Tech Brief

Actifio GO

Backup and Disaster Recovery-as-a-Service for Google Cloud

Deploying Actifio GO in a Google Cloud Project

Last updated on January 2, 2023

Actifio GO is a Google Cloud based backup and disaster recovery solution for Google Cloud workloads. Actifio does not support the fresh deployment of Actifio GO to backup workloads in Google Cloud Platform (GCP). It is evolved to Google Backup and DR service. If you are planning to backup workloads running in GCP, deploy Google Backup and DR service.

This tech brief includes:

Planning an Actifio GO Deployment

When planning to deploy Actifio GO for your enterprise, consider:

Overview of the Deployment Process

When you deploy Actifio GO, you:

1. Plan the deployment and collect information and resources required in Actifio GO Deployment Prerequisites.
2. Deploy Actifio GO for Backups:
    a. Deploy the Actifio Global Manager (AGM) into your Google Cloud Project. This is the management plane. Note that the fresh deployments of AGM is not supported.

b. Deploy one or more Actifio Sky data mover appliances to do the work managed by AGM. You can deploy Actifio Sky only if you already have the AGM deployed.

Overview of the Service Architecture

Actifio GO for GCP is a SaaS service. The service architecture comprises several components that together deliver the service. The key components of the service are:

Actifio Global Manager: This is the management plane that resides in the Actifio Cloud. Each tenant of the service gets a dedicated and isolated management plane that connects to the customer's cloud Sky.

Actifio Sky: Actifio Sky is the data mover built with Actifio patented Virtual Data Pipeline (VDP) technology. Actifio Sky has the smarts to efficiently capture, move and manage the lifecycle of data within your enterprise.

Actifio Connector: Actifio connectors are light pieces of software that call the application native APIs to efficiently capture data from production applications in an incremental forever fashion. Additionally, Actifio connectors also provide the application awareness at the time of recovery.

The service management plane is provisioned within the Actifio Cloud in a dedicated VPC. This management VPC is peered to a VPC within your GCP account to establish network connectivity between the management plane and your account.

Standalone project deployment

The simplest installation topology is to add Actifio GO to an existing or new stand-alone project. In this topology the installation automation peers the automatically created AGM Project/VPC to a VPC that you nominate in a stand-alone project that you nominate.

You then deploy backup appliances into that standalone project. The deployer uses the CIDR range you supply to create the subnet in the AGM project. You cannot manage this subnet, but because it is peered to your VPC, it needs to use an IP address range that is not used in the network that you do manage (your VPCs and subnets).

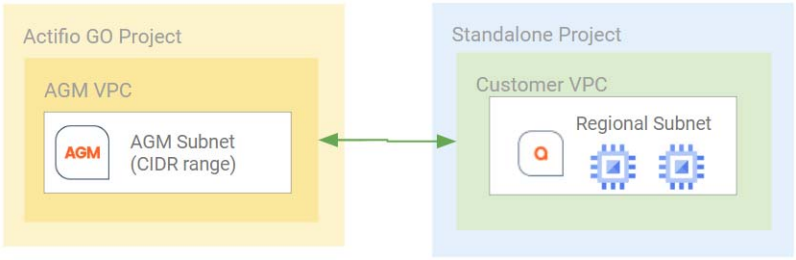

Shared VPC deployment

Google Cloud provides ultimate flexibility in configuring networks within the cloud. The most common architecture leverages shared VPCs. In this architecture, the network setup is configured in a central project also known as the host project. Virtual Private Cloud (VPCs) are shared with client projects that are called service projects.

An organization might have multiple shared VPCs which in turn may have multiple service projects within their GCP account.

Deploying Actifio GO in a shared VPC configuration

Actifio GO leverages Google Cloud VPC network peering technology to establish connectivity between a management plane in a Google managed project and your Google Cloud account. In other words, a VPC networking peer is established between a Google created VPC in a Google managed project (where your management console is housed) and a customer managed VPC in a customer project that you nominate.

There are several design considerations to do this:

- It is a best practice to deploy compute elements into service projects rather than their host projects.
- It is not possible to establish a VPC peer with a service project, since the service project does not 'own' the VPC being peered to. The initial peering must be to a host project, even if all backup appliances (and the data sources they are going to protect) are going to be deployed into service projects.
- It is common practice for Cloud customers to place all resources belonging to a specific department/use case within a single project for billing and chargeback purposes.
- Every time you deploy a new backup appliance, if the targeted project is not peered to the Google managed project/VPC where the management console is housed, then this peering will be automatically provisioned for you. If the targeted project is a service project then the host project details will also need to be supplied to the installer.
- Backup appliances need to be located so that agent-based backup traffic does not generate unnecessary network charges.
- For agentless backup of Compute Engine Instances, the location of the Backup Appliance is not a major design consideration. A single backup appliance can manage agentless Compute Engine Instance backups across multiple projects as well as multiple geos, regions and zones.

