

DEPLOYMENT GUIDE

INFOBLOX vDISCOVERY FOR GOOGLE CLOUD

TABLE OF CONTENT

- OVERVIEW 3**
 - Introduction.....3
 - Prerequisites3
 - Basic Workflow.....3
- ENABLING GOOGLE CLOUD FOR vDISCOVERY 3**
 - Service Account3
- INFOBLOX vDISCOVERY TASK7**
 - Create a vDiscovery Task.....7
 - Run the vDiscovery Task.....11
- vDISCOVERY DATA..... 11**
 - Cloud Network Automation11
 - IPAM..... 13

OVERVIEW

INTRODUCTION

Infoblox vDiscovery provides enhanced visibility of your networks and virtual machines (VMs), 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, bringing all this data under a single pane of glass. In this guide, you will be introduced to Infoblox vDiscovery for Google Cloud.

PREREQUISITES

The following are prerequisites for Infoblox vDiscovery with Google Cloud:

- Valid subscription and login to Google Cloud.
- 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 Google Cloud:

- Sign in to the **Google Cloud 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 GOOGLE CLOUD FOR vDISCOVERY

SERVICE ACCOUNT


To enable the connection from vDiscovery to Google Cloud and for it to work properly, you must use a service account with the appropriate permissions assigned to it. In Google Cloud, 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 Google Cloud 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**.

1 Service account details

Service account name
gcp-guide

Display name for this service account

Service account ID *
gcp-guide

Email address: gcp-guide@my-first-project-277818.iam.gserviceaccount.com

Service account description

Describe what this service account will do

CREATE AND CONTINUE

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

2 Grant this service account access to project (optional)

Grant this service account access to My First Project so that it has permission to complete specific actions on the resources in your project. [Learn more](#)

Select a role 

IAM condition (optional)  [+ ADD IAM CONDITION](#) 

[+ ADD ANOTHER ROLE](#)

[CONTINUE](#)

- Select **Project** → **Viewer**.
- Click **CONTINUE**.

Type to filter

- Project
- Access Approval
- Android Management
- App Engine
- AutoML
- BigQuery
- Billing
- Binary Authorization

[MANAGE ROLES](#)

[+ ADD ANOTHER ROLE](#)

Viewer
Read access to all resources.

[CONTINUE](#) [CANCEL](#)

- Click **DONE**.

✓ Service account details

✓ Grant this service account access to project (optional)

3 Grant users access to this service account (optional)

Grant access to users or groups that need to perform actions as this service account. [Learn more](#)

Service account users role 

Grant users the permissions to deploy jobs and VMs with this service account

Service account admins role 

Grant users the permission to administer this service account

[DONE](#) [CANCEL](#)

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

Service accounts for project "My First Project"

A service account represents a Google Cloud service identity, such as code running on Cloud. Organization policies can be used to secure service accounts and block risky service accounts.

Filter Enter property name or value

☐	Email	Status
☐	211704872751-compute@developer.gserviceaccount.com	Enabled
☐	gcp-guide@my-first-project-277818.iam.gserviceaccount.com	Enabled
☐	newview@my-first-project-277818.iam.gserviceaccount.com	Enabled

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

gcp-guide

DETAILS **PERMISSIONS** **KEYS** **METRICS**

Keys

Service account keys could pose a security risk if compromised. [Learn more about setting organization policies for service accounts on Google Cloud here](#).

Add a new key pair or upload a public key certificate from an existing key.

Block service account key creation using [organization policies](#). [Learn more about setting organization policies for service accounts](#).

ADD KEY ▾

Create new key

Upload existing key

Key creation date	Key expiration date

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

Create private key for "gcp-guide"

Downloads a file that contains the private key. Store the file securely because this key can't be recovered if lost.

Key type

☒ JSON

Recommended

☐ P12

For backward compatibility with code using the P12 format

CANCEL CREATE

14. If prompted, complete any steps to save the resulting file.
15. Click **CLOSE**.

Private key saved to your computer

 my-first-project-277818-e3792c280f37.json allows access to your cloud resources, so store it securely. [Learn more best practices](#)


CLOSE

INFOBLOX vDISCOVERY TASK


Infoblox vDiscovery can work with or without the 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 IP address management (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**.


 The screenshot shows the 'vDiscovery Job Wizard > Step 1 of 5' form. It has a title bar and a close button (X). The form contains three fields:

- *Job Name: A text input field with 'GCP-1' entered.
- *Member: A dropdown menu showing 'infoblox.localdomain' with 'Select' and 'Clear' buttons next to it.
- Comment: A large text area for additional information.

 At the bottom, there are four buttons: 'Cancel', 'Previous', 'Next', and 'Save & Close' (with a dropdown arrow).

8. In the **Server Type** menu, select **Google Cloud**.
9. Click on the **Upload** button.

vDiscovery Job Wizard > Step 2 of 5

*Server Type: GCP

*Service Account File: Upload

Cancel Previous Next Save & Close

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**.

Upload

File: C:\fakepath\my-first-project-277818-e3792c2 Select Upload

Close

12. Click **Next**.

vDiscovery Job Wizard > Step 2 of 5

*Server Type: GCP

*Service Account File: superb-flag-230804-fdabccbdb233.json Upload

Cancel Previous Next Save & Close

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)".

vDiscovery Job Wizard > Step 3 of 5

If a network view is not automatically detected...

For public IP addresses, use:

☒ This network view: default

☐ The tenant's network view (if it does not exist, create a new one)

For private IP addresses, use:

☒ This network view: default

☐ 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**.

The screenshot shows the NIOS vDiscovery configuration window. On the left is a sidebar with tabs: General, Endpoint, Network View (selected), Data Consolidation, and Schedule. The main area is titled 'Basic' and contains a 'Filtering Options' section. In this section, 'Enable Filter' is checked. Below it, a dropdown menu labeled '*Network' is set to 'EXCLUDE'. Underneath is a table with a header 'Networks' and one entry: '172.16.0.0/16'. To the right of the table are a plus sign (+) and a trash icon. At the bottom of the window are 'Cancel' and 'Save & Close' buttons.

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**.

vDiscovery Job Wizard > Step 4 of 5

When inserting discovered data into NIOS

- ☒ Merge the discovered data with existing data
- ☒ Update discovered data for managed objects
- ☒ For every newly discovered IP address, create:
 - ☒ Host
 - ☐ A & PTR Record

The DNS name will be computed from the formula: For example, .

Select the DNS view to which the DNS records are being added:

- ☐ Use this DNS view for public IPs:
- ☐ Use this DNS view for private IPs:

If you did not select either option, all DNS records are created in the DNS view to which the zone belongs. No updates are performed if the zone is associated with multiple DNS views.

When discovered data is associated with managed data

- ☒ Auto-consolidate properties for managed tenants

parameters: vm_id, vm_name, discovered_name, tenant_id, tenant_name, subnet_id, subnet_name, network_id, network_name, vport_name, ip_address, ip_address_octet1 or 1, ip_address_octet2 or 2, ip_address_octet3 or 3, ip_address_octet4 or 4. Note that it does not support IPv6 addresses. For example, when you enter and the discovered vm_name = XYZ, the DNS name for this IP becomes XYZ.corp100.com. When you enter here and the discovered name for the IP is ip-172-31-1-64.us-west-1.compute.internal, the DNS name for this IP is ip-172-31-1-64.us-west-1.compute.internal.

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**.

vDiscovery Job Wizard > Step 5 of 5

☐ Enable

☒ Once **Schedule once**

☐ Hourly

☐ Daily

☐ Weekly

☐ Monthly

Start Date


Start Time

Time Zone

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 (EAs), 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/adaptor. You can drill down to review all subnets and VMs that have been discovered under that tenant.

Actions	Mgmt Platform	Name	ID	VMs	Networks	Created	Last updated
	GCP	my-first-project...	my-first-project...	2	7	2023-08-29 13:00	2023-08-29 13:00

- **VPCs:** This is not used for Google Cloud 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.

Actions	Tenant	Cloud Usage	Owned By	Network View	Mgmt P
	my-first-project-277818	Used by cloud	Grid	default	GCP
	my-first-project-277818	Used by cloud	Grid	default	GCP
	my-first-project-277818	Used by cloud	Grid	default	GCP
	my-first-project-277818	Used by cloud	Grid	default	GCP
	my-first-project-277818	Used by cloud	Grid	default	GCP

Actions	Mgmt Platform	VM Name	VM ID	IP Address	VM Avail Zone	Networks	VM VPC	VM Tenant
	GCP	instance-1	879716824462...	34.106.131.164	us-west3-a	2	None	my-first-proj
	GCP	instance-1	879716824462...	10.0.5.2	us-west3-a	2	None	my-first-proj
	GCP	instance-2	233140693685...	10.0.5.3	us-west3-a	2	None	my-first-proj
	GCP	instance-2	233140693685...	34.106.83.184	us-west3-a	2	None	my-first-proj

- **VMs:** A global overview of all VMs 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.

infoblox Dashboards Data Management Cloud Smart Folders Grid Administration

Tenants VPCs Networks VMs Cloud Platform Members

All Cloud VMs by IP Address

Quick Filter: None Filter On Show Filter

Actions	Mgmt Platform	VM Name	VM ID	IP Address	VM Avail Zone	Networks	VM VPC	VM Tenant
[Icon]	GCP	instance-1	879716824462...	34.106.131.164	us-west3-a	2	None	my-first-projec
[Icon]	GCP	instance-1	879716824462...	10.0.5.2	us-west3-a	2	None	my-first-projec
[Icon]	GCP	instance-2	233140693685...	10.0.5.3	us-west3-a	2	None	my-first-projec
[Icon]	GCP	instance-2	233140693685...	34.106.83.184	us-west3-a	2	None	my-first-projec

IPAM

IPAM 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 Dashboards Data Management Cloud Smart Folders Grid Administration

IPAM VLANs Super Host DNS File Distribution

default Network View

Quick Filter: None Filter On Show Filter Toggle flat view

Network	Cloud Usage	Owned By	IPAM Utilization	Discovery Engine	CMP Type
10.0.1.0/24	Used by cloud	Grid	0.0%	vDiscovery	GCP
10.0.2.0/24	Used by cloud	Grid	0.0%	vDiscovery	GCP
10.0.3.0/24	Used by cloud	Grid	0.0%	vDiscovery	GCP
10.0.5.0/24	Used by cloud	Grid	0.7%	vDiscovery	GCP
34.106.0.0/16	Used by cloud	Grid	0.0%	vDiscovery	GCP
192.168.3.0/25	Used by cloud	Grid	0.0%	vDiscovery	GCP
192.168.3.128/25	Used by cloud	Grid	0.0%	vDiscovery	GCP



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