DEPLOYMENT GUIDE

Infoblox vDiscovery for GCP

Table of Contents

Overview

Introduction

Infoblox vDiscovery provides enhanced visibility of your networks and virtual machines, and automatic creation of DNS records for discovered IP addresses.

With Infoblox vDiscovery, you will find an easy to deploy and cost-effective solution that enables visibility, reporting and automation of your network and VM resources across multiple cloud platforms, including Google Cloud Platform, or GCP, bringing all this data under a single pane of glass. In this guide, you will be introduced to Infoblox vDiscovery for GCP.

Prerequisites

The following are prerequisites for Infoblox vDiscovery with GCP:

- Valid subscription and login to GCP.
- Ability to create (or obtain) the key for a service account.
- TCP port 443 access from the Infoblox appliance that will run vDiscovery.
- Be able to resolve and access common resources from the Infoblox appliance that will run vDiscovery, such as:
  - accounts.google.com
  - oauth2.googleapis.com
  - www.googleapis.com
  - gserviceaccount.com

Note: The Cloud Network Automation (CNA) license in NIOS is optional.

Basic Workflow

The following bullet points outline the basic steps involved with creating a vDiscovery task for GCP:

- Sign in to the **GCP Console** (https://console.cloud.google.com/).
- Create/review your service account that will be used by vDiscovery and verify that the appropriate role(s) is assigned.
- Generate/obtain the key for the service account that will be used by vDiscovery.
- In NIOS, navigate to either the **Data Management** or **Cloud tab**.

- Create the vDiscovery task.
- Run the vDiscovery task.

Enabling GCP for vDiscovery

Service Account

To enable the connection from vDiscovery to GCP and for it to work properly, you must use a service account with the appropriate permissions assigned to it. In GCP, this is done with roles and can be accomplished using the primitive role, predefined roles or custom roles. If in doubt, use the 'Viewer' (primitive) role, as is described below using a new account as an example.

To create a service account:

1. In the GCP Console, expand the navigation menu and navigate to **IAM & admin** → **Service accounts**.

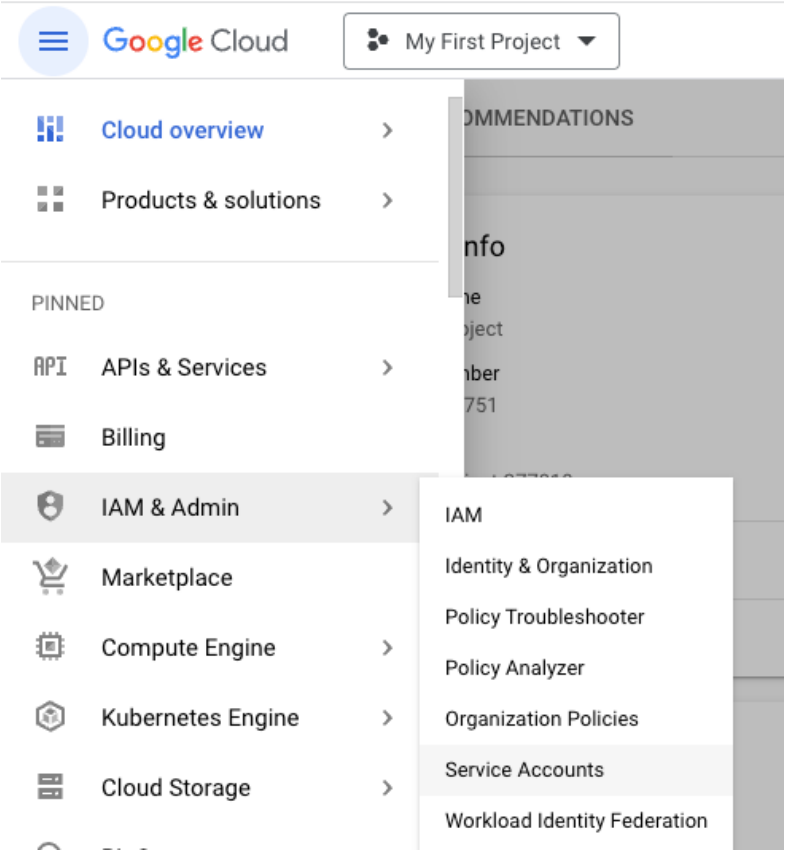

2. Click **CREATE SERVICE ACCOUNT**.

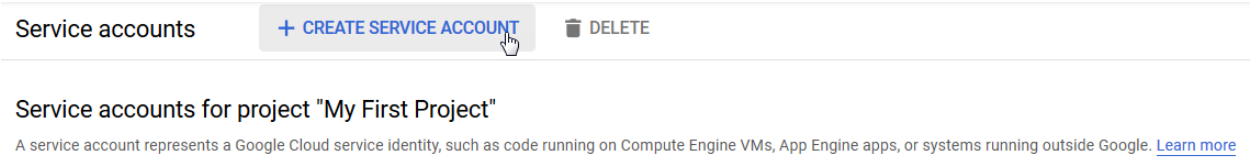

3. Enter a display name for your service account.
4. Review the Service account ID.
5. Optional: Enter a Service account description.
6. Click **CREATE AND CONTINUE**.