Given these considerations, there are several possible deployment topologies that you may deploy into. We will discuss two of them here:

- New dedicated backup project
- Existing host projects

New dedicated backup project

In this topology you deploy backup appliances into a new dedicated project. This architecture provides an isolation between Actifio GO service-side components and the production network. This architecture is also desirable from a billing and chargeback perspective because the Backup project in the diagram below provides a single billing report for all compute/storage and network resources consumed by the backup processes.

Note that in Google Cloud, projects and VPCs are global constructs. Hence this architecture allows customers to deploy backup appliances in any region within the backup VPC. For instance, in the example below, Host Project 1 might be used by a US based team with resources in us-east region and Host Project 2 might be used by a EU based team with resources in eu-west region. To ensure the backup appliances are closest to production projects, the backup appliance for project1 should be deployed in a subnet that is in us-east and the backup appliance for project2 should be deployed in a subnet that is in eu-west.

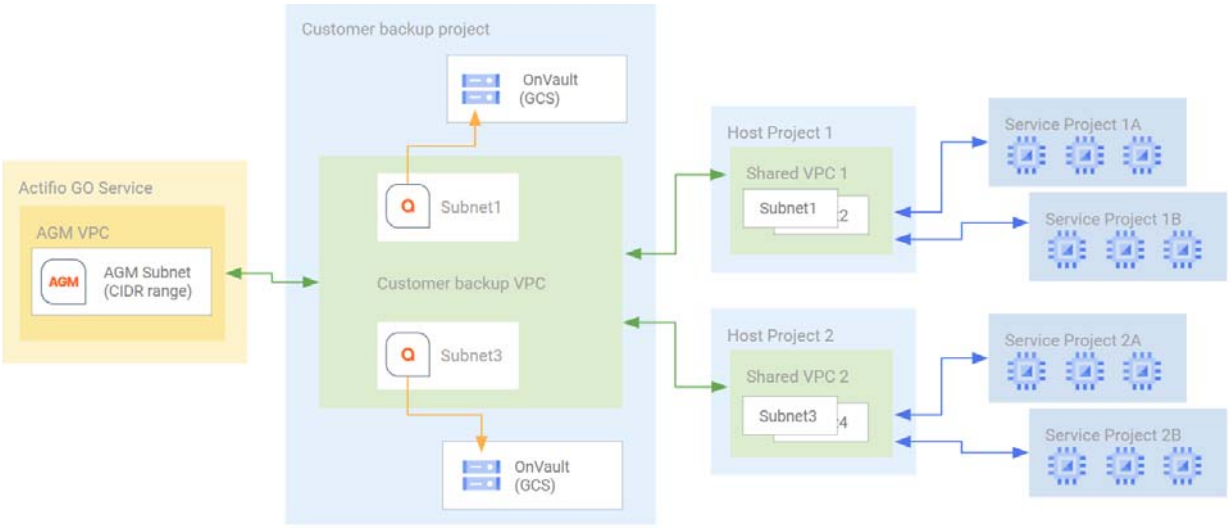

Existing host projects

Many organizations already have an existing Google Cloud deployment and may not wish to add additional projects and VPC peerings.

In this topology the installation automation peers the AGM VPC with an existing host project and deploys the backup appliances into existing service projects. This architecture avoids the requirement to add additional projects or additional VPC peers.

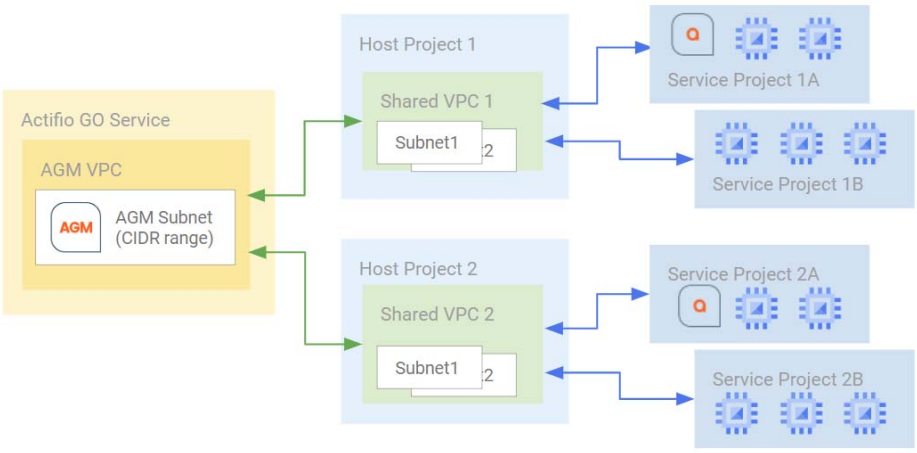

Additional Considerations: Peering and Performance

In addition to the Sky placement as discussed above, it is important to consider CIDR conflicts and performance considerations.

