SEAGATE

Lyve Cloud Object Storage Resources Guide

LYVE™ Cloud

Contents

Apostrophe CMS

Apostrophe CMS can be used with Lyve Cloud.

Configure Apostrophe CMS

1. Install Required Modules:

   npm install @aws-sdk/client-s3

2. Configure Lyve Cloud credentials. It's recommended that you use environment variables (.env) to define credentials.

   - LC_ACCESS_KEY_ID=<access-key>
   - LC_SECRET_ACCESS_KEY=<secret-key>
   - LC_ENDPOINT=<s3-endpoint>for examplehttps://s3.us-east-1.lyve.seagate.com
   - LC_BUCKET_NAME=<bucket>

   where:

   - <bucket> is the name of your Lyve Cloud bucket.
   - <access-key> is your access key.
   - <secret-key> is your secret key.
   - <s3-endpoint> is the appropriate Lyve Cloud S3 endpoint URL, for example, https://s3.us-east-1.lyve.seagate.com.

Sample command for listing S3 objects in the bucket

```
require('dotenv').config();
const AWS = require('aws-sdk');

// Configure AWS credentials and region
AWS.config.update({
accessKeyId: process.env.LC_ACCESS_KEY_ID,
secretAccessKey: process.env.LC_SECRET_ACCESS_KEY,
region: process.env.LC_REGION,
endpoint: process.env.LC_ENDPOINT
});

const s3 = new AWS.S3();
// List objects in a specific S3 bucket
s3.listObjectsV2({ Bucket: process.env.LC_BUCKET_NAME }, (err, data) => {
if (err) {
console.error('Error listing objects:', err);
} else {
const objects = data.Contents;
```

```
console.log(objects);
}
});
```

Result example

```
[
  {
    Key: 'Pic-1.png',
    LastModified: 2024-11-05T12:17:38.261Z,
    ETag: '"651b690ab5db93496d9ca412e8d7823a"',
    ChecksumAlgorithm: [],
    Size: 1121955,
    StorageClass: 'STANDARD'
  },
  {
    Key: 'Pic-2.png',
    LastModified: 2024-11-05T12:17:28.824Z,
    ETag: '"03685db85f0533211f4f865768e0100d"',
    ChecksumAlgorithm: [],
    Size: 479053,
    StorageClass: 'STANDARD'
  },
  {
    Key: 'Pic-3.png',
    LastModified: 2024-11-05T12:17:33.578Z,
    ETag: '"7b71081d8b7b21d6ee712692cc0d335e"',
    ChecksumAlgorithm: [],
    Size: 645935,
    StorageClass: 'STANDARD'
  }
]
```

Prerequisites

Before configuration, ensure that the following requirements have been met:

Operating system	Windows Server 2019 or later
	Red Hat Enterprise Linux (RHEL) 8.4 or later
Recommended hardware	
Sofware	Ensure required licenses and packages are available.
Browser compatibility	Chrome recommended.
Region information	Your region information is required to complete the configuration.

Download Commvault software

Navigate to the Commvault Store and download the latest version of Commvault Express for your operating system. Make sure that the downloaded file is the correct release.

Install Commvault

> **i** In the instructions that follow, Commvault is shown being installed on the C: drive. Note, however, that installing on the C: drive is not recommended. Choose another drive in your system for better performance and future scalability.

1. Launch the executable and follow the onscreen instructions to install Commvault. Complete the installation, making sure that all dependencies are resolved.
2. Reboot the system to finalize the setup.
3. After rebooting, launch the Commvault management console and configure the initial settings.
4. The following confirmation is displayed once the installation and configuration are completed:

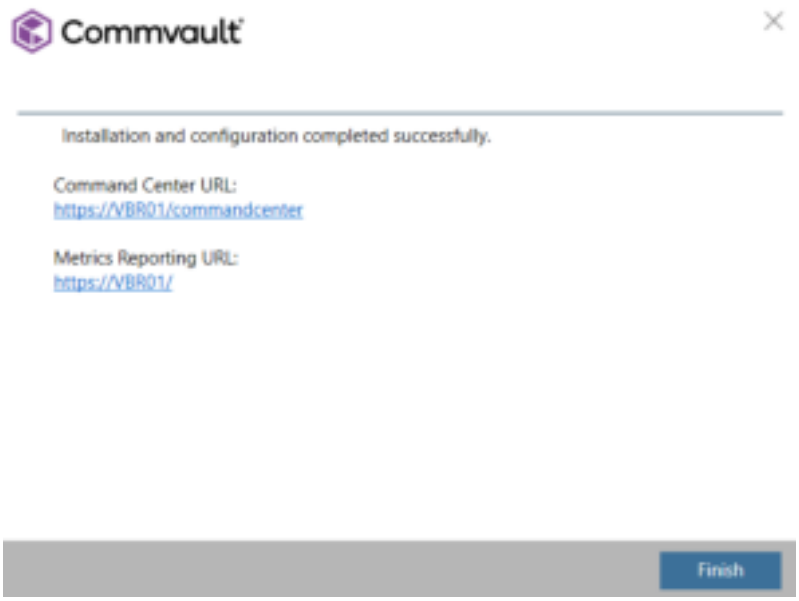

5. Select **Finish**.

Create a new database

Create a new MS SQL database (DB) for backup and recover operations (see Getting Started with SQL Server).

DB operations use a lot of system resources. Make sure that the system running the Commvault Command Center has a high-performance CPU and a large amount of local memory.

The following screenshot illustrates a typical configuration for Commvault on a Windows 2020 server with moderate performance:

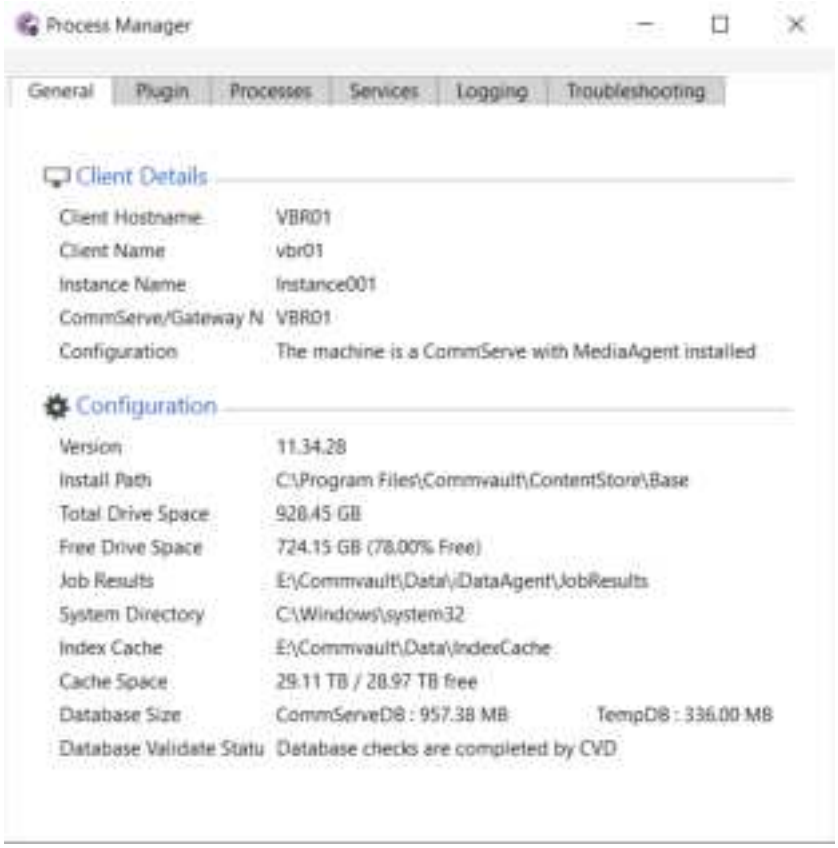

Set up an account

> When setting up your backup operation user or general user account, use Chrome to avoid known issues with Firefox.

1. Open your Chrome browser and create a new backup operator or general user account.
2. Once the account has been created, configure the necessary permissions and roles.
3. Log in to the management console using the newly created account.

> If the login fails, ensure that the correct account permissions have been assigned to the user.

Add cloud storage

1. Launch the Commvault Command Center using the link provided after installation.

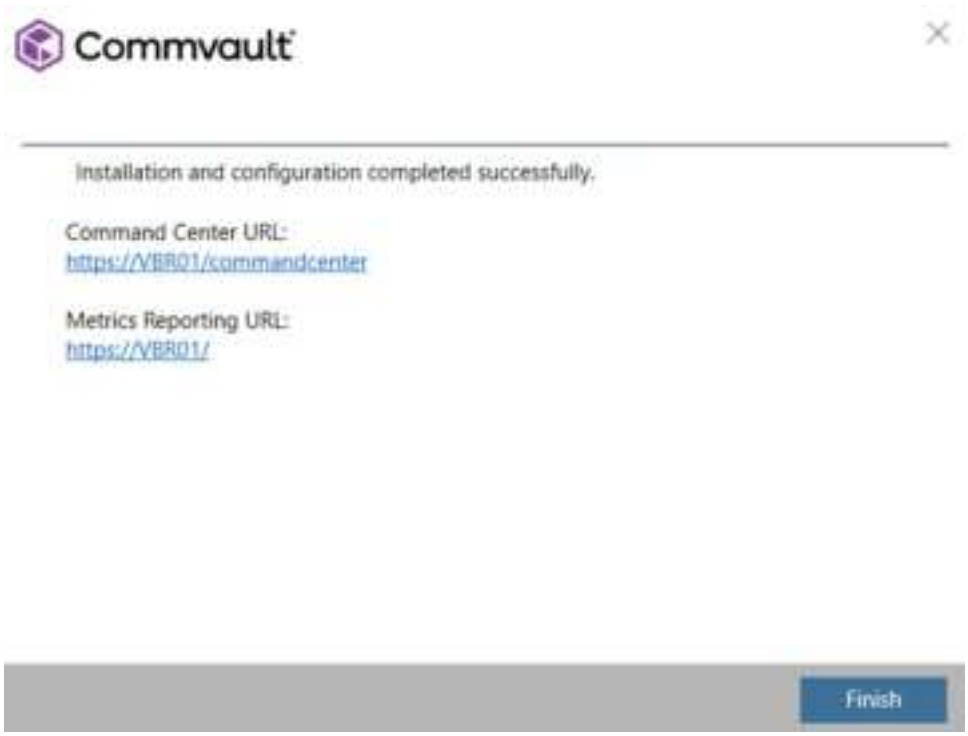

2. Select **Add storage**, and then select **LET'S GET STARTED**.

3. Select the appropriate object storage service (for example, AWS S3, Microsoft Azure Blob Storage, and so on).
4. Specify the following:

Type	S3 Compatible Storage
Name	Enter the S3 account name
Service host	https://s3.sv15.lyve.seagate.com

5. Select the Edit icon next to 'Credentials' and enter your credential name, access key ID, and secret access key.

> **i** If your bucket is configured with a governance lock, you may encounter issues. Ensure compliance with your organization's policies for object-level governance.

6. Edit your cloud object bucket information. Ensure that the correct region information is selected.
7. Select **Save**.

Validate backup and recovery

Run a test backup

1. After setting up your cloud storage and DB instances, run a small test backup to ensure the system is working as expected (see **Validate cloud upload and download** below).
2. Verify the backup data integrity and recovery processes.

Validate cloud upload and download

1. Install a S3 browser on your local system.
2. Launch the S3 browser and verify that content is uploaded to Lyve Cloud successfully.
3. Install a S3 browser on a Windows client. (This can be a free version.)