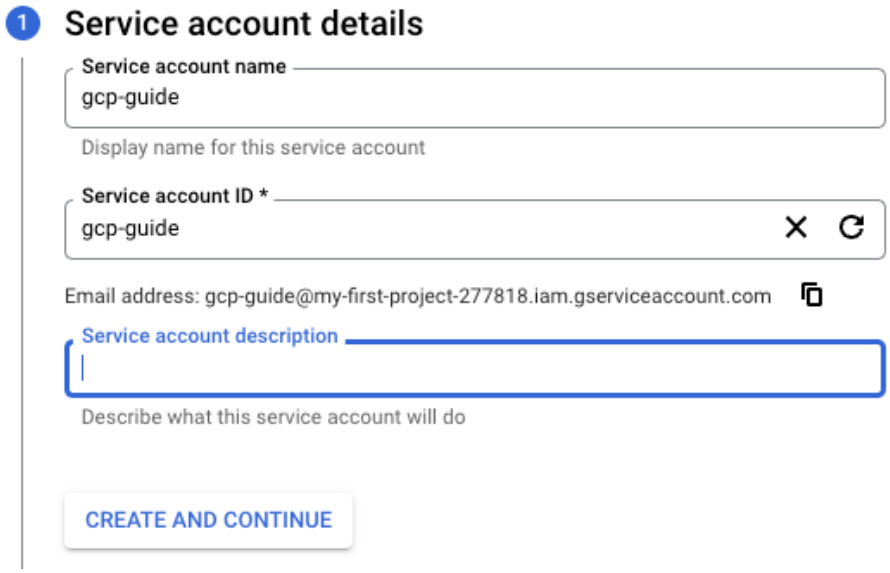

7. Expand the **Select a role** menu.

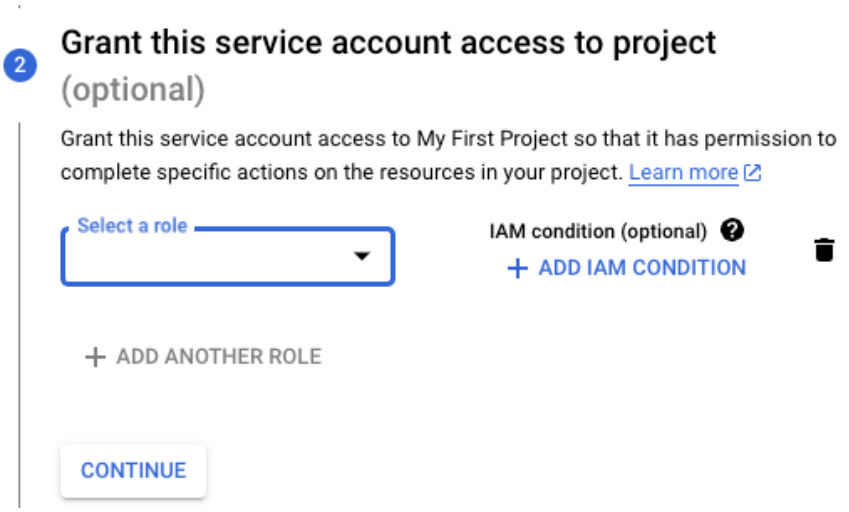

8. Select **Project** → **Viewer**.

9. Click **Continue**.

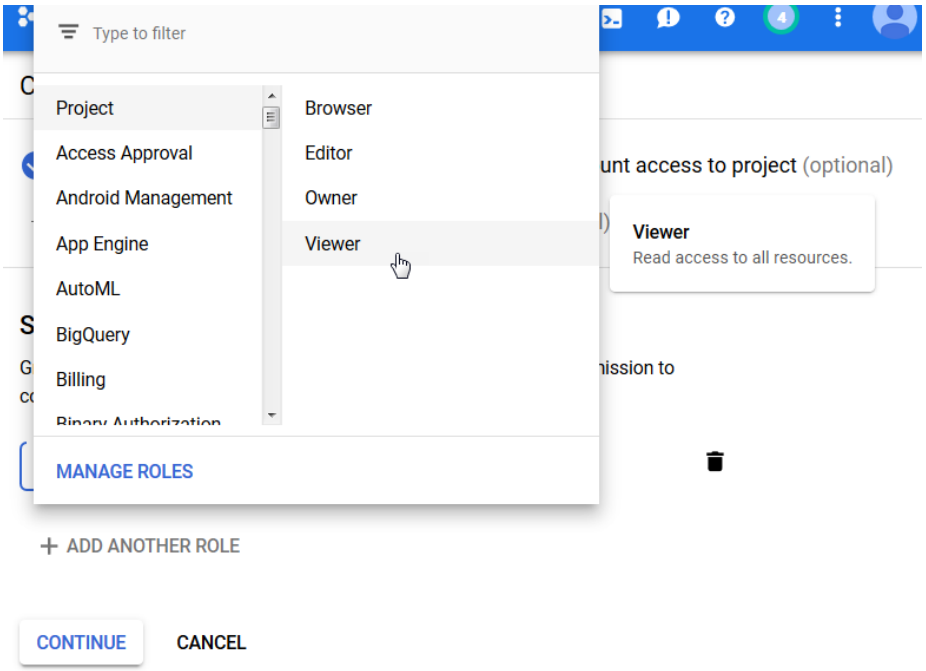

10. Click **DONE**.

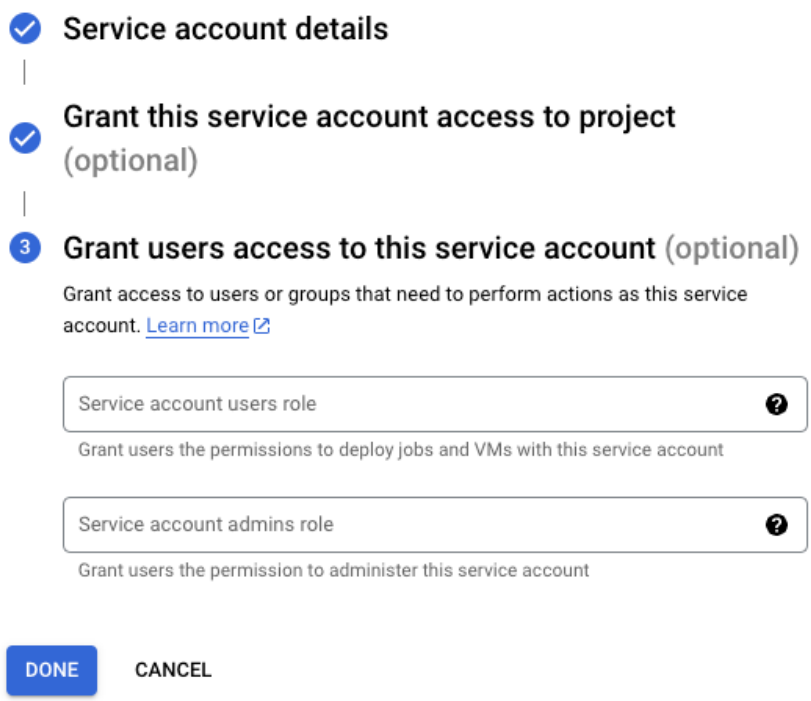

11. On the Service Accounts page, click on the new service account.

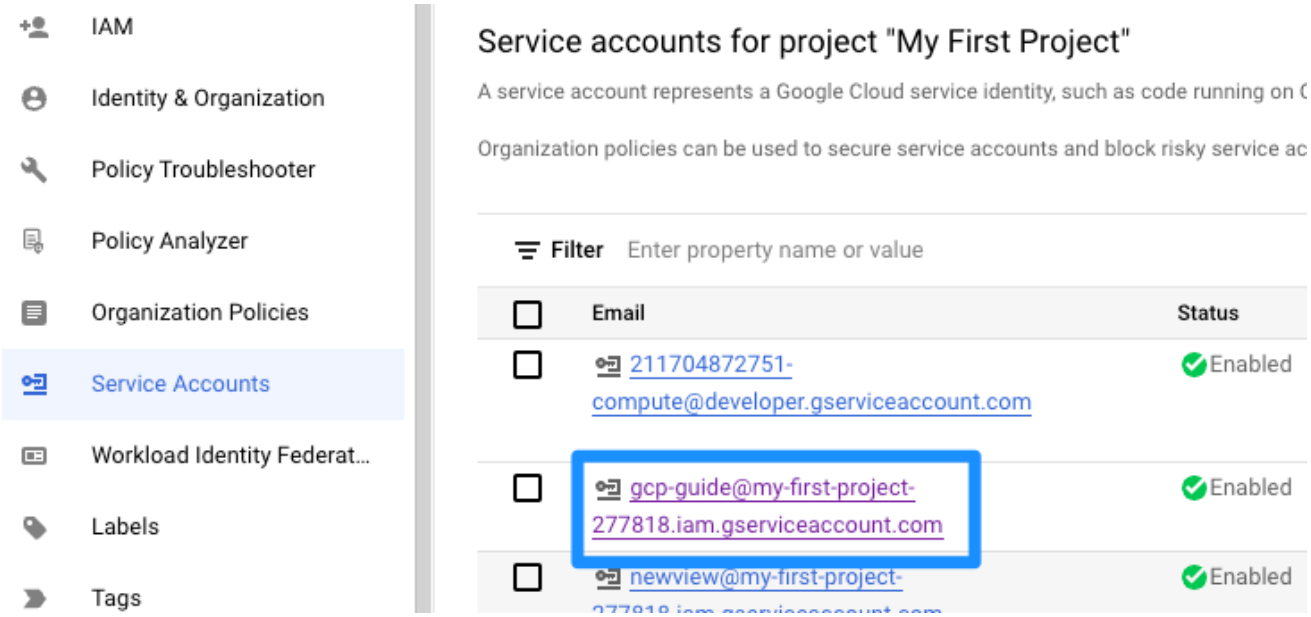

12. On the **Keys** tab, Click **ADD KEY** → **Create new key**.

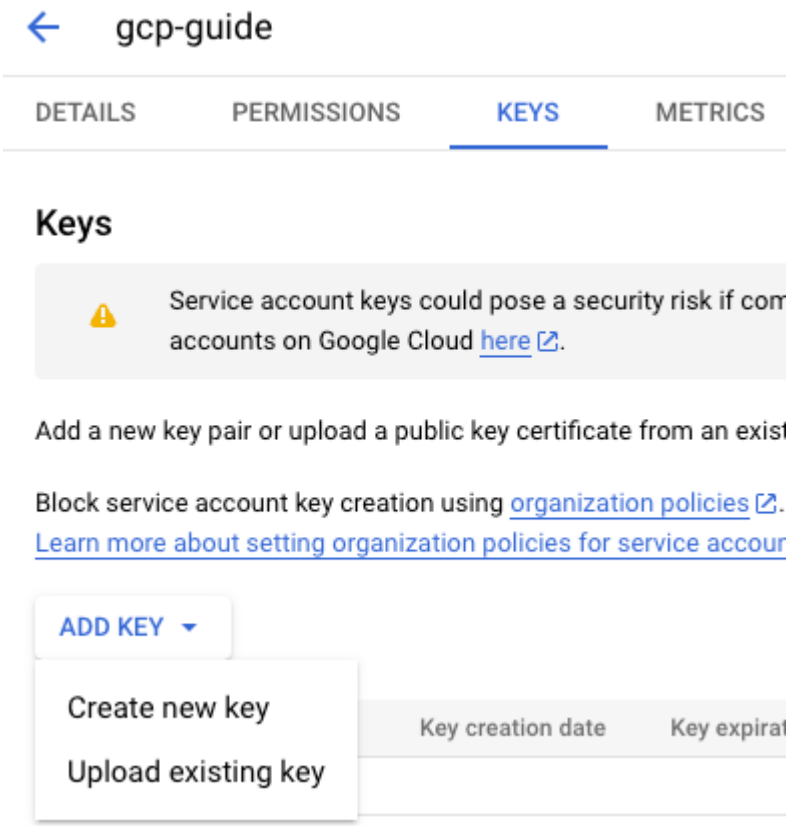