CIDR conflicts and VPC peering

VPC peering imposes certain constraints on the deployment architecture. Two VPCs cannot be peered if they have subnets using identical CIDR ranges. In the architecture above, if the Sky VPC is being peered to two production VPCs that have the subnets with identical or overlapping CIDR ranges, the VPC peering operation will fail. In such circumstances, it becomes necessary to separate the Sky deployments into separate projects to create isolation.

Performance considerations

Data traffic from production hosts flow to Actifio Sky using the Google backbone. The region and zone where the Sky is placed has an impact on the network egress costs incurred by the customer. For instance, if you have production assets in multiple regions within a VPC, it's best to deploy one Sky per region or zone in order to keep network costs low. Customers of Actifio GO service can deploy as many data movers as needed without additional license cost. The deployment wizard makes it easy to deploy new data movers.

Actifio GO Deployment Prerequisites

Before you begin the deployment of the Actifio GO, ensure that these prerequisites are reviewed and met:

- Actifio Service Account Email and Roles
- AGM and Sky Appliance VM Network Requirements
- OnVault Storage Bucket

Actifio Service Account Email and Roles

The Google Cloud Project in which the Actifio GO Sky will be deployed will use this Actifio Service account email address: **sky-launcher-195802@sky-launcher-195802.iam.gserviceaccount.com**.

You need these permissions to run a provided script for the installation.

- iam.roles.create
- iam.roles.get
- iam.roles.update
- resourcemanager.projects.getIamPolicy
- resourcemanager.projects.setIamPolicy

Click **Copy** to copy the provided script and execute the commands from Google Cloud Shell, which will create one or two IAM Roles depending on the type of project. The new role(s) will be assigned service email address and the correct permissions. After the installation is complete, this role is no longer required; it can be removed.

The IAM Role for a Standalone Project is ActifioGO. It is different from the two IAM roles for a Host & Service Project (ActifioGOHost and ActifioGOService).

Required policy overrides

During provisioning of Actifio GO, there are several Org policies that will (if implemented) prevent Actifio GO from being implemented. Five policy overrides are needed for success. These are detailed in Appendix A: Required Policy Overrides on page 14.

AGM and Sky Appliance VM Network Requirements

Identify these details for the Google Cloud Project in which the Sky appliance is to be deployed:

CIDR Range: CIDR IP range (valid private IP range) that is globally unique (to avoid VPC peering and duplicate IP address conflicts) in Google Cloud project to be used for allocation of IP address for AGM appliance.

Mask: Subnet Mask to be applied on CIDR range for selection of IP address to be applied on AGM

VPC: Name of VPC to be used for installation of Sky in GCP project

Subnet: Name of subnet to be used for installation of Sky appliance

Region: Google Cloud Region to be used for installation of Sky appliance

Zone: Google Cloud Zone to be used for installation of Sky appliance

Note: *The VPC and Subnet must exist prior to deployment. The wizard will not create these automatically in the GCP Project. Application-consistent backups for VMware VMs require VMware Tools to be enabled.*

OnVault Storage Bucket

Actifio OnVault enables GCS storage for storing backups. Configuration of an OnVault pool is an optional step that can be skipped as part of initial deployment and performed after a successful deployment; see ***Configuring Actifio OnVault*** at https://docs.actifio.com/Actifio-GO/.

To configure an OnVault pool, you will require the following details:

Bucket ID: ID of the Vault Bucket

Service account: Service account associated with the OnVault bucket

KeyFile: Bucket Authentication key file in P12 format

Note: For an OnVault pool, you will require to create a storage bucket configured to store the data in the desired region. You can use any of the storage classes i.e., Standard, Nearline, Coldline, and Archive storage. Make sure that the access control type has been set to "Uniform" (recommended) / "Fine grained" on the GCS bucket and private Google access has been enabled on the subnet.

Deploying Actifio GO for Backups

Deploying Actifio GO for backups involves:

1. Deploying Actifio Global Manager (AGM)
2. Deploying the Sky Appliance(s)

Deploying Actifio Global Manager (AGM)

Actifio GO does not support fresh deployment of AGM. You can deploy the Sky appliances only if you already have an AGM deployed.

Deploying the Sky Appliance(s)

Actifio Sky is responsible for running all backup processes. Sky is a data mover VM that is deployed close to where the production applications are running.