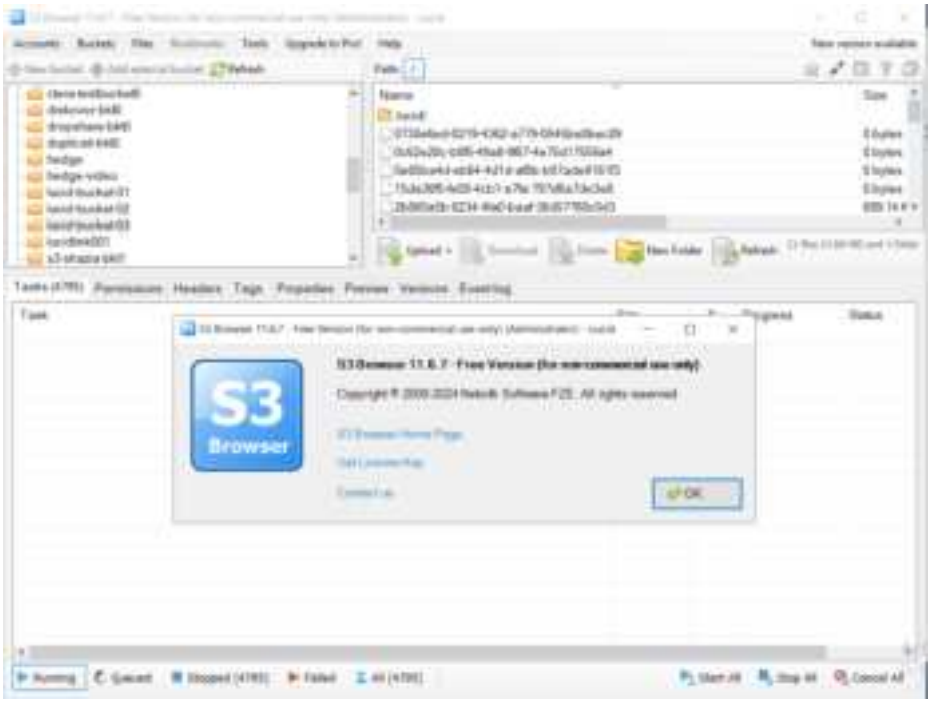

4. Create a Cloud user account.

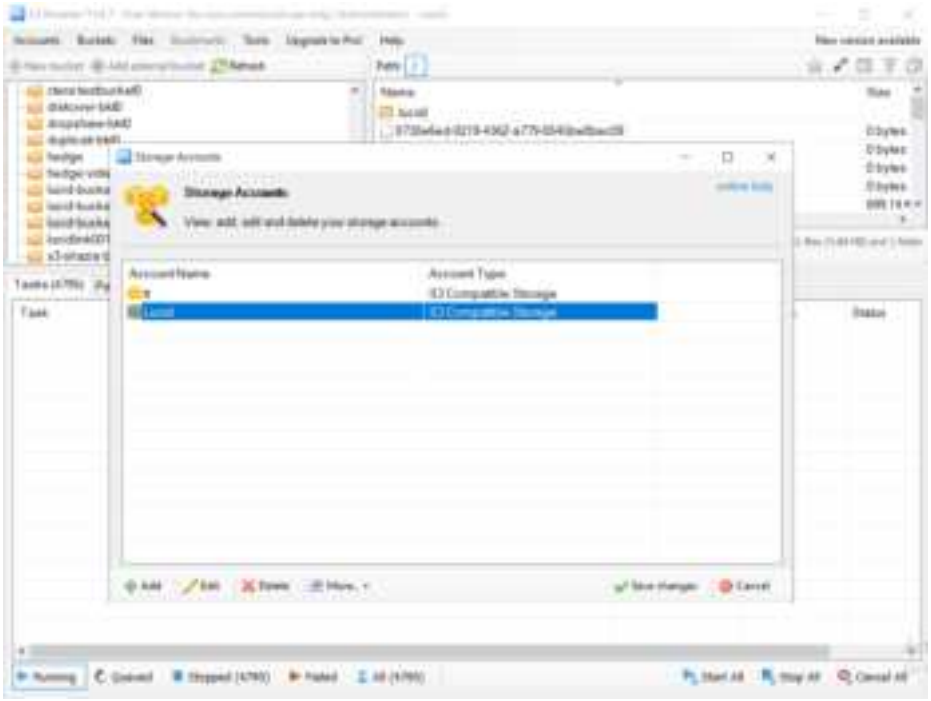

5. Enter the cloud user credentials.

6. Test the cloud user credentials. If you can see the contents stored in the bucket, the cloud operation is working properly (see below).

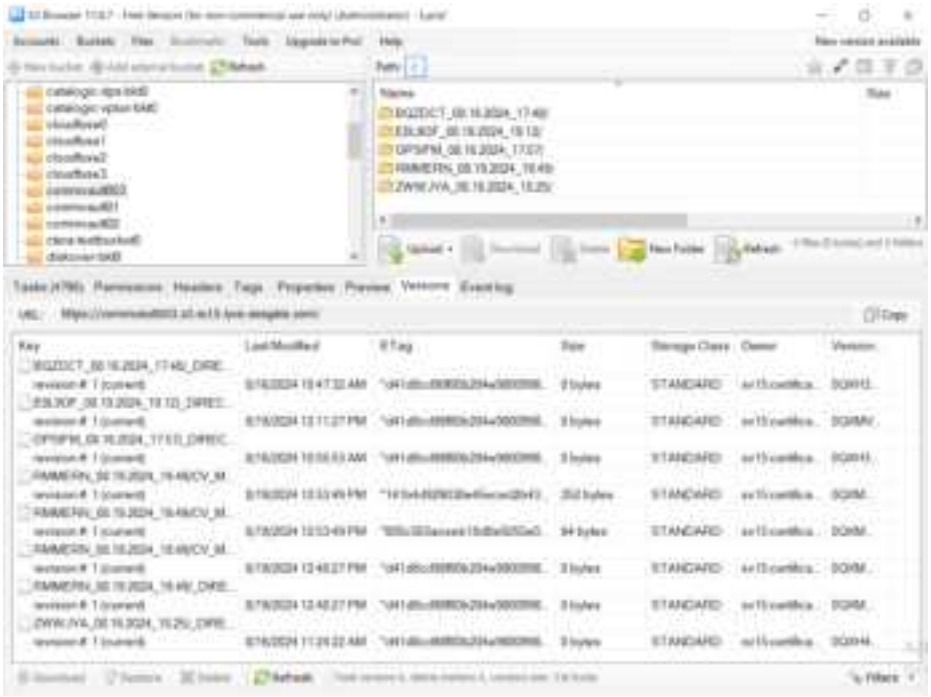

Review snapshots

1. Use the snapshot feature to capture point-in-time backups for validation.

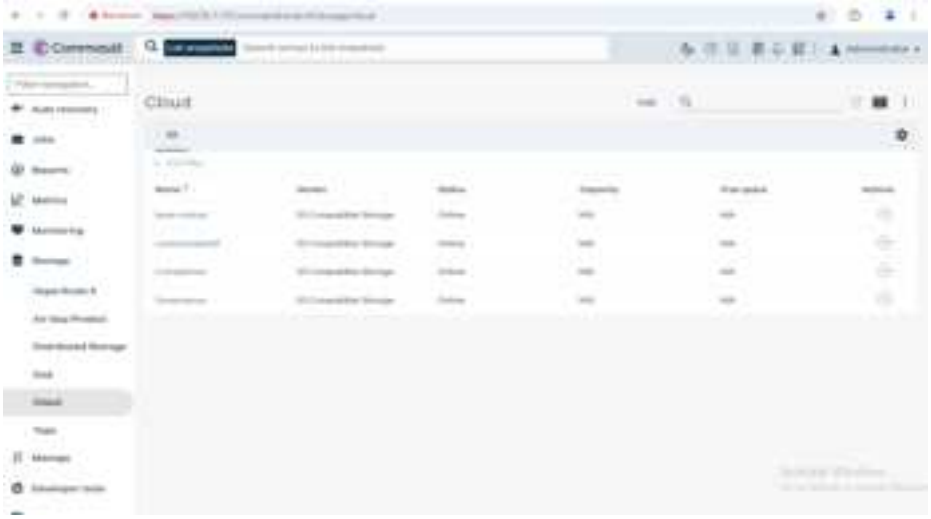

2. Mount the snapshot to ensure that it reflects the current state of your backup environment.

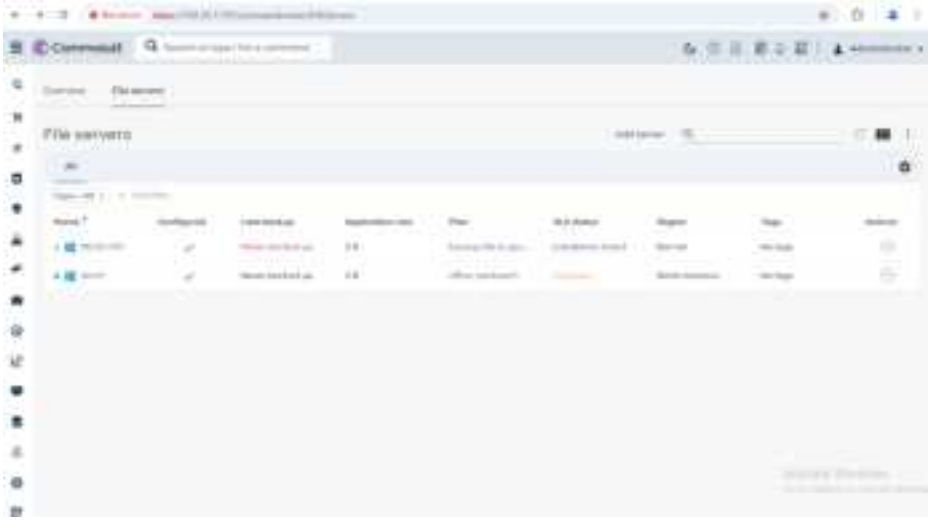

3. Configure the server.

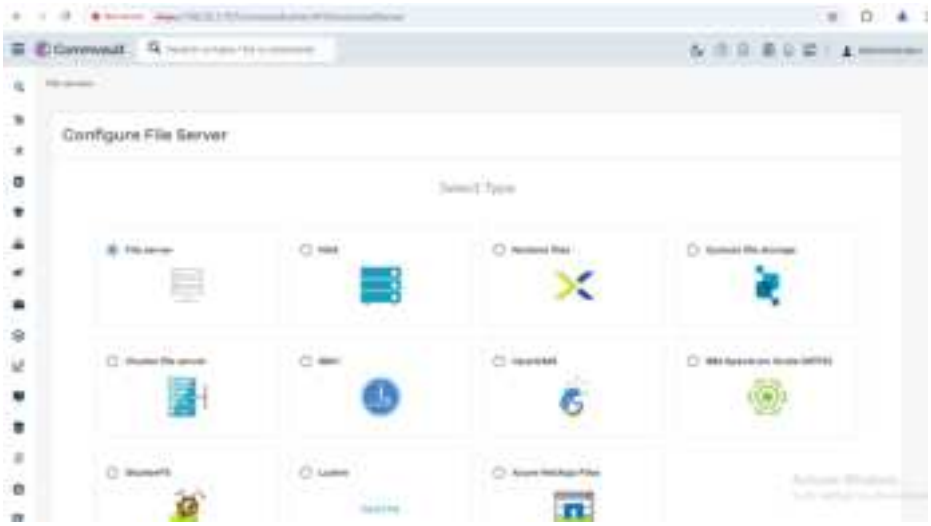

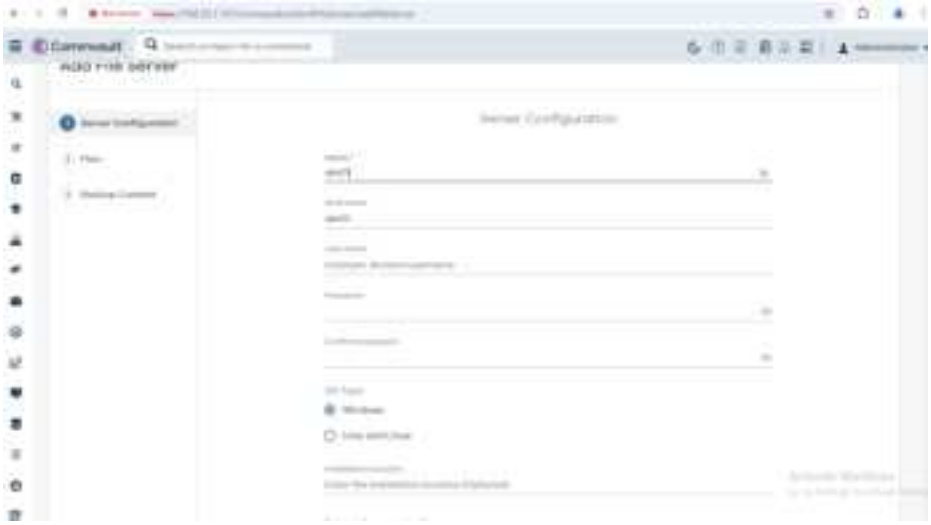

4. Select a backup plan.

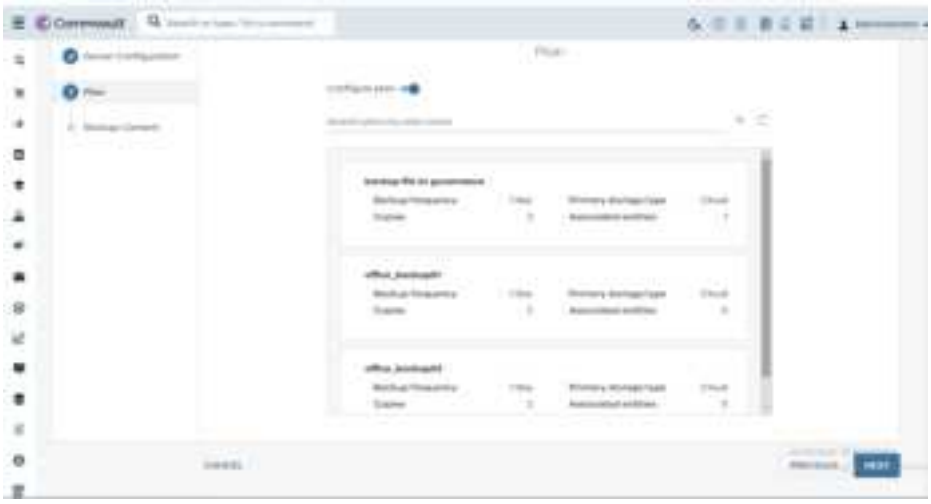

5. Select your backup content and add the plan.

Troubleshooting

Cloud Storage Fails Without Region Info

Ensure that the correct region information was entered during cloud configuration. Incorrect or missing data will prevent the storage bucket from being added.

Governance Lock and Compliance Issues

If governance lock compliance is required on the object level, ensure that the configuration matches your organization's regulatory requirements.

FileBrowser

Version 14.5 of FileBrowserGO and FileBrowser for Business can be used with Lyve Cloud Object Storage by following the instructions below.

App Store links

- https://itunes.apple.com/us/app/filebrowser-for-business/id854618029
- https://itunes.apple.com/us/app/filebrowsergo/id335493278

Set up Lyve Cloud Object Storage in FileBrowserGO and FileBrowser for Business

1. Choose the Amazon S3 connector type when adding a new storage location.

2. Paste the Lyve Cloud access key/secret key settings into the appropriate boxes.
3. Select the region or just leave it blank, as the region will be specified by the URL.
4. Enter the appropriate URL for the region you will be using, for example, https://s3.<region>.sv15.lyve.seagate.com

where:

<region> is the appropriate Lyve Cloud S3 endpoint URL, for example, us-east-1.

5. Set the display name to one appropriate for your connection.
6. Select **Save**. You will see a list of buckets, or objects in your bucket.

