COHESITY

Version 1.3

March 2024

CloudArchive Incremental with Periodic Full to Incremental Forever — Migration Guide

ABSTRACT

This guide provides comprehensive details about how to migrate archival operations from CloudArchive Incremental with Periodic Full to Incremental Forever archival format.

Table of Contents

CloudArchive Incremental with Periodic Full to Incremental Forever Migration Process Overview 4

Benefits of Migrating to CloudArchive Incremental Forever 4

CloudArchive Incremental Forever Migration Workflows 4

Migrate All or Selected CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever (Automated Migration) 4

Planned Migration of CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever (Expedited Migration) 4

CloudArchive Incremental with Periodic Full to Incremental Forever Use Cases . 6

Prerequisites 6

Migrate CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever 6

Migrate All CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever 7

Migrate a Selected CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever 8

Planned (expedited) Migration of CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever 9

Edit Cluster Configuration using Iris CLI 9

Update External Target Configuration using API 10

Migration Considerations 12

Staggered Full Archive in CloudArchive Incremental Forever 12

CloudArchive Incremental with Periodic Full Target with LCM Enabled 13

CloudArchive Incremental with Periodic Full Archiving to Multiple Targets 13

CloudArchive Incremental with Periodic Full Archiving to NFS/ S3 Compatible Target Using NFS Protocol 13

CloudArchive Incremental with Periodic Full Archiving to Remote S3-Compatible Target..... 13

CloudArchive Incremental with Periodic Full Archiving to External Target with WORM/Object Lock Enabled 14

CloudArchive Incremental with Periodic Full vs CloudArchive Incremental Forever API Cost 14

Space Reclamation in CloudArchive Incremental Forever 14
Frequently Asked Questions (FAQ) 15
Validate Migration Process 17
Technical Support and Resources 18
Appendix 19
Appendix A—Prepare an Existing External Target for CloudArchive Incremental Forever (AWS) 19
Appendix B—Register a New External Target for CloudArchive Incremental Forever (AWS) 22
Appendix C—Prepare an Existing External Target for CloudArchive Incremental Forever (Azure) 25
Appendix D—Register a New External Target for CloudArchive Incremental Forever (Azure) 33
Your Feedback 40
About the Author 40
Document Version History 40

Figures

Figure 1: CloudArchive Incremental Forever Migration Workflows 5
Figure 2: CloudArchive Incremental with Periodic Full to Incremental Forever Decision Tree 6
Figure 3: CloudArchive Incremental with Periodic Full to Incremental Forever Migration Process 7
Figure 4: Migrate Selected CloudArchive Incremental with Periodic Full Protection Groups to CloudArchive Incremental Forever 8
Figure 5: Staggered Archival Workflow 12

Tables

Table 1: Migration Workflow 5

CloudArchive Incremental with Periodic Full to Incremental Forever Migration Process Overview

CloudArchive Incremental Forever is an improved archival feature first introduced in Cohesity version 6.6 that helps you protect your workloads with greater speed, storage efficiency, and lower cost.

Cohesity recommends migrating all your supported CloudArchive workloads to Incremental Forever archival format.

Benefits of Migrating to CloudArchive Incremental Forever

- **Reduced Cluster Load —** CloudArchive Incremental Forever eliminates the need for periodic full archival of data and metadata. It lowers the load on the cluster from long-running periodic full archival runs.
- **Lower network bandwidth and faster archives** — Compared to CloudArchive Incremental with Periodic Full, Incremental Forever significantly reduces network bandwidth and provides faster archival of data with its forever incremental archival format.
- **Higher storage efficiency with reduced storage cost** — CloudArchive Incremental Forever offers unparallel storage efficiency and reduces storage cost with improved external target-level deduplication (CloudArchive Incremental with Periodic Full provides deduplication at the protection group level), shorter space reclamation cycles, and intelligent life cycle management(optional).

CloudArchive Incremental Forever Migration Workflows

This document provides complete information on the different workflows, such as migrating all or selected CloudArchive Incremental with Periodic Full Protection Groups to CloudArchive Incremental Forever and planning an expedited migration of all CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever archival format.

Migrate All or Selected CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever (Automated Migration)

This migration workflow offers two options: you can migrate either all or a selection of supported CloudArchive Incremental with Periodic Full workloads to Incremental Forever. By selecting this workflow, the workloads will be migrated to Incremental Forever and take effect during the next scheduled periodic full archival.

Planned Migration of CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever (Expedited Migration)

With this migration workflow, you can migrate all supported CloudArchive Incremental with Periodic Full workloads to Incremental Forever immediately without needing to wait for the next periodic full archival cycle. The migration process will take effect during the next scheduled archival run.

Table 1: Migration Workflow

MIGRATION WORKFLOW	EFFORT	MIGRATION TIME
Automated — Migrate all or selected workloads	Update external target configuration to Incremental forever.	Migration happens during the next periodic full archival.
Expedited — Planned migration of workloads	Need to update the external target configurations using CLI / API.	Immediate. Migration happens during the next archival schedule.

Figure 1: CloudArchive Incremental Forever Migration Workflows

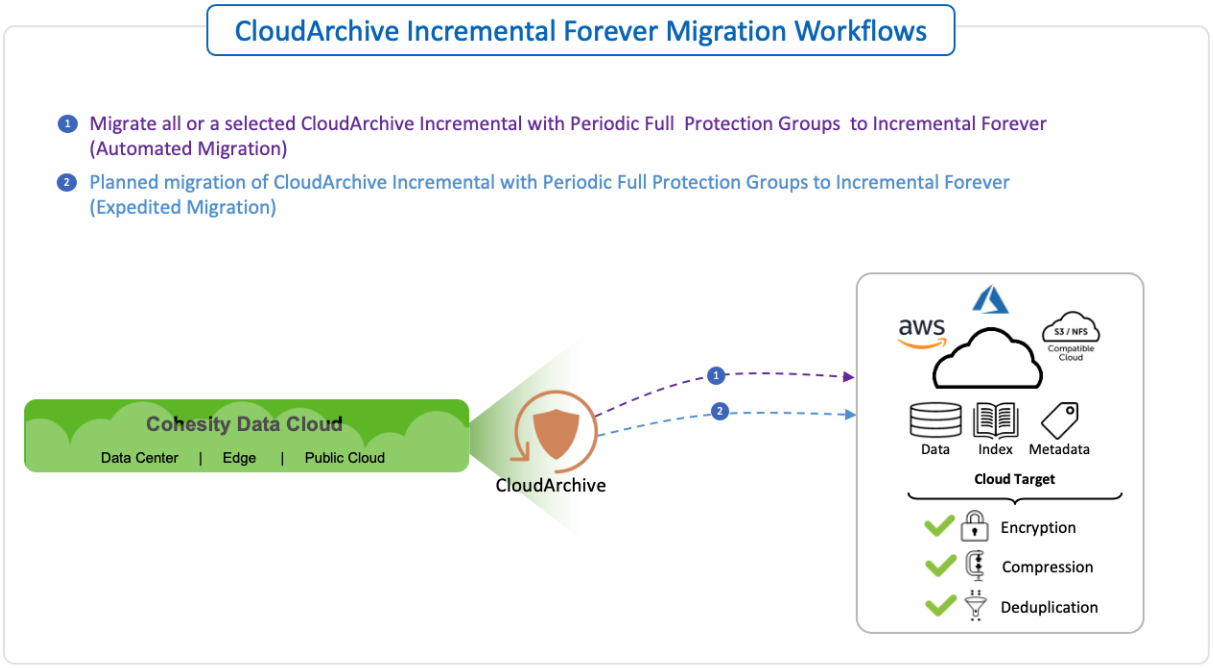

CloudArchive Incremental with Periodic Full to Incremental Forever Use Cases

Below are the everyday use cases where customers want to move their backup to CloudArchive Incremental Forever archival format.

Figure 2: CloudArchive Incremental with Periodic Full to Incremental Forever Decision Tree

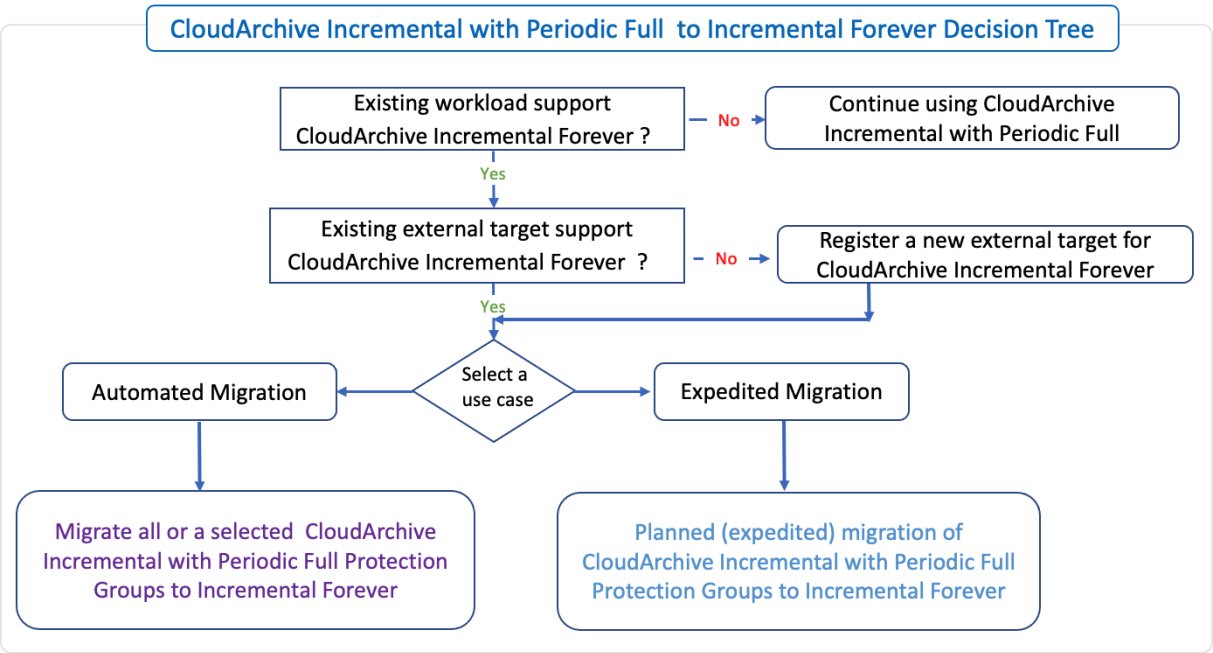

Prerequisites

- Review the following documentation to ensure that the workload and the external targets are supported for CloudArchive Incremental Forever.
  - Workload support matrix
  - External target support matrix

Migrate CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever

This section explains how to migrate either all or a portion of your CloudArchive Incremental with Periodic Full Protection Groups to the Incremental Forever format. If you are currently running archival jobs in the CloudArchive Incremental with Periodic Full format and wish to transition to the Incremental Forever format on the existing external target, these steps will guide you.

Migrate All CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever

1. The recommended method for migration is switching the current external target settings from "Incremental with Periodic Full archival" to "Incremental Forever."

> **IMPORTANT NOTE**:
> Registering a new target in Incremental forever format and updating the existing policy with the new target triggers an immediate full archival of all Protection Groups. This action can result in significant network congestion and potentially slow down other backup activities within the cluster. To prevent this, please adhere to the following precautions:
>
> - Ensure that the existing policy is connected to no more than four Protection Groups simultaneously.
> - Attach the next set of four Protection Groups to the policy after the archival of the initial four Protection Groups is completed.
> - By managing the attachment of Protection Groups in batches of four, you can mitigate the risk of potential network congestion and ensure a smoother initial full archival process in CloudArchive Incremental Forever format.

2. Once you prepare the existing external target for CloudArchive Incremental Forever, the jobs will automatically upload a new full in CloudArchive Incremental Forever format during the next full upload.

3. Prepare the existing external target for CloudArchive Incremental Forever—
   Appendix A—Prepare an existing external target for CloudArchive Incremental Forever (AWS)
   Appendix C—Prepare an existing external target for CloudArchive Incremental Forever (Azure)

> **NOTE**: Archives will continue to be in CloudArchive Incremental with Periodic Full format till the next full archive happens.

Figure 3: CloudArchive Incremental with Periodic Full to Incremental Forever Migration Process

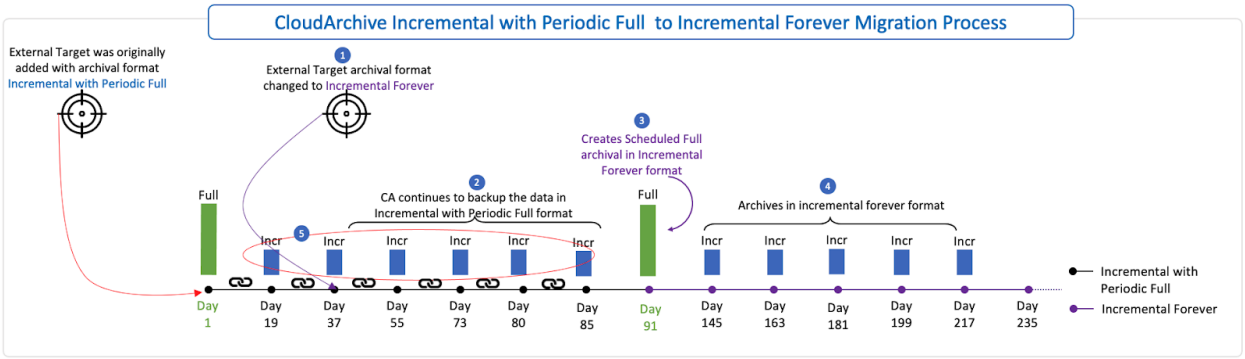

4. After the reference full, incremental backups in CloudArchive Incremental Forever format will now be referencing the new reference full backup.
5. Existing archived data in CloudArchive Incremental with Periodic Full format will be deleted based on the retention period specified in the policy. If required, you can also recover the data from the CloudArchive Incremental with Periodic Full archivals.

Migrate a Selected CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever

If a customer has multiple workloads and only partial workloads support CloudArchive Incremental Forever and the rest support CloudArchive Incremental with Periodic Full, Cohesity still recommends using CloudArchive Incremental Forever to reduce overall costs for storing the CloudArchive data.

1. Prepare the existing external target for CloudArchive Incremental Forever—

   Appendix A—Prepare an existing external target for CloudArchive Incremental Forever (AWS)

   Appendix C—Prepare an existing external target for CloudArchive Incremental Forever (Azure)

2. To migrate partial Protection Groups (not all), after changing the external target format to Incremental Forever archival, only the sources that support CloudArchive Incremental Forever format will be migrated to CloudArchive Incremental Forever format during the next periodic full.
3. The remaining Protection Groups (sources that doesn't support CloudArchive Incremental Forever format) will continue to archive in CloudArchive Incremental with Periodic Full format to the same target.
4. When the remaining job's sources start supporting CloudArchive Incremental Forever format, those will also get migrated to CloudArchive Incremental Forever format during the next reference full backup.