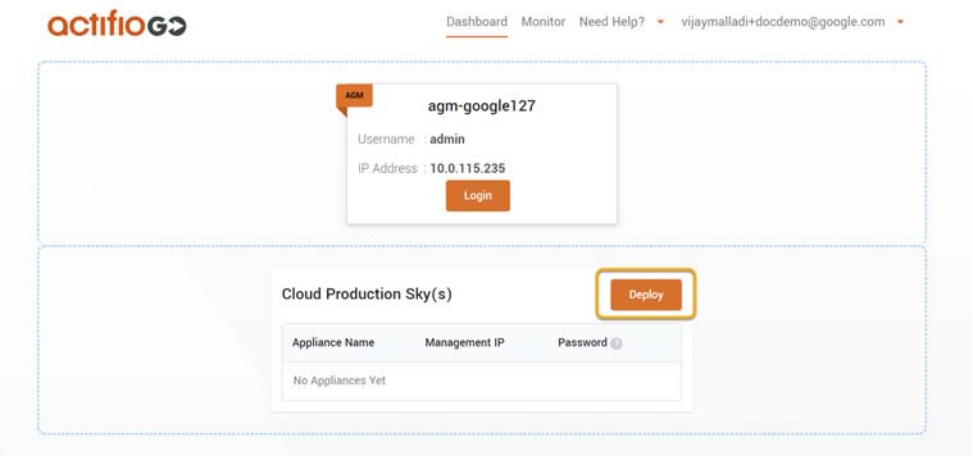

Deploying a Sky Appliance

1. Return to the Dashboard.
2. In the Cloud Production Sky(s) box, select **Deploy**.

Note: *Refer to this video to understand the Sky deployment process: https://youtu.be/Jsy91pe_3n8*

---

3. Enter the name of the backup project created to host Actifio Sky. Click **Validate** to verify the permissions.
   The wizard will proceed to the next step if the role validation is successful.
4. The next step collects the input required to install Sky. All text entries must be lower case, and hyphens (-) are OK.
   Sky can be deployed in a Standalone Project or in a Service Project. Slide the toggle to set which project type to use for the deployment. By default it is Standalone Project. To deploy the Sky on Service Project, slide the toggle to the right.
   Depending on the project type you see one or two scripts below. You will copy and execute the scripts.

Deploy Sky on Service project: Select whether to deploy this Sky as a Standalone project or a Sky as a Service project.

- **Standalone Project Id**: Find your Project ID in the Home screen of your project in Google Cloud Console. Then copy the script below the setting and execute it in Google Cloud Shell to create the IAM role.

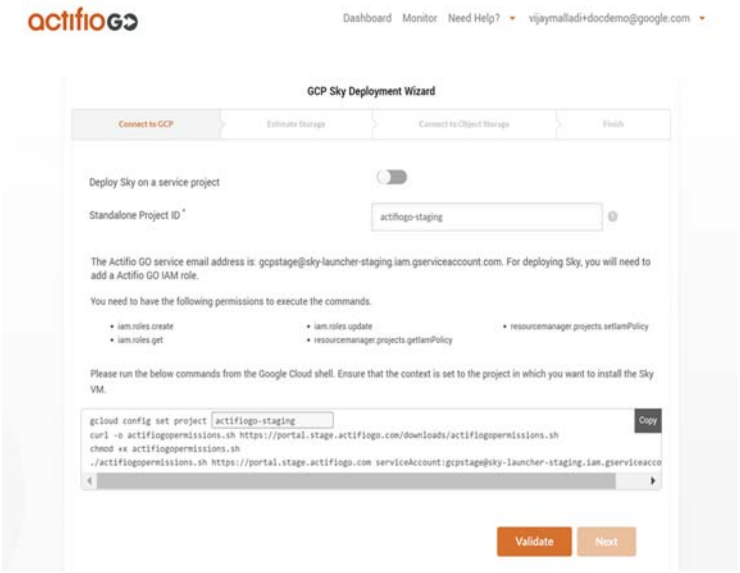

- **Sky as a Service project**: Slide the toggle to the right. Two new fields replace the Standalone Project ID field.
  Host Project ID: This is normally the same project where you have shared VPCs.
  Service Project ID: The ID of a project attached to the host project by a Shared VPC Admin.
  Then copy the two scripts below the setting and execute both of them in Google Cloud Shell to create the IAM role.

