



Cloud Import Service User Manual & Reference Guide



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 Lyve Mobile with Cloud Import	5
Cloud Import Process Overview	5
Security and Lyve Mobile with Cloud Import	5
Key terms	5
2 IP Address Access	7
Required IP addresses	7
3 File Naming Guidelines	8
4 File Size Limitations	11
5 Create a Cloud Import Project	12
Adding a new import destination	14
From the Import Plans page	15
From the Project Details page	15
6 Importazione in Amazon S3	16
Prerequisiti	16
• Suggestioni	16
• Esempio di criterio di autorizzazione IAM Amazon	17
Completamento dei prerequisiti	17
• Raccolta di informazioni	17
• Creare un criterio di autorizzazione IAM per il bucket.	17
• Creare un ruolo IAM che riconosce il servizio di importazione Lyve come attendibile.	18
Configurazione del piano di importazione	19
• Verifica dei dati di accesso	19
• Verifica e inoltro del piano di importazione	21
• Come invitare un altro utente a configurare un piano di importazione	21
Linee guida per la denominazione	21
Procedure consigliate	23
Risoluzione dei problemi	23
7 Import to Google Cloud Storage	24
Prerequisites	24
• Seagate authorizations	24
• Recommendations	26
Configure your import plan	27
• Enter credentials	27
• Review and submit your import plan	28
• Inviting another user to configure an import plan	28
Naming guidelines	28
Best practices	29

Troubleshooting	29
8 Import to IBM Cloud	30
Prerequisites	30
• Seagate authorizations	30
• Recommendations	32
Configure your import plan	33
• Enter credentials	33
• Review and submit your import plan	34
• Inviting another user to configure an import plan	34
Naming guidelines	34
Troubleshooting	35
9 Import to Microsoft Azure Blob Storage	36
Prerequisites	36
• Recommendations	36
Configure your import plan	36
• Enter credentials	36
• Review and submit your import plan	38
• Inviting another user to configure an import plan	38
Naming guidelines	38
Best practices	39
Troubleshooting	39
10 .Import to OVHcloud	40
Prerequisites	40
• Recommendations	40
Configure your import plan	41
• Enter credentials	41
• Review and submit your import plan	42
• Inviting another user to configure an import plan	42
Naming guidelines	42
Best practices	43
Troubleshooting	43
11 .Import to Seagate Lyve Cloud	44
Prerequisites	44
• Recommendations	44
Configure your import plan	44
• Enter credentials	44
• Review and submit your import plan	46
• Inviting another user to configure an import plan	46
Naming guidelines	46
Best practices	48
Troubleshooting	48
12 .Import to Wasabi	49

Prerequisites	49
• Seagate authorizations	49
• Recommendations	50
Configure your import plan	50
• Enter credentials	50
• Review and submit your import plan	52
• Inviting another user to configure an import plan	52
Naming guidelines	52
Troubleshooting	52

13 .Invite Another User to Configure an Import Plan 54

14 .Move Data to a Lyve Mobile Array 55

15 .Send a Lyve Mobile Array to a Seagate Import Site 56

Issues preventing successful validation	57
---	----

16 .Track Import Status 58

17 .Confirm Import Completion 59

Lyve Mobile with Cloud Import

Lyve Mobile with Data Transfer from Seagate® is a high-capacity edge storage solution that enables businesses to aggregate, store, move, and activate their data. Scalable and modular, this integrated solution eliminates network dependencies so you can transfer mass data sets in a fast, secure, and efficient manner. With our new cloud import option, your data can be saved securely on the device and imported to the cloud destination of your choice.

The solutions are delivered as a service—you order and pay only for the devices you need, when you need them. Take a right-sized approach to your data transfer needs with flexible service plan options designed to optimize your budget. Adapt to changing project needs by adjusting your subscription at any time.

Cloud Import Process Overview

1. Sign in to Lyve Management Portal. If you do not have an account, register at lyve.seagate.com.
2. Create a Lyve Mobile with Cloud Import project.
3. Configure the import plan for your project.
4. Move data onto your Lyve Mobile Array(s).
5. Send Mobile Array(s) to a Seagate import site.
6. After completion of the import, verify your files in your cloud destination and confirm the import in Lyve Management Portal.
7. Device(s) are cryptographically erased. A confirmation document detailing the erasure is sent.

Security and Lyve Mobile with Cloud Import

You should always utilize best practices of ensuring encrypted data transfer protocols between Lyve Mobile and your cloud provider. Seagate provides a highly secure data center and network architecture that is built to meet the requirements of most security-sensitive organizations. Third-party agencies also regularly review and test the security of our systems, architecture, and processes. When storing your cloud destination credentials, all your information is transmitted and stored with industry standard encryption and access can only be requested by your device.

However, ensuring your data is protected is a shared responsibility that requires you to follow your organization's security policies, maintain the sensitivity of your data, and align with applicable laws and regulations.

Key terms

Import destination—An import destination is a cloud and region where your data will be imported to.

Import plan—An import plan is tied to a project and contains the details which Seagate uses to import your data to your specified import destination. These details include credentials required to authenticate

access to your cloud destination's resources and services.

IP Address Access

If a firewall or IP restrictions are configured by your organization, you must list Seagate's Cloud Import services' IP address(es) as an allowed source.

Required IP addresses

i **Important**—If these IP addresses are not listed as allowed sources, Seagate cannot import your data.

Region	IP address(es) to allow
North America	20.253.219.114 38.100.210.253 38.104.105.74 67.200.249.173
Europe	91.242.219.6 149.5.114.11 185.212.46.0/29 193.242.211.16/28

File Naming Guidelines

Seagate follows general S3 file naming conventions.

! Folder names cannot contain forward slash / characters.

Safe characters	
Alphanumeric characters	
0-9	numerals
a-z	lowercase letters
A-Z	uppercase letters
Special characters	
*	asterisk
!	exclamation point
-	hyphen
(parenthesis (open)
)	parenthesis (close)
.	period
'	single quote
_	underscore

Characters to avoid	
&	ampersand
	ASCII characters <ul style="list-style-type: none">• ASCII ranges 00-1F hex (0-31 decimal) and 7F (127 decimal)• non-printable ASCII (128-255 decimal characters)

@	at sign
\	backslash
^	caret
:	colon
,	comma
{	curly brace (left)
}	curly brace (right)
\$	dollar sign
=	equal sign
/	forward slash
`	grave
<	greater-than symbol
>	less-than symbol
%	percent sign
	pipe or vertical bar
+	plus sign
#	pound character
?	question mark
"	quotation mark
;	semi-colon
	space - sequences with spaces, especially multiple spaces, may be lost
[square bracket (left)
]	square bracket (right)



Be sure to check the file naming guidelines for your specific cloud destination:

- [Naming guidelines for Amazon S3](#)
- [Naming guidelines for Google Cloud Storage](#)
- [Naming guidelines for IBM Cloud](#)
- [Naming guidelines for Microsoft Azure Blob Storage](#)
- [Naming guidelines for OVHcloud](#)
- [Naming guidelines for Seagate Lyve Cloud](#)
- [Naming guidelines for Wasabi](#)

File Size Limitations


In general, Seagate does not recommend importing individual files larger than 5TB. Please refer to your cloud providers file size limitations and best practices.

Create a Cloud Import Project



You must be a master account user or a sub user to create a project.

1. Go to lyve.seagate.com and sign in. Enter a verification code to continue to Lyve Management Portal.
2. On the Lyve Services page, click **Get Started** on the Lyve Mobile with Cloud Import card.



Lyve Mobile with Cloud Import
Simple, secure, efficient data upload to any major public S3 cloud destination. Capture your data in the field, send the Mobile Array to Seagate, and easily track the status of your import.

[Get Started →](#)





3. On the Service Plan page, review the Lyve service selected.
4. Review the associated rate plan selection for this service.



A one-time import service charge is applied to each storage device in a project with a payment due at checkout.

5. Check the box to acknowledge that you have read and understand our [Cloud Import Service User Manual & Reference Guide](#). Click **Continue**.
6. Add products to your subscription by clicking the Plus (+) icon or entering a quantity to the field below a product.

