Seagate 2303us Deploy Parsec Labs with Lyve Cloud User Guide

Solution Integration Guide

Deploy Parsec Labs
with Lyve Cloud
Store and move massive volumes of data—affordably.

Contents

Challenge

A proven data-protection strategy used by most companies is the 3-2-1 backup rule, which states that you should have at least three copies of your data, on two different types of media, with at least one copy stored offsite.

Solution

Parsec Labs and Seagate Lyve® Cloud can inexpensively satisfy third copy and media diversity or provide a

fourth copy of file data as a failsafe.
And where traditional backup solutions move data into a proprietary format, data copied by Parsec Labs into Seagate Lyve Cloud is accessible using the standard S3 protocol.
Seagate Lyve Cloud is a simple, trusted, and efficient object storage solution for mass data. Predictable capacity-based pricing with no hidden fees for egress or API calls reduces TCO, so you'll never be surprised by your cloud bill. Put your data to work with verifiable trust and ease of use at scale from the global leader in data management.

Parsec Labs is the latest generation of high-performing, petabyte-scale data mobility offerings developed for the largest companies in America. Having proven itself at scale and in extreme use cases, Parsec Labs' data mobility offer brings unparalleled technical and financial results to the larger market.
To reap the benefits of this frictionless storage-only cloud, customers only need Lyve Cloud as their storage platform of choice on Parsec's centralized interface. Together, Lyve Cloud and Parsec provide a worry-free user experience for mass enterprise data mobility.

Benefits Summary

- Intelligent Data Mobility: Fluidly move data to where it's needed, when it's needed.
- Cost Efficient: Lyve Cloud with Parsec offers a budget-smart solution that scales to exabyte capacity with no vendor lock-in resulting from egress and S3 API fees. Lyve Cloud's transparent pricing structure allows enterprises to pay only for the storage needed.
- Expandable: Lyve Cloud and Parsec Labs enable petabyte-scale data mobility spanning the cloud and on-premise systems.

Deploying Lyve Cloud with Parsec Labs

Deployment Prerequisites

- Configured Lyve Cloud storage account
- Configured Parsec Labs account

Configuration Overview
The configuration for Lyve Cloud with Parsec Labs consists of three simple tasks.

- Task #1: Create and deploy a bucket and permissions to configure Lyve Cloud with Parsec Labs.
- Task #2: Create a new cloud storage source and target on the Parsec Labs account using information from Lyve Cloud.
- Task #3: Create cloud replication jobs using Seagate Lyve Cloud and Parsec Labs for fail-safe data protection.

Task #1: Deploy Lyve Cloud Bucket and Permissions

Step 1: Create Bucket
Go to the Bucket section of the Lyve Cloud console and select Create Bucket.

