ARISTA

Deployment Guide

VeloCloud SD-WAN Google Cloud Platform Virtual Edge

Services

Arista.com

Arista Networks

Headquarters	Support	Sales
5453 Great America Parkway Santa Clara, CA 95054 USA +1-408-547-5500	+1-408-547-5502 +1-866-476-0000	+1-408-547-5501 +1-866-497-0000
www.arista.com/en/	support@arista.com	sales@arista.com

© Copyright 2025 Arista Networks, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos, and service marks ("Marks") displayed in this documentation are the property of Arista Networks in the United States and other countries. Use of the Marks is subject to the Arista Networks Terms of Use Policy, available at www.arista.com/en/terms-of-use. Use of marks belonging to other parties is for informational purposes only.

Contents

Chapter 1: Google Cloud Platform Virtual Edge Deployment Guide..... 1

Chapter 2: Google Cloud Virtual Edge Deployment Overview..... 2

Chapter 3: Virtual Edge Deployment on GCP..... 3

Chapter 4: Prepare GCP Environment..... 5
4.1 Create a VPC Network..... 5
4.2 Create Inbound Firewall Rules..... 7
4.3 Create Routes in a VPC Network..... 9
4.3.1 Add a Branch-to-Branch Route in a VPC Network..... 10

Chapter 5: Provision an Edge..... 12

Chapter 6: Deploy a Virtual Edge from GCP Marketplace..... 18

Chapter 7: Deploy a Virtual Edge with GCP Deployment Manager..... 21
7.1 Enable Deployment Manager..... 22

Chapter 8: Verify Edge Activation..... 25

Appendix A: References..... 26
A.1 Related Documents..... 26

Chapter 1

Google Cloud Platform Virtual Edge Deployment Guide

This document provides instructions for deploying Virtual VeloCloud Edge on Google Cloud Platform (GCP).

Chapter 2

Google Cloud Virtual Edge Deployment Overview

More customers are moving workload to Public Cloud infrastructure and expect to extend SD-WAN from remote sites to public cloud to guarantee SLA. There are multiple options, leveraging distributed VeloCloud Gateways to establish IPSec towards public cloud private network or deploy virtual edge directly on Google Cloud Platform (GCP).

For small branch deployment that demand throughput less than 1G, single virtual edge can be deployed in the Private GCP network. For larger data center deployments that demand multi-gig throughput, hub clustering can be deployed.

Note: In the VeloCloud Hub clustering design, a Layer 3 Instance is leveraged on the LAN side to run BGP between hubs in the cluster and the Layer 3 Instance for route distribution in LAN. Since the GCP router does not support dynamic routing protocol, a third-party virtual router is required in the GCP infrastructure.

This document illustrates a basic topology and a high-level workflow to deploy a virtual Edge (vVCE) on GCP. For more information, see Virtual Edge Deployment on GCP.

Prerequisites

- GCP account and login information.
- Familiarity with GCP networking concepts. For more information, see https://cloud.google.com/vpc/docs/overview.
- Orchestrator target and admin account to login.

GCP Machine Types

Bandwidth throughput and the number of network interfaces need to be considered when sizing the VeloCloud SD-WAN Virtual Edge. The minimum number of network interfaces required is three (GE1, GE2, GE3).

Table 1: Bandwidth Throughput

Throughput	30 Mbps	50 Mbps	100 Mbps	200 Mbps	400 Mbps	1 Gbps
vCPU	2	2	2	2	4	4
Memory	4 GB	4 GB	4 GB	8 GB	8 GB	8 GB

Table 2: Machine Type

Machine Type	vCPUs	Memory (Gb)	Max NICs
n1-standard-4	4	15	4
n1-standard-8	8	30	8

Chapter 3

Virtual Edge Deployment on GCP

Describes the Virtual Edge deployment on the Google Cloud Virtual Private Cloud (VPC) with three VPC networks: Management VPC (10.0.2.x/24), Public VPC (10.0.0.x/24), and Private VPC (10.0.1.x/24), each for a subnet connected to the Edge as shown in the following topology diagram.