Figure 4: Migrate Selected CloudArchive Incremental with Periodic Full Protection Groups to CloudArchive Incremental Forever

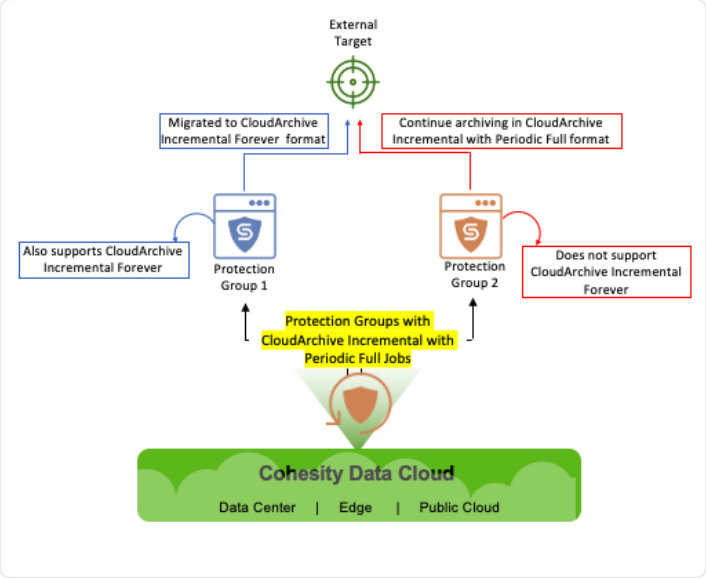

NOTE: If you don't want to wait till the next scheduled full, then you may have to plan for an expedited migration approach.

Planned (expedited) Migration of CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever

This section explains how to perform a planned (expedited) migration of CloudArchive Incremental with Periodic Full Protection Groups to Incremental Forever format. You can perform planned migration by editing the cluster configuration using CLI or by using Cohesity APIs.

Edit Cluster Configuration using Iris CLI

To switch your CloudArchive Incremental with Periodic Full to Incremental Forever immediately, you need to edit the cluster config and set the full archive interval to 0.

1. Run the command **`iris_cli vault ls`** to get the ***<vault_id>*** of the external target.

```
admin@127.0.0.1> vault ls
VAULT ID                        : 12675205
VAULT NAME                      : Vaultaws
VAULT TYPE                      : Amazon
AWS ACCESS KEY ID               : AKIAXH75V3UCROMCOKS5
AWS REGION                      : us-east-1
BUCKET NAME                     : sarancav1tocav2target
ENCRYPTION ENABLED              : true
COMPRESSION                     : LOW COMPRESSION
DEDUPLICATION ENABLED           : true
INCREMENTAL ARCHIVE             : true
FOREVER INCREMENTAL ARCHIVE     : false
FULL ARCHIVAL INTERVAL DAYS     : 90
VAULT USAGE TYPE                : Archival
admin@127.0.0.1>
```

2. Alternatively, you can also find the vault ID from UI. Select **Infrastructure** > **External Targets** and select the target, and the vault ID will be displayed in the URL.

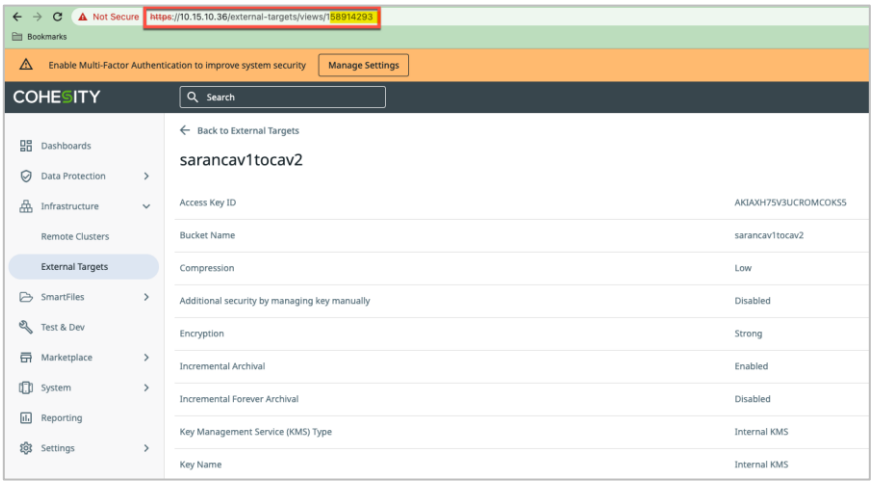

3. Run the below command.

```
iris_cli vault edit id=<vault_id> full-archive-interval-days=<updated_vaule>
Eg: iris_cli vault edit id=4606572076 full-archive-interval-days=0
```

4. Run the command `iris_cli vault ls` to validate the changes.
5. To validate the change from UI, select **Infrastructure** > **External Targets** and select the target and ensure that **Incremental Forever Archival** is now set to **Enabled**.

   Cohesity will create a full archival in the CloudArchive Incremental Forever format during the next scheduled run, and the subsequent archival will be in the Incremental Forever archival format.

Update External Target Configuration using API

You can also use Cohesity API to update the external target and the cluster configuration to switch your CloudArchive Incremental with Periodic Full Protection groups to Incremental Forever.

1. Select the help icon and select REST API to open the Cohesity REST API page.

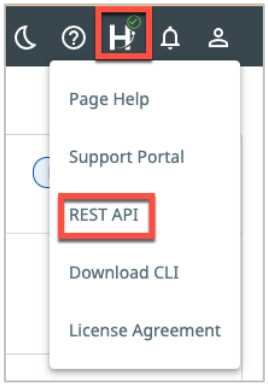

2. Get AccessToken using the API **POST/public/accessTokens**.

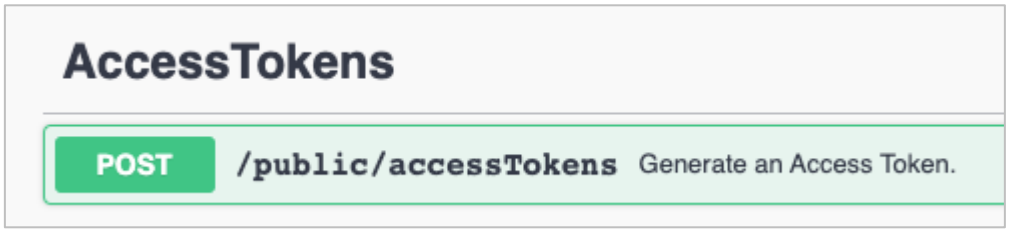

3. Use the **`GET/public/vaults`** API to find the vault ID.

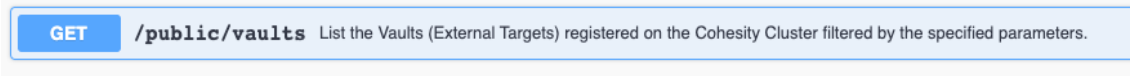

4. Use the API **`PUT/public/vaults/{id}`** to update the vault settings.

```
"isForeverIncrementalArchiveEnabled": true,
"fullArchiveIntervalDays": 0
```

5. Run the **`GET/public/vaults`** API to validate the changes.

Migration Considerations

Review the following migration considerations before you proceed with the migration process.

Staggered Full Archive in CloudArchive Incremental Forever

- During the scheduled or expedited full archive, to avoid long-running tasks and network congestion, Cohesity takes full backups in CloudArchive Incremental Forever format in a staggered way.
  - Example: If there are six jobs (1 job = 1 Protection Group) and the next scheduled periodic full is on day 91, Cohesity takes a full backup of the first four jobs at a time. Once the first job is completed, Cohesity will pick up the remaining jobs and create full archival in CloudArchive Incremental Forever format. Due to this process, if the first job takes two days to complete, the full archival of the next scheduled job will start on day 93 instead of day 91. See the flow chart below to comprehend the staggered archival flow.

Figure 5: Staggered Archival Workflow

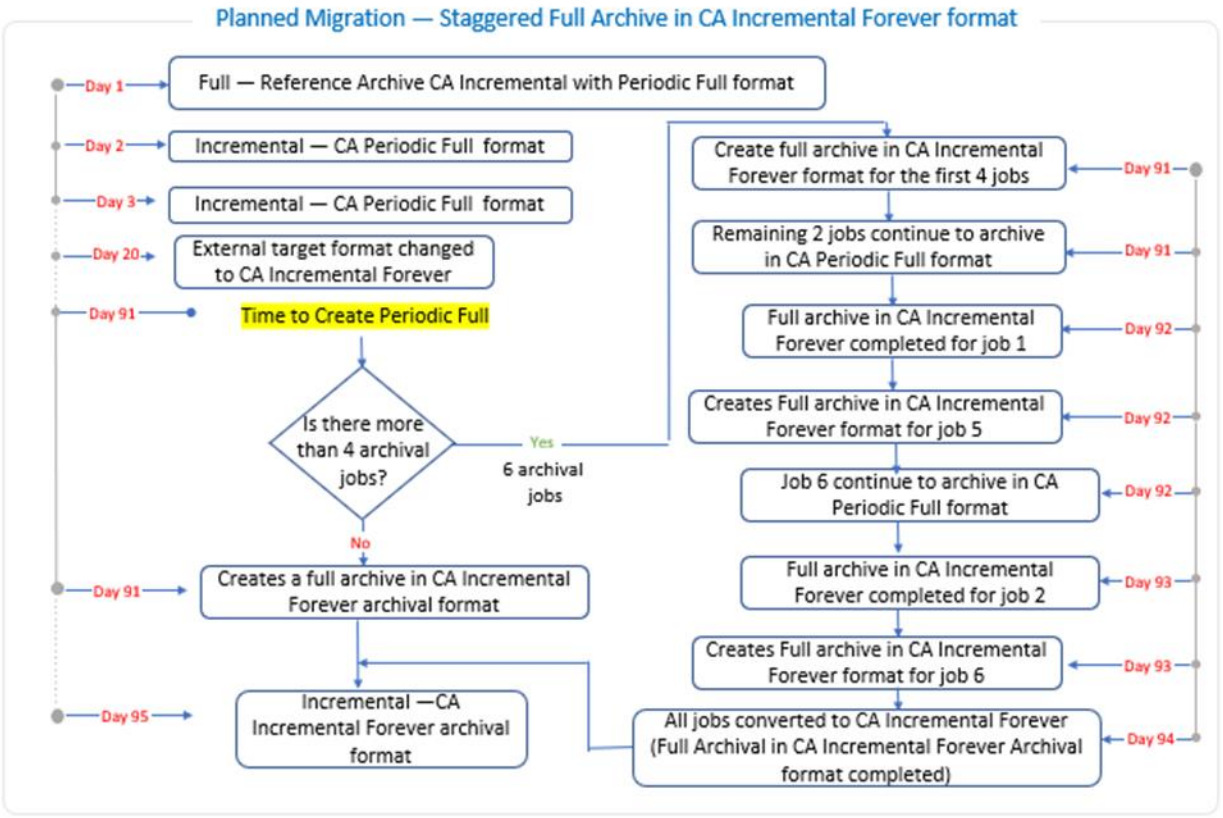

CloudArchive Incremental with Periodic Full Target with LCM Enabled

- If you are archiving CloudArchive Incremental with Periodic Full jobs to an external target with Life Cycle Management (LCM) enabled at the service provider end, Cohesity will not allow changing the archival format to Incremental Forever (CloudArchive Incremental Forever) archival format.
- You must disable LCM to prepare the target for CloudArchive Incremental Forever. Using service providers LCM is not supported with CloudArchive Incremental Forever.
- Cohesity supports policy-based LCM (Data Movement) with CloudArchive Incremental Forever, and it is recommended to set LCM through Cohesity policy and not at the external target.

CloudArchive Incremental with Periodic Full Archiving to Multiple Targets

- Customers who have segregated their external targets due to the use of multiple archive schedules with CloudArchive Incremental with Periodic Full (which is a recommended approach to reduce external target capacity consumption) should create a new archive using a single external target for Incremental Forever. The reason for this is that Incremental Forever now offers enhanced deduplication capabilities from the Storage Domain to the external target level, eliminating the need to segregate external targets for better space efficiency.

CloudArchive Incremental with Periodic Full Archiving to NFS/ S3 Compatible Target Using NFS Protocol

- Cohesity recommends utilizing the S3 protocol for archiving with CloudArchive Incremental Forever if the external target supports both NFS and S3 protocols, as it can deliver better performance.

CloudArchive Incremental with Periodic Full Archiving to Remote S3-Compatible Target

- CloudArchive Incremental Forever does not support remote S3-compatible targets. When registering an S3-compatible target with Incremental Forever archival format, ensure that the target is hosted in the same datacenter.
- For efficient space reclamation, Cohesity periodically downloads data from the S3-compatible targets, and using a remote S3-compatible target will create network congestion with frequent data download.

CloudArchive Incremental with Periodic Full Archiving to External Target with WORM/Object Lock Enabled

- CloudArchive Incremental Forever does not support external targets with WORM/Object Lock enabled. Before migrating to the Incremental Forever format, ensure that the WORM/Object Lock is disabled on the external target. Otherwise, the archival operations will fail with permission denied error.

CloudArchive Incremental with Periodic Full vs CloudArchive Incremental Forever API Cost

- After migrating the jobs to CloudArchive Incremental Forever, you will notice relatively higher API costs for CloudArchive Incremental Forever. This is because CloudArchive Incremental Forever writes data in small chunks for efficient space reclamation.
- Despite its relatively higher API cost, choosing CloudArchive Incremental Forever archival format can result in significant cost savings in the long term due to its efficient space reclamation and deduplication capabilities.
- CloudArchive Incremental Forever provides 2-3x more cost savings than CloudArchive Incremental with Periodic Full.

Space Reclamation in CloudArchive Incremental Forever

- Space reclamation is automatically handled by the source Cohesity cluster. Once an archive has expired and is safe to remove, Cohesity will delete the archive using S3 APIs.
- CloudArchive Incremental Forever also uses serverless Lambda functions in AWS and Function App in Azure to achieve cost-efficient space reclamation.
- CloudArchive Incremental Forever is intelligent enough to calculate the savings before executing the space reclamation process. If there are no savings, then it waits for more chunks to expire so that the Lambda/Azure Function App cost is not overshooting the storage cost.

Frequently Asked Questions (FAQ)

This section covers the answers to the frequently asked questions related to CloudArchive Incremental with Periodic Full to CloudArchive Incremental Forever migration.

1. The customer currently archives to Legacy Glacier. Is there a way to migrate the CloudArchive Incremental with Periodic Full Protection Groups to CloudArchive Incremental Forever?

   Ans: No. Legacy Glacier is not supported. Customers have to register a new target and create new archival in CloudArchive Incremental Forever format and let the existing data in Legacy Glacier expire as per the retention period.