13. Select **JSON** and click **CREATE**.

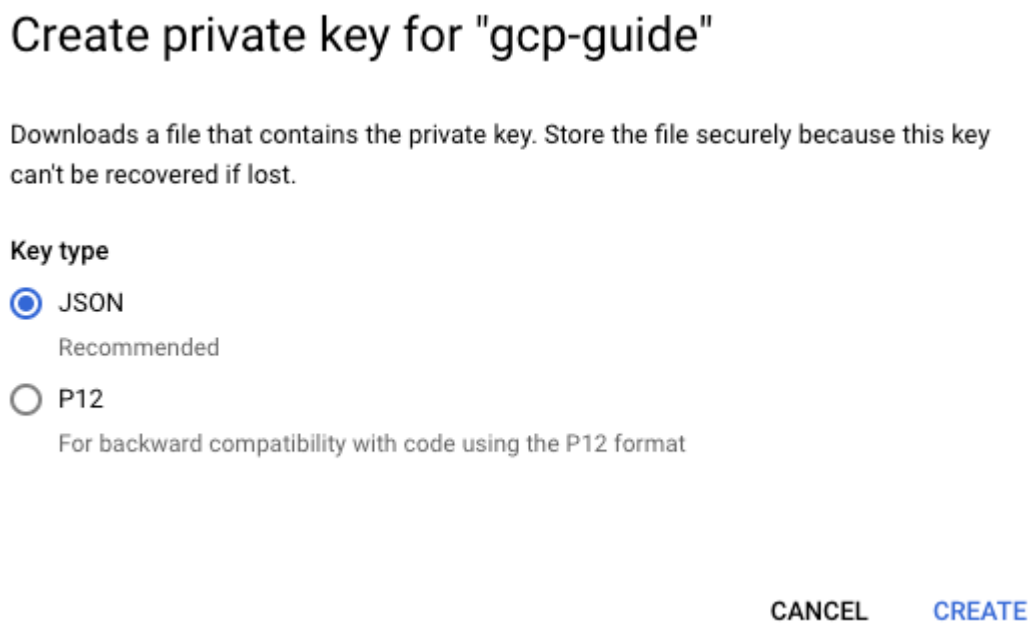

14. If prompted, complete any steps to save the resulting file.

15. Click **Close**.

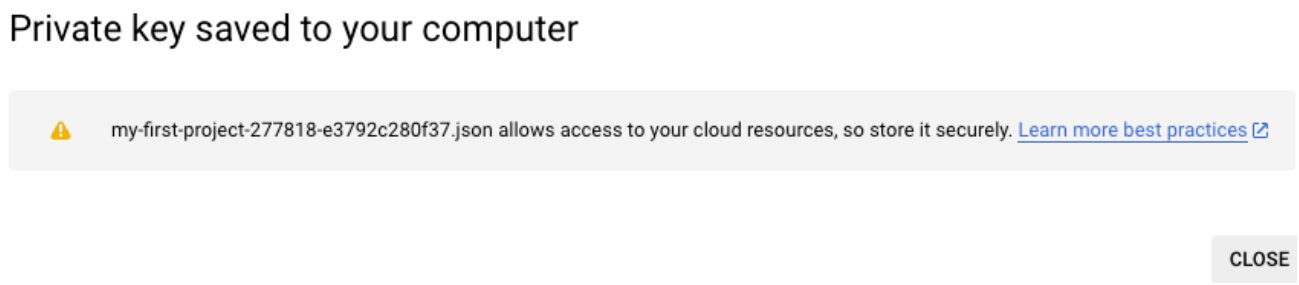

Infoblox vDiscovery Task

Infoblox vDiscovery can work with or without the Cloud Network Automation (CNA) license. CNA provides enhanced visibility for your cloud resources, greatly extending your searching, reporting and monitoring capabilities. When deployed without CNA, vDiscovery will help you keep your IPAM data up to date.

