

NetWitness[®] Platform

Version 12.5

Google Cloud Platform Installation Guide

Contact Information

NetWitness Community at <https://community.netwitness.com> contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to <https://www.rsa.com/en-us/company/rsa-trademarks>. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability. This software is subject to change without notice and should not be construed as a commitment by RSA.

It is advised not to deploy third-party repos or perform any change to the underlying NetWitness Operating System that is not part of the supported NetWitness version. Any such change outside of the NetWitness approved image may result in a service or functionality conflict and require a reimage of the NetWitness system to bring NetWitness back to an optimized functional state. In the event a third-party repo is deployed, or other non-supported change is made by the customer without NetWitness approval, the customer takes full responsibility for any system malfunction until the issue can be remediated through troubleshooting efforts or a reimage of the service.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Miscellaneous

This product, this software, the associated documentations as well as the contents are subject to NetWitness' standard Terms and Conditions in effect as of the issuance date of this documentation and which can be found at <https://www.netwitness.com/standard-form-agreements/>.

© 2024 RSA Security LLC or its affiliates. All Rights Reserved.

September, 2024

Contents

Google Cloud Platform Installation Overview	4
GCP Deployment Scenarios	5
NetWitness Full Stack VPC Visibility	5
Hybrid Deployment	5
GCP Deployment	6
Checklist	6
Prerequisites	6
Find NetWitness Platform GCP Images	6
Establish gcloud Environment	7
Create an Instance using Google Cloud SDK Shell	7
Create a Firewall Rule	8
Connect to VM Instance using SSH	10
Method 1: SSH - Open in Browser Window	10
Method 2: SSH - Cloud Shell	11
Installation Tasks	11
Tasks - Install 12.5 on the NetWitness Server (NW Server) Host and Component Hosts	12
Storage Configurations	18
Create a Disk	18
Configure Hosts (Instances) in NetWitness Platform	21
Configure Packet Mirroring	21
Task 1 - Create Instance Group	21
Task 2 - Create Packet Mirroring Policy	22
Task 3 - Create Load Balancer	27
Task 4 - Verify Packet Mirroring in GCP	31
GCP Instance Configuration Recommendations	32
Packet Concentrator	35

Google Cloud Platform Installation Overview

Google Cloud Platform (GCP) instances have the same functionality as the NetWitness hardware and virtual hosts. NetWitness recommends that you perform the following tasks when you set up your GCP environment.

Before you can deploy NetWitness in GCP, you need to:


- Review the recommended compute and memory specifications needed for each NetWitness instance.
- Get familiar with the NetWitness Storage Guide to understand the types of drives and volumes needed to support NetWitness instances. For more information, see [Storage Guide for NetWitness® Platform 12.5.0.0](#).
- Make sure that you have a NetWitness Throughput license.

GCP Deployment Scenarios

The following diagrams illustrate some common GCP deployment scenarios.


NetWitness Full Stack VPC Visibility

This diagram shows all NetWitness components (full stack) deployed in GCP.



Hybrid Deployment

This diagram shows the Log Decoder and Concentrator deployed in GCP with all other NetWitness components deployed on premises.



GCP Deployment

This topic contains the rules and high-level tasks you must follow to deploy NetWitness components in the GCP.

Checklist

Step	Description
1	Prerequisites
2	Find NetWitness Platform GCP Images
3	Establish gcloud Environment
4	Create an Instance using Google Cloud SDK Shell
5	Create a Firewall Rule
6	Connect to VM Instance using SSH
7	Installation Tasks
8	Configure Hosts (Instances) in NetWitness Platform
9	Configure Packet Mirroring

Prerequisites

You need the following items before you begin the integration process:

- Access to GCP console.
- Google Cloud SDK Toolkit installed.
- Network routability (and proper GCP firewall rules) for the instances to transfer data to the NetWitness.

Find NetWitness Platform GCP Images

There are two types of NetWitness GCP images. The following description guides you on which image is appropriate for your deployment:

- **Lite Image:** `rsa-nw-12-5-0-0-21645-lite`
If you have an active NW Server or NetWitness software repository, use the Lite image. This image has a small footprint and does not contain the NetWitness software packages. When you set up NetWitness using the Lite image, it will require the IP address of the NW Server or NetWitness

software repository. This image is available publicly, and you can deploy it without needing to contact NetWitness Customer Support.

- **Full Image:** `rsa-nw-12-5-0-0-21645-full`

For fresh or new deployments, where an NW Server is deployed for the first time, use the Full image. This image contains the entire NetWitness software library and other files required to complete the installation.

To get access to the Full image, open an NetWitness Customer Support case

(<https://community.netwitness.com/t5/support-information/how-to-contact-netwitness-support/tap/563897>) to get granted the proper image permissions.

Note: For more information on GCP image deployment, see the blog post [Running NetWitness in Google Cloud](#).

Establish gcloud Environment

1. Download and install the Google Cloud SDK Toolkit (<https://cloud.google.com/sdk>).
2. Run the following gcloud commands to log in and set the proper project:

```
gcloud auth login
gcloud config set project <project name>
```

Create an Instance using Google Cloud SDK Shell

1. Determine what machine type is appropriate for selection. For more information, see [GCP Instance Configuration Recommendations](#).
2. Using the image names listed under [Find NetWitness Platform GCP Images](#), create a GCP Instance by running the command:

```
gcloud compute instances create <instance name> --image <netwitness image name> --image-project <rsa project> --machine-type <machine type> --zone <gcp zone> --network <gcp project network> --subnet <gcp project subnet> --no-address
```

For example:

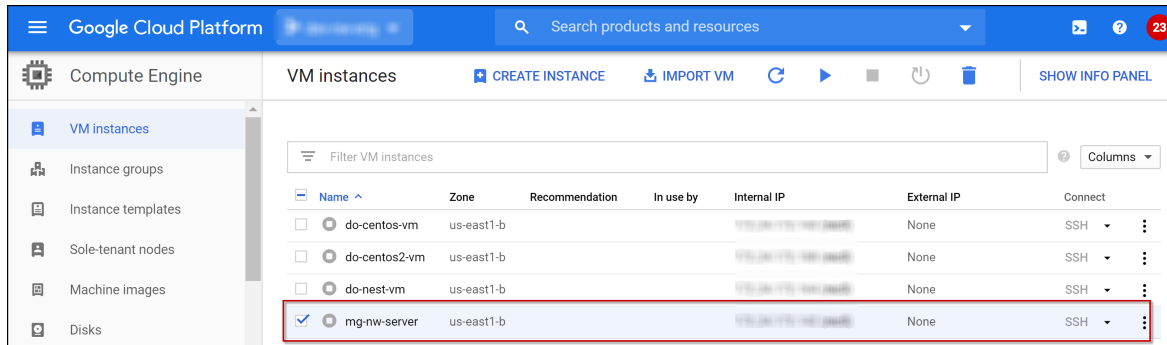
```
gcloud compute instances create nw-server --image rsa-nw-12-5-0-0-21645-full --image-project nw-onprem-images-prod --machine-type n1-standard-8 --zone us-west1-c --network rsa-network --subnet rsa-subnet --no-address
```

Note:

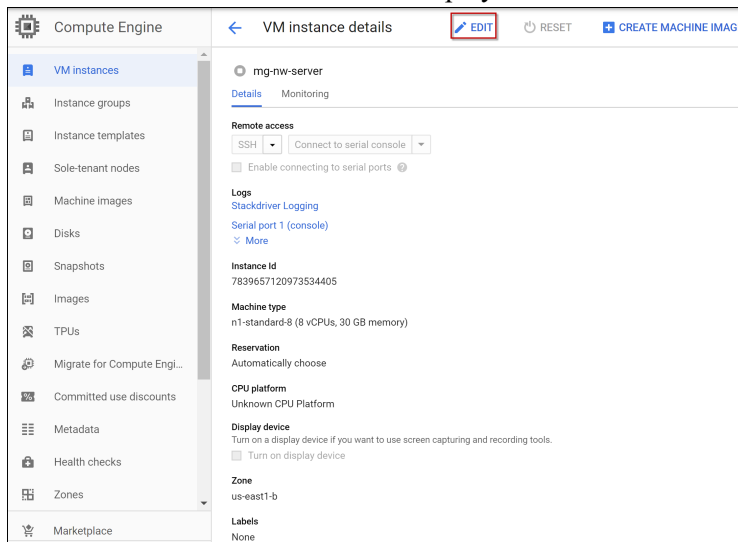
- The network and subnet values may vary based on the setup.
- The simple example is shown above, but there are many other options available. For more information, see the [Quickstart: Creating a New Instance Using the Command Line](#) section in the GCP documentation.

