Morris & Opazo

Centralize and automate your data on AWS Backup

We Love the Cloud

aws PARTNER Advanced Tier Services

- Public Sector
- Immersion Day
- Solution Provider
- AWS Lambda Delivery
- Amazon Kinesis Delivery
- Amazon Redshift Delivery
- Amazon API Gateway Delivery
- Data & Analytics Services Competency
- Amazon EC2 for Windows Server Delivery

morrisopazo.com

contacto@morrisopazo.com

AMAZON BACKUP

A service that significantly simplifies the centralized management and automation of your data protection. Learn how to back up your services automatically and how to restore them in just a couple of clicks.

Whether you are using cloud services or on-premises components, AWS Backup is capable of generating backups of these as a safe and stable point with which to recover them easily and quickly in case the service in question suffers any inconvenience.

AWS BACKUP

Morris & Opazo

One of the biggest needs customers had when working in depth with AWS was to have one place where they can manage, configure, and regulate all their backup activity across all accounts and resources in their organization.

In the past, this was only possible if the client made backups on their own and restored their services manually, so when they had to do the same with several databases, instances, storage devices, etc., it became an extremely tedious task and in general impractical in terms of time and money used.

AWS Backup was born to provide a solution to the needs of customers in a secure, centralized and cost-effective way.

AWS BACKUP

Morris & Opazo

Through the following example you will learn how to make your backups automatically periodically according to your needs, as well as how to restore them in case of an emergency.

In essence, the process consists of creating a Security Backup Plan in which we will define the rules that it must follow and the permissions that it will have to carry out these actions and in case of wanting to make a restoration, select the point to which we want to return and define if We make changes at the component level and availability of our service.

ELEMENTS TO BACKUP

Morris & Opazo

Currently AWS Backup natively supports the following services:

Amazon EC2, Amazon RDS, DocumentDB, VSS, Amazon Aurora, Amazon Neptune, Amazon S3, Amazon EFS, VMware, Amazon EBS, FSx (Luster & WFS), Dynamo DB, Volume Gateway.

Being the newest in this list VMware and Amazon S3, we will use the latter to restore a deleted object

BACKUP

Morris & Opazo

First, we locate the Bucket or service to use and the elements it contains.