Basic Topology

Figure 3-1: Basic GCP Topology

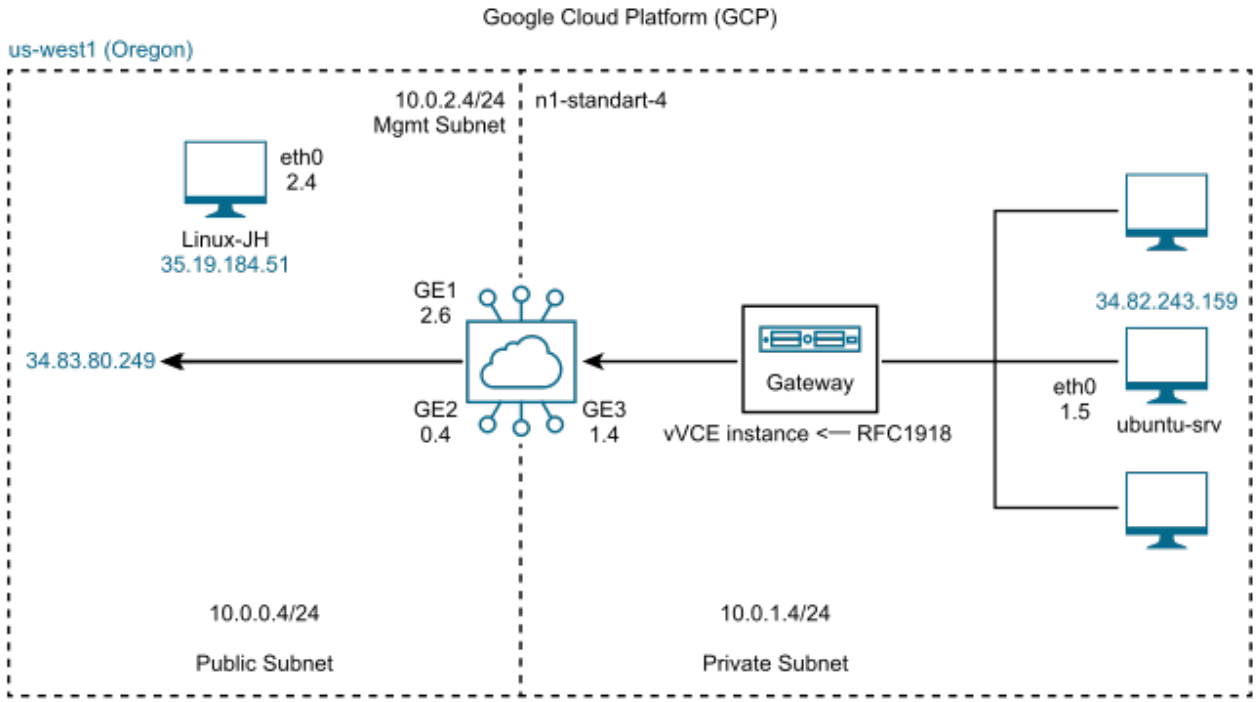

The Virtual Edge routes between the two subnets. The Public VPC Routes forward all offnet traffic to the Internet Gateway. The Gateway Router in the Private subnet will forward all traffic to the LAN facing interface on the Virtual Edge (GE3). In this example, a default route is used to forward “ALL” traffic from the workloads but is not necessary. RFC1918 summarization or specific branch/hub prefixes can be used to narrow what is sent to the Virtual Edge. For example, if the workload in the Private Subnet still needs to be accessible via SSH from publicly sourced IPs then the VPC Router could be configured to point the default route (0.0.0.0/0) to Internet Gateway and RFC1918 summarization to Virtual Edge.

High-Level Workflow

To deploy a VeloCloud SD-WAN Virtual Edge on Google Cloud Platform, perform the following steps:

1. Prepare the GCP environment:

a. Create three Virtual Private Cloud (VPC) networks (MGMT VPC network, Public VPC network, and Private VPC network), each for a subnet connected to the Edge (n1-standard-4) as shown in the topology diagram.

   1. MGMT subnet for console/management access to the Edge through Management Interface GE1.
   2. Public subnet for Internet access from the Edge through WAN-side Interface GE2.
   3. Private subnet for LAN-side device access through LAN-side Interface GE3.

   For steps, see Create a VPC Network.

b. Create inbound firewall rules for VPC networks: Management, Private, and Public. For steps, see Create Inbound Firewall Rules.

c. Add a new default route (0.0.0.0/0) entry in route table of Private VPC network pointing to the Edge, with the next hop IP address as the Edge's GE3 interface IP.

   For steps, see Create Routes in a VPC Network.

2. Provision an Edge on the VeloCloud Orchestrator as follows:

   a. Create an Edge of type **Virtual Edge**.

   b. Change GE1 interface to **Routed** from **Switched**, and deactivate **WAN Overlay** and **NAT Direct Traffic**.

   c. Change GE2 interface to **Routed** from **Switched**, and activate **WAN Overlay** and **NAT Direct Traffic**.

   d. Deactivate **WAN Overlay** and **NAT Direct Traffic** for GE3 interface, which will be the next hop for devices connected to Private Subnets (LAN devices).

   For additional information, see Provision an Edge.

3. Deploy the Virtual Edge. You can deploy the Virtual Edge by using one of the following methods:

   - Deploy a Virtual Edge from GCP Marketplace
   - Deploy a Virtual Edge with GCP Deployment Manager

4. Verify if the Virtual Edge is up in the Orchestrator.

Chapter 4

Prepare GCP Environment

Before deploying a Virtual Edge on Google Cloud Platform (GCP), you must prepare the GCP environment by completing the following steps:

- Create a VPC Network
- Create Inbound Firewall Rules
- Create Routes in a VPC Network

4.1 Create a VPC Network

Ensure you have a Google account and access/login information to the Google Cloud Platform (GCP) Console.

You can choose to create an Automatic mode or Custom mode Virtual Private Cloud (VPC) network. Automatic mode networks create one subnet in each Google Cloud region automatically when you create the network. For Custom mode VPC networks, you have to create a network and then create subnets that you want within a region. You can create subnets when you create the network or you can add subnets later, but you cannot create instances in a region that has no subnet defined.

1. Log into the GCP Console.
2. Select **VPC Networks**.
   The **VPC Networks** page appears.
3. Select **Create VPC network**.

The **Create a VPC network** page appears.