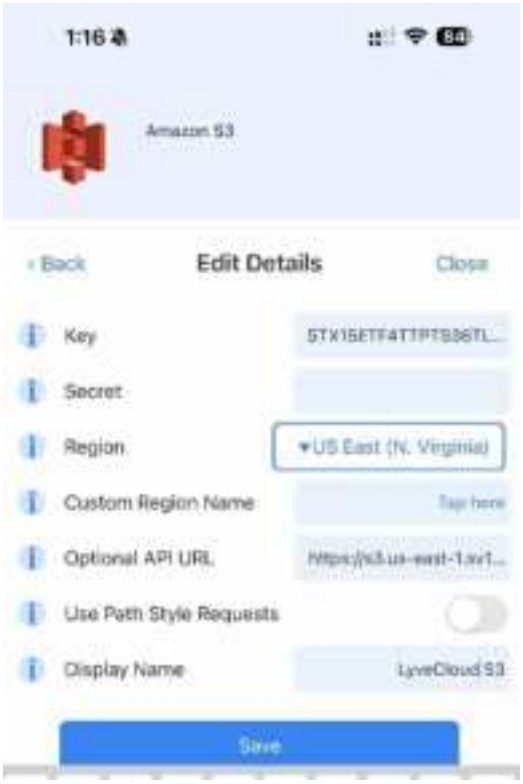

All file management features are available using FileBrowser, including copy, move, delete, rename, and create folder. Lyve Cloud Object Storage supports drag and drop. The iOS Files app can now be used to open documents in Lyve Cloud.

> Note—It's not possible to create buckets using FileBrowser.

IBM Aspera with HSTS

The following instructions are intended to provide a seamless and secure deployment of IBM Aspera with HSTS in a Windows or Linux (RedHat) environment, ensuring optimized performance for moving data to Lyve Cloud Object Storage. Adjustments may be made as needed based on specific organization policies and infrastructure.

Pre-Deployment

Network Configuration

Ports

Ensure the following are accessible:

- **Aspera TCP/UDP Port:** 33001 (default for data transfer)
- **SSH Port:** 22 (keep enabled initially to avoid accidental lockouts)

Firewall Rules

Permit inbound and outbound traffic on necessary ports. Verify the port status:

```
# netstat -na |grep 33001
```

SSH Security

Configure SSH Daemon

Open the sshd_config file and set:

- AllowTcpForwarding no
- AllowAgentForwarding no
- PubkeyAuthentication yes
- PasswordAuthentication yes
- port 33001
- port 22

> **Note**—For initial test-run we do not recommend to disable TCP port 22 because it is defaulted for ssh connection on most of the ssh client and you may lock yourself out in ssh accidently if you disable it.

Restart SSH Service

Apply changes:

```
# systemctl restart sshd.service
```

Verify that port 33001 is now listening:

```
# netstat -na |grep 33001
tcp        0      0 0.0.0.0:33001           0.0.0.0:*               LISTEN
tcp        0      0 10.0.10.242:33001       125.20.120.90:52332     ESTABLISHED
tcp        0      0 10.0.10.242:33001       125.20.120.90:57615     ESTABLISHED
tcp        0      0 10.0.10.242:33001       14.194.8.182:62305      ESTABLISHED
tcp6       0      0 :::33001
```

HSTS Installation and Verification

Installation

Install HSTS on RedHat using:

```
# rpm -Uvh ibm-aspera-hsts-<version>-linux-64-release.rpm
```

To verify the installation:

- HSTS Version: # ascp -A
- License Validation: # cat /etc/aspera-license

User and Environment Setup

User Creation

Create and configure a default user (svcAspera):

```
# useradd svcAspera
# su svcAspera
# sudo chsh -s /bin/aspshell svcAspera
```

SSH Key Setup

Establish user SSH credentials:

```
# sudo mkdir /home/<username>/.ssh
# sudo chmod 700 /home/<username>/.ssh
# sudo touch /home/<username>/.ssh/authorized_keys
# sudo chmod 600 /home/<username>/.ssh/authorized_keys
```

Verify ssh and network ports are open for Aspera and HSTS:

```
# systemctl restart sshd.service
# netstat –ha |grep 33001
# default TCP/UDP ports Aspera uses.
```

Configuring Transfer Parameters

Global Transfer Settings

To enforce security policies on data transfers, configure the following:

```
# sudo asconfigurator -x "set_node_data;authorization_transfer_in_value,deny"
# sudo asconfigurator -x "set_node_data;authorization_transfer_out_value,deny"
```

User-specific Parameters

Set specific transfer permissions and bandwidth for user=svcAspera:

```
# sudo asconfigurator -x
"set_user_data;user_name,svcAspera;authorization_transfer_in_value,allow"
# sudo asconfigurator -x
"set_user_data;user_name,svcAspera;transfer_in_bandwidth_flow_target_rate_default,96000"

# sudo asconfigurator -x "set_node_data;authorization_transfer_in_value,deny"
# sudo asconfigurator -x "set_node_data;authorization_transfer_out_value,deny"

# sudo asconfigurator -x
"set_user_data;user_name,test_aspera;authorization_transfer_in_value,allow"
user_name: svcAspera
# sudo asconfigurator -x
"set_user_data;user_name,test_aspera;authorization_transfer_out_value,allow"
success
user_name: svcAspera
# sudo asconfigurator -x "set_user_data;user_name,test_aspera;absolute,s3:// s3.us-west-
1.sv15.lyve.seagate.com/aspera-test-bucket"
success
user_name: svcAspera
```

```
# sudo asconfigurator -x
"set_user_data;user_name,test_aspera;transfer_in_bandwidth_flow_target_rate_cap,unlimited"
success
user_name: svcAspera
# sudo asconfigurator -x
"set_user_data;user_name,test_aspera;transfer_out_bandwidth_flow_target_rate_cap,unlimited"
success
user_name: svcAspera
# sudo asconfigurator -x
"set_user_data;user_name,test_aspera;transfer_in_bandwidth_flow_target_rate_default,96000"
success
user_name: svcAspera
# sudo asconfigurator -x
"set_user_data;user_name,test_aspera;transfer_out_bandwidth_flow_target_rate_default,96000"
success
user_name: svcAspera
```

Reference

ascp is IBM Aspera's version of the UNIX tool – scp, and async is IBM Aspera's version of the UNIX tool rsync. Many of the arguments are similar but the internals of how the files are transferred are very different. More information about the tools can be found at:

- **ascp:** https://www.ibm.com/docs/en/ahts/4.4.x?topic=atfcl-ascp-command-reference
- **async:** https://www.ibm.com/docs/en/ahts/4.4.x?topic=ra-async-command-reference

Enable Trapd for Object Storage (Optional)

Trapd is the Aspera service that enables it to write to the object storage (including Hadoop Distributed File System (HDFS). Trapd is supported in HSTS for Linux 64-bit and Windows 64-bit. HSTS should be in close proximity to the object storage to minimize latency.

Note—Trapd is disabled by default in HSTS.

- **Enable Trapd:** # /opt/aspera/bin/astrap-config.sh enable
- **Disable Trapd (if switching to another storage configuration):** # /opt/aspera/bin/astrap-config.sh disable

When switching from virtual node to pvcl_cloud or use another option, link the pvcl_cloud kernel to ensure Trapd is indeed disable.

```
#/opt/aspera/bin/astrap-config.sh disable
#ln -s /opt/aspera/lib/pvcl/libpvcl_cloud.so /opt/aspera/lib/libpvcl_cloud.so
```

Use a text editor to open the s3.properties file (/opt/aspera/etc/trapd/s3.properties) and make the following changes:

```
s3service.https-only=true
s3service.s3-endpoint=s3.sv15.lyve.seagate.com        ## for example
s3service.s3-endpoint-https-port=443
s3service.disable-dns-buckets=true
s3service.use-path-style-url=true
```

Restart the trapd service:

```
# sudo systemctl restart asperatrapd
```

S3 Configuration for Object Storage

Edit the s3.properties file (/opt/aspera/etc/trapd/s3.properties) to configure your S3 endpoint for secure, path-style URL usage:

```
s3service.https-only=true
s3service.s3-endpoint=<s3-endpoint>
s3service.s3-endpoint-https-port=443
s3service.use-path-style-url=true
```

where <s3-endpoint> is your S3 endpoint URL.

Restart trapd to apply these settings:

```
# sudo systemctl restart asperatrapd
```

IBM Aspera Desktop Client Setup (Windows)

Installation

1. Install the IBM Aspera Desktop Client on Windows using the provided installer, for example, **IBMAsperaDesktopClient-ML-4.4.3.891-win-v143-64-release.exe**. Proceed through the setup wizard steps.
2. When the installation is complete, launch Aspera Desktop Client on Windows.
3. Configure your S3 Connection parameters:

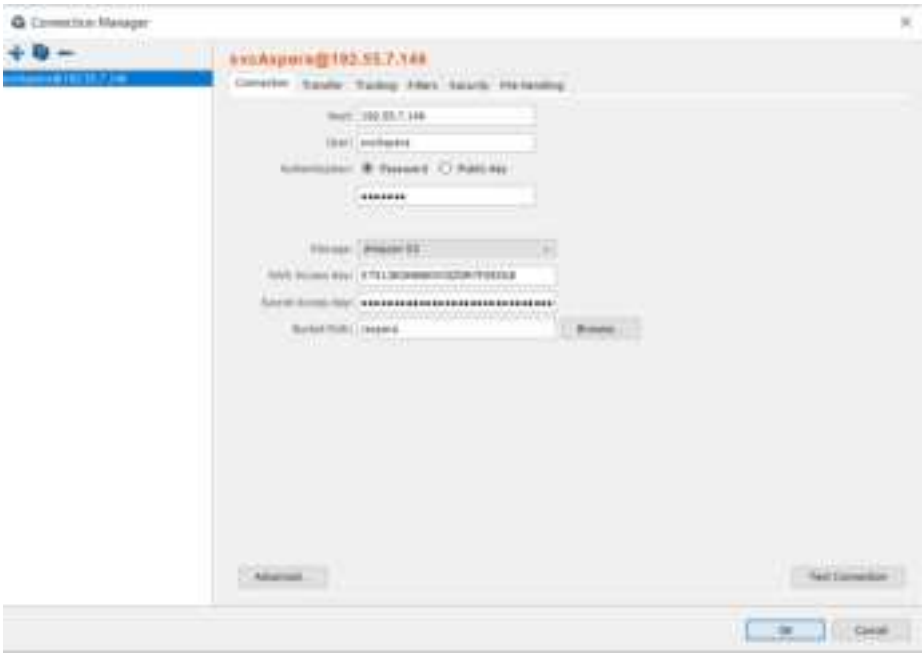

When using the default user:

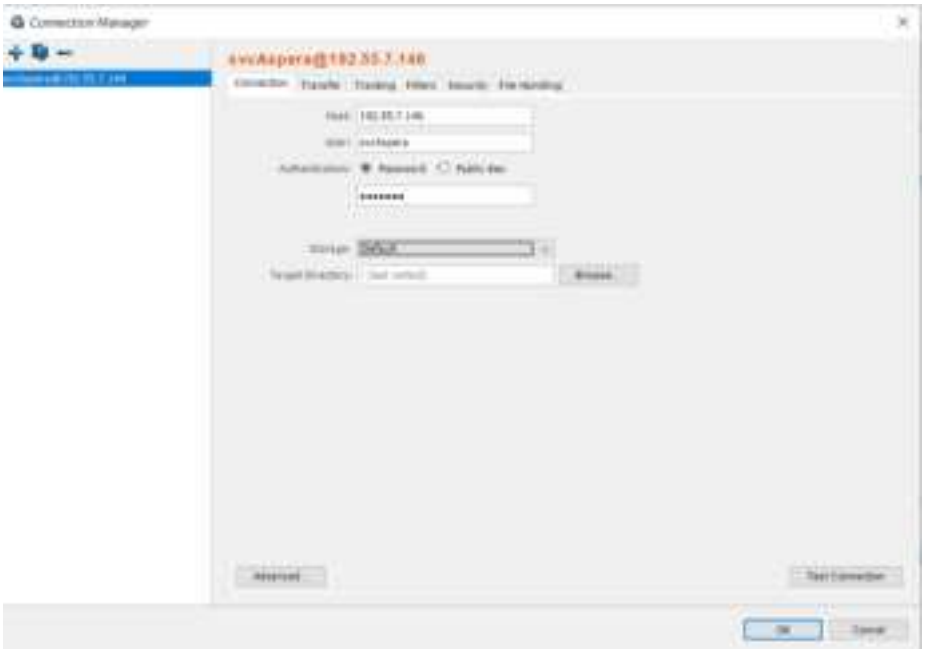

The user credential and S3 endpoint access URL can be configured on either HSTS server side or alternatively on the Windows client side.

To configure the user credential on the Windows client side, the user credentials and S3 endpoint URL information need to be commented out in aspera.conf.

For configuration on the HSTS, the user credential and S3 access endpoint information are added in aspera.conf, following this format:

```
s3://<access-key>:<secret-key>@<s3-endpoint>
```

where:

- <access-key> is your access key.
- <secret-key> is your secret key.
- <s3-endpoint> is the appropriate Lyve Cloud S3 endpoint URL, for example, us-east-1.