Create a vDiscovery Task

1. Login to the Infoblox Grid Manager GUI.

2. Switch to the Cloud or **Data Management** → **IPAM** tab.

3. Expand the **vDiscovery** menu and select **Discovery Manager**.

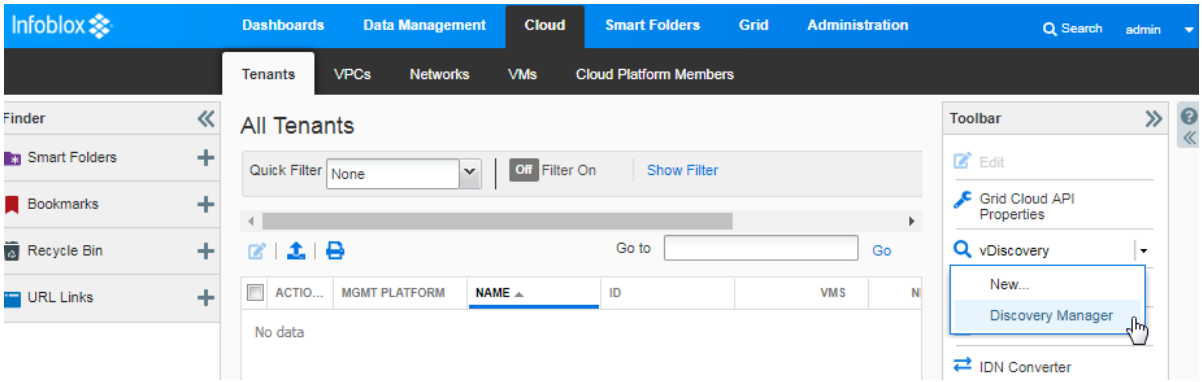

4. Click on the **+** (**Add**) button.

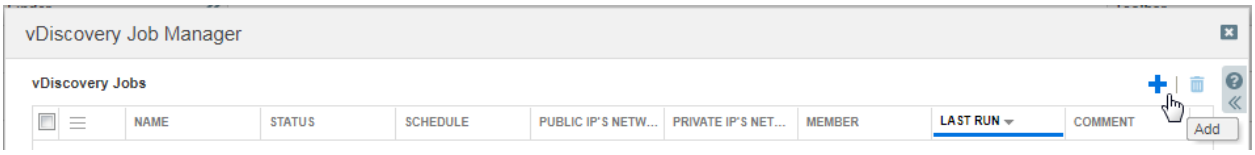

5. Enter a descriptive name.

6. Click **Select** to assign a Grid member to the vDiscovery task.

7. Click **Next**.

8. In the **Server Type** menu, select **GCP**.

9. Click on the **Upload** button.

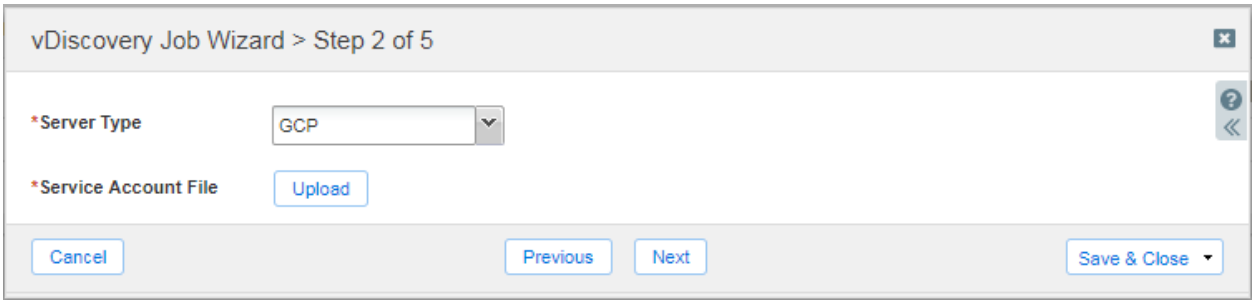

Note: The file name for the service account file must be unique.

10. Click **Select** and choose your service account key file.

11. Click **Upload**.

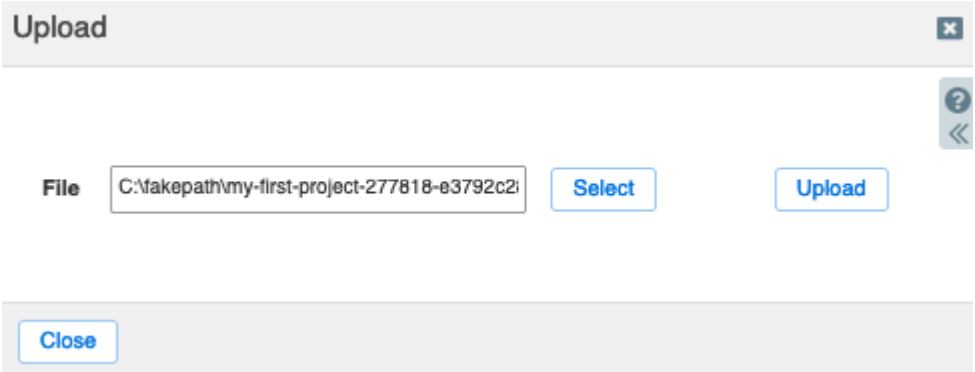

12. Click **Next**.

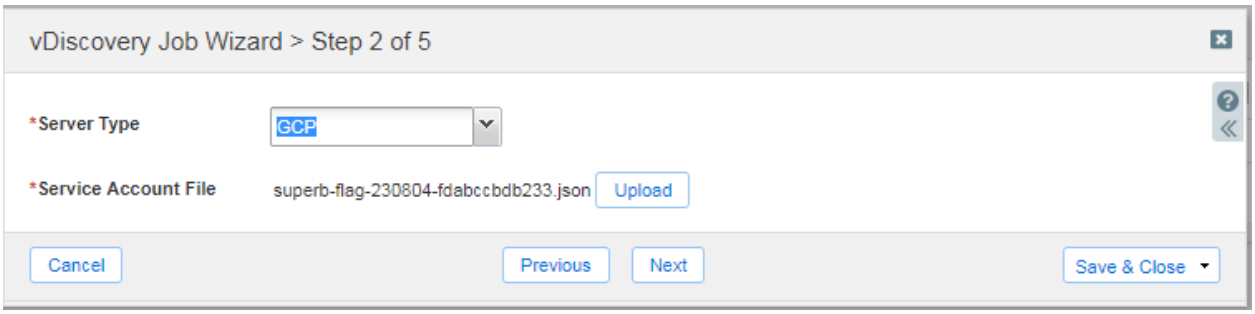

13. Review the configuration available for Network Views.

*Note: The most common cause for vDiscovery to fail to import any data is a "Sync Error" due to overlapping/conflicting address space. To account for any address space conflicts that are encountered during the vDiscovery process or with your existing IPAM data, you may need to select the option to use "**The tenant's network view (if it does not exist, create a new one)**".*

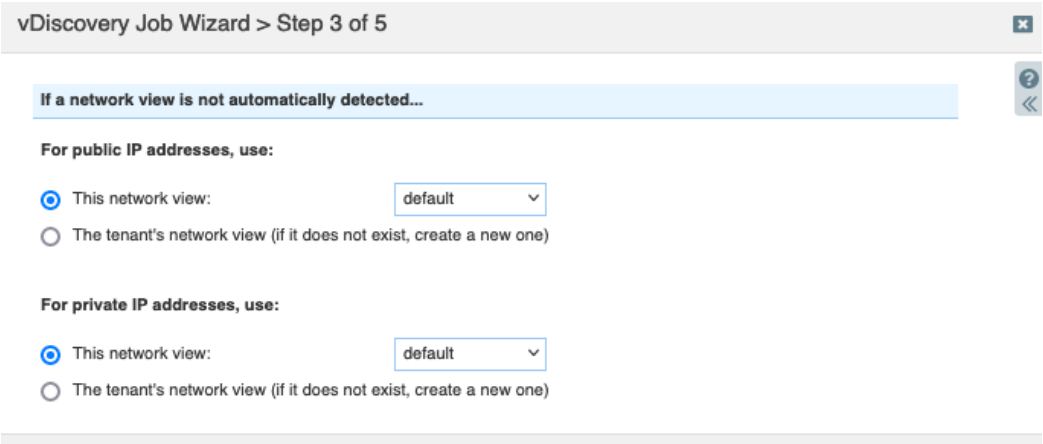

14. Optional: NIOS 8.6.3 and later add the ability to limit which virtual networks are discovered using Include or Exclude filters. To filter which networks are discovered:

    a. Select **Enable Filter**.

    b. Use the Network dropdown to select **INCLUDE** or **EXCLUDE**. If you select include, only networks that you list in the filter are discovered. If you select exclude, all networks except those listed in the filter are discovered.

    c. Click the **+**.

    d. Enter a network in CIDR format, for example **172.16.0.0/16**.

    e. Click the **+** again to add additional CIDRs.

15. Click **Next**.

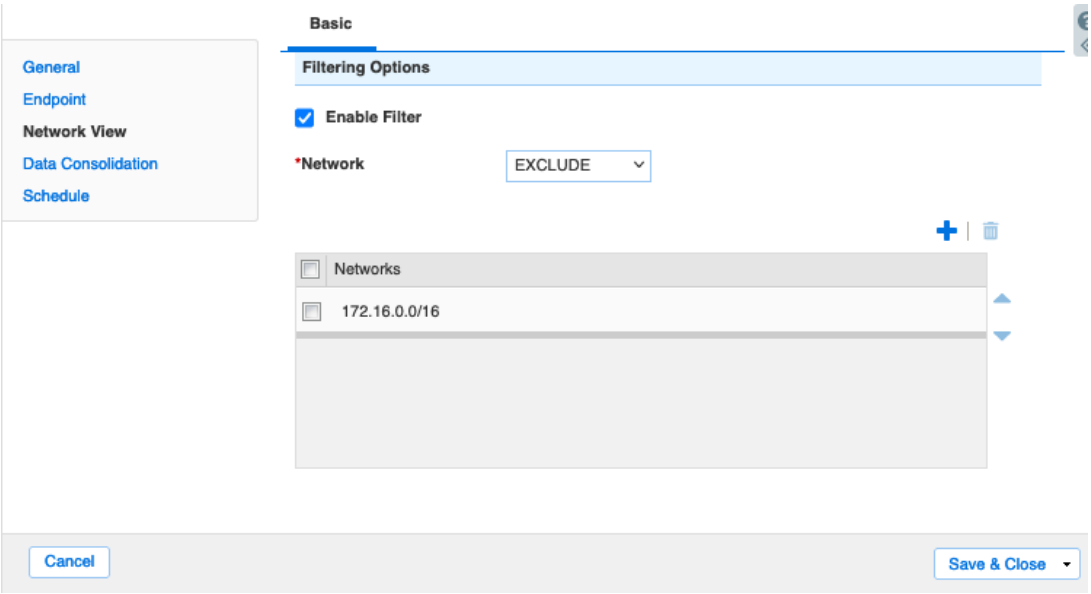

16. Review the options for handling discovered data. For detailed information on each of the options, refer to NIOS Documentation.

17. For automatic creation of DNS records, enable the option "**For every newly discovered IP address, create:**".

    a. Select the desired DNS record object type. If in doubt, stick with the default (Host) option. For zones integrated with the Microsoft Management feature, use the A & PTR Record option.

    b. The name for DNS records that are created is controlled with a macro, with the most commonly used macro being **${vm_name)**. In the text box, type the desired macro, followed by the zone that you want to use. Example: **${vm_name}.mycompany.com**.

The zone must be created separately from the vDiscovery task, though this can be done after the vDiscovery has already been created. If vDiscovery runs before the zone is created, any discovered objects will be marked as 'unmanaged' until the zone is created, and it runs again.

If a different format is desired for the DNS record name, a full list of available macros can be found in the Help panel. To view this, click on the question mark at the top-right hand corner of the window and scroll down to the section titled "**The DNS name will be computed from the formula**".

18. Click **Next**.

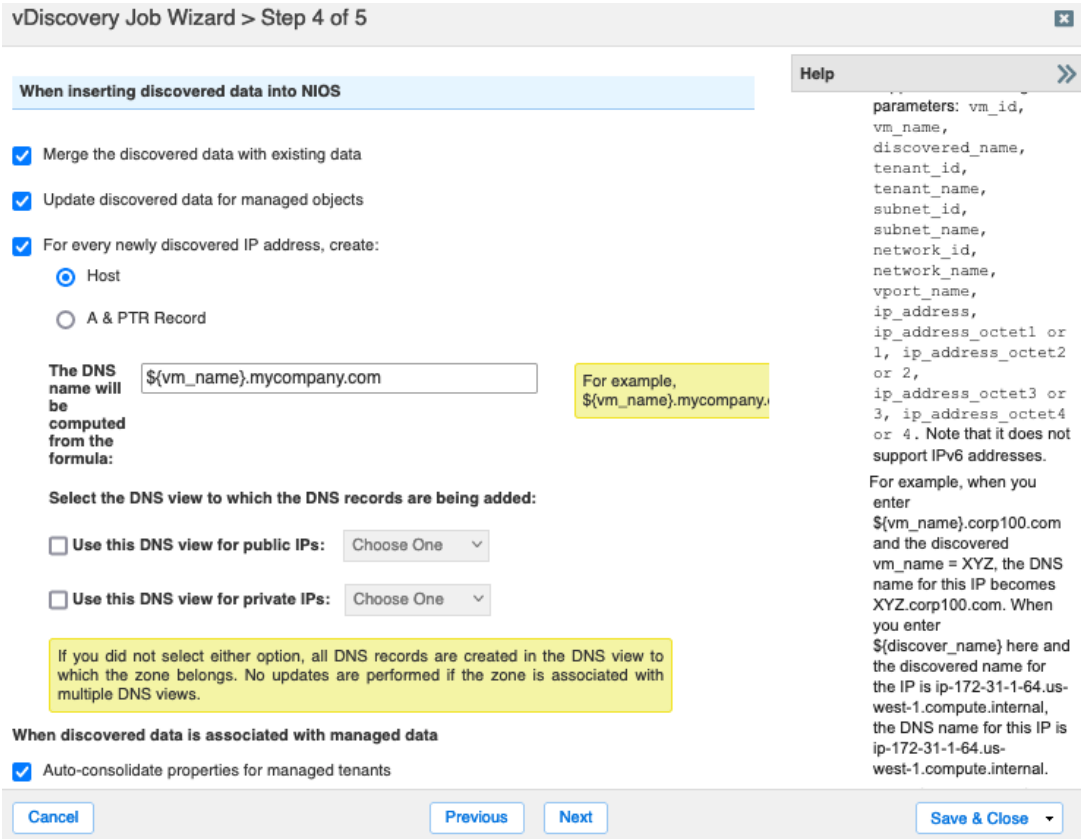

19. Optional: Configure a schedule to automatically run the vDiscovery task.

Note: The scheduler enables you to run the vDiscovery task as frequently as once an hour. If this must be run more frequently, this can be accomplished using the API. Refer to the Infoblox REST API guide for examples and guidelines on this process.