2. The customer currently archives to tape. What options are available for the customer to migrate the CloudArchive Incremental with Periodic Full Protection Groups to CloudArchive Incremental Forever?

   Ans: No, CloudArchive Incremental Forever does not support tapes.

3. Does CloudArchive Incremental with Periodic Full to CloudArchive Incremental Forever migration also migrate the existing data to CloudArchive Incremental Forever Archival format?

   Ans: No, Migration will convert the next scheduled archivals to CloudArchive Incremental Forever format. Existing data (data already archived in CloudArchive Incremental with Periodic Full format) will still be in CloudArchive Incremental with Periodic Full format and will not be migrated to CloudArchive Incremental Forever format.

4. If a file is deleted from the S3 bucket, can Cohesity re-sync?

   Ans: Not currently. Files in the S3 buckets should not be deleted or modified. See the note mentioned in the CloudArchive Guide below:

   > **IMPORTANT**: Customers should never manually edit, change, or delete Cohesity Platform archives directly in cloud object storage.

5. I have upgraded my Cohesity version to 6.8. What happens to the existing CloudArchive Incremental with Periodic Full jobs?

   Ans: Existing CloudArchive Incremental with Periodic Full jobs will continue as it is. Cohesity version upgrade will not convert CloudArchive Incremental with Periodic Full Protection Groups to CloudArchive Incremental Forever.

6. Do I need a new external target, or can I reuse my existing external target?

   Ans: You can use either the existing target or a new target for CloudArchive Incremental Forever. Ensure that the existing CloudArchive Incremental with Periodic Full target supports CloudArchive Incremental Forever from the support matrix.

7. Do I need to make any changes to the existing jobs to switch from CloudArchive Incremental with Periodic Full to CloudArchive Incremental Forever?

   Ans: Once you prepare the external target for CloudArchive Incremental Forever, the jobs will automatically upload a new full in CloudArchive Incremental Forever format during the next scheduled periodic full upload. There is no manual intervention required.

8. I have migrated CloudArchive Incremental with Periodic Full to CloudArchive Incremental Forever. Can I still recover data from the existing CloudArchive Incremental with Periodic Full archives?

   Ans: Yes, existing data is retained in the CloudArchive Incremental with Periodic Full format and can be recovered or deleted based on the retention.

9. I have prepared the target for CloudArchive Incremental Forever. However, I also have a workload getting archived to the same bucket, which is not supported for CloudArchive Incremental Forever. Will this create an issue?

   Ans: No, the supported CloudArchive Incremental Forever workloads will archive to the target in CloudArchive Incremental Forever format and the existing CloudArchive Incremental with Periodic Full workload will continue archiving to the same bucket in the CloudArchive Incremental with Periodic Full format. When the workload supports CloudArchive Incremental Forever (Cohesity version upgrade), it will automatically get converted to CloudArchive Incremental Forever format.

10. Customer wants to migrate to CloudArchive Incremental Forever. However, their security policy does not allow running CloudFormation Template (CFT) in their environment. Can the customer still get the workloads migrated to CloudArchive Incremental Forever format?

    Ans: Cohesity uses CFT to efficiently perform space reclamation using serverless Lambda functions. If CFT is not enabled, Cohesity will not be able to perform efficient space reclamations, which could significantly reduce the cloud cost. If you still want to enable CloudArchive Incremental Forever without CFT, contact support to prepare the target for CloudArchive Incremental Forever format without the CFT requirements.

Validate Migration Process

To ensure the migration from CloudArchive Incremental with Periodic Full to CloudArchive Incremental Forever is successful, follow these guidelines:

1. Log in to the siren page. **http://<cluster-ip>/siren/v1/remote** and select **Icebox**.

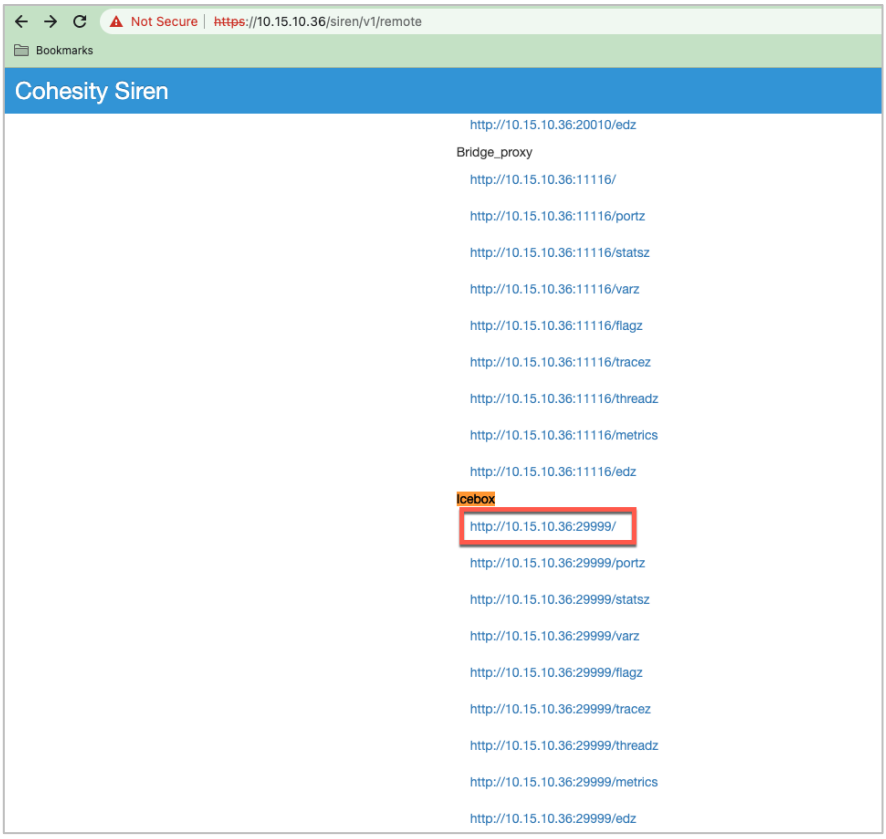

2. Review the Completed archival jobs. The jobs migrated to CloudArchive Incremental Forever will now have the type as **kCloudDomainArchiveJob**.

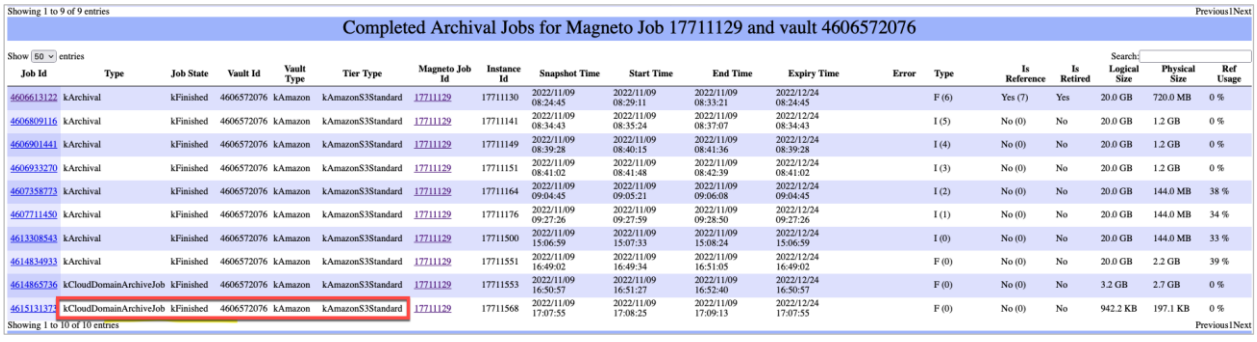

Showing 1 to 9 of 9 entries — Previous 1 Next

Completed Archival Jobs for Magneto Job 17711129 and vault 4606572076

Show 50 entries — Search:

Job Id	Type	Job State	Vault Id	Vault Type	Tier Type	Magneto Job Id	Instance Id	Snapshot Time	Start Time	End Time	Expiry Time	Error	Type	Is Reference	Is Retired	Logical Size	Physical Size	Ref Usage
4606613122	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711130	2022/11/09 08:24:45	2022/11/09 08:29:11	2022/11/09 08:33:21	2022/12/24 08:24:45		F (6)	Yes (7)	Yes	20.0 GB	720.0 MB	0 %
4606809116	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711141	2022/11/09 08:34:43	2022/11/09 08:35:24	2022/11/09 08:37:07	2022/12/24 08:34:43		I (5)	No (0)	No	20.0 GB	1.2 GB	0 %
4606901441	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711149	2022/11/09 08:39:28	2022/11/09 08:40:15	2022/11/09 08:41:36	2022/12/24 08:39:28		I (4)	No (0)	No	20.0 GB	1.2 GB	0 %
4606933270	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711151	2022/11/09 08:41:02	2022/11/09 08:41:48	2022/11/09 08:42:39	2022/12/24 08:41:02		I (3)	No (0)	No	20.0 GB	1.2 GB	0 %
4607358773	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711164	2022/11/09 09:04:45	2022/11/09 09:05:21	2022/11/09 09:06:08	2022/12/24 09:04:45		I (2)	No (0)	No	20.0 GB	144.0 MB	38 %
4607711450	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711176	2022/11/09 09:27:26	2022/11/09 09:27:59	2022/11/09 09:28:50	2022/12/24 09:27:26		I (1)	No (0)	No	20.0 GB	144.0 MB	34 %
4613308543	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711500	2022/11/09 15:06:59	2022/11/09 15:07:33	2022/11/09 15:08:24	2022/12/24 15:06:59		I (0)	No (0)	No	20.0 GB	144.0 MB	33 %
4614834933	kArchival	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711551	2022/11/09 16:49:02	2022/11/09 16:49:34	2022/11/09 16:51:05	2022/12/24 16:49:02		F (0)	No (0)	No	20.0 GB	2.2 GB	39 %
4614865736	kCloudDomainArchiveJob	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711553	2022/11/09 16:50:57	2022/11/09 16:51:27	2022/11/09 16:52:40	2022/12/24 16:50:57		F (0)	No (0)	No	3.2 GB	2.7 GB	0 %
4615131373	kCloudDomainArchiveJob	kFinished	4606572076	kAmazon	kAmazonS3Standard	17711129	17711568	2022/11/09 17:07:55	2022/11/09 17:08:25	2022/11/09 17:09:13	2022/12/24 17:07:55		F (0)	No (0)	No	942.2 KB	197.1 KB	0 %

Showing 1 to 10 of 10 entries — Previous 1 Next

Technical Support and Resources

- Cohesity Support Portal provides you access to a robust, on-demand, and detailed knowledge base, along with high-quality services to boost your experiences with Cohesity products.
- Cohesity Product Documentation provides you access to the latest product documentation to support your deployment of Cohesity products including technical guides and third-party software support matrix for Cohesity Data Protection.
- Cohesity Developer Portal provides you ready-to-use integrations with the automation and orchestration tools of your choice to streamline operations.

Appendix

Appendix A—Prepare an Existing External Target for CloudArchive Incremental Forever (AWS)

1. Select **External Targets** and click the **Edit** icon to change the archival format of the external target.

2. Currently the Archival Format is selected as Incremental with Periodic Full. From the drop-down list, select the Archival Format as Incremental Forever.

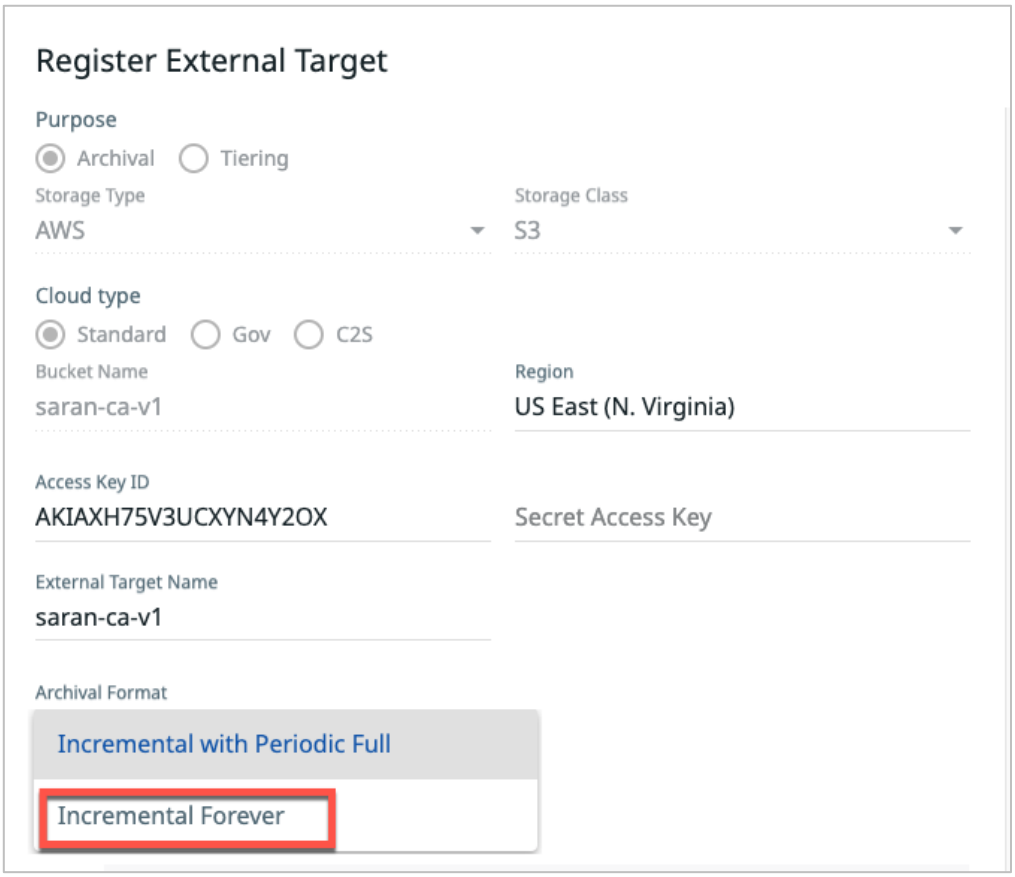

3. Once you change the Archival Format to Incremental Forever, download the Cloud Formation Template from the admonition link shown below.

> **NOTE**: To register an AWS S3 target on an AWS Cloud Edition, there is no requirement to deploy a lambda function. Simply changing the Archival format to “Incremental forever” should be sufficient.