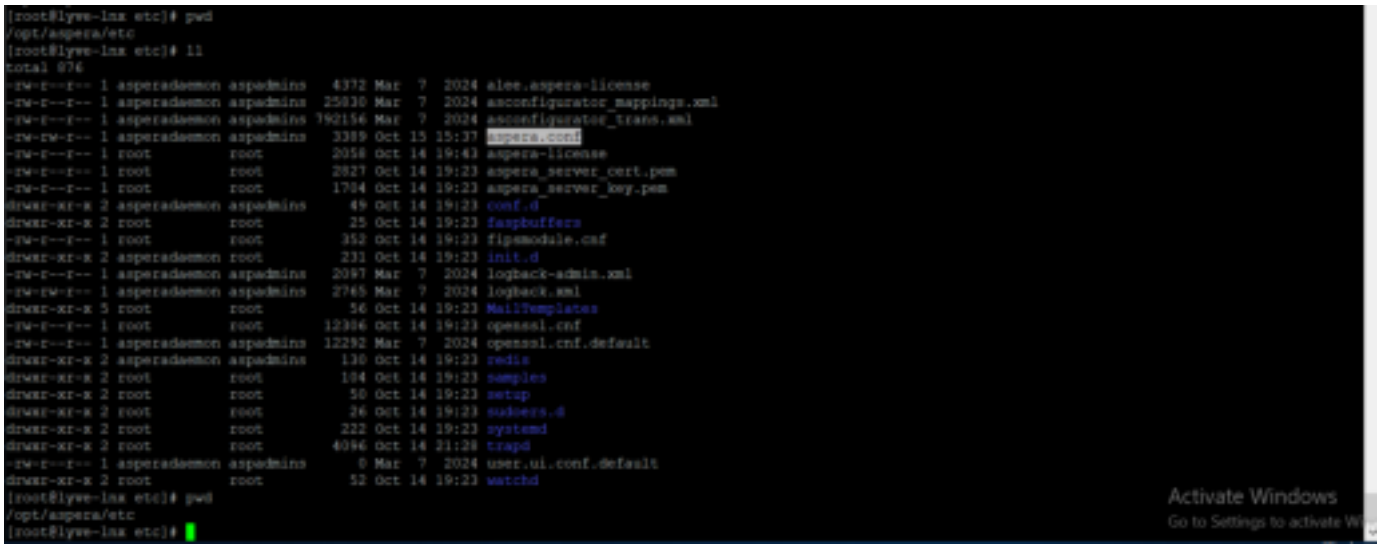

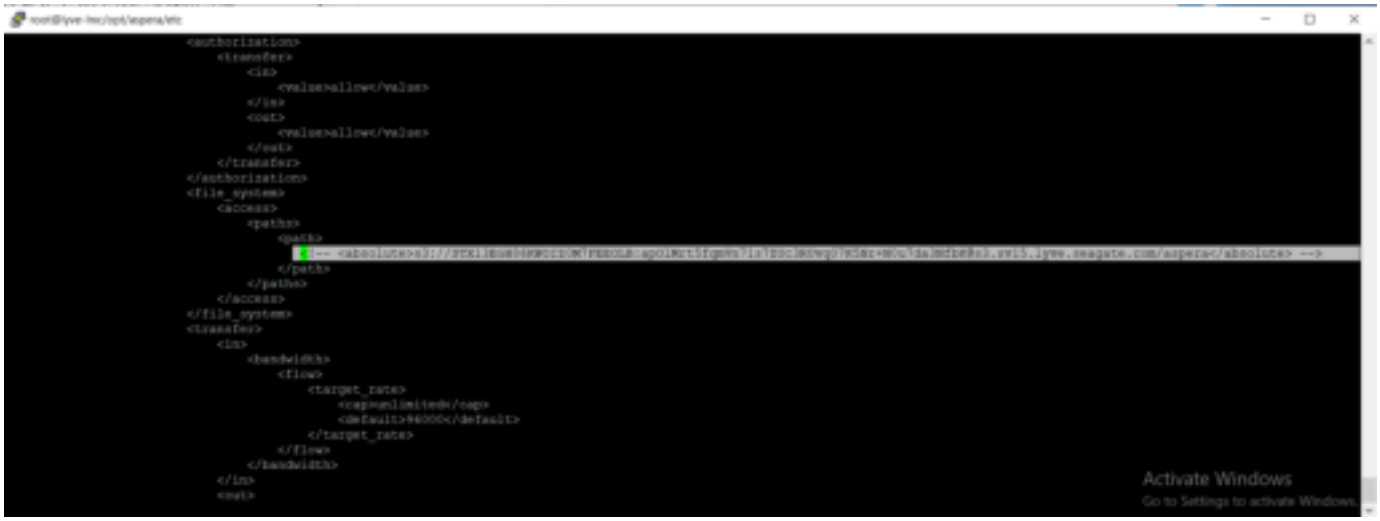

1. Ensure the TCP/UDP ports are open.

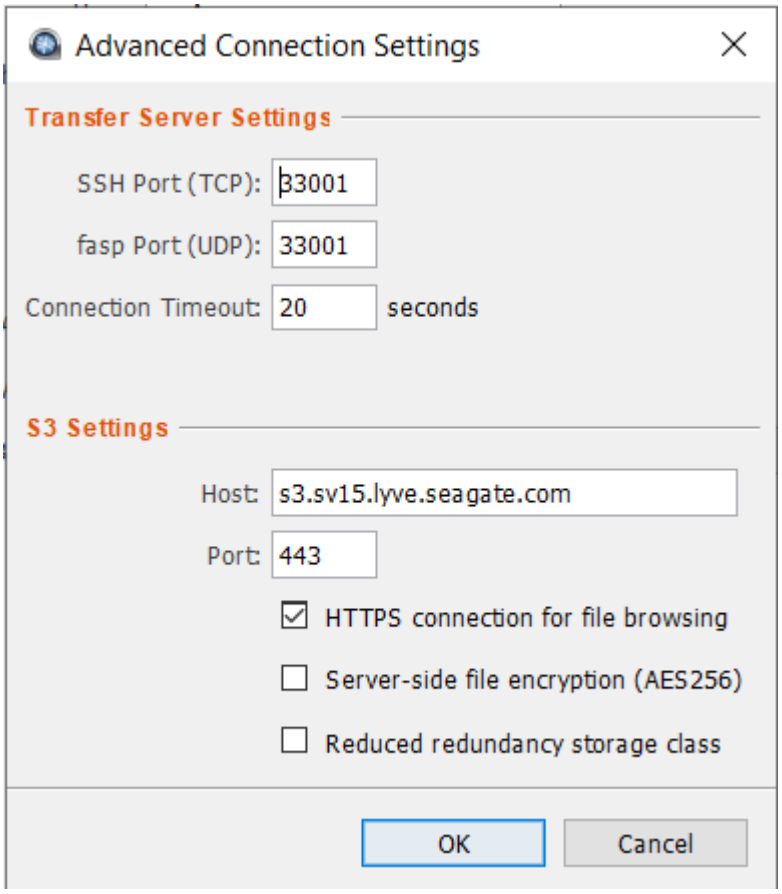

2. Set up initial transfer speed parameters.

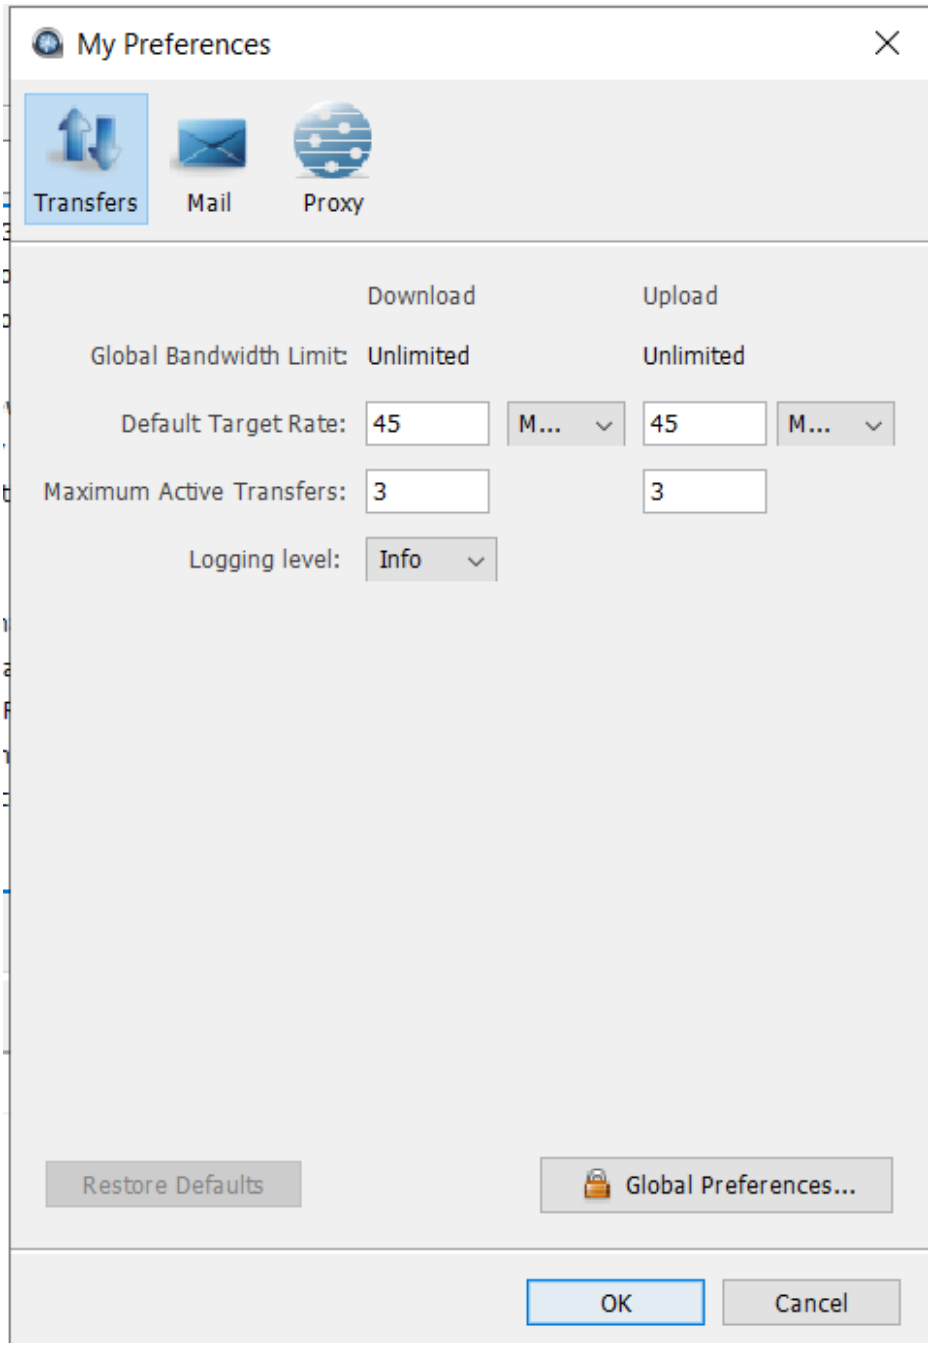

3. Test the setup from an Aspera Windows Client.
4. Launch Connection Manager.
5. Send files to and receive files from S3 cloud storage.

