

Deploy Parsec Labs with Lyve Cloud

Store and move massive volumes of data—affordably.

Contents

02... Challenge

02... Solution

03... Deploying Lyve Cloud with Parsec Labs

03... Task #1: Deploy Lyve Cloud Bucket and Permissions

05... Task #2: Deploy Parsec with Lyve Cloud

08... Task #3: Create Cloud Replication Jobs Using Parsec Labs and Seagate Lyve Cloud for Fail-Safe Data Protection

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*.



The screenshot shows the Lyve Cloud console interface. On the left, a sidebar menu lists various sections: Home, Marketplace, STORAGE (highlighted), IDENTITY & ACCESS, and YOUR ACCOUNT. Under STORAGE, 'Buckets' is selected. A modal dialog titled 'Create Bucket' is open in the center. It contains the following fields and options:


- Bucket Name:** A text input field with placeholder text 'Enter Bucket Name'. Below it are three validation rules:
 - ✓ Length must be between 3 and 63
 - ✓ Only lowercase characters, numbers, and "-" allowed
 - ✓ Start with a number or a lower-case character. The bucket name can be an alpha-numeric string (cannot end with "-").
- Region (Metro):** A dropdown menu with placeholder text 'Enter Preferred Region'.
- Object Immutability:** A toggle switch labeled 'ON'.
- Note:** Object versioning is automatically enabled in this mode.
- Set Duration:** A text input field with '30' and a dropdown menu with 'Day(s)'.
- Footer:** A green 'Create' button.

In the background, the 'Buckets' section of the console is visible, showing a table with columns for Name and a '+ Create' button in the top right corner.



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.

The screenshot shows the 'Storage Systems' configuration page. At the top, there are tabs for 'Storage Systems', 'Storage Assets', and 'Global Defaults'. Below the tabs, there's a section for 'Add Storage System' with a search bar. A dropdown menu shows 'NETAPP-CDOT'. Below this, the 'Add Local Storage Connection' form is visible. It includes fields for 'Connection Label' (CDOT SMB - SMB - 172.22.15.124), 'Protocol' (SMB), 'Connection Address' (172.22.18.120), 'Login' (labibackup.admin), and 'Password' (masked with asterisks). There's also a checkbox for 'Asset Defaults' (Writable). At the bottom right of the form are 'Submit' and 'Cancel' buttons. Below the form, there's a message 'No Storage Connections Configured for System' and a search bar for 'Search Storage...'. At the bottom, there's a section for 'EMC Unity' with 'Edit', 'Refresh', and 'Remove' buttons.

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.

The screenshot shows the 'Cloud Storage' configuration page. At the top, there are tabs for 'Hosts', 'Buckets', and 'Catalogs'. Below the tabs, there's a section for 'New S3 Access Host'. It includes fields for 'Label' (Seagate Lyve Cloud), 'Host Address*' (s3.us-east-1.lyveccloud.seagate.com:443), and 'Type*' (S3 Compatible). Below these fields, there's a red box around the 'Advanced Options' link. The 'Advanced Options' section includes various settings: 'Multipart Preferred Size' (64 (Default) MB), 'Segment Threshold Size' (5242880 (Default) MB), 'Segment Size' (5242880 (Default) MB), 'Segment Suffix' (.p0 (Default)), 'Long Path Alias Mode' (use default setting (as necessary)), 'Segments with Chars (Default: false)', 'Ignore Empty Directories (Default: false)', 'No Longpath Alias Index (Default: false)', 'Use HTTP', 'Use Legacy Authentication (v2)', 'Ignore Certificate Expiration', 'Allow Self-Signed Certificates', 'Disabled', and 'Ignore Certificate Chain'. At the bottom right of the form are 'Submit' and 'Cancel' buttons. A footnote at the bottom left states: '* Indicates a required field for S3 Host creation'.



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.

The screenshot shows the 'Add Account' form in the Seagate Lyve Cloud interface. The form is titled 'Add Account' and has a search bar labeled 'Search Accounts...'. Below the search bar, there is a section for 'Add Account' with the following fields:

- Label:** A text input field containing 'NetApp SMB Critical Share Repo Buckets'.
- Access Key*:** A text input field containing 'ZP1WMSPYAPWMQRBK'.
- Secret Key*:** A text input field containing a masked string '*****'.