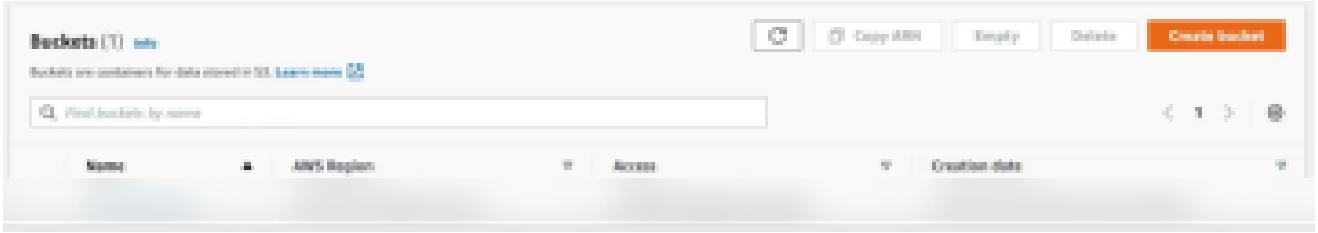

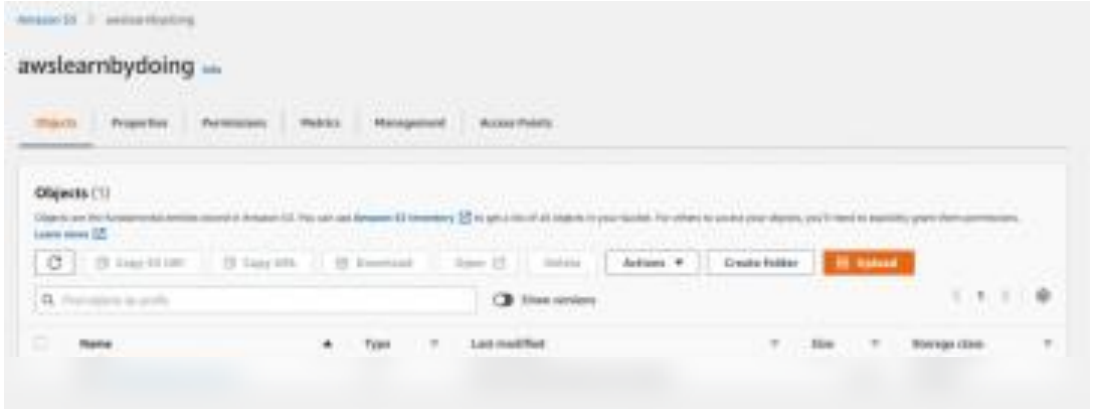

BACKUP

Morris & Opazo

aws

We make sure that versioning is activated

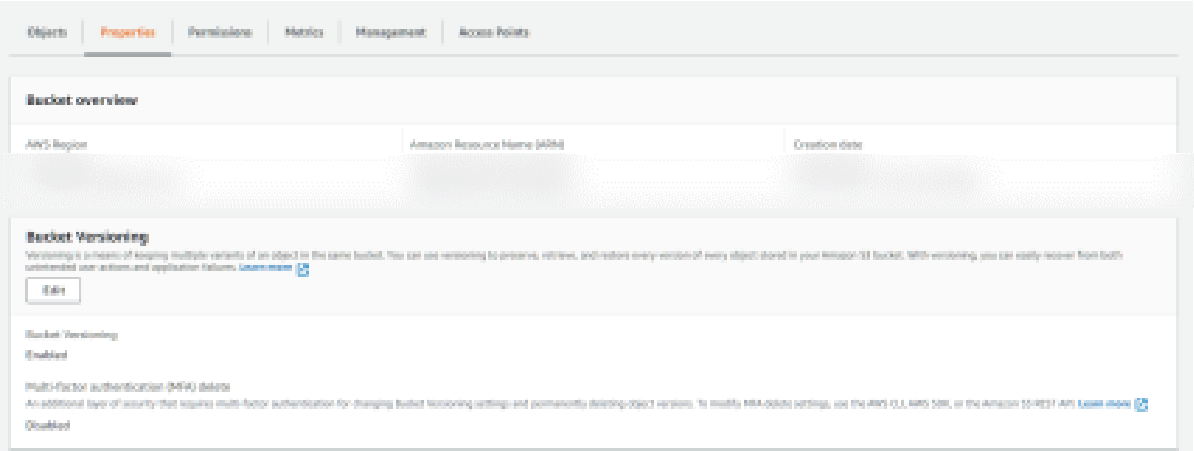

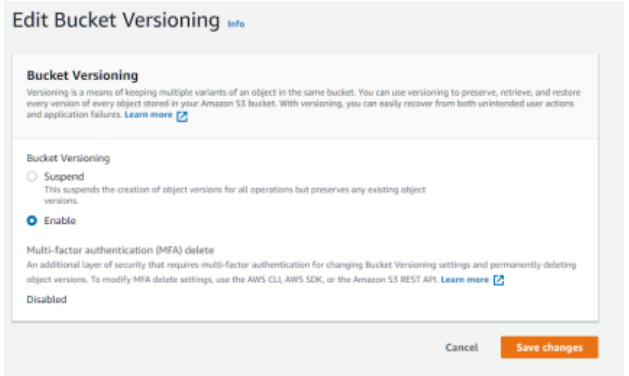

PERMITS

Morris & Opazo

aws

In order for our Backup Plan to carry out all the necessary actions, we must grant the pertinent permissions to the services on which it is going to act, for this we will create an IAM role that we will later assign to our Backup Plan. If we are going to use a specific service, AWS also provides us with a role called “Default Role” that includes (in the case of S3) the following permissions:

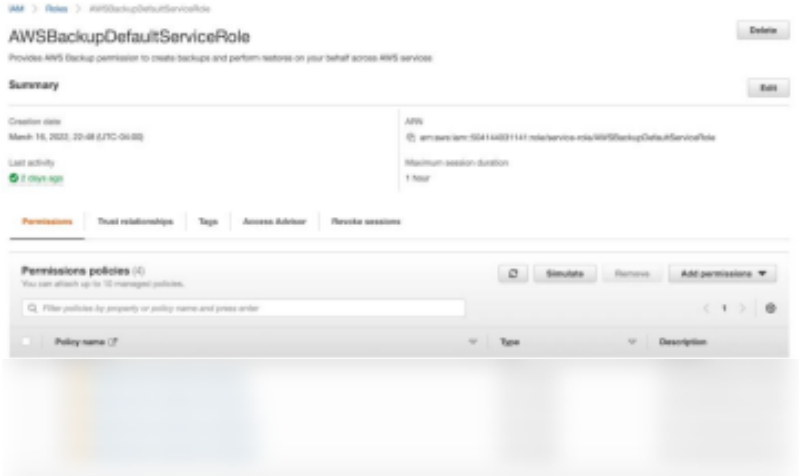

BACKUP PLAN

Morris & Opazo

To create our Backup Plan we go to our AWS Backup dashboard and select “Create Backup Plan”