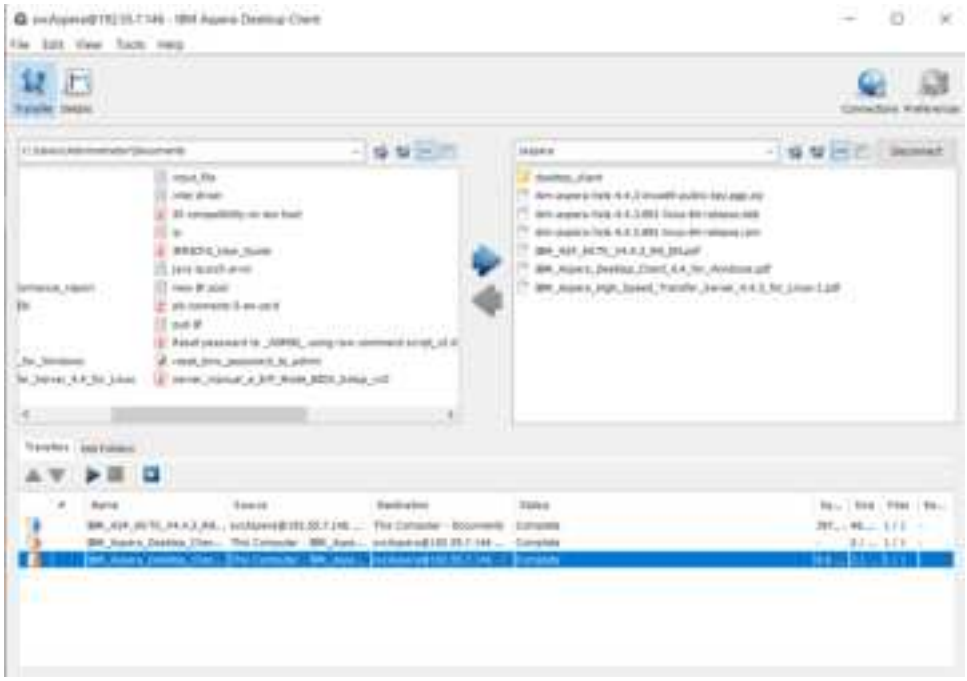

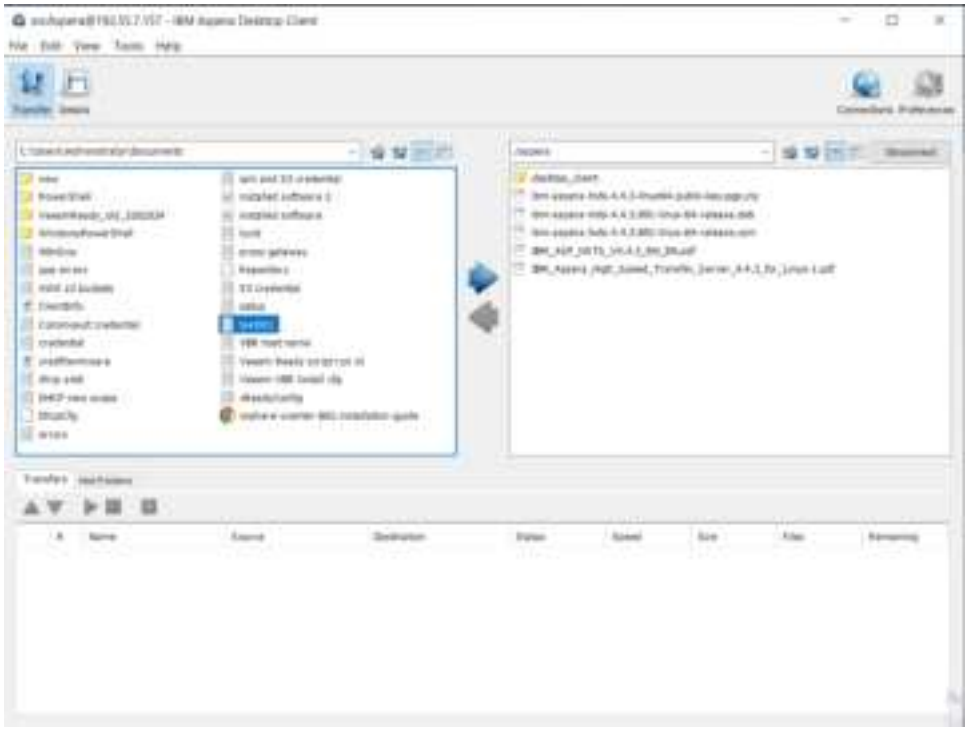

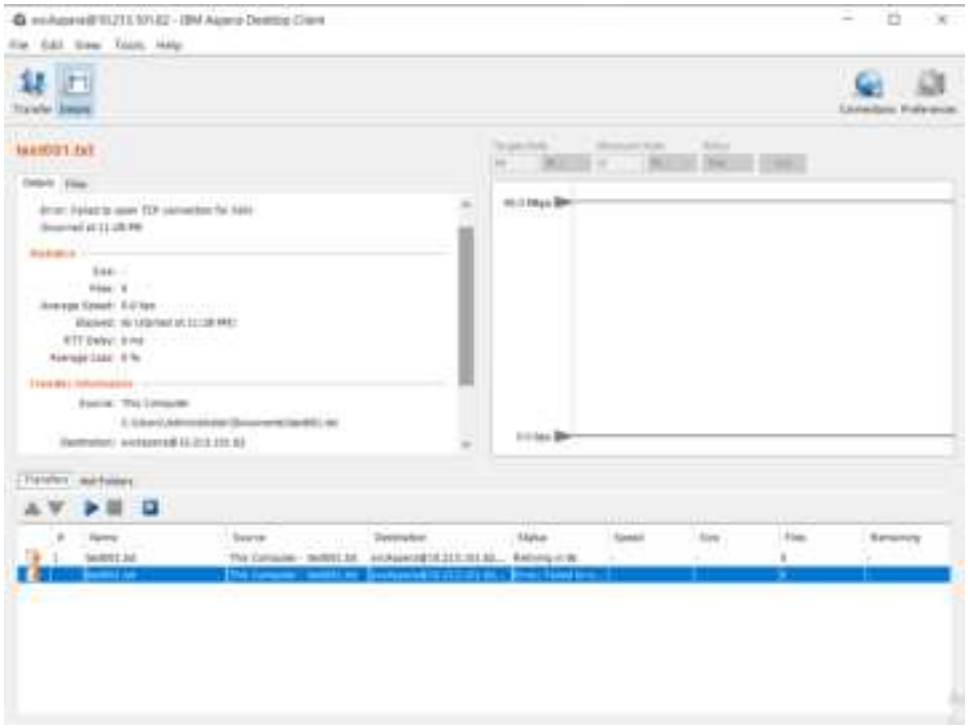

Prerequisites

Before configuration, ensure that the following requirements have been met:

- Fully licensed Windows 2020 server.
- Active internet connection.
- Configurator must have administrative privileges on the computer hosting the LucidLink application.
- Seagate S3 cloud storage bucket access to use for your filespace.

Install LucidLink app

1. Open your browser and navigate to theLucidLink download page.
2. Download the appropriate installer for your operating system.
3. Open the folder where you receive downloads and launch the installer.
4. Follow the onscreen instructions to complete the installation.

Create a cloud storage instance

1. Open the LucidLink app.
2. Log in with your account credentials. If you don't have a LucidLink account, register and create an account by selecting the option on the login screen.
3. Create a cloud, or map the filespace to a cloud service provider. (For Seagate Lyve Cloud, selec**Other Cloud**.)

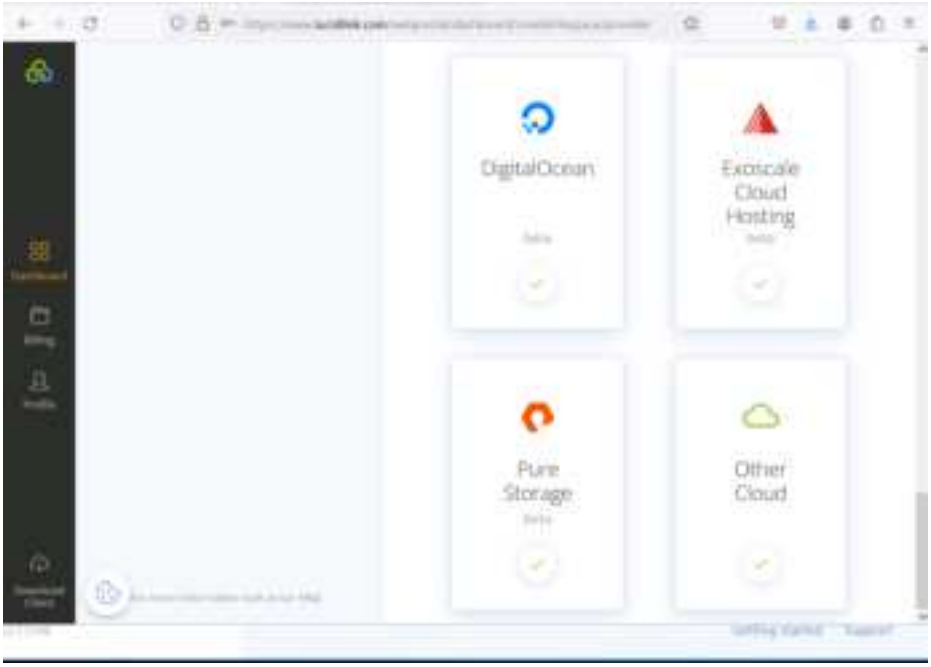

4. Select **Create a filespace**.
5. Enter a name that identifies the filespace within your domain.

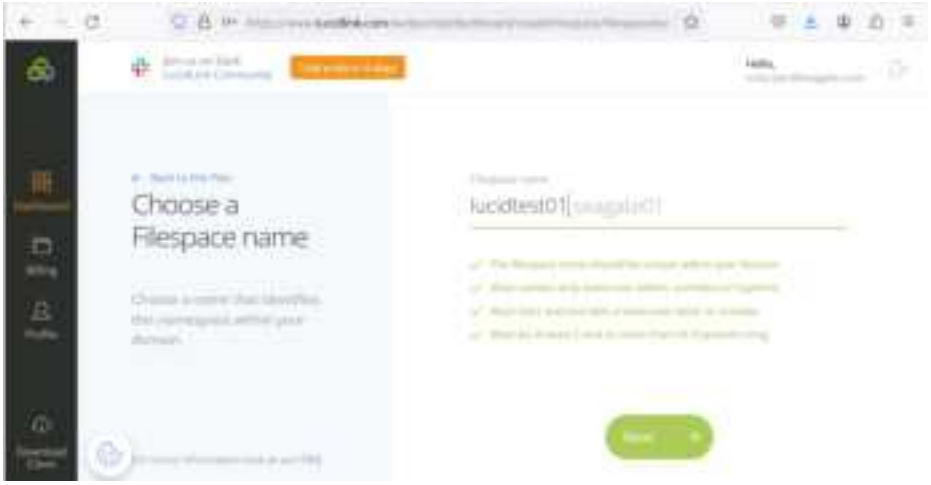

6. Select a cloud storage region.

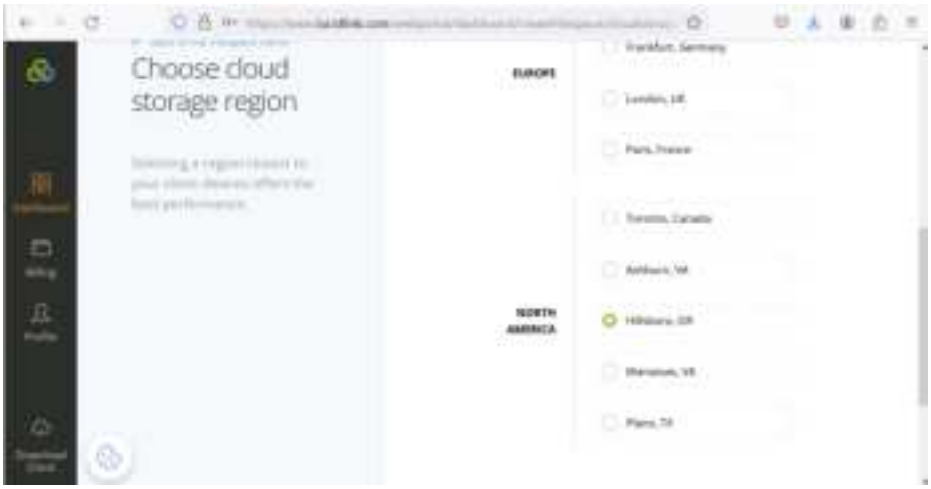

7. (Optional) If you want to change the block size for your filespace, select **configure advanced settings**. (For recommended block sizes, see Choosing a LucidLink Block Size.)
8. Verify your choices and select **Create**.

