



# Enterprise-scale data storage and protection



Seagate Lyve Cloud with IBM Spectrum Protect



## About This Guide

This paper details the steps and best practices to deploy Seagate's Lyve Cloud, as a new cloud tier to an existing IBM Spectrum Protect Server.

## Audience

This paper is written for storage and backup administrators familiar with administering and managing backup environments.

# CONTENTS

- 3 CHALLENGE SUMMARY
- 3 SOLUTION SUMMARY
- 4 BENEFITS OF USING SEAGATE LYVE CLOUD WITH IBM SPECTRUM PROTECT
- 5 DEPLOYING SEAGATE LYVE CLOUD WITH IBM SPECTRUM PROTECT
- 6 TASK 1: CREATE A LYVE CLOUD SERVICE ACCOUNT
- 7 TASK 2: ADD A CLOUD CONTAINER POOL TO IBM SPECTRUM PROTECT

## Challenge Summary

From exponential data growth and new workloads, to the rise of cloud, and the risk of cyberattack, businesses must respond to a fast-changing landscape of threats to business continuity. Data is vulnerable, and as ransomware attacks hit the headlines, there has never been a more urgent time to move beyond back-up with key capabilities to drive simplicity, security, and compliance.

Organizations need storage and data protection that seamlessly scales and extends across the entire IT infrastructure; gives control over cost and performance, avoids complexity, and improves resiliency against cyber threats. Seagate Lyve Cloud with IBM Spectrum Protect combines Seagate's best-in-class data storage performance with market-leading IBM multi-workload data protection and management. The combined solution delivers predictable cloud storage economics, data mobility, cyber resiliency, and an easy-to-use gateway for efficient ingestion and backup retrieval from Lyve Cloud.

## Solution Summary

The Lyve Cloud with IBM Spectrum Protect solution addresses exponential data growth, and the need for a solution that scales seamlessly, while delivering a simplified and more cost-efficient approach to data storage protection.

The joint solution enables advanced data protection, securing applications and workloads on hybrid, multi-cloud architectures, centralizes storage management, and reduces cost. Seagate Lyve Cloud's data storage focuses on security, flexibility, and predictable economics. Integrated with IBM Spectrum Protect's built-in cloud integration and data efficiency capabilities, it offers security-rich and cost effective data protection and back-up in the cloud.

## Benefits of using LyveCloudwithIBM Spectrum Protect

Seagate Lyve Cloud and IBM Spectrum Protect is the easy choice for accelerating cloud adoption and modernizing data protection. With IBM Spectrum Protect and Seagate Lyve Cloud, storage admins get access to:

**Simplified Backup and Recovery**—Lyve Cloud with IBM Spectrum Protect simplifies backup and recovery with fast, simple, and flexible recovery that minimizes data loss. Rapid recovery with no need for data hydration ensures access to data is restored almost immediately. Utilize Seagate's best-in-class Lyve Cloud Storage Object, S3-compatible data storage for always on availability to facilitate easy back-up and reliable retrieval without wait.

**Business Continuity**—Ensure business continuity with Seagate Lyve Cloud's high-availability cloud storage management. Easy-to-understand, simple pricing with no API charges or egress fees empowers customers to back up and move all back-up data in a single repository for compliance without lock-in concerns.

**Multi-Workload Protection**—Lower operational cost with flexible, scalable storage options, unifying and simplifying data protection for file servers, virtual environments, applications, and data sent from API or S3.

**Scalable Performance**—Scale up to support massive data growth with flexible, scalable, storage options and the capacity to manage billions of objects per backup server, with less disruption and complexity, as backup workloads grow.

**Storage Efficiency**—Drive exceptional storage efficiency with incremental "forever" backup version protection, compression, deduplication, and policy-based life cycle management.

**Cyber Resiliency**—Protect your storage investment with encryption, always-on data monitoring, secure communication, and proactive security notifications.

**Hybrid, Multi-cloud Performance**—Manage complexity across hybrid environments through a unified platform that simplifies management, and gives you control of your data wherever it resides. Back-up to cloud or back-up in the cloud with security-rich built-in cloud integration.



# Deploying Seagate Lyve Cloud with IBM Spectrum Protect

## Deployment Prerequisites

Lyve Cloud Storage Account, this includes:

1. Obtain Access and Secret Keys for the storage account.
  - Have the ability to read/write/list and create buckets and objects and the ability to delete objects
2. IBM Spectrum Protect Account
  - Follow IBM Spectrum Protect Best Practices for your workload and environment

## Configuration Overview

The configuration for Lyve Cloud with IBM Spectrum Protect is divided into 2 simple tasks.

- Task 1: Create a Lyve Cloud Service Account. For more information, see the Lyve Cloud reference guide: [Lyve Cloud Quick Start Guide](#)
- Task 2: Add a cloud container storage pool to IBM Spectrum Protect. For more information, see the Spectrum Protect reference guide: [Configuring a Cloud Container Storage Pool](#)

## Task 1: Create Lyve Cloud Service Account

It is assumed that a Lyve Cloud Storage Account has been created and configured. This consists of the following steps:

- Set up an S3 Bucket
- Create Bucket Permissions (*Do not enable compliance mode/immutability - this feature is not currently supported by IBM Spectrum Protect.*)
- Create a Service Account - establishing access key, secret key, and associated URL for your account; see example below

### Create Service Account

In the Lyve Cloud Console, select “Service Accounts.” On the Service Accounts page, select the “Create Service Account” button.

The screenshot displays the Lyve Cloud console interface. On the left sidebar, the 'Service Accounts' menu item is highlighted with a red circle 1. The main content area shows the 'Create Service Account' form with a red circle 2 highlighting the 'Create Service Account' button. The form includes a 'Service Account Name' field with the value 'company-app1-key' and a 'Select Permissions' section where 'Master permission' is selected. A 'New Service Account Created' dialog box is open, showing the 'Access' key '2LAJ00UZK00DATAL' and the 'SECRET' key 'AOAG2DM3I0IEUWG1VT3L3YUATIF0VIJ5'. A red circle 3 highlights the dialog box. A yellow callout box indicates that the user must copy or download the secret information before closing the dialog. The dialog also includes a 'Copy' button and 'Download as' options for CSV and JSON.

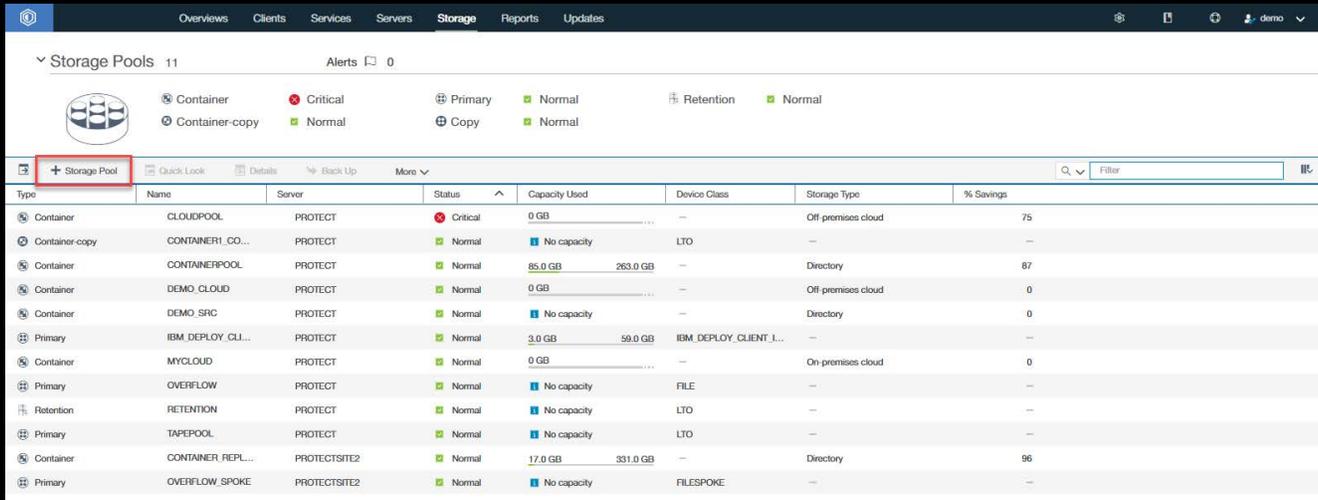
*Note: You must have at least one bucket with at least one associated permission before you can establish the credentials needed to add Lyve Cloud, as a new Cloud Tier*



## Task 2: Add a Cloud Tier on IBM Spectrum Protect Server

### Step 2: Add Storage Pool

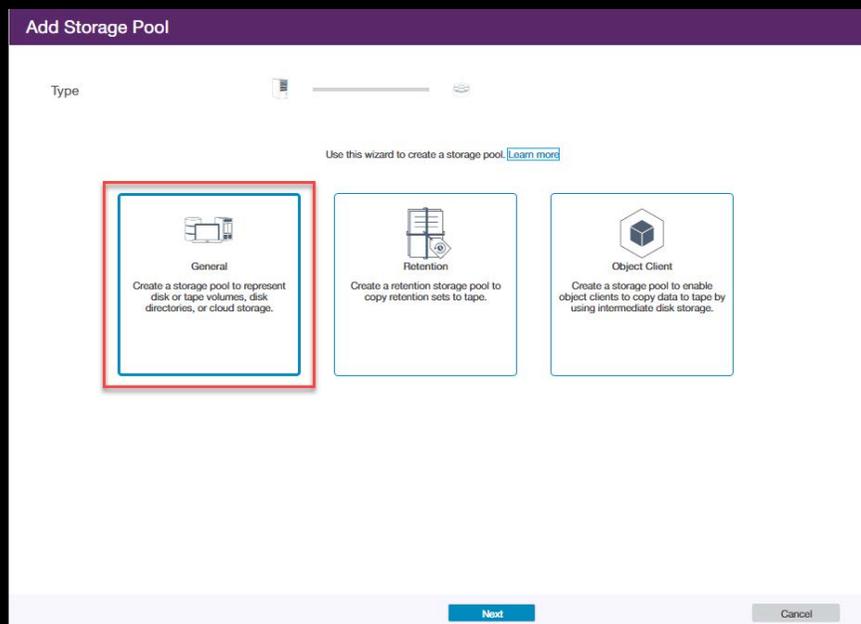
On the "Storage Pool" page of the IBM Spectrum Protect Operations Center, select "+ Storage Pool" to start the storage pool wizard, and complete the steps to create a storage pool.



| Type           | Name              | Server       | Status   | Capacity Used      | Device Class           | Storage Type       | % Savings |
|----------------|-------------------|--------------|----------|--------------------|------------------------|--------------------|-----------|
| Container      | CLOUDPOOL         | PROTECT      | Critical | 0 GB               | ---                    | Off-premises cloud | 75        |
| Container-copy | CONTAINER1_CO...  | PROTECT      | Normal   | No capacity        | LTO                    | ---                | ---       |
| Container      | CONTAINERPOOL     | PROTECT      | Normal   | 85.0 GB / 263.0 GB | ---                    | Directory          | 87        |
| Container      | DEMO_CLOUD        | PROTECT      | Normal   | 0 GB               | ---                    | Off-premises cloud | 0         |
| Container      | DEMO_SRC          | PROTECT      | Normal   | No capacity        | ---                    | Directory          | 0         |
| Primary        | IBM_DEPLOY_CLI... | PROTECT      | Normal   | 3.0 GB / 59.0 GB   | IBM_DEPLOY_CLIENT_L... | ---                | ---       |
| Container      | MYCLOUD           | PROTECT      | Normal   | 0 GB               | ---                    | On-premises cloud  | 0         |
| Primary        | OVERFLOW          | PROTECT      | Normal   | No capacity        | FILE                   | ---                | ---       |
| Retention      | RETENTION         | PROTECT      | Normal   | No capacity        | LTO                    | ---                | ---       |
| Primary        | TAPEPOOL          | PROTECT      | Normal   | No capacity        | LTO                    | ---                | ---       |
| Container      | CONTAINER_REPL... | PROTECTSITE2 | Normal   | 17.0 GB / 331.0 GB | ---                    | Directory          | 96        |
| Primary        | OVERFLOW_SPOKE    | PROTECTSITE2 | Normal   | No capacity        | FILESPOKE              | ---                | ---       |

### Step 3: Create Storage Pool

In the "Add Storage Pool Wizard", select the type as "General" to configure a cloud container storage pool.



Add Storage Pool

Type

Use this wizard to create a storage pool. [Learn more](#)

- General**  
Create a storage pool to represent disk or tape volumes, disk directories, or cloud storage.
- Retention**  
Create a retention storage pool to copy retention sets to tape.
- Object Client**  
Create a storage pool to enable object clients to copy data to tape by using intermediate disk storage.

Next Cancel

✓ Select "General" to configure a cloud container storage pool.



## Task 2: Add a Cloud Tier on IBM Spectrum Protect Server

### Step 4: Create Cloud Container Identity

At the "Identity" step of the wizard, specify a name for the storage pool and the server, as an identifier for the Storage Pool.

The screenshot shows the 'Add Storage Pool' wizard at the 'Identity' step. The title bar is purple and says 'Add Storage Pool'. Below it, a progress bar shows 'PROTECT' selected. The main content area has the heading 'Identity' and a sub-heading 'PROTECT'. Below that, it says 'Create a storage pool to store client data. Learn more'. There are three input fields: 'Name' with the value 'S3-CLOUD', 'Server' with a dropdown menu showing 'PROTECT', and 'Description' which is empty. A tooltip points to the 'Name' field with the text 'The pool name must be unique for the selected server. Learn more'. Another tooltip points to the 'Server' dropdown with the text 'Select server from drop down'. At the bottom, there are 'Back', 'Next', and 'Cancel' buttons.

✓ Specify the Name and Server, as an identifier for the Storage Pool.

### Step 5: Specify Storage Pool Type

At the "Type" step of the wizard, select off-premises cloud to configure a cloud-container storage pool in Lyve Cloud.

The screenshot shows the 'Add Storage Pool' wizard at the 'Type' step. The title bar is purple and says 'Add Storage Pool'. Below it, a progress bar shows 'PROTECT' and 'S3-CLOUD' selected. The main content area has the heading 'Type' and a sub-heading 'PROTECT S3-CLOUD'. Below that, it says 'Choose the type of pool that best supports your business goals. Learn more'. There is a blue information icon followed by the text 'To copy data from an existing directory-container pool, cancel the wizard, select the pool, and click More > Add Container-copy Pool.' There are two main sections: 'Container-based storage' and 'Traditional volume-based storage'. Under 'Container-based storage', there are three radio buttons: 'Directory' (File-based storage on disk with optional copy pools), 'On-premises cloud' (Object-based storage that is managed by internal IT staff in your data center. For example, IBM Cloud Object Storage and other certified S3 providers), and 'Off-premises cloud' (Storage in vendor-managed repositories, using IBM Cloud, OpenStack Swift, Amazon S3, or Microsoft Azure). The 'Off-premises cloud' option is selected and highlighted with a blue box. Under 'Traditional volume-based storage', there are three radio buttons: 'Disk (primary)' (Storage on disk or in a mountable deduplicating appliance), 'Tape (primary)' (Storage on tape or in a deduplicating VTL), and 'Tape (copy)' (Copies of primary storage on tape or in a VTL). At the bottom, there are 'Back', 'Next', and 'Cancel' buttons.

✓ Select "Off-premises Cloud" to configure a cloud-container storage pool in Lyve Cloud

## Task 2: Add a Cloud Tier on IBM Spectrum Protect Server

### Step 6: Add Credentials

At the "Credentials" step of the wizard, enter connection information.

The screenshot shows the 'Add Storage Pool' wizard at the 'Credentials' step. The progress bar indicates that the 'CREDENTIALS' step is active. The 'PROTECT' and 'S3-CLOUD' options are visible. The instructions state: 'Select the cloud type and enter connection information for accessing the cloud. [Learn more](#)'. The form fields include: Pool type (Off-premises cloud), Encryption (checked), Cloud type (Amazon - S3 API), Access key ID, Secret access key, Bucket name, Region, URL, and Cloud storage class (S3 Standard, S3 Standard, S3 Intelligent-Tiering). Two red boxes highlight the 'Secret access key' field with the text 'Add Lyve Cloud Access Key and Secret Key' and the 'Cloud storage class' dropdown with the text 'Select cloud storage class based on your use case.' The bottom navigation bar contains 'Back', 'Next', and 'Cancel' buttons.

✓ Specify credentials for connection information

### Step 7: Add Local Storage

At the "Local Storage" step of the wizard, specify existing file system directories for disk storage.

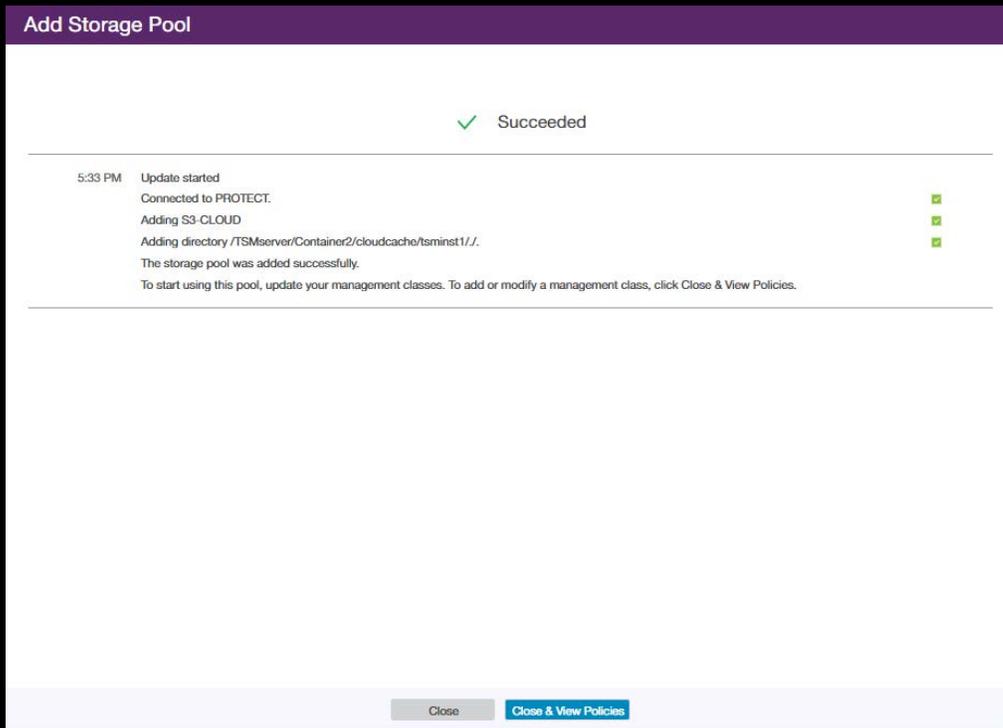
The screenshot shows the 'Add Storage Pool' wizard at the 'Local Storage' step. The progress bar indicates that the 'LOCAL STORAGE' step is active. The 'PROTECT' and 'S3-CLOUD' options are visible. The instructions state: 'Specify one or more existing directories where S3-CLOUD can temporarily store data before it is transferred to the cloud. Local storage is not required if the pool is only used as a tiering target. If data is backed up directly to the pool, local storage is required and can improve performance. [Learn more](#)'. The 'Directories' section has a text input field with a plus sign to its right. The bottom navigation bar contains 'Back', 'Add Storage Pool', and 'Cancel' buttons.

✓ Specify existing file system directories for Local Storage

## Task 3: Add a Cloud Tier on IBM Spectrum Protect Server

### Step 6: Finish Working with Wizard

At the Summary step of the wizard, review configuration.



Review the "Storage Pool Configuration." Close the wizard, and view "Policies" to start using the new data pool.

## Summary

The Lyve Cloud with IBM Spectrum Protect solution delivers enterprise-scale data storage and protection with frictionless movement, storage, backup, and fast recovery, in one limitlessly, scalable package. Easy-to-understand, simple pricing frees customers from lock-ins and egress fees, and allows enterprises to meet the challenge of storing and protecting all their data.



## Ready to Learn More?

For more information on Lyve Cloud,  
visit: <http://www.seagate.com/lyvecloud>

For more information on IBM Spectrum Protect visit:  
<http://www.ibm.com/products/data-protection-and-recovery>