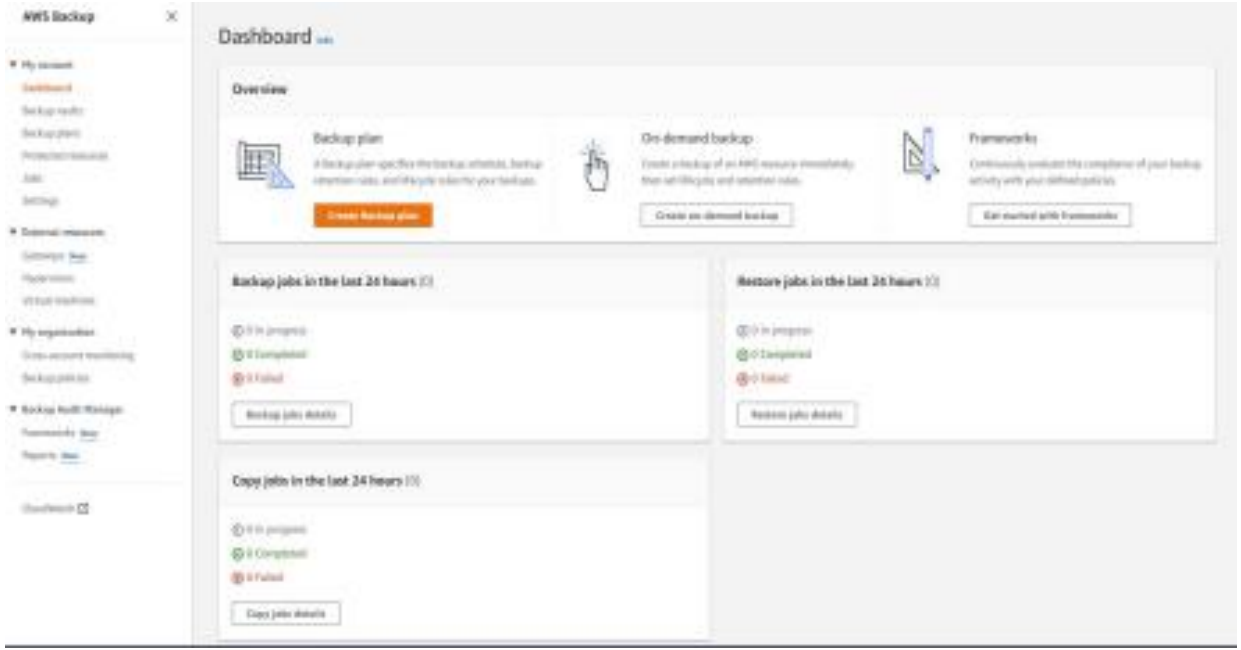

BACKUP PLAN

Morris & Opazo

Once here we have the possibility to choose a pre-made template or create one according to our needs. In this case we will create a new plan