Figure 4-1: Create a VPC network

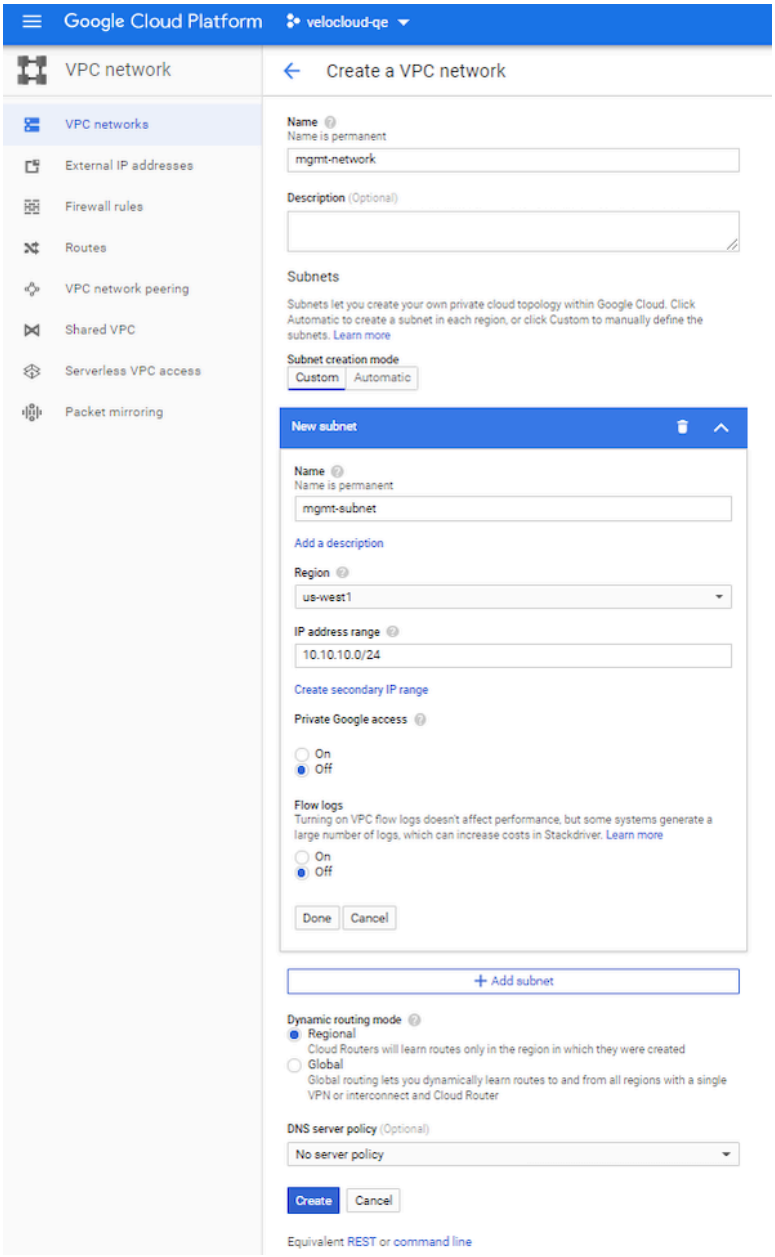

4. In the **Name** textbox, enter a unique name for the VPC network.
5. Under **Subnets**, choose **Custom** or **Automatic** as the **Subnet creation mode**. If you choose **Custom**, then in the **New subnet** area, specify the following configuration parameters for a subnet:
   a. In the **Name** textbox, enter a unique name for the subnet.
   b. From the **Region** drop-down menu, select a region for the subnet.
   c. In the **IP address range** textbox, enter an IP address range.
   d. To define a secondary IP range for the subnet, select **Create secondary IP range**.

e. **Private Google access**: Choose whether to activate Private Google Access for the subnet when you create it or later by editing it.

f. **Flow logs**: Choose whether to activate VPC flow logs for the subnet when you create it or later by editing it.

g. Select **Done**.

6. To add more subnets, select **Add subnet** and repeat the steps in Step 5. You can also add more subnets to the network after you have created the network.
7. Choose the **Dynamic routing mode** for the VPC network.
8. Select **Create**.

   The VPC network and subnet are created.

Create Inbound Firewall Rules

4.2 Create Inbound Firewall Rules

- Ensure you have a Google account and access/login information to the Google Cloud Platform (GCP) Console.
- Ensure you have created the VPC networks.
- Review the firewall rule components and ensure to become familiar with firewall configuration components as used in Google Cloud.

Firewall rules are defined at the network level, and only apply to the network where they are created. To create inbound firewall rules for a VPC network, perform the following steps:

1. Log into the GCP Console.
2. Select **VPC Networks**.

   The **VPC Networks** page appears.
3. Select the VPC network for which you want to add firewall rules.

   The **VPC network details** page for the selected VPC network appears.

4. Go to the **Firewall rules** tab and select **Add firewall rule**. The **Create a firewall rule** page appears.

Figure 4-2: Create a firewall rule

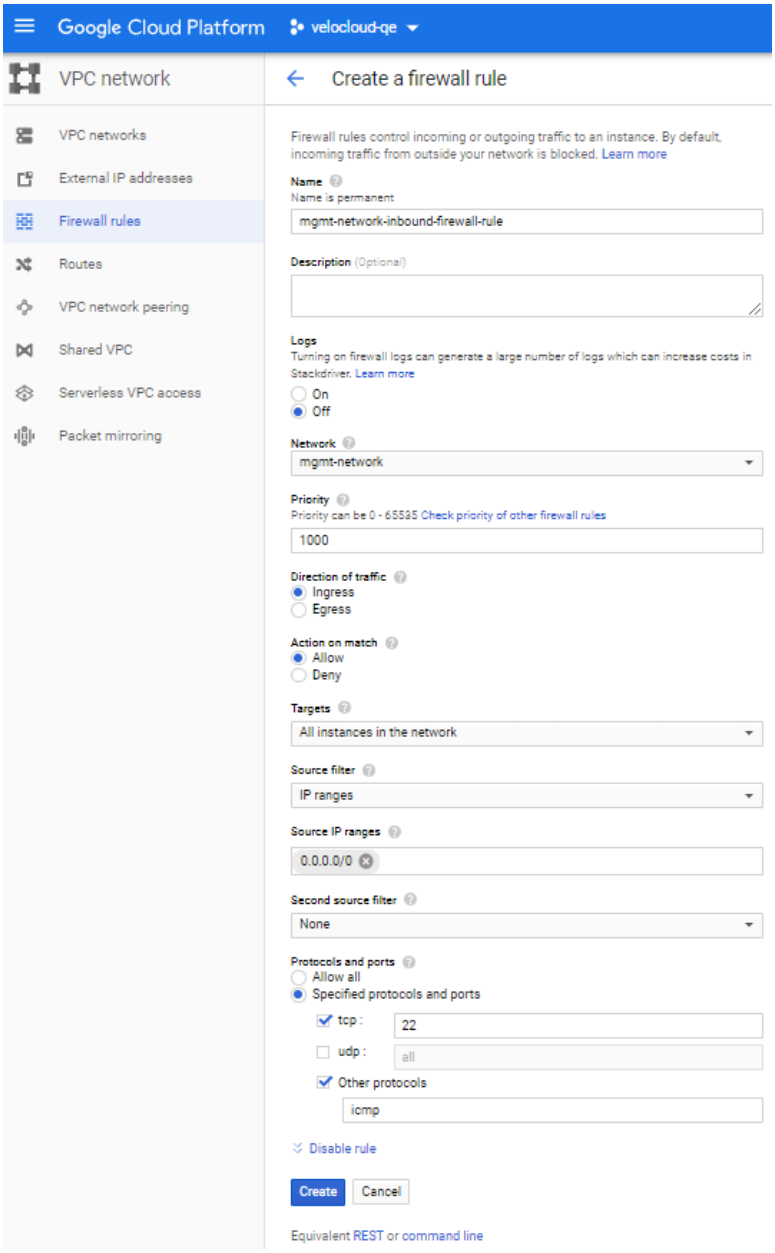

5. In the **Name** textbox, enter a unique name for the firewall rule.
6. Optionally, you can activate firewall logging by selecting **On** under **Logs**. By default, firewall logging is deactivated.
7. For **Direction of traffic**, choose **ingress**.
8. For **Action on match**, choose **Allow** or **Deny**.
9. From the **Targets** drop-down menu, select the targets for the rule:
   - If you want the rule to apply to all instances in the network, choose **All instances in the network**.
   - If you want the rule to apply to select instances by network (target) tags, choose **Specified target tags**, then type the tags to which the rule should apply into the **Target tags** textbox.

