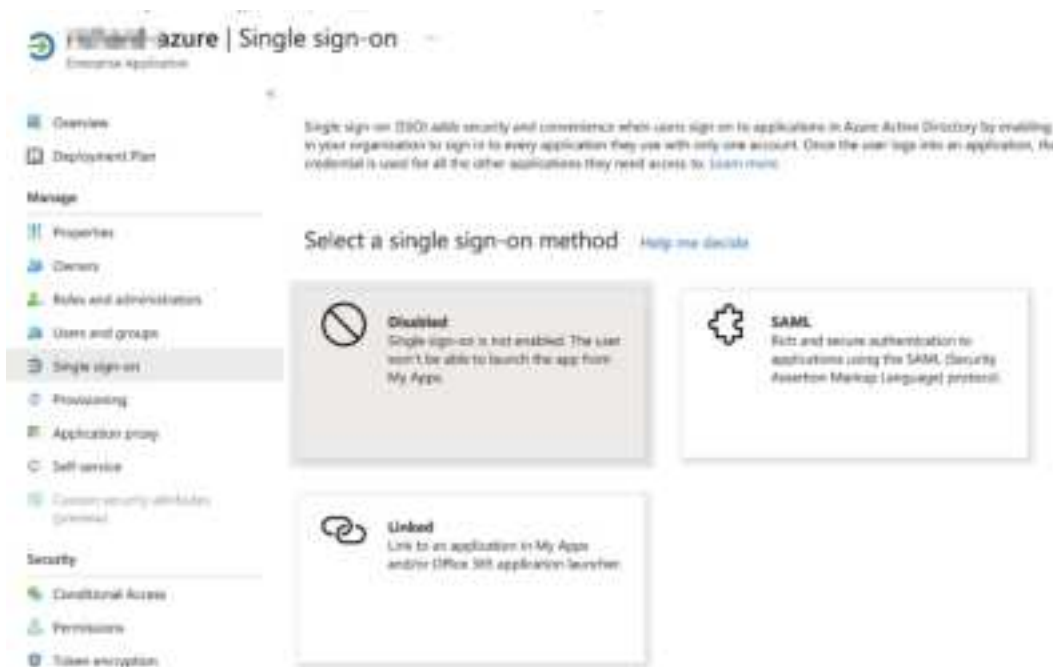
What's the name of your app?

What are you looking to do with your application?

- ☐ Configure Application Proxy for secure remote access to an on-premises application
- ☐ Register an application to integrate with Azure AD (App you're developing)
- ☒ Integrate any other application you don't find in the gallery (Non-gallery)

Create

6. In the navigation pane, select **Single sign-on**, and then click the **SAML**.



In the **Basic SAML Configuration** section, click **Edit**:

Basic SAML Configuration		Edit
Identifier (Entity ID)	https://mycompany:443/webconsole	
Reply URL (Assertion Consumer Service URL)	https://mycompany:443/webconsole	
Sign on URL	Optional	
Relay State	Optional	
Logout Url	Optional	

Enter the following:


- Identifier (Entity ID): urn:lvvecloud:<tenant short name>:saml.

For example: urn:lyvecloud:seagate-saml

- Reply URL (Assertion Consumer Service URL): https://auth.lyve.seagate.com/login/callback?connection=&tenant_short_name&-saml .

For example: <https://lyvecloud-sandbox.us.auth0.com/login/callback?connection=seagate3-saml>

7. In the SAML Signing Certificate section, click Download against Federation Metadata XML.



This XML metadata file is used to upload for Federation configuration on Lyve Cloud console. For more information, see [Configure Lyve Cloud as a service provider](#).

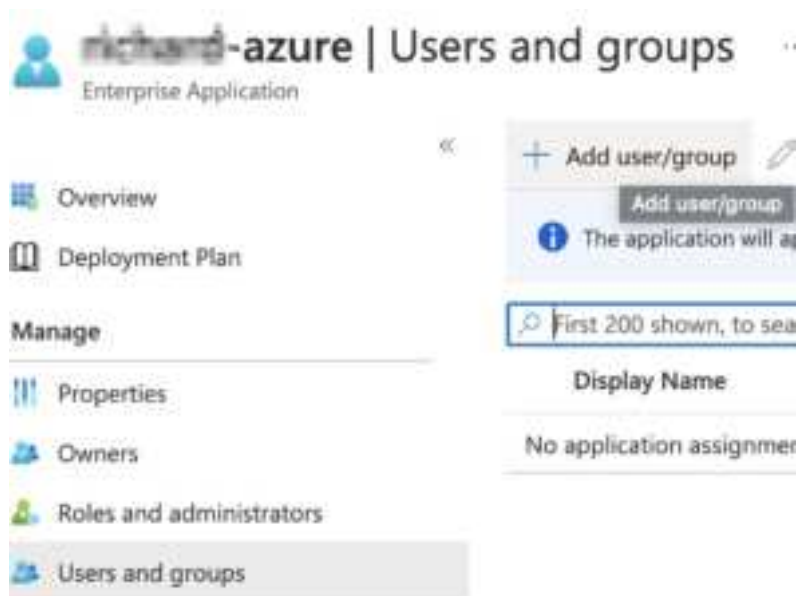
How can a Microsoft Azure user log in to Lyve Cloud?

Before you login, ensure the user must be an existing AD user, so users can utilize SSO via Azure AD.

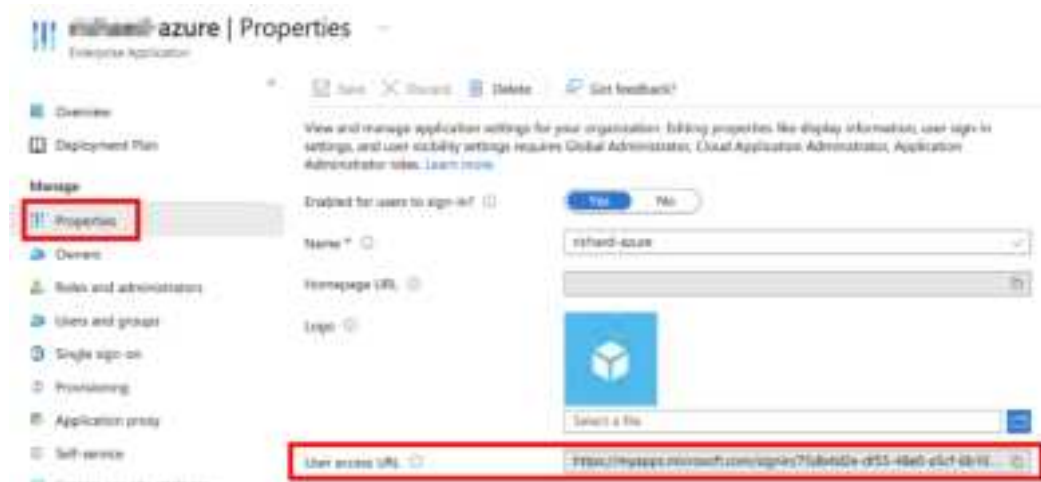
To login to Lyve Cloud user as an Azure user:

1. Log in to the Azure account as the administrator.

Select Users and groups and then select Add user/group to add the existing Azure user. This user is granted permission to use Lyve Cloud application.



2. Select **Properties** and copy the **User access URL**. You must provide this URL to the user to login Lyve Cloud using Azure credentials.



Managing notification recipients

The notification feature allows people to receive service notices and other important Lyve Cloud information via email. Add any number of notification recipients, even if they are not registered as Lyve Cloud users. Any user that is added as an Administrator, Storage Administrator, or Auditor is automatically added to the recipient list.

Once you add someone to the notification list, they start receiving service announcements and product-related emails.

Recipients can stop receiving these notifications in the following ways:

- An administrator removes the recipient from the notification list.
- The recipient selects the **Unsubscribe** link in any notification email.
- An administrator disables the registered Lyve Cloud user.

Adding a recipient

To add a recipient to receive notifications:

1. On the left-hand menu, select **Notifications**.
2. On the Recipients page, select **Add New Recipient**.
3. Enter the **First Name**, **Last Name**, and **Email**, and then select **Add**.

Editing a recipient

You can only edit data for recipients who are not registered as users in Lyve Cloud.

To edit a recipient:

1. On the left-hand menu, select **Notifications**.
2. On the Notifications page, select **Edit**.
3. Edit the user information and select **Edit**.

The recipient's list on the Notifications page displays the changes.

Removing a recipient

To remove a recipient:

1. On the left-hand menu, select **Notifications**.
2. On the Notifications page, select **Remove** next to the recipient's name.
3. Select **Yes, Remove** in the confirmation.

The recipient is now removed from the list.



Note—If the recipient selects the Unsubscribe link in the notification email and accepts the confirmation, the recipient is removed from the list.

Video: Lyve Cloud - How to add a notification request

Seagate on Vimeo: [Lyve Cloud - How to add a notification request](#)

Administrator's Guide - Sub-Account Management

The Sub-accounts feature allows you to maintain a multi-level account structure. Lyve Cloud provisions the master account, and the master account provisions Sub-accounts. These Sub-accounts are managed by and connected to the master account. Each sub-account can provision and manage its own storage, users, and account settings.

Large organizations can use Sub-accounts to support different departments. It can be used by Service Providers to manage their independent implementations or by storage resellers to aggregate the usage of different Sub-customers.



Note—Some items may not be available to some users, depending on the permissions applied by your administrator.

Master accounts

Lyve Cloud provisions the master account for an organization. A master account manages the entire lifecycle of its Sub-accounts, including creating, disabling, and deleting a Sub-account. The master account has the details for each Sub-accounts usage and is responsible for billing all its Sub-accounts. The master account can create a limited number of Sub-accounts. The administrator can request to increase the limit by contacting support at support.lyvecloud@seagate.com. For more information, see [Creating a support ticket](#).

The master account can view all Sub-account statistics on the dashboard, and this helps the organization identify its own and its customer's sub-accounts.

The following image displays the dashboard of the master account.

- Home
- Marketplace
- STORAGE
- Regions
- Buckets
- Permissions
- Service Accounts
- IDENTITY & ACCESS
- Users
- MFA
- Federated Login
- Notification Recipients
- YOUR ACCOUNT
- Settings
- Billing
- Sub-accounts
- Support

Home

General Overview

Sub-accounts

5

Buckets

75

Average Usage ⓘ

100 TB

Estimated Cost (USD) ⓘ


2,850.76

General Report

This month


Download

Daily Average Usage ⓘ



Usage Report

Usage by Accounts ⓘ



Accounts Summary

Account ID	Users	Service Accounts	Buckets	Average Usage ⓘ ↓	Created On
sub-account-1	4	2	26	221.5 TB 23%	YYYY-MM-DD
sub-account-2	4	2	17	191 TB 20%	YYYY-MM-DD
sub-account-3	4	2	9	87.5 TB 13%	YYYY-MM-DD
sub-account-4	4	2	11	94.5 TB 13%	YYYY-MM-DD
sub-account-5	4	2	9	45.5 TB 7%	YYYY-MM-DD
master-account	4	2	9	0 B 0%	YYYY-MM-DD

Rows per page 10

1-6 of 6

The master account can view a list of all Sub-accounts. When the master account administrator selects a Sub-account name, it displays the Sub-account dashboard. For more information, see [Understanding sub-account dashboard](#).

Using support

Reseller's sub-accounts can create and manage support tickets through the Lyve Cloud console, while managed service providers (MSPs) are responsible for creating and managing support tickets for their sub-accounts. While creating a support ticket, you must include the organization name, phone number to contact and details concerning the issue. To create a support ticket, see [Creating a support ticket](#).

If you do not have a support option available, please contact your administrator with the required details of the issue.

Managing Sub-accounts

Role-based access to manage Sub-account

The following table describes access to manage Sub-account features based on the user's role:

Actions	Admin	Storage Admin	Auditor (Read only)
Sub-account menu	✓	×	✓
Create Sub-account	✓	×	×
Disable Sub-account	✓	×	×
Delete Sub-account	✓	×	×
View Sub-account details	✓	×	✓
View all storage usage	✓	✓	✓

Creating a Sub-account

Based on the user type, the fields are available for the administrator to create a Sub-account (see the table above).

The number of Sub-accounts per master account is limited. If you need additional Sub-accounts beyond the limit, please create a support ticket to request an increase to the limit.



Note—Some users may not see all the attributes for the Sub-account based on the attributes set for the master account.

To create a Sub-account:

1. On the left-hand menu, select **Sub-accounts**.
2. On the Sub-accounts management page, select **Create Sub-account**.
3. Enter the following:

- **Organization Information**

- **Organization Name:** Specify the name of the organization. The name is chosen when your first create a sub-account and is provisioned for your organization. The length must be between 3-64 characters.
- **Enable Lyve Cloud Support:** This option displays the support menu option in the left navigation pane for the Sub-account.
- **Street Address:** Specify the street address of the organization.
- **City:** Specify the city name of the organization.
- **State:** Specify the state of the organization.
- **Zip Code:** Specify the zip code.
- **Country/Region:** Specify the country or region of the organization.
- **Phone Number:** Specify the contact number.

The screenshot shows a web form titled "Organization Information" with a subtitle "An organization represents a legal entity that a sub-account belongs to. Learn more". The form contains several input fields: "Organization Name" (a single-line text box), "Enable Lyve Cloud Support" (a checkbox that is checked), "Street Address" (a single-line text box), "Street Address 2 (Optional)" (a single-line text box), "City" and "State" (two side-by-side single-line text boxes), "Zip Code" and "Country/Region" (two side-by-side single-line text boxes, with a dropdown arrow on the right of the Country/Region box), and "Phone Number" (a single-line text box).

- **Account Information**

- **Account Name:** Specify the name of the account. You specify an account name when you create a new account, while an account name uniquely identifies an account within your organization. Only alphabet letters, numbers, spaces, dash (-), ' (single quote), (_) underscore, (,) comma, (&) ampersand and (.) periods are allowed. The length must be between 3-64 characters.
- **Account ID:** Account ID is unique across all the Lyve Cloud accounts. A single URL is used to access the Lyve Cloud console authenticated by the account ID. The account id is used in the URL, and the URL is account specified in the following format: `https://<account_ID>.console.lyvecloud.seagate.com`. You cannot change the account ID once it is created and is generated immediately.

Account Information
An account provides the basis for creating, enabling and using all Lyve Cloud services. [Learn more](#)

Account Name

Account ID 

- **Admin Details:** Specify the administrator's details managing the Sub-account. Once the account is created, the administrator will receive an email with a link to create a password for the new Lyve Cloud account. If this user exists in the master account, the administrator will receive an email with a link to the new account.
 - First Name: Specify the name of the administrator.
 - Last Name: Specify the last name of the administrator.
 - Email: Specify the email address of the administrator. After registering, the administrator will receive all notifications on the registered email.
 - Confirm Email: Re-enter the email to verify the email mentioned in the Email field.

Admin Details
This admin will be the only person who will have access to the account when it is first created. [Learn more](#)

First Name

Last name

Email

Confirm Email

4. Click **Create**.

After the sub-account is created, it is listed on the Sub-accounts page, and you can view the status of the sub-account. Initially, the status is Being Provisioned, which will be changed to Enabled or Provision failed. For more information, see [Listing Sub-accounts](#).

If the provisioning fails, you can try to create an account again by choosing the option Try again from the ellipses or contact support at support.lyvecloud@seagate.com.

Use Dismiss to remove the failed sub-account from the list.

LYVE Cloud Help John Doe

- Home
- Marketplace
- STORAGE
 - Regions
 - Buckets
 - Permissions
 - Service Accounts
- IDENTITY & ACCESS
 - Users
 - MFA
 - Federated Login
 - Notification Recipients
- YOUR ACCOUNT
 - Settings
 - Billing
 - Sub-accounts**
 - Support

Sub-accounts

Create and manage sub-accounts. [Learn more](#)

[Create Sub-accounts](#) [Download](#)

Account Name	Account ID	Organization Name	Created On	Status
lyve-lab	lyve-lab	Seagate		Provision failed

3 out of 99 sub-accounts created

Rows per page: 10 [Try Again](#) [Dismiss](#)

When the total number of sub-accounts with the status **Being provisioned**, **Enabled** and **Disabled** reaches the total sub-account limit, you cannot create a new Sub-account, and a warning message is displayed.

LYVE Cloud Help John Doe

- Home
- Marketplace
- STORAGE
 - Regions
 - Buckets
 - Permissions
 - Service Accounts
- IDENTITY & ACCESS
 - Users
 - MFA
 - Federated Login
 - Notification Recipients
- YOUR ACCOUNT
 - Settings
 - Billing
 - Sub-accounts**
 - Support

Sub-accounts

Create and manage sub-accounts. [Learn more](#)

[Create Sub-accounts](#) [Download](#)

You have reached the sub-account limit for your account. Please contact support to increase the sub-account limit.

Account Name	Account ID	Organization Name	Created On	Status
lyve-lab	lyve-lab	Seagate		Being provisioned
lyve-rock	lyve-rock	Seagate	YYYY-MM-DD HH:MM:SS	Enabled

1 out of 2 sub-accounts created

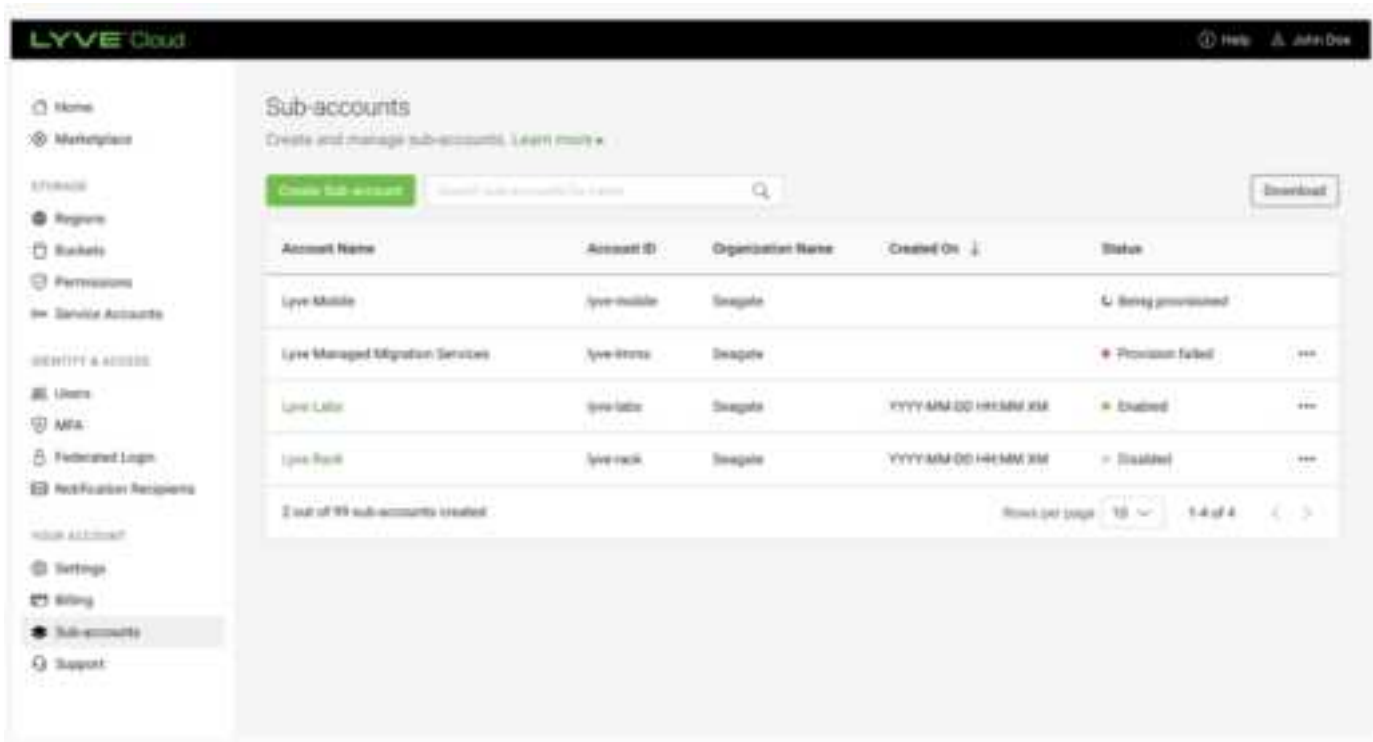
Rows per page: 10 1-2 of 2

i Note—You can create a new Sub-account even if the status is **Being provisioned** for an already created Sub-account.

Only the admin who created the sub-account receives an email notifying whether the sub-account was successfully created or failed.

Listing Sub-accounts

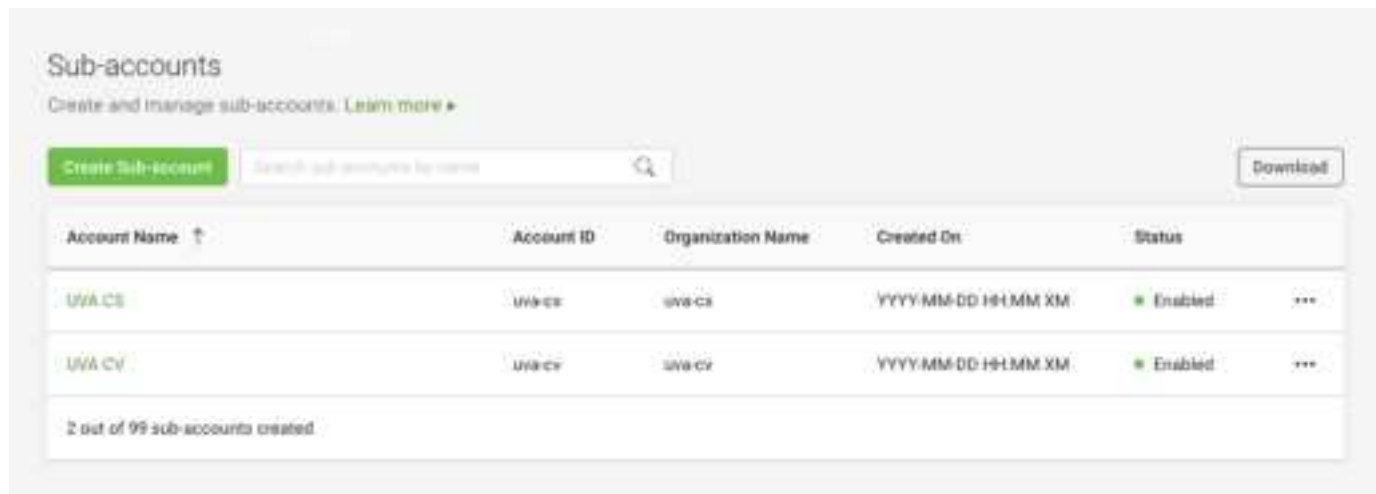
Sub-accounts are listed under the Sub-accounts management page in the master account. You can search the Sub-accounts by name.



Column Name	Description
Account Name	Displays the name of the account.
Account ID	Displays the account ID of the Sub-account. This is a unique ID across Lyve Cloud.
Organization Name	Displays the name of the organization.
Created On	Displays the date and time when the Sub-account is created.
Status	<i>Sub-account UI updates</i> Displays the status of the Sub-account. The status can be: <ul style="list-style-type: none">• Being provisioned: Specifies the sub-account creation is still in progress.• Enabled: Specifies the sub-account creation is successful.• Disabled: Specifies the sub-account is disabled.• Provision failed: Specifies the sub-account creation has failed.
Trial	Displays the remaining number of days for the trial to expire.

Exporting the Sub-account list

By selecting Download, the administrator of the master account can export the Sub-account list in the CSV file.



Disabling/enabling a Sub-account

Disabling a Sub-account restricts all account users from accessing the console and all its service accounts from accessing storage. After the account is disabled, the master account is not billed for that period.

You must contact the administrator if the option to disable an account is unavailable.

You can enable a Sub-account if it is disabled. You can enable the account with all data access and its users without restrictions. Enabling a disabled account permits all active users to access the console, and the service accounts are re-enabled.

To disable a Sub-account:

1. On the left-hand menu, select **Sub-accounts**.
2. Select the ellipsis for the account to disable, and then select **Disable**.
3. Select **Yes**, in the Disable Sub-account confirmation dialog.
4. Verify the Sub-account name to disable, and select **Confirm**.

A confirmation message is displayed on the Sub-account list page.

Alternatively, you can disable a Sub-account from the Sub-accounts dashboard from the Settings tab.

Sub-accounts >

UVA CS

uva-cs

Current Usage	Month-to-date Usage	Buckets	Created on	Created by
50TB DoD +10%	50TB MoM +10%	100	YYYY-MM-DD HH:MM XM	John Doe

Usage **Settings**

Organization Information

Organization Name:
UVA CS

Admin Details

First Name	Last Name	Email
John	Doe	john.doe@uva-cs.com

Danger Zone

Disable Sub-account

[Enable / Disable sub-account](#) [Learn more](#) ➤

[Disable](#)

Delete Sub-account

Sub-account can only be deleted at a minimum of 30 days after the account has been disabled. [Learn more](#) ➤

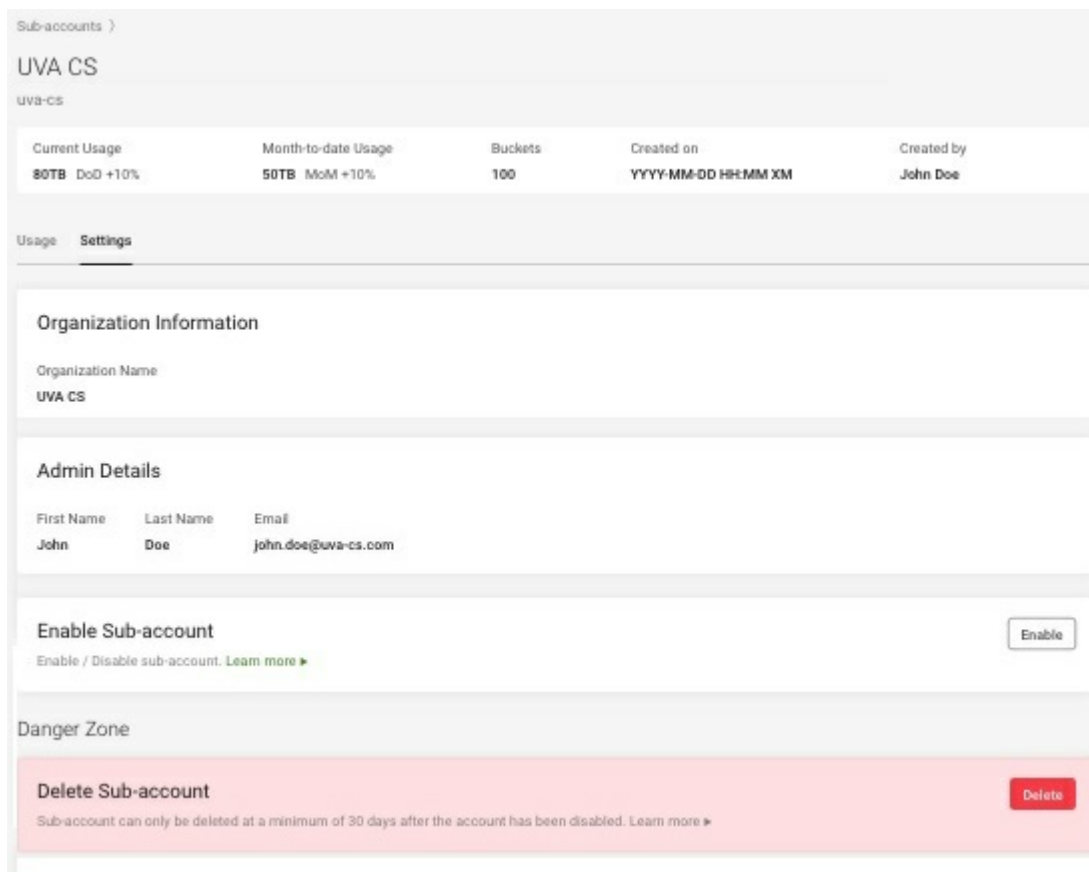
[Delete](#)

Deleting a Sub-account

You can delete a Sub-account if it is disabled for at least 30 days. You can confirm the status of the disabled account on the Sub-account detail page. Once the Sub-account is deleted, all the data, including buckets, objects, service accounts, users, etc., are also deleted and cannot be restored.

To delete a Sub-account:

1. On the left-hand menu, select **Sub-accounts**.
2. Select the Sub-account to delete.
3. Select the **Settings** tab on the Sub-account detail page and then select **Delete**.



Video: Lyve Cloud - How to use the Sub-Account feature of Lyve Cloud

Seagate on Vimeo: [Lyve Cloud - How to use the Sub-Account feature of Lyve Cloud](#)

Understanding the Sub-account usage dashboard

Sub-account usage is monitored from the Sub-account usage dashboard.

The Sub-account dashboard displays statistics of the storage system, including usage, the number of buckets, daily average usage etc. After selecting a Sub-account on the left navigation pane, you can view the dashboard. Initially, there is no data on the dashboard. Once you create Sub-accounts and the accounts begin storing data in the buckets, the dashboard displays essential details in different sections. A graphical view of daily average usage with the date range filter displays the storage usage trend.

There are two tabs on the dashboard displaying the Usage information and Settings for the Sub-account.

- **Month-to-date Usage:** Displays the average usage of the account from the beginning of the month until the current date.
- **Buckets:** Displays the total number of buckets created in the Sub-account.
- **Created on:** Displays the date and time of Sub-account creation.
- **Created by:** Displays the administrator's name who created the Sub-account.
- **Usage:** The Usage tab is selected by default.

- Reports

- Daily Average Usage: Displays the daily average from a series of four usage snapshots within 24 hours of data stored in all the buckets.
- This month is a default selection that displays the daily average usage trend for the current month to date.
 - Selecting the Last 6 months shows the usage trend of the previous six months. Each data point displays the monthly average for that month.
 - Selecting a Custom time range allows you to choose the monthly time range, and the data points display the monthly average usage.
 - Date range selection: Select a current month, last six months, or custom time range to view usage trends.

Use the Date range selection to choose the length of time of the report. Download the usage data in CSV format by selecting Download. This report shows the Date, Region Name, Bucket Name, Usage(byte), and Usage (GB) in the excel sheet.

- **Settings:** Click the tab to display the details of the Sub-account.
 - Organization Information: Displays the organization name and address of the Sub-account.
 - Admin Details: Displays the administrator details of the Sub-account.
 - Danger Zone: This section allows you to disable and delete the Sub-account. You must be very careful before you perform any of these actions. For more information, see [Disabling/enabling a Sub-account](#) and [Deleting a Sub-account](#).



The screenshot shows a web form titled "Organization Information". Below the title is a note: "An organization represents a legal entity that a sub-account belongs to. [Learn more](#)". The form contains several input fields: "Organization Name", "Street Address", "Street Address 2 (Optional)", "City", "State", "Zip Code", "Country/Region" (a dropdown menu), and "Phone Number". There is also a checkbox labeled "Enable Live Chat Support" which is currently checked.

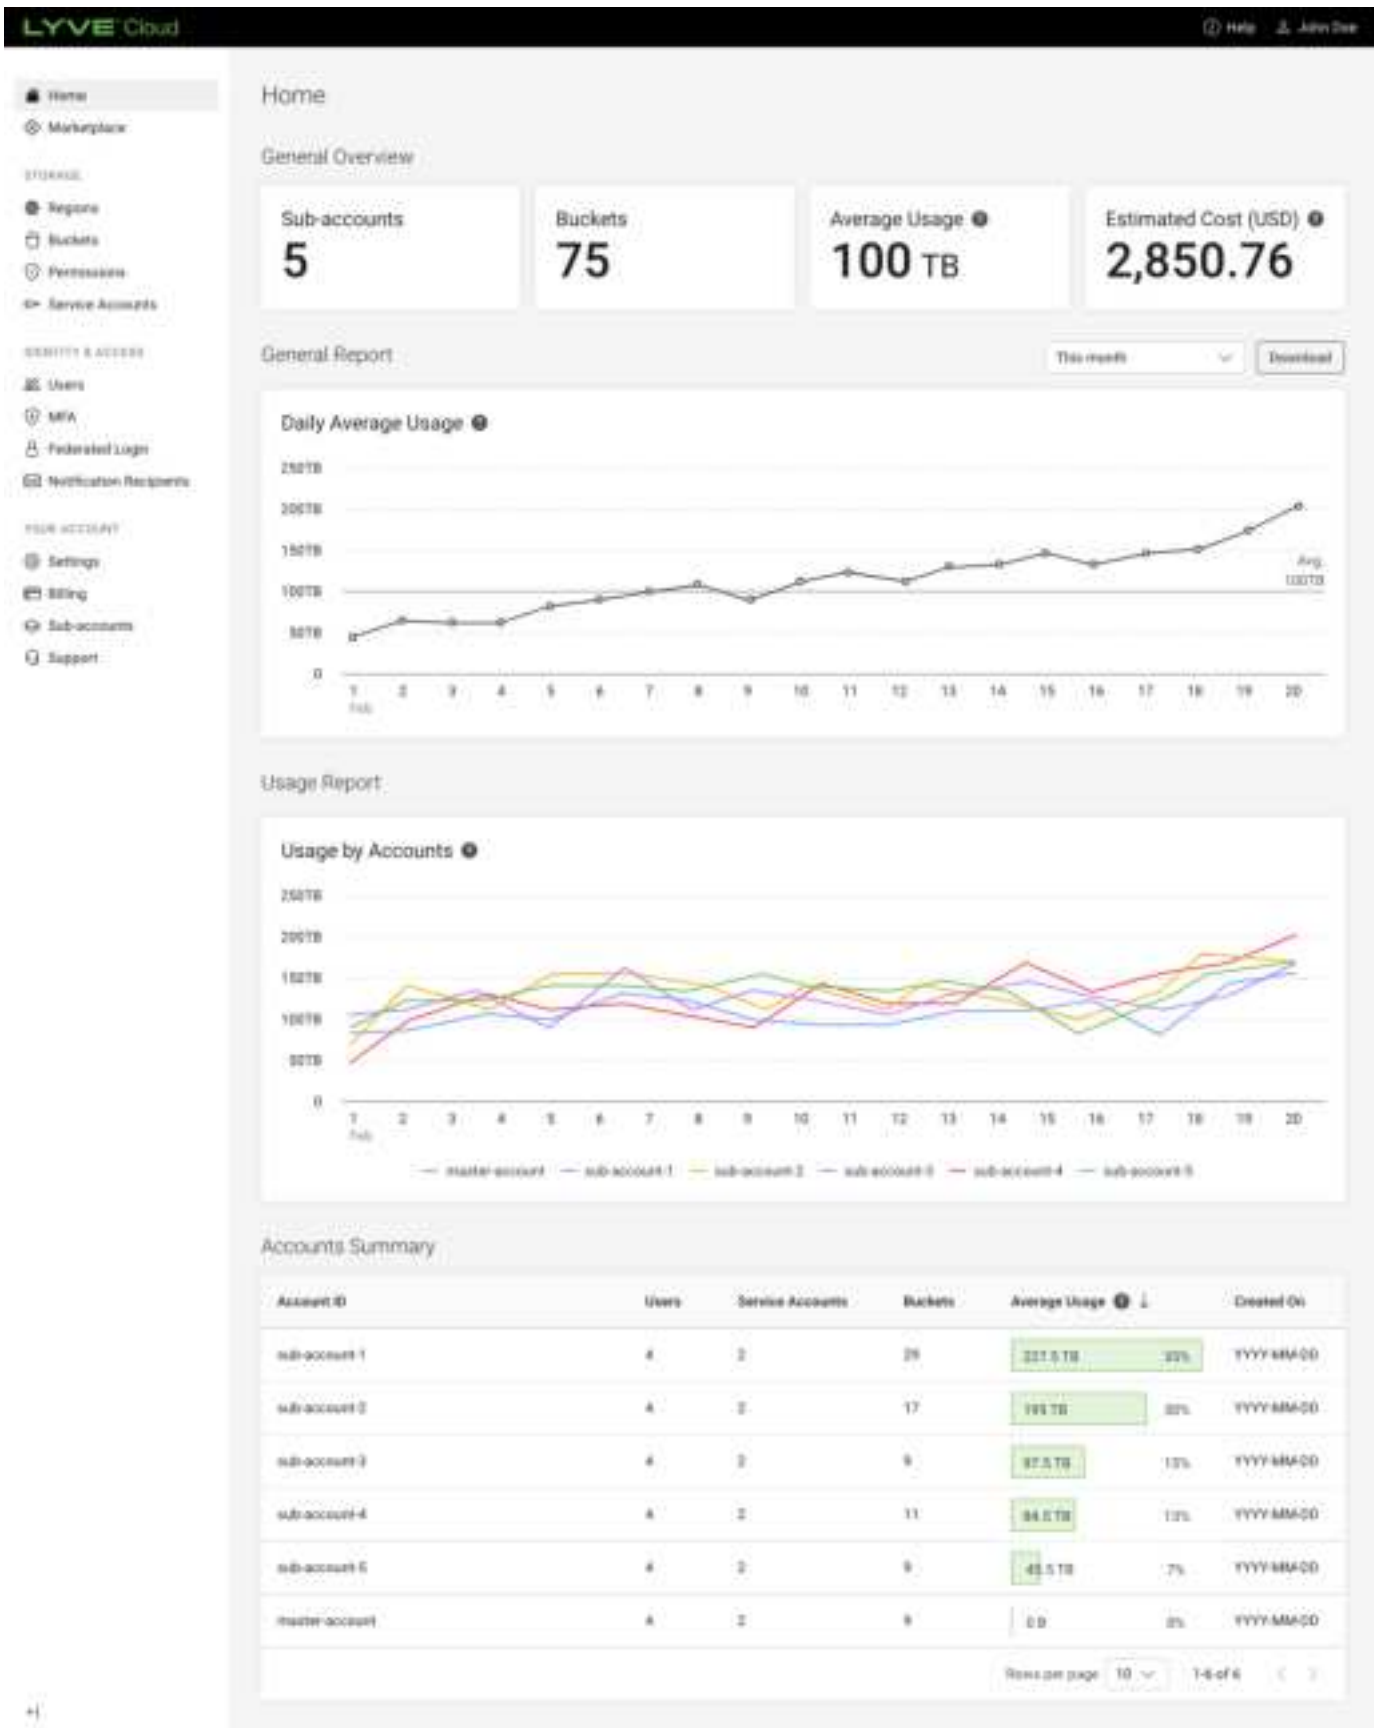
Understanding Sub-accounts billing

Master account billing

The master account administrator can view the usage total, the number of buckets, and the storage trends of the buckets in that account. The Home page displays the details of the master account that is combined with Sub-accounts. It shows the number of buckets, total usage from the beginning of the

month to date, and a graphical representation of unlimited usage.

The Sub-account usage graph displays the usage of each Sub-account on the same chart. The graph shows different colour lines for separate Sub-accounts.



Invoicing master account



Note—The Unit Price (9.00) shown in the invoice is just a placeholder.

Billing for a disabled Sub-account

The master account is not charged for the usage of the disabled Sub-account. If the account is disabled during the period of usage calculation, then the usage is calculated till the date the account is disabled.

For example, If the Sub-account is disabled on June 15, and the usage is calculated for June, the Sub-account is billed from June 1 to June 15.

If you re-enable the Sub-account, the account is billed on a pro-rata basis. For example, If the account is enabled on August 10, then the account is billed on a pro-rata basis from August 10 till the end of August.

Connecting S3 Clients

This guide explains how to configure third-party clients including AWS CLI, Cyberduck, Mountain Duck, or S3 Browser to manage data in Lyve Cloud. Lyve Cloud is an S3-compatible storage service for data-intensive applications such as data backup and analytic workloads, that leverage multi-petabyte data lakes. You can also use any other compatible third-party client to copy and move files, manage files and folders, and synchronize folders between Lyve Cloud and your local storage, once you've established a connection.

This document and its subtopics provide instructions about how to:

- Configure Cyberduck, Mountain Duck, S3 Browser, or rclone to connect to Lyve Cloud.
- Upload, download, and delete files, and how to best manage files and folder.
- Disconnect from the server.

Using Cyberduck

Use Cyberduck to connect with Lyve Cloud and transfer your data. For more information, review the [Cyberduck Tutorial](#) and [Cyberduck Quick Reference Guide](#).

Prerequisites

- [Download](#) and Install Cyberduck.
- Register the S3 (HTTPS) profile for preconfigured settings. For more information, see [Generic S3 profiles](#).

To enable and register the S3 (HTTPS) profile:

1. Select Edit and then Preferences. For Mac users, select Cyberduck and then Preferences.
2. In the Profiles tab, choose S3 (HTTPS) from the connection profiles list.

Alternatively, to register the S3 (HTTPS) profile:

1. Open [S3 \(HTTPS\) profile](#) connection profile file. Copy the file contents into a notepad/any text editor.
2. Save the notepad file name with .cyberduckprofile extension and change the Save as type to All Files.

- You need both the access key and secret key for each account you plan to connect with Cyberduck. For more information, see [Creating service accounts](#).




Note—Consult your organization's policies and the EULA policies of the software before downloading 3rd party applications.

Connect Cyberduck to Lyve Cloud

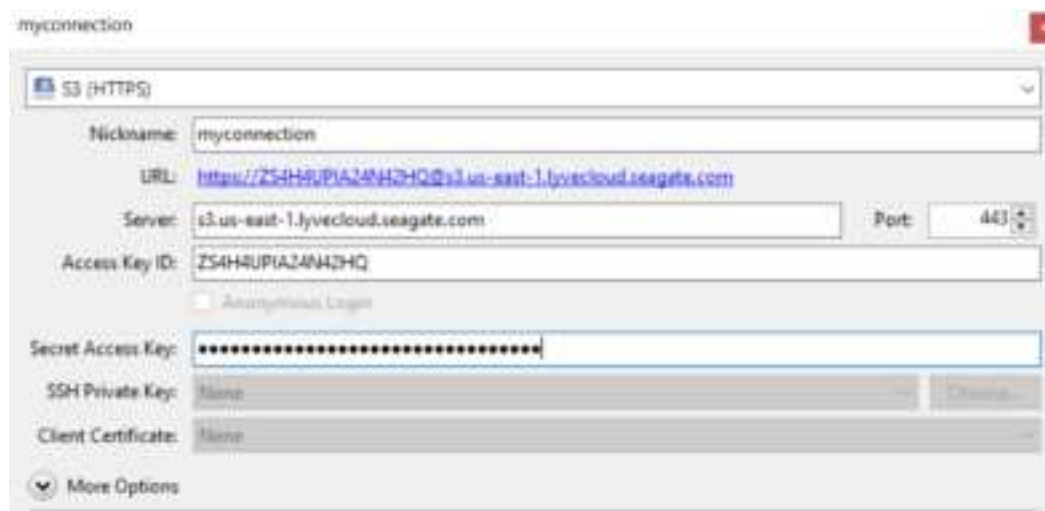
Bookmarks store the details of the connection to easily re-connect to the server.

To connect Cyberduck to Lyve Cloud:

1. In Cyberduck, select **Bookmark|New Bookmark**. Mac users: Select+ in the bottom left to add a new bookmark.



2. Select the **S3 (HTTPS)** protocol from the list.
3. Enter the following mandatory details to add your connection to Lyve Cloud:

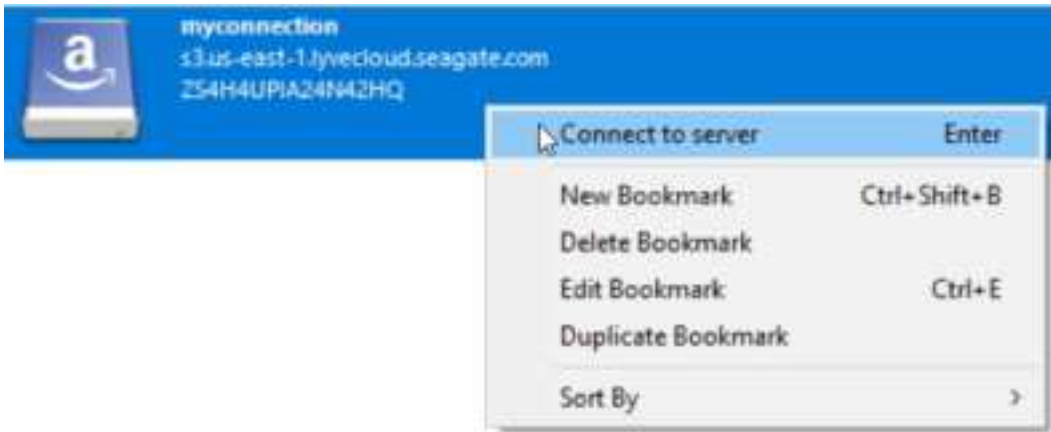


Field Name	Description
Nickname	Enter a name for the bookmark.
URL	Displays the URL once you enter the server and access key in the following format:http:// <Access Key ID><Server>
Server	Enter Lyve Cloud S3 endpoint. For more information see S3 API endpoints . Lyve Cloud supports only region-specific S3 endpoints. To access buckets created in different regions in the S3 client, add an endpoint connection for each of the regions. Tip—Copy the URL without https://

Field Name	Description
------------	-------------

Port	Enter 443 as the port number to access the server.
Access Key ID	<p>Enter your access key, a private key for authentication to connect a bucket created in Lyve Cloud.</p> <p>The access key is displayed when you create a new service account in Lyve Cloud. A service account contains bucket credentials for Lyve Cloud to access a bucket.</p> <p>For more information, see Creating service accounts.</p>
Secret Access Key	<p>Enter your secret key, a private key password for authentication to connect a bucket created in Lyve Cloud.</p> <p>The secret key displays when you create a new service account in Lyve Cloud.</p> <p>For more information, see Creating service accounts.</p>

- The bookmark is displayed once you close the window. Right-click the bookmark and select **Connect to Server**.



- Select **Continue** to establish the connection.
 - View the buckets available in the created bookmark once the connection is established.
 - On the Cyberduck client, the **Disconnect** button is displayed in the top right corner. A green dot appears to the right of active bookmarks, signifying an established connection. If no connection is established, the Disconnect icon is greyed out.

6. Right-click the bucket or select Actions to perform various operations or actions. For more information, see [Managing data](#).

Video: Lyve Cloud - How to connect Cyberduck to Lyve Cloud

Seagate on Vimeo:[Lyve Cloud - How to connect Cyberduck to Lyve Cloud](#)

Managing data


Perform various actions once a connection is established between Cyberduck and Lyve Cloud. For more information, see [Cyberduck Help](#).

Uploading data to a bucket

i Note—The object name can contain any of these special characters like @, #, *, \$, %, &, !, ?, , , ;, ', ", |, +, =, <, >, ^, (,), {, }, [,] and alphanumeric characters like 0-9, a-z, A-Z . However, using any of these characters can cause issues due to limiting factors of S3 client SDK.

To upload data to a bucket:

1. Select and right-click the bucket and select Upload and choose the file to upload.
2. After the file transfer progress status is complete, you can view the file.



3. Accept the certification installation and select Continue.



Note—Certificate installation is prompted only when the certificate is installed for the first time.

Downloading data to local storage

To download data to local storage:

1. Expand the bucket where files are available.
2. Right-click the file to download and select one of the following options:

- **Download:** Download a file to the predefined path.

The Transfers dialog displays the connection status. Select **Continue** in the Download dialog. You can view the remote file location and the local file location, but you cannot change the download path. Once the download is complete, the Transfers dialog displays the status.

- **Download As:** Download a file in the required format.

Select **Save as Type** from the list, and select **Save**.

- **Download To:** Download the file to a specific location.

Select the download folder in the Browse to Folder dialog, or create a new folder. Select **Save**.

Deleting data from the bucket

To delete data from the bucket:

1. Expand the bucket from which to delete data.
2. Right-click the data file, and select **Delete**.
3. Select **Delete** in the Confirmation prompt.

Creating a new folder

To create a new folder:

1. In Cyberduck, select **Select** and open a bucket.
2. Right-click inside the bucket and select **New Folder**.
3. Enter the folder name in the Create New Folder screen.

Deleting a folder

To delete a folder:

1. Navigate into the bucket from which to delete the folder.
2. Right click the folder and select **Delete**.
3. Select **Delete** in the confirmation prompt.

Disconnecting Cyberduck from Lyve Cloud

To disconnect Cyberduck from Lyve Cloud:

1. Open Cyberduck to view all the available bookmarks or connections.
2. Select **Disconnect** in the top-right corner on the Cyberduck client.



Note—A green dot indicates an active bookmark and an established connection. However, if the connection is not established, the Disconnect icon is greyed out.

Using S3 Browser

Use [S3 Browser](#) to connect with Lyve Cloud and manage your data transfer. For more information on S3 Browser see, [S3 Browser Help](#).

Prerequisites

- You will need the access key and secret key for each account you'll be using to connect with S3 browser. For more information, see [Creating service accounts](#).



Note—Consult your organization's policies and the EULA policies of the software before downloading 3rd party applications.

Connect S3 Browser to Lyve Cloud

To connect S3 Browser to Lyve Cloud:

1. Open S3 Browser and select **Accounts**, then select **Add New Account**.
2. Enter the following mandatory details:

Add New Account


Add New Account
[online help](#)

Enter new account details and click Add new account

Account Name:

Assign any name to your account.

Account Type:

S3 Compatible Storage

Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Access Key ID:

Required to sign the requests you send to Amazon S3. see more details at <https://js3browser.com/keys>

Secret Access Key:

Required to sign the requests you send to Amazon S3. see more details at <https://js3browser.com/keys>

☐ **Encrypt Access Keys with a password:**

Turn this option on if you want to protect your Access Keys with a master password.

☒ **Use secure transfer (SSL/TLS)**

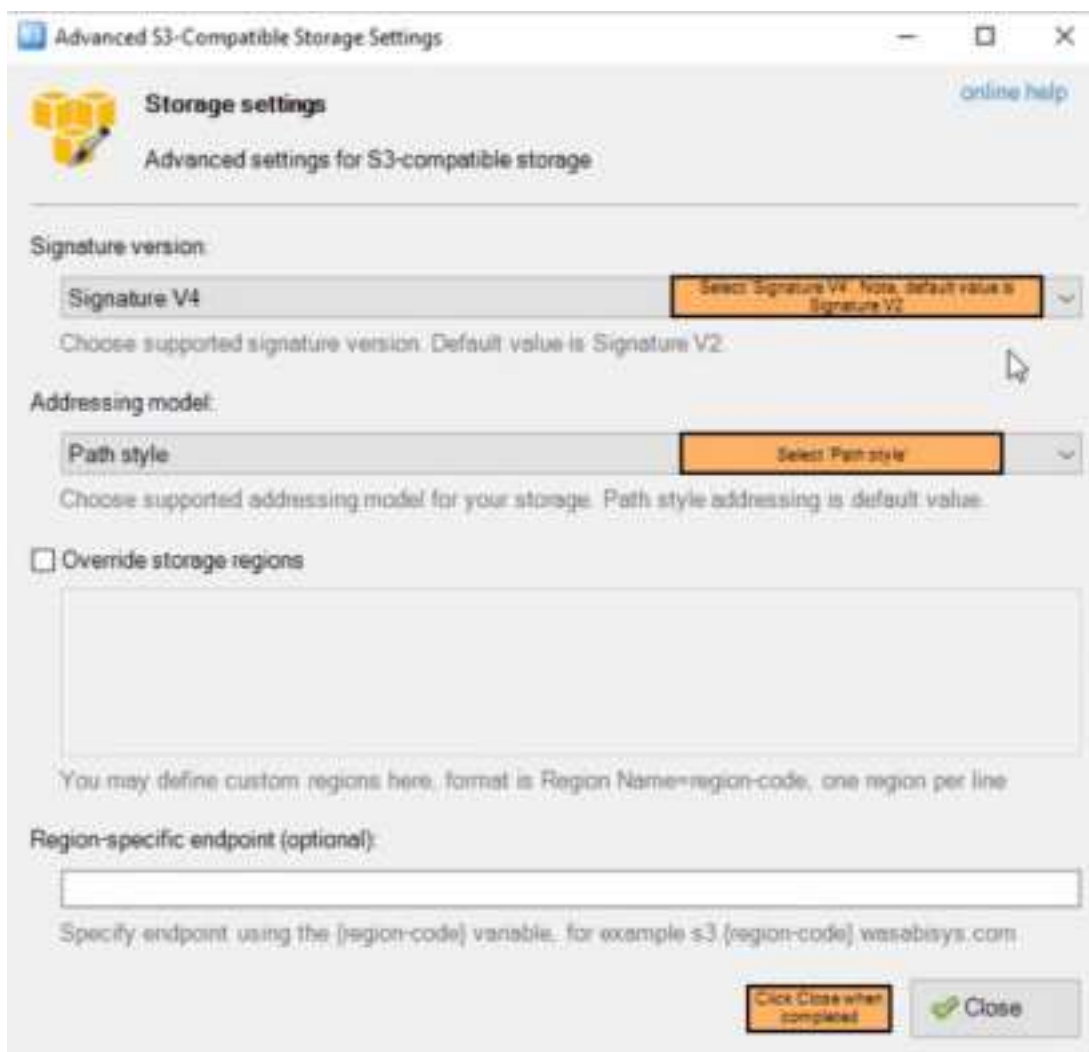
If checked, all communications with the storage will go through encrypted SSL/TLS channel.

[Advanced S3-compatible storage settings](#)
[Select Advanced S3-compatible storage settings and proceed with further instructions](#)
[Cancel](#)

Field Name	Description
Account name	Enter an account name.
Account type	Enter the account type. Select S3 Compatible Storage from the list.
REST Endpoint	Enter Lyve Cloud S3 endpoint. For more information see S3 API endpoints . Currently, Lyve Cloud supports only region-specific S3 endpoints. To access buckets created in different regions in the S3 client, add an endpoint connection for each of the regions.
Access Key	Enter your access key. The access key displays when you create a new service account in Lyve Cloud. A service account contains the bucket credentials for the Lyve Cloud bucket. For more information, see Creating service accounts .
Secret Key	Enter your secret key. The secret key displays when you create a new service account in Lyve Cloud. For more information, see Creating service accounts .

Field Name	Description
Use secure transfer (SSL/TLS)	Select this option to ensure all communication with the storage passes through encrypted SSL/TLS.
Advanced S3 Compatible storage settings	Select the signature version and addressing model in the advanced settings. For more information, see Advanced S3 compatible storage settings below.

Advanced S3 compatible storage settings



Advanced S3-Compatible Storage Settings

Storage settings
Advanced settings for S3-compatible storage

[online help](#)

Signature version:
Signature V4
Select Signature V4. Note, default value is Signature V2
Choose supported signature version. Default value is Signature V2.

Addressing model:
Path style
Select Path style
Choose supported addressing model for your storage. Path style addressing is default value.

☐ Override storage regions

You may define custom regions here, format is Region Name=region-code, one region per line

Region-specific endpoint (optional):

Specify endpoint using the {region-code} variable, for example s3-{region-code}.wasabisys.com

Click Close when completed Close

Field Name	Description
Signature Version	Select Signature V4 . For more information, see S3 Browser's help documentation .
Addressing Model	Path Style is selected by default and is the recommended setting. For more information, see S3 Browser's help documentation .

Video: Lyve Cloud - How to use S3 Browser with Lyve Cloud

Seagate on Vimeo: [Lyve Cloud - How to use S3 Browser with Lyve Cloud](#)

Managing data

Perform various actions once the connection between S3 Browser and Lyve Cloud is established. For more information, see [S3 Browser's help documentation](#).

Uploading data to a bucket

i Note—The object name can contain any of these special characters like @, #, *, \$, %, &, !, ?, , , ;, ', ", |, +, =, <, >, ^, (,), {, }, [,] and alphanumeric characters like 0-9, a-z, A-Z . However, using any of these characters can cause issues due to limiting factors of S3 client SDK.

To upload data to a bucket:


1. Select the bucket, and then select **Upload**.
2. Select **Upload file(s)** or **Upload folder(s)**.
3. Select the file, and then select **Open**.

Downloading data to local storage

To download data to your local machine:

1. Select the bucket where the data file is available.
2. Select the folder or file(s) to download, and then select **Download**.

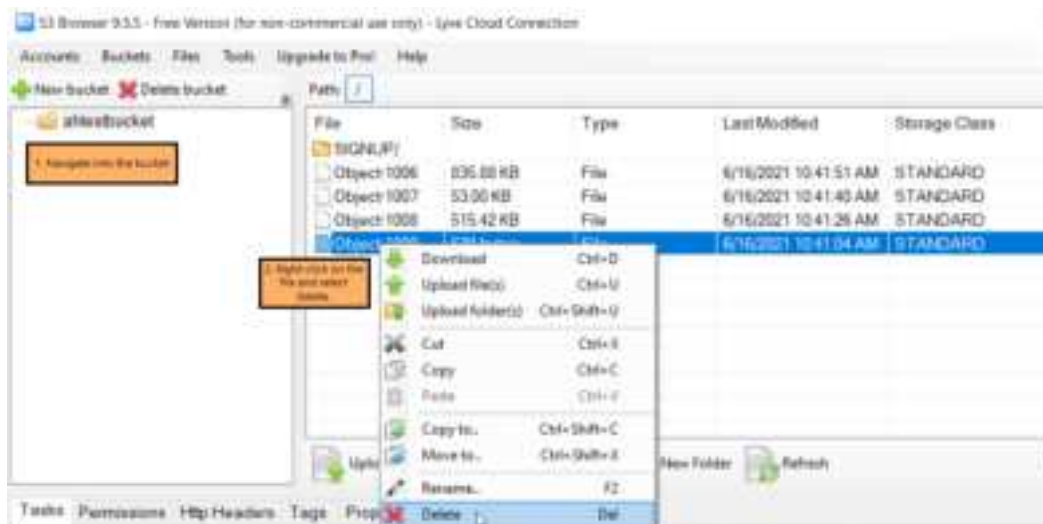
The Tasks tab displays the upload or download progress.




Deleting data from a bucket

To delete data from a bucket:

1. Navigate into the bucket, select the file from the right pane, and select Delete.



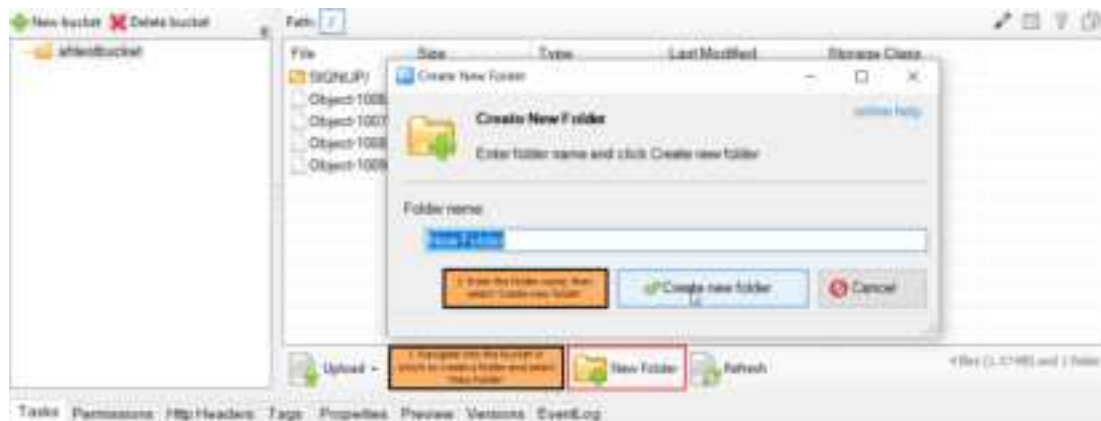
2. Select Yes in the Confirm File Delete dialog.



Creating a new folder

To create a new folder:

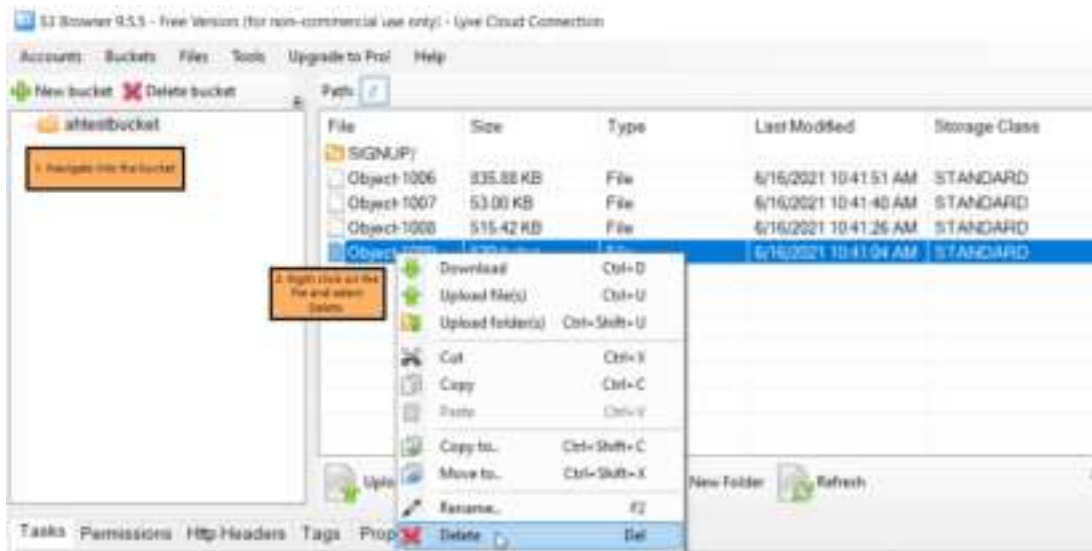
1. Navigate into the bucket in which to create a folder.
2. Select **New Folder**. Enter a folder name, and then select **Create New folder**.




Deleting a folder

To delete a folder:

1. Open the bucket, right-click the folder to delete, then select **Delete**.




2. Select **Yes** in the Confirm File Delete dialog.



Disconnecting S3 Browser from Lyve Cloud

To disconnect S3 browser:

1. Open the S3 Browser to view all available connections.
2. Select **Accounts** from the menu.
3. Select **Manage Accounts**. Select the account name, and then select **Delete**.



4. Select **Save Changes**.



Using Mountain Duck

Use Mountain Duck to mount your Lyve Cloud storage as a disk in the Windows File Explorer or Mac OS Finder, and manage your files through a familiar interface. For more information on Mountain Duck see, [Mountain Duck Help](#)

Prerequisites

- You will need the access key and secret key for each account you'll be using to connect with S3 browser. For more information, see [Creating service accounts](#).

i **Note**—Consult your organization's policies and the EULA policies of the software before downloading 3rd party applications.

i **Note**—Differences between Windows and macOS are mentioned specifically. Otherwise, instructions are similar for both operating systems.

Connect Mountain Duck to Lyve Cloud

To connect Mountain Duck to Lyve Cloud:

1. Download [Mountain Duck's S3 \(HTTPS\) profile](#) for preconfigured settings. For more information, see their [Generic S3 profiles](#) documentation.
2. Open the downloaded file with Mountain Duck. The New Connection dialog appears.

i **Note**—The dialog for macOS does not have a Drive Letter field.

The screenshot shows the 'New Connection' dialog in Mountain Duck for an S3 (HTTPS) profile. The title bar says 'lc-demo'. The profile type is 'S3 (HTTPS)'. The fields are as follows:

- Nickname: lc-demo
- Labels: (empty)
- URL: https://s3.us-east-1.lyvecloud.seagate.com
- Server: s3.us-east-1.lyvecloud.seagate.com
- Port: 443
- Access Key ID: AEJKASWFDRGHVDFG
- Secret Access Key: (masked with dots)
- SSH Private Key: none
- Client Certificate: none
- Path: (empty)
- Encoding: Default
- Timezone: UTC
- Connect Mode: Smart Synchronization
- Options: Read only (checked)

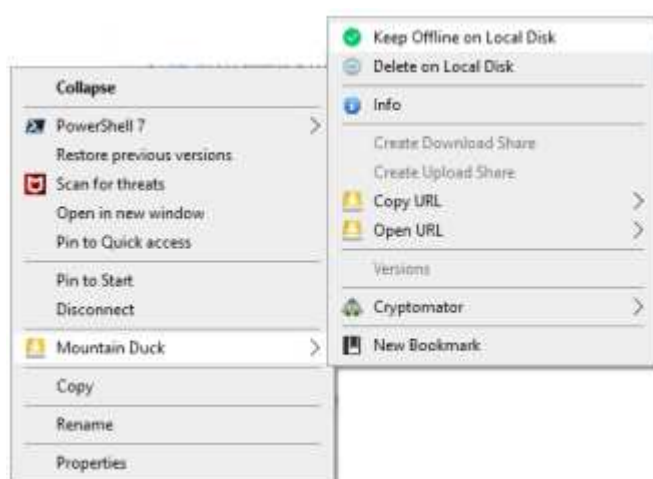
Buttons at the bottom: Delete, Connect.

3. Enter the following information:

Field Name	Description
------------	-------------








Field Name	Description
Nickname	Enter a unique name. This will be the name of your connection bookmark.
Server	Enter the endpoint based on the region. To select endpoint see the section called "S3 API endpoints" . Currently, Lyve Cloud supports only region-specific S3 endpoints. To access buckets created in different regions in the S3 client, add an endpoint connection for each of the regions.
Port	This should populate as 443. If not, enter that port number.
Access Key ID	Enter your access key ID. For more information, see the section called "Creating service accounts" .
Secret Access Key	Enter your secret key. For more information, see the section called "Creating service accounts" .
Drive Letter(Windows only)	Enter a drive letter so that Mountain Duck always uses that same letter for the mounted drive.

1. Once the fields have been filled, select **OK**. This creates your connection bookmark.
2. Select the Mountain Duck icon in the Windows system tray or the macOS' menu bar.
3. Select your connection bookmark, and then select **Connect**. When the connection is created, a notification appears.
4. Select the new drive and right-click or Ctrl+click in the File Explorer or Finder to bring up the context menu. Right-click in the folder to get the context menu for macOS.
5. Select **Mountain Duck** and select **Keep offline on local disk** to sync all the data to local drive.



When connected, the drive and folder contents display their sync status. Look for a circle in the lower-left corner of the folder or file icon. Once mounted, all the files are stored on your local drive.

icon	Meaning
------	---------

icon	Meaning
	In Progress. Synchronization is in progress for this item.
	In sync. This item is selected to be synced, and the content will always be available offline.
	Sync error. This item cannot be synchronized.
	Up to date. This item is synced and up to date.
	Ignored. The file is available in its temporary location and never synced to cloud or remote storage.
	Paused. The sync on that item is paused.
	Online only. This item is available in the cloud but can be opened and edited when you have an active connection to the server.

To learn more about various Mountain Duck options and sync modes, see Mountain Duck's [Help documentation](#).

Video: Lyve Cloud - How to mount Lyve Cloud as a local drive on Windows with Mountain Duck

Seagate on Vimeo: [Lyve Cloud - How to mount Lyve Cloud as a local drive on Windows with Mountain Duck](#)

Managing data

Once your Lyve Cloud storage is mounted as a drive, managing your files works much the same as working in any other network drive. Many of these operations may only be performed once a given bucket has synced with the local drive. Learn more about Mountain Duck's [user interface documentation](#).

i **Note**—The object name can contain any of these special characters like @, #, *, \$, %, &, !, ?, , , ;, ' , " , | , + , = , < , > , ^ , (,) , { , } , [,] and alphanumeric characters like 0-9, a-z, A-Z . However, using any of these characters can cause issues due to limiting factors of S3 client SDK.

Copying data

To copy data:

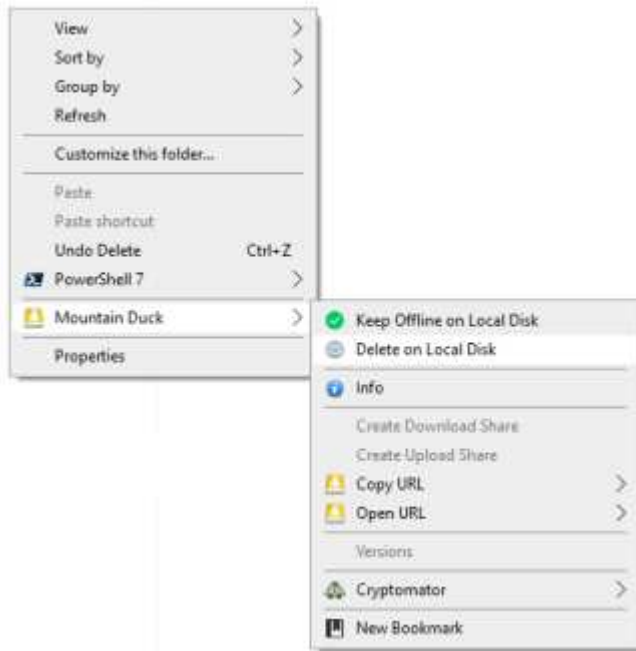
1. Select the data from your source drive and copy it.
2. Navigate to your destination and paste your data.

You can also drag and drop the data from one folder to another. If the bucket's or your service account's permissions do not allow you to write to that bucket, an error message appears.

Deleting bucket data

To delete bucket data:

1. Navigate to the synced bucket from which you want to delete the data.
2. Right-click or Ctrl+click the file or folder inside the bucket to bring up the local context menu.
3. Select **Mountain Duck**, and then select **Delete on Local Disk**.



4. Right-click or Ctrl+click the file and select **Delete**, and then select **Yes** in the confirmation box to delete the object permanently from the bucket.

Creating a folder in a bucket

To create a folder in a bucket:

1. Navigate into the bucket where you want to create a new folder.
2. Right-click or Ctrl+click the bucket and select **New Folder**.
3. Type a new folder name, and then select **Enter**.

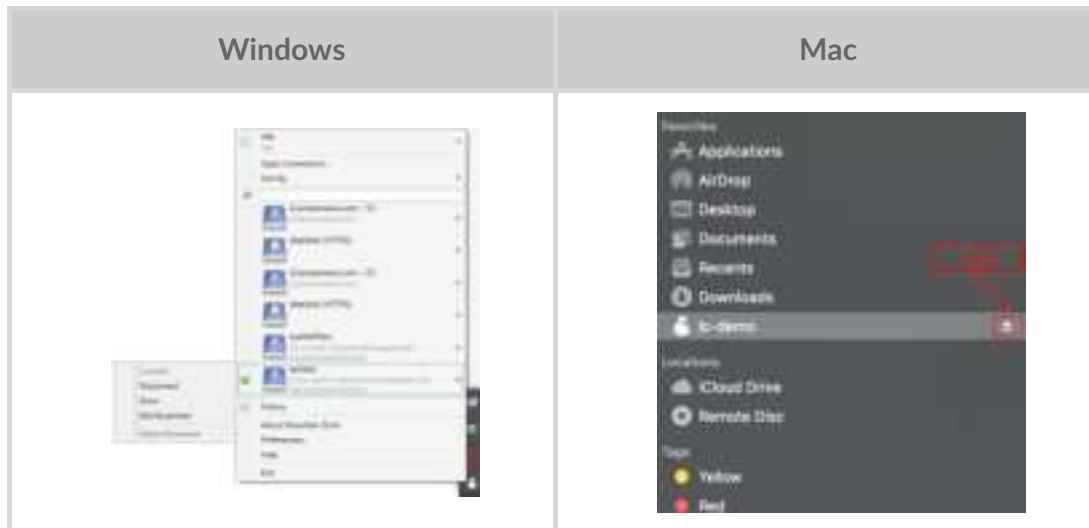
The new folder immediately begins syncing with its Lyve Cloud destination.

Disconnecting Mountain Duck from Lyve Cloud

To disconnect Mountain Duck from Lyve Cloud:

1. In the Windows system tray or the macOS menu bar, select the Mountain Duck client icon. The menu lists all of your connection bookmarks.
2. **Windows:** Select a connection bookmark, and select **Disconnect**. **macOS:** Select the Eject icon next to the connection bookmark, or click a connection and select **Eject**.

A notification pop-up appears when the connection is broken.



Using Rclone

Use rclone to connect with Lyve Cloud and manage your files from the command line, or mount the cloud storage as a drive.

Prerequisites

You will need the access key and secret key for each account you'll be using to connect with Rclone. For more information, see [Creating service accounts](#).



Note—Consult your organization's policies and the EULA policies of the software before downloading 3rd-party applications.

Connecting to Lyve Cloud from Linux

Installing Rclone

1. [Download Rclone](#) for Linux, then extract the rclone binary to your desired location.
2. To use Rclone, open a terminal window and navigate to the directory where you saved the executable.

Configuring rRclone to connect to Lyve Cloud

To configure a remote connection with Rclone:

1. Run **rclone config** to setup and select **n** for a new remote.

No remotes found - make a new one
n) New remote

```
s) Set configuration password
q) Quit config
n/s/q> n
```

2. Enter a name for the configuration.

```
name> <Name>
```

3. Select **s3** storage.

```
Type the storage to configure.
Choose a number from below, or type in your own value
1 /Fichier
  \ (fichier)
2 /Akamai NetStorage
  \ (netstorage)
3 / Alias for an existing remote
  \ (alias)
4 / Amazon Drive
  \ (amazon cloud drive)
5 / Amazon S3 Compliant Storage Providers including AWS, Alibaba, Ceph, Digital Ocean, Dreamhost, IBM COS, Lyve Cloud, Minio, RackCorp, SeaweedFS, and Tencent COS
  \ (s3)
6 / Backblaze B2
  \ (b2) [snip]
46 / seafilehttp Connection
  \ (seafile)
Storage> 5
```

4. Choose **Lyve Cloud** as the storage provider.

```
Choose the S3 provider.
Choose a number from below, or type in your own value
Press Enter for the default (")
1 / Amazon Web Services (AWS) S3
  \ (AWS)
2 / Alibaba Cloud Object Storage System (OSS) formerly Aliyun
  \ (Alibaba)
3 / Ceph Object Storage
  \ (Ceph)
4 / Digital Ocean Spaces
  \ (DigitalOcean)
5 / Dreamhost DreamObjects
  \ (Dreamhost)
6 / IBM COS S3
  \ (IBMCOS)
7 / Seagate Lyve Cloud
  \ (LyveCloud)
8 / Minio Object Storage
```

```

\ (Minio)
9 / Netease Object Storage (NOS)
  \ (Netease)
10 / RackCorp Object Storage
  \ (RackCorp)
11 / Scaleway Object Storage
  \ (Scaleway)
12 / SeaweedFS S3
  \ (SeaweedFS)
13 / StackPath Object Storage
  \ (StackPath)
14 / Storj (S3 Compatible Gateway)
  \ (Storj)
15 / Tencent Cloud Object Storage (COS)
  \ (TencentCOS)
16 / Wasabi Object Storage
  \ (Wasabi)
17 / Any other S3 compatible provider
  \ (Other)
Provider>7

```

5. Enter **false** to enter your credentials.

Get AWS credentials from the runtime (environment variables or EC2/ECS meta data if no env vars).
 Only applies if access_key_id and secret_access_key is blank. Enter a boolean value (true or false).
 Please Enter for the default ("false"). Choose a number from below, or type in your own value

```

1 / Enter AWS credentials in the next step
  \ "false"
2 / Get AWS credentials from the environment (env vars or IAM)
  \ "true"
env_auth>false

```

6. Enter your **access key** and **secret key**.

AWS Access Key ID.
 Leave blank for anonymous access or runtime credentials.
 Enter a string value. Press Enter for the default ("")

```

access_key_id> <access key>

```

AWS Secret Access Key (password)
 Leave blank for anonymous access or runtime credentials.
 Enter a string value. Press Enter for the default ("")

```

secret_access_key> <secret key>

```

7. Leave the **region** blank.

Region to connect to.
 Leave blank if you are using an S3 clone and you don't have a region.
 Enter a string value. Press Enter for the default ("")
 Choose a number from below, or type in your own value

```
1 / Use this if unsure.
| Will use v4 signatures and an empty region.
\ ()
2 / Use this only if v4 signatures don't work.
| E.g. pre Jewel/v10 CEPH.
\ "other-v2-signature"
region> <>
```

8. Specify the endpoint for Lyve Cloud. For more information about endpoints, see [S3 API endpoints](#).

```
Endpoint for S3 API.
Required when using an S3 clone.
Choose a number from below, or type in your own value.
Press Enter to leave empty.
1 / Seagate Lyve Cloud US East 1 (Virginia)
\ (s3.us-east-1.lyvecloud.seagate.com)
2 / Seagate Lyve Cloud US West 1 (California)
\ (s3.us-west-1.lyvecloud.seagate.com)
3 / Seagate Lyve Cloud AP Southeast 1 (Singapore)
\ (s3.ap-southeast-1.lyvecloud.seagate.com)
endpoint> 1
```

9. Press **Enter** to skip the location constraint as there is no location constraint

```
Location constraint - must be set to match the Region.
Leave blank if not sure. Used when creating buckets only.
Enter a string value.
Press Enter for the default ("")location constraint>
```

10. Choose default ACL (private).

```
Canned ACL used when creating and or storing or copying objects.
This ACL is used for creating objects and if bucket_acl isn't set, for creating buckets too.
Note that this ACL is applied when server-side copying objects as S3
It doesn't copy the ACL from the source but rather writes a fresh one.
Enter a string value. Press Enter for the default ("")
Choose a number from below, or type in your own value
1 / Owner gets FULL_CONTROL.
| No one else has access rights (default).
\ (private)
2 / Owner gets FULL_CONTROL.
| The ALLUsers group gets READ access.
\ (public-read)
3 / Owner gets FULL_CONTROL.
| The ALLUsers group gets READ and WRITE access.
[snip]
acl>1
```

11. Select **n** to save the default advanced configuration.

```
Edit advanced config? (y/n)
y) Yes
n) No (default)
y/n>n
```

12. Review the displayed configuration and accept to save the **remote** and then quit. The config file should look like this:

```
NAME]
type = s3
Provider = LyveCloud
env_auth = false
access_key_id = xxx
secret_access_key = yyy
region = us-west-1
endpoint = s3.us-east-1.lyvecloud.seagate.com
acl = private
```

13. Click **y** to confirm the configuration.

```
y) Yes this is OK (default)
e) Edit this remote
d) Delete this remote
y/e/d>y
```

14. Type **q** to quit the configuration, else select any of the following to edit, delete, rename, copy, Set configuration password.

```
Current remotes:

Name      Type
====      ====
ashrcl    s3

e) Edit existing remote
n) New remote
d) Delete remote
r) Rename remote
c) Copy remote
s) Set configuration password
q) Quit config
e/n/d/r/c/s/q>q
```

Video: Lyve Cloud - How to configure Rclone

Seagate on Vimeo:[Lyve Cloud - How to configure Rclone](#)

Mounting Lyve Cloud as a drive

Prior to mounting Lyve Cloud as a drive, connect and test the connection by viewing the contents of one or more buckets using the `rclone ls` command.

To mount Lyve Cloud as a drive, use this command where:

- **remote** is the name of the remote server
- **path/to/files** is the exact path to the bucket(s)
- **path/to/local/mount** is the local directory:

```
rclone mount remote:path/to/files /path/to/local/mount
```

For more information, see [Rclone's mount command documentation](#).

Managing data

Here are several of the more commonly-needed commands for viewing and managing your data from the command line. See the [Rclone docs](#), including information on [global flags](#), for additional information.

Viewing information about your buckets and directories

There are three list commands with easily readable output available: `ls`, `lsd`, and `lsl`.

1. To list all data in a certain bucket, where **remote** is the name of the remote and **path** is the name of the bucket:

```
rclone ls remote:path [flags]
```

2. To list the directories in a certain remote and see the **total directory size, modification time, and number of objects** in the directories:

```
rclone lsd remote:path [flags]
```

Or

```
rclone lsd remote: [flags]
```

3. To list all objects in a certain remote and see **modification time, size and path** where path is the remote path beginning with the bucket name. Any of the filtering options can be applied to this command.

```
rclone lsl remote:path [flags]
```

4. Learn more about [ls](#), [lsd](#), and [lsl](#).

Video: Lyve Cloud - How to use Rclone list commands

Seagate on Vimeo: [Lyve Cloud - How to use Rclone list commands](#)

Uploading data to a bucket

i Note—The object name can contain any of these special characters like @, #, *, \$, %, &, !, ?, , , ; , ' , " , | , + , = , < , > , ^ , (,) , { , } , [,] and alphanumeric characters like 0-9, a-z, A-Z . However, using any of these characters can cause issues due to limiting factors of S3 client SDK.

1. To upload data into a bucket, use either of these commands:

```
rclone copy C:/path/to/filename remote:path [flags]
rclone copy filename remote:path
```

2. You can also use *rclone copy* to copy a file or directory to a new location and rename the directory at the same time. Neither of these commands deletes the file from the source, and neither of these commands will copy unchanged files.
3. Learn more about [copy](#) and [copyto](#).

Video: Lyve Cloud - How to use Rclone Copy-to and Copy-sync Commands

Seagate on Vimeo: [Lyve Cloud - How to use Rclone Copy-to and Copy-sync Commands](#)

Downloading data to local storage

This is the same as copying, but the source path is something in the remote or in a bucket, while the destination path is on your local storage.

```
rclone copy remote:path C:/path/to/filename [flags]
```

Delete bucket data

- To delete files in a certain path from a certain bucket:

```
rclone delete remote:path [flags]
```

- You can use flags to delete only files with certain characteristics:

For example, to delete files that are over 100MB:

```
rclone --min-size 100MB delete remote:path
```

To delete only a specific file:

```
rclone deletefile remote:path [flags]
```

Learn more about rclone [delete](#) and [deletefile](#).

Creating a new folder

To create a new folder:

- For example, to create a file named blue in the current location:

```
rclone mkdir blue
```

Create folders in other paths, or with other permissions, by setting the proper [flags](#). Learn more about [mkdir](#).

Deleting a folder

- To delete an empty folder.

For example, to delete a file named blue in the current location:

```
rclone rmdir blue
```

- Add flags to delete folders in locations other than the current directory.

For example, to delete an empty folder named blue that contains other empty folders:


```
rcclone rmdirs blue
```

The folders must be empty for `rmdir` or `rmdirs` to work.

Learn more about `rmdir` and `rmdirs`.

Copying data

1. First, connect rclone to Lyve Cloud. For step by step instructions, see [Configuring rclone to connect to Lyve Cloud](#). Once the configuration is complete, the `rclone.config` file must be updated as:

```
[REMOTE NAME]type = s3provider = Otherenv_auth = falseaccess_key_id = XXXXXXXXXXXsecret_access_key = YYY  
YYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYendpoint = https://s3.us-east-1.lyvecloud.seagate.comacl = privateregion = us-east  
-1
```

2. Start copying data from the existing cloud provider buckets to the buckets created in Lyve Cloud. To copy data:

```
$ rclone copy SOURCE REMOTE:[SOURCE BUCKET] <TARGET REMOTE>[TARGET BUCKET>]<PREFIX>
```

- A. To copy all the data, including prefixes, from the source bucket to the target bucket:

```
$ rclone copy SRT:[ SB ] TRT:[TB]
```

- B. To copy all objects with a prefix to the target:

```
$ rclone copy SRT:[ SB ]/mypath1 TRT:[TB]/mypath1
```

Deleting a remote from Lyve Cloud

To delete a remote from Lyve Cloud, delete the remote's name using this command, changing `REMOTE_NAME` to the name of the remote to disconnect:

```
rcclone config delete REMOTE_NAME [flags]
```

Learn more about `config delete`.



Note—Rclone will also disconnect whenever you shut down your computer.

Migrating Data

To migrate data:

1. Set the Source and the Target as remote. For more information, see [the section called “Configuring rclone to connect to Lyve Cloud”](#).
2. Sync the Source and Target remote using rclone sync command.

```
rclone sync <source remote name>:path <target remote name>:path
```

3. Once the source and the target are synced, all the data from the source is copied, removed or migrated to the target remotely.
4. You can use the following flags in the command to check the status of the sync/copy/migration.
 - `--progress` Displays the real-time transfer progress.
 - `--interactive`: Enables interactive mode and displays interactive for every action taken.

For more information on RClone, see <https://rclone.org/s3/>.

Video: Lyve Cloud - Use Rclone Delete and Purge commands

Seagate on Vimeo: [Lyve Cloud - Use Rclone Delete and Purge commands](#)

Frequently Asked Questions

General

What is Lyve Cloud?

How do I sign up for Lyve Cloud?

What are the S3 service URL's for Lyve Cloud's regions?

Is Lyve Cloud S3 compatible?

What kind of workloads or applications are ideal for Lyve Cloud services?

Can I store my data in Lyve Cloud and leverage resources and applications from other cloud providers?

How will I be notified of Lyve Cloud service updates?

Can I change my registered email address?

Can I reset my password?

How does Lyve Cloud assure data is consistent and all applications have a single source of truth?

How does Lyve Cloud manage data integrity and protect against data corruption?

Does Lyve Cloud guarantee permanent deletion of data?

Does Lyve Cloud support multipart uploads?

What level of uptime does Lyve Cloud support?

Are there a minimum and maximum sizes for stored objects?

What kind of data can I store in Lyve Cloud?

What versions of TLS / SSL does Lyve Cloud support?

Does Lyve Cloud support WORM/Write-Once-Read-Many for data immutability?

Does Lyve Cloud support management and data access auditing?

Can I connect Lyve Cloud with third-party tools?

Which operating systems environments are supported by the Lyve Cloud console?

Does Lyve Cloud support strong consistency?

Q:	What is Lyve Cloud?
A:	Lyve Cloud is a simple, trusted, and efficient on-demand solution for mass capacity storage. Predictable economics, verifiable trust, and ease of use make Lyve Cloud the right choice for storing your data.
Q:	How do I sign up for Lyve Cloud?
A:	To sign up for Lyve Cloud services, contact our sales team at sales.lyvecloud@seagate.com .or <i>Self-subscribe Lyve Cloud services using your credit card</i>
Q:	What are the S3 service URL's for Lyve Cloud's regions?

A:	<p>The S3 service URL's for Lyve Cloud:</p> <ul style="list-style-type: none"> • Eastern region (Northern Virginia) is https://s3.us-east-1.lyvecloud.seagate.com • Western region (California) is https://s3.us-west-1.lyvecloud.seagate.com • Asia Pacific region (Singapore) is https://s3.ap-southeast-1.lyvecloud.seagate.com • Central region (Oklahoma) is https://s3.us-central-1.lyvecloud.seagate.com • Central region (Texas) is https://s3.us-central-2.lyvecloud.seagate.com • Europe region (London) is https://s3.eu-west-1.lyvecloud.seagate.com <p>The management console URL is <a href="https://<account-id>.console.lyvecloud.seagate.com">https://<account-id>.console.lyvecloud.seagate.com </p>
Q:	Is Lyve Cloud S3 compatible?
A:	Yes, Lyve cloud is S3-compatible object storage that uses S3 API, and it works with other S3-compatible applications.
Q:	What kind of workloads or applications are ideal for Lyve Cloud services?
A:	Lyve Cloud is designed to handle all of your data storage needs for workloads such as backup, disaster recovery and big data analytics.
Q:	Can I store my data in Lyve Cloud and leverage resources and applications from other cloud providers?
A:	Yes. Lyve Cloud is co-located with Azure, Amazon, and other cloud providers to facilitate high-speed data transfer. Lyve Cloud lets you keep control of your data while leveraging computing and application resources from other cloud providers.
Q:	How will I be notified of Lyve Cloud service updates?
A:	The Lyve Cloud support team sends notifications regarding any system updates. Updates are sent to the email addresses of all registered Lyve Cloud users.

Q:	Can I change my registered email address?
A:	Currently, you cannot change the registered email address from within your Lyve Cloud account. For more information, contact the Lyve Cloud support team at support.lyvecloud@seagate.com .
Q:	Can I reset my password?
A:	Yes, you can reset your password, select the Forgot Password link on the login screen and follow the instruction. You can also contact your Lyve Cloud admin to reset the password.
Q:	How does Lyve Cloud assure data is consistent and all applications have a single source of truth?
A:	Lyve Cloud uses a robust data consistency model for writing an object to the disk. Therefore, the data read by any application is identical, providing a single source of truth.
Q:	How does Lyve Cloud manage data integrity and protect against data corruption?
A:	Lyve Cloud maintains data integrity as a core mission. At the physical layer, Lyve Cloud addresses data integrity by using Seagate's enterprise-class hard disks and cloud-scale data durability and self-healing technologies. The erasure coding algorithm ensures 11 x 9's of data durability, and the self-healing technology can be configured to protect against silent data corruption or bit rot by performing an integrity check of all objects within a bucket at least once a year.
Q:	Does Lyve Cloud guarantee permanent deletion of data?
A:	To permanently delete data, client applications should use S3 API calls to delete all objects and buckets that the application creates.
Q:	Does Lyve Cloud support multipart uploads?
A:	Lyve Cloud supports multipart uploads where the object part or chunk can range from 5 MB to 5 GB, with a maximum number of 10,000 parts per an object. This enables the efficient upload of large files and recovery from transmission errors.
Q:	What level of uptime does Lyve Cloud support?
A:	Lyve Cloud architecture implements redundancy at the hardware and software stack. All critical software components are designed to tolerate multiple failures. Lyve Cloud makes all reasonable efforts to maintain monthly uptime of 99.9% with applicable service credits.
Q:	Are there a minimum and maximum sizes for stored objects?
A:	There is no minimum object size, and the maximum object size for a single PUT operation is 5 TB. Lyve Cloud recommends using multipart uploads for files larger than 100 MB.
Q:	What kind of data can I store in Lyve Cloud?

A:	Lyve Cloud allows you to store any data in any format as long as it complies with Lyve Cloud's Terms and Conditions.
Q:	What versions of TLS / SSL does Lyve Cloud support?
A:	Lyve Cloud is compatible with TLS 1.2.
Q:	Does Lyve Cloud support WORM/Write-Once-Read-Many for data immutability?
A:	Lyve Cloud supports S3 object lock, which includes WORM/data immutability. This can be enabled at the bucket level.
Q:	Does Lyve Cloud support management and data access auditing?
A:	After enabling the Audit Log feature, Lyve Cloud keeps an audit log of all user access to your Lyve Cloud console and all client application's S3 operations. Users can designate a bucket to receive these audit logs from Lyve Cloud.
Q:	Can I connect Lyve Cloud with third-party tools?
A:	<p>Yes. You can connect to Lyve Cloud using either CyberDuck, Rclone, or any S3 browser. Using these tools, create folders in buckets or upload files.</p> <p>Consult your organization's policies and the EULA policies of the software before downloading 3rd party applications.</p> <p>To connect to Lyve Cloud with third-party tools, see Connecting S3 clients.</p>
Q:	Which operating systems environments are supported by the Lyve Cloud console?
A:	Lyve Cloud console supports Windows, Linux, and Mac operating systems. Lyve Cloud provides high standards of data protection and security. To ensure customer data is not compromised, we may not support clients running operating systems or application versions that are not actively supported by the relevant ISV or an open-source community.
Q:	Does Lyve Cloud support strong consistency?
A:	The Lyve Cloud S3 platform provides strong read-after-write consistency. This means that all GET , PUT , and LIST operations in S3 are consistent, ensuring that what you write is exactly what you will read, and that the results of a LIST operation will be a precise representation of the contents in the bucket. This feature is particularly useful for applications such as data lakes.

Buckets

[Does Lyve Cloud have any object naming limitations?](#)

[Do buckets have a maximum storage limit?](#)

[How do I manage my buckets using API?](#)

How do I utilize multipart upload and identify which parts of data have successfully been uploaded?
 When should I use ListMultipartUploads API?
 How can a customer confirm the encryption status of their objects?
 Can I check bucket utilization using commands?
 How can I create a bucket with object immutability (lock) enabled with a set duration using S3 API?
 How many buckets can you create in Lyve Cloud?

Q:	Does Lyve Cloud have any object naming limitations?
A:	The object name can contain any of these special characters like @, #, *, \$, %, &, !, ?, , , , ;, ' , " , , +, =, <, >, ^, (,), {, }, [,] and alphanumeric characters like 0-9, a-z, A-Z . However, using any of these characters can cause issues due to limiting factors of S3 client SDK .
Q:	Do buckets have a maximum storage limit?
A:	No. There is no storage limit for data stored in a single bucket in Lyve Cloud.
Q:	How do I manage my buckets using API?
A:	Lyve Cloud does not recommend managing buckets through API's. We recommend managing buckets directly from the Lyve Cloud console or from your data management platform.
Q:	How do I utilize multipart upload and identify which parts of data have successfully been uploaded?
A:	<p>Lyve Cloud supports multipart uploads of objects up to 5TBs. Using this method, large files are broken into smaller pieces for a more efficient upload. The pieces are then put back together at the end of the process. A multipart upload consists of three steps:</p> <ol style="list-style-type: none"> 1. Multipart upload initiation : When a multipart upload is initiated, AWS S3 will return an upload ID, which is a unique identifier for the multipart upload. This upload ID is a required field for all upload parts, list parts, complete upload or stop upload commands. 2. Parts upload: In this step, you can specify the part numbers which you would like to upload. A part number uniquely identifies a part and its place in the uploading object. For each part upload, please save the part number and ETag value. These will be needed to complete step 3. 3. Multipart upload completion : In this step, S3 will complete the upload by piecing the parts in order based on the part number. After a successful complete request, the individual parts will no longer exist. <p>For complete instructions, see Uploading and copying objects using multipart upload If an object is above 5TB in size, then the multipart upload completion command will not succeed. If a multi-part upload fails, the upload can resume with the part of the upload which failed. To view which parts of an upload succeeded, use the ListParts command. This will return all uploaded parts with their size and each one's part number. For more information, see: ListParts .</p>
Q:	When should I use ListMultipartUploads API?

A:	List multipart upload lists the in-progress multipart uploads that are initiated but are not yet completed or aborted. This API allows writing code that will the uploads that are not completed successfully on time. Lyve Cloud is performing this automatically, by cancelling all pending multi parts after 24 hours.
Q:	How can a customer confirm the encryption status of their objects?
A:	<p>Lyve Cloud enforces standard TLS 1.2 with 256-bit advanced encryption standard (AES) Galois/Counter Mode (GCM)—otherwise known as AES-256-GCM—to establish secure communications to the customer in transit and at rest. As an authenticated encryption algorithm, GCM provides proven security of the symmetric-key cryptographic cipher that has wide adoption for its performance. Seagate storage hardware is validated by Federal Information Processing Standards (FIPS) 140-2/3, which directly aligns with the Lyve Cloud focus on security and performance. To view the objects encryption status, please follow the steps below.</p> <p>Pre-requisites</p> <ul style="list-style-type: none"> Download a command-line tool such as AWS CLI <p>Open your command line application (Command Prompt for PC, Terminal for Mac) and use the following command.</p> <pre>--profile profile name --endpoint URL s3api head-object--bucket bucket name--key file name</pre> <pre>C:\Users\515515>aws --profile Kevin --endpoint https://s3.us-east-1.lyvecloud.seagate.com s3api head-object --bucket brawleytest --key Goals.docx</pre> <p>Result</p> <pre>bytes 20900 application/vnd.openxmlformats-officedocument.wordprocessingml.document "34162b3bec92a8334bd9fca388477f85" Mon, 06 Dec 2021 21:10:00 GMT AES256 bytes 20900 application/vnd.openxmlformats-officedocument.wordprocessingml.document "34162b3bec92a8334bd9fca388477f85" Mon, 06 Dec 2021 21:10:00 GMT AES256 METADATA 20210828T175628Z 1b6a154b13045741f4b61ab07ed55567754f44aa6796cc250b2a506c6c83a11a METADATA 20210828T175628Z 1b6a154b13045741f4b61ab07ed55567754f44aa6796cc250b2a506c6c83a11a</pre> <p>Note—The encryption is shown here as AES256 which is highlighted in bold.</p>
Q:	Can I check bucket utilization using commands?

A:

The content and data quantity in a bucket can be viewed through the following commands.

```
--profile profile name --endpoint URL s3 ls --summarize --human-readable --recursive s3://bucket
```

```
C:\Users\515515>aws --profile sv15 --endpoint https://s3.us-west-1.lyvecloud.seagate.com s3 ls --summarize --human-readable --recursive s3://mybuck
```

Result

```
2021-07-03 22:06:34    6 Bytes my-test-file.txt
2021-07-03 22:07:48   12 Bytes my-test-file1.txt
2021-07-03 22:29:33   11 Bytes my-test-file2.txt
2021-07-01 00:46:18  531 Bytes service-acounts.txt
Total Objects: 4 Total Size: 560 Bytes
```

Note—This command will list all contents in the bucket.

Q:

How can I create a bucket with object immutability (lock) enabled with a set duration using S3 API?

<p>A:</p>	<p>1. Create a bucket and enable object immutability using the following command:</p> <pre>aws s3api create-bucket --bucket &lt;bucket name> --object-lock-enabled-for-bucket --profile &lt;profile name> --endpoint &lt;endpoint></pre> <pre>aws --profile va3 --endpoint https://s3.us-east-1.lyvecloud.seagate.com s3api create-bucket --bucket mybucket --object-lock-enabled-for-bucket</pre> <p>Note—You can enable object immutability only while creating a bucket.</p> <p>Response</p> <pre>{ "Location": http://s3.us-east-1.lyvecloud.seagate.com/my-bucket-with-object-lock }</pre> <p>2. Set the duration using the following command:</p> <pre>aws s3api --profile &lt;profile name> put-object-lock-configuration --bucket &lt;bucketname> --object-lock-configuration &lt;value> --endpoint &lt;endpoint></pre> <pre>aws --profile va3 --endpoint https://s3.us-east-1.lyvecloud.seagate.com s3api create-bucket --bucket mybucket --object-lock-enabled-for-bucket</pre> <p>Response</p> <pre>aws s3api --profile lcprd put-object-lock-configuration --bucket mybucket --object-lock-configuration ObjectLockEnabled="Enabled",Rule={DefaultRetention={Mode="COMPLIANCE",Days=1}} --endpoint https://s3.us-east-1.lyvecloud.seagate.com</pre>
<p>Q:</p>	<p>How many buckets can you create in Lyve Cloud?</p>
<p>A:</p>	<p>You can create up to 100 buckets per account. If you need additional buckets, you can create a maximum of 1000 buckets. To create additional buckets, Create a support ticket.</p>

Service Account

What if I lose service account credentials?

<p>Q:</p>	<p>What if I lose service account credentials?</p>
------------------	--

A:	<p>As a reminder, Service accounts allow applications to authenticate and access Lyve Cloud buckets and objects. The access key and secret key (credentials) are generated when you create a service account. This information must be saved at the time of account creation and cannot be recovered afterwards.</p> <p>If necessary you can edit the service account, however, editing does not generate credentials.</p> <p>If you lose or misplace your service account credentials and believe you need new credentials, you must create a new service account. You may also delete the old service account.</p>
-----------	--

User and Roles

Can I switch admin roles in Lyve Cloud?

Q:	Can I switch admin roles in Lyve Cloud?
A:	<p>Yes. If you are an administrator and would like to give/transfer your role to another user, you must first create a new user and assign their role as an administrator. The new administrator must edit your role to Storage Admin or Auditor. For information on how to add and edit users, please visit Administrator's Guide - Identity and Access Management (IAM).</p>

Self Service

Who is eligible for the Lyve Cloud free tier?

How do I sign up for the Lyve Cloud free tier?

I cannot register for Lyve Cloud services using a credit card and get the following screen during registration.

I cannot access the registration link and receive the error The Site Can't be Reached.

After registration, the Lyve Cloud site fails to load with a Registration failed message.

When accessing the registration link, an unexpected error occurred.

How do I update or edit my credit card information?

When can I see my invoice for this month?

I did not receive the email with the confirmation code. How do I proceed?

I did not receive the email with a confirmation of account creation. How do I proceed?

How do I cancel my Lyve Cloud account?

When will my Credit Card be charged after the cancellation?

Why does my credit card provider show a transaction as "pending" from Lyve Cloud on my customer bill after I have registered for the free trial?

Q:	Who is eligible for the Lyve Cloud free tier?
A:	Anyone not previously signed up for Lyve Cloud is eligible for a free trial.

Q:	How do I sign up for the Lyve Cloud free tier?
A:	You can sign up for the Lyve Cloud Free Tier using https://signup.lyvecloud.seagate.com/product/250TB
Q:	I cannot register for Lyve Cloud services using a credit card and get the following screen during registration. [INSERT faq-cannot-register-with-credit-card-01.png]
A:	The Lyve Cloud subscription has reached the limit. You can wait for a period of time and try again. If you are still facing issues, please contact our support team at support.lyvecloud@seagate.com .
Q:	I cannot access the registration link and receive the error The Site Can't be Reached .
A:	You can perform the following actions and retry: <ul style="list-style-type: none"> • Check that your internet connection is working or restarting your router or modem often resolves this problem. • Disable your firewall and antivirus software. • Clear your browser cache. • Check your DNS server settings.
Q:	After registration, the Lyve Cloud site fails to load with a Registration failed message. [INSERT faq-registration-failed-01.jpg]
A:	The registration could have failed because the credit card authorization was incomplete or the credit card was declined. If the authorization is successful and still the registration fails, please contact our support team at support.lyvecloud@seagate.com .
Q:	When accessing the registration link, an unexpected error occurred. [INSERT faq-unexpected-error-01.jpg]
A:	Please contact our support team at support.lyvecloud@seagate.com .
Q:	How do I update or edit my credit card information?
A:	To update or edit credit card information such as the expiration date, see Managing payment method .Managing payment method
Q:	When can I see my invoice for this month?
A:	Your invoice is generated on the first of every month for the previous month. To view and download invoices, see Understanding billing .Understanding billing
Q:	I did not receive the email with the confirmation code. How do I proceed?

A:	Please check your inbox for an email entitled Lyve Cloud verification code . Be sure to also check your spam folder. If you still haven't received the email, please contact our support team support.lyvecloud@seagate.com .
Q:	I did not receive the email with a confirmation of account creation. How do I proceed?
A:	Please check your inbox for an email entitled You are invited to Lyve Cloud! . Be sure to also check your spam folder. If you still haven't received the email, please contact our support team support.lyvecloud@seagate.com .
Q:	How do I cancel my Lyve Cloud account?
A:	Lyve Cloud customers can close their own accounts. For more information, see Cancel Lyve Cloud service . Cancel Lyve Cloud service
Q:	When will my Credit Card be charged after the cancellation?
A:	When your account is cancelled, the billing period ends on the cancellation date. The bill is generated, and the credit card is charged on the cancellation date at the end of the day. For more information, see Calculating payment . Calculating payment
Q:	Why does my credit card provider show a transaction as "pending" from Lyve Cloud on my customer bill after I have registered for the free trial?
A:	Lyve Cloud only authorizes charges during registration and does not execute an actual charge to your credit card. Some credit card companies elect to note this authorization on customer statements as a pending charge. Not all companies follow this practice. While you may see a charge from Lyve Cloud as pending, you will not be charged for Lyve Cloud service during the free trial period. Please speak with your credit card company about concerns with this notation on your customer bill. You may also contact Lyve Cloud support for further assistance.

Sub-accounts

What happens to all the sub-accounts in the master account once the master account is disabled?

Will billing stop once a sub-account is disabled?

Will all users of a sub-account be disabled when a sub-account is disabled?

Are service accounts disabled after a sub-account is disabled?

Q:	What happens to all the sub-accounts in the master account once the master account is disabled?
A:	Once a master account is disabled, all associated sub-accounts are disabled.
Q:	Will billing stop once a sub-account is disabled?

A:	Yes, billing is stopped once a sub-account is disabled. For example, If the sub-account is disabled on June 15, the usage calculation for June month is pro-rata. The sub-account is billed from June 1 to June 15.
Q:	Will all users of a sub-account be disabled when a sub-account is disabled?
A:	Yes, all sub-account users are disabled when a sub-account is disabled.
Q:	Are service accounts disabled after a sub-account is disabled?
A:	Yes, all service accounts are disabled when a sub-account is disabled.

Billing

[What storage classes does Lyve Cloud offer?](#)

[How much does it cost to use Lyve Cloud?](#)

[How is the capacity utilization metric calculated?](#)

[What is the cost impact of using Lyve Cloud's immutability feature?](#)

[What is the cost impact of using Lyve Cloud's audit logging features?](#)

[I'm not planning to use the object immutability feature but I may be overwriting a file with the same name multiple times. What is the cost impact of this?](#)

[Are there any charges for egress or API requests in Lyve Cloud?](#)

[Can I get an extension of a free trial?](#)

[I am currently using the free trail, how can I become a paid customer?](#)

[How is Lyve Cloud billing information collected, and how do I get billed?](#)

[Where can I get a copy of my invoice?](#)

Q:	What storage classes does Lyve Cloud offer?
A:	Lyve Cloud offers a single storage tier and a single consistent price based on the amount of data stored.
Q:	How much does it cost to use Lyve Cloud?
A:	The Lyve Cloud service is based on a simple pricing model. Lyve Cloud charges a monthly fee based on the average capacity of data stored during a given month. You are not charged for S3 API calls or data egress.
Q:	How is the capacity utilization metric calculated?
A:	Lyve Cloud measures the total storage consumption four times a day. Storage consumption number recorded is the average of the four instances per day. The monthly average usage is calculated by taking the average of all daily records for a given month. For example, if the daily average for the first 10 days of the month is 10 TB, the average for next 10 days is 20 TB and the average for last 10 days of the month is 30 TB, then the total consumption for the month is 20 TB $((10+20+30)/3=20)$.
Q:	What is the cost impact of using Lyve Cloud's immutability feature?

A:	Once object immutability is switched on, bucket versioning is automatically enabled. Updating a file creates a new version of the objects being stored. This process results in an increase in storage usage and cost.
Q:	What is the cost impact of using Lyve Cloud's audit logging features?
A:	Enabling audit logging creates log files for buckets or console activity. These log files are stored and treated like all other billable storage.
Q:	I'm not planning to use the object immutability feature but I may be overwriting a file with the same name multiple times. What is the cost impact of this?
A:	Let us consider an example to answer this question. On day 1, you store a file called "file1.txt". On day 2, you then overwrite "file1.txt" with a new copy of "file1.txt" while not changing the file name. The original copy of "file1.txt" is, therefore, overwritten and the cost is impacted only by changes in the file size of the latest "file1.txt" that may have occurred.
Q:	Are there any charges for egress or API requests in Lyve Cloud?
A:	There are no egress or API charges in Lyve Cloud.
Q:	Can I get an extension of a free trial?
A:	Yes, please use the extend option on your trial banner or contact your Lyve Cloud sales representative for an extension.
Q:	I am currently using the free trail, how can I become a paid customer?
A:	Please contact your Lyve Cloud sales representative to become a paid customer.
Q:	How is Lyve Cloud billing information collected, and how do I get billed?
A:	Invoices reflect a calculation of the average monthly storage usage for the previous month. Invoices are available on the first day of every month. You then have an agreed time period to complete payment.
Q:	Where can I get a copy of my invoice?
A:	You can view and download a copy of your invoice from the Lyve Cloud console. For more information see Understanding billing . Understanding billing

Security

[How is my data protected?](#)

[Does Lyve Cloud mine my data?](#)

[What happens to my data if I am no longer a Lyve Cloud customer?](#)

[What authentication mechanisms are supported?](#)

How secure are Lyve Cloud datacenters?
 Which OTP applications can be used for MFA login?
 Can I use CORS with Lyve Cloud?
 How can I change my authentication method?
 Why must I use a mobile phone to set up MFA?
 My mobile device with authenticator app is lost or stolen, what do I do?
 Can Email be used as the 2nd method of Auth for MFA?
 Can organizations use their own SAML with MFA?
 I have lost my recovery code, how do I login to the Lyve Cloud console?
 How can I change my Multi-Factor Authentication (MFA) Phone Number?
 Can I change my authenticator app?
 Can I register multiple Lyve Cloud accounts in an authenticator app for MFA?

Q:	How is my data protected?
A:	<p>Lyve Cloud protects customer data in transit by using:</p> <ul style="list-style-type: none"> • Transport layer security (TLS) for data in-flight TLS 1.2 (AES-256-GCM) • Encryption for data at rest <p>The data is always encrypted at rest using one of two server-side methods:</p> <ul style="list-style-type: none"> • Encryption with a client-provided key (part of S3 request headers) - SSE-C • Encryption with an S3 managed encryption keys - SSE-S3 <p>All data, regardless of whether it is encrypted or not on the client-side (SSE-C), is encrypted using AES 256-bit encryption at rest. The keys are never shared and can be rotated based on the customer's security policy. Data in flight is encrypted using TLS 1.2, and client applications can only connect using HTTPS protocol. Lyve Cloud follows industry best practices for design and security models.</p> <p>Contact sales.lyvecloud@seagate.com for a complete overview of security analysis conducted by a third party.</p>
Q:	Does Lyve Cloud mine my data?
A:	No, Lyve Cloud does not mine any customer data. All data stored in Lyve Cloud is encrypted. We strongly recommend that customers use client-side encryption for complete data protection.
Q:	What happens to my data if I am no longer a Lyve Cloud customer?
A:	Lyve Cloud does not make any secondary copies of the data. To permanently delete the data, the client application should use the S3 API calls to delete all objects and buckets it created. Once this is complete, customers can email support.lyvecloud@seagate.com to request that their account information be permanently deleted. This ensures that any remaining customer information is removed from the Lyve Cloud cluster.
Q:	What authentication mechanisms are supported?

A:	<p>Access to the Lyve Cloud admin portal is supported by multiple authentication schemes, including:</p> <ul style="list-style-type: none"> • Multi-factor authentication using either SMS OTP (One Time Password) or an authenticator mobile app • Federated login using the customer's IDP login flow
Q:	How secure are Lyve Cloud datacenters?
A:	<p>Lyve Cloud prioritizes a secure and protected infrastructure.</p> <ul style="list-style-type: none"> • A dedicated staff manages and protects each site 24×7, year-round. • Each site is equipped with security cameras to monitor inside the data center and the surrounding area. • Facilities are unmarked so as to not draw attention from outside. • Building access is controlled using biometric measures.
Q:	Which OTP applications can be used for MFA login?
A:	<p>Lyve Cloud supports the use of third-party authenticator apps as verification methods for MFA logins. You can use any authenticator app that generates temporary codes based on the time-based one-time password. There are many free and paid authenticator apps to choose from. Widely-used options include Google Authenticator, Microsoft Authenticator, DUO, Authy, Okta Verify, Auth0 Guardian, OneLogin Protect, and Oracle Authenticator.</p>
Q:	Can I use CORS with Lyve Cloud?
A:	<p>Currently, Lyve Cloud does not support Cross-Origin Resource Sharing (CORS), nor does it support hosting static websites using custom domains or anonymous access to public buckets.</p>
Q:	How can I change my authentication method?
A:	<p>You must contact your administrator to reset MFA for the user. To reset MFA, see Resetting MFA For An Individual IAM User. After resetting MFA, you must again enroll in MFA. For more information, see Enrolling in MFA.</p>
Q:	Why must I use a mobile phone to set up MFA?
A:	<p>Your device is unique to you. This helps to ensure that your account can only be accessed by the person in possession of your phone. Even if someone has your Lyve Cloud credentials, they will not be able to access your Lyve Cloud account without your mobile phone.</p>
Q:	My mobile device with authenticator app is lost or stolen, what do I do?

A:	To change your phone number you must contact the administrator of your account to reset MFA. For more information, see Resetting MFA For An Individual IAM User . After the administrator resets MFA, you must again enroll in MFA on your new device. For more information, see Enrolling in MFA .
Q:	Can Email be used as the 2nd method of Auth for MFA?
A:	No, email is not supported as an MFA method. We only support the authenticator apps and SMS. This is because email credentials can be easily compromised or reset. With a mobile device, it is more difficult to get the SMS code or use the authenticator app to access your account than it is to access an email account.
Q:	Can organizations use their own SAML with MFA?
A:	Yes, organizations can use their own SAML with MFA. Lyve Cloud MFA always applies to password users even if federated login is enabled for the account.
Q:	I have lost my recovery code, how do I login to the Lyve Cloud console?
A:	You must contact your administrator to reset MFA. To reset MFA for the user, see Resetting MFA For An Individual IAM User . After resetting MFA, you must again enroll in MFA. For more information, see Enrolling in MFA .
Q:	How can I change my Multi-Factor Authentication (MFA) Phone Number?
A:	An administrator must reset the MFA for a user to change the associated phone number. To reset MFA for the user, see Resetting MFA For An Individual IAM User . After resetting MFA, you must again enroll for MFA. For more information, see Enrolling in MFA .
Q:	Can I change my authenticator app?
A:	Yes, you can change the authenticator app by installing the preferred authenticator app. You must contact your administrator to reset MFA for the user. To reset MFA, see Resetting MFA For An Individual IAM User . After resetting MFA, you must again enroll in MFA. For more information, see Enrolling in MFA .
Q:	Can I register multiple Lyve Cloud accounts in an authenticator app for MFA?
A:	<p>You may use different authenticator apps for different Lyve Cloud accounts. If you are required to use the same authenticator app for multiple Lyve Cloud accounts, refer to the MFA application's help section to learn how to add multiple accounts. Follow the steps based on the desired authenticator app. Refer to few of the commonly used authenticator apps:</p> <ul style="list-style-type: none"> • Google • Microsoft • Oracle

Support

How can I contact Support?

What can I expect after a support ticket is raised or an email is sent?

What details are required when contacting support through email?

When can I contact Lyve Cloud Support?

When will the support team contact me after I create a support ticket?

Where to find help?

Q:	How can I contact Support?
A:	<p>There are two ways to contact Lyve Cloud Support.</p> <ol style="list-style-type: none">1. Create a support ticket from the Lyve Cloud console. For more information see the section called "Creating a support ticket".2. Email support.lyvecloud@seagate.com
Q:	What can I expect after a support ticket is raised or an email is sent?
A:	<p>Communication</p> <p>Once you submit a ticket, you will get an auto-generated email from support.lyvecloud@seagate.com that includes your ticket number.</p> <p>Your ticket is added to a queue and is processed in the order that it is received. However, support times do vary depending on issue severity.</p> <p>If you need to contact support for any reason regarding your ticket, you can reply to the auto-generated email, or you can leave a comment in the portal under your ticket details.</p> <p>SLA</p> <p>We will use reasonable efforts to provide a workaround within 24 hours. Afterwards, we will do a root cause analysis and fix the issue to ensure it does not occur again.</p> <p>Escalation</p> <p>Lyve Cloud uses an internal escalation process to ensure tickets are resolved quickly and that they receive the right technical expertise. Following this process, there may be cases where customers learn their ticket has been escalated to another part of the support organization.</p> <p>Resolution</p> <p>When the ticket is resolved, our support team will send out a post-support survey to assess the effectiveness of our support program. Your feedback is highly valuable to us.</p>
Q:	What details are required when contacting support through email?

A:	<p>When emailing support, you must include as much information as possible. Please include the following in your request:</p> <ul style="list-style-type: none"> • Your account ID. It can be found on your Lyve Cloud URL - <account_id>.console.lyvecloud.seagate.com • Summary and description of the issue you are experiencing • Environment (i.e. us-east-1 (Virginia), us-west-1 (California), ap-southeast-1 (Singapore), etc.) • URL Link (the section called "S3 API endpoints") • Attachment with an error screen, if any
Q:	When can I contact Lyve Cloud Support?
A:	Lyve Cloud support teams are staffed around the globe, working 24/7, 365 days a year.
Q:	When will the support team contact me after I create a support ticket?
A:	A new ticket will trigger an auto-response acknowledgement that you will receive within 30 minutes. Our support team will revert you within 4 hours. Seagate will use reasonable efforts to provide a workaround within 24 hours. Later, we will do a root cause analysis and fix the issue to ensure it does not occur again.
Q:	Where to find help?
A:	Before contacting support, it may be quick and convenient to find the information you need in our documentation. The Lyve Cloud Documentation provides several resources to answer your questions. Resources include our Quick Start Guide, Video tutorials, Connecting to S3 clients, Known Issues, FAQs, and much more.

Partner

[Where can I view my customers' storage activity?](#)

[Where can I view my bill?](#)

[How can I calculate storage cost per customer?](#)

Q:	Where can I view my customers' storage activity?
A:	Lyve Cloud partners can view the monthly storage activity of their customers by viewing the Usage Trend dashboard on the Lyve Cloud console. Total number of users, buckets, service accounts, average usage, and support tickets can also be viewed on the Customer distributions dashboard on the Lyve Cloud console.
Q:	Where can I view my bill?
A:	Your bill is sent to the email address registered with the partner account.

Q:	How can I calculate storage cost per customer?
A:	Invoices feature monthly storage used by customers to let partners determine to price for each customer. Lyve Cloud Partner Portal provides usage by customers that can be downloaded from the Home page dashboard. This enables partners to determine to price for each of their customers.

HIPAA

[What is HIPAA?](#)

[Is Seagate Lyve Cloud HIPAA compliant?](#)

[What measures does Seagate use to ensure its service is HIPAA compliant?](#)

[Where can I learn more about HIPAA?](#)

Q:	What is HIPAA?
A:	<p>The Health Insurance Portability and Accountability Act (HIPAA) is a US law enacted in 1996. HIPAA is a comprehensive set of standards that regulate the protection and use of protected health information (PHI) in the healthcare industry. The law applies to healthcare providers, health plans, and healthcare clearinghouses that transmit health information in electronic form.</p> <p>HIPAA establishes national standards for the privacy and security of PHI and requires that these entities implement appropriate administrative, physical, and technical safeguards to protect the confidentiality, integrity, and availability of PHI. HIPAA also sets rules for the use and disclosure of PHI and gives individuals certain rights with respect to their PHI, such as the right to access, inspect, and receive a copy of their PHI.</p> <p>In summary, HIPAA is a comprehensive set of standards and rules that regulate the protection and use of PHI in the healthcare industry and establishes national standards for the privacy and security of PHI.</p>
Q:	Is Seagate Lyve Cloud HIPAA compliant?
A:	Yes, Seagate's Lyve Cloud storage service has been audited by trusted third-party certifiers and is certified to comply with all HIPAA standards for the secure handling and protection of PHI.
Q:	What measures does Seagate use to ensure its service is HIPAA compliant?

<p>A:</p>	<p>Seagate takes numerous measures to ensure that our service is compliant with HIPAA standards and provides a highly secure and robust data storage service. Such measures include:</p> <ul style="list-style-type: none"> • Lyve Cloud is certified to the ISO/IEC 27001 Standard, as well as attested to the SOC 2 Type 2 Standard. We are also HIPAA compliant. <ul style="list-style-type: none"> • ISO 27001 is the world's leading information security standard, providing control requirements to create an Information Security Management System (ISMS). • SOC 2 is an extremely popular form of cybersecurity audit, used by a rapidly growing number of organizations to demonstrate they take cybersecurity and privacy seriously. • HIPAA Security Rule addresses the safeguarding of ePHI through the application of administrative, physical, and technical safeguards. Compliance is required by all covered organizations defined by HIPAA and the Office of Civil Rights (OCR) audit guidelines and assessment standards. • Two-factor authentication for users to ensure secure data custodians. • Standardized TLS 1.2 data encryption. • Authentication and authorization within every data transaction, using both account access keys and a cryptographic signature. • Secure data deletion with Secure Erase. When a customer ends their service with Lyve Cloud, their data is securely cryptographically erased as the SSE-C key is only available to the customer in the API. • Enacting proper staff security awareness training and regulation policies, including procedures for authorizing access to PHI, as well as security incident response. • Entering into and documenting appropriate business associate contracts with covered entities using Lyve Cloud service. • Periodic, thorough risk analyses of current business practices/audits, with an aim to identify potential security risks and test current security and contingency policies. • Safe logging and monitoring procedures ensure the recording of all login attempts, individuals who access electronic PHI, as well as any security incidents and mitigation responses involved. • Extensive resilience testing and contingency backup architecture that prevents data loss in the case of a datacenter going down.
<p>Q:</p>	<p>Where can I learn more about HIPAA?</p>
<p>A:</p>	<p>For more information regarding HIPAA compliance with Seagate's Lyve Cloud storage service, please see Lyve Cloud Compliance, or visit Lyve Cloud HIPAA Business Associate Addendum. Lyve Cloud Compliance</p>

Release Notes

The Lyve Cloud Release Notes provide information about the latest Lyve Cloud feature release. Information that doesn't apply to the release will be archived.

Release 5.7

February 03, 2023

Features

Registering for Lyve Cloud using a credit card, AKA Lyve Cloud Self-Service

Customers can use a credit card to use Lyve Cloud Self-Service and register for a Lyve Cloud account. The account is charged based on their usage. The invoice is generated, and the credit card is charged on the first of every month for the previous month's actual usage.

Registration Process

During registration, the customer completes a form using the following link:

<https://signup.lyvecloud.seagate.com/product/250TB>

The requested information from the customer includes:

- Account Information
- Contact Information
- Billing Information
- Payment details

For more information, see [Registering Lyve Cloud using a credit card](#). Registering Lyve Cloud using a credit card

Credit Card Authorization

The customer specifies payment information while registering for Lyve Cloud, and the credit card is authorized.

Trial Period

Customers are offered a free 30-day trial. Customers are charged a pro-rated fee for the calendar month in which the trial ends. The pro-rated fee is based on the number of days the account is active, between the day after the trial ends and the end of the month.

Billing and Charging

Customers are committing to a minimum usage of 250TB monthly, and the bill for a full month of use is 250 TB/month * \$9 / TB = \$2,250/month.

- if actual usage \leq 250 TB, the customer is charged for 250TB
- if actual usage $>$ 250 TB, the customer is charged for actual usage multiplied by \$9

The customer bill generates, and the credit card is charged for the previous month's usage on the first day of every month.

For more information, see [Understanding billing for customers signing up with Credit Card](#) and [Calculating Payment](#). Understanding billing for customers signing up with Credit Card Calculating payment

Cancelling the Lyve Cloud account

- Accounts are closed immediately, and the customer is not charged if an account is cancelled within the trial period.
- Customers are charged based on the number of days of usage if the account is cancelled after the trial period ends.

For more information, see [Cancel Lyve Cloud service](#). Cancel Lyve Cloud service

Release 1.2.7 (Data Mover)

February 03, 2023

Features

Lyve Cloud Data Mover is a file copy utility that can replicate on-premises file systems to Lyve Cloud and back to on-premise. For more information, see [Lyve Cloud Data Mover](#).

You must download the OVA files for Data Mover Portal and Data Mover Worker from the Lyve Cloud console. For more information, see [Install Data Mover Portal](#) and [Install Data Mover Worker](#).

Release 5.6

January 15, 2023

Features

Lyve Cloud Account API is a set of methods for Lyve Cloud programmatic management. It can be accessed via <https://api.lyvecloud.seagate.com/v2>. The account API credentials can be generated only by account

administrators. For more information, see [Using Account API](#).

The APIs provide programmatic access to the Lyve Cloud to perform administrative tasks. The API uses a JSON structure and may include Boolean, Number, and String data types. The following are included in this initial release.

- Authentication: It authenticates to the service and manages the session token.
- Permission: It allows access control for the buckets.
- Service Account: It allows control for S3 API credentials.
- Usage: It provides information about the account's storage usage.

For more information, see [Lyve Cloud Account API](#).

Release 5.5

October 12, 2022

Features

- Lyve Cloud Compute by Zadara allows users to rent virtual servers on which to run their own applications adjacent to Lyve Cloud S3-compatible storage. It provides a web service through which a user can create a virtual machine image and boot a virtual machine. Users can create, launch, and terminate instances as needed, paying for active servers. Lyve Cloud Compute also provides control over autoscaling groups, virtual private networks, block storage and other related resources. It can be managed using a reach UI, as well as CLI and EC2-compatible API.

For more information, see [Lyve Cloud Compute](#).

- Lyve Cloud Analytics by Iguazio is a production-centric data analytics solution that accelerates data engineering, data science, and data governance, all while encouraging development and machine learning operations best practices. This architecture is built on a foundation of data storage, computing, monitoring, and security. The platform's tight integration provides users with a unified approach that reduces the technical debt associated with modern analytic and machine learning systems.

For more information, see [Lyve Cloud Analytics](#).

Release 5.4

August 13, 2022

Features

The Sub-accounts feature allows you to maintain a multi-level account structure and replaces the partner portal. The sub-accounts is available for Managed Service Providers (MSP) and Resellers. If you are not an MSP or Reseller and still want to use sub-accounts, contact the Lyve Cloud support team

at support.lyvecloud@seagate.com.

Lyve Cloud provisions the master account, and the master account provisions Sub-accounts. These Sub-accounts are managed by and connected to the master account. Each sub-account can provision and manage its storage, users, and account settings. For more information, see [Administrator's Guide - Sub-Account Management](#).

Enhancements

The service account expiration enhances the security of the service account. You can now set the expiration duration for a service account. When a service account expires, the secret credentials associated also expire. For more information, see [Service account settings](#).

You can Clone to create a duplicate Service Account, where it creates new access and secret keys.

Release 5.3

January 15, 2022

Features

- A new Start Here panel located on the right side of the Header pane is available for quick access to basic tasks, training videos, and other helpful resources to get started with Lyve Cloud.
- Global Account Management allows customers to create buckets in two different regions or create service accounts to access buckets in different regions. Currently, there are two regions called California (us-west-1) and Virginia (us-east-1). This provides simplified management of multiple regions on Lyve Cloud console and the ability to increase redundancy and availability. For more information see [Understanding Global Accounts](#).
- We are introducing a Marketplace that provides more information about partner solutions, including Backup and Recovery, Surveillance and Compute that are certified with Lyve Cloud.
- Billing features:
 - View the estimated bill for next month's payment cycle. For more information see [Understanding the home page dashboard](#).
 - View and download average storage use reports from the previous months or years since the beginning of account creation. For more information, see [Understanding the home page dashboard](#).
 - See the Billing tab recently added to the main navigation to track and download invoices.

Enhancements

- Compliance Mode has now been renamed to Object Immutability. For more information, see [Using object immutability](#).
- Instant copy functionality of the S3 endpoint from the bucket details page. For more information, see [Supported S3 API calls](#).
- Support for S3 Select capability, which allows the retrieval of a subset of data from an object and can

improve data transfer efficiency.

Release features overview video

[Lyve Cloud - Feature Release 5.3](#)

Troubleshooting Guides

Experiencing issues? Choose from the topics below to follow step-by-step guides to help you get back on track.

Login issue

Experiencing issues logging into Lyve Cloud? Are you unable to reach the console login page or having issues loading the page that is preventing you from logging in? Are you having trouble reaching your S3 buckets?

Problems accessing the Lyve Cloud login page or your S3 buckets are usually caused by two issues:

- The Lyve Cloud IP addresses are blocked by your organization's router or firewall.
- Certain services used by Lyve Cloud to enable secure login are blocked by your organization's firewall rules or your router's access control list.

Solution

1. Ensure Lyve Cloud domains are allowed.

Make sure the following URLs are added to the allow list of your firewall:

- Lyve Cloud domains:

Lyve Cloud Console URL <account_ID>.console.lyvecloud.seagate.com

Region	Endpoint	Virtual hosted-style endpoint
US-East-1	<code>https://s3.us-east-1.lyvecloud.seagate.com</code>	<code>https://[bucket_name].s3.us-east-1.lyvecloud.seagate.com</code>
US-West-1	<code>https://s3.us-west-1.lyvecloud.seagate.com</code>	<code>https://[bucket_name].s3.us-west-1.lyvecloud.seagate.com</code>
US-Central-1	<code>https://s3.us-central-1.lyvecloud.seagate.com</code>	<code>https://[bucket_name].s3.us-central-1.lyvecloud.seagate.com</code>

Region	Endpoint	Virtual hosted-style endpoint
US-Central-2	<code>https://s3.us-central-2.lyvecloud.seagate.com</code>	<code>https://[bucket_name].s3.us-central-2.lyvecloud.seagate.com</code>
EU-West-1	<code>https://s3.eu-west-1.lyvecloud.seagate.com</code>	<code>https://[bucket_name].s3.eu-west-1.lyvecloud.seagate.com</code>
AP-Southeast-1	<code>https://s3.ap-southeast-1.lyvecloud.seagate.com</code>	<code>https://[bucket_name].s3.ap-southeast-1.lyvecloud.seagate.com</code>

Additional content for authentication:

Lyve Cloud Console	Documentation Portal	Lyve Cloud Status
<code>cdn.auth0.com</code>	<code>help.lyvecloud.seagate.com</code>	<code>Status.lyvecloud.seagate.com</code>
<code>cdn.jsdelivr.net</code>	<code>*.cloudfront.net</code>	<code>*.cloudfront.net</code>
<code>maxcdn.bootstrapcdn.com</code>	<code>userfiles-kb.s3.amazonaws.com</code>	<code>www.google.com</code>
<code>code.jquery.com</code>	<code>fonts.gstatic.com</code>	<code>www.gstatic.com</code>
<code>ajax.googleapis.com</code>		<code>fonts.gstatic.com</code>
<code>fonts.googleapis.com</code>		<code>polyfill.io</code>
<code>fonts.gstatic.com</code>		

2. Verify Port 443 access to the domains listed in step 1.

Check to verify that Port 443 is not blocked by your router or firewall. If the port is blocked, consult your IT department to enable access.

3. Reload the Lyve Cloud page to verify changes made in steps 1 and 2 are allowing login access.

If you continue to experience issues logging into Lyve Cloud after confirming access to the port and domains recommended in this document, please contact your Seagate Customer Success Manager for further assistance.

Resources

Partner deployment solutions and additional resources can be found in [Lyve Cloud S3 Storage Resources Guide](#).

Lyve Cloud Compliance

This Lyve Cloud Compliance Overview document provides a summary of the key certifications and compliance requirements that organizations should consider when selecting and implementing cloud services. From HIPAA to ISO 27001 to SOC 2, this guide covers the most widely recognized standards and best practices for cloud security and privacy. It explains why these certifications and requirements are critical to ensure the protection of sensitive information in the cloud.

HIPAA

Lyve Cloud has a HIPAA Compliant report.

Seagate built Lyve Cloud to be the industry leader for the Healthcare sector with core tenants around: Resilience, Compliance, Performance and Value. Lyve Cloud, according to Health and Human Services HHS.org, is a No View SaaS provider. Being HIPAA compliant, we always ensure complete protection of your data under our care. Lyve Cloud uses a standards-based approach which produces the highest level of compliance and security in the market.

For the cost of an archive, we are delivering a Hot Tier of object storage in a high availability HIPAA-compliant offering. Lyve Cloud is new and more modern than most traditional providers and, as such was built for security upfront with fewer vulnerabilities.

Lyve Cloud has a HIPAA Compliant report

A HIPAA (Health Insurance Portability and Accountability Act) compliant cloud is expected to meet certain standards to ensure the protection of sensitive health information. The following are some of the key expectations of a HIPAA-compliant cloud: A HIPAA (Health Insurance Portability and Accountability Act) compliant cloud is expected to meet certain standards to ensure the protection of sensitive health information. The following are some of the key expectations of a HIPAA-compliant cloud:

- **Security:** The cloud provider must have strong security measures in place to protect the confidentiality, integrity, and availability of electronically protected health information (ePHI). This includes encryption, access controls, and audit logs.
- **Privacy:** The cloud provider must have strict privacy policies and procedures in place to ensure that ePHI is only accessed by authorized individuals. This includes limiting access to ePHI, conducting background checks on employees, and training employees on privacy and security.
- **Compliance:** The cloud provider must comply with all HIPAA regulations, including the HIPAA Security Rule, the HIPAA Breach Notification Rule and associated policies, procedures, and documentation.
- **Business Associate Agreement:** The cloud provider must sign a Business Associate Agreement (BAA) with its clients to ensure that they understand and agree to comply with HIPAA regulations. Business Associate Agreement: The cloud provider must sign a Business Associate Agreement (BAA) with its clients to ensure that they understand and agree to comply with HIPAA regulations.
- **Disaster Recovery and Business Continuity:** The cloud provider must have a disaster recovery and business continuity plan in place to ensure that ePHI can be recovered in the event of a disaster or data loss.

- **Monitoring and Auditing:** The cloud provider must regularly monitor and audit their systems and processes to ensure that they comply with HIPAA regulations and to identify and address any potential security or privacy breaches.
- **Technical Support:** The cloud provider must provide technical support to their clients to ensure that they can effectively use the cloud and resolve any issues that may arise. Technical Support: The cloud provider must provide technical support to their clients to ensure that they can effectively use the cloud and resolve any issues that may arise.

It is important to note that the responsibility for ensuring HIPAA compliance does not rest solely with the cloud provider. The entity that uses the cloud, known as the covered entity, must also ensure that they comply with HIPAA regulations.

ISO 27001

ISO 27001 is an international standard for information security management that outlines a set of security controls and best practices for protecting sensitive information. A cloud service that holds an ISO 27001 certificate is expected to meet the following expectations:

Information Security Management System (ISMS): The cloud service must have an ISMS in place that is designed to manage and protect sensitive information. The ISMS should cover all aspects of information security, including but not limited to access control, incident management, risk management, and business continuity.

- **Security Controls:** The cloud service must have a comprehensive set of security controls in place to protect sensitive information. These controls should include access controls, encryption, firewalls, and intrusion detection systems. Security Controls: The cloud service must have a comprehensive set of security controls in place to protect sensitive information. These controls should include access controls, encryption, firewalls, and intrusion detection systems.
- **Risk Management:** The cloud service must have a robust risk management process in place to identify, assess, and mitigate potential risks to sensitive information. This includes conducting regular security assessments and risk analyses and implementing appropriate controls to mitigate identified risks.
- **Data Privacy:** The cloud service must have strict data privacy policies and procedures in place to ensure that sensitive information is protected and only accessed by authorized individuals. Business Continuity and Disaster Recovery: The cloud service must have a business continuity and disaster recovery plan in place to ensure that sensitive information is protected in the event of a disaster or data loss.
- **Monitoring and Auditing:** The cloud service must regularly monitor and audit its systems and processes to ensure that they are compliant with ISO 27001 and to identify and address any potential security or privacy breaches.
- **Technical Support:** The cloud service must provide technical support to their clients to ensure that they can effectively use the cloud and resolve any issues that may arise.
- **Continual Improvement:** The cloud service must have a continuous improvement process in place to ensure that its security controls and processes are updated and improved over time to stay ahead of emerging threats and risks.

A cloud service with an ISO 27001 certificate is expected to have a comprehensive and robust approach to information security management that covers all aspects of information security, including but not limited to risk management, data privacy, business continuity, and monitoring and auditing.

Type 2 SOC 2

Lyve Cloud has a Type 2 SOC 2 Attestation report

A SOC 2 attestation is a third-party assessment of a cloud service provider's controls related to the security, availability, processing integrity, confidentiality, and privacy of the information processed by the service. Type 2 SOC 2 attestation specifically refers to an assessment of the cloud service provider's controls over a period of time, typically six months or more.

A cloud service provider that has received a Type 2 SOC 2 attestation is expected to meet the following expectations:

- **Information Security:** The cloud service provider must have a comprehensive information security program in place to protect sensitive information and meet the requirements of the SOC 2 standard. This includes implementing security controls such as access controls, encryption, and firewalls and regularly monitoring and auditing their systems.
- **Availability:** The cloud service provider must have a robust availability program in place to ensure that their services are available to their clients when needed. This includes implementing redundant systems and processes, monitoring the availability of their services, and having a disaster recovery plan in place. Availability: The cloud service provider must have a robust availability program in place to ensure that their services are available to their clients when needed. This includes implementing redundant systems and processes, monitoring the availability of their services, and having a disaster recovery plan in place.
- **Processing Integrity:** The cloud service provider must have controls in place to ensure the integrity of the information processed by their services. This includes implementing controls such as data validation, error checking, and audit trails.
- **Confidentiality:** The cloud service provider must have strict confidentiality policies and procedures in place to ensure that sensitive information is protected and only accessed by authorized individuals.
- **Privacy:** The cloud service provider must have a comprehensive privacy program in place to meet the requirements of relevant privacy regulations and standards, such as GDPR or HIPAA. Privacy: The cloud service provider must have a comprehensive privacy program in place to meet the requirements of relevant privacy regulations and standards, such as GDPR or HIPAA.
- **Monitoring and Auditing:** The cloud service provider must regularly monitor and audit their systems and processes to ensure that they comply with SOC 2 and to identify and address any potential security or privacy breaches.
- **Technical Support:** The cloud service provider must provide technical support to their clients to ensure that they can effectively use their services and resolve any issues that may arise.
- **Continual Improvement:** The cloud service provider must have a continuous improvement process in place to ensure that their security controls and processes are updated and improved over time to stay ahead of emerging threats and risks.

A cloud service provider with a Type 2 SOC 2 attestation is expected to have a comprehensive and robust approach to information security, privacy, and availability and to have controls in place to ensure the integrity and confidentiality of the information processed by their services.

Summary

Having key certifications and meeting compliance requirements is important for Lyve Cloud to build customer trust and confidence. Certifications demonstrate compliance with legal and regulatory requirements, improve security, help manage risks, and can provide a competitive advantage.

Lyve Cloud has key international certifications and attestations and is constantly expanding this list based on customer feedback.

Lyve Cloud Limitations

Bucket and object limitations

The table lists the limitations of buckets, objects, service accounts and permissions and their associated actions.

Item	Specification
Maximum number of buckets	100 buckets per account.. Note —To create additional buckets (up to 1000 buckets), raise a Support Ticket. For more information, see Creating a support ticket .
Maximum number of objects per bucket	no-limit
Maximum object size (multipart upload)	5 TiB
Minimum object size	0 B
Maximum object size per PUT operation	5 GiB
Maximum length for bucket names	63
Maximum length for object names	1024. Note —if an object name exceeds 255 characters in length, it must contain one or more "/" characters. Furthermore, any segment between two "/" characters should not exceed 255 characters in length.

Multipart upload limitation

The table lists the multipart upload limitations.

Item	Specification
Maximum number of parts per upload	10,000

Item	Specification
Part size range	5MiB to 4GiB. The last part can be 0 B to 5 GiB
Maximum number of parts returned per list parts request	10,000
Maximum number of objects returned per list objects request	1,000
Maximum number of multipart uploads returned per list multipart uploads request	1,000

Incompatible S3 API calls

The table lists S3 API calls that are incompatible.

S3 API	Description
<code>ListMultipartUploads</code>	This command does not list active multipart uploads using a bucket name. Add a prefix and the prefix value must be an object name to list the object multipart in the bucket.
<code>GetBucketLocation</code>	This command returns the bucket location, regardless of a Signature v4 region different from the S3 endpoint region.
<code>ListBuckets</code>	<ul style="list-style-type: none"> This command lists all the buckets, regardless of a Signature v4 region different from the S3 endpoint region. This command does not list the bucket Owner DisplayName.
<code>ListObjectsV2</code>	<ul style="list-style-type: none"> It returns owner information even if you do not pass the <code>fetch owner</code> parameter. It returns Key Count as 0 instead of 1 when the prefix does not have an object. Using two consecutive slashes (//) in a prefix path results in an incorrect sequence of objects in a bucket. Returns Key Count as 0 instead of 1 when the prefix does not have an object.
<code>PutObjectRetention</code>	It returns an HTTP 204 message on successful execution instead of HTTP 200.
<code>HeadObject</code>	Returns incorrect HTTP code if an object does not exist.

S3 API	Description
<code>CompleteMultipartUpload</code>	The Entity tag (ETag) value of an object received in response does not match with AWS S3.
<code>Putobject</code>	<ul style="list-style-type: none"> Uploading an object where the object key contains a new line character (0x0A) returns 503 as an incorrect error code. Returns a different error message, whereas AWS returns a 400 error code when you specify the following parameters: <ul style="list-style-type: none"> <code>--content-length</code>: Greater than the actual content length, Lyve Cloud does not return any error code. <code>--sse-customer-key</code>: Less than or greater than 32 bytes, LyveCloud returns a 403 error code. <code>--object-lock-legal-hold-status</code>: A random string is passed, Lyve Cloud allows uploading objects despite incorrect <code>--object-lock-legal-hold-status</code>. <code>--tagging</code>: For multiple tagset with invalid separators, Lyve Cloud allows multiple tagset with invalid separators, but it splits the value from first=. <code>--metadata</code>: LyveCloud returns capitalized keys when you pass lowercase letters.
<code>UploadPartCopy</code>	Using "/" (slash) as a prefix to the source bucket returns a 403 error. For example, to copy the object <code>logs/january.txt</code> from bucket <code>bucket1</code> , use the header <code>x-amz-copy-source path: bucket1/logs/january.txt</code> instead of the <code>x-amz-copy-source path: /bucket1/logs/january.txt</code> .
<code>GetObject</code>	This command returns a content range of the first part number irrespective of the part number specified of the object.



Note—In Lyve Cloud, the custom metadata is stored in a different sequence than AWS S3. An object with a prefix that matches the name of another object is not supported. For example, an object named `/A/B` (where **A** is a prefix and **B** is the object name) is not supported if its prefix matches another object, `/A`.