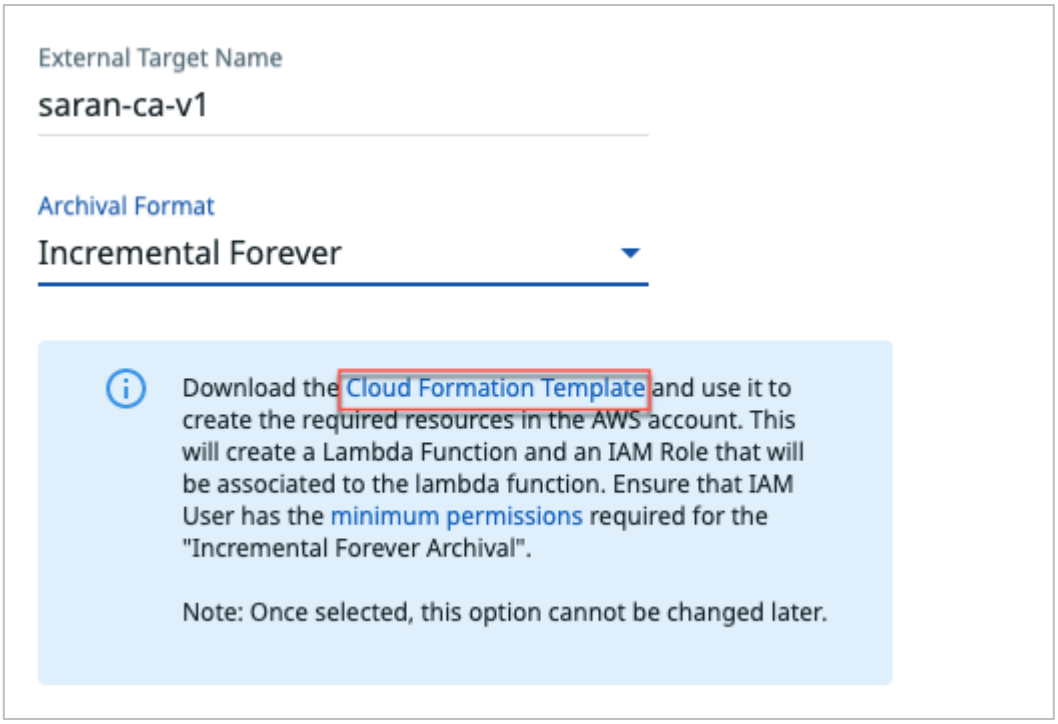

> **NOTE**: Review the minimum permission requirements under Cloud Formation Template section in the documentation. Additionally, make sure to grant the minimum permissions required for CloudArchive and CloudRetrieve.

4. Log into AWS Console.
5. Navigate to **Services** > **CloudFormation > Create Stack**.

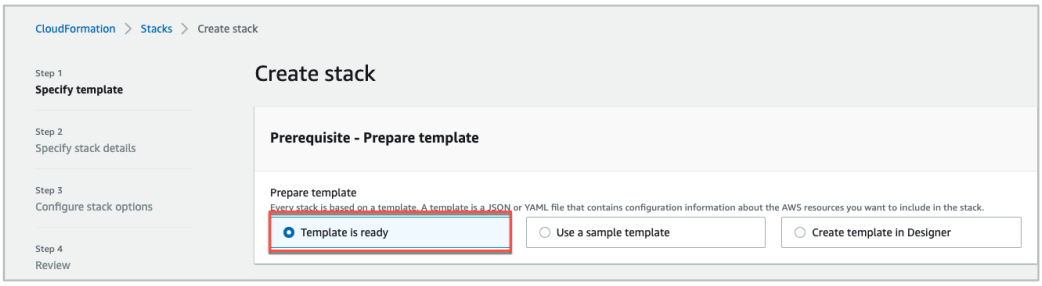

6. Under Specify template, select Upload a template file, select the downloaded CFT file, and then click **Next**.

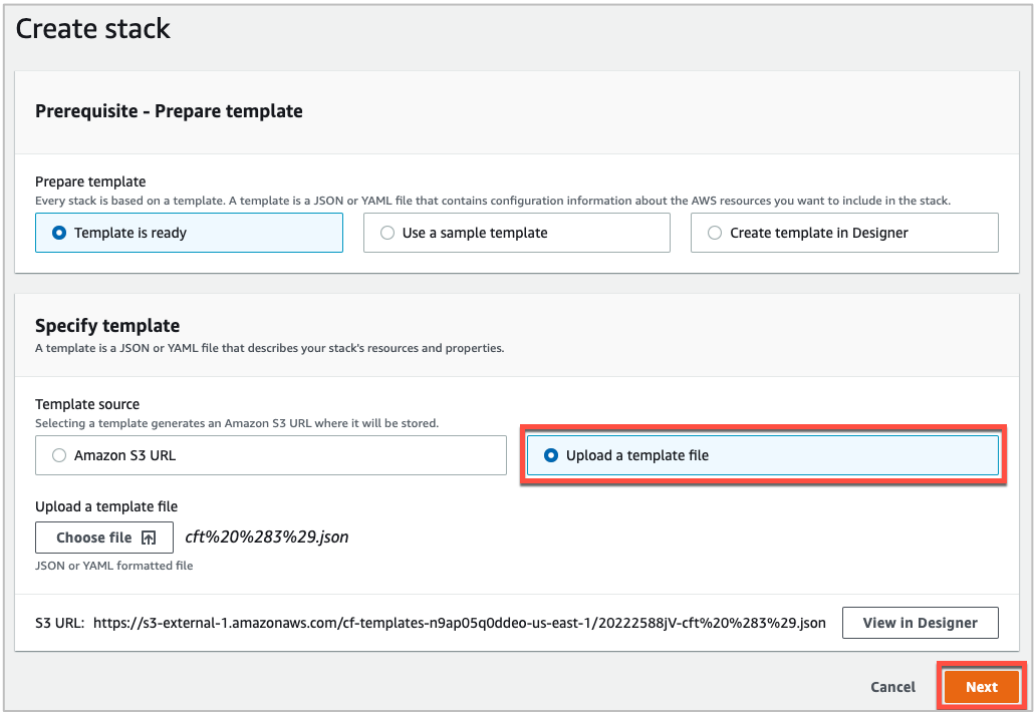

7. In **Configure stack options**, click **Next,** select the checkbox to acknowledge the creation of IAM resources, and then click **Create Stack**.

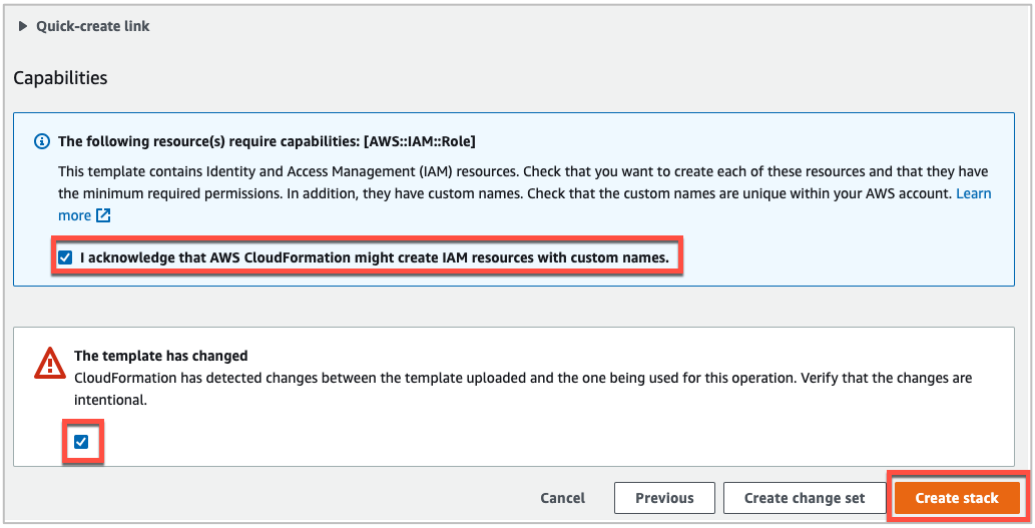

8. Once the stack is completed, go back to the Edit external target UI and enter the **Secret Access Key** and then click **Register** to update the external target archival format to Incremental Forever.

Appendix B—Register a New External Target for CloudArchive Incremental Forever (AWS)

1. Select **External Targets > Add External Target**.
2. Select the Purpose as **Archival** and select and enter the required fields.
   - Storage Type: **AWS**
   - Storage Class: S3 Cloud type: **Standard**
   - Bucket Name: <Enter the name of the bucket>
   - Region: <Select the region where the bucket is created>
   - Access Key ID: <Enter the Access Key>
   - Secret Access Key: <Enter the secret Key>
   - External Target Name:<Provide a name for the target>
   - Archival Format: **Incremental Forever**

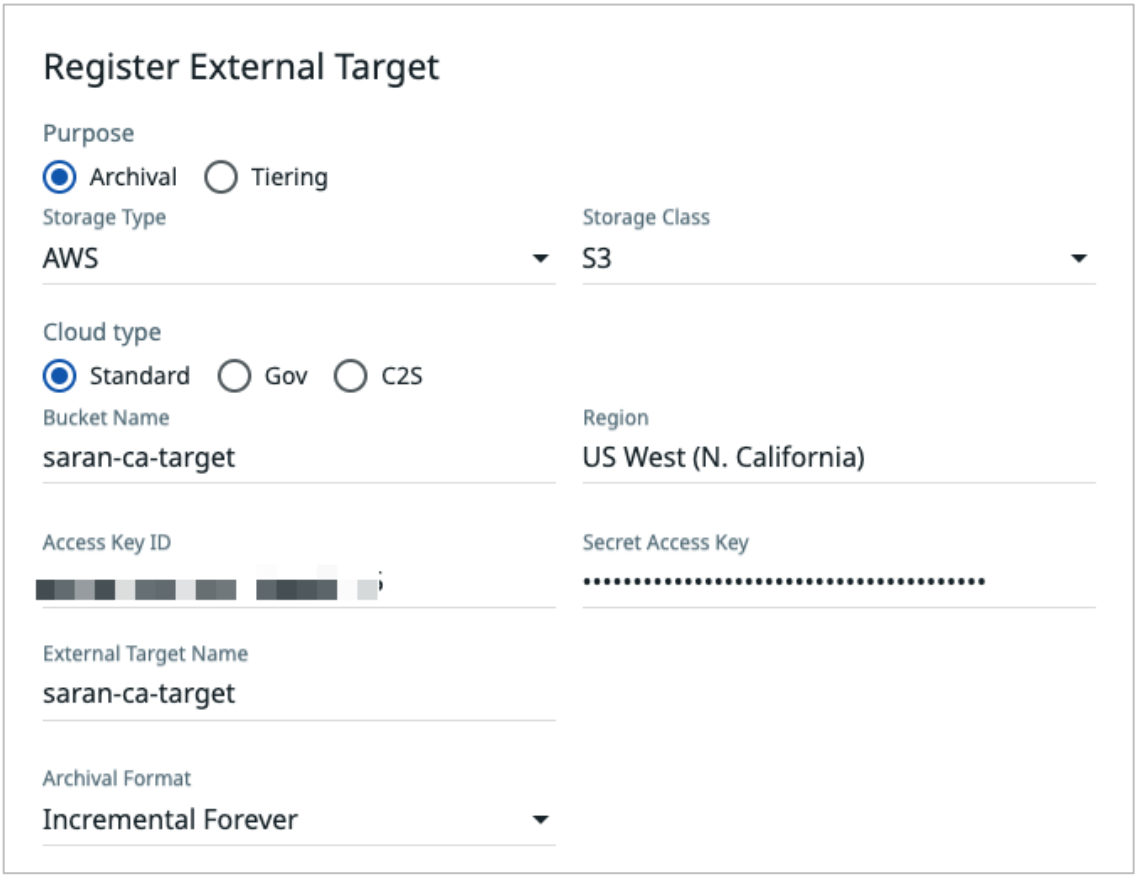

3. Once you change the Archival Format to Incremental Forever, download the Cloud Formation Template from the admonition link shown below.

> **NOTE**: To register an AWS S3 target on an AWS Cloud Edition, there is no requirement to deploy a lambda function. Simply changing the Archival format to “Incremental forever” should be sufficient.

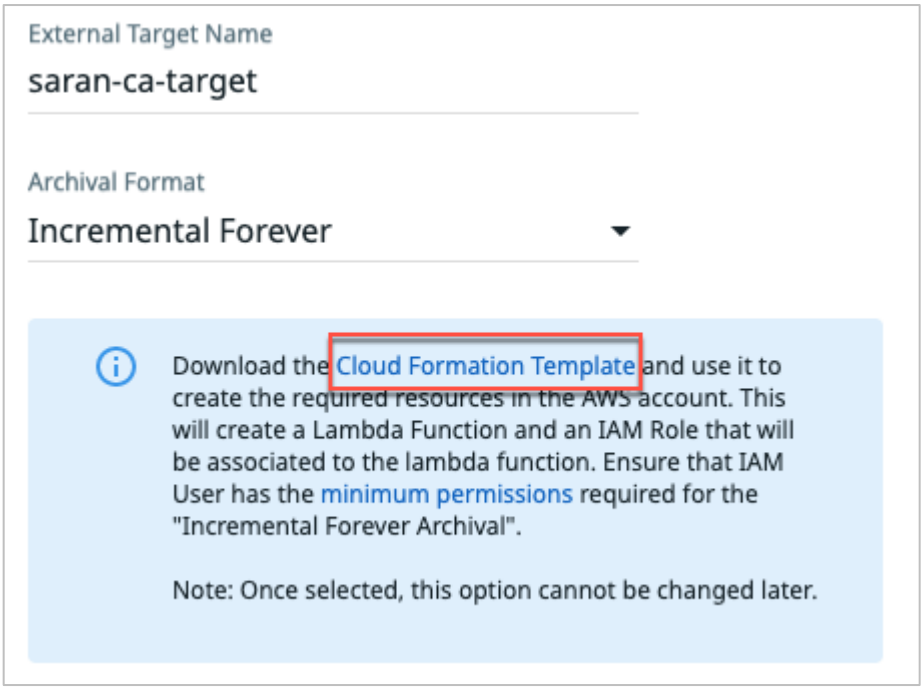

> **NOTE**: Review the minimum permission requirements under Cloud Formation Template section in the documentation. Additionally, make sure to grant the minimum permissions required for CloudArchive and CloudRetrieve.

4. Log into AWS Console.
5. Navigate to **Services** > **CloudFormation > Create Stack**.

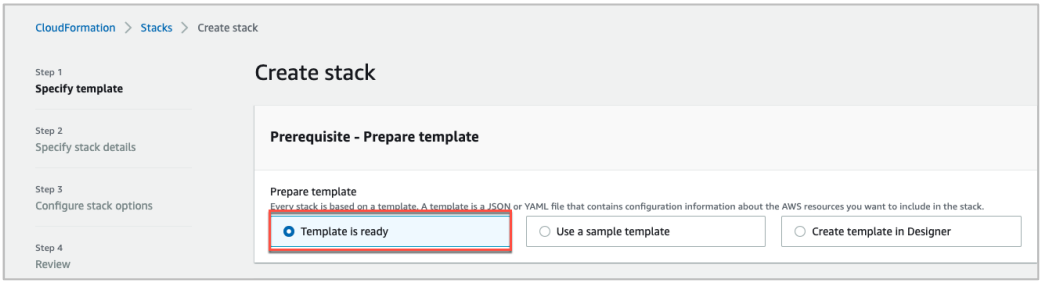

6. Under **Specify template**, select **Upload a template file,** select the downloaded CFT file, and then click **Next**.