- If you want the rule to apply to select instances by associated service account, choose **Specified service account**, indicate whether the service account is in the current project or another one under **Service account scope**, and choose or type the service account name in the **Target service account** field.

10. From the **Source** filter drop-down menu, select **IP ranges**.
11. In the Source IP ranges textbox, enter the CIDR blocks to define the source for incoming traffic by IP address ranges. Use 0.0.0.0/0 for a source from any network.
12. Define the **Protocols and ports** to which the rule will apply:
    a. Select **Allow all** or **Deny all**, depending on the action, to have the rule apply to all protocols and ports.
    b. Define specific protocols and ports:
       - Select **tcp** to include the TCP protocol and ports. Enter all or a comma delimited list of ports, such as 20-22, 80, 8080.
       - Select **udp** to include the UDP protocol and ports. Enter all or a comma delimited list of ports, such as 67-69, 123.
       - Select **Other protocols** to include protocols such as ICMP, VCMP, SNMP, and so on as per the requirement.
13. Optionally, you can create the firewall rule, but not enforce it by setting its enforcement state to deactivated. Select **Disable rule**, then select **Disable**.
14. Select **Create**.
    The firewall rules are created for the selected VPC network.

Create Routes in a VPC Network

4.3 Create Routes in a VPC Network

- Ensure you have a Google account and access/login information to the Google Cloud Platform (GCP) Console.
- Ensure you have created VPC networks.

Describes how to add a new default route in a Private Virtual Private Cloud (VPC) network pointing to an Edge as illustrated in topology diagram.

1. Log into the GCP Console.
2. Select **VPC Networks**.
   The **VPC Networks** page appears.
3. Select the VPC network (Private VPC network) for which you want to add a new default route.
   The **VPC network details** page appears.
4. Go to the **Routes** tab and then delete the default route that was created during the VPC network creation.
5. Select **Add route**.

The **Create a route** page appears.

Figure 4-3: Create a route

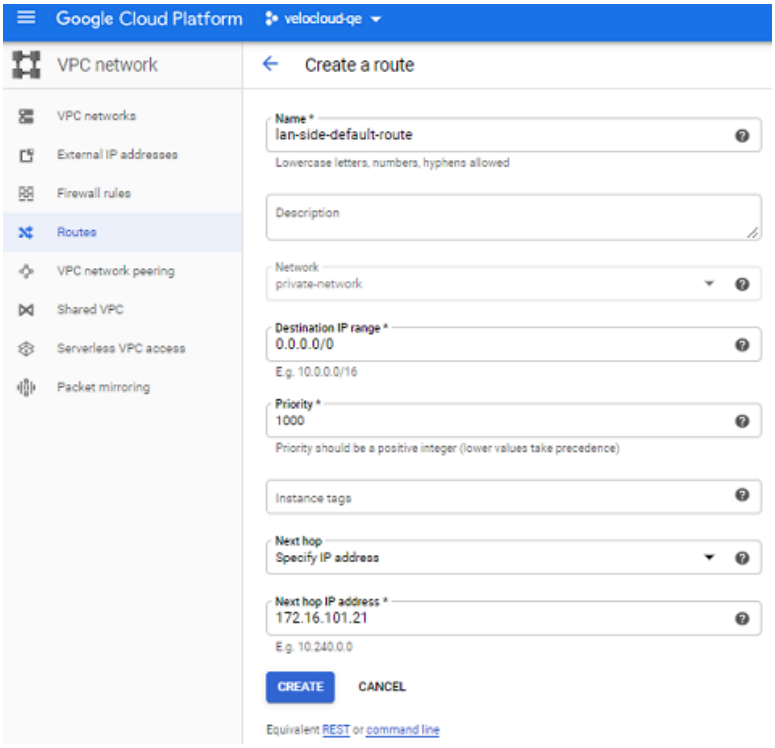

6. Enter the following details:
   a. In the **Name** textbox, enter a unique name for the route entry.
   b. In the **Destination IP range** textbox, specify the new default route (for example, 0.0.0.0/0).
   c. In the **Priority** textbox, specify a priority for the route. A priority is only used to determine routing order if routes have equivalent destinations.
   d. From the **Next hop** drop-down menu, select **Specify IP address**.
   e. In the **Next hop IP address** textbox, enter the IP address of the edge interface in the selected VPC network.
7. Select **Create**.
   A route entry is added in the route table of the selected VPC network.

4.3.1 Add a Branch-to-Branch Route in a VPC Network

- Ensure you have a Google account and access/login information to the Google Cloud Platform (GCP) Console.
- Ensure you have created VPC networks.

Describes how to add a branch-to-branch route in a Public Virtual Private Cloud (VPC) network pointing to an Edge as illustrated in Single-Arm Topology.

1. Log into the GCP Console.

2. Select **VPC Networks**.
   The **VPC Networks** page appears.
3. Select the VPC network (Public VPC network) for which you want to add a branch-to-branch route.
   The **VPC network details** page appears.
4. Go to the **Routes** tab and select **Add route**.
   The **Create a route** page appears.