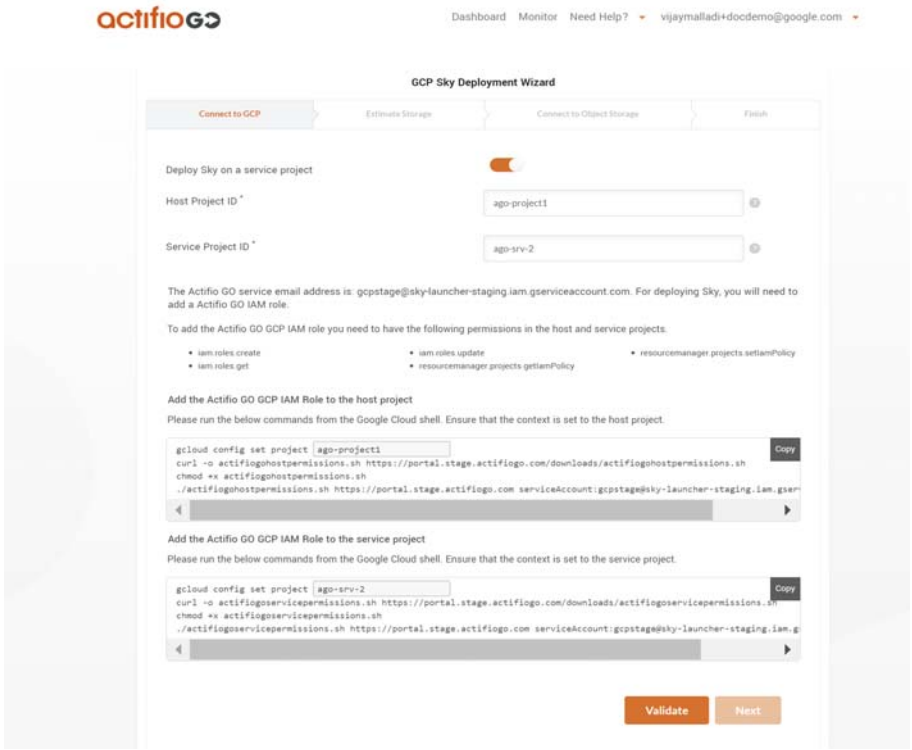

5. Click Validate to validate the information and move to the network information summary page:
   - **CIDR Range**: AGM must be deployed within this range for communication with the Sky.
   - **Appliance Name**: Name for the Sky VM
   - **Network Tags** (optional): Tags that you may want to apply for this VM.
   - **VPC/Subnet**: VPC/Subnet where the Sky will reside, lowercase only, no spaces or special characters.
   - **Sky Region/Sky Zone**: Region/Zone where the Sky will reside.
   - **Primary and Secondary DNS** (optional): Enter the details of the DNS server, if you intend to use a custom DNS server for name resolution.

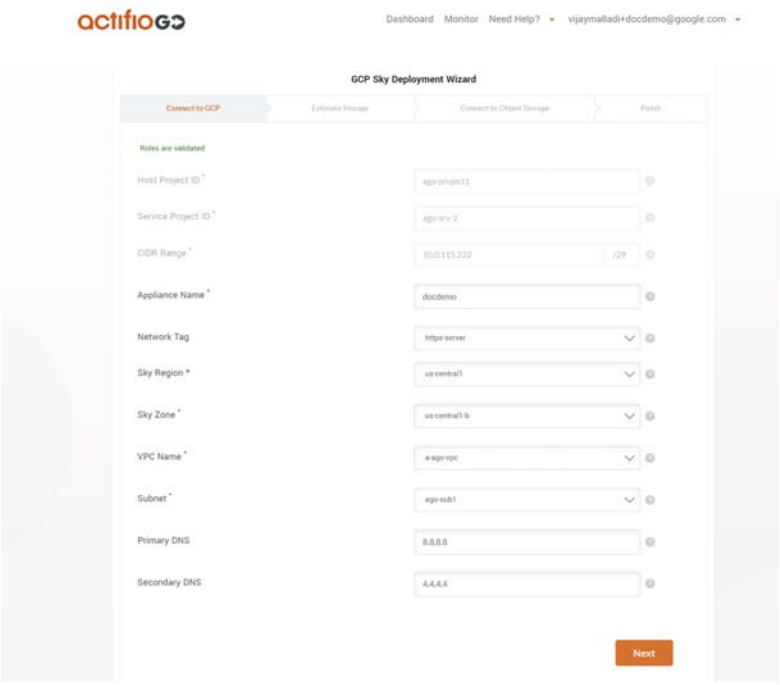

6. Click **Next** to proceed to storage estimation. In this step, the wizard helps you estimate storage required for the snapshot pool. This is a simple estimator to get started. More storage can be added later.
   - **GCE VMs and Agentless GCVE VMs** use PD snapshots, so they do not need a large snapshot pool. If this is all you need, select this option.

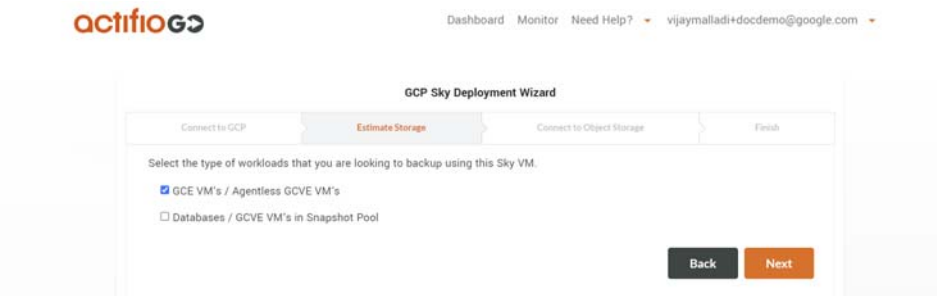

   - Databases and GCVE VM images do not use PD snapshots, they require the images to be captured in a snapshot pool. This requires storage; select **Databases/GCVE VMs in Snapshot Pool** and enter information about application size, change information, and how long you want to retain the images in the snapshot pool. You will get a snapshot pool between 4TB and 64TB. The estimator will determine the snapshot pool size required.
     If your needs are projected to be greater than 64TB, you will get a 64TB pool and then you must contact customer support to add additional pools to meet your requirements.