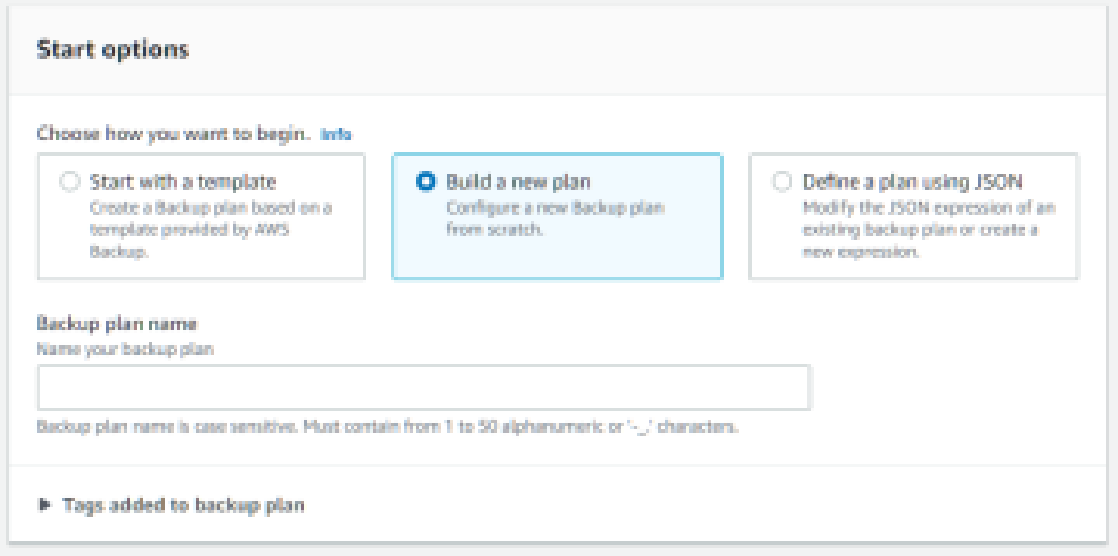

RULES

Morris & Opazo

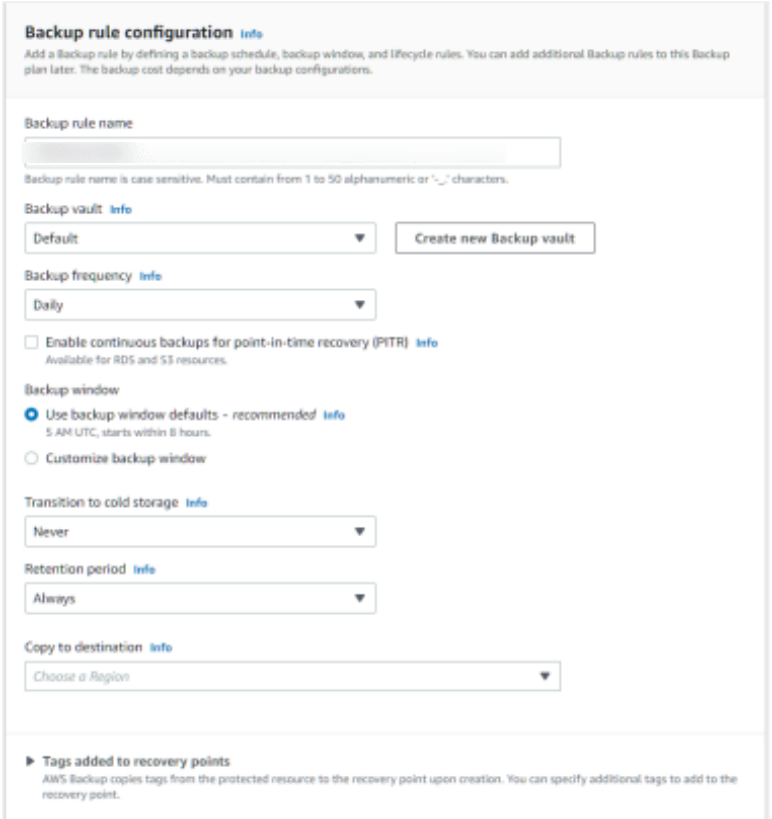

The rules define how often backups are generated, what time they are done, how long the data is retained, and where it is stored. With cold storage we can also choose whether our data is going to be stored long-term and choose when our data is moved to this type of storage.

VAULT

Morris & Opazo

Vault refers to the space that AWS Backup gives us to store our backups, we can use the "Default Vault" or create our own and organize them as we please. We have 2 ways to do it, the first would be to create the vault previously from the Backup Vaults tab and the second is to create it when creating our Backup Plan, for the purposes of this blog we will use the latter.

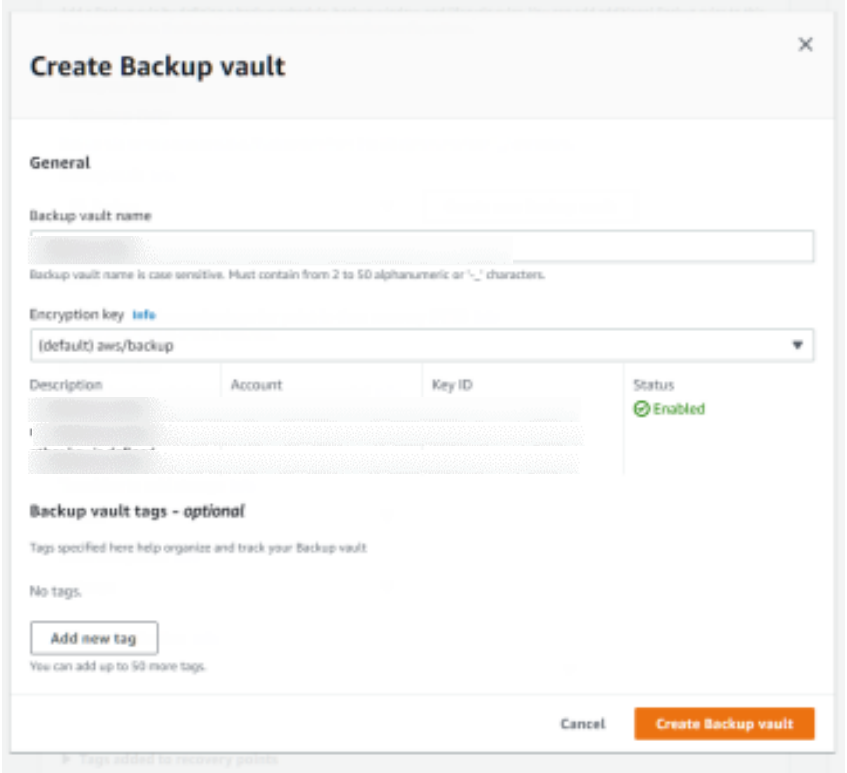

CREATION

Morris & Opazo

aws

Once the rules and the Vault where we will store our backups have been defined, we proceed to create the Backup Plan.