   Figure 4-4: Create a route

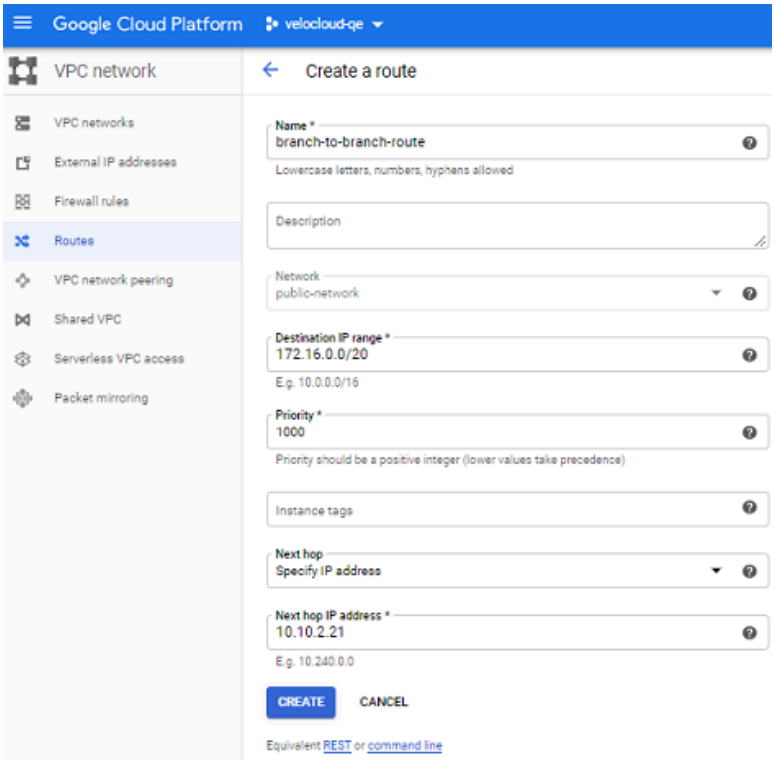

5. Enter the following details:
   a. In the **Name** textbox, enter a unique name for the route entry.
   b. In the **Destination IP range** textbox, specify the IP address of a branch in the enterprise network, for example 172.16.0.0/20.
   c. In the **Priority** textbox, specify a priority for the route. A priority is only used to determine routing order if routes have equivalent destinations.
   d. From the **Next hop** drop-down menu, select **Specify IP address**.
   e. In the **Next hop IP address** textbox, enter the IP address of the edge interface in the selected VPC network.
6. Select **Create**.
   A route entry is added in the route table of the selected VPC network.

Chapter 5

Provision an Edge

Ensure you have the Orchestrator host name and admin account to login.

To provision an Edge, perform the following steps:

1. Log into the Orchestrator application as Admin user, with your login credentials.
2. Go to **Configure** > **Edges**.
3. In the **Edges** screen, select **Add Edge**.
   The **Provision an Edge** screen displays.

   Figure 5-1: Provision an Edge

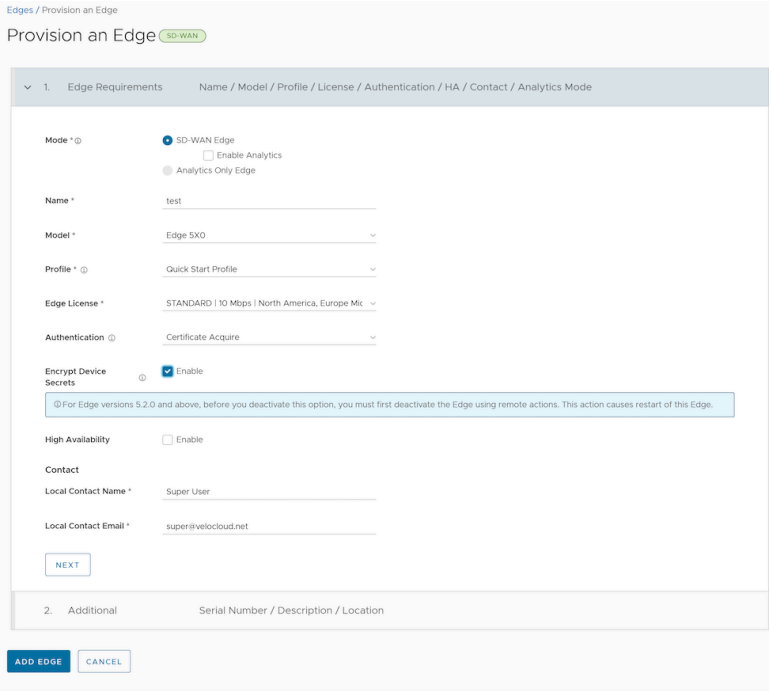

4. You can configure the following options:

Table 3: Option Descriptions

Option	Description
Mode	**SD-WAN Edge** mode selected by default.
Name	Enter a unique name for the Edge.
Model	Select **Virtual Edge** from the menu.
Profile	Select **Quick Start Profile** from the menu. **Note:** If an **Edge Staging Profile** is displayed as an option due to Edge Auto-activation, it indicates that this Profile is used by a newly assigned Edge, but has not been configured with a production Profile.
Edge License	Select an Edge license from the menu. The list displays the licenses assigned to the Enterprise, by the Operator.