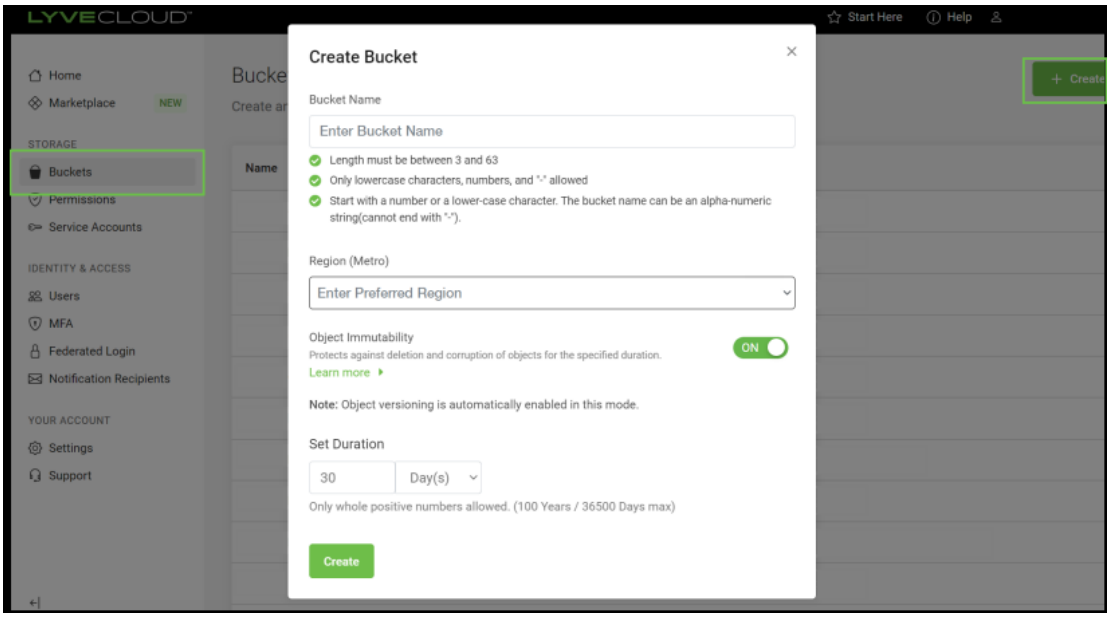

Step 2: Create Permissions
Go to the Permissions section of the Lyve Cloud console and select Create Bucket Permission.

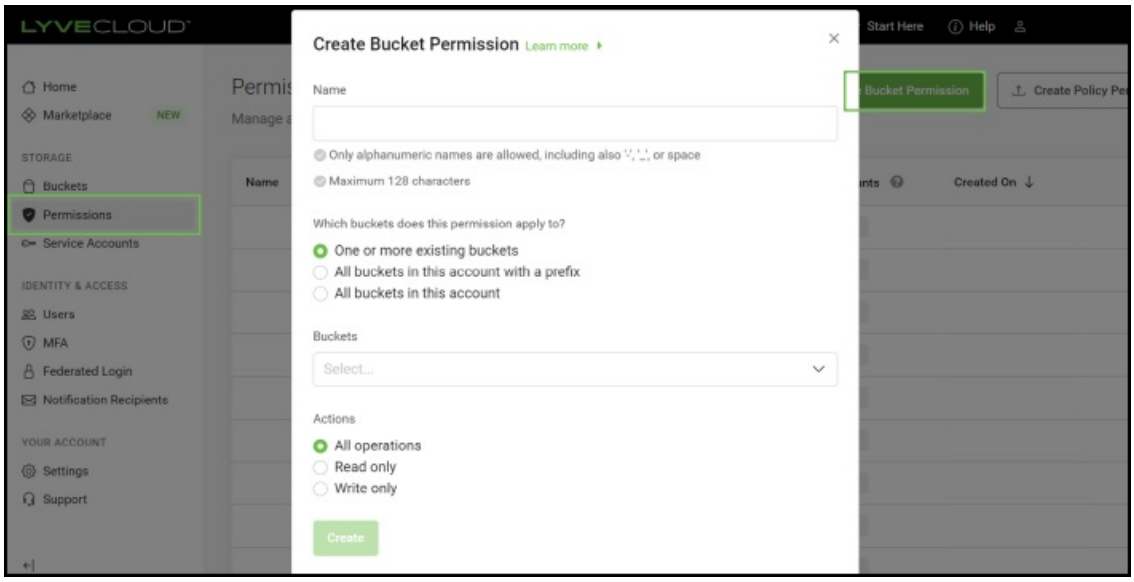

Note: You must select all buckets in this account with a prefix. This will allow Parsec Labs to create sub-buckets in Lyve Cloud.

Step 3: Create Service Account
Go to the Service Accounts section of the Lyve Cloud console and select Create Service Account.