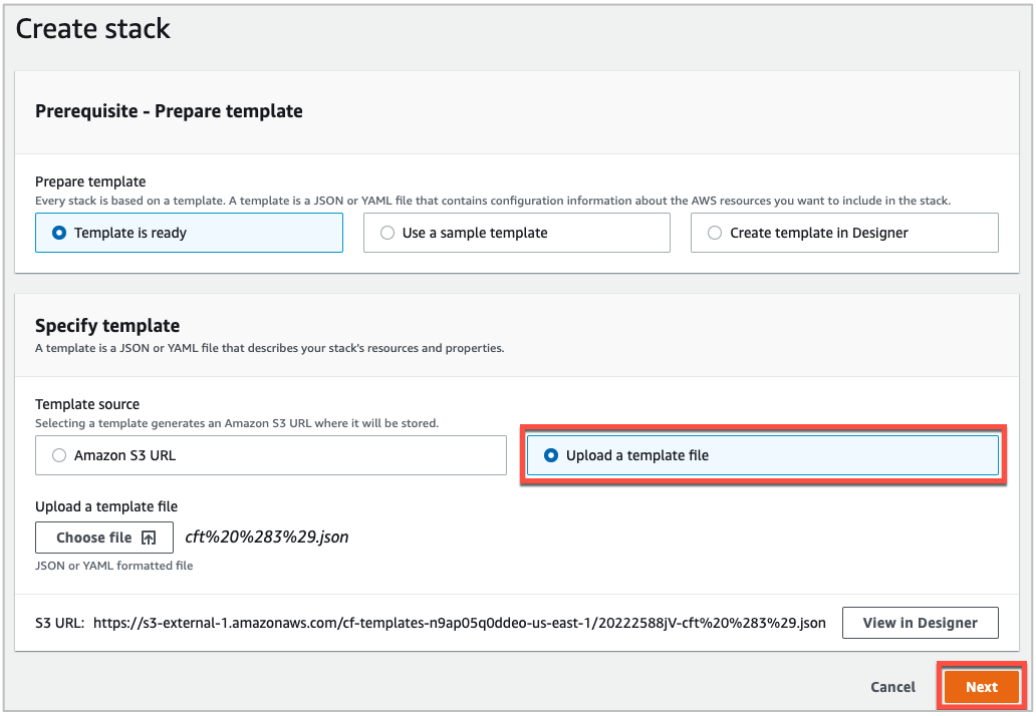

7. In the **BucketName** field, enter a **Stack name** and enter the name of the existing target.

8. In **Configure stack options**, click **Next,** select the checkbox to acknowledge the creation of IAM resources, and then click **Create Stack**.

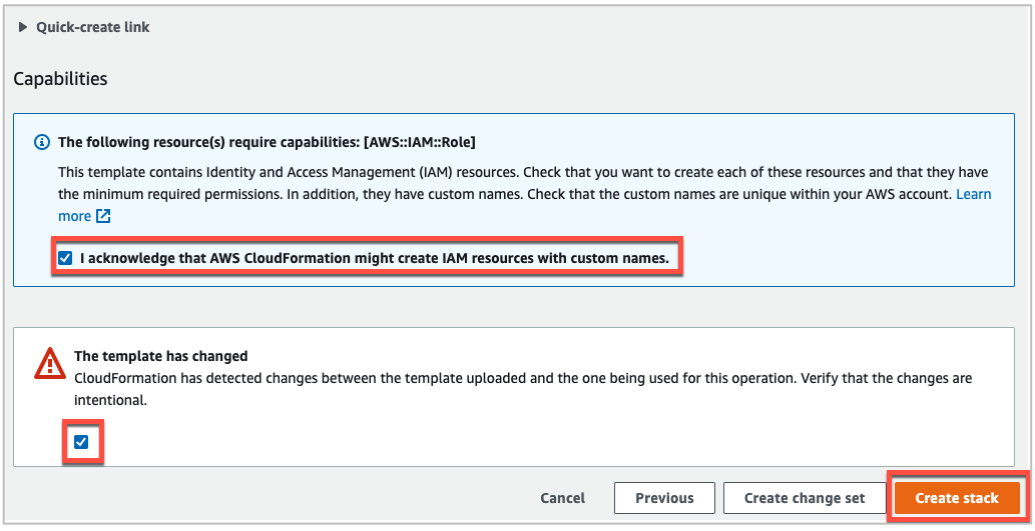

9. Once the stack is completed, go back to the Add external target UI and then click **Register** to register the external target in the Incremental Forever archival format.

Appendix C—Prepare an Existing External Target for CloudArchive Incremental Forever (Azure)

1. Select **External Targets** and the **Edit** icon to change the archival format of the external target.

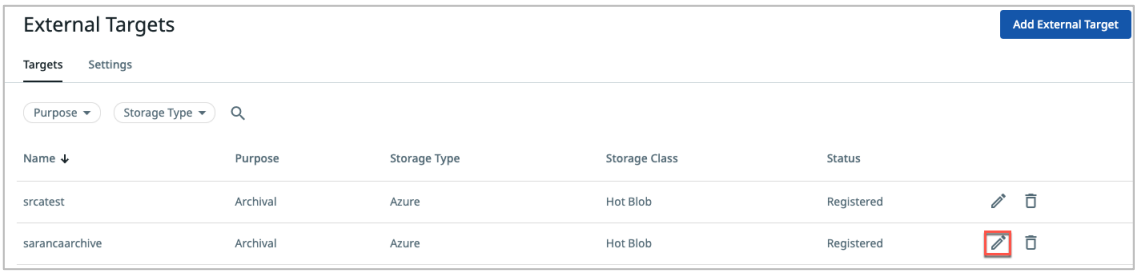

2. Currently, the Archival Format is selected as Incremental with Periodic Full. From the drop-down list, select the Archival Format as **Incremental Forever**.

> **NOTE**: To register an Azure Blob target on an Azure Cloud Edition, there is no requirement to deploy the Azure Function App. Simply changing the Archival format to “Incremental forever” should be sufficient.

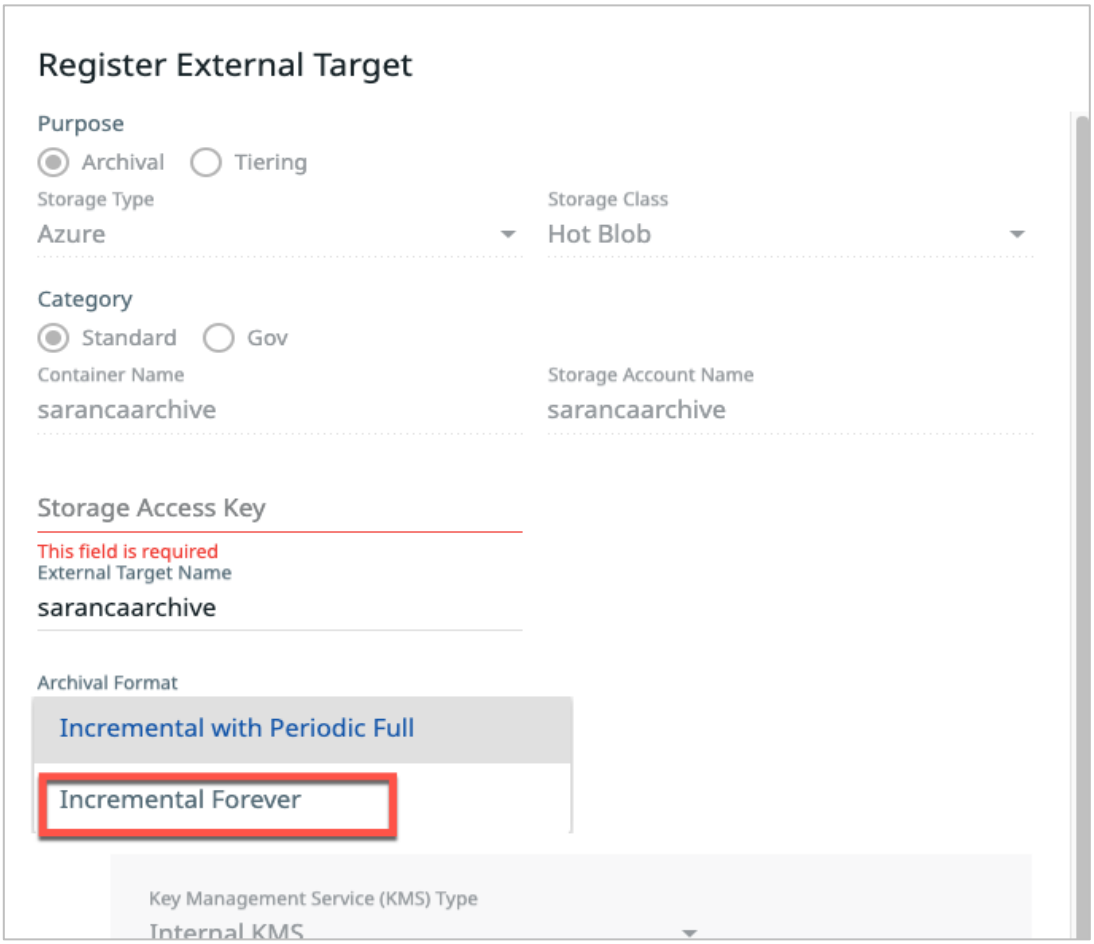

3. Once you change the Archival Format to Incremental Forever, download the **Azure Resource Template** (ART) from the admonition link shown below.