Option	Description
Authentication	Select the mode of authentication from the menu: • **Certificate Deactivated**: Edge uses a pre-shared key mode of authentication. **Warning:** This mode is not recommended for any customer deployments. • **Certificate Acquire** - Selected by default and recommended for all customer deployments. With **Certificate Acquire** mode, certificates issue at the time of Edge activation and renew automatically. The Orchestrator instructs the Edge to acquire a certificate from the certificate authority of the Orchestrator by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Edge uses the certificate for authentication to the Orchestrator and for establishment of VCMP tunnels. **Note:** After acquiring the certificate, the option can be updated to **Certificate Required**, if needed. • **Certificate Required**: This mode is only appropriate for customer enterprises that are "static". A static enterprise is defined as one where no more than a few new Edges are likely to be deployed and no new PKI oriented changes are anticipated. **Important: Certificate Required** has no security advantages over **Certificate Acquire**. Both modes are equally secure and a customer using **Certificate Required** should do so only for the reasons outlined in this section. **Certificate Required** mode means that no Edge heartbeats accepted without a valid certificate. **CAUTION:** Using this mode can cause Edge failures in cases where a customer is unaware of this strict enforcement. With this mode, the Edge uses the PKI certificate. Operators can change the certificate renewal time window for Edges by editing the Orchestrator's System Properties. For more information, contact your Operator. **Note:** • With the Bastion Orchestrator feature enabled, the Edges that are to be staged to Bastion Orchestrator should have the authentication mode set to either **Certificate Acquire** or **Certificate Required**. • When an Edge certificate is revoked, the Edge is deactivated and needs to go through the activation process. The current QuickSec design checks certificate revocation list (CRL) time validity. The CRL time validity must match the current time of Edges for the CRL to have impact on new established connection. To implement this, ensure the Orchestrator time is updated properly to match with the date and time of the Edges.
Encrypt Device Secrets	Select **Enable** to allow the Edge to encrypt the sensitive data across all platforms. This option is also available on the Edge **Overview** page. **Note:** For Edge versions 5.2.0 and above, before you deactivate this option, you must first deactivate the Edge using remote actions. This causes restart of the Edge.
High Availability	Select **Enable** to apply High Availability (HA). Edges can be installed as a single standalone device or paired with another Edge to provide High Availability (HA) support.
Local Contact Name	Enter the name of the site contact for the Edge.
Local Contact Email	Enter the email address of the site contact for the Edge.

5. Enter all the required details and select **Next** to configure the following additional options:

   Note: The **Next** button is activated only when you enter all the required details.

   Table 4: Additional Option Descriptions

Option	Description
Serial Number	Enter the serial number of the Edge. If specified, the Edge must display this serial number on activation. **Note:** When deploying virtual VeloCloud SD-WAN Edges on AWS Edges, make sure to use the instance ID as the serial number for the Edge.
Description	Enter an appropriate description.
Location	select the **Set Location** link to set the location of the Edge. If not specified, the location is auto-detected from the IP address when the Edge is activated.

6. Select **Add Edge**. The Edge is provisioned, and the activation key is displayed on the top of the page. Make a note of the activation key to use it for launching the Edge from the AliCloud Console.

   Note: The activation key expires in one month if the Edge device is not activated against it.

7. Configure Virtual Edge interfaces. The following steps are explained considering Topology A.
   a. Go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
   b. Select the link to an Edge or select the **View** link in the **Device** column of the Edge. The configuration options for the selected Edge are displayed in the **Device** tab.
   c. Go to the **Interface Settings** area.
   d. In the **Connectivity** category, expand **Interfaces**. Different types of Interfaces available for the selected Edge are displayed.
   e. Select the **Override Interface** check box, and then update configurations of Virtual Edge interfaces **GE1 Interface**, **GE2 Interface**, and **GE3 Interface** as follows:
      - Change **GE1** interface capability to **Routed**, and deactivate **WAN Overlay** and **NAT Direct Traffic**.
      - Change **GE2** interface capability to **Routed** and ensure **WAN Overlay** and **NAT Direct Traffic** are activated.

- For **GE3** interface, deactivate **WAN Overlay** and **NAT Direct Traffic**, which will be the next hop for devices connected to Private VPC subnets (LAN devices).

Figure 5-2: Virtual Edge

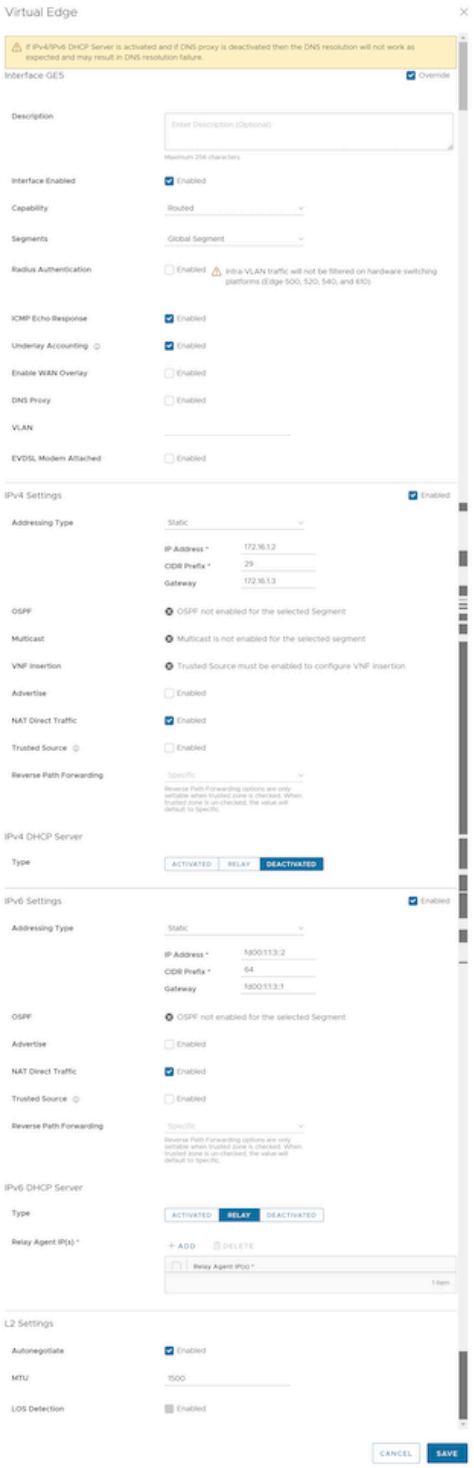

For additional information, see the topic *Configure Interface Settings for Edges* in the *Arista VeloCloud SD-WAN Administration Guide*.

8. Select **Save**.
   A Virtual Edge is provisioned on the Orchestrator.

Deploy the Virtual Edge on GCP. You can deploy the Virtual Edge by using one of the following methods:

- Deploy a Virtual Edge from GCP Marketplace
- Deploy a Virtual Edge with GCP Deployment Manager

Chapter 6

Deploy a Virtual Edge from GCP Marketplace

Ensure you have a Google account and access/login information to the Google Cloud Platform (GCP) Console.