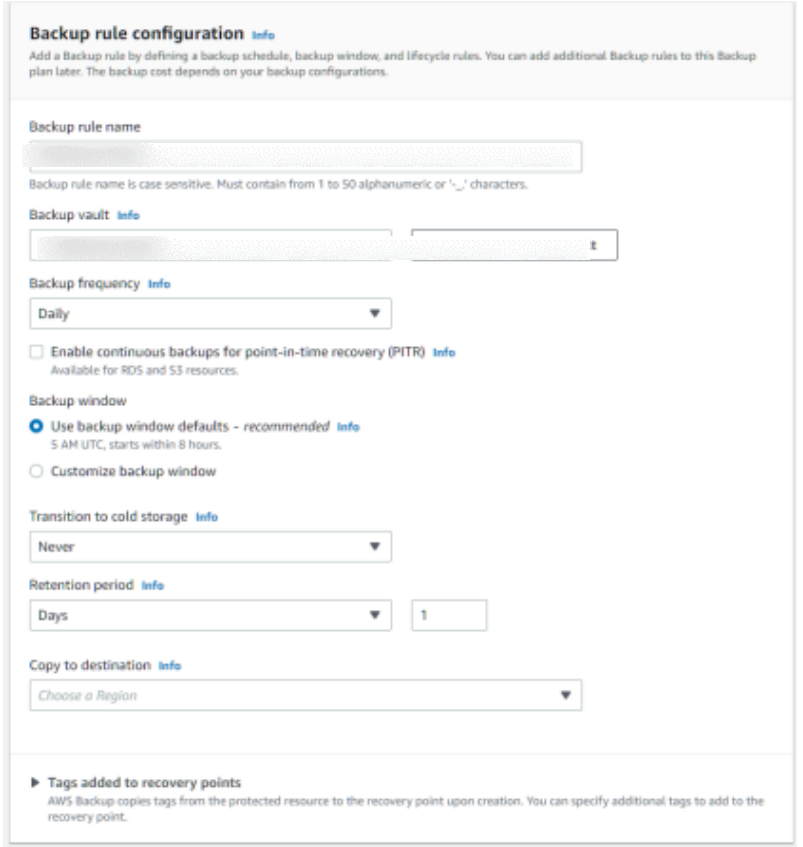

RESOURCE ALLOCATION

Morris & Opazo

After creating our Backup plan we will assign the corresponding resources so that you can work with them.

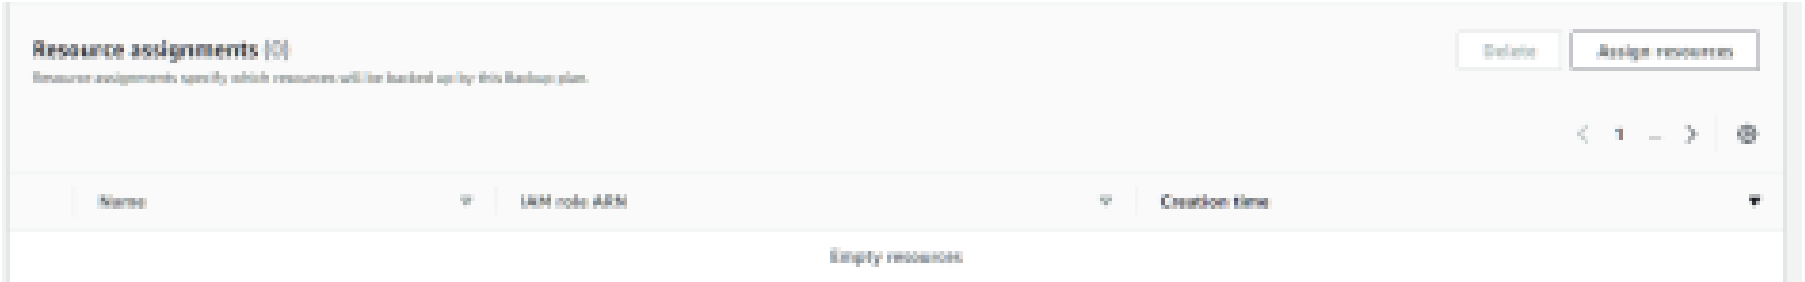

We select our bucket and wait for the first job to be created automatically (we can also configure ourselves the time we want it to be done)

JOBS

Morris & Opazo

We can monitor the effects of this in the Jobs tab, which once completed would look something like this (keep in mind that depending on the service, the creation process could take more or less time):

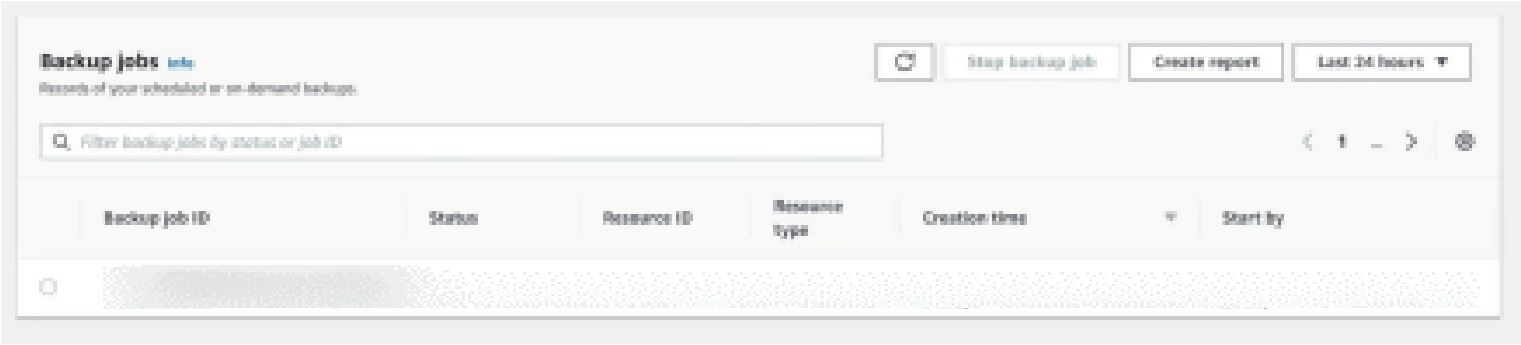

Now we have our backup created.