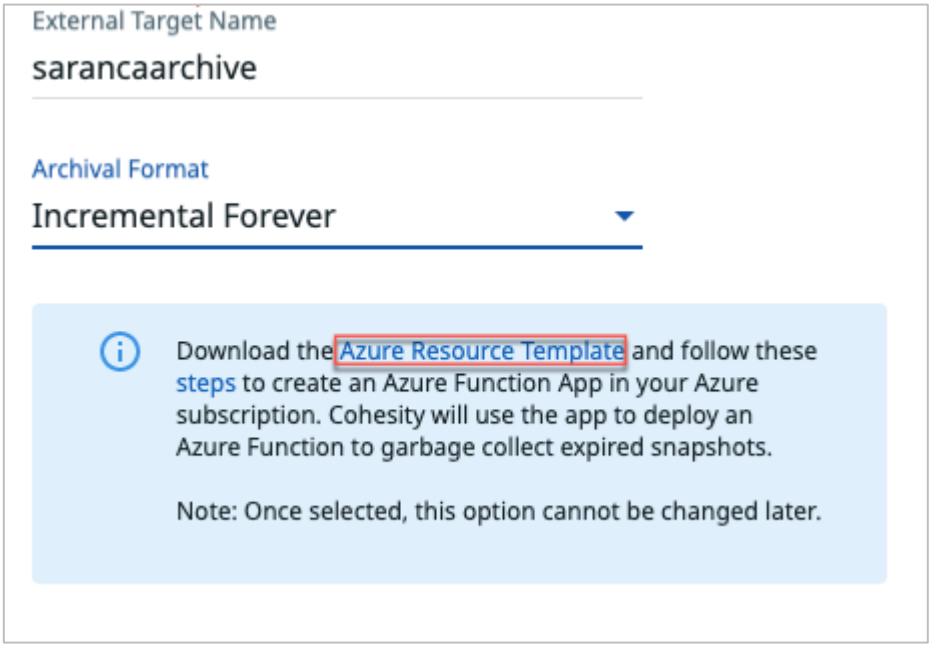

4. Log into Azure Portal.
5. Select **All services** > **Templates** > **+ Create**.

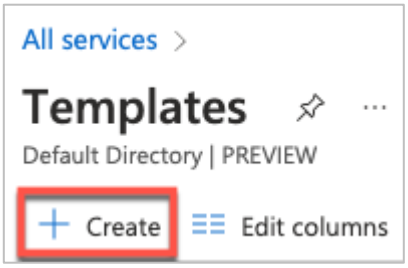

6. Enter a **Name** and **Description** for the template and click **Next:ARM Template**.

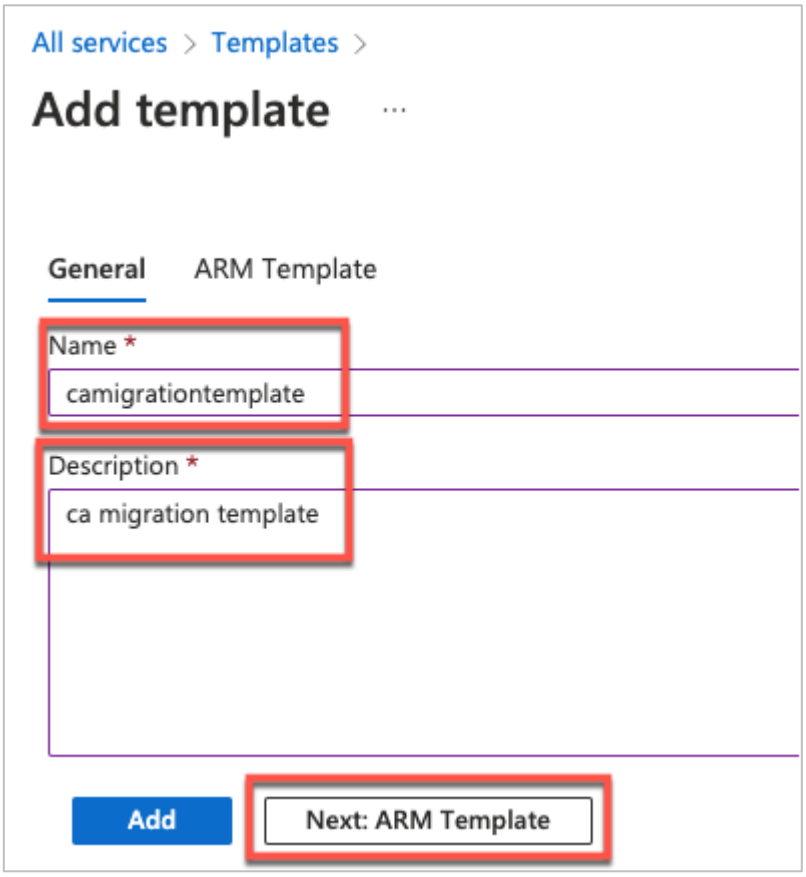

7. Unzip the downloaded ART folder and copy the content from the file **template.json**.

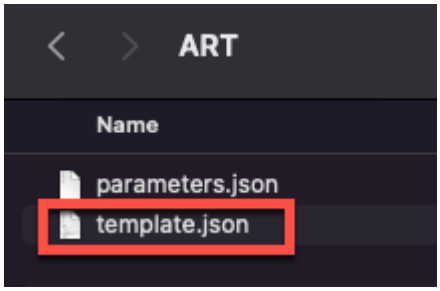

8. In the ARM Template section, replace the content from the template.json file and click **Add** to save the template.

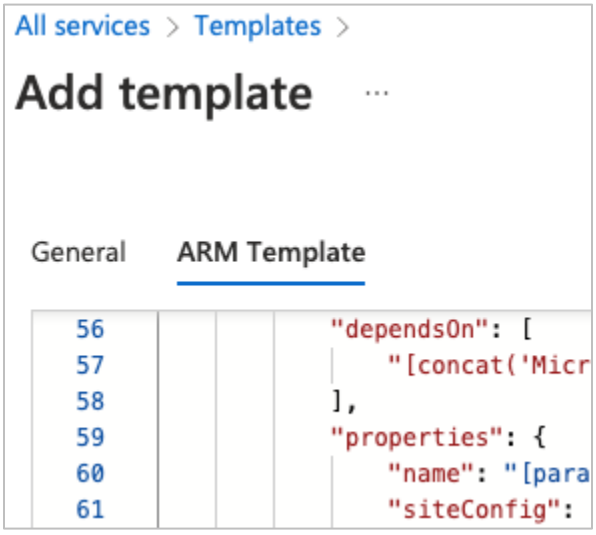

9. In the Templates, select the newly created template and click **Deploy**.

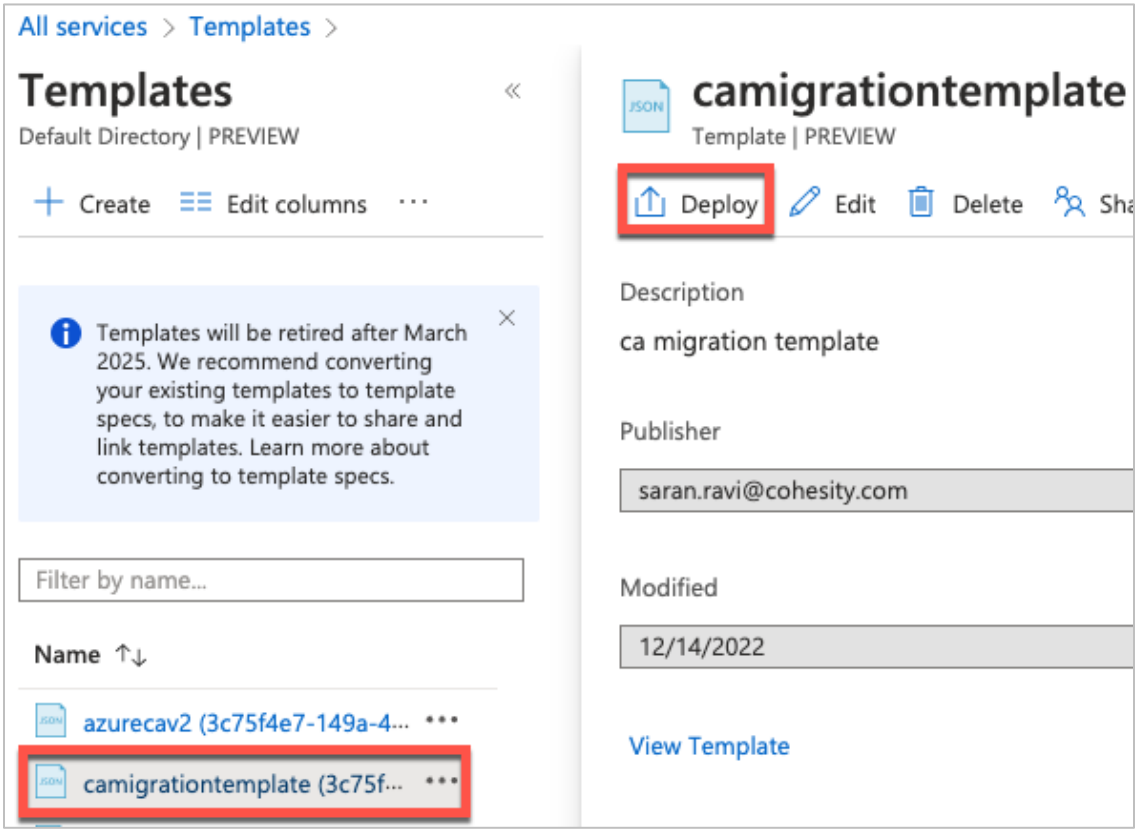

10. From the ART folder, copy the content from the file **parameters.json**.

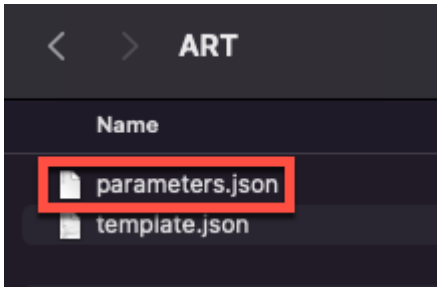

11. In the **Custom Deployment** page, click **Edit parameters**.

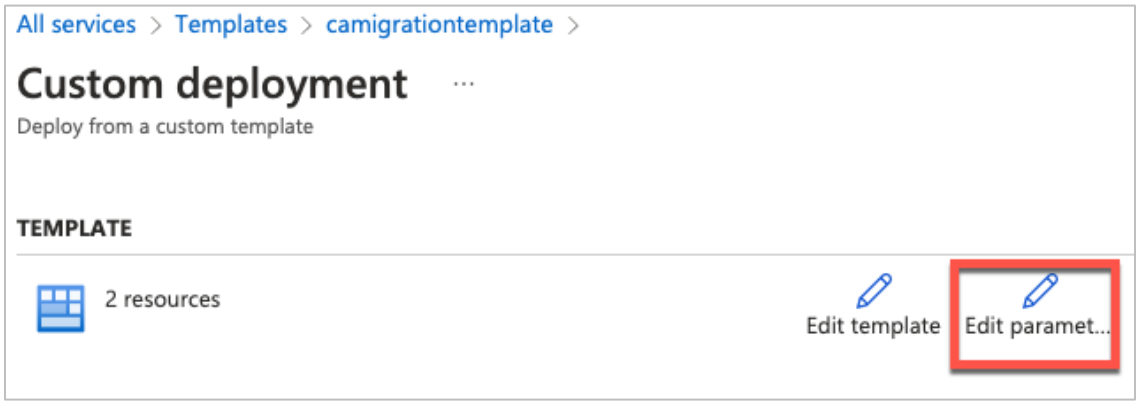

12. Replace the Edit Paramters with the content from paramters.json file and click **Save**.

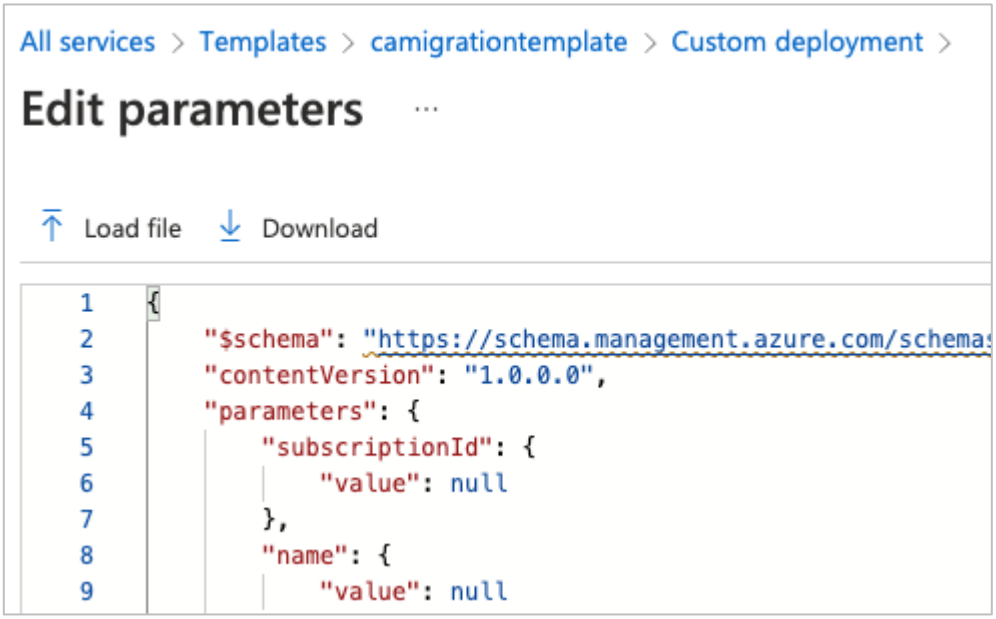

13. In the Custom deployment page, enter following details:

- **Resource group** — *<Name of the resource group where the storage account is created>*
- **Subscription Id** — *<Azure subscription Id>*
- **Name** — *<name for the azure function app>*
- **Location** — *<location of the storage account>*

- **Hosting Plan name** — **Consumption**
- **Server Farm Resource Group** — *<Resource group where the function app will get deployed>*
- **Storage Account Name** — *<Name of the storage account that you want to convert to CloudArchive Incremental Forever format>*

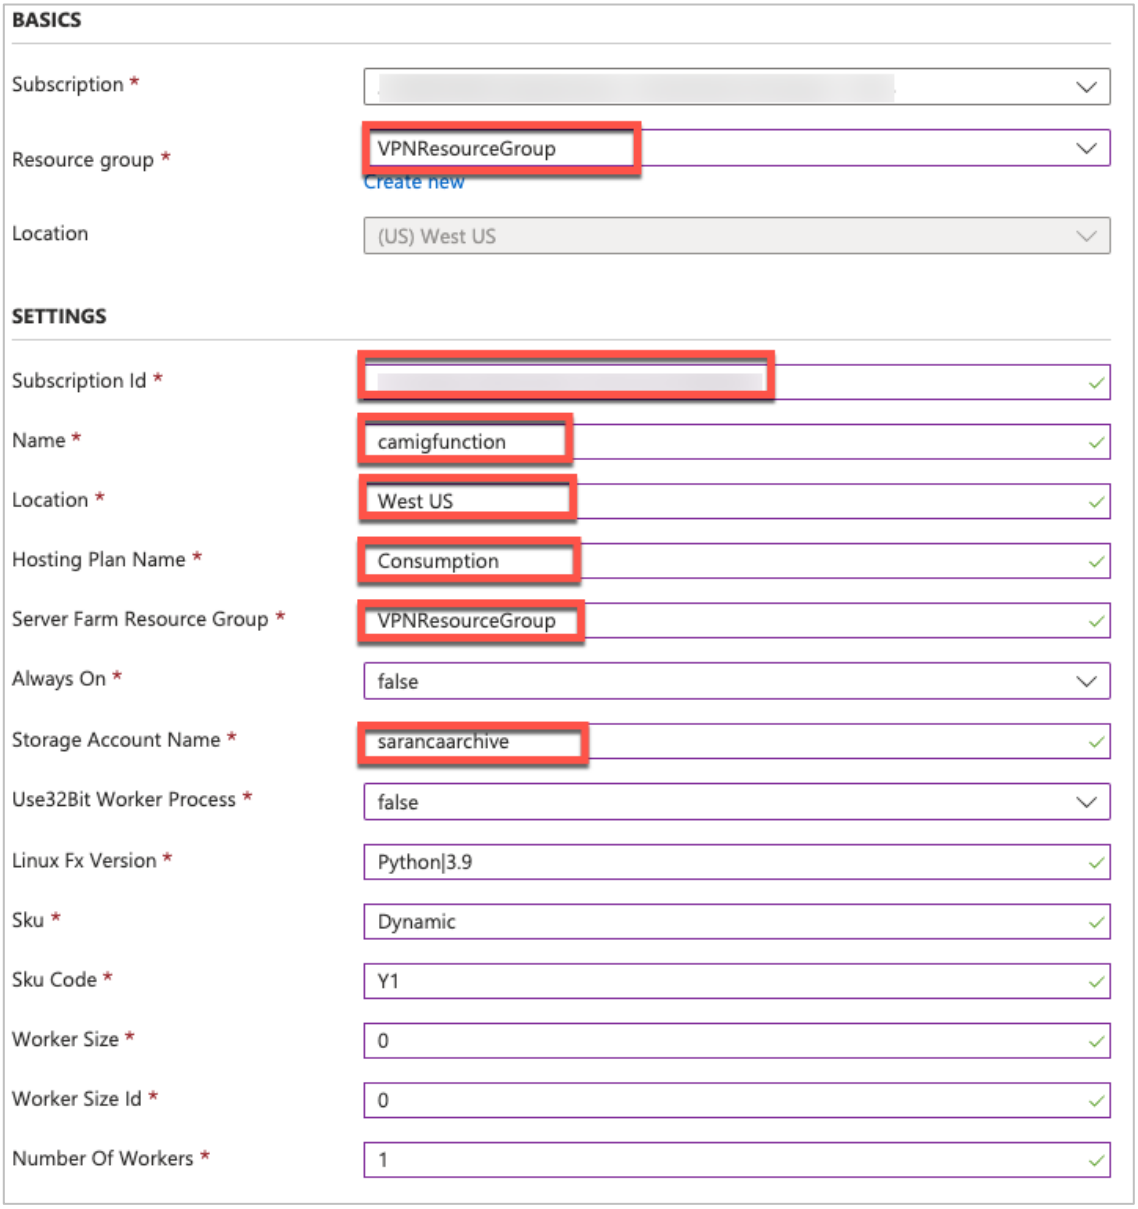

NOTE: Rest of the field values are auto-populated from the parameters file.