In the Google Cloud Console, you can create and deploy a Virtual Machine (VM) instance using a boot disk image, a boot disk snapshot, or a container image. To create and deploy a Virtual Machine (VM) instance using a boot disk image, perform the following steps:

1. Log into the GCP Console.
2. VeloCloud SD-WAN is now available in the GCP Marketplace. Do a search for VeloCloud SD-WAN in the Marketplace to get started.

   Figure 6-1: Launching SD-WAN

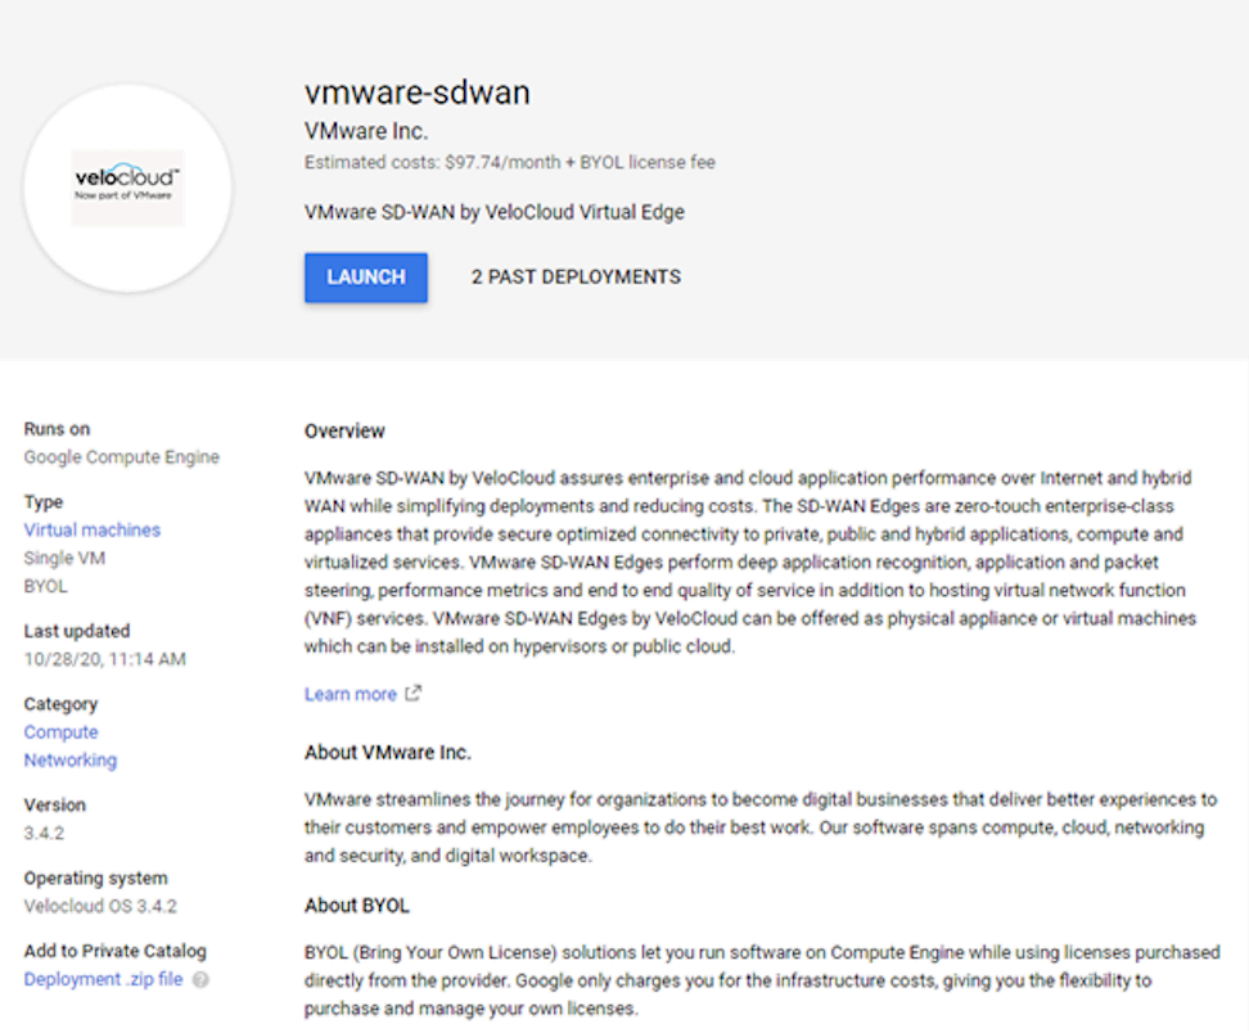

3. Select **LAUNCH**.

The **Create an instance** page appears.