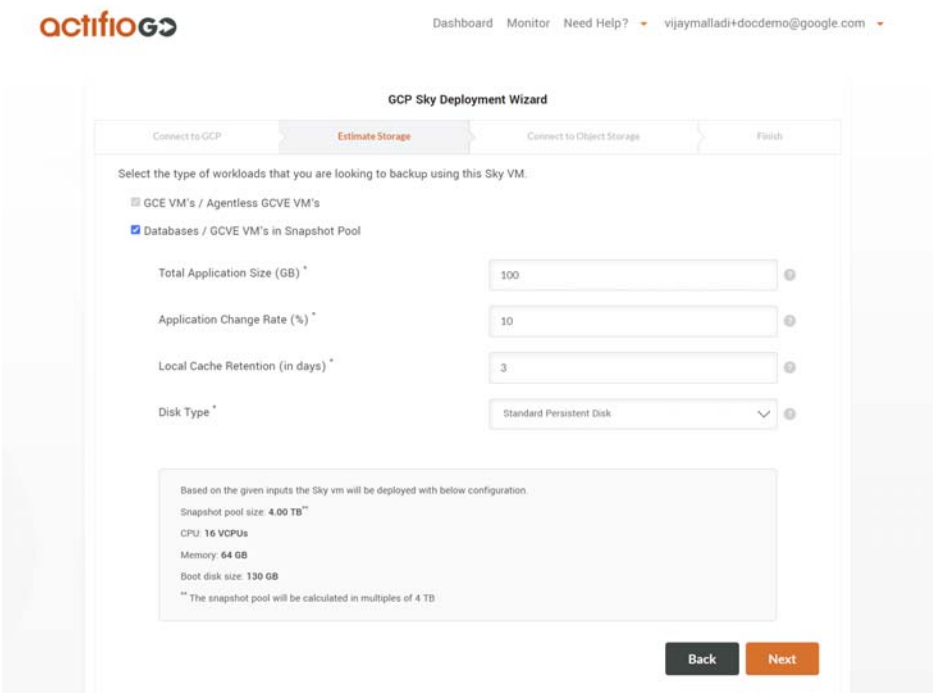

7. Optionally, you can provision a GCS bucket for Long Term Data Retention used by AGM as an OnVault Pool. This requires the creation of a bucket and credentials to the bucket in P12 format. Steps a-d may be skipped if you are not going to configure the GCS bucket at this time.

   a. Create a storage bucket configured to store the data in the desired region. You can use any of the storage classes i.e., Standard, Nearline, Coldline, and Archive storage. Make sure that the access control type has been set to "Uniform" (recommended) / "Fine grained" on the GCS bucket and private Google access has been enabled on the subnet.

   b. Create a service account that will be used to access the bucket. Create a new role that contains the following 6 rights:

      - storage.buckets.get
      - storage.objects.create
      - storage.objects.delete
      - storage.objects.get
      - storage.objects.list
      - storage.objects.update

   c. Add your service account user to this role for your bucket.

   d. Return to the GO Deployment Wizard and select the Pool Type as below:

   - Select the Google Cloud Storage option only if the "Uniform" access control type has been enabled on your storage bucket. This option supports all four Google object storage classes i.e., Standard, Nearline, Coldline, and Archive Storage.
   - Select the Google Coldline Storage or Google Nearline Storage option only if the "Fine grained" access control type has been enabled on your storage bucket.

   e. Specify bucket name, desired OnVault pool name, and the service account (in email syntax), as well as the key file in P12 format for the service account created in step b. Then click **Next**.

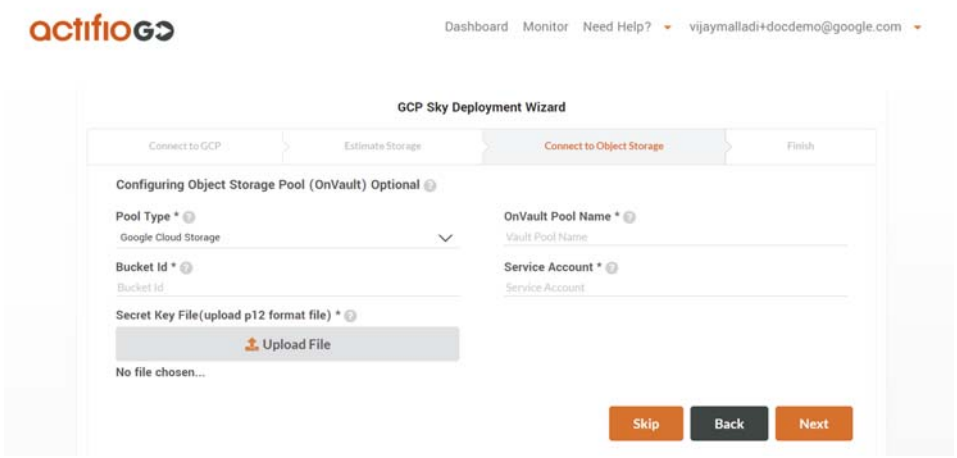

8. The last step provides an opportunity to verify the input parameters one last time before starting the installation. Check especially for upper-case characters (all entries must be lower case, and hyphens (-) are OK.)
If everything is good, click **Start Deployment**. Deployment can take 30 minutes or more.

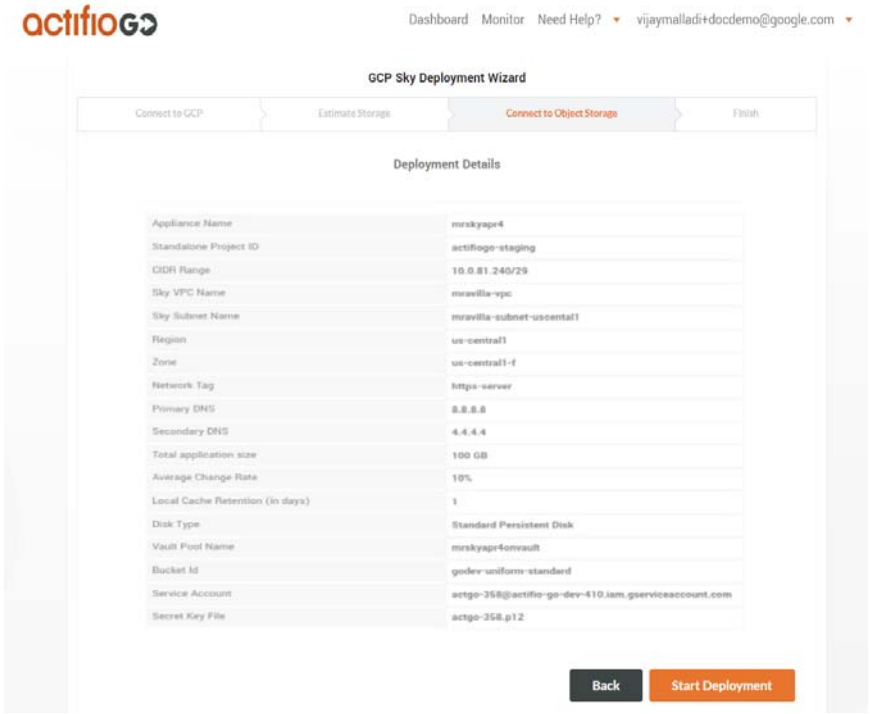

9. When the deployment is finished, you see a summary page with three actions:
    - **Add another Sky** brings you back to Step 2. Only select this to add another Sky.
    - **Access AGM** launches the AGM.