14. Select the agree terms checkbox and then click **Purchase** to deploy the function app.

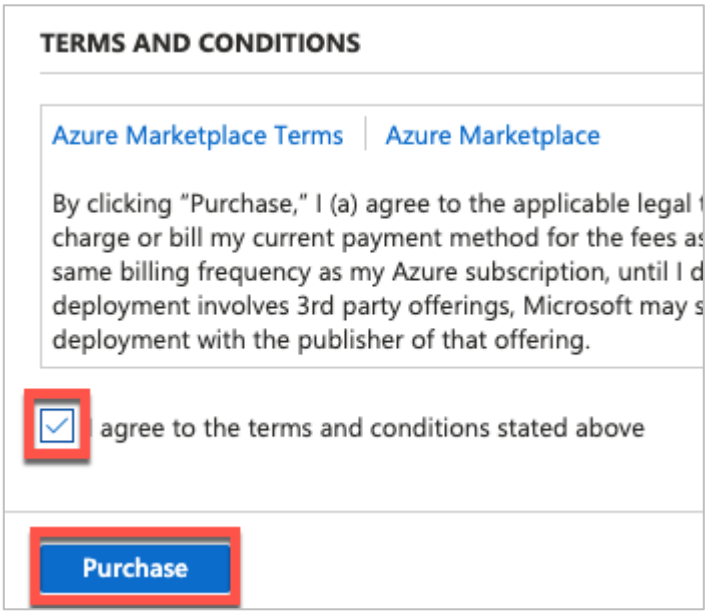

15. Navigate to **All services** > **Function App** and select the function.

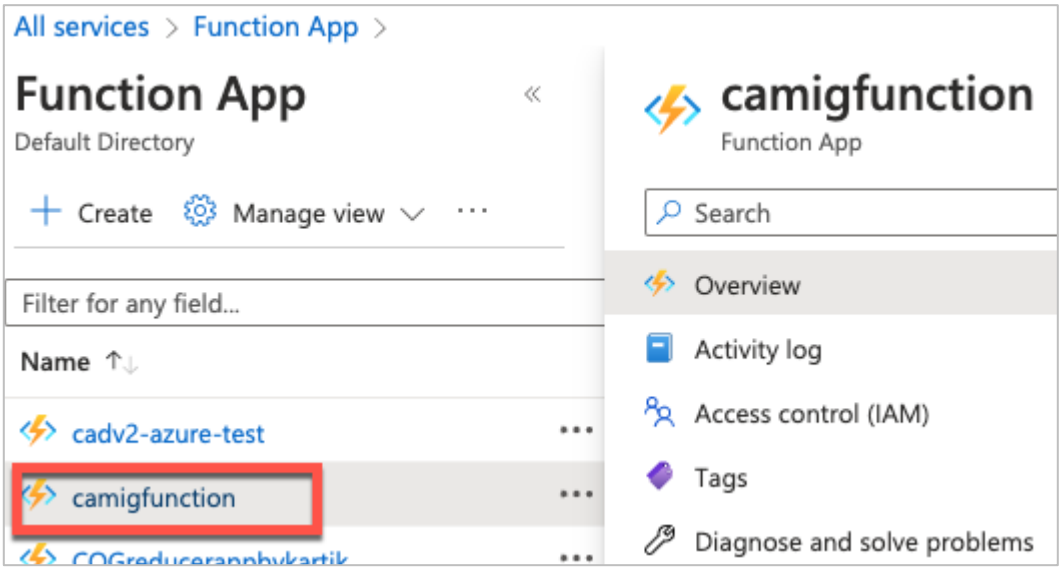

16. Click **Get publish profile** to download the profile.

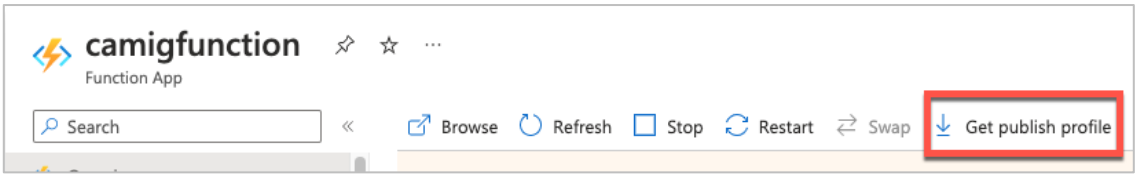

17. Open the downloaded file in a text editor and copy the value of **userPWD**.

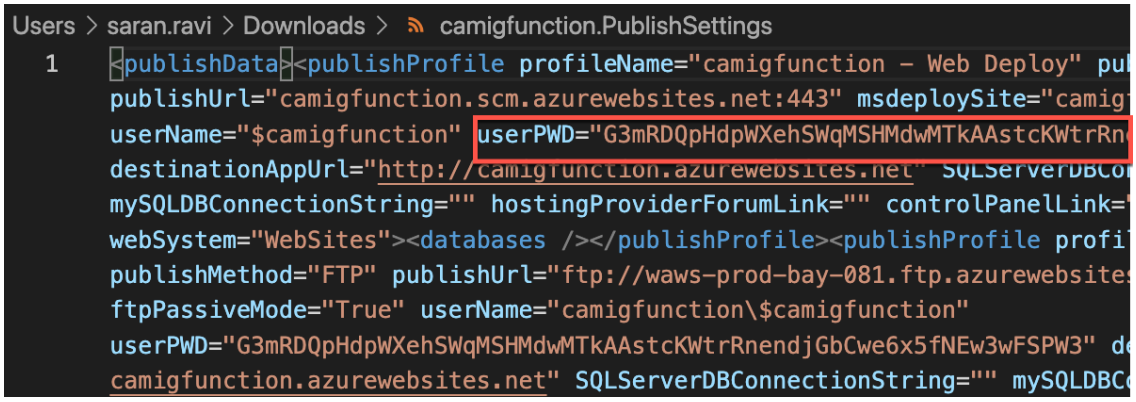

18. In the Edit external Target page, enter the **Storage Access Key** of the storage account, the function app name in the **Function App Name** field, and userPWD value in the **Function App Deployment Key** field. Click **Save** to register the target in CloudArchive Incremental Forever format.

NOTE: Incremental Forever Archival is now set as Enabled.

Appendix D—Register a New External Target for CloudArchive Incremental Forever (Azure)

1. Select **External Targets > Add External Target.**

   Select the Purpose as **Archival** and select and enter the required fields.

   - Storage Type: **Azure**
   - Storage Class: **Hot blob** Cloud type: **Standard**
   - Container Name: <*Enter the name of the container*>
   - Storage Account Name: <*Enter the name of the Storage Account*>
   - Storage Access Key: <*Enter the storage access key*>
   - External Target Name:<*Provide a name for the target*>
   - Archival Format: **Incremental Forever**

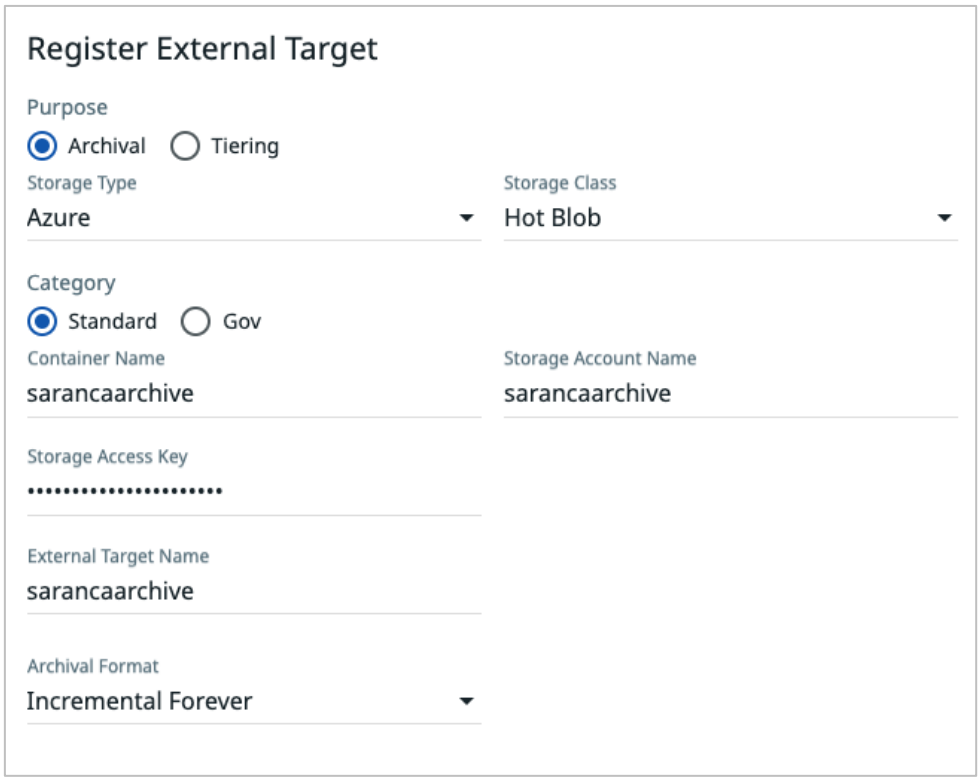

2. Once you change the Archival Format to Incremental Forever, download the **Azure Resource Template** (ART) from the admonition link shown below.

> **NOTE**: To register an Azure Blob target on an Azure Cloud Edition, there is no requirement to deploy the Azure Function App. Simply changing the Archival format to “Incremental forever” should be sufficient.

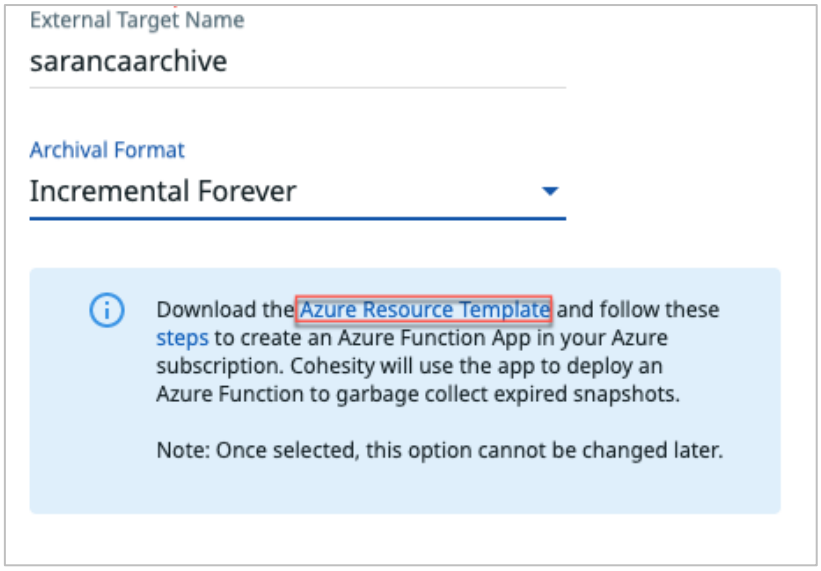

3. Log into Azure Portal.
4. Select **All services** > **Templates** > **+ Create**.

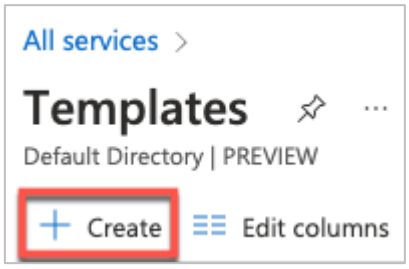

5. Enter a **Name** and **Description** for the template and click **Next:ARM Template**.

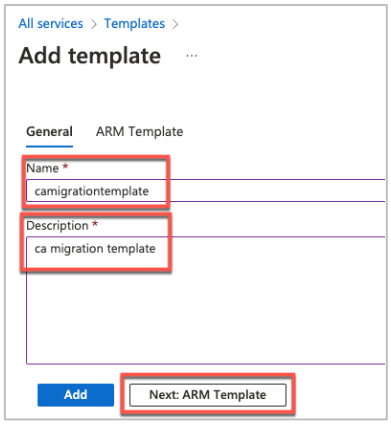

6. Unzip the downloaded ART folder and copy the content from the file **template.json**.

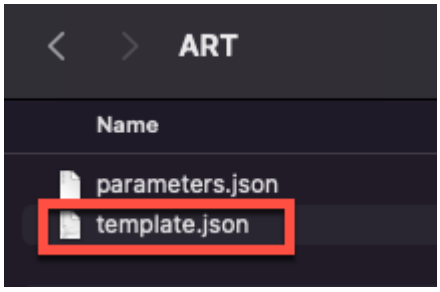

7. In the ARM Template section, replace the content from the template.json file and click **Add** to save the template.

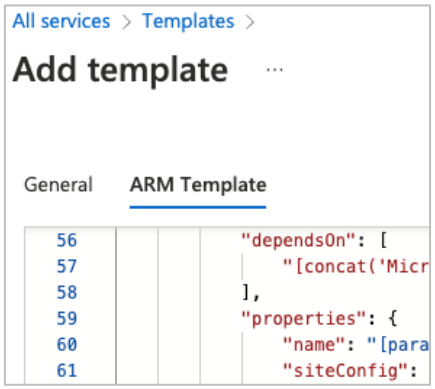

8. In the Templates, select the newly created template and click **Deploy**.

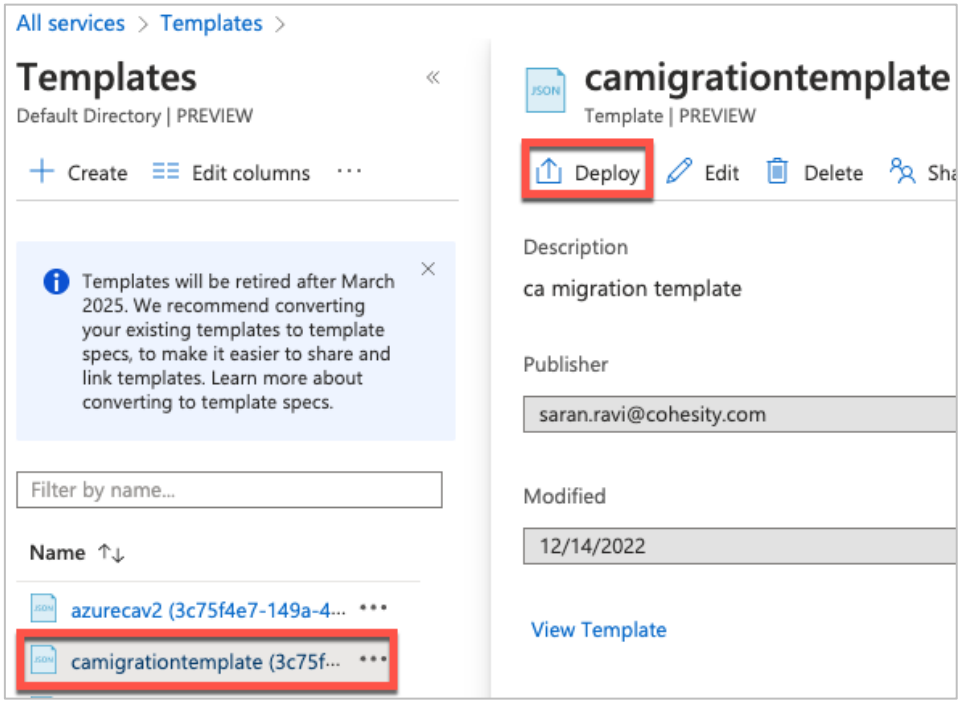

9. From the ART folder, copy the content from the file **parameters.json**.

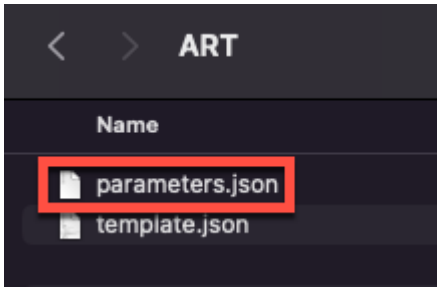

10. In the **Custom Deployment** page, click **Edit parameters**.

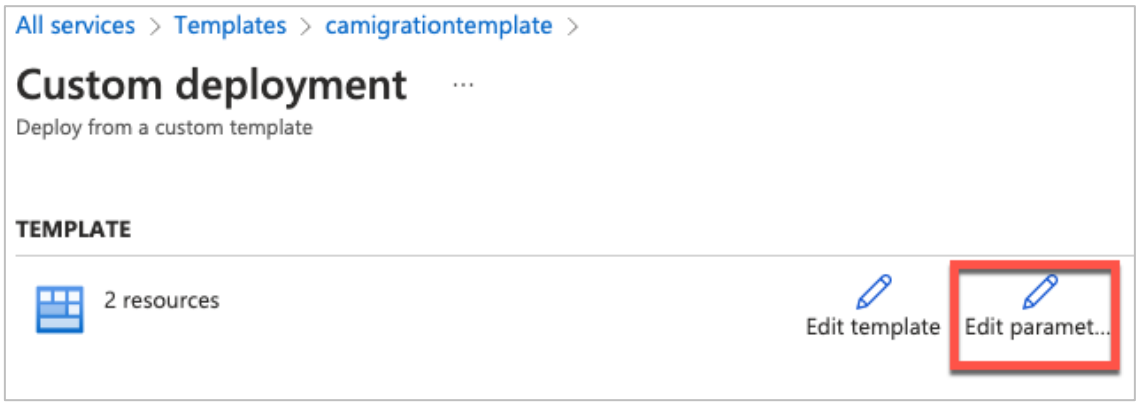

11. Replace the Edit Paramters with the content from paramters.json file and click **Save**.

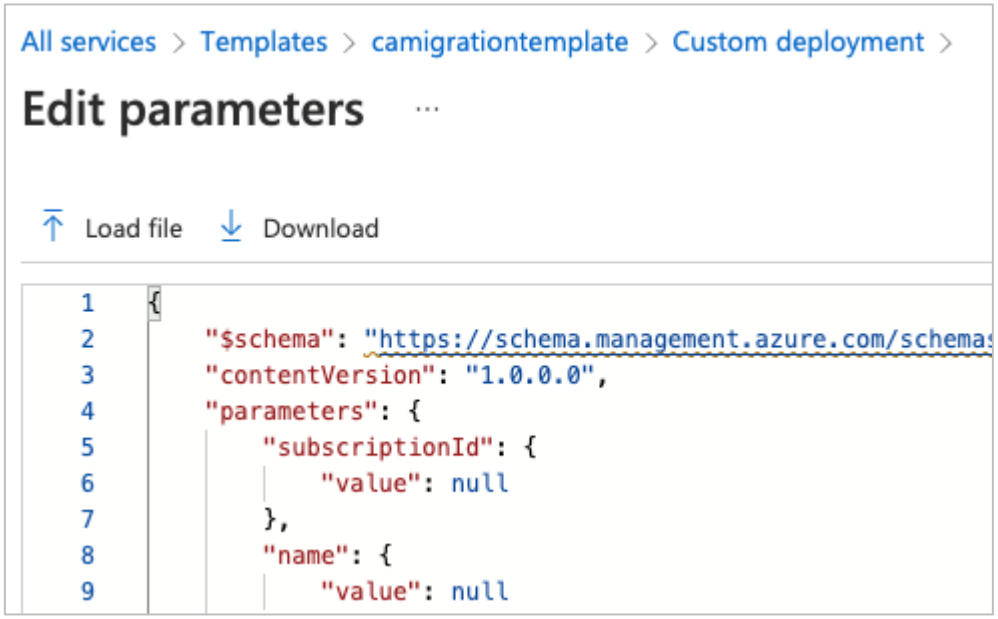

12. In the Custom deployment page, enter following details:

- **Resource group** — *<Name of the resource group where the storage account is created>*
- **Subscription Id** — *<Azure subscription Id>*
- **Name** — *<name for the azure function app>*
- **Location** — *<location of the storage account>*
- **Hosting Plan name** — **Consumption**
- **Server Farm Resource Group** — *<Resource group where the function app will get deployed>*
- **Storage Account Name** — *<Name of the storage account which you want to convert to CloudArchive Incremental Forever format>*

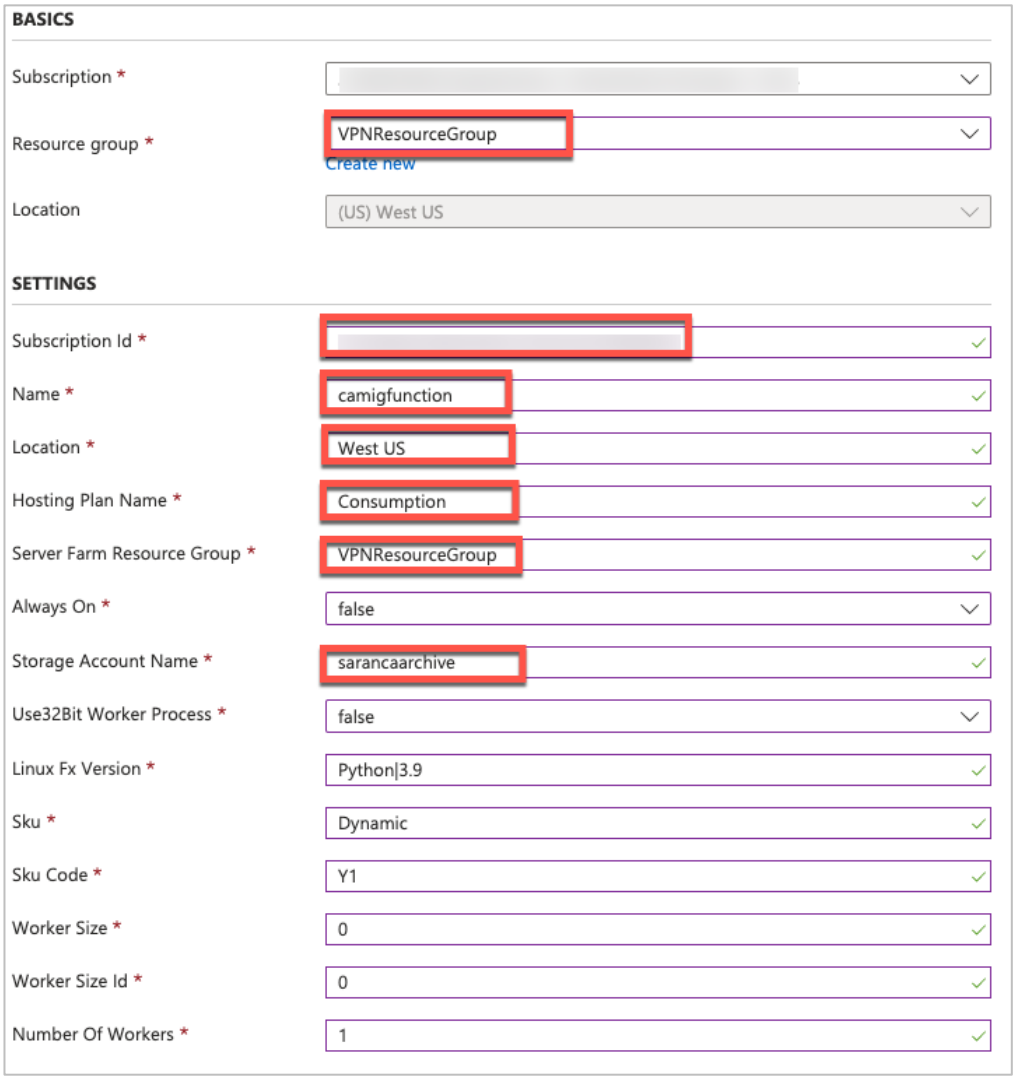

NOTE: Rest of the field values are auto-populated from the parameters file.

13. Select the agree terms checkbox and then click **Purchase** to deploy the function app.

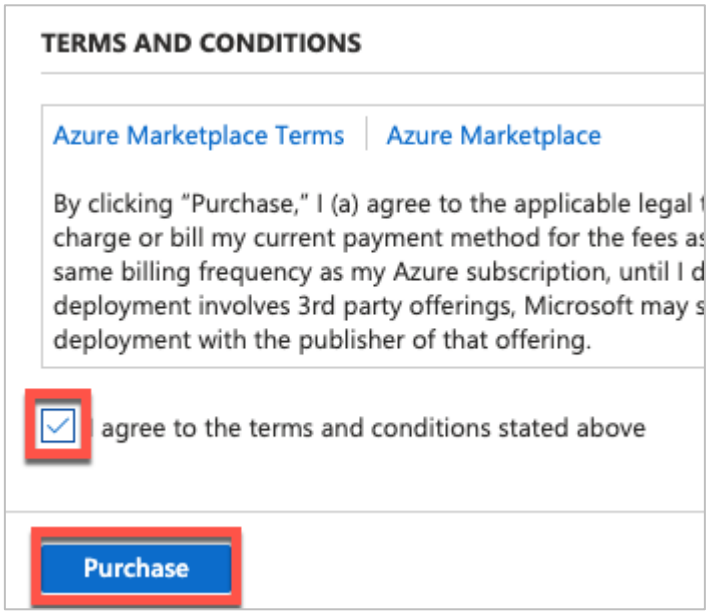

14. Navigate to **All services** > **Function App** and select the function.

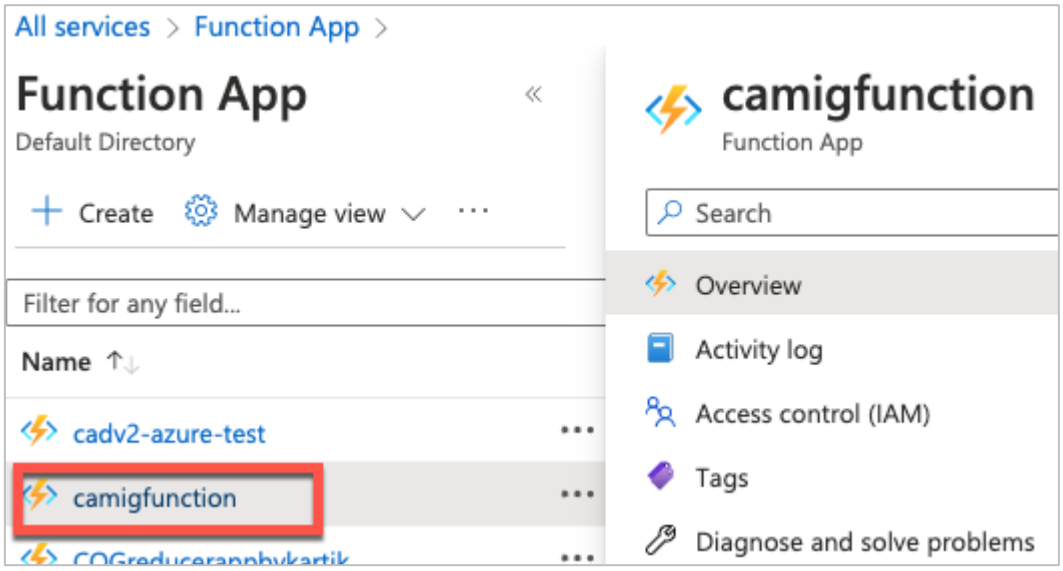

15. Click **Get publish profile** to download the profile.

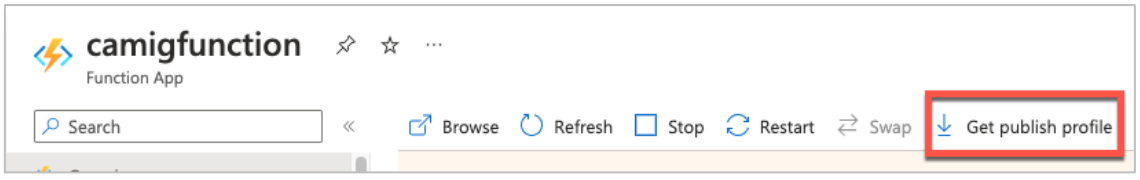

16. Open the downloaded file in a text editor and copy the value of **userPWD**.

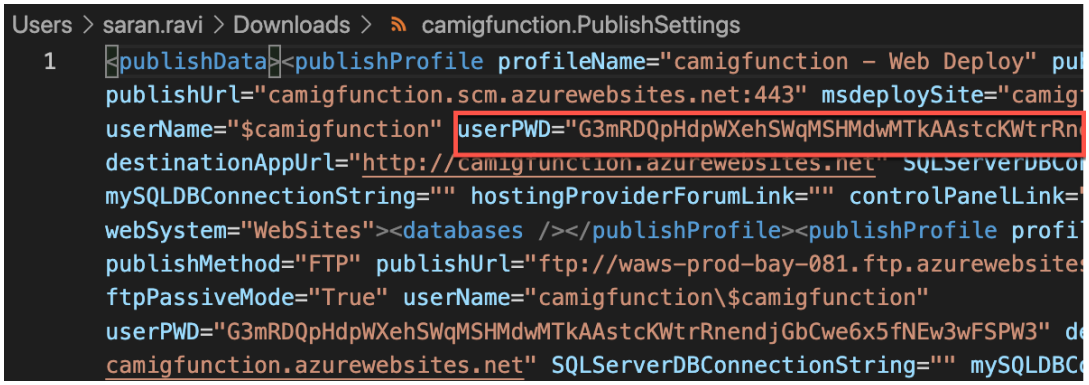

17. In the Add external Target page, enter the **Storage Access Key** of the storage account, the function app name in the **Function App Name** field, and userPWD value in the **Function App Deployment Key** field. Click **Register** to register the target in CloudArchive Incremental Forever format.

NOTE: Incremental Forever Archival is now set as Enabled.

Your Feedback

Was this document helpful? Send us your feedback!

About the Author

Saran Ravi is a Staff Technical Solutions Engineer at Cohesity. In his role, Saran focuses on Cloud and Kubernetes.

Other essential contributors include:

- Praveen Yarlagadda, Technical Director, Engineering
- Anirudh Kumar, Staff 2 Engineer, Engineering
- Adaikappan Arumugam, Director Product Solutions

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.3	Mar 2024	Minor updates
1.2	Dec 2023	Content updates
1.1	June 2023	Minor updates
1.0	Mar 2023	First full release

ABOUT COHESITY

Cohesity is a leader in AI-powered data security and management. Aided by an extensive ecosystem of partners, Cohesity makes it easier to protect, manage, and get value from data – across the data center, edge, and cloud. Cohesity helps organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring for malicious behavior, and rapid recovery at scale. Cohesity solutions are delivered as a service, self-managed, or provided by a Cohesity-powered partner. Cohesity is headquartered in San Jose, CA, and is trusted by the world's largest enterprises, including six of the Fortune 10 and 42 of the Fortune 100.

Visit our website and blog, follow us on Twitter and LinkedIn and like us on Facebook.