3. To modify the machine type, region, zone, or other configurations:

- a. Go to **Google Cloud Platform > Compute Engine > VM Instance** view to find the VM Instance.
- b. Select the instance and click **STOP**.



- c. Click the Instance name (**mg-nw-server**).
- d. Click **EDIT** to modify the settings according to your preference and click **Save**.
The **VM instance details** view is displayed.



- e. SSH to the newly-created instance using the default NetWitness credentials.

Create a Firewall Rule

Every VPC network has two implied firewall rules that permit outgoing connections and block incoming connections. These rules allow egress traffic to everywhere "Implied IPv4 allow egress rule" and block incoming traffic from everywhere "Implied IPv4 deny ingress rule". These rules are not shown in the Google Cloud console. Firewall rules that you create can override these implied rules.

To allow RDP and SSH access to all VM instances in your network, you must create a fire wall rule.

To create a firewall rule

1. Log in to the Google Cloud Console.
2. Click **Firewall rule**.
3. Click **Create firewall rule**.

The Create a firewall rule view is displayed.

The screenshot shows the 'Create a firewall rule' form in the Google Cloud Console. The form is titled 'Create a firewall rule' and is part of the 'VPC network' section. It includes the following fields and options:

- Name:** A text input field with a placeholder 'Name *' and a note 'Lowercase letters, numbers, hyphens allowed'.
- Description:** A text input field with a placeholder 'Description'.
- Logs:** A section with a note 'Turning on firewall logs can generate a large number of logs which can increase costs in Cloud Logging. Learn more'. It has two radio buttons: 'On' and 'Off' (selected).
- Network:** A dropdown menu with a red border and a red error message 'Network is required'.
- Priority:** A text input field with a placeholder '1000' and a note 'Priority can be 0-65536'.
- Direction of traffic:** Two radio buttons: 'Ingress' (selected) and 'Egress'.
- Action on match:** Two radio buttons: 'Allow' (selected) and 'Deny'.
- Targets:** A dropdown menu with the selected option 'All instances in the network'.
- Source filter:** A dropdown menu with the selected option 'IPv4 ranges'.
- Source IPv4 ranges:** A text input field with a placeholder 'Source IPv4 ranges *'.
- Second source filter:** A dropdown menu with the selected option 'None'.
- Protocols and ports:** Two radio buttons: 'Allow all' and 'Specified protocols and ports' (selected). Under 'Specified protocols and ports', there are three checkboxes: 'TCP' (checked), 'UDP', and 'Other'. Below these is a text input field for 'Ports' with a placeholder 'Ports' and a note 'E.g. 20, 50-60'. Below that is another text input field for 'Protocols' with a note 'Separate multiple protocols by commas, e.g. ah, tcp'.

At the bottom of the form, there is a 'DISABLE RULE' link, a 'CREATE' button, a 'CANCEL' button, and an 'EQUIVALENT COMMAND LINE' dropdown menu.

4. Fill in the details to configure a firewall rule:
 - **Name:** Enter a name for firewall rule.
 - **Description:** Enter the description for firewall rule.
 - **Logs:** By default, **Off** is selected.
 - **Network:** Select a network from the drop-down list.
 - **Priority:** Enter the priority range. Lower integers indicate higher priorities.

Note: Priority range can be between 0 to 65536.

- **Direction of traffic:** Select **Ingress**.
- **Action on match:** Select **Allow**.
- **Targets:** Select **All instances in the network** from the drop-down list.

- **Source filter:** Select **IPv4 ranges** from the drop-down list.
 - **Source IPv4 ranges:** Enter **35.235.240.0/20**.
 - **Second source filter:** Enter any secondary source filter if available.
 - **Protocols and ports:** Select **specified protocols and ports** and select **TCP** and enter 22,3389 to allow both RDP and SSH.
5. Click **Create**.
- Source IPv4 **35.235.240.0/20** range contains all IP addresses that IAP uses for TCP forwarding in GCP.

Connect to VM Instance using SSH

There are different ways to connect to the VM instance using SSH.

1. [Method 1: SSH - Open in Browser Window](#)
2. [Method 2: SSH - Cloud Shell](#)


Method 1: SSH - Open in Browser Window

Perform the following steps to access SSH - Open in Browser Window.

1. Log in to the Google Cloud Console.
2. Search **Identity-Aware Proxy** in the console.

Note: To enable the Cloud Identity-Aware Proxy API, you must have **IAP-secured Tunnel User** and **Service Usage Admin** permissions.

3. Enable **Cloud Identity-Aware Proxy API**.



This will enable the **SSH** button in the row of newly created VM instance.

4. Go to the **VM instances** page.
5. In the row of an instance, select **Open in browser window** from the **SSH** drop-down menu.

<input type="checkbox"/>	Status	Name ↑	Zone	Recommendations	In use by	Internal IP	External IP	Network	Connect
<input type="checkbox"/>	✓	admin-server		Save \$419 / mo		(nic0)		shared-vpc-pontus	SSH
<input type="checkbox"/>	✓	concentrator		Save \$498 / mo					

Open in browser window

A new window opens and connects to the VM instance.

Method 2: SSH - Cloud Shell

Perform the following steps to access SSH - Cloud Shell.

1. Log in to the Google Cloud console page.
2. On the top right corner of the console, click **Activate Cloud Shell**.



The cloud shell page is displayed.

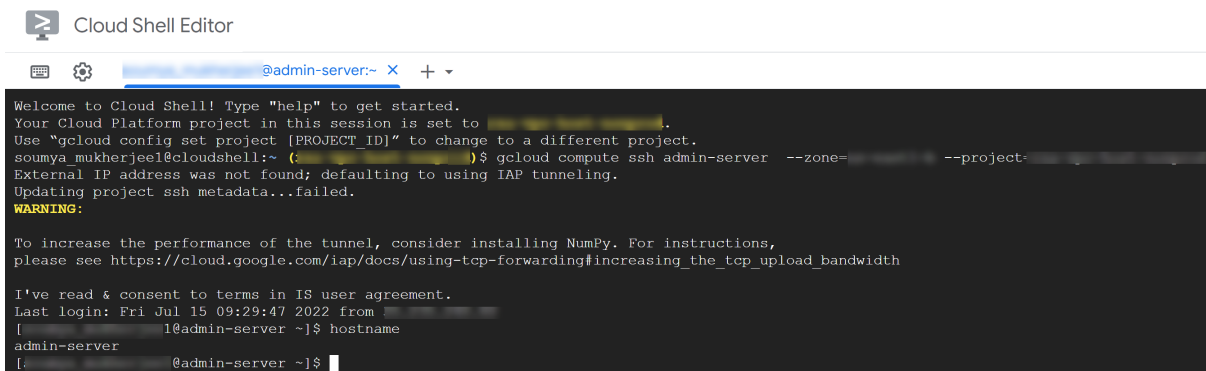
3. Run the following command:

```
- gcloud compute ssh <instance-name> --zone=<zone-name> --project=<project-name>
```

A confirmation pop-up is displayed to Authorize Cloud Shell.

4. Click **Authorize**.

This will connect to the VM instance through SSH.



Installation Tasks

Before you begin the installation tasks make sure you open the firewall ports. For more information on the lists of all the ports in a deployment, see the [Network Architecture and Ports](#) topic in the *Deployment Guide for NetWitness Platform 12.5*.