20. Click **Save & Close**.

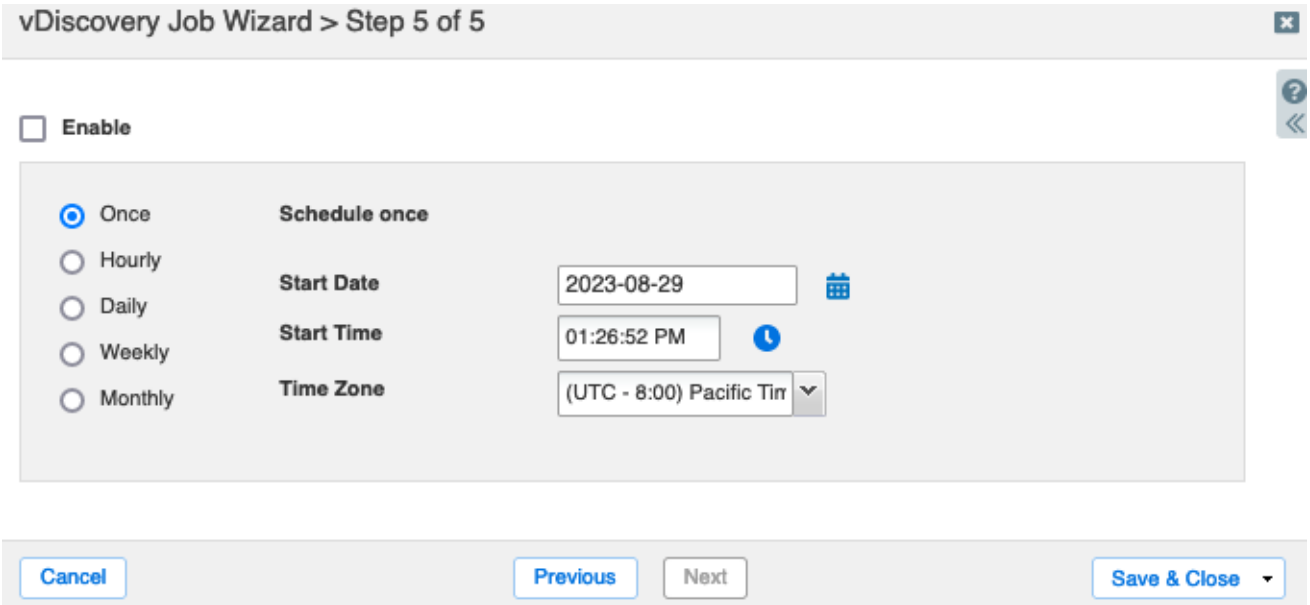

Run the vDiscovery Task

To manually start the vDiscovery task:

1. In the vDiscovery Job Manager, click on the gear wheel and select **Start**.

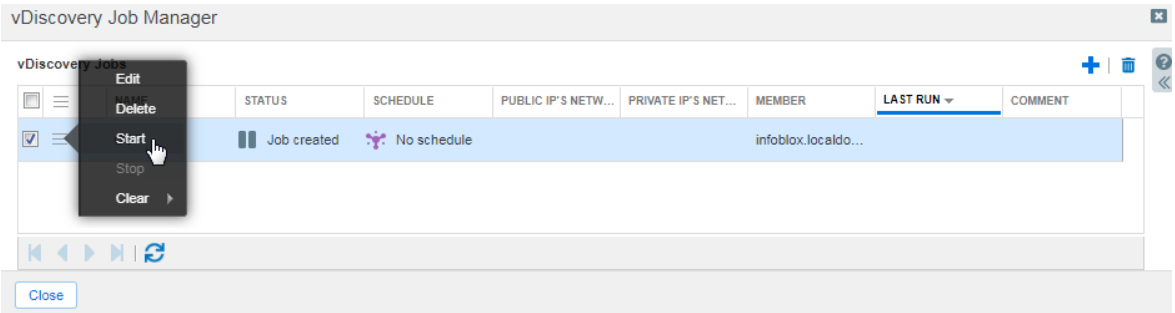

2. Click **Yes** to start the vDiscovery job.

3. Click the **Refresh** button at the bottom of the window until the Status shows **Job completed**.

Note: The status may show the vDiscovery task completed but with warnings. This can happen if objects are skipped, including if the name for a VM is in an invalid format (vDiscovery does not support dotted VM names), for any instances that have been terminated, or if the zone configured in the vDiscovery task cannot be found.

4. Click **Close**.

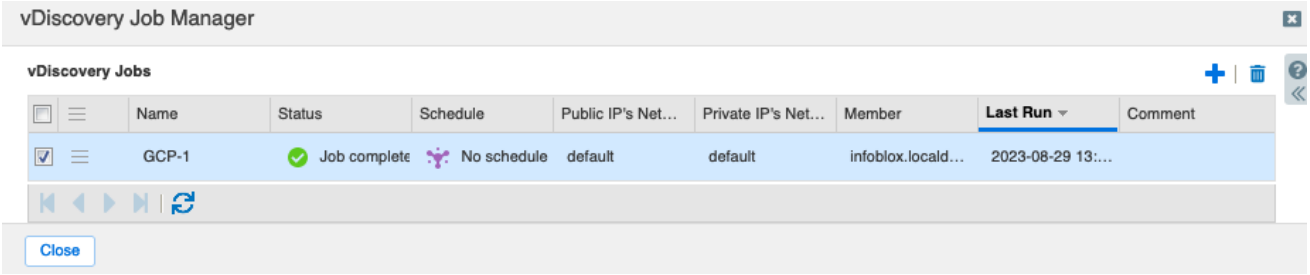

vDiscovery Data

Data collected by vDiscovery can be tracked through Data Management (IPAM, DHCP and DNS) and if the CNA license is installed, additional details will be found under the Cloud tab. Objects created by vDiscovery will automatically include metadata in their properties or extensible attributes (EA's), a useful addition that enables you to easily identify, locate and report on your resources deployed in the cloud.

Cloud Network Automation

When the CNA license is installed, you will find the Cloud tab in your Grid Manager GUI. With the Cloud tab come four additional tabs and each of these provide different perspectives for viewing your cloud data, making it easy to see what is running in your cloud environment based on different parameters.