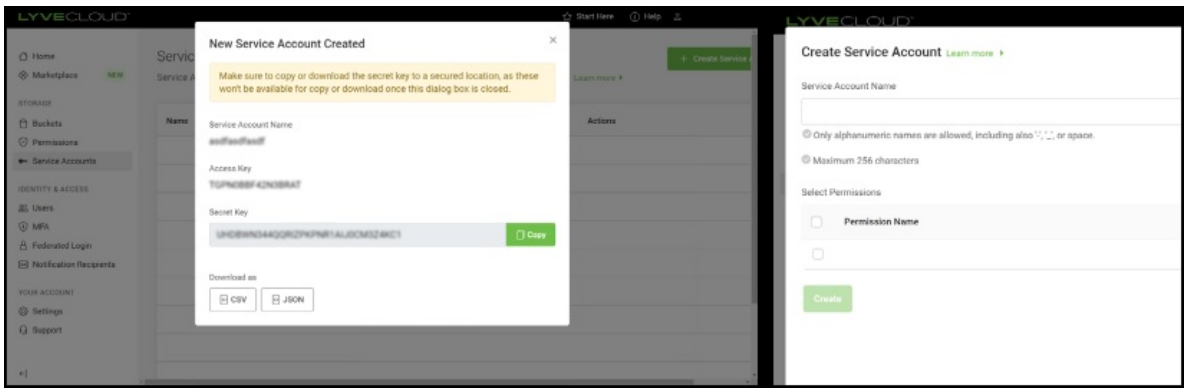

Task #2: Deploy Parsec with Lyve Cloud

Step 1: Enter the Parsec Job Console

Begin on the job console with the main menu on the side. Allocate a storage subsystem to migrate data from. Go to Storage and select an existing array or add a new one, providing a system name. For this exercise, select Local and add a new local NAS filer.

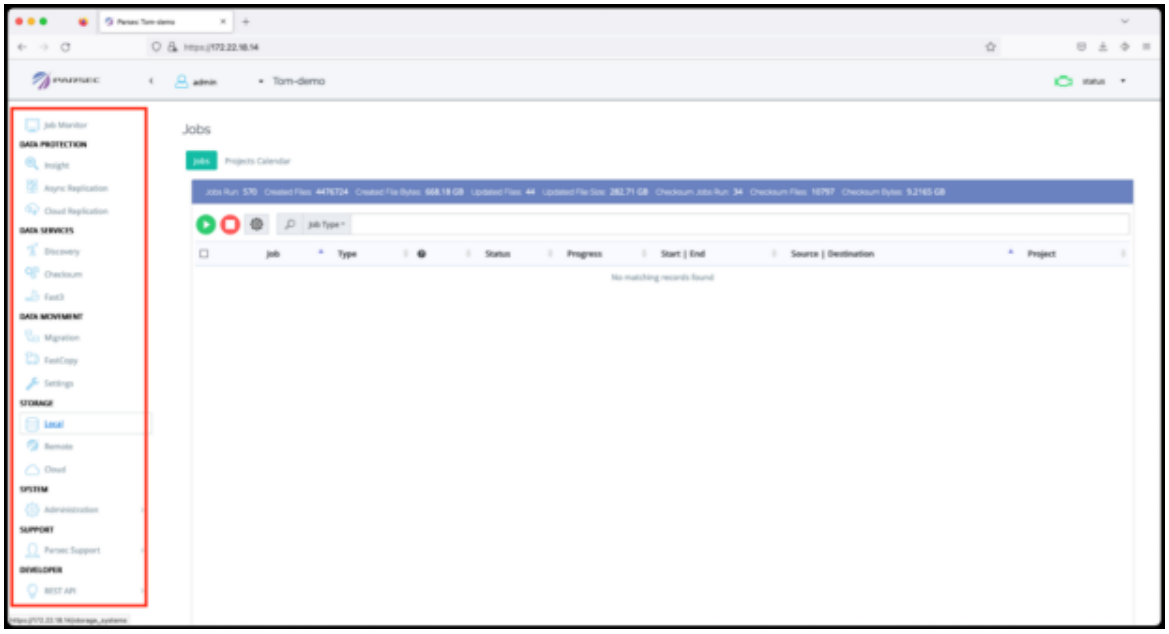

On the next screen, select Add Storage System and type the name of the system you intend to migrate or copy data from. In this scenario, Net App CDOT was added to the source list.