Caution: Do not proceed with the installation until the ports on your firewall are configured.

Tasks - Install 12.5 on the NetWitness Server (NW Server) Host and Component Hosts

Complete the following steps to install 12.5 on NW Server host and other component hosts. Steps that are specific to the NW Server host or to component hosts are noted.

Caution: If you want to install the Endpoint Relay Server, do not run the `nwsetup-tui` script. Follow the instructions in [\(Optional\) Installing and Configuring Relay Server](#) in the *Endpoint Configuration Guide for NetWitness Platform 12.5*.

1. Log in to the host with the `root` credentials and run the `nwsetup-tui` command to set up the host. This initiates the `nwsetup-tui` (Setup program) and the EULA is displayed.

Note: Use the following options to navigate the Setup prompts.

- 1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, and use the Tab key to move to and from commands (such as `<Yes>`, `<No>`, `<OK>`, and `<Cancel>`). Press **Enter** to register your command response and move to the next prompt.
- 2.) The Setup program adopts the color scheme of the desktop or console you use to access the host.
- 3.) If you specify DNS servers during the Setup program (`nwsetup-tui`) execution, they **MUST** be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` script to proceed. Any misconfigured DNS servers cause the Setup program to fail. If you need to reach a DNS server after setup that is unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see "(Optional) Task 1 - Re-Configure DNS Servers Post 12.5" in the "Post Installation Tasks" section in this guide.

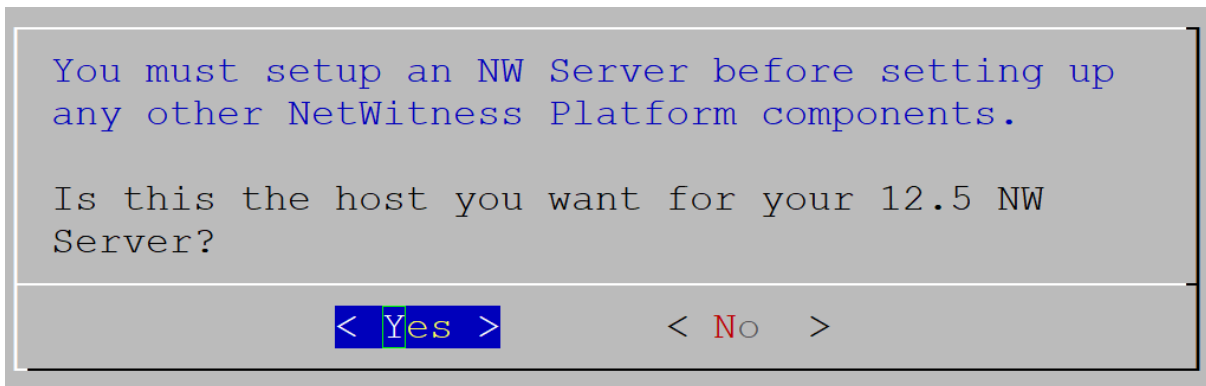
If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Update Repository** prompt in step 10 (the DNS servers are not defined so the system cannot access the external repo).

```
By clicking "Accept", you (the "Customer") hereby agree, on behalf of your company or
organization, to be bound by the terms and conditions of the End User License Agreement
(the "EULA") located at
https://www.rsa.com/content/dam/rsa/PDF/shrinkwrap-license-combined.pdf with RSA Security
LLC ("RSA", or appropriate affiliate entity in the relevant jurisdiction). In addition,
Customer hereby agrees and acknowledges that, if Customer chooses to host its data with any
third party or in a public cloud environment, RSA has no responsibility for the storage or
protection of any Customer data or for any associated security breach notifications. The
terms herein and in the EULA shall supersede any relevant terms in any other agreement
between the Customer and RSA. For customers of the RSA NetWitness® products, all data
analyzed in connection herewith shall be at a cost to Customer based on RSA's then current
through-put pricing model.
```

`<Accept >`

`<Decline>`

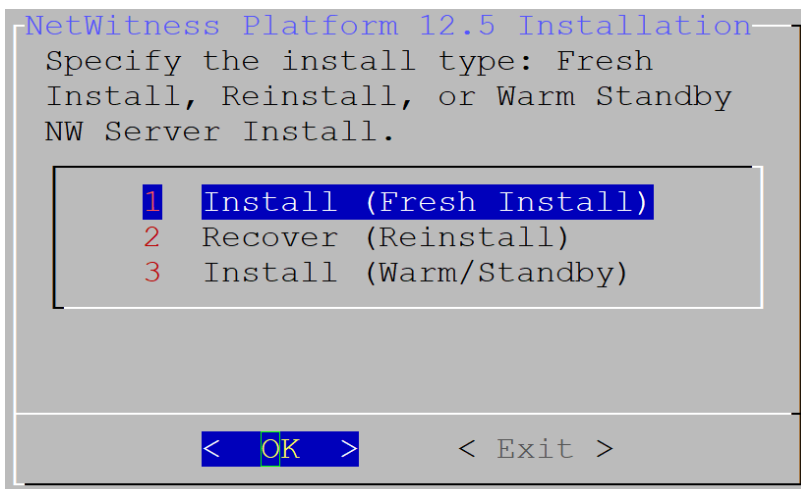
2. Tab to **Accept** and press **Enter**.
The **Is this the host you want for your 12.5 NW Server** prompt is displayed.



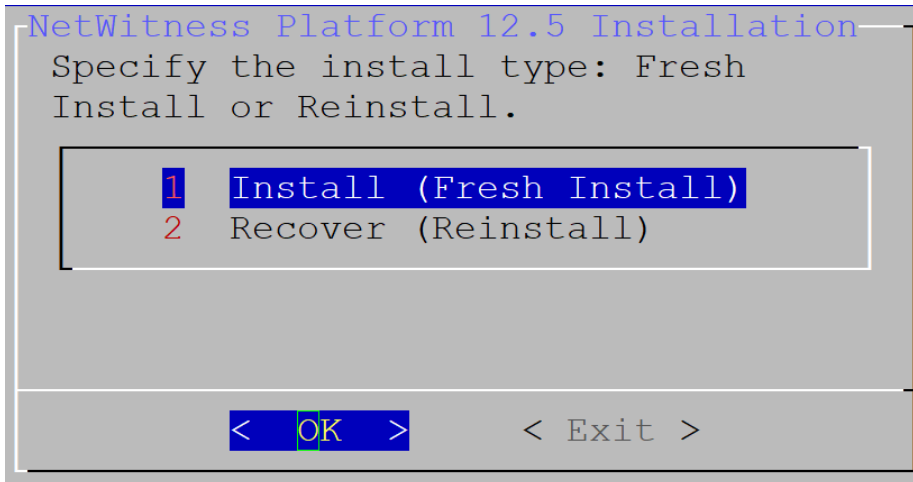
3. Tab to **Yes** and press **Enter** to install NW Server.
Tab to **No** and press **Enter** to install Component Hosts.

Caution: If you choose the wrong host for the NW Server and complete the setup, you must restart the setup program (step 2) and complete all the subsequent steps to correct this error.

4. The **Install** or **Recover** prompt is displayed.
NetWitness Server Host prompt

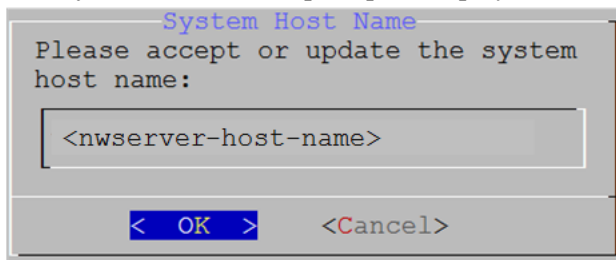


Component Hosts prompt



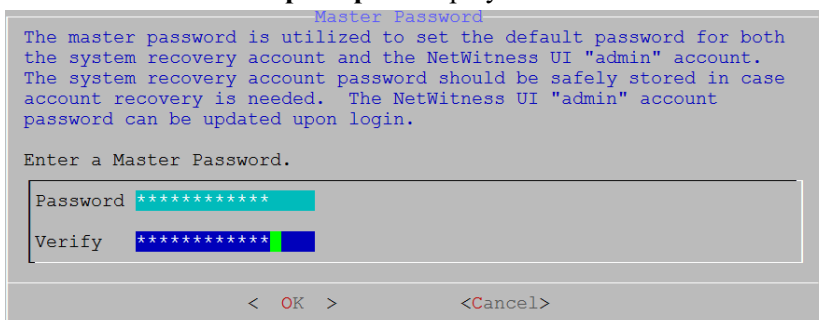
Press **Enter**. By default, the **Install (Fresh Install)** option is selected.