RESTORE

Morris & Opazo

For the purposes of this example we will simulate the loss of an object in our S3 bucket (To edit the access rules see “*”)

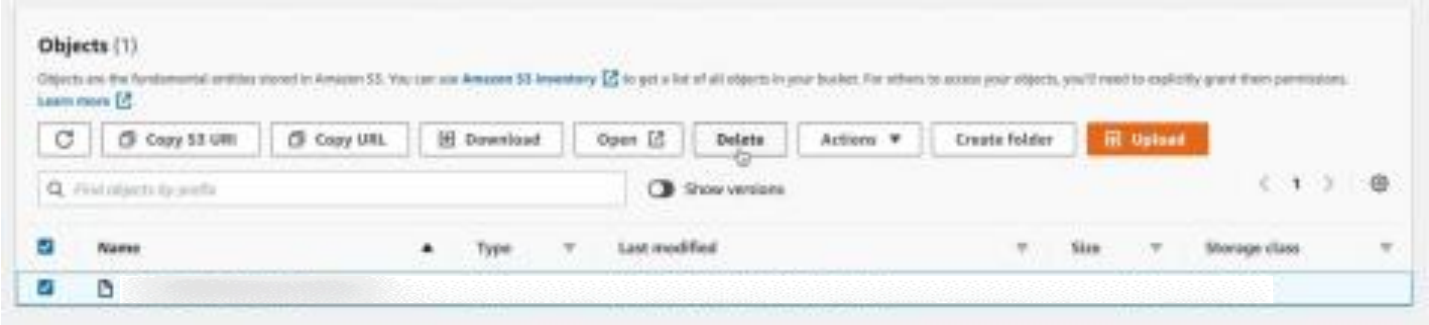

RESTORE

Morris & Opazo

Now that there is nothing in our bucket we will try to restore the object in question using the recovery point from our Backup Vault created by our Backup plan, we select the Vault that we have created (S3_backup)

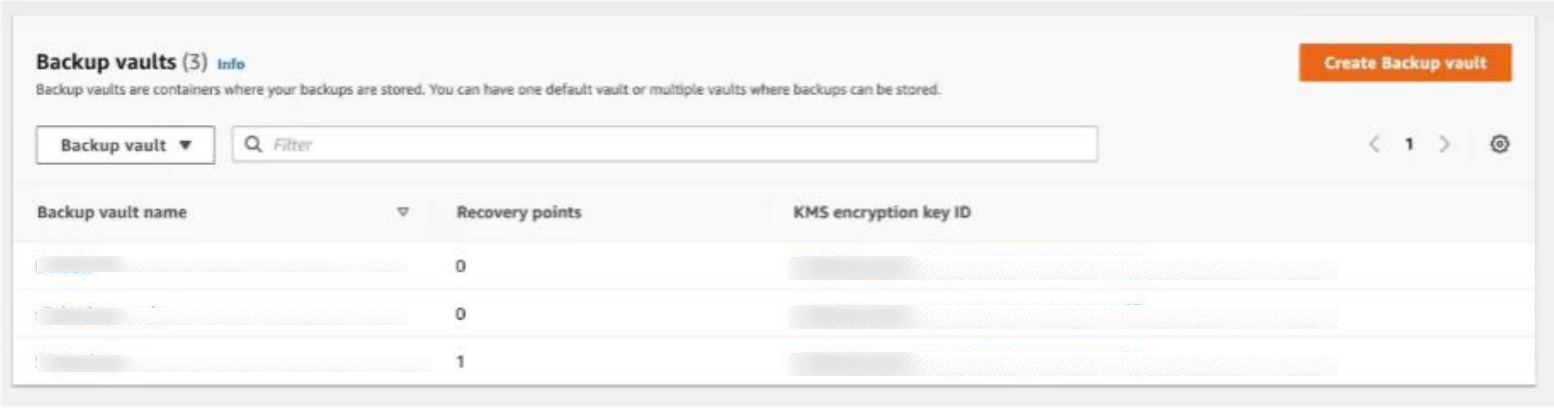

RESTORE

Morris & Opazo

From here we choose the point to which we want to return where our object was still inside the bucket, once selected we display the actions button and press "Restore"

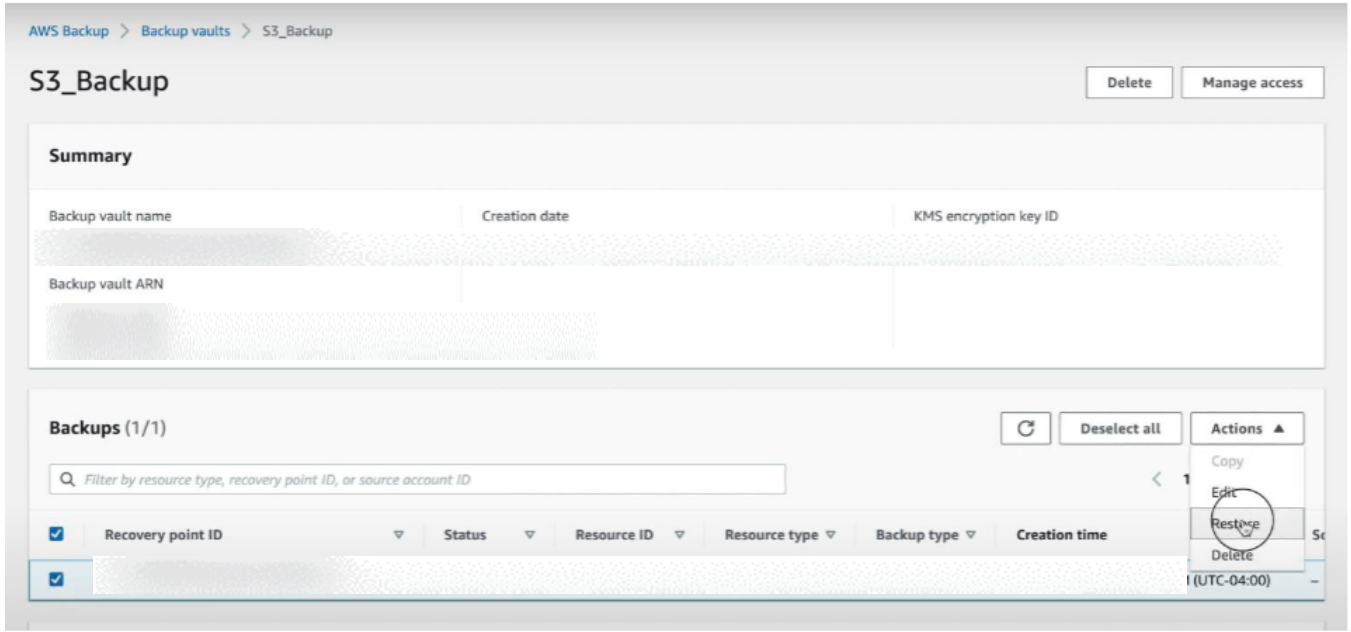

RESTORE

Morris & Opazo

It is very important to remember that for this process to take effect in our S3 bucket we must grant the necessary permissions to both the same bucket and the Backup plan, so we must enable the ACLs and thus allow them to interact with each other.
For this we go to the permissions tab of our bucket and go down to the object ownership section (Object Ownership) and press edit:

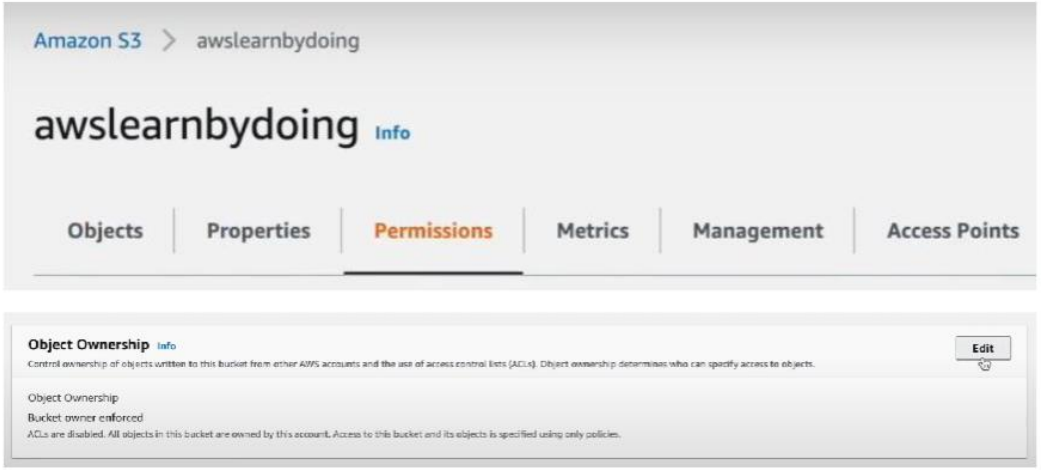

RESTORE

Morris & Opazo

We enable the ACLs and save

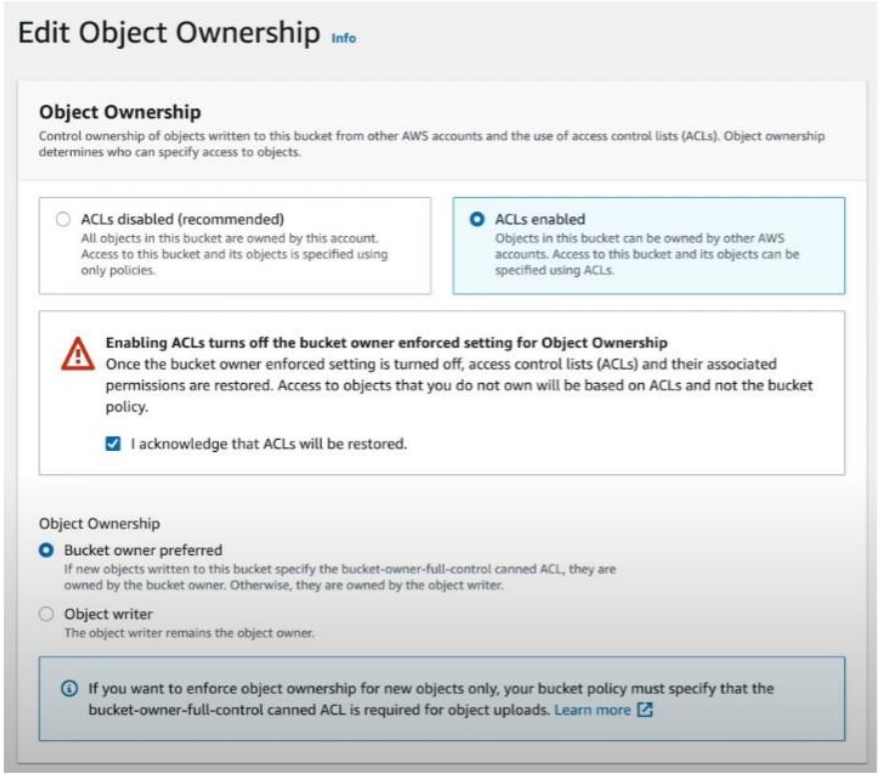

RESTORE

Morris & Opazo

To confirm that our restoration has been successful we go to the Restore Jobs tab under the Jobs section in AWS Backup

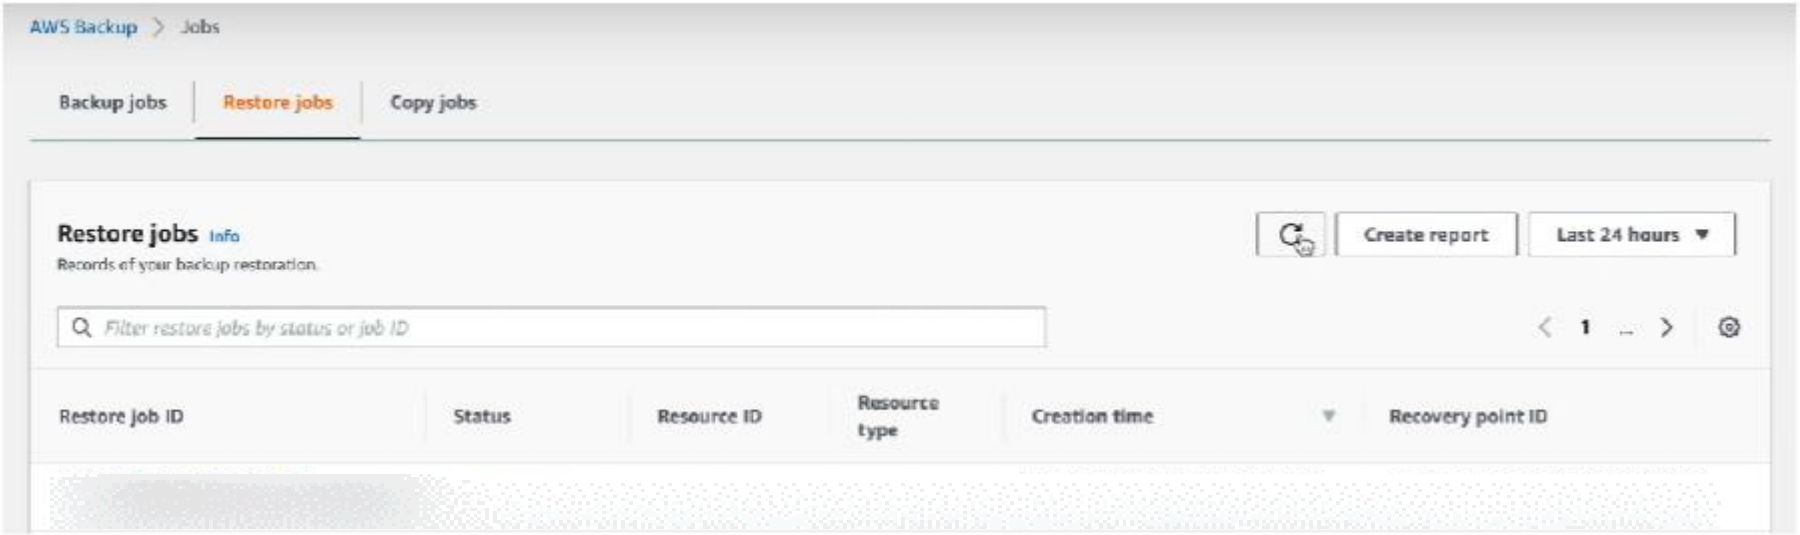

RESTORE

Morris & Opazo

aws

Keep in mind that when doing a restore you can not only go back to a previous state of your resource or service, but you can also change configurations as needed, such as instances for a database or change the availability zone, even change of region on certain occasions (review documentation according to the service to confirm availability). And that's how we fully automatically backed up a service and learned how to restore it to a previous point at our convenience.

Try it for yourself and see all the features that AWS Backup has for you and your business!

Morris & Opazo

Expertos en Tecnología de la Nube

We Love the Cloud

aws PARTNER Advanced Tier Services

- Public Sector
- Immersion Day
- Solution Provider
- AWS Lambda Delivery
- Amazon Kinesis Delivery
- Amazon Redshift Delivery
- Amazon API Gateway Delivery
- Data & Analytics Services Competency
- Amazon EC2 for Windows Server Delivery

Morrisopazo.com

contacto@Morrisopazo.com