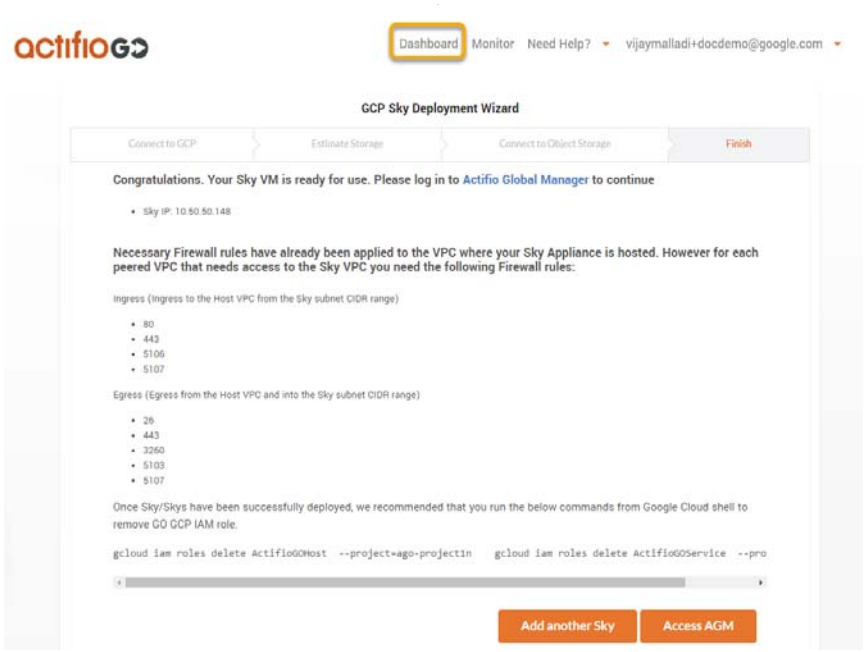

Sky Deployed for a Standalone Project

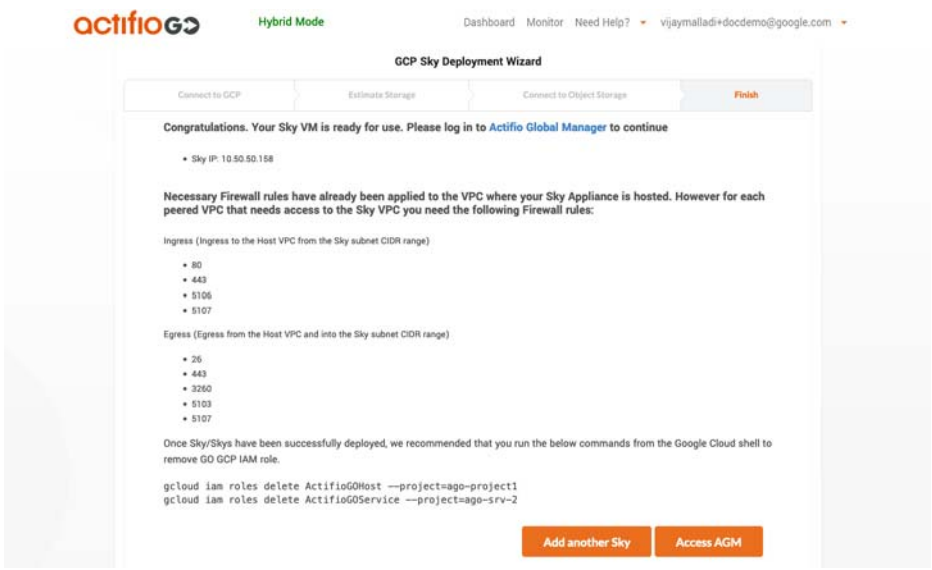

Sky Deployed for a Service Project

Disaster Recovery Orchestration

Actifio GO can carry out disaster recoveries by doing individual workload recoveries with an externally scripted orchestration. You can use AGM APIs to trigger recoveries and run custom scripts on servers where data has to be recovered. Sample scripts are available on https://github.com/Actifio/AGMPowerLib/blob/main/README.md.

These videos offer an overview of the API process:

https://www.youtube.com/watch?v=H89nlhgDIOk&list=PLS5jq0z48j6C-35A9nJfY8wzy5d087cQ9&index=26&t=141s

https://www.youtube.com/watch?v=gAkR3yUnmFM&list=PLS5jq0z48j6C-35A9nJfY8wzy5d087cQ9&index=27&t=3s

Next Steps

The steps above configure Actifio GO for backup and DR. Production workloads can be onboarded for backup. Please review the documentation section for more information on specific workloads.

- Documentation Site: **https://docs.actifio.com/Actifio-GO**
- Protecting and Recovering GCE Instances: https://docs.actifio.com/Actifio-GO/PDFs/GCEInstancesBackupAndRecovery.pdf
- Protecting and Recovering Microsoft SQL Databases and Instances: https://docs.actifio.com/Actifio-GO/PDFs/MS-SQL-Server_BackupAndRecovery.pdf
- SAP HANA DBA Guide: https://docs.actifio.com/Actifio-GO/PDFs/SAP-HANA_BackupsInActifioGO.pdf
- Oracle DBA Guide: https://docs.actifio.com/Actifio-GO/PDFs/DBAOracle.pdf
- Getting Started Guide: https://docs.actifio.com/Actifio-GO/PDFs/Introducing.pdf
- Network Administrator's Guide: https://docs.actifio.com/Actifio-GO/PDFs/NetworkConfiguration.pdf

Appendix A: Required Policy Overrides

During provisioning of Actifio GO, there are several Org policies that will (if implemented) prevent Actifio GO from being implemented. These five policy overrides are needed for success.

Organizational Policies

The following IAM Organizational policies may need to be modified:

1. Enable cross domain sharing.

   `https://console.cloud.google.com/iam-admin/orgpolicies/iam-allowedPolicyMemberDomains`

   Allow the following Domain while running the deployment: `C016sljgj`

2. Allow VPC Peering

   `https://console.cloud.google.com/iam-admin/orgpolicies/compute-restrictVpcPeering`

   Allow the following Org while running the deployment (use this syntax): `under:organizations/136083516469`

3. Allow non Shielded VMs

   `https://console.cloud.google.com/iam-admin/orgpolicies/compute-requireShieldedVm`

   Turn enforcement off while deploying Sky Appliances.

4. Allow trusted image project

   `https://console.cloud.google.com/iam-admin/orgpolicies/compute-trustedImageProjects`

   Add the following project while running the deployment: `projects/sky-launcher-195802`

5. Enable service account key generation

   `https://console.cloud.google.com/iam-admin/orgpolicies/iam-disableServiceAccountKeyCreation`

   Allow service key creation, at least while the keys are being generated.