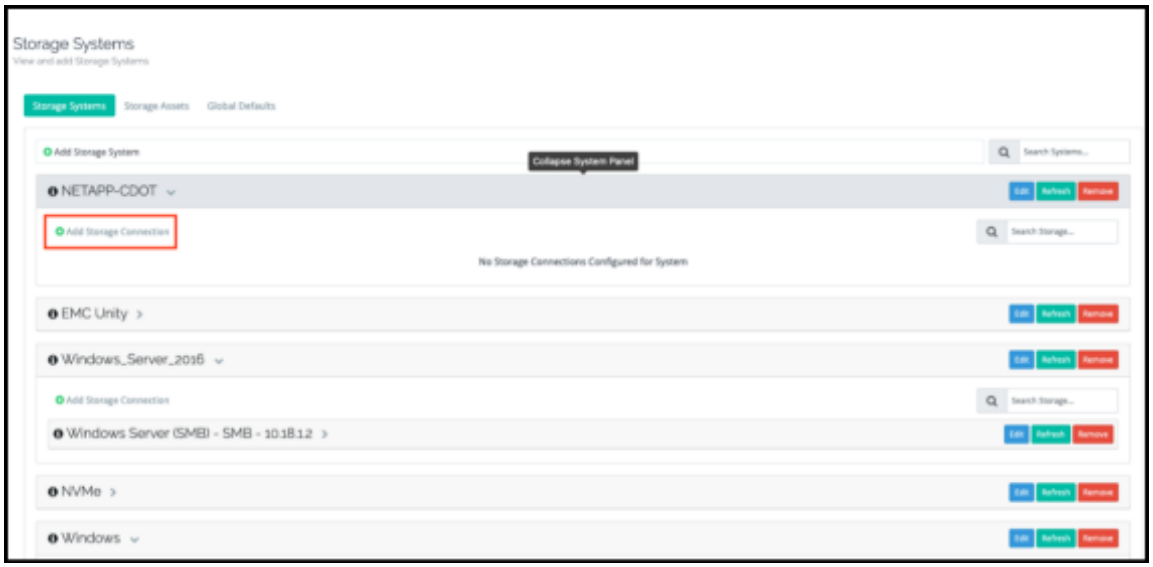

Step 2: Add a Storage Connection

Connect to the filer by clicking Add Storage Connection. This will connect to the filer management interface and be protocol specific. In the case of NetApp, it's a connection to SVM.

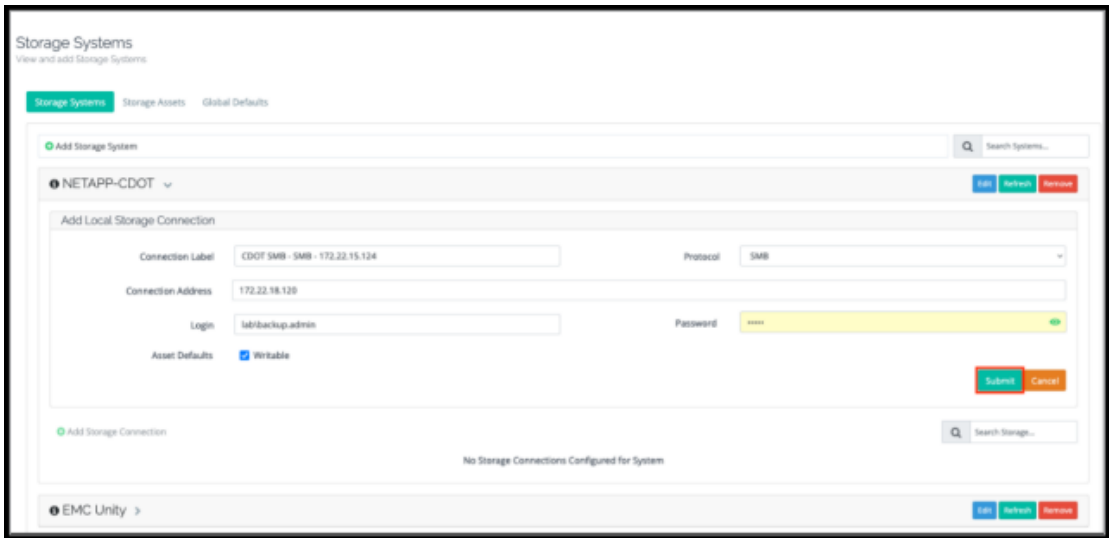

Provide a name (connection label) and select the connection protocol and IP address or fully qualified domain name.
For SMB connections, provide credentials for a member of the domain backup operators group. For NFS exports, the Parsec appliance IP must be on the export list.
Select Submit. At this point, the shares will be detected automatically and you'll see them listed on the screen.

Step 3: Add a Cloud Storage Host
On the main menu, go into Storage and select Cloud. Under Hosts, select Add Cloud Storage Host.

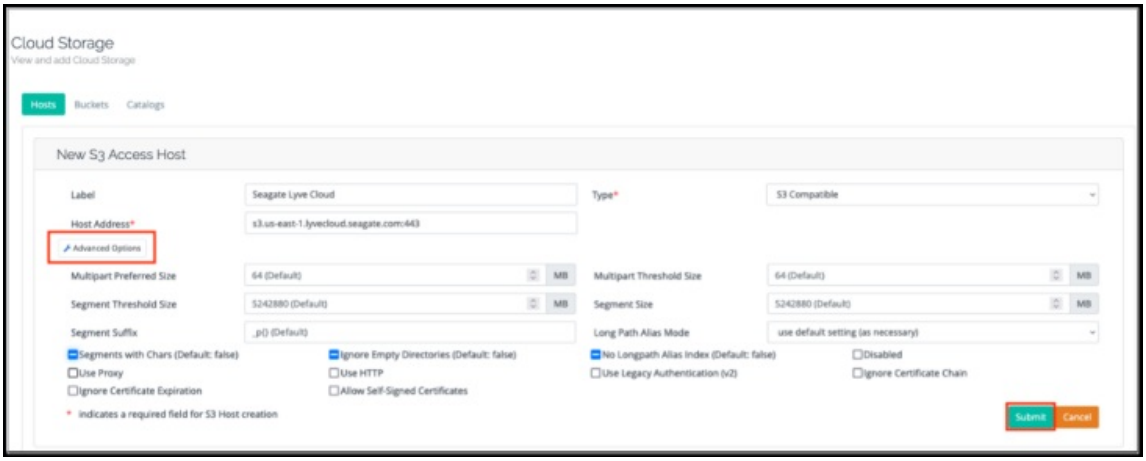

Select Advanced Options to gain access to long-path aliasing, adjusting multipart segment size, and configuring a proxy address.
Submit the configuration and hit Continue. After the host is added, select Add Account under the host name.

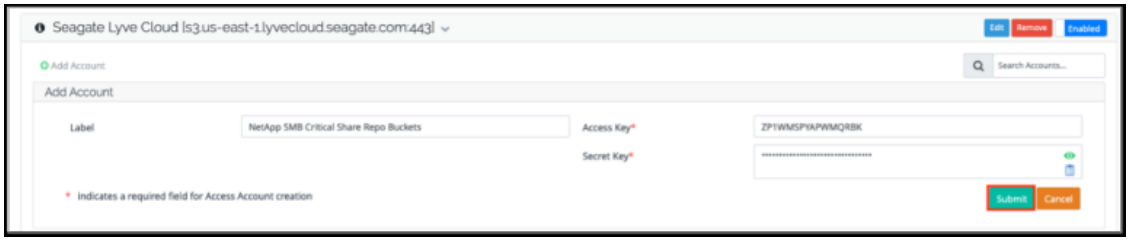

On the next screen, provide a label (any name) for the account and enter the access and secret keys. Click Submit. When the confirmation screen appears, click Continue.

After the account is created, select, Rescan for Buckets.
This completes the process of adding source and target. Now you can start protecting data with Lyve Cloud and Parsec.

Task #3: Create Cloud Replication Jobs Using Parsec Labs and Seagate Lyve Cloud for Fail-Safe Data Protection
Step 1: Create a job in Parsec to replicate a file share to Seagate Lyve Cloud S3 bucket. On the main menu under Data Protection, select Cloud Replication.
On the next screen, select Create New Project.

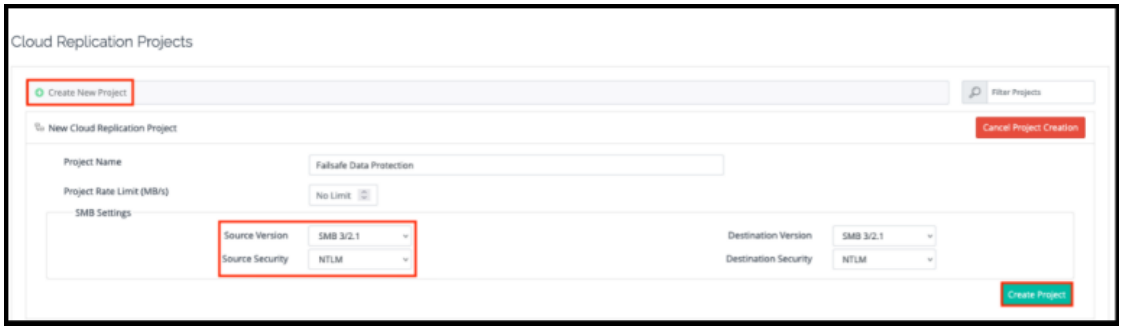

We'll call this project, Failsafe Data Protection. If the source share is an SMB share, select the SMB version (SMB 1, 2, 2.1, and 3) and security style (NTLM or Kerberos).

Step 2: Create a Project
Select Create Project.

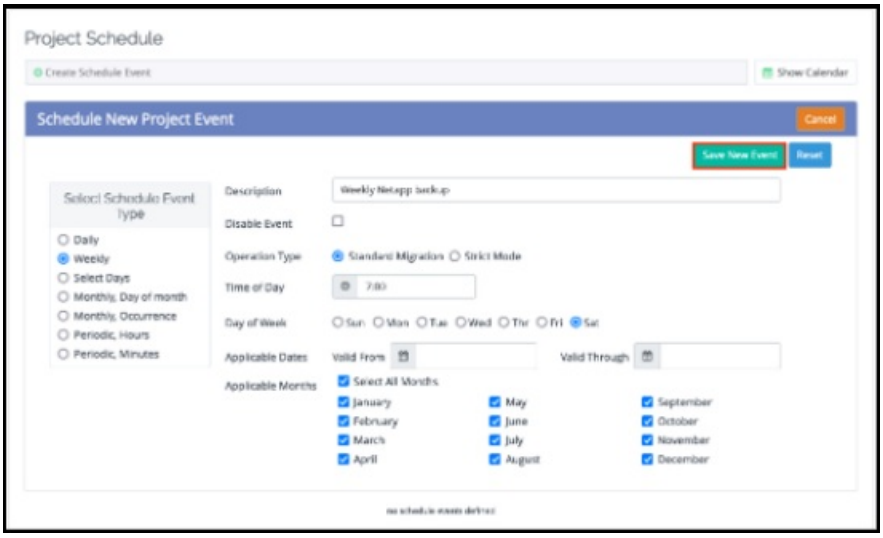

Schedules are set at the project level. Here we'll schedule jobs within our project to run every Saturday beginning at 7 a.m. and set a duration.
Within our new project, we'll create a job by choosing Create Job.

Give the job a name and select the source datastore by choosing in the source box.

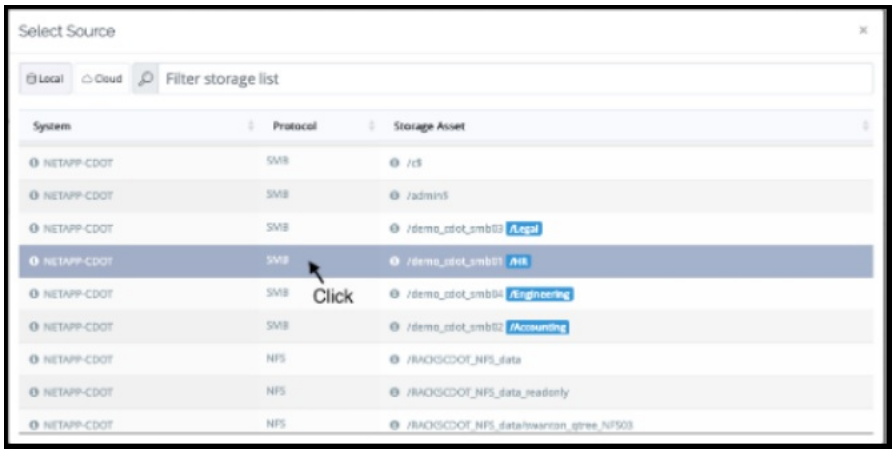

For this example, we'll select the share that's tagged /HR.

After selecting the source, the select destination dialog will automatically appear. Choose the desired destination to select it.
After selecting the source share and destination S3 bucket, you have the option to create include and exclude expressions based on certain metadata criteria. For instance, you may not want to include certain data or exclude from users, etc.

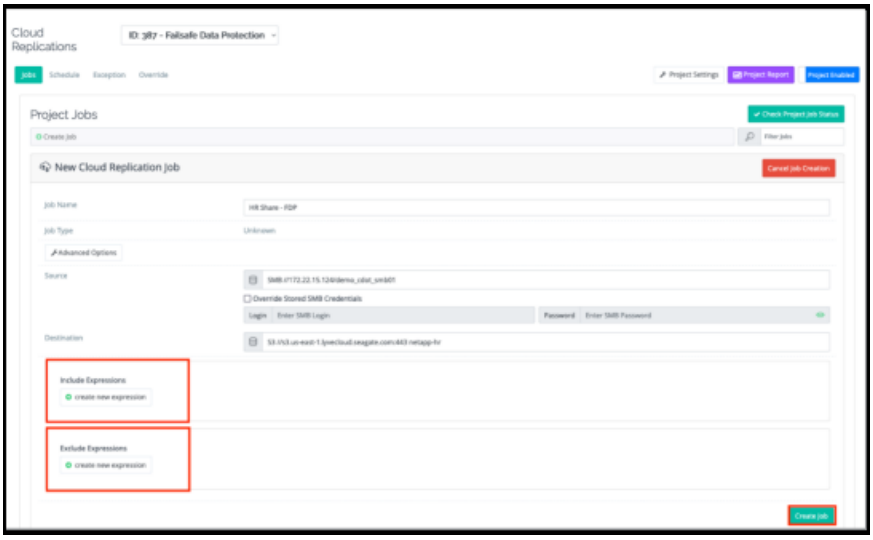

The job will now appear in the job monitor.

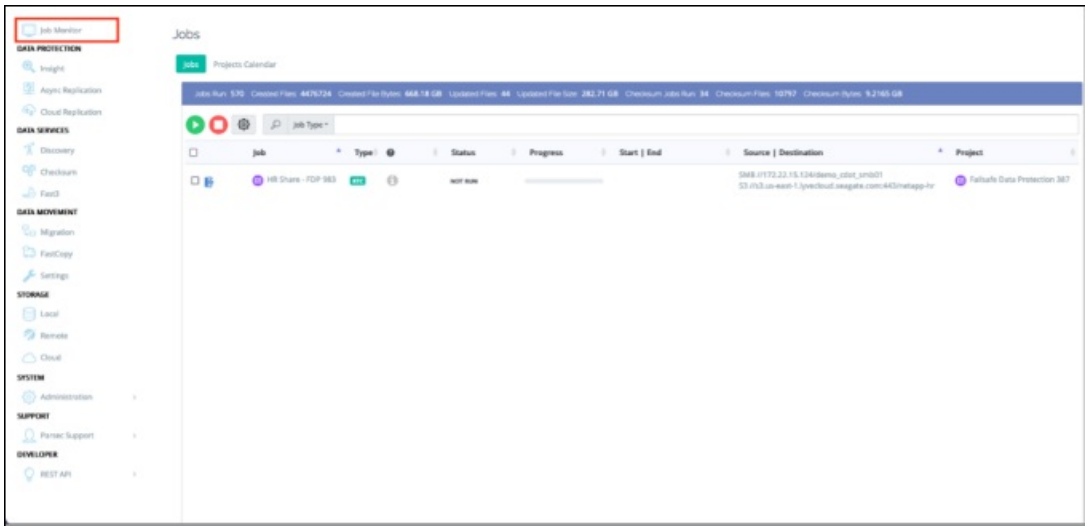

Even if a job is scheduled to run automatically, you can run it manually by selecting the left checkbox and clicking the green arrow.

After the job is completed and the share data has been replicated to the Lyve Cloud S3 bucket, you'll see the results show up in the job monitor menu.

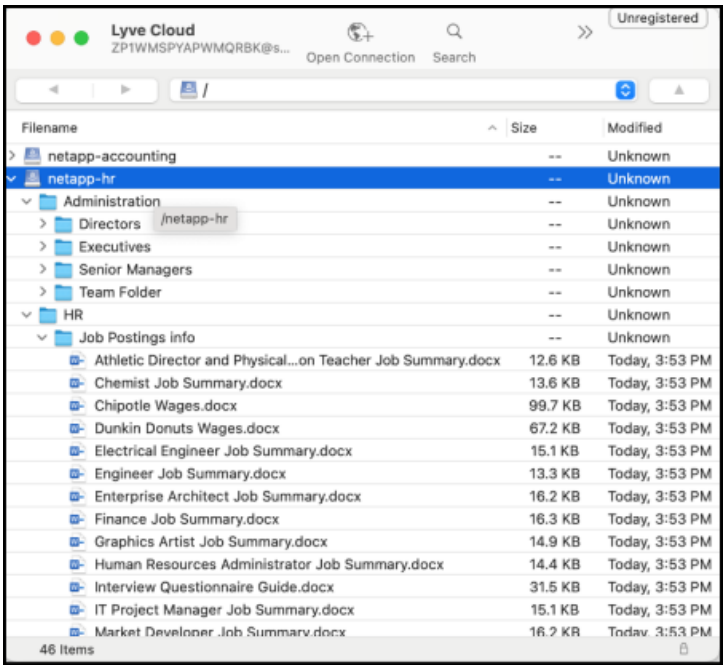

Our job has been completed and the share data has been replicated to the Lyve Cloud S3 bucket.

Conclusion

Enterprises today are overwhelmed with data, so the ability to easily search, sort, and move data between onpremise storage and the cloud is paramount for a high-functioning organization. With scalability comes cost

and the need to hit your target budget. This requires petabyte-scale solutions, offered at a price you can both predict and afford. Seagate Lyve Cloud and Parsec Labs deliver.

Ready to Learn More?
For more information on Parsec Labs, visit: **www.parseclabs.com**
For more information on Lyve Cloud, visit: **www.seagate.com**

seagate.com

Documents / Resources

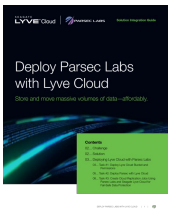	**Seagate 2303us Deploy Parsec Labs with Lyve Cloud** [pdf] User Guide 2303us Deploy Parsec Labs with Lyve Cloud, 2303us, Deploy Parsec Labs with Lyve Cloud, Lyve Cloud

References

- **The Leader in Mass Data Storage Solutions \| Seagate US**
- **Home \| Parsec Labs Data Migration for hyper scalers**
- **The Leader in Mass Data Storage Solutions \| Seagate US**
- **User Manual**