Lyve Devices:

			
Lyve Mobile Array (96TB HDD)	Lyve Mobile Array (46TB SSD)	Lyve Mobile Array (60TB HDD)	Lyve Mobile Array (92TB SSD)
Product Details ↗	Product Details ↗	Product Details ↗	Product Details ↗
− 1 +	− 0 +	− 0 +	− 0 +

Click **Continue**.

7. Select the RAID level for each Mobile Array in your project. RAID options are RAID 0 and RAID 5 (default).

Select Configuration

Please select the RAID level for the Lyve Mobile Array. Configuration settings only apply to Lyve Mobile Array products. Any Lyve Mobile accessories or shuttles are not configurable.

Device	RAID Level
 Lyve Mobile Array (96 TB HDD)	RAID 0



Configuration settings are only displayed for configurable devices.

Click **Continue**.

8. Fill in project details:

- Project name
- Project start and end date
- Contact information
- Shipping information

Click **Continue**.

9. Select an import destination from the table. If there are no previous import destinations listed, or you need to add a new one, see [Adding a new import destination](#) below.



All devices within a project use the same import destination. If you want some devices to use another cloud or region, you'll need to create a separate project.

Ensure that the correct import destination in the table is selected and that the information is accurate. Click **Continue**.



Adding a new import destination

Click **Add Import Destination** and enter the following destination details:

- Import destination name and (optionally) description.
- Cloud destination and region.

Each import destination must have a unique name, along with a unique cloud and region combination. You will configure the bucket/container and credentials later.

Click **Save Import Destination**.

10. Review your service plan, project details, shipping details, import destination, and devices. If you have a valid promo code, enter it in the order summary section.
11. Submit your order or request a quote from Seagate.



To receive a quote, click on the link at the bottom of the page.

12. Once you submit your cloud import project, you'll see a confirmation page. Click the **Go to Plans** button to start your plan configuration. **Note**—It may take a minute for your project to appear on the Import Plans page.) If you click **Back to Projects**, you can complete your import plan configuration at a later time. An import plan will remain in draft status until the details are configured and submitted.



You'll only see the confirmation screen above if you **submit a project**. If you **request a quote** or **register a deal**, you can configure your import plan after the quote/deal has been approved. There are two other ways to start configuring your import plan:

From the Import Plans page

1. Go to lyve.seagate.com and sign in. Enter a verification code to continue to Lyve Management Portal.
2. On the Lyve Services page, go to the Lyve Mobile with Cloud Import card. Click on the **Manage Projects** dropdown and select **View Plans**.
3. In the Import Plans table, locate the project that you would like to configure the plan for and click the **Configure Plan** button in the Actions column.

Note—If you requested a quote or submitted a deal, you will be able to configure the plan once your project is approved and active.

From the Project Details page

1. Go to lyve.seagate.com and sign in. Enter a verification code to continue to Lyve Management Portal.
2. On the Lyve Services page, go to the Lyve Mobile with Cloud Import card. Click on the **Manage Projects** dropdown and select **View Projects**.
3. Click on the cloud import project that you would like to configure a plan for.
4. Click **Configure Plan**.

For details on configuring your import plan for your specific cloud destination, see the following:

- [Import to Amazon S3](#)
- [Import to Google Cloud Storage](#)
- [Import to IBM Cloud](#)
- [Import to Microsoft Azure Blob Storage](#)
- [Import to OVHcloud](#)
- [Import to Seagate Lyve Cloud](#)
- [Import to Wasabi](#)

Importazione in Amazon S3

Prerequisiti

Per poter configurare e inviare il piano di importazione, è necessario completare i seguenti passaggi in modo che il servizio di importazione Lyve possa accedere in modo sicuro al bucket Amazon S3 specificato per importare i dati.

Abbonamento AWS: configurare un [account AWS](#).

Bucket Amazon S3: configurare un bucket dedicato per l'importazione. Per ulteriori informazioni, vedere [Creazione di un bucket](#).

Autorizzazioni di Seagate: creare un ruolo IAM e i relativi criteri. Per ulteriori informazioni, vedere [Concessione dell'accesso ad account AWS appartenenti a terzi](#)

Per eseguire l'importazione, Seagate **richiede** le seguenti autorizzazioni:

- s3:AbortMultipartUpload
- s3:CreateBucket
- s3>DeleteObject
- s3:GetAccelerateConfiguration
- s3:GetBucketLocation
- s3:GetObject
- s3:GetObjectAttributes
- s3:ListBucket
- s3:ListBucketMultipartUploads
- s3:ListMultipartUploadParts
- s3:PutObject

! **Importante:** la mancata concessione a Seagate delle autorizzazioni di cui sopra comporterà l'esito negativo del piano di importazione.

Suggerimenti

Seagate raccomanda **fortemente** di attenersi alle seguenti procedure consigliate:

- Creare un bucket specifico destinato al progetto di importazione Lyve.
- Bloccare ogni accesso pubblico al bucket.
- Verificare che il controllo delle versioni del bucket sia disabilitato.
- Verificare che la crittografia su server sia abilitata.

- Creare un criterio di autorizzazione IAM.
- Creare un ruolo IAM che riconosce il servizio di importazione Lyve come attendibile, associando il criterio IAM creato.
- Disabilitare o eliminare il ruolo una volta portato a termine il progetto di importazione sul cloud.
- Disabilitare o eliminare il criterio una volta portato a termine il progetto di importazione sul cloud.









Esempio di criterio di autorizzazione IAM Amazon

```
{ "Version": "2012-10-17", "Statement": [ { "Sid": "LyveMobilePolicyTemplate", "Effect": "Allow", "Action": [ "s3:PutObject", "s3:GetObject", "s3:ListBucketMultipartUploads", "s3:AbortMultipartUpload", "s3:GetObjectAttributes", "s3:CreateBucket", "s3:ListBucket", "s3:GetAccelerateConfiguration", "s3>DeleteObject", "s3:GetBucketLocation", "s3:ListMultipartUploadParts" ], "Resource": [ "arn:aws:s3:::[bucketname]", "arn:aws:s3:::[bucketname]/*" ] } ] }
```

Completamento dei prerequisiti

Raccolta di informazioni

1. Fare clic su **Import Plans** (Piani di importazione) nella barra di navigazione. Nella tabella dei piani di importazione, individuare il progetto per cui si desidera configurare il piano e fare clic sul pulsante **Configure Plan** (Configura piano) nella colonna delle azioni. **Nota:** se è stato richiesto un preventivo o inoltrato un accordo, sarà possibile configurare il piano una volta che il progetto è approvato e attivo.

Project	Devices	Cloud Destination	Action
✓ 	 2	 Amazon S3 ⓘ	
✓ 	 1	 Amazon S3 ⓘ	

Verrà visualizzata la procedura di configurazione guidata del piano:

2. Copiare le informazioni richieste, prendendo nota dell'**ID dell'account S3 di Lyve** e dell'**ID esterno**, in quanto saranno necessari per i passaggi successivi. Se è necessario aggiornare prima la destinazione di importazione, fare clic sul passaggio **Import Destination** (Importa destinazione) e modificare.

Creare un criterio di autorizzazione IAM per il bucket.

1. Accedere alla console AWS.
2. Accedere al servizio IAM.

3. Selezionare i criteri.

Fare clic sul pulsante **Create policy** (Crea criterio).

4. Fare clic sulla scheda **JSON**.

5. Copiare lo script JSON riportato di seguito:

```
{ "Version": "2012-10-17", "Statement": [ { "Sid": "LyveMobilePolicyTemplate", "Effect": "Allow", "Action": [ "s3:PutObject", "s3:GetObject", "s3:ListBucketMultipartUploads", "s3:AbortMultipartUpload", "s3:GetObjectAttributes", "s3:CreateBucket", "s3:ListBucket", "s3:GetAccelerateConfiguration", "s3>DeleteObject", "s3:GetBucketLocation", "s3:ListMultipartUploadParts" ], "Resource": [ "arn:aws:s3:::{bucketname}", "arn:aws:s3:::{bucketname}/*" ] } ] }
```

6. Incollare il testo copiato nell'editor JSON.

7. Sostituire **{bucketname}** con il nome del bucket in cui si desidera importare i dati.

8. Fare clic sul pulsante **Next: Tags** (Avanti - Tag).

9. Aggiungere i tag (facoltativo) e fare clic sul pulsante **Next: Review** (Avanti - Verifica).

10. Nella pagina **Review policy** (Esame criterio) assegnare il nome **CriterioAccessoLyveMobile** al criterio.

11. Fare clic su **Create policy** (Crea criterio).

Creare un ruolo IAM che riconosce il servizio di importazione Lyve come attendibile.