Figure 6-2: Create an instance

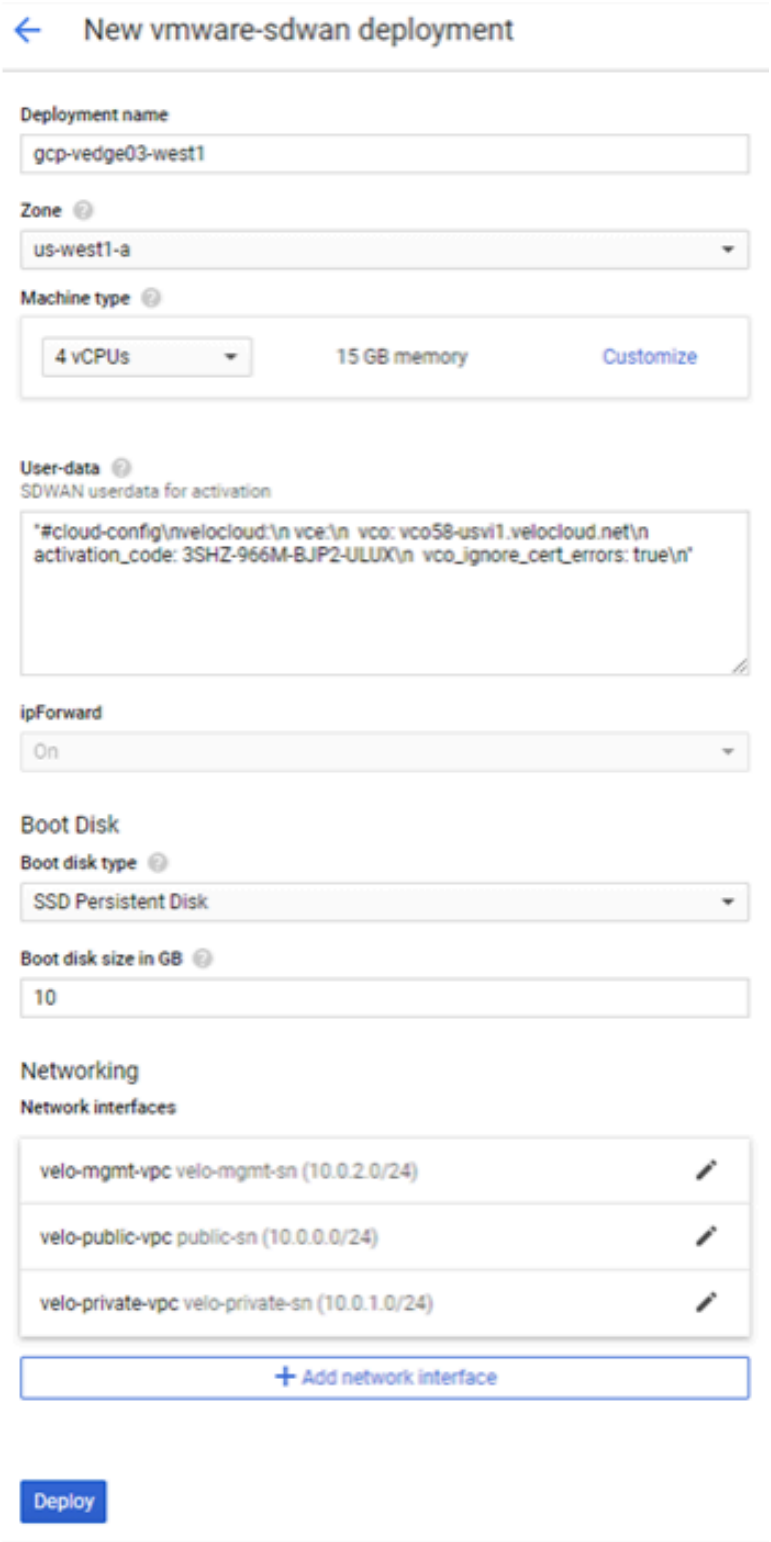

4. In the **Deployment name** textbox, enter a unique name for your instance.
5. From the **Zone** drop-down menu, select a region where the VPC networks are created.

6. Select a machine configuration for your instance. From the **Machine type** drop-down menu, select an option based on the topology configured.
7. In the **User-data** field, provide the cloud-init information in the following sample format for activating the virtual Edge against the target VeloCloud Orchestrator.

   Example: Sample cloud-init user-data

   ```
   "#cloud-config\nvelocloud:\n vce:\n vco: vco58-usvi1.velocloud.net\n activation_code:
    3SHZ-966M-BJP2-ULUX\n vco_ignore_cert_errors: false\n"
   ```

   It is very important that the format is correct to ensure it is processed otherwise activation will silently fail (i.e., no error event in the Orchestrator).

   The ***#cloud-config*** needs to be encapsulated with quotes or GCP will throw an error during the launch time. The quotes are already included in the **User-data** field. You can simply insert the cloud-config between the quotes or cut and paste the sample cloud-init user-data, replacing all text in the **User-data** field.

   Note: The newline (\n) characters and spacing must be exact for the parsing to work. It is best to cut and paste the example above into a Notepad and just replace values as needed without modifying the spacing.

   SSH Keys are managed at the project level. See https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys.
8. IP Forwarding allows the interfaces on the virtual Edge to process packets not destined for the local interface's MAC address. IP Forwarding is activated by default and this is required for proper routing to work and cannot be changed.
9. In the **Boot disk** area, the disk type and size of the Boot image should be left at their defaults: SSD Persistent Disk and 10 GB, respectively.
10. In the **Networking** area, add interfaces for the configured VPC networks, as follows:
    a. Under **Network interfaces**, select the **+ Add Network Interface** icon.
    b. From the **Network** drop-down menu, select the network to which you want to add an interface.
    c. Configure External IP as follows:
       - For Management network- Select **None** for **External IP**.
       - For Public network- Select **External IP: Ephemeral** as this interface needs to map to a public Internet IP.
       - For Private network- Select **None** for **External IP**.
    d. Select **Done**.
    e. To add another interface, select **Add network interface** and repeat the above steps from **b** to **d**.
11. Select **Deploy**.

    A virtual edge instance is created, and the Compute Engine automatically starts the Edge instance after it is created.

Verify Edge Activation

Chapter 7

Deploy a Virtual Edge with GCP Deployment Manager

To deploy a VeloCloud SD-WAN Virtual Edge on Google Cloud Platform using a Deployment Manager, perform the following steps:

1. Activate the Cloud Deployment Manager API in GCP. For steps, see Enable Deployment Manager.
2. Provision a Edge on the Orchestrator as follows:
   a. Create an edge of type **Virtual Edge** and make a note of the activation key that will be displayed on the top of the screen once the edge is provisioned.
   b. Configure a VLAN IP address (use 169.254.0.1 /24) for the edge. Do not activate **Advertise** and **DHCP**.
   c. Configure Virtual Edge interfaces as follows:
      - Change GE2 interface capability to **Routed** from **Switched** and activate **WAN Overlay** and **DHCP Addressing**.
      - For GE3 interface, deactivate **WAN Overlay** and **NAT Direct Traffic** as this interface will be used for the LAN-side gateway.

   For additional information, see Provision an Edge.

   Note: The VeloCloud Orchestrator needs the Device Settings configured first before edge activation. If this step is missed, the virtual edge activates but then goes offline a few minutes later.

3. Deploy the GCP image by creating the VPC networks first and then deploying the DM template with the relative reference for each interface. CLOUD-INIT is also used in the template to supply Orchestrator target and activation key for the virtual Edge.
   a. Create three Virtual Private Cