

FEBRUARY - MAY 2017



PROOF OF CONCEPT AND CASE STUDY

IBM Spectrum Protect and Backing up to Object Storage in the Cloud

Andrew Wojnarek | Sr. Systems Engineer

ATS Innovation Center, Malvern PA

© Copyright Advanced Technology Services Group 2017



Introduction

The ATS Group is a business partner specializing in the full IBM enterprise stack, and IBM Spectrum Protect in particular.

We've helped customers set up IBM Spectrum Protect environments that backup petabytes of data in a way that's robust and actually works at scale. There has been an explosion of the cloud as well as a downpour of people talking cloud migration – and there's nothing easier than moving your backups to the cloud.

In this PoC we set up the following:

- AIX 7.2
- Power 8 S822
- IBM Spectrum Protect 8.1
- 3 clients: AIX, Linux and Windows

The goal of this PoC was the following:

- Install IBM Spectrum Protect.
- Figure out how IBM Spectrum Protect tiers backups to a cloud.
- Simulate a customer's environment.



The Hybrid Cloud Backup Architecture

One of the biggest factors driving IT spending and business innovation is cost. Reduction in cost is a key motivator behind evaluations and new purchases. This is where protecting your data in a cloud can be very beneficial. Backing up and protecting an organizations data is one of the top use cases for public cloud services. Hybrid cloud solutions are sought after as the next step, since many organizations already have investments in data centers and can utilize their existing infrastructure while leveraging the power of the cloud.

When your organization is considering that next tape library upgrade or extension, or ordering another round of tapes – first consider a cloud based DR solution. We have found that a cloud DR solution combines the best of disk/tape solutions:

- 1. Long distance data retention**
- 2. Access to data using a modern interface for better analytics / modeling**
- 3. Reduced capital purchases**

IBM Spectrum Protect fits into this hybrid cloud backup architecture perfectly. Its modern data protection fits enterprise needs and supports storage targets for Disk, Tape, object storage and cloud. Its virtualization is optimized and application aware. IBM Spectrum Protect is also adopting more native cloud support each quarter (recently AWS and Azure), as well as private cloud adoption. Another great feature that can fit well in a cloud computing environment is the deduplication and incremental backups. It allows your data to be backed up more efficiently, and thus save you more money.



Description

We successfully got the standard stack up and running quickly, running IBM Spectrum Protect on AIX 7.x. We established some dummy hosts and backed up some data we created. We also built a filesystem called /cloudcache:

```
/dev/cloudcachelv 254720.00 20305.72 93% 2356 1% /cloudcache
```

We then created a Storage Pool that is connected to **Object Storage** out in the cloud:

The screenshot shows a configuration wizard with two main options: Container-based storage and Traditional volume-based storage. Under Container-based storage, 'Off-premises cloud' is selected. Under Traditional volume-based storage, 'Disk (primary)' is selected. A message at the top says: 'Choose the type of pool that best supports your business goals. Learn more' and 'To copy data from an existing directory-container pool, cancel the wizard, select the pool, and click More > Add Container-copy Pool.'

After entering a username, APIkey and the URL, it **connected** immediately:

The screenshot shows a summary bar for the storage pool configuration. It includes: 'Container' (selected), 'STGAWS.02', 'SP81_POC', 'Normal' (checked), and '0 GB'.



What we learned is that the general process for backing up to the cloud **follows these steps:**

- 1 A client backs up directly to IBM Spectrum Protect via the filesystem /cloudcache.
- 2 This cache fills up and empties at scheduled times during the day.
- 3 The system copies the data up to Object Storage.
- 4 Data is drained from the filesystem.

This process is roughly what happens daily as you utilize IBM Spectrum Protect and Object Storage. When you create the Storage Pool, the option to encrypt the data at rest is already set.

This is the way the data looks like as it sits **on Object Storage:**

<input type="checkbox"/>	Name	Last modified	Size	Storage class
<input type="checkbox"/>	000000000000014a.dcf	Jan 18, 2017 8:03:54 AM	15.6 MB	Standard

Since the S3 endpoint is over HTTPS, the data leaves IBM Spectrum Protect encrypted, and sits on disk inside the Object Storage bucket encrypted. In this way, data is continually encrypted and protected throughout the process.



With Object Storage, one of the neat things is the control you have over pricing. It's so easy to compare between vendors – as you can see to the right, we easily compared Amazon's Glacier Storage (top) to IBM Cloud Object Storage (bottom).

	Standard Storage	Standard - Infrequent Access Storage †	Glacier Storage
First 50 TB / month	\$0.023 per GB	\$0.0125 per GB	\$0.004 per GB
Next 450 TB / month	\$0.022 per GB	\$0.0125 per GB	\$0.004 per GB
Over 500 TB / month	\$0.021 per GB	\$0.0125 per GB	\$0.004 per GB

Monthly storage volume (TB)	Standard	Vault	Cold Vault
0-499.99	\$0.0300 GB/Month	\$0.0200 GB/Month	\$0.0110 GB/Month
500+	\$0.0260 GB/Month	\$0.0180 GB/Month	\$0.0110 GB/Month

It's very simple to put in all of the characteristics of your backup strategy and get a clear picture on cost per month.

Here is what to expect from the ACTLOG showing the 'Cloud Sweeping':

Roll off to cloud bucket gvicaix61a.sp81.poc.stgaws.01:

Once roll off completes local cache is emptied.

- ANR0984I Process 25 for Local to Cloud Transfer started in the BACKGROUND at 08:02:23.
- ANR3832I Local to Cloud Transfer process 25 for Storage Pool STGAWS.01 started.
- ANR3833I Local to Cloud Transfer process 25 for Storage Pool STGAWS.01 has completed.
- ANR0986I Process 25 for Local to Cloud Transfer running in the BACKGROUND processed 487 items for a total of 99,252,589 bytes with a completion state of SUCCESS at 08:04:33.



We found several advantages of this type of setup:

- Customers who currently back up directly to tape will find this process very desirable.
- Customers who are considering moving a function of a legacy environment to the cloud can benefit from this style of backup.
- Customers who may not want to invest the capital expenses in a tape library + tapes + offsite storage can save costs through this kind of cloud backup.

At the same time, it's important to understand that there is no 'tier' to the cloud. Currently, customers back up to a cache, and it drains. Tiering would consist of backing up to IBM Spectrum Protect, and putting a policy in place that directs data to be moved from primary storage to the cloud after a 30-day period.

It's also helpful to keep in mind:

- You need to be especially conscious of your organization's bandwidth. In our office test lab, we quickly saturated the network, causing some congestion for our VoIP system. For reference, IBM claims that the ideal cloud object storage environment would support:
 - Backup performance up to 16TB/hour.
 - Recovery objectives up to 3.0TB/hour.
- We are anxiously awaiting an enhancement to IBM Spectrum Protect to be released in late 2017. This enhancement will provide tiering capabilities where policies are in place to automatically tier to Cloud Object Storage based on management class settings.
- You can monitor your IBM Spectrum Protect environment with Galileo Performance Explorer®. Galileo monitors performance and configuration, as well as provides actionable insight about the environment to ensure it's running optimally to meet your backup and restore service level objectives.



Result

In conclusion, we were able to successfully build an IBM Spectrum Protect server and backup a handful of different operating systems to the cloud. We discovered several benefits, including how much control you have over the entire process, especially as each individual server is not connecting directly to the cloud provider. This allows you to commit your backups to a cache, and have it slowly trickle up to the cloud. With the incoming feature of tiering in IBM Spectrum Protect, and with the existing functionality to back directly up to the cloud, this is a rockstar of an edition to one of our favorite products.



THEATSGROUP.COM