1. Nella barra laterale fare clic su **Roles** (Ruoli). Fare clic sul pulsante **Create role** (Crea ruolo).

2. Nella pagina di **selezione dell'entità ritenuta attendibile**, selezionare **Custom trust policy** (Criteri di attendibilità personalizzati).

3. Copiare il criterio di attendibilità fornito di seguito:

```
{ "Version": "2012-10-17", "Statement": [ { "Effect": "Allow", "Principal": { "AWS": "arn:aws:iam::{accountid}:root" }, "Action": "sts:AssumeRole", "Condition": { "ForAnyValue:StringEqualsIfExists": { "sts:ExternalId": [ "{externalid}" ] } } ] }
```

4. Incollare il testo copiato nell'editor JSON.
5. Sostituire `{accountid}` con il valore copiato per l'ID account S3 di Lyve. Sostituire `{externalid}` con il valore copiato per l'ID esterno..
6. Nella pagina **Add permissions** (Aggiungi autorizzazioni), aggiungere il criterio **CriterioAccessoLyveMobile** creato in precedenza e fare clic su **Next** (Avanti).

Se si devono configurare più piani di importazione, aggiungere l'ID esterno di ciascun piano separato da una virgola (,). Ad esempio:

```
{ "Version": "2012-10-17", "Statement": [ { "Effect": "Allow", "Principal": { "AWS": "arn:aws:iam::{accountid}:root" }, "Action": "sts:AssumeRole", "Condition": { "ForAnyValue:StringEqualsIfExists": { "sts:ExternalId": [ "{firstexternalid}", "{secondexternalid}", "{thirdexternalid}" ] } } ] }
```

7. Fare clic su **Next** (Avanti) per uscire dall'editor JSON.
8. Nella pagina **Add permissions** (Aggiungi autorizzazioni), aggiungere il criterio **CriterioAccessoLyveMobile** creato in precedenza e fare clic su **Next** (Avanti).
9. Nella pagina **Name, review, and create** (Assegna nome, esamina e crea), immettere il nome del ruolo in **Role name**, ad esempio `RuoloAccessoLyveMobile`.
10. Verificare le informazioni immesse per l'entità di fiducia e per le autorizzazioni:

Verificare quanto segue:

- A. che per AWS sia stato immesso il valore copiato dall'ID dell'account S3 di Lyve;
- B. che per "sts:ExternalID" sia stato immesso il valore copiato dall'ID esterno;
- C. in **Policy name** (Nome criterio) sia stato immesso il nome del criterio creato in precedenza (**CriterioAccessoLyveMobile**).

Configurazione del piano di importazione



Una volta completati i prerequisiti, continuare con il passaggio successivo per immettere i dati di accesso.



Occorre convalidare i dati di accesso e inviare il piano prima di poter scaricare le etichette per la spedizione del reso.

Verifica dei dati di accesso

1. In Lyve Management Portal fare clic su **Import Plans (Importa piani)**. Fare clic su **Configure Plan (Configura piano)** accanto alla destinazione cloud Amazon S3.

Project	Devices	Cloud Destination	Action
✓ 113	⚠ 2	 Amazon S3 ⓘ	Configure Plan
✓ 113	⚠ 1	 Amazon S3 ⓘ	Configure Plan

2. Assicurarsi di aver completato i prerequisiti e fare clic su **Next (Avanti)**.

3. Immettere i dati di accesso:

- A. Immettere l'ID dell'account AWS.
- B. Immettere il nome del bucket. Il nome deve corrispondere esattamente, inclusa la distinzione tra lettere maiuscole e minuscole.
- C. Confermare il nome del ruolo IAM creato nei passaggi preliminari.
- D. Confermare l'ID esterno creato nei passaggi preliminari.
- E. (Facoltativo) Immettere un nome per la cartella.



A ogni dispositivo di memorizzazione nel progetto verrà destinata una cartella nel bucket. Il numero di serie del dispositivo verrà aggiunto automaticamente al nome della cartella al momento dell'importazione.

- Se si specifica un nome, Seagate lo utilizzerà per creare la cartella o le cartelle nel bucket per conto dell'utente. **(Consigliato)**
- Se si lascia vuoto questo campo, Seagate creerà una o più cartelle per i file e utilizzerà il numero di serie del dispositivo come nome.
- In alternativa, se nel bucket è già presente una cartella in cui si vuole importare i file, specificare il nome della cartella
- **Importante:** assicurarsi che il criterio dei bucket non blocchi la creazione delle cartelle. Se si specifica un nome per una nuova cartella da creare, assicurarsi che il nome segua le [linee guida per la denominazione](#).

4. Fare clic su **Validate (Convalida)**.



Se la convalida ha esito negativo, verificare l'ID dell'account AWS, il nome del bucket, il criterio di autorizzazione IAM e il ruolo IAM, quindi provare nuovamente a convalidare.

5. Fare clic su **Next** (Avanti).

Verifica e inoltro del piano di importazione

1. Verificare la destinazione di importazione e i dati di accesso. Se è necessario apportare modifiche, fare clic sull'icona di modifica.
2. Selezionare la casella per confermare di aver letto e compreso le informazioni contenute in questa guida di riferimento.
3. Fare clic su **Submit Plan** (Invia piano).

Come invitare un altro utente a configurare un piano di importazione

Se un altro membro dell'organizzazione deve configurare il piano di importazione per un progetto, è possibile invitarlo a farlo in Lyve Management Portal. Vedere [Invitare un altro utente a configurare un piano di importazione](#).

Linee guida per la denominazione



I nomi delle cartelle non possono contenere barre (/).

Caratteri sicuri	
Caratteri alfanumerici	
0-9	numeri
a-z	lettere minuscole
A-Z	lettere maiuscole
Caratteri speciali	
*	asterisco
!	punto esclamativo

-	trattino
(parentesi (aperta)
)	parentesi (chiusa)
.	punto
'	virgoletta singola
_	sottolineatura

Caratteri da evitare

&	e commerciale
	<p>Caratteri ASCII</p> <ul style="list-style-type: none"> • Intervalli ASCII 00-1F esadecimali (0-31 decimali) e 7F (127 decimali) • ASCII non stampabili (128-255 caratteri decimali)
@	chiocciola
\	barra rovesciata
^	accento circonflesso
:	due punti
,	virgola
{	parentesi graffa (sinistra)
}	parentesi graffa (destra)
\$	simbolo del dollaro
=	segno di uguale
/	barra
`	accento grave
<	simbolo maggiore di
>	simbolo minore di
%	segno di percentuale

	pipe o barra verticale
+	segno più
#	cancelletto
?	punto interrogativo
"	virgolette doppie
;	punto e virgola
	spazio (le sequenze con spazi, soprattutto con più spazi, potrebbero andare perdute)
[parentesi quadra (sinistra)
]	parentesi quadra (destra)

Procedure consigliate

Consultare questi articoli della knowledge base:

- [Procedure consigliate per la gestione delle chiavi di accesso AWS](#)
- [Procedure di sicurezza consigliate per Amazon S3](#)
- [Procedure consigliate per il controllo degli accessi](#)
- [Creazione di backup di Amazon S3](#)
- [Ripristino di dati S3](#)

Risoluzione dei problemi

Consultare questo articolo della knowledge base:

- [Risoluzione dei problemi](#)

Import to Google Cloud Storage

Prerequisites

Before you can configure and submit your import plan, make sure to complete the following steps so that Lyve Import Service can securely access your specified Google Cloud Storage bucket to import your data:

Google Cloud subscription—Set up an [Google Cloud account](#).