These tabs include:

- **Tenants**: A global overview of all data through a single discovery source. This may correspond to an individual vDiscovery task or plugin/adapter. You can drill down to review all subnets and VMs that have been discovered under that tenant.

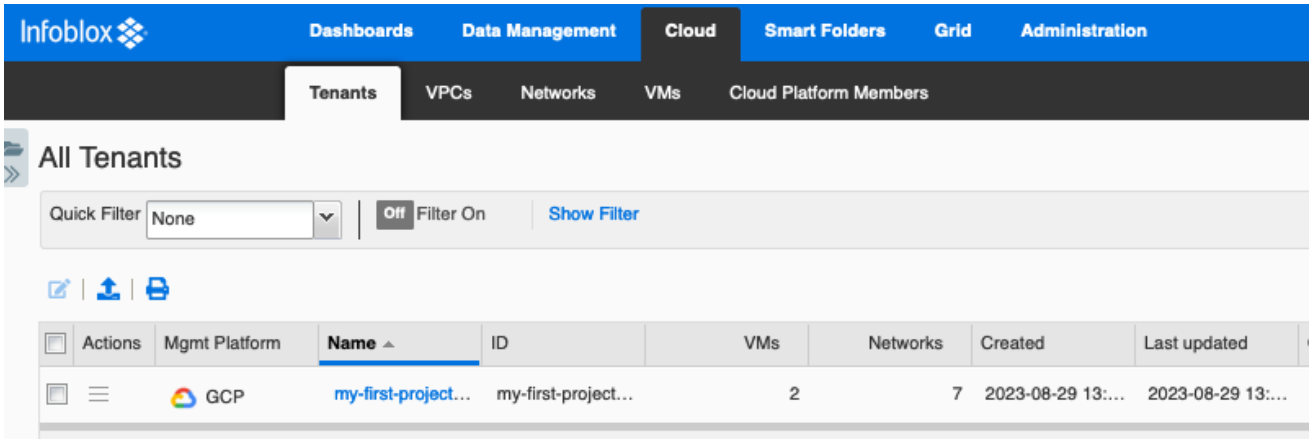

- **VPCs**: This is not used for GCP but will display any discovered AWS VPCs and Azure vNets. You can drill down to review all subnets and VMs that have been discovered under an individual VPC/vNet.
- **Networks**: A global overview of all subnets that have been discovered. Easily jump to IPAM or other perspectives to view additional details for a subnet. Searches, Smart Folders and reports can also leverage the metadata stored as EAs for each subnet.