5. The **System Host Name** prompt is displayed.



Press **Enter** if you want to keep this name. If not, edit the host name, tab to **OK**, and press **Enter** to change it.

6. The **Master Password** prompt is displayed.



Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

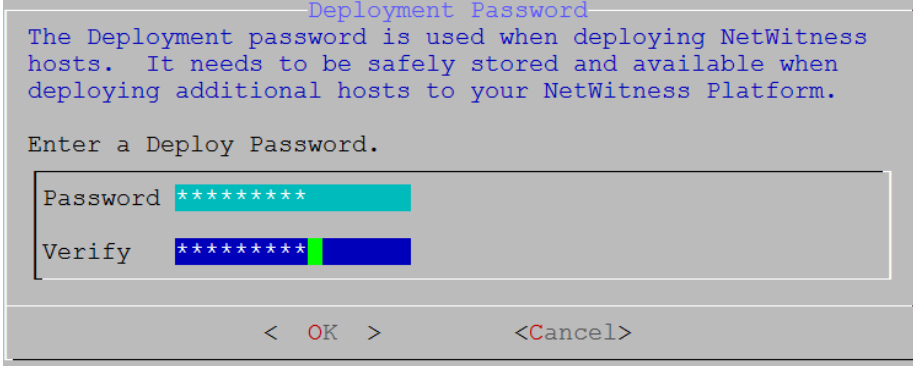
The following list of characters are supported for Master Password and Deployment Password:

- Symbols : ! @ # % ^ + ,
- Numbers : 0-9
- Lowercase Characters : a-z
- Uppercase Characters : A-Z

No ambiguous characters are supported for Master Password and Deployment Password (for example: space { } [] () / \ ' " ` ~ ; : . < > -).

Caution: This step 6 is applicable only for NW Server Host.


7. The **Deployment Password** prompt is displayed.



Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.

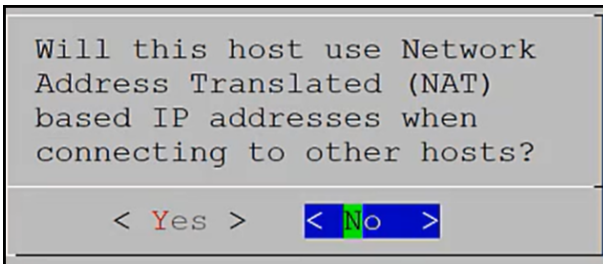
Caution: This step 7 is applicable for both NW Server and Component Hosts.

8. The setup program finds a valid IP address for this host and the following prompt is displayed.



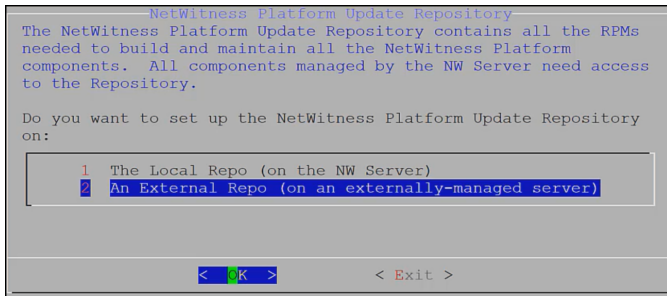
Tab to **No** and press **Enter**, if you want to use this IP and avoid changing your network settings.

9. The Use Network Address Translation (NAT) prompt is displayed.




- For the NW Server, tab to **No** and press **Enter**.
- For Component Hosts, if this host requires the use of NAT-based addresses to communicate with the NW Server, tab to **Yes**. Otherwise, tab to **No** and press **Enter**.

10. The **NetWitness Platform Update Repository** prompt is displayed.



- For NW Server Host: Select **2 An External Repo (on an externally-managed server)**. Tab to **OK** and press **Enter**.
- For Component Hosts: Select **1 The Local Repo (on the NW Server)**. Tab to **OK** and press **Enter**.

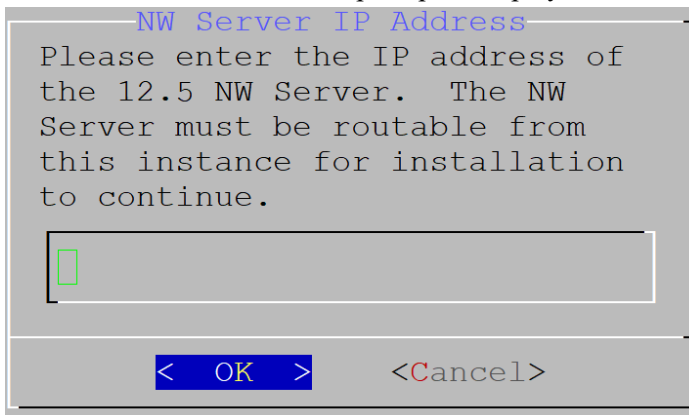
11. The **NetWitness Platform 12.5 External Update Repo URL** prompt is displayed.



Enter the base URL of the NetWitness Platform external repo, tab to **OK** and press **Enter**.

Caution: This step 11 is applicable only for NW Server Host.


12. The **NW Server IP Address** prompt is displayed.



Enter the IP address of 12.5 NW Server host, tab to **OK** and press **Enter**.


Caution: This step 12 is applicable only for Component Hosts.

13. The **Disable Firewall** prompt is displayed.



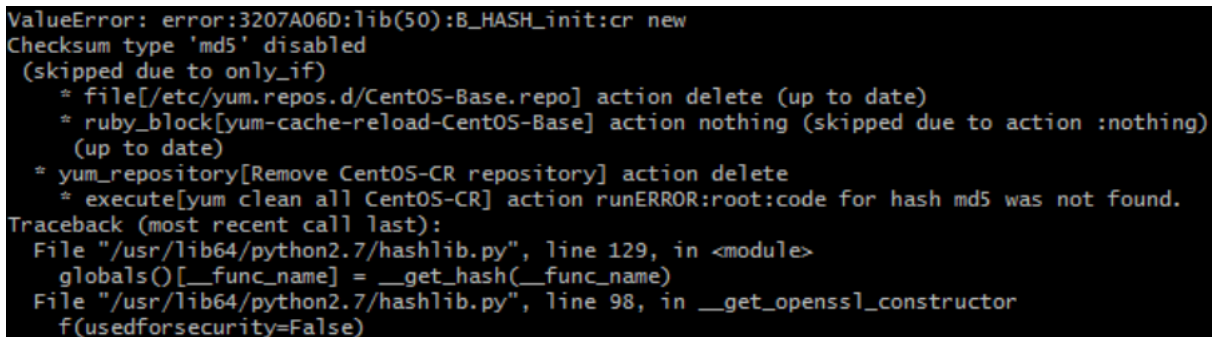
Tab to **No** and press **Enter**.

14. Press **Enter** to install 12.5 on the NW Server.
The **Start Install/Upgrade** prompt is displayed.



15. When **Installation complete** is displayed, you have installed the 12.5 NW Server on this host.

Note: Ignore the hash code errors similar to the errors shown in the following screen shot that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.



Note: If you want to perform a silent installation using CLI, see the [Silent Installation Using CLI](#) topic in the *Physical Host Installation Guide*.