Google Cloud project—Set up a Google Cloud project. To learn more, see [Creating and managing projects](#). Note—Make sure that billing is enabled for your Cloud project. To learn more, see [Verify the billing status of your projects](#).

Google Cloud Storage bucket—Set up a dedicated bucket for your import. To learn more, see [Create buckets](#).

IP address access—If configured by your organization, list Seagate's IP address(es) as an allowed source. See [IP Address Access](#).

Seagate authorizations—See below.

Seagate authorizations

Seagate requires permissions to read, write, and list to your bucket to perform the import. Hash-based message authentication code (HMAC) keys using an access ID and secret are required to authenticate requests to your cloud resources. To generate the HMAC keys, follow the steps below after creating your bucket:

1. Using the Google Cloud console, go to the Cloud Storage **Buckets** page and click **Settings**.
2. Click the **Interoperability** tab. Click **Create A Key For A Service Account**.

The Interoperability API allows Google Cloud Storage to interoperate with tools written for other cloud storage systems. This enables you to [run migrations to Cloud Storage](#) and to authenticate both user and service accounts using keyed-hash message authentication codes (HMAC). [Learn more](#)

Request endpoint

Make sure the request endpoint in the tools or libraries you use with other cloud storage systems (e.g., Amazon S3) uses the Cloud Storage URI.

Storage URI
`https://storage.googleapis.com`

Service account HMAC

Use access keys with your organization's Cloud Platform service accounts when you don't want to tie HMAC authentication to specific user accounts. Recommended for production workloads. [Learn more](#)

- Each service account can use up to five keys.
- Note that keys must be deactivated before they can be deleted.
- Grant your service accounts the required permissions for their intended operations – typically this is the IAM Storage Object Admin role.

Access keys for service accounts

This project doesn't have any service account HMAC keys.

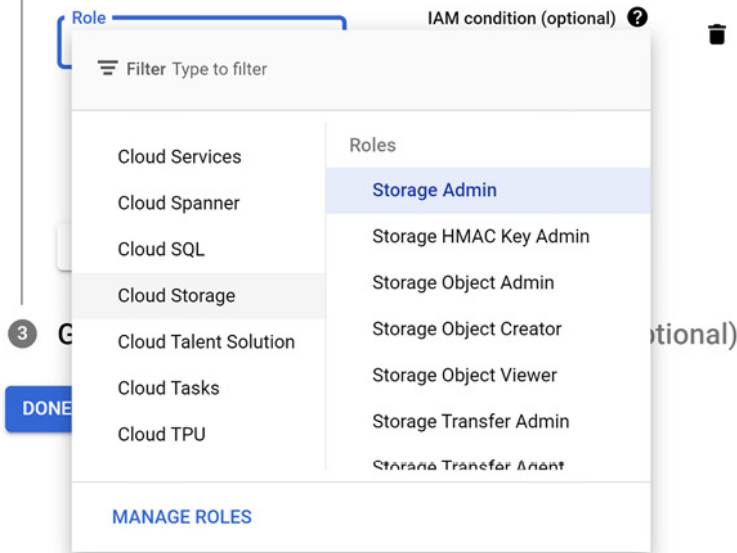
[+ CREATE A KEY FOR A SERVICE ACCOUNT](#)

3. Select the service account you want the HMAC key to be associated with, or click **Create New Account** to create a new service account.
4. If creating a new service account, select **Storage Admin** for the role.

✓ **Service account details**

2 **Grant this service account access to project (optional)**

Grant this service account access to jr-project so that it has permission to complete specific actions on the resources in your project. [Learn more](#)



5. Add an IAM condition with the following selections:

- **Condition type** = Type
- **Operator** = is
- **Resource Type** = storage.googleapis.com/Bucket.

Click **Save**.

6. Record the service account HMAC key.
7. Navigate to the Cloud Storage **Buckets** page and locate the bucket to which you want to assign access for your import. Click the Bucket overflow menu χ and select **Edit Access**.
8. Click **Add Principal**.
9. Enter the email address of the service account the HMAC keys are associated with **Note**—You can find the service account email in the IAM console.
10. Select the **Storage Admin** role and click **Save**.

i To learn more, see [HMAC keys](#).

Recommendations

Seagate **strongly** recommends the following best practices:

- Create a bucket dedicated to your Lyve Import project.
- When creating your bucket, select “Region” for location type.
- Block all public access for your bucket.
- Disable or delete the HMAC key after the cloud import project has ended.

i **Important note on file sizes**—In general, Seagate does not recommend importing individual files larger than 5TB. Please refer to your cloud provider’s file size limitations and best practices.

Configure your import plan

Enter credentials

1. In Lyve Management Portal, click **Import Plans** in the navigation bar, and then click **Configure Plan**.
2. Enter your Access ID and Secret. (If you need to update the import destination first, click on the **Import Destinations** step and edit.)

i Characters you enter in these fields are always masked.

3. Enter your bucket name. The name is case-sensitive and must match exactly.
4. (Optional) Enter a name for your folder.

i Each storage device in your project will have a designated folder in your bucket. The device’s serial number will be automatically appended to the folder name at the time of import.

- Provide a name for Seagate to use to create the folder(s) in your bucket on your behalf. **(Recommended)**
- If you leave this field blank, Seagate will create a folder(s) for your files and will use the device’s serial number as its name.
- Alternatively, if you have an existing folder within your bucket that you would like to import your files to, provide the name of this folder.
- **Important**—Make sure that your bucket policy does not block folder creation. If you are providing a name for a new folder to be created, ensure that the name follows the Naming Guidelines.

5. Validate Credentials.



If the validation fails, check that the access ID, secret, and bucket name entered are accurate, and then revalidate.

6. Check the box to confirm that you have read and understand the details in [IP Address Access](#).
7. Click **Next**.

Review and submit your import plan

1. Review your import destination and credential details.
2. Check the box to confirm that you've read and understand the information in this reference guide.
3. Click **Submit Plan**.

Inviting another user to configure an import plan

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so in Lyve Management Portal. See [Invite Another User to Configure an Import Plan](#).

Naming guidelines

Bucket naming guidelines:

- Bucket names can only contain lowercase letters, numeric characters, dashes, underscores, and dots. Spaces are not allowed. Names containing dots require verification.
- Bucket names must start and end with a number or letter.
- Bucket names must contain 3-63 characters. Names containing dots can contain up to 222 characters, but each dot-separated component can be no longer than 63 characters.
- Bucket names cannot be represented as an IP address in dotted-decimal notation (for example, 192.168.5.4).
- Bucket names cannot begin with the `goog` prefix.
- Bucket names cannot contain `google` or close misspellings, such as `g00gle`.

Object naming guidelines:

- Object names can contain any sequence of valid Unicode characters, of length 1-1024 bytes when UTF-8 encoded.
- Object names cannot contain [Carriage Return or Line Feed characters](#).
- Object names cannot start with `.well-known/acme-challenge/`.
- Objects cannot be named `|` or `..`.

Avoid the Following in Object Names:

- [Control characters](#) that are illegal in XML 1.0 (`#x7F-#x84` and `#x86-#x9F`): these characters cause XML listing issues when you try to list your objects.
- The `#` character: Google Cloud CLI commands interpret object names ending with `<numeric string>` as version identifiers, so including `#` in object names can make it difficult or impossible to perform operations on such versioned objects using the `gcloud` CLI.
- The `[`, `]`, `*`, or `?` characters: `gcloud` storage and `gsutil` interpret these characters as wildcards, so

including them in object names can make it difficult or impossible to perform wildcard operations with those tools.

- Sensitive or personally identifiable information (PII): object names are more broadly visible than object data. For example, object names appear in URLs for the object and when listing objects in a bucket.

To learn more, see [Object Naming Requirements](#).

Best practices

See the following knowledge base article:

- [Best Practices for Cloud Storage](#)

Troubleshooting

See the following knowledge base articles:

- [Support](#)
- [Resources](#)

Import to IBM Cloud

Prerequisites

Before you can configure and submit your import plan, make sure to complete the following steps so that Lyve Import Service can securely access your specified IBM Cloud bucket to import your data:

IBM Cloud subscription—Set up an [IBM Cloud Platform account](#).

Object Storage instance—Set up a storage instance. To learn more, see [Choosing a plan and creating an instance](#).

IBM Cloud bucket— Set up a dedicated bucket for your import. To learn more, see [Create some buckets to store your data](#).


IP address access—If configured by your organization, list Seagate’s IP address(es) as an allowed source. See [IP Address Access](#).

Seagate authorizations —See below.

Seagate authorizations

Seagate requires permissions to read, write, and list to your bucket to perform the import. Hash-based message authentication code (HMAC) keys using an access ID and secret are required to authenticate requests to your cloud resources. To generate the HMAC keys, follow the steps below after creating your bucket:

1. In your Object Storage instance, click the **Service credentials** tab.
2. Click the **New Credential** button.



3. Name the credential and make the following selections:
 - **Role** = None
 - **Service ID** = Auto Generated
 - **Include HMAC Credential** = On

Create Credentials ×

Name:
test-credentials

Role:
None

Select Service ID (Optional)
Auto Generated

Include HMAC Credential
 On

Cancel Add

Click the **Add** button. Once added, you can expand the credentials to view the values for the access key ID and secret access key.

i When these credentials are created, the underlying service ID has access to any bucket in your instance (if it was automatically generated). To limit access to a specific bucket or subset of buckets, you will need to edit the access policy of the service ID tied to these credentials.

Proceed through the steps below to edit the access policy for the service ID:

1. Navigate to the IAM console by clicking **Manage > Access (IAM)**. Click **Service IDs** in the side panel. Click on the service ID you want to edit.
2. Under **Access policies**, locate the role with the access policy you want to edit. Click the **Actions** icon and select **Edit**.

Access policies

Based on your assigned role, you can click the role to view or edit the policy.


Service: All Role: All Search Assign access +

Service	Resources	Role	Conditions	Last permit
Cloud Object Storage	serviceInstance string equals Test Storage	Writer		--

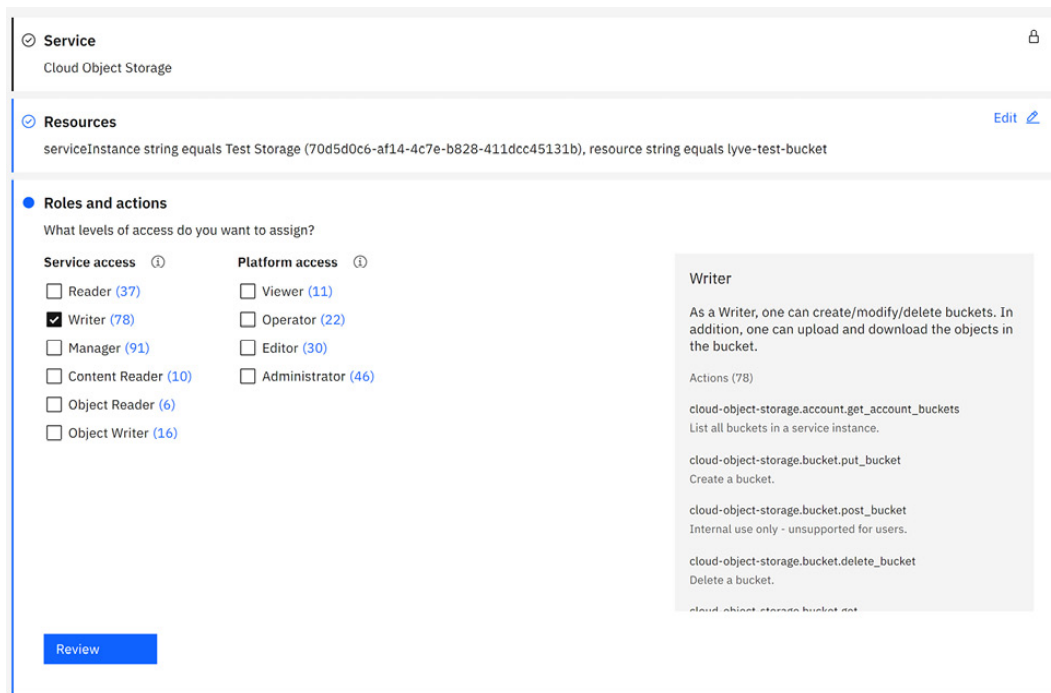
Items per page: 100 1-1 of 1 item 1

- Details
- Edit**
- Remove

3. Click on the **Resources** tab and select **Edit**. Select **Specific resources** and add conditions to scope access to specific resources.



- Click **Next** to continue to the **Roles and actions** tab. In the **Service access** column, assign the **Writer** role. Click **Review**.



- Click **Save**.



To learn more, see [Assigning access to an individual bucket](#)

Recommendations

Seagate **strongly** recommends the following best practices:

- Create a bucket dedicated to your Lyve Import project.
- When creating your bucket, select “Regional” for resiliency.
- Block all public access for your bucket.
- Disable or delete the HMAC key after the cloud import project has ended.

i **Important note on file sizes**—In general, Seagate does not recommend importing individual files larger than 5TB. Please refer to your cloud provider’s file size limitations and best practices.

Configure your import plan

Enter credentials

1. In Lyve Management Portal, click **Import Plans** in the navigation bar, and then click **Configure Plan**.
2. Enter your Access Key ID and Secret Access Key. (If you need to update the import destination first, click on the **Import Destinations** step and edit.)

i Characters you enter in these fields are always masked.

The screenshot displays the 'Enter Credentials' step of the import plan configuration. On the left, a vertical progress bar shows three steps: 'Import Destinations' (completed with a checkmark), 'Enter Credentials' (current step, highlighted in green), and 'Review & Submit' (next step). The main content area is titled 'Access Key ID' and contains the following sections:

- Access Key ID:** A text input field for the Access Key ID.
- Secret Access Key:** A text input field for the Secret Access Key.
- Bucket:** A text input field for the bucket name. Below the field, a note states: "Use an existing bucket/container that is designated for this project. Input must match your bucket/container name exactly and is case-sensitive."
- Folder:** A text input field for the folder name. Below the field, a note states: "Provide the name of an existing folder or provide a name for Seagate to use to create a new folder. Each device in your project will have its own folder and can be identified by the serial number which will be automatically appended to the folder name. See [Naming Guidelines](#)." Below this is an "Optional" section with a "Validate Credentials" button.

At the bottom of the form, there is a checkbox labeled "I have read and understand the following information:" followed by a link to the "IP Address Access Guide". Below the checkbox are two buttons: "Next" and "Cancel".

3. Enter your bucket name. The name is case-sensitive and must match exactly.
4. (Optional) Enter a name for your folder.



Each storage device in your project will have a designated folder in your bucket. The device's serial number will be automatically appended to the folder name at the time of import.

- Provide a name for Seagate to use to create the folder(s) in your bucket on your behalf. **(Recommended)**
- If you leave this field blank, Seagate will create a folder(s) for your files and will use the device's serial number as its name.
- Alternatively, if you have an existing folder within your bucket that you would like to import your files to, provide the name of this folder.
- **Important**—Make sure that your bucket policy does not block folder creation. If you are providing a name for a new folder to be created, ensure that the name follows the Naming Guidelines.

5. Click **Validate Credentials**.



If the validation fails, check that the access ID, secret, and bucket name entered are accurate, and then revalidate.

6. Check the box to confirm that you have read and understand the details in [IP Address Access](#)
7. Click **Next**.

Review and submit your import plan

1. Review your import destination and credential details.
2. Check the box to confirm that you've read and understand the information in this reference guide.
3. Click **Submit Plan**.

Inviting another user to configure an import plan

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so in Lyve Management Portal. See [Invite Another User to Configure an Import Plan](#).

Naming guidelines

Bucket naming guidelines:

- Must be unique across the whole IBM Cloud Object Storage system.
- Do not use any personal information (any part of a name, address, financial or security accounts or SSN)
- Must start and end in alphanumeric characters (3 to 63)
- Characters allowed: lowercase, numbers and nonconsecutive dots and hyphens
- Avoid using these characters: `/ \ " ? < > |`. This will not cause issues with IBM Cloud Object Storage but may cause issues with your applications.

Object naming guidelines:

- Object keys can be up to 1024 characters in length, and it's best to avoid any characters that might be problematic in a web address. For example, `?`, `=`, `<`, and other special characters might cause unwanted behavior if not URL-encoded.

Troubleshooting

See the following knowledge base articles:

- [FAQ](#)
- [Support](#)

Import to Microsoft Azure Blob Storage

Prerequisites

Before you can configure and submit your import plan, make sure to complete the following steps so that Lyve Import Service can securely access your specified Azure container to import your data:

Azure subscription—Set up an [Azure free account](#).

Azure storage account—Set up an Azure storage account. To learn more, see [Create an Azure storage account](#).

Azure container—Set up a dedicated container for your import. To learn more, see [Create a container](#).

Seagate authorizations—Ensure that Seagate is authorized to read, write, and list to an existing container.

IP address access—If configured by your organization, list Seagate's IP address(es) as an allowed source. See [IP Address Access](#).

Additionally, see [How to configure the Azure Storage Firewall](#)

Recommendations

Seagate recommends creating a container dedicated to your Lyve project.



Important note on file sizes—In general, Seagate does not recommend importing individual files larger than 5TB. Please refer to your cloud provider's file size limitations and best practices.

Configure your import plan

Enter credentials


The storage account name and storage account key are required to authenticate requests to your cloud resources. Access keys can be found in the Azure portal under the Security + networking tab in your storage account.

1. In Lyve Management Portal, click **Import Plans** in the navigation bar, and then click **Configure Plan**.

1. In Lyve Management Portal, click **Import Plans** in the navigation bar, and then click **Configure Plan**.
2. Enter your storage account name and storage account key. (If you need to update the import destination first, click on the Import Destination step and edit.)



Characters you enter in security-related fields are always masked.



Storage Account Name

Storage Account Key

Container

Use an existing bucket/container that is designated for this project. Input must match your bucket/container name exactly and is case-sensitive. See [Prerequisites](#).

Folder

Provide the name of an existing folder or provide a name for Seagate to use to create a new folder. Each device in your project will have its own folder and can be identified by the serial number which will be automatically appended to the folder name. See [Naming Guidelines](#).

Optional

I have read and understand the following information:
[IP Address Access Guide](#)



Important—If you rotate your account access keys at any point, make sure that the import plan for your Lyve project is updated with the new credentials.

3. Enter your container name.
4. (Optional) Enter a name for your folder.

Each storage device in your project will have a designated folder in your container. The device's serial number will be automatically appended to the folder name at the time of import.



- Provide a name for Seagate to use to create the folder(s) in your container on your behalf. (Recommended)
- If you leave this field blank, Seagate will create a folder(s) for your files and will use the device's serial number as its name.
- Alternatively, if you have an existing folder within your container that you would like to import your files to, provide the name of this folder.
- **Important**—Make sure that your container policy does not block folder creation. If you are providing a name for a new folder to be created, ensure that the name follows the [Naming Guidelines](#).

5. Click **Validate Credentials**.



If the validation fails, check that the storage account name, storage account key, and container entered are accurate, then revalidate. If you have a key expiration policy in place, ensure that the access key provided has not expired. See [Manage storage account access keys](#).

6. Check the box to confirm that you have read and understand the details in [IP Address Access](#).

7. Click **Next**.

Review and submit your import plan

1. Review your import destination and credential details. If you need to make changes to the credentials provided, click the Edit icon.
2. Check the box to confirm that you've read and understand the information in this reference guide.
3. Click **Submit Plan**.

Inviting another user to configure an import plan

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so in Lyve Management Portal. See [Invite Another User to Configure an Import Plan](#).

Naming guidelines

Note the following naming guidelines:

- Every folder within a container must have a unique name.
- A folder name can contain any combination of characters.
- For blobs in Azure Storage, a folder name must be at least one character long and cannot be more than 1,024 characters long.
- Folder names are case-sensitive.
- Reserved URL characters must be properly escaped.
- Avoid folder names that end with a dot . , a forward slash / , or a sequence or combination of the two.

For additional information on naming folders, see [Naming and Referencing Containers, Blobs, and Metadata](#).

Best practices

See the following knowledge base articles:

- [Security recommendations for Blob storage](#)
- [Best practices for monitoring Azure Blob Storage](#)

Troubleshooting

See the following knowledge base articles:

- [Monitor, diagnose, and troubleshoot Microsoft Azure Storage](#)
- [Troubleshoot Azure RBAC](#)
- [Azure Blob Storage FAQ](#)
- [Microsoft Q&A question page](#)
- [Azure Storage on Stack Overflow](#)

Import to OVHcloud

Prerequisites

Before you can configure and submit your import plan, make sure to complete the following steps so that Lyve Import Service can securely access your specified OVHcloud container to import your data:

OVHcloud subscription—Set up an [OVHcloud account](#).

OVH Public Cloud project—Set up a OVHcloud Public Cloud project. To learn more, see [Creating your first OVHcloud Public Cloud project](#).

OVHcloud container—Set up a dedicated object container for your import. To learn more, see [Object Storage - Creating a bucket](#).

OVHcloud Public Cloud instance—Set up an instance. To learn more, see [Creating an instance](#).

IP address access—If configured by your organization, list Seagate's IP address(es) as an allowed source. See [IP Address Access](#).

Seagate authorizations—Seagate requires permissions to read, write, and list to your container to perform the import. Hash-based message authentication code (HMAC) keys using an access ID and secret are required to authenticate requests to your cloud resources. To generate the HMAC keys, select an existing user or create a new user to link to your container. Set Read and write access to your container for this user. Once the user has been created and added to your container, you will see the credentials.

To learn more, see [Object Storage - Identity and access management](#)

Recommendations

Seagate **strongly** recommends the following best practices:

- Create a container dedicated to your Lyve Import project.
- Block all public access for your container.
- Disable or delete the HMAC key after the cloud import project has ended.



Important note on file sizes—In general, Seagate does not recommend importing individual files larger than 5TB. Please refer to your cloud provider's file size limitations and best practices.

Configure your import plan

Enter credentials

1. In Lyve Management Portal, click **Import Plans** in the navigation bar, and then click **Configure Plan**.
2. Enter your Access Key and Secret Access Key. (If you need to update the import destination first, click on the **Import Destinations** step and edit.)



Characters you enter in these fields are always masked.



Access Key

Secret Access Key

Container

Use an existing bucket/container that is designated for this project. Input must match your bucket/container name exactly and is case-sensitive.

Folder

Provide the name of an existing folder or provide a name for Seagate to use to create a new folder. Each device in your project will have its own folder and can be identified by the serial number which will be automatically appended to the folder name. See [Naming Guidelines](#).

Optional

I have read and understand the following information:

IP Address Access Guide [↗](#)

3. Enter your container name. The name is case-sensitive and must match exactly.
4. (Optional) Enter a name for your folder.

i Each storage device in your project will have a designated folder in your container. The device's serial number will be automatically appended to the folder name at the time of import.

- Provide a name for Seagate to use to create the folder(s) in your container on your behalf. **(Recommended)**
- If you leave this field blank, Seagate will create a folder(s) for your files and will use the device's serial number as its name.
- Alternatively, if you have an existing folder within your container that you would like to import your files to, provide the name of this folder.
- **Important**—Make sure that your container policy does not block folder creation. If you are providing a name for a new folder to be created, ensure that the name follows the Naming Guidelines.

5. Click **Validate Credentials** .

i If the validation fails, check that the access, secret access key, and container name entered are accurate, and then revalidate.

6. Check the box to confirm that you have read and understand the details in [IP Address Access](#).
7. Click **Next**.

Review and submit your import plan

1. Review your import destination and credential details.
2. Check the box to confirm that you've read and understand the information in this reference guide.
3. Click **Submit Plan**.

Inviting another user to configure an import plan

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so in Lyve Management Portal. See [Invite Another User to Configure an Import Plan](#).

Naming guidelines

Container naming guidelines:

- Must be between 3 and 63 characters long.
- Must begin and end with lower case alphanumeric characters (a to z and 0 to 9).
- Must be unique within the same OVHcloud region.
- May contain the following punctuation marks: `!` and `-`.
- Must not contain multiple punctuation marks in a row (`!!` or `--` or `!-` or `--`).
- Must not look like an IP address (for example, 192.168.1.1).

Best practices

See the following knowledge base article:

- [Best practices](#)

Troubleshooting

See the following knowledge base articles:

- [Object Storage - Technical Limitations](#)
- [FAQ Public Cloud OVHcloud](#)

Import to Seagate Lyve Cloud

Prerequisites

Before you can configure and submit your import plan, make sure to complete the following steps so that Lyve Import Service can securely access your specified Lyve Cloud bucket to import your data:

Lyve Cloud account—Work directly with a [Lyve Cloud Expert](#) to create your Lyve Cloud account.

Lyve Cloud bucket—Set up a bucket for your import. To learn more, see [Managing buckets](#).

Bucket permissions—To learn more, see [Managing bucket access permissions](#)

Seagate authorizations—Ensure that Seagate is authorized to read, write, and list to an existing bucket.

Service account—To learn more, see [Managing service accounts](#).

IP address access—If configured by your organization, list Seagate's IP address(es) as an allowed source. See [IP Address Access](#).

Recommendations

Seagate recommends creating a bucket dedicated to your Lyve project.

i **Important note on file sizes**—In general, Seagate does not recommend importing individual files larger than 5TB. Please refer to your cloud provider's file size limitations and best practices.

Configure your import plan


Enter credentials

An access key and secret key are required to authenticate requests to your cloud resources. To learn more, see [Manage service account](#).

1. In Lyve Management Portal, click **Import Plans** in the navigation bar, and then click **Configure Plan**.
2. Enter your access key and secret key. (If you need to update the import destination first, click on the Import Destination step and edit.)



Characters you enter in security-related fields are always masked.



Access Key ID

Access Key ID

Secret Access Key

Secret Access Key

Bucket

Use an existing bucket/container that is designated for this project. Input must match your bucket/container name exactly and is case-sensitive.

See [Prerequisites](#).

Bucket

Folder

Provide the name of an existing folder or provide a name for Seagate to use to create a new folder. Each device in your project will have its own folder and can be identified by the serial number which will be automatically appended to the folder name. See [Naming Guidelines](#).

Folder

Optional

Validate Credentials

I have read and understand the following information:

IP Address Access Guide

Next

Cancel

3. Enter your bucket name.
4. (Optional) Enter a name for your folder.

Each storage device in your project will have a designated folder in your bucket. The device's serial number will be automatically appended to the folder name at the time of import.



- Provide a name for Seagate to use to create the folder(s) in your bucket on your behalf. (Recommended)
- If you leave this field blank, Seagate will create a folder(s) for your files and will use the device's serial number as its name.
- Alternatively, if you have an existing folder within your bucket that you would like to import your files to, provide the name of this folder.
- **Important**—Make sure that your bucket policy does not block folder creation. If you are providing a name for a new folder to be created, ensure that the name follows the [Naming Guidelines](#).

5. Click **Validate Credentials**.



If the validation fails, check that the access key, secret key, and bucket name entered are accurate, then revalidate.

6. Check the box to confirm that you have read and understand the details in [IP Address Access](#).
7. Click **Next**.

Review and submit your import plan

1. Review your import destination and credential details. If you need to make changes to the credentials provided, click the Edit icon.
2. Check the box to confirm that you've read and understand the information in this reference guide.
3. Click **Submit Plan**.

Inviting another user to configure an import plan

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so in Lyve Management Portal. See [Invite Another User to Configure an Import Plan](#).

Naming guidelines

Safe characters	
Alphanumeric characters	
0-9	numerals
a-z	lowercase letters
A-Z	uppercase letters
Special characters	
*	asterisk
!	exclamation point
-	hyphen
(parenthesis (open)
)	parenthesis (close)
.	period
'	single quote
_	underscore

Characters to avoid	
&	ampersand
	ASCII characters <ul style="list-style-type: none"> • ASCII ranges 00–1F hex (0–31 decimal) and 7F (127 decimal) • non-printable ASCII (128–255 decimal characters)
@	at sign
\	backslash
^	caret
:	colon
,	comma
{	curly brace (left)
}	curly brace (right)
\$	dollar sign
=	equal sign
/	forward slash
`	grave
<	greater-than symbol
>	less-than symbol
%	percent sign
	pipe or vertical bar
+	plus sign
#	pound character
?	question mark
"	quotation mark
;	semi-colon
	space - sequences with spaces, especially multiple spaces, may be lost
[square bracket (left)

]	square bracket (right)
---	------------------------

Note the following additional requirements:

- An object name matching a prefix is not supported. For example, an object with the name /A/B, where A is a prefix and B is the object name, should not be imported with another object named A.
- A standalone period . in the prefix folder is not supported.
- A standalone period . as an object name is not supported.

Best practices

See the following knowledge base article:

- [Frequently asked Questions](#)

Troubleshooting

See the following knowledge base articles:

- [Troubleshooting Guide](#)
- [Release Notes](#)

Import to Wasabi

Prerequisites

Before you can configure and submit your import plan, make sure to complete the following steps so that Lyve Import Service can securely access your specified Wasabi bucket to import your data:

Wasabi subscription—Set up a [Wasabi account](#).

Wasabi bucket—Set up a dedicated bucket for your import. To learn more, see [Working with Buckets and Objects](#).

IP address access—If configured by your organization, list Seagate’s IP address(es) as an allowed source. See [IP Address Access](#).

Seagate authorizations—See below.

Seagate authorizations

Seagate requires permissions to read, write, and list to your bucket to perform the import. Hash-based message authentication code (HMAC) keys using an access key ID and secret access key are required to authenticate requests to your cloud resources. To generate the HMAC keys, follow the steps below after creating your bucket:

1. Using Wasabi’s console, create a policy.
2. Copy the provided JSON script below to paste into your policy document:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::yourbucketname",
        "arn:aws:s3:::yourbucketname/*"
      ]
      "Condition": {}
    }
  ]
}
```

3. Replace `yourbucketname` with your actual S3 bucket name. Click **Create Policy**.
4. Create a user with programmatic access and attach the policy you created to this user. To learn more,

see [Creating a User Account and Access Key](#)

5. Record the Access Key and Secret Key that are generated in a safe place.

Recommendations

Seagate **strongly** recommends the following best practices:

- Create a bucket dedicated to your Lyve Import project.
- Block all public access for your bucket.
- Disable or delete the HMAC key after the cloud import project has ended.

i **Important note on file sizes**—In general, Seagate does not recommend importing individual files larger than 5TB. Please refer to your cloud provider's file size limitations and best practices.

Configure your import plan

Enter credentials

1. In Lyve Management Portal, click **Import Plans** in the navigation bar, and then click **Configure Plan**.
2. Enter your Access Key ID and Secret Access Key. (If you need to update the import destination first, click on the **Import Destinations** step and edit.)

i Characters you enter in these fields are always masked.



Access Key ID

Secret Access Key

Bucket

Use an existing bucket/container that is designated for this project. Input must match your bucket/container name exactly and is case-sensitive.

Folder

Provide the name of an existing folder or provide a name for Seagate to use to create a new folder. Each device in your project will have its own folder and can be identified by the serial number which will be automatically appended to the folder name. See [Naming Guidelines](#).

Optional

I have read and understand the following information:
IP Address Access Guide [↗](#)

3. Enter your bucket name. The name is case-sensitive and must match exactly.
4. (Optional) Enter a name for your folder.

i Each storage device in your project will have a designated folder in your bucket. The device's serial number will be automatically appended to the folder name at the time of import.

- Provide a name for Seagate to use to create the folder(s) in your bucket on your behalf. **(Recommended)**
- If you leave this field blank, Seagate will create a folder(s) for your files and will use the device's serial number as its name.
- Alternatively, if you have an existing folder within your bucket that you would like to import your files to, provide the name of this folder.
- **Important**—Make sure that your bucket policy does not block folder creation. If you are providing a name for a new folder to be created, ensure that the name follows the Naming Guidelines.

5. Click **Validate Credentials**.



If the validation fails, check that the access ID, secret, and bucket name entered are accurate, and then revalidate.

6. Check the box to confirm that you have read and understand the details in [IP Address Access](#).
7. Click **Next**.

Review and submit your import plan

1. Review your import destination and credential details.
2. Check the box to confirm that you've read and understand the information in this reference guide.
3. Click **Submit Plan**.

Inviting another user to configure an import plan

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so in Lyve Management Portal. See [Invite Another User to Configure an Import Plan](#).

Naming guidelines

Bucket naming guidelines:

- The name must be unique across all existing bucket names in Wasabi.
- A bucket name must:
 - Be a valid DNS-compliant name
 - Begin with a lowercase letter or number, and
 - Consist of 3 to 63 lowercase letters, numbers, periods, and/or dashes.
 - The name cannot contain underscores, end with a dash, have consecutive periods, or use dashes adjacent to periods.
 - The name cannot be formatted as an IP address (for example, 123.45.678.90).

Characters to avoid:

- % (percent)
- < (less than symbol)
- > (greater than symbol)
- \ (backslash)
- # (pound sign)
- ? (question mark)
- Certain file names may have non-ASCII characters that are 4 byte UTF8 characters (such as emojis). Wasabi does not support these characters and will return a 400 error message to an application that tries to write a file with 4 byte UTF characters in the file name. We recommend renaming the affected files, if possible.

Troubleshooting

See the following knowledge base articles:

- [FAQs](#)
- [Troubleshooting](#)

Invite Another User to Configure an Import Plan

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so in Lyve Management Portal.

1. Go to lyve.seagate.com and sign in. Enter a verification code to continue to Lyve Management Portal.
2. Click on the Account icon in the navigation bar and select **User Management**.
3. Click on the **Quick Action** dropdown menu and select **Add User**. Alternatively, click **Add User** in one of the user role descriptions.
4. Provide details for the new user such as their name and contact information.
5. Select the project(s) that you would like to assign to this user.
6. Select the import destination(s) that this user can view and add to import plans.

The invited user will receive an email with a link to register their account. The user will need to complete the account registration in order to configure the plan.

Move Data to a Lyve Mobile Array



Lyve Mobile Array can be used as direct-attached storage. See the [Lyve Mobile Array user manual](#).



Lyve Mobile Array can also support connections via Fibre Channel, iSCSI and Serial Attached SCSI (SAS) connections using the Lyve Rackmount Receiver. For details, see the [Lyve Rackmount Receiver user manual](#).



For high-speed mobile data transfers, connect Lyve Mobile Array using the Lyve Mobile PCIe Adapter. See the [Lyve Mobile Mount and PCIe Adapter user manual](#) or [Lyve Mobile Mount and PCIe Adapter - Front Loader user manual](#).

Send a Lyve Mobile Array to a Seagate Import Site

Send your Lyve Mobile Array for cloud import after you have completed the following:

- Created a cloud import project in Lyve Management Portal
- Configured the import plan with your cloud service credentials
- Moved data to your Lyve Mobile Array



Sending a Lyve Mobile Array to a Seagate import site is a separate process from other types of hardware returns:

Returning a Lyve Mobile accessory— If you need to return an accessory, see [How do I return a device?](#) in the **Lyve Management Portal User Manual**.

Returning a Lyve Mobile Array for another reason—There may be situations in which you need to return a device for non-import reasons, for example, the device needs to be replaced. Contact your sales representative or customer success manager or, use the Lyve Virtual Assistant in Lyve Management Portal.

1. Go to lyve.seagate.com and sign in. Enter a verification code to continue to Lyve Management Portal.
2. On the Lyve Services page, go to the Lyve Mobile with Cloud Import card. Click on the **Manage Projects** dropdown and select **View Plans**.
3. From the Import Plans table, click on a project to expand the row and view the Mobile Arrays within that project.
4. Find the Mobile Array you would like to send to a Seagate import site and click **Send for Import**. This will initiate a cloud validation process which may take several minutes.



A cloud validation process is initiated. The process may take several minutes.

5. Once the validation is complete, click **Print Shipping Label** to view and print your shipping label in a separate browser tab.

! Issues preventing successful validation

The validation may fail if there were issues with the following:

Security credentials provided in the import plan—You are prompted to edit the credentials and revalidate the import plan. In the Actions column, click the **Edit Credentials** button for the applicable project.

Connecting to the cloud destination—In the Actions column, click the **Troubleshoot** button for the applicable project. Follow the onscreen instructions.

Track Import Status

The status of your import plan(s) can be tracked in Lyve Management Portal.

1. Go to lyve.seagate.com and sign in. Enter a verification code to continue to Lyve Management Portal.
2. On the Lyve Services page, go to the Lyve Mobile with Cloud Import card. Click on the **Manage Projects** dropdown and select **View Plans**.
3. Find your project in the Import Plans table and click to expand it.
4. View the status of the import plan in the Status column. You can also hover over the infographic icons to see more information about a status.



Projects with the orange warning icon contain import plan(s) that require one or more actions to be taken. Click the project to expand it and view the action(s) required.

Once a device has been received at the import site, you can track the status of the import by doing the following:

1. Click on a project to expand it.
2. Click the **Check Import Status** button next to a device.

Import Plans

If a different member of your organization needs to configure the import plan for a project, you can invite them to do so from the User Management tab in the navigation bar.

Search by Project Name	Cloud Destination			
Project	Devices	Cloud Destination	Action	
✓ 11/8	1	Seagate Lyve Cloud ⓘ	Edit Credentials	
✓ demo	3	Seagate Lyve Cloud ⓘ	Edit Credentials	
^ 12345	2	Seagate Lyve Cloud ⓘ	Edit Credentials	
Product Name	Serial Number	Status	Tracking No	Action
Lyve Mobile Array (46TB SSD)	NB2060B2	Completed ⓘ	1ZXXXXXXXXXXXXXXXXX	Check Import Status
Lyve Mobile Array (92TB SSD)	-	Plan Accepted ⓘ	-	

Confirm Import Completion

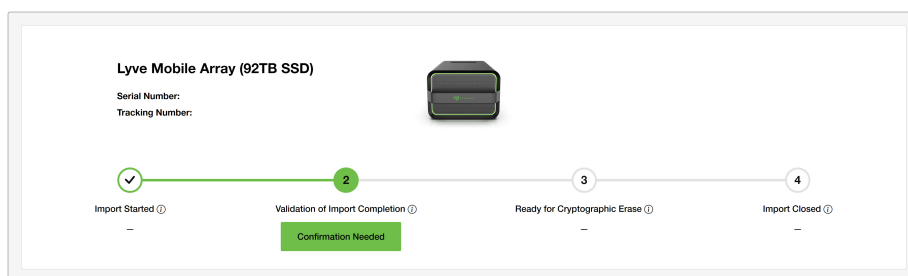
Upon completion of your cloud import, verify that your files have been successfully imported to your cloud destination.

! **Important**—Ensure that all your files have been successfully imported to your cloud destination. **If there's an issue with your import, do not confirm it.** Contact your sales representative or use the [Lyve Support Center](#) to report the issue.


After verifying the files in your cloud destination, confirm the import in Lyve Management Portal.

! **Important**—Confirmation of the import plan is required. Once you confirm the import in Lyve Management Portal, Seagate will purge the AES encryption key used to write data to the drive, making the data irretrievable. This erasure follows NIST SP 800-88 r1 standards.

1. [Go to lyve.seagate.com](https://lyve.seagate.com) and sign in. Enter a verification code to continue to Lyve Management Portal.
2. On the Lyve Services page, go to the Lyve Mobile with Cloud Import card. Click on the **Manage Projects** dropdown and select **View Plans**.
3. Find your project in the Import Plans table and click to expand it.
4. Find the import plan for the applicable device and click **Check Import Status**.
5. Click **Confirmation Needed**.



6. Check the box to confirm your understanding of the device erasure.
7. Click **Confirm Import**.



After the device has been cryptographically erased, Seagate will send a certificate confirming the erasure of the device.