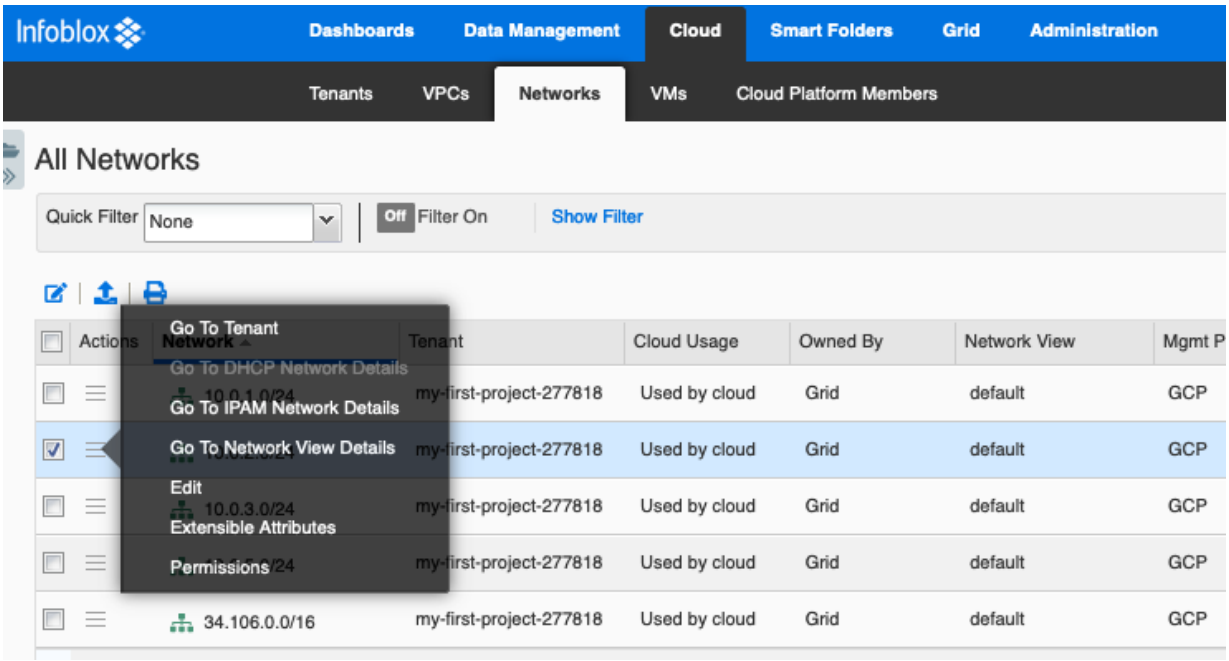

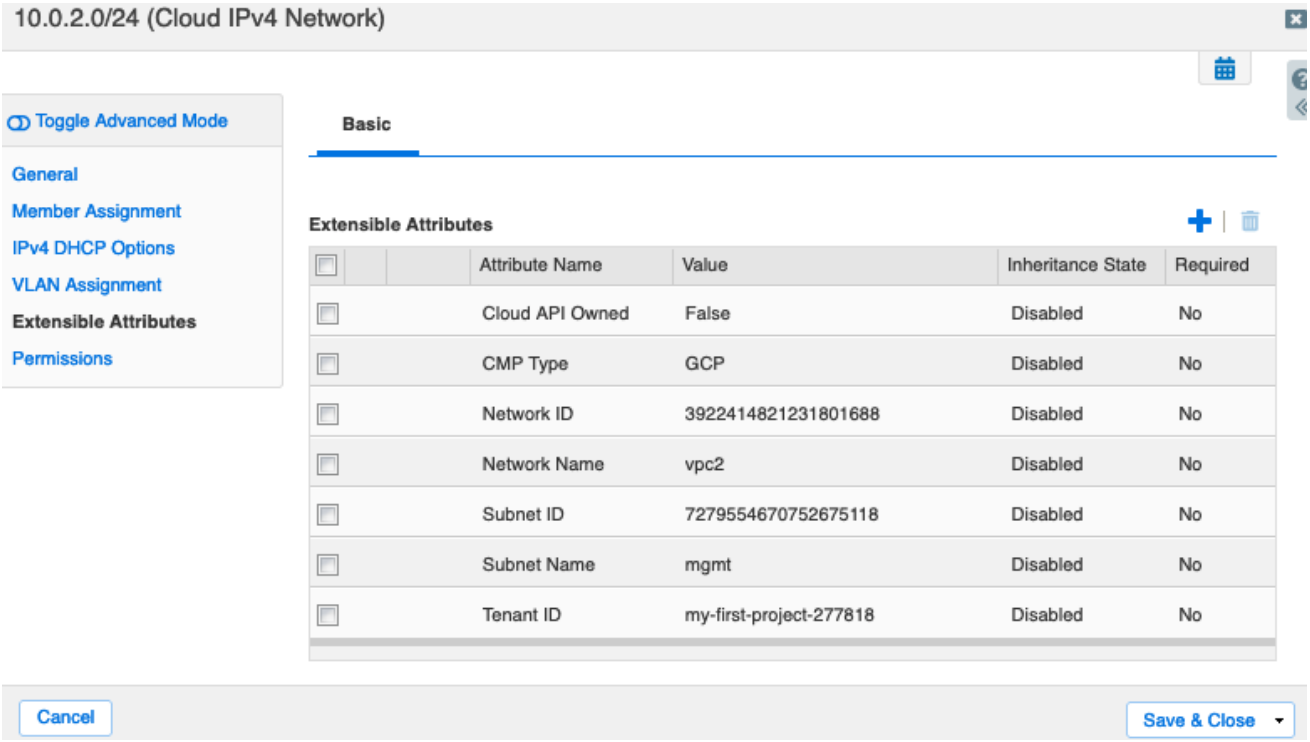

- **VMs**: A global overview of all virtual machines that have been discovered and displayed per IP address. Metadata is stored in the properties for each VM, and you can readily jump to other perspectives to view and manage additional resources, including any DNS records that may have been created for the VM.

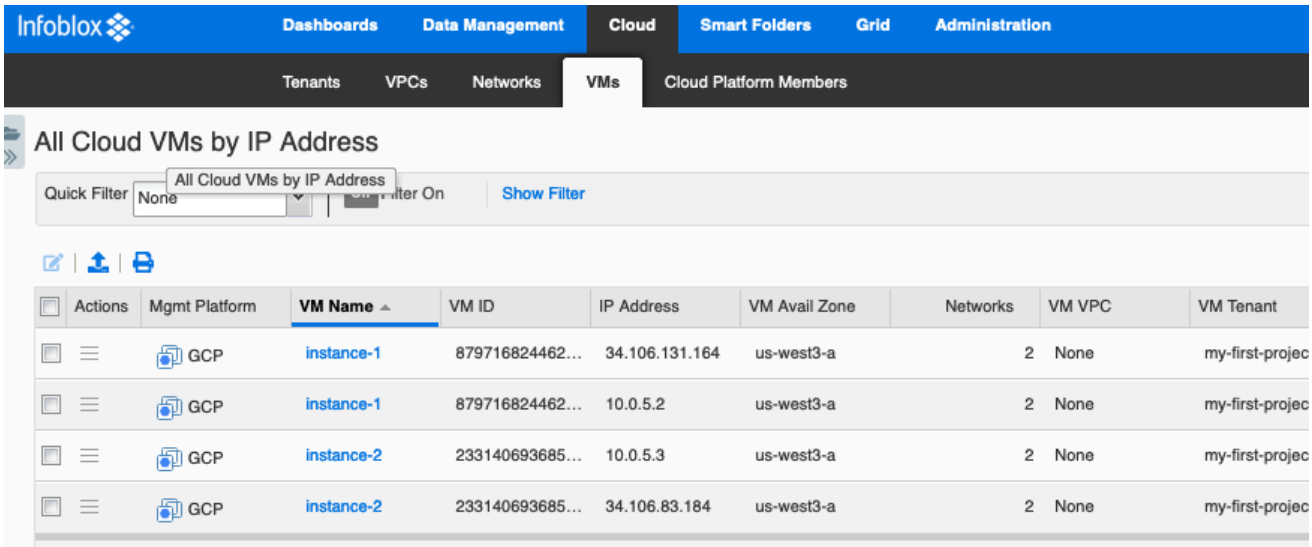

IPAM

IPAM, or IP Address Management, provides an easy view of all data from an IP address perspective. If you are looking for an object based on its IP address, this can be one of the easiest ways to drill down and see everything there is for that IP, including all objects that are associated with it.

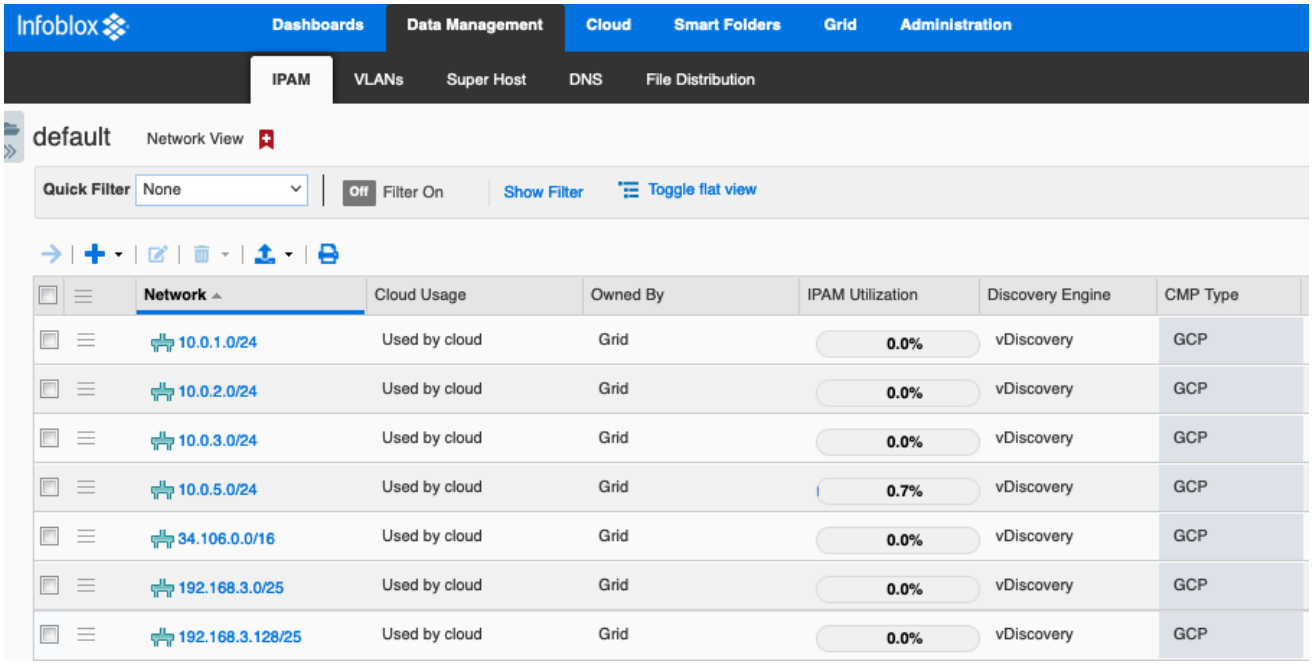

Infoblox unites networking and security to deliver unmatched performance and protection. Trusted by Fortune 100 companies and emerging innovators, we provide real-time visibility and control over who and what connects to your network, so your organization runs faster and stops threats earlier.

Corporate Headquarters
2390 Mission College Blvd, Ste. 501
Santa Clara, CA 95054
+1.408.986.4000
www.infoblox.com