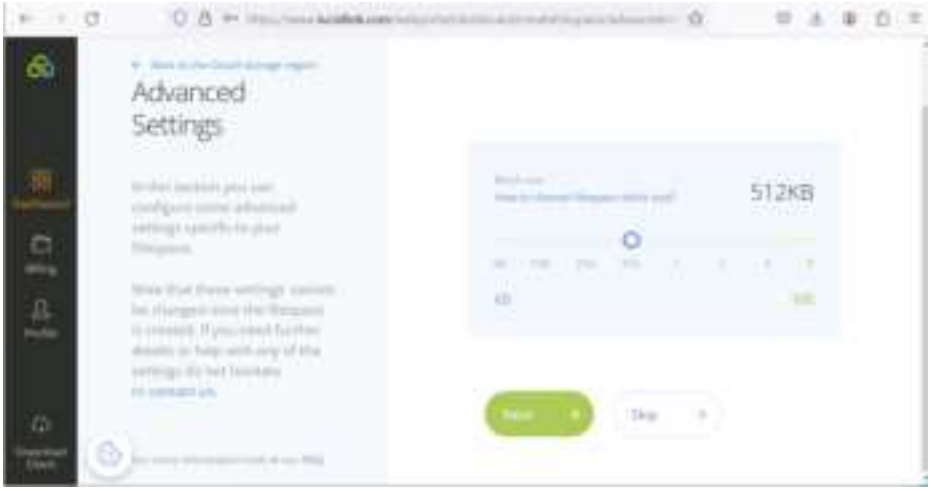

Initialize and connect the filespace

1. After your filespace has been set up, select the **Initialize** button on your filespace's card.

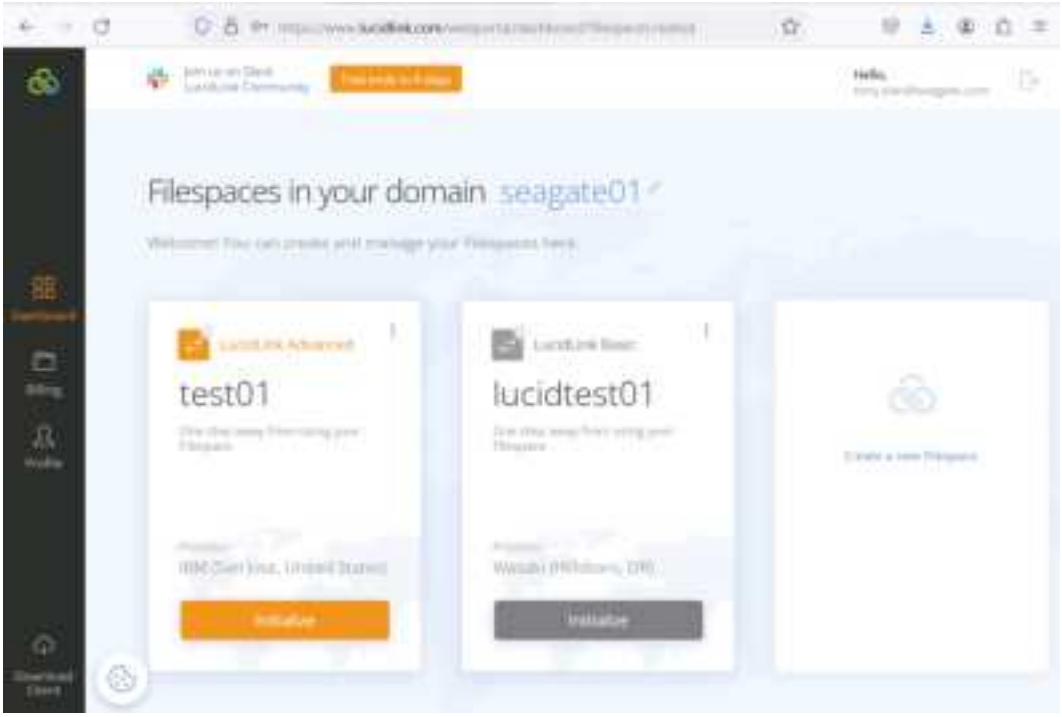

2. Enter your Lyve Cloud user account credentials in the LucidLink app, and then select **Next**.

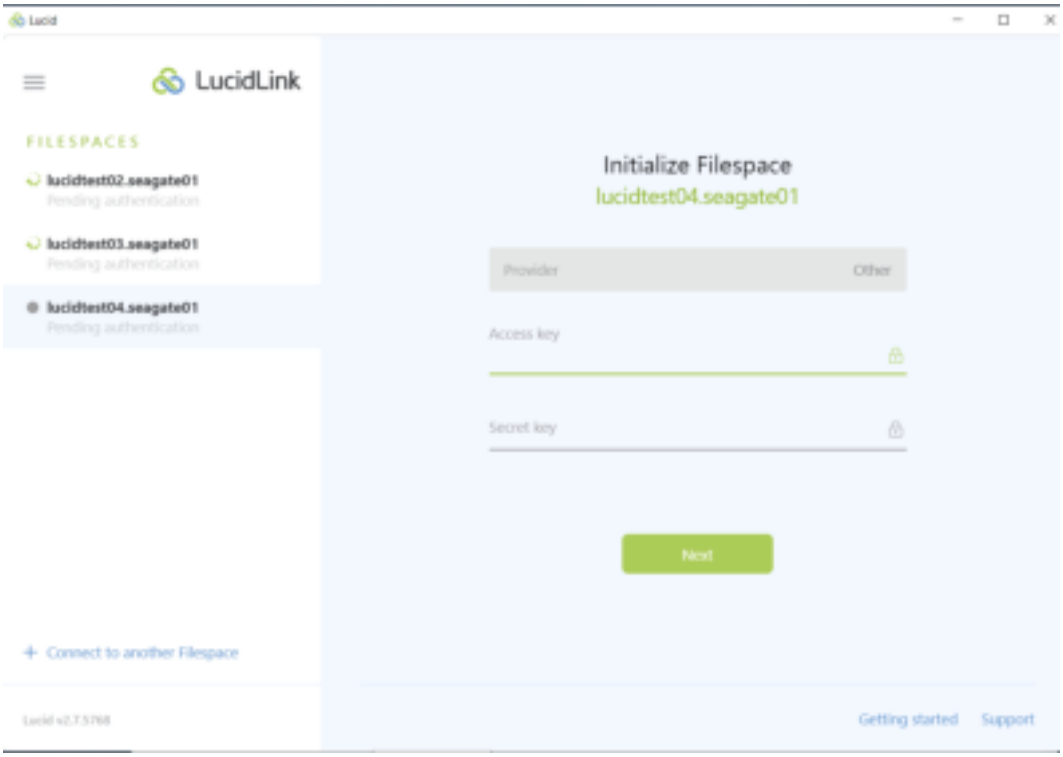

3. Create a root password

> **i** Due to the robust security mechanism enforced by LucidLink, there is no password recovery option. LucidLink cannot recover a root user password. The root user can change all other user passwords at any time.

Once the filespace has been created, you must connect to it make it accessible as a local drive on your system.

1. Enter the Filespace name in the following format:

   <filespace-name>.<domain>

   Select **Continue**.

2. To connect to a filespace for the first time, enter your root credentials.

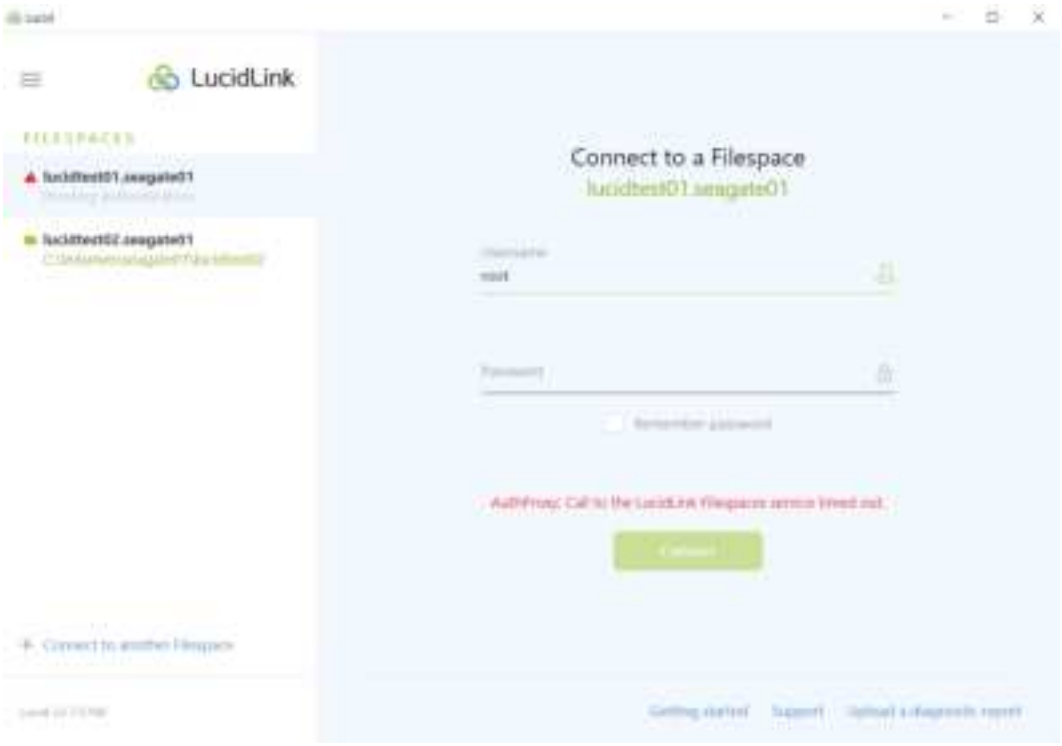

Select **Connect**.

Filespace management

You can view the current status of your filespace from the LucidLink control panel.

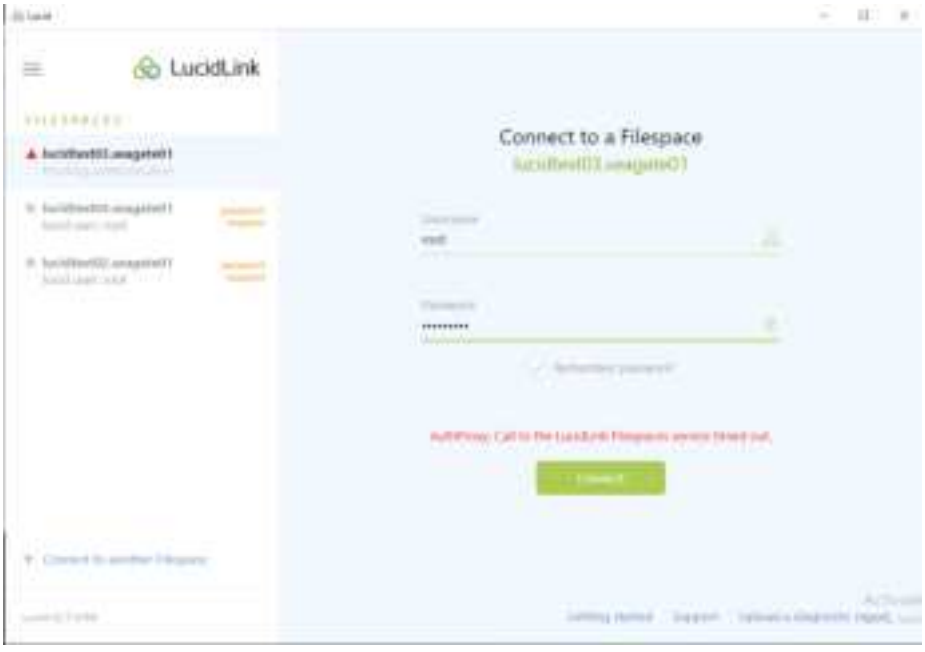

Once connected, the user will have an access to the control panel.

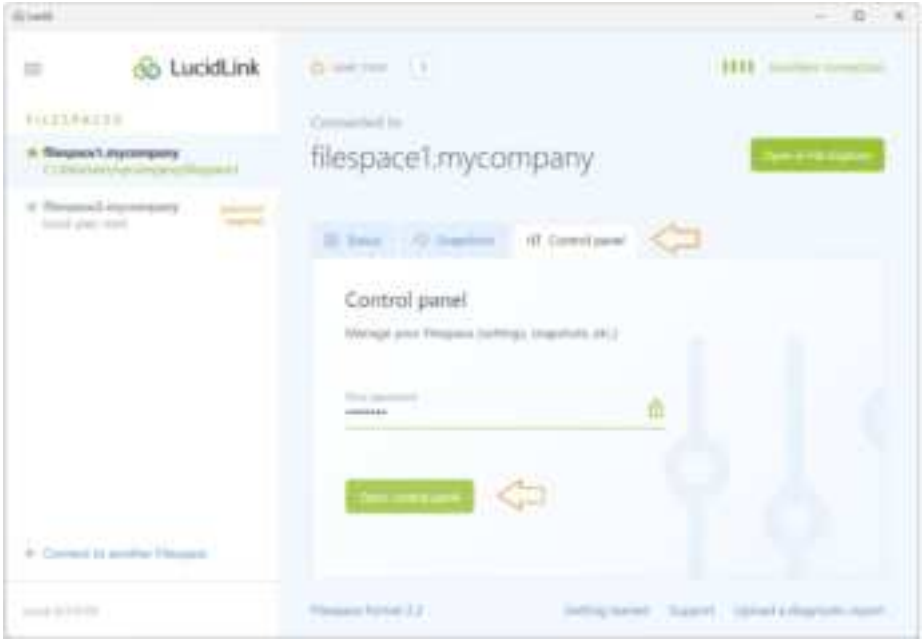
LucidLink
filespace1.mycompany
Control panel

MinIO Client

MinIO Client is a S3 compatible client that allows you to connect to Lyve Cloud Object Storage and perform operations on your Lyve Cloud buckets.

Configure MinIO

To configure MinIO Client for use with with Lyve Cloud:

1. Download Minio Client and set executable permissions:

   - wget https://dl.minio.io/client/mc/release/linux-amd64/mc
   - chmod +x mc

2. Configure a Lyve Cloud alias using your Lyve Cloud credential:

   mc alias set minio https://s3.<region>.sv15.lyve.seagate.com <access-key> <secret-key>

   where:

   - <region> is the appropriate Lyve Cloud S3 endpoint URL, for example, us-east-1.
   - <access-key> is your account access key.
   - <secret-key> is your account secret key.

Sample commands

List buckets

```
./mc ls minio
[2024-10-02 16:41:40 CEST]     0B acronis-cloudseed/
[2024-08-28 17:48:54 CEST]     0B alluxio-test-bkt0/
[2024-10-23 09:29:33 CEST]     0B ansible-bucket/
[2024-10-24 09:47:35 CEST]     0B apache-bucket-1/
```

List objects in specific bucket

```
./mc ls minio/minio-bucket
[2024-11-07 10:29:42 CET]   541B 0R0PJTMQ9PA5V8W1GYEDECRVBF v1 PUT clean.sh
[2024-11-07 10:29:42 CET]   131B 0R0PJTMQ9KEY65Y989H8AVEG0E v1 PUT config.json
[2024-11-07 10:29:42 CET]     0B 0R0PJTMHV4BABGHPSSBR476JXA v1 PUT mytext.txt
```

Upload object to bucket

./mc cp tesh.bash minio/minio-bucket
tesh.bash: **3.02 KiB / 3.02 KiB**

2.29 KiB/s 1s

Restic

Restic is validated for use with Lyve Cloud Object Storage.

Configure Restic

To configure Restic for use with Lyve Cloud:

1. Export the variables:

   - export RESTIC_REPOSITORY="s3:s3.<region>.sv15.lyve.seagate.com/<bucket>"
   - export AWS_ACCESS_KEY_ID="<access-key>"
   - export AWS_SECRET_ACCESS_KEY="<secret-key>"

   where:

   - <region> is the appropriate Lyve Cloud region, for example, us-east-1.
   - <bucket> is the name of your Lyve Cloud bucket.
   - <access-key> is your access key.
   - <secret-key> is your secret key.

2. Initialize the repository:

   - Run command: restic init
   - Enter password for new repository

Restic Basic Commands

Backup

```
restic backup /Users/660186/desktop/PicFolder
```

Result

```
enter password for repository:
repository 91500045 opened (version 2, compression level auto)
created new cache in /Users/660186/Library/Caches/restic
no parent snapshot found, will read all files
[0:00]          0 index files loaded

Files:          4 new,     0 changed,     0 unmodified
```

Dirs: 4 new, 0 changed, 0 unmodified
Added to the repository: 2.154 MiB (2.121 MiB stored)

processed 4 files, 2.149 MiB in 0:19
snapshot 69ea64c2 saved

Restore

restic restore latest --target /Users/660186/documents/Restic

Result

enter password for repository:
repository 91500045 opened (version 2, compression level auto)
[0:00] 100.00% 1 / 1 index files loaded
restoring snapshot 69ea64c2 of [/Users/660186/desktop/PicFolder] at 2024-11-15 09:36:30.23865 +0100 CET by 660186@sgs-u660186l320 to /Users/660186/documents/Restic
Summary: Restored 8 files/dirs (2.149 MiB) in 0:04

Check

restic check

Result

using temporary cache in /var/folders/56/5mhgst5d00xdtvlg3nx1gmcrmrpk1z/T/restic-check-cache-4112886137
create exclusive lock for repository
enter password for repository:
repository 91500045 opened (version 2, compression level auto)
created new cache in /var/folders/56/5mhgst5d00xdtvlg3nx1gmcrmrpk1z/T/restic-check-cache-4112886137 load indexes
[0:00] 100.00% 1 / 1 index files loaded
check all packs
check snapshots, trees and blobs
[0:00] 100.00% 1 / 1 snapshots

S3cmd

S3cmd is validated for use with Lyve Cloud Object Storage.

Configure S3cmd

To start using S3cmd, you'll need to configure it with Lyve Cloud Object Storage.

1. Run the configuration command:

   ```
   s3cmd –configure
   ```

2. Enter your S3 credentials:

Access Key	Enter your access key.
Secret Key	Enter your secret key.
Default Region	Enter the default region for your bucket, for example,us-east-1.
S3 Endpoint	Enter s3.<region>.sv15.lyve.seagate.com, where <region> is the appropriate Lyve Cloud region, for example,s3.us-east-1.sv15.lyve.seagate.com.
host_bucket	Enter %(bucket)s.s3.<region>.sv15.lyve.seagate.com, where <region> is the appropriate Lyve Cloud region, for example,%(bucket)s.s3.us-east-1.sv15.lyve.seagate.com.

3. Test the configuration to ensure everything is set up correctly.

Sample commands

List buckets

```
s3cmd ls
```

Result Example

```
2024-10-02 14:41  s3://acronis-cloudseed
2024-08-28 15:48  s3://alluxio-test-bkt0
```

2024-10-23 07:29 s3://ansible-bucket
2024-10-24 07:47 s3://apache-bucket-1
2024-10-23 11:20 s3://apache-spark-bucket
2024-11-05 08:33 s3://apostrophecms-bucket
2024-02-22 19:51 s3://aspera
2024-03-04 05:44 s3://aspera-test-bucket
2024-10-01 04:45 s3://aspera01
2024-10-08 20:22 s3://bacula-bkt-0
2024-10-25 19:46 s3://calamu-bkt0

Upload a file

s3cmd put /Users/660186/documents/test-user.sh s3://s3cmd-bucket

Result Example

WARNING: Module python-magic is not available. Guessing MIME types based on file extensions.

upload: '/Users/660186/documents/test-user.sh' -> 's3://s3cmd-bucket/test-user.sh' [1 of 1]

 130 of 130   100% in   2s   56.21 B/s done

Download a file

s3cmd get s3://s3cmd-bucket/config.json /Users/660186/desktop/local-config.json

Result Example

download: 's3://s3cmd-bucket/config.json' -> '/Users/660186/documents/local-config.json' [1 of 1]

 131 of 131   100% in   0s   271.77 B/s done

Sync a local folder to an S3 bucket

s3cmd sync /Users/660186/desktop/PicFolder s3://s3cmd-bucket

Result Example

WARNING: Module python-magic is not available. Guessing MIME types based on file extensions.
upload: '/Users/660186/desktop/PicFolder/Pic-1.png' -> 's3://s3cmd-bucket/PicFolder/Pic-1.png' [1 of 3]
 1121955 of 1121955   100% in   14s   73.53 KB/s done
upload: '/Users/660186/desktop/PicFolder/Pic-2.png' -> 's3://s3cmd-bucket/PicFolder/Pic-2.png' [2 of 3]
 479053 of 479053   100% in   5s   83.06 KB/s done
upload: '/Users/660186/desktop/PicFolder/Pic-3.png' -> 's3://s3cmd-bucket/PicFolder/Pic-3.png' [3 of 3]
 645935 of 645935   100% in   17s   35.10 KB/s done

Done. Uploaded 2246943 bytes in 38.5 seconds, 56.98 KB/s.

S3FS

S3FS (fuse) is certified for use with Lyve Cloud Object Storage.

Configuration instructions are provided below. See also:How to Mount S3 Bucket on CentOS and Ubuntu using S3FS.

Configure S3FS

Your Lyve Cloud access key and secret key are required to configure S3FS.

1. Replace the AWS_ACCESS_KEY_ID and AWS_SECRET_ACCESS_KEY with your actual Lyve Cloud access key and secret key values.

   ```
   $  <your-password-file> AWS_ACCESS_KEY_ID:AWS_SECRET_ACCESS_KEY
   ```

2. Make sure that the file has proper permission:

   ```
   $ chmod 600 <your-password-file>
   ```

3. Run the following command to mount s3fs:

   ```
   s3fs-bucket <your-local-folder> -o passwd_file=<your-password-file> -o url=https://s3.
   <region>.sv15.lyve.seagate.com
   ```

   where:

   - <your-password-file> is a file containing your S3 access and secret keys.
   - <your-local-folder> is the folder containing your files where you want to sync with S3 bucket.
   - <region> is the appropriate Lyve Cloud region, for example, us-east-1.

Once the mountpoint is working, sync all objects from s3fs-bucket to mount folder, for example:

```
# ls  -l s3-bucket
total 2197
-rw-r-----. 1 root root 1121955 Nov 19 12:43 Pic-1.png
-rw-r-----. 1 root root  479053 Nov 19 12:43 Pic-2.png
-rw-r-----. 1 root root  645935 Nov 19 12:43 Pic-3.png
-row-r--r--. 1 root root    1465 Nov 19 09:01 README.md
```

s5cmd

s5cmd is validated for use with Lyve Cloud Object Storage.

Configure S3cmd

To start using s5cmd, you'll need to configure it with your Lyve Cloud S3 storage service. Once s5cmd is installed, run the following commands to insert access keys into a credentials file:

1. Make sure you have the AWS directory to hold your credentials file in:

   ```
   $ mkdir ~/.aws
   ```

2. Create and edit the credentials file to contain your access key pair:

   ```
   $ vim ~/.aws/credentials
   ```

   ```
   [default]
   aws_access_key_id = <access-key>
   aws_secret_access_key = <secret-key>
   ```

Basic Commands

> **Note**—The Lyve Cloud --endpoint-url must appear before the command in order for the command to run.

List buckets

```
s5cmd --endpoint-url=https://s3.<region>.sv15.lyve.seagate.com ls
```

where:

- <region> is the appropriate Lyve Cloud region, for example, us-east-1.

Create bucket

```
s5cmd --endpoint-url=https://s3.<region>.sv15.lyve.seagate.com mb s3://s5cmd-bucket
```

where:

- <region> is the appropriate Lyve Cloud region, for example␣us-east-1.

Upload file

s5cmd --endpoint-url=https://s3.<region>.sv15.lyve.seagate.com cp envpod.yaml s3://s5cmd-bucket cp envpod.yaml s3://s5cmd-bucket/envpod.yaml

where:

- <region> is the appropriate Lyve Cloud region, for example␣us-east-1.

Download file

s5cmd --endpoint-url=https://s3.<region>.sv15.lyve.seagate.com cp s3://s5cmd-bucket/export_repo.sh /Users/660186/desktop/export.sh

where:

- <region> is the appropriate Lyve Cloud region, for example, us-east-1.

Sync a local folder to an S3 bucket:

s5cmd --endpoint-url=https://s3.<region>.sv15.lyve.seagate.com sync my-sync-folder s3://s5cmd-bucket
cp my-sync-folder/Pic-2.png s3://s5cmd-bucket/my-sync-folder/Pic-2.png
cp my-sync-folder/Pic-3.png s3://s5cmd-bucket/my-sync-folder/Pic-3.png
cp my-sync-folder/Pic-1.png s3://s5cmd-bucket/my-sync-folder/Pic-1.png

where:

- <region> is the appropriate Lyve Cloud region, for example, us-east-1.

Signiant Media Shuttle

Signiant Media Shuttle is a hybrid solution that includes host-based software SDCX and the web-based Media Shuttle user interface.

Pre-deployment Requirements

Before you begin the deployment workflow, ensure that your production environment meets Signiant's system requirements.

SDCX Installation

1. Download the Signiant app.

2. Install and launch SDCX server on Windows Server. It's recommended that the user restart the SDCX services after the initial installation.

3. Register the SDCX server

4. Restart the SDCX services

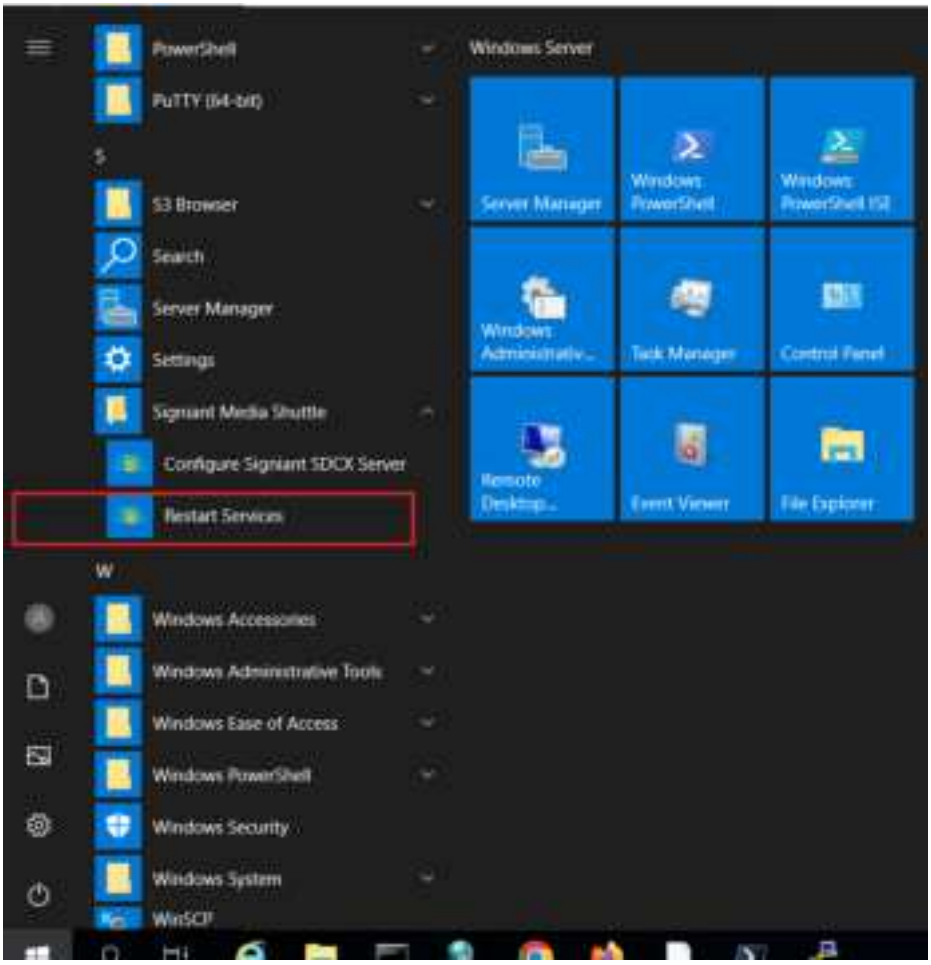

Launch SDCX services

1. Once you've launch SDCX server, select **Assign Media Shuttle portals to this server** at the localhost:8080/spring/install

> **Note**—The SDCX server is the one you installed on the host server.

2. Enter the Media Shuttle credentials you received from Signiant Support.

3. Create a portal at Media Shuttle. Select **Save Changes**.

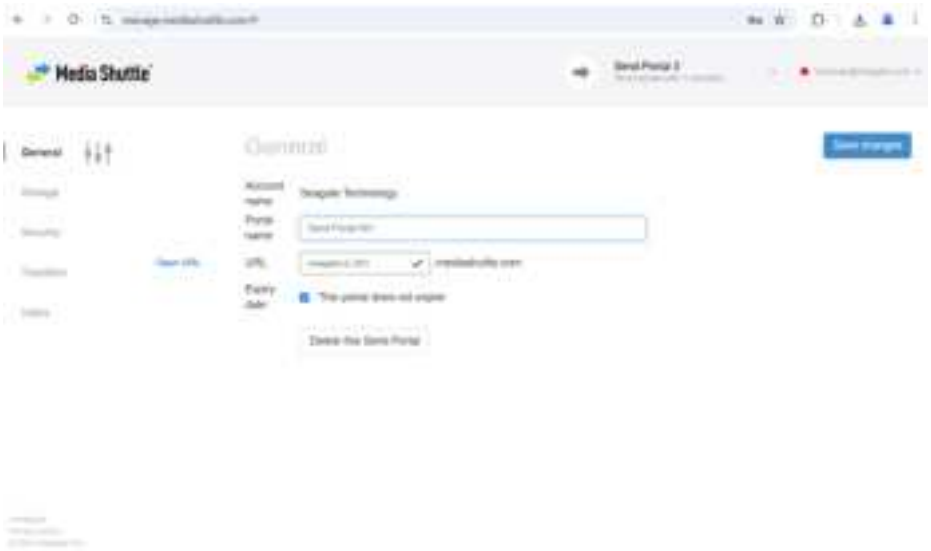

4. Add a S3 Compatible storage.

5. Create a S3 user credential in Signiant, using the Lyve Cloud credentials provided for the user account. In the example below, bucket signiant01 is created for a S3 account user at S3-us-west-1.sv15.lyve.seagate.com.

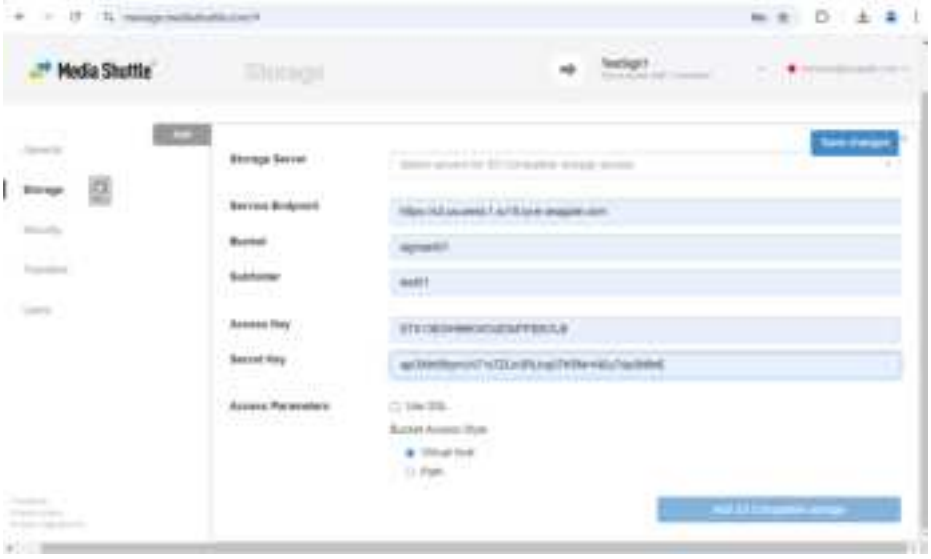

6. Assign the cloud storage to the local host. In the example below, signiant01 is assigned to hostD.

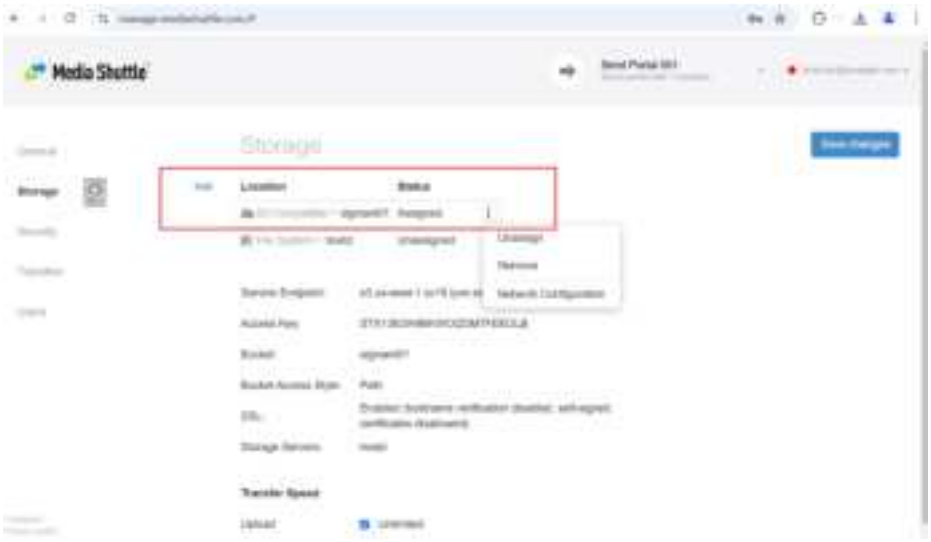

7. Send the portal that you created in the steps above.

8. Select the Arrow icon to notify the web service.

Object Upload and Download

Use Media Shuttle to upload or download file objects to and from Lyve Cloud object storage. In the following example, the cloud upload/download user interface is ready for further operations.

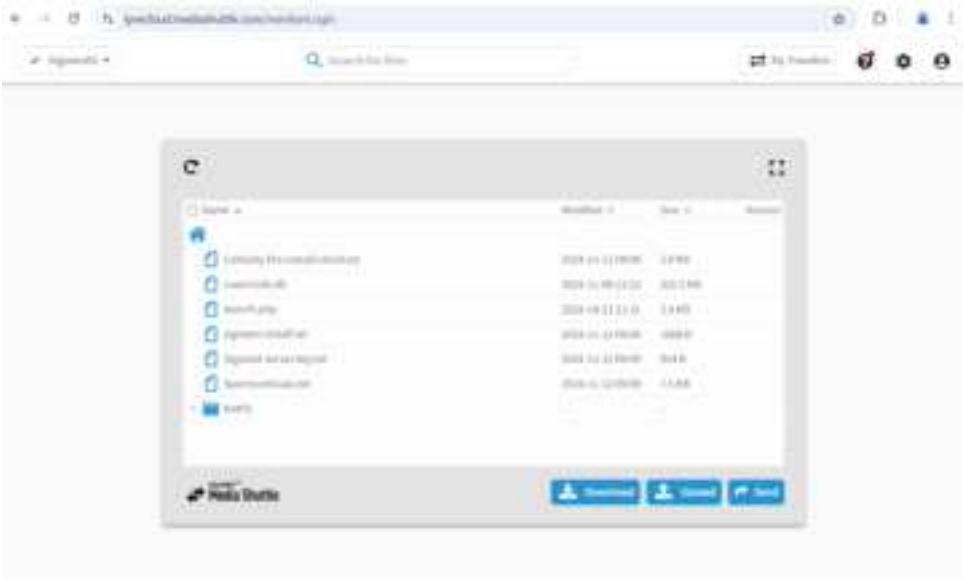

Upload example:

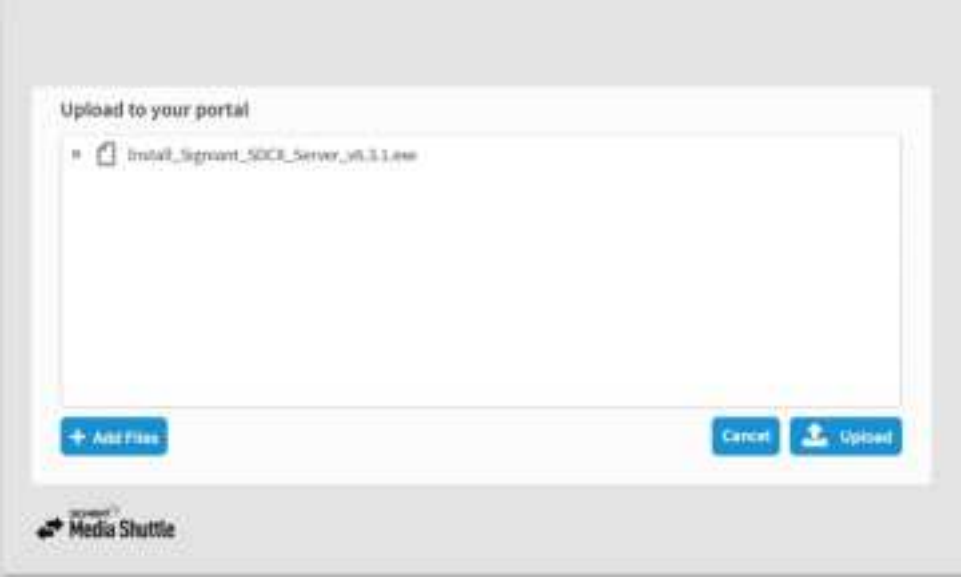

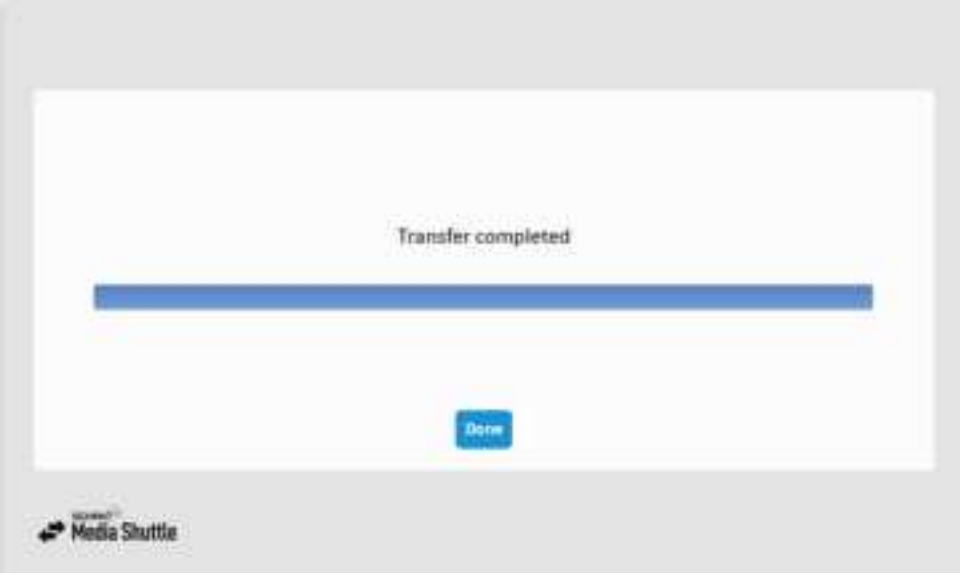

When uploading is completed, the registered user will receive an email alert from the Signiant service portal, notifying them that an upload has been received.

Download example:

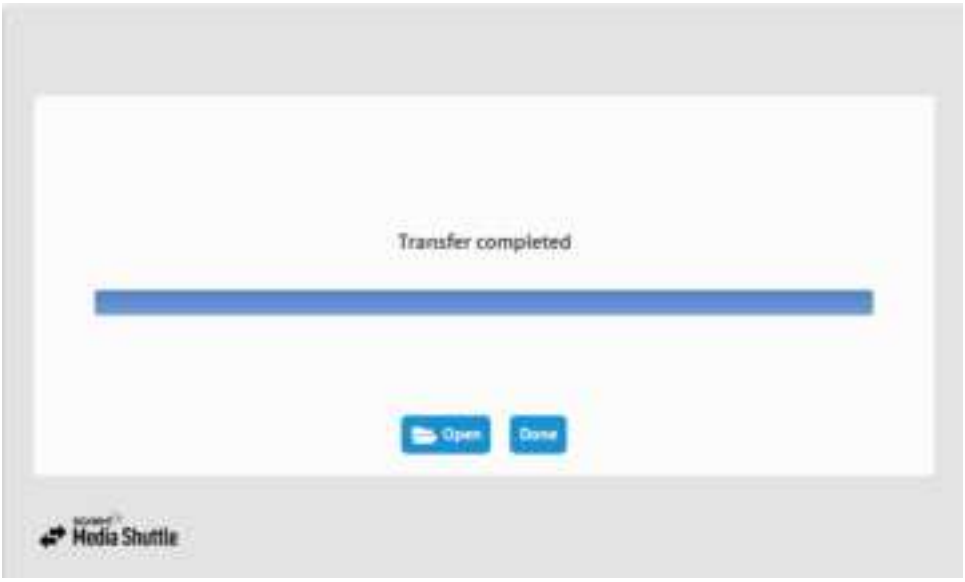


Signiant01
Your content has been downloaded by tony.tian@seagate.com.
This delivery contains the following content:
Install_Signiant_SDCX_Server_v6.3.1.exe
Total delivery size: 599.3 MB
Privacy Policy
91 Hartwell Avenue,
Lexington, MA 02421, USA
mediashuttle.com

Transmit

Transmit is validated for use with Lyve Cloud Object Storage.

Configure Transmit

1. Download transmit from the following link.
2. Add Amazon S3 Server. We are using Transmit 5.10.6.

- **Protocol:** Amazon S3
- **Address:** <region>.sv15.lyve.seagate.com
- **Port:** 443 (default)
- **Access Key ID:** <access-key>
- **Secret:** <secret-key>
- **Remote Path:** Accept default or specify another

where:

- <region> is the appropriate Lyve Cloud region, for example, us-east-1.
- <access-key> is your access key.
- <secret-key> is your secret key.

3. Drag and drop to upload or download files.

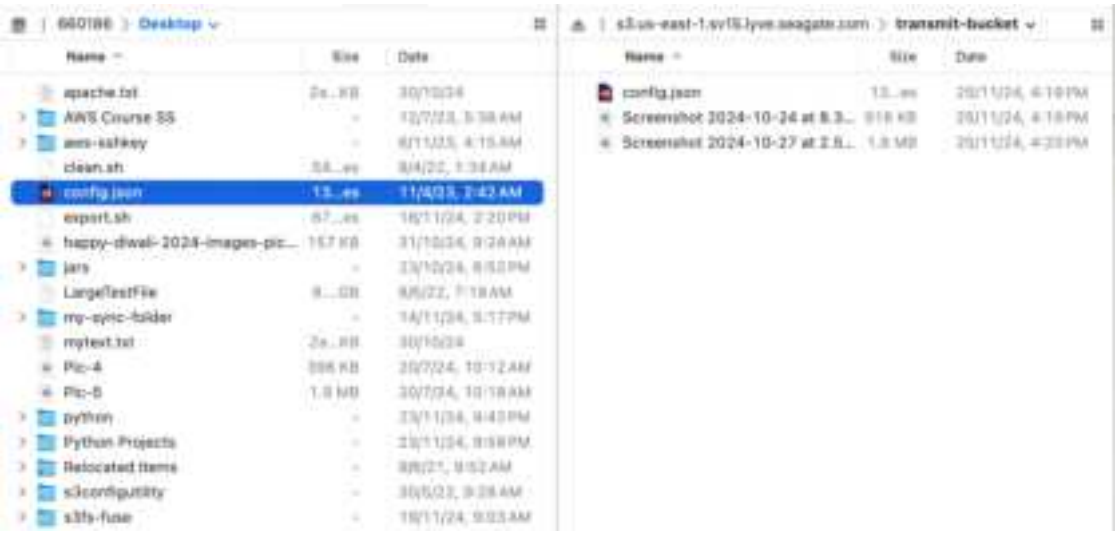
660186 Desktop
Name
Size
Date
apache.txt
AWS Course SS
aws-sshkey
clean.sh
config.json
export.sh
happy-diwali-2024-images-pic...
jars
LargeTestFile
my-sync-folder
mytext.txt
Pic-4
Pic-5
python
Python Projects
Relocated Items
s3configutility
s3fs-fuse
s3.us-east-1.sv15.lyve.seagate.com
transmit-bucket
Name
Size
Date
config.json
Screenshot 2024-10-24 at 8.3...
Screenshot 2024-10-27 at 2.5...