Storage Configurations

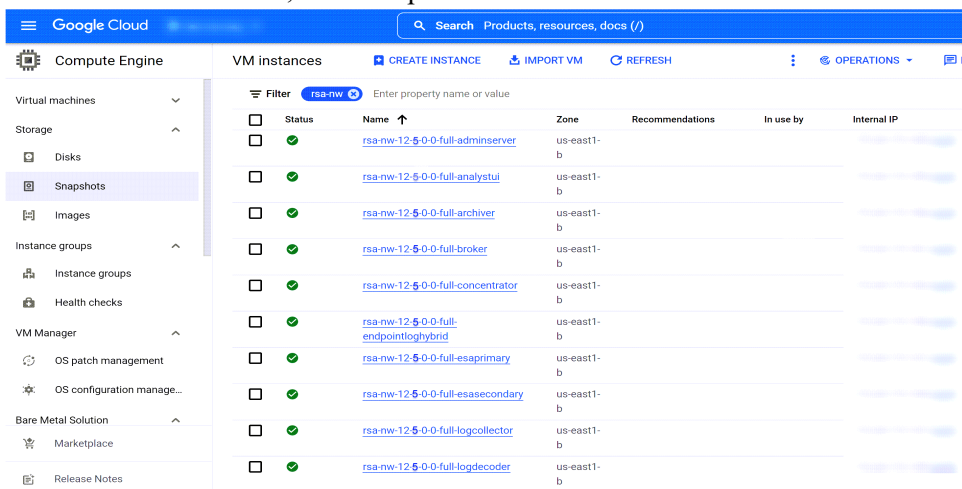
For storage allocations of all host types, see the [Prepare Virtual or Cloud Storage](#) topic in the *Storage Guide for NetWitness® Platform 12.5*.


Create a Disk

To understand which volumes are required to support the instance, see the [Storage Requirements](#) topic in the *Storage Guide for NetWitness® Platform 12.5*.

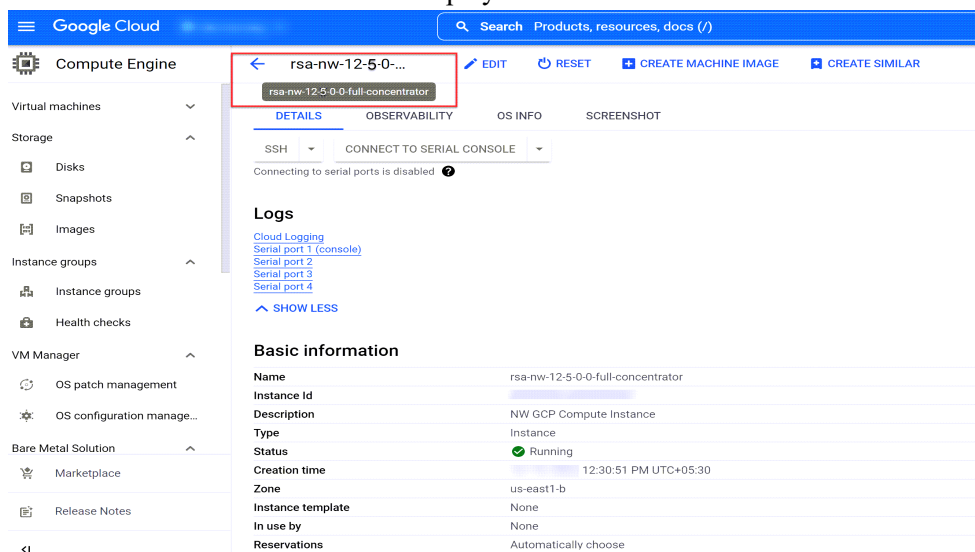
To create a disk

1. Go to **Google Cloud Platform > Compute Engine > VM Instance** to identify the VM instance.
2. Click the instance name, for example **rsa-nw-12-5-0-0-full-concentrator**.

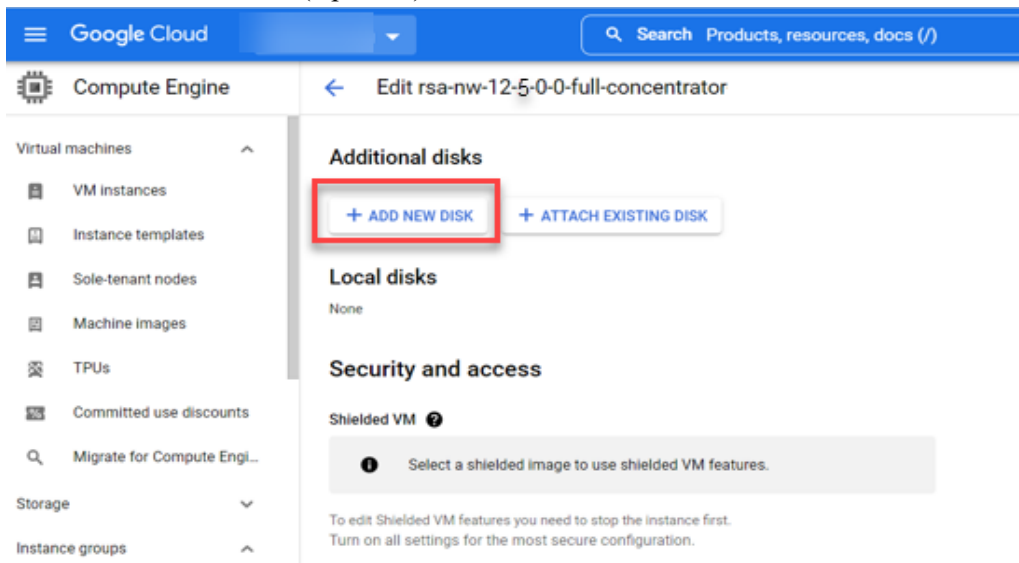


3. Click  to make the required changes.

The VM instance details view is displayed.



4. Under **Additional disks** (Optional), click [+ Add new disk](#) to add a new disk.



5. The **New disk** dialog is displayed.

The screenshot shows the 'New disk' dialog in the Google Cloud Platform console. The dialog is titled 'New disk (full-concentrator, Blank, 500 GB)'. It contains the following fields and options:

- Name:** Name is permanent. Input field: full-concentrator.
- Description (Optional):** Input field.
- Type:** Standard persistent disk (selected from a dropdown menu).
- Snapshot schedule:** No schedule (selected from a dropdown menu). A warning message says 'Create snapshot schedules to automatically back up your data.' with a 'Dismiss' button and a link to 'Learn more about creating snapshot schedules'.
- Source type:** Blank disk (selected), Image, Snapshot.
- Mode:** Read/write (selected), Read only.
- Deletion rule:** Keep disk (selected), Delete disk.
- Size (GB):** 500.
- Estimated performance:**

Operation type	Read	Write
Sustained random IOPS limit	375.00	750.00
Sustained throughput limit (MB/s)	60.00	60.00
- Encryption:** Data is encrypted automatically. Select an encryption key management solution.
 - Google-managed key (No configuration required)
 - Customer-managed key (Manage via Google Cloud Key Management Service)
 - Customer-supplied key (Manage outside of Google Cloud)
- Device name:** Used to reference the device for mounting or resizing. Input field: Based on disk name (default), full-concentrator.

At the bottom, there is a message: 'You're creating an unformatted disk. Format the disk after you attach it to your VM instance. [Formatting and mounting a zonal persistent disk](#)'. Below this is a 'Save changes to complete adding this disk.' message and 'Done' and 'Cancel' buttons.

Enter values in the following fields:

- Name:** Enter a unique name.
- (Optional) Description:** Enter the required additional information.
- Type:** Select **Standard persistent disk** from the drop-down menu.
- Snapshot schedule:** By default, **No schedule** option is selected.

- e. Under **Source type** > **Blank disk**.
 - **Mode**: By default, **Read/write** option is selected.
 - **Deletion rule**: By default, **Keep disk** option is selected.
 - f. **Size (GB)**: Enter the required size.
 - g. Under **Encryption**: By default, **Google-managed key** is selected.
 - h. **Device name**: By default, **Based on disk name (default)** is selected.
 - i. Click **Done**.
6. Click **Save**.

Note: To add another disk to the instance type, repeat steps 1 to 6.

Configure Hosts (Instances) in NetWitness Platform

Configure individual hosts and services as described in NetWitness [Host and Services Getting Started Guide](#). This guide also describes the procedures for applying updates and preparing for version upgrades.

Configure Packet Mirroring


You must complete the following tasks to create packet mirroring in GCP.

- [Task 1 - Create Instance Group](#)
- [Task 2 - Create Packet Mirroring Policy](#)
- [Task 3 - Create Load Balancer](#)
- [Task 4 - Verify Packet Mirroring in GCP](#)

Task 1 - Create Instance Group

Complete the following steps to create the Instance group:

1. In the Google Cloud Console, go to the **Instance groups** page.
2. Click **Create Instance Group**.



3. Click **New Unmanaged Instance Group** from the left panel.

Google Cloud

Search Products, resources, docs (/)

← Create Instance Group

Set up a group of load balancing VMs. [Learn more](#)

New managed instance group (stateless)
Automatically manage groups of VMs that do stateless serving and batch processing.

New managed instance group (stateful)
Automatically manage groups of VMs that have persistent data or configurations (such as databases or legacy applications).

New unmanaged instance group
Manually manage groups of load balancing VMs.

Name *
instance-group-2
Name is permanent

Description

Location

Region * us-east1 **Zone *** us-east1-b

Network and instances
Select instances that reside in a single zone, VPC network, and subnet.

Network * us-east1

Subnetwork * us-east1

VM instances
decoder

Select VMs
decoder

Port mapping
To send traffic to instance group through a named port, create a named port to map the incoming traffic to a specific port number, then go to "HTTP load balancing" to create a load balancer using this instance group.

CREATE CANCEL EQUIVALENT COMMAND LINE

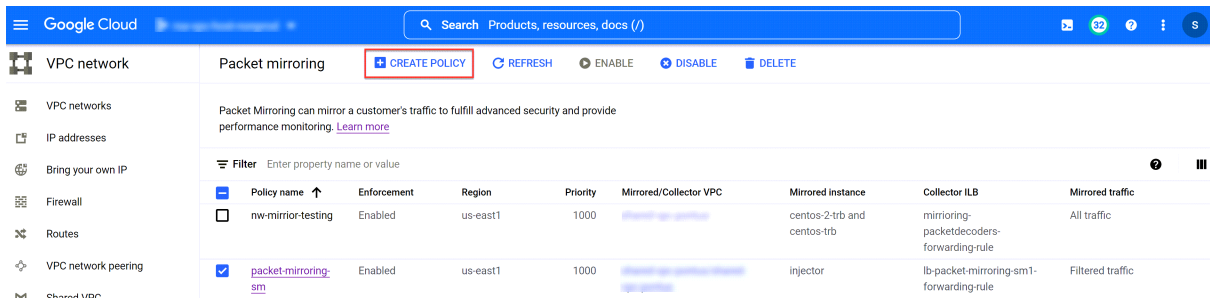
- Fill in the details to configure the unmanaged instance group:
 - Name:** Enter a name for the unmanaged instance group.
 - Description:** Enter the description for instance group.
 - Under **Location:** Select a region from the **Region** drop-down list and select a zone from the **Zone** drop-down list.
 - Network:** Select a network from the drop-down list.
 - Subnetwork:** Select a subnetwork from the drop-down list.
 - VM Instance:** Select the required decoder from the **Select VMs** drop-down list. The mirrored traffic will be sent to this VM instance.
 - Port Mapping:** By default, the port is selected.
- Click **Create**.

Task 2 - Create Packet Mirroring Policy

In the following procedure, Mirrored source and destination VM instances are in the same VPC network or subnetwork. You can also set both Mirrored source and destination VM instances on a different VPC network or subnetwork. For more information, see <https://cloud.google.com/vpc/docs/using-packet-mirroring>.

Complete the following steps to Create Packet Mirroring Policy.

1. In the Google Cloud console, go to the **Packet Mirroring** page.
2. Click **Create Policy**.



3. Under **Define policy overview**, enter the following details:

1 Define policy overview

Policy name * ?
 Lowercase, no spaces.

Region *

Policy priority

Policy enforcement

Enabled
 Disabled

- **Name:** Enter a name for the policy.
 - **Region:** Select the region from the drop-down list.
 - Under **Policy enforcement**, select **Enabled** and click **Continue**.
4. Under **Select VPC network**, select **Mirrored source and collector destination are in the same VPC network**.

✓ Define policy overview

2 Select VPC network

Select the VPC network or networks where your mirrored and collector instances are located. You can only select networks that you have permissions to use.

If the mirrored and collector instances are in the same network, select **Mirrored source and collector destination are in the same VPC network**. If they are in different networks that are peered, select **Mirrored source and collector destination are in separate, peered VPC networks**. [Learn more](#)

Mirrored source and collector destination are in the same VPC network

Network *

Mirrored source and collector destination are in separate, peered VPC networks

CONTINUE

- Select the required network from the **Network** drop-down list.
- Click **Continue**.

5. Under **Select mirrored source**, select **Select individual instances**.

Select individual instances

Select one or more VM instances as your mirrored source

Instance selection table

Filter Enter property name or value

VM instance name	Project	Zone
<input type="checkbox"/> admin-server	rsa-vpc-host-nonprod	us-east1-b
<input type="checkbox"/> adminserver-dnd	dev-nw-eng	us-east1-b
<input type="checkbox"/> centos-2-trib	gcp-nw-integration	us-east1-b
<input type="checkbox"/> centos-trib	gcp-nw-integration	us-east1-b
<input type="checkbox"/> concentrator	rsa-vpc-host-nonprod	us-east1-b
<input type="checkbox"/> decoder	rsa-vpc-host-nonprod	us-east1-b
<input type="checkbox"/> decoder-19302	dev-nw-eng	us-east1-b
<input type="checkbox"/> do-nest-vm	dev-nw-eng	us-east1-b
<input type="checkbox"/> fraseb-test1	rsa-vpc-host-nonprod	us-east1-b
<input checked="" type="checkbox"/> injector	rsa-vpc-host-nonprod	us-east1-b
<input type="checkbox"/> ml-qe-do-adminserver	gcp-nw-integration	us-east1-b
<input type="checkbox"/> ml-qe-do-concentrator	gcp-nw-integration	us-east1-b
<input type="checkbox"/> ml-qe-do-decoder	gcp-nw-integration	us-east1-b
<input type="checkbox"/> rsa-nw-adminserver	gcp-nw-integration	us-east1-b
<input type="checkbox"/> rsa-nw-packet-decoder	gcp-nw-integration	us-east1-b
<input type="checkbox"/> sa-19302	dev-nw-eng	us-east1-b

SELECT CANCEL

Select the required instances from the **Instance selection table** and click **Continue**.

Note: All traffic will be mirrored from this instance.

6. Under **Select collector destination**, click **create new L4 internal load balancer** and follow [Task 3 - Create Load Balancer](#) to complete configuration of load balancer.