A red asterisk indicates a required field for Access Account creation. At the bottom right, there are 'Submit' and 'Cancel' buttons.

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.

The screenshot shows the 'NetApp SMB Critical Share Repo Buckets' page in the Seagate Lyve Cloud interface. The page has a search bar labeled 'Search Buckets...'. Below the search bar, there is a section for 'NetApp SMB Critical Share Repo Buckets [ZP1WMSPYAPWMQRBK]' with 'Edit', 'Remove', and 'Enabled' buttons. Below this, there is a section for 'Automatically Detected Buckets' with a 'Rescan for buckets' button highlighted in red. The 'Automatically Detected Buckets' section lists three buckets:

Bucket Name	Edit	Enabled
netapp-accounting	Edit	Enabled
netapp-hr	Edit	Enabled
netapp-legal	Edit	Enabled

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.

The screenshot shows the 'Cloud Replication Projects' page. At the top, there is a 'Create New Project' button highlighted with a red box. Below this is the 'New Cloud Replication Project' form. The 'Project Name' field is set to 'Failsafe Data Protection'. The 'Project Rate Limit (MB/s)' is set to 'No Limit'. Under 'SMB Settings', the 'Source Version' and 'Source Security' fields are highlighted with a red box; 'Source Version' is set to 'SMB 3/2.1' and 'Source Security' is set to 'NTLM'. The 'Destination Version' is set to 'SMB 3/2.1' and 'Destination Security' is set to 'NTLM'. A 'Create Project' button is at the bottom right, and a 'Cancel Project Creation' button is at the top right.

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.

The screenshot shows the 'Project Schedule' page. At the top, there is a 'Create Schedule Event' button. Below this is the 'Schedule New Project Event' form. The 'Description' field is set to 'Weekly Netapp backup'. The 'Disable Event' checkbox is unchecked. The 'Operation Type' is set to 'Standard Migration'. The 'Time of Day' is set to '7:00'. The 'Day of Week' is set to 'Sat'. The 'Valid From' and 'Valid Through' fields are empty. The 'Applicable Months' section shows a grid of months from January to December, all of which are checked. A 'Save New Event' button is at the top right, and a 'Cancel' button is at the bottom right.




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.

Cloud Replications

ID: 387 - Fallsafe Data Protection

JobsScheduleExceptionOverride

Project SettingsProject ReportProject Enabled

Project Jobs

Create jobEnter job

New Cloud Replication Job

Cancel Job Creation

Job NameHR Share - FDP

Job TypeUnknown

Advanced Options

SourceSMB //172.22.15.124/demo_cdot_smb01

Override Stored SMB Credentials

LoginEnter SMB LoginPasswordEnter SMB Password

DestinationS3 //s3.us-east-1.lyvecloud.seagate.com/443/netapp-hr

Include Expressions

Create new expression

Exclude Expressions

Create new expression

Create Job

The job will now appear in the job monitor.

Job Monitor

DATA PROTECTION

Insight

Async Replications

Cloud Replications

DATA SERVICES

Discovery

Checksum

Fast3

DATA MOVEMENT

Migration

FastCopy

Settings

STORAGE

Local

Remote

Cloud

SYSTEM

Administration

SUPPORT

Parsec Support

DEVELOPER

REST API

Jobs


Projects Calendar

Jobs Run: 570 Created Files: 4476724 Created File Bytes: 668.18 GB Updated Files: 44 Updated File Size: 282.71 GB Checksum Jobs Run: 34 Checksum Files: 10797 Checksum Bytes: 9.2165 GB


Job Type

Job	Type	Status	Progress	Start	End	Source	Destination	Project
HR Share - FDP 983	etc	NOT RUN				SMB //172.22.15.124/demo_cdot_smb01	S3 //s3.us-east-1.lyvecloud.seagate.com/443/netapp-hr	Fallsafe Data Protection 387

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 on-premise 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

© 2023 Seagate Technology LLC. All rights reserved. Seagate, Seagate Technology, and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Lyve is either a trademark or registered trademark of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. Seagate reserves the right to change, without notice, product offerings or specifications. SC8.1-2303US