✓ Define policy overview

✓ Select VPC network

✓ Select mirrored source

4 Select collector destination

Select an L4 internal load balancer that balances traffic across your collector instances (the backend instances), which collect all the mirrored traffic. The load balancer must have a forwarding rule specifically for packet mirroring. [Learn more](#)

Collector project [SELECT PROJECT](#)

Collector destination *

You can also [create new L4 internal load balancer](#)

[CONTINUE](#)

5 Select mirrored traffic

[SUBMIT](#) [CANCEL](#)

7. Once the load balancer is created, click **Refresh** in the **Collector destination** drop-down menu.

Define policy overview
 Select VPC network
 Select mirrored source
 4 Select collector destination

Select an L4 internal load balancer that balances traffic across your collector instances (the backend instances), which collect all the mirrored traffic. The load balancer must have a forwarding rule specifically for packet mirroring. [Learn more](#)

Collector project [SELECT PROJECT](#)

Collector destination *

Filter | Type to filter

Yc

lb-packet-mirroring-sm1-forwarding-rule (lb-packet-mirroring-sm1)

[REFRESH](#)

5 S

[SUBMIT](#) [CANCEL](#)

8. Select the newly created load balancer and **Continue**.
9. Do one of the following:
 - Select either both ingress (incoming) and egress (outgoing) traffic to be mirrored from the source VM instance by enabling **Mirror all traffic (default)** which is same as **Allow both ingress and egress traffic** under **Traffic direction**.
 - Select **Mirror filtered traffic** and select **Allow egress traffic only** to mirror only the outgoing traffic from the source VM instance and send them to decoder.
 - Select **Mirror filtered traffic** and select **Allow ingress traffic only** to mirror only the incoming traffic from from the source VM instance you want and send them to decoder.

← Create policy

✓ Select mirrored source

✓ Select collector destination

5 Select mirrored traffic

Specify the traffic to mirror. By default, all ingress and egress is mirrored. If you want to reduce the amount of mirrored traffic, add filters to mirror only certain traffic. [Learn more](#)

Mirror all traffic (default)

Mirror filtered traffic

Protocol filters

Allow all protocols

Allow specific protocols

IP range filters

Allow all IP ranges

Allow specific IP ranges

Traffic direction

Allow both ingress and egress traffic

Allow ingress traffic only

Allow egress traffic only

SUBMIT CANCEL

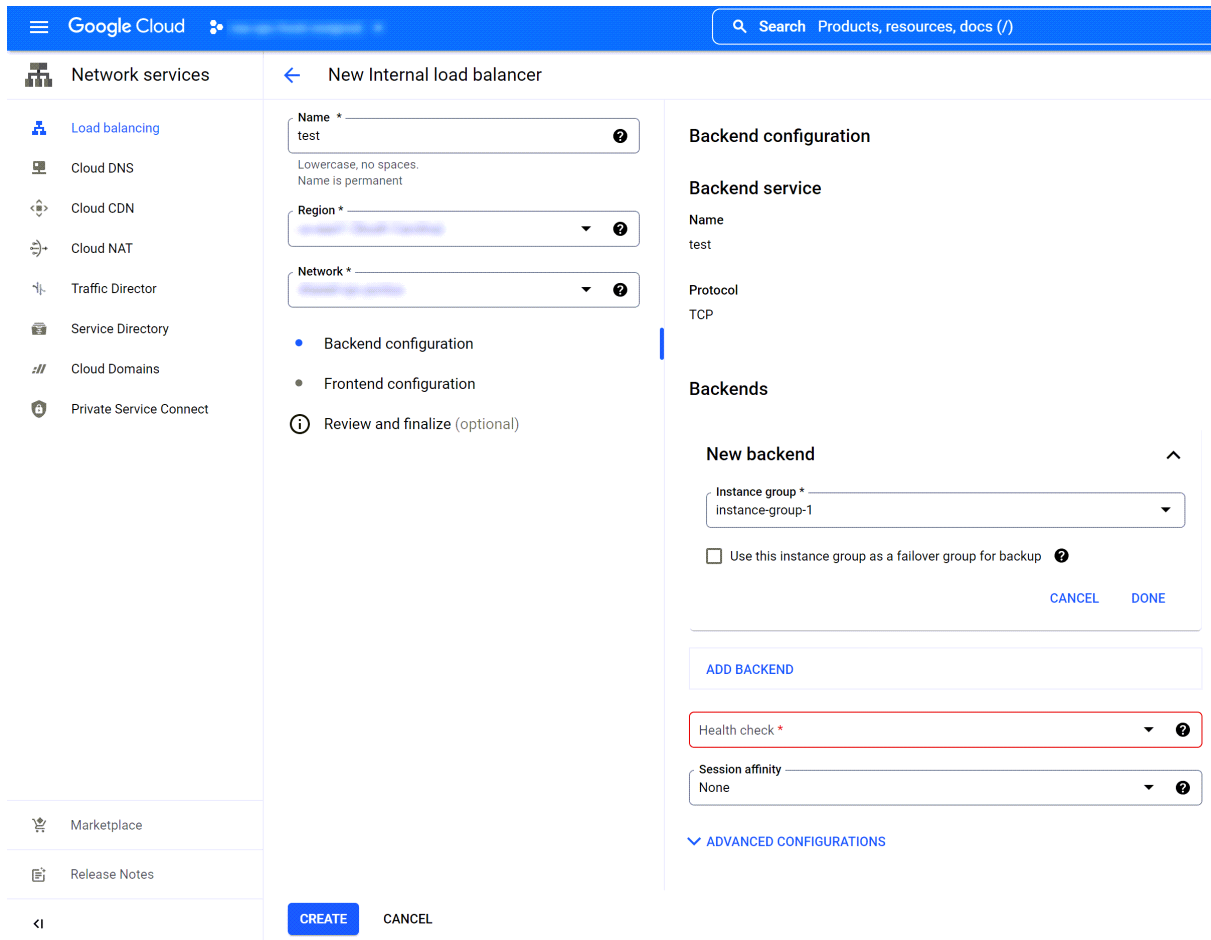
10. Click **Submit**. The policy will be created successfully.

Task 3 - Create Load Balancer

Complete the following steps to create the Load Balance:

1. Click **Load balancing**.
2. Click **+ Create Load Balancers**.
3. Enter the following details to configure load balancer:
 - **Name**: Enter a name for the load balancer.
 - **Region**: Select a region from the drop-down list.
 - **Network**: Select a network from the drop-down list.

4. Click **Backend configuration**.



5. Under **New backend**, select the instance group created in [Task 1 - Create Instance Group](#) from the **Instance group** drop-down list.
6. Perform the following steps to create a Health Check:

Health Check

Name * ?
Name is required

Description

Scope

Global
 Regional

Region
us-east1 (South Carolina) ?

Protocol ?
Port * ?

Proxy protocol ?

Request ? Response ?

Logs

On
Turning on Health check logs can increase costs in Cloud Logging.
 Off

Health criteria

Define how health is determined: how often to check, how long to wait for a response, and how many successful or failed attempts are decisive

Check interval * seconds ?
Timeout * seconds ?

Healthy threshold * consecutive successes ?

Unhealthy threshold * consecutive failures ?

SAVE CANCEL

- **Name:** Enter a name for the health check.
 - (Optional) **Description:** Enter the description for the health check.
 - Under **Scope:** By default, **Regional** option is selected.
 - **Region:** Select the region from the drop-down list.
 - **Protocol:** Select the **TCP** protocol from the drop-down list.
 - **Port:** Enter the port number **80**.
 - **Proxy Protocol:** By default, **None** option is selected.
 - **Logs:** By default, **Off** option is selected.
 - Under **Health Criteria**, retain the default values.
 - Click **Save**.
7. Click **Frontend Configuration**.
 8. Under **New Frontend IP and port**, configure the following details:


The screenshot displays the 'New Internal load balancer' configuration interface in the Google Cloud console. The interface is organized into three main sections:

- Left Sidebar:** Contains navigation links for 'Network services' and various services like 'Load balancing', 'Cloud DNS', 'Cloud CDN', 'Cloud NAT', 'Traffic Director', 'Service Directory', 'Cloud Domains', and 'Private Service Connect'.
- Main Configuration Area:**
 - Name:** A text input field containing 'test'. Below it, a note states: 'Lowercase, no spaces. Name is permanent.'
 - Region:** A dropdown menu with a blurred selection.
 - Network:** A dropdown menu with a blurred selection.
 - Configuration Steps:** A list of steps: 'Backend configuration', 'Frontend configuration' (which is selected with a blue checkmark), and 'Review and finalize (optional)' (with an information icon).
- Right Sidebar (Frontend configuration):**
 - New Frontend IP and port:** A section header.
 - Name (Optional):** A text input field.
 - Description:** A text input field.
 - Protocol:** A dropdown menu set to 'TCP'.
 - Subnetwork:** A dropdown menu set to 'dev-nw-eng'.
 - Internal IP:**
 - Purpose:** Radio buttons for 'Non-shared' (selected) and 'Shared'.
 - IP address:** A dropdown menu set to 'Ephemeral (Automatic)'.
 - Ports:** Radio buttons for 'Single', 'Multiple', and 'All' (selected).
 - Global access:** Radio buttons for 'Disable' (selected) and 'Enable'.
 - Service label (Optional):** A text input field.
 - Packet Mirroring:** A checkbox labeled 'Enable this load balancer for Packet Mirroring' which is checked.
 - SHOW LESS:** A link to collapse the configuration options.


At the bottom of the main configuration area, there are two buttons: 'CREATE' (highlighted in blue) and 'CANCEL'.

- **(Optional) Name:** Enter a name.
- **Description:** Provide the description.
- Under **Protocol**, in the **TCP** section, select a **subnetwork** from the drop down list.
- Under **Internal IP**, in the **Purpose** section, by default **Non-shared** is selected.
- **IP address:** By default, **Ephemeral (Automatic)** is selected.
- Under **Ports**, select **All** option.
- Under **Global Access**, by default **Disable** is selected.
- Under **Packet Mirroring**, select **Enable this load balancer for Packet Mirroring**.
- (Optional) Under **Review and Finalize**, review all the configured details.
- Click **Create**.

- The Load Balancer will be created. The Health Check may show unhealthy status but it does not impact packet mirroring policy.



Task 4 - Verify Packet Mirroring in GCP

- Go to the Packet Decoder service and click  > **View** > **Config** page and set the **Capture Interface Selected** parameter to the following value:
packet_mmap_eth0 (bpf)

Decoder Configuration	
Name	Config Value
Adapter	
Berkeley Packet Filter	
Capture Interface Selected	packet_mmap_eth0 (bpf)

- SSH to the source mirrored instance.
- Run the following command `curl <yahoo.com>` on source mirrored instance.
- SSH to the Packet Decoder.
- Run the following command `tcpdump` on Packet Decoder and verify if the traffic has been mirrored.

```
[root@decoder soumya.mukherjee]# tcpdump | grep HTTP
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes
11:13:09.450642 IP injector.us-east1-b.c.rsa-vmc-host-nonprod.internal.41188 > media-router-fp74.prod.media.vip.gql.yahoo.com.http: Flags [P.], seq 0:73, ack 1, win 222, options [nop,nop,TS val 867058835, ecc 201727156], length 73: HTTP: GET / HTTP/1.1
11:13:19.884829 IP metadata.google.internal.http > decoder.us-east1-b.c.rsa-vmc-host-nonprod.internal.35222: Flags [P.], seq 914507579:914510133, ack 671653804, win 65535, length 2554: HTTP: HT
TP/1.1 200 OK
11:13:19.885086 IP metadata.google.internal.http > decoder.us-east1-b.c.rsa-vmc-host-nonprod.internal.35222: Flags [P.], seq 3844548539:3844551093, ack 2506033132, win 65535, length 2554: HTTP:
HTTP/1.1 200 OK
11:13:19.885338 IP decoder.us-east1-b.c.rsa-vmc-host-nonprod.internal.35222 > metadata.google.internal.http: Flags [P.], seq 1:228, ack 2554, win 65320, length 227: HTTP: GET /computeMetadata/v
1/?recursive=true&alt=json&wait_for_change=true&last_etag=c0bfa909c63f335&timeout_sec=60 HTTP/1.1
11:13:19.885701 IP decoder.us-east1-b.c.rsa-vmc-host-nonprod.internal.35228 > metadata.google.internal.http: Flags [P.], seq 1:229, ack 2554, win 65320, length 228: HTTP: GET /computeMetadata/v
1/?recursive=true&alt=json&wait_for_change=true&timeout_sec=60&last_etag=c0bfa909c63f335 HTTP/1.1
```

- Log in to the NetWitness Platform to verify on Packet Decoder.
- Go to **Investigate** > **Events** and select **Concentrator** from query profile drop-down menu.
- To verify, click search will filter out the domains name based on their configuration.

GCP Instance Configuration Recommendations

Note: These recommendations can be used as a baseline for 12.5.0.0 and adjusted as needed.

Instance compute, and memory utilization will vary depending on content applied, ingestion rates, and the number of running queries.

This topic contains the minimum GCP instance configuration settings recommended for the NetWitness virtual stack components.

- Compute Engine Instance:
 - Minimum instance type - **n2-standard-32** is the minimum instance type required for any NetWitness component image so that it can function.
 - Machine type adjustments - You must adjust machine types according to your ingestion rate, content and parsers, dashboard reports, scheduled reports, investigations, and active users.
 - All the components were integrated.
 - The Log stream includes a Log Decoder, Concentrator, and Archiver.
 - The Endpoint Hybrid stream includes an Endpoint Server, Concentrator, and Log Decoder.
 - Respond receives alerts from the Reporting Engine, and Event Stream Analysis.
 - The background load includes reports, charts, alerts, investigation, and respond.

- Persistent Disk (Storage)

For performance recommendations, recommended storage allocation per NetWitness host, and input/output operations per second, see the "[Storage Requirements](#)" topic in the *Storage Guide for NetWitness® Platform 12.5.0.0*.

The following table displays the specification recommendations for NetWitness GCP instances.

- RAID's were not configured because single SSD disks provided the required IO/s, and no scaling issues were found.

IMPORTANT: The recommended configurations can handle up to 15,000 requests per second (EPS). However, if your system is under a heavier load, you can increase the memory size to accommodate more requests.

Virtual Log Collector (VLC)

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-4	4	16 GB
10,000	n2-standard-4	4	16 GB

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
15,000	n2-standard-4	4	16 GB

Archiver

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-4	4	16 GB
10,000	n2-standard-4	4	16 GB
15,000	n2-standard-4	4	16 GB

Broker

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-4	4	16 GB
10,000	n2-standard-4	4	16 GB
15,000	n2-standard-4	4	16 GB

Log Concentrator

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-8	8	32 GB
10,000	n2-standard-8	8	32 GB
15,000	n2-standard-8	8	32 GB

Note: The memory can be increased to handle the query load on the concentrator. This includes queries on the **Investigate** page, alert generation rules, and more. You can also adjust the maximum number of concurrent queries that can be run on the concentrator, based on the load.

Event Stream Analysis (ESA)

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
9,000	n2-standard-8	8	32 GB
18,000	n2-standard-16	16	64 GB
30,000	n2-standard-32	32	128 GB

Log Decoder

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
5,000	n2-standard-8	8	32 GB
10,000	n2-standard-16	16	32 GB
15,000	n2-standard-32	32	32 GB

Packet Decoder

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
500 Mbps	n2-standard-custom	12	32 GB
1 Gbps	n2-standard-custom	24	64 GB

Packet Concentrator

Compute Engine Instance			
EPS	Machine Type	Virtual CPU's	Memory
500 Mbps	n2-standard-8	8	32 GB

NetWitness Endpoint Hybrid

Compute Engine Instance			
Agents	Machine Type	Virtual CPU's	Memory
15,000 agents	n2-standard-48	48	192 GB

New Health and Wellness

Compute Engine Instance		
Machine Type	Virtual CPU's	Memory
n2-standard-4	4	16 GB

NetWitness Server and Co-Located Components

Compute Engine Instance		
Machine Type	Virtual CPU's	Memory
n2-standard-16	16	128 GB

Note: Extra memory is necessary for efficiently managing the workload of queries, including generating reports, charts, alerts, and lists on the Reporting Engine.

Analyst UI

Compute Engine Instance		
Machine Type	Virtual CPU's	Memory
n2-standard-8	8	32 GB

UEBA

Compute Engine Instance		
Machine Type	Virtual CPU's	Memory
n2-standard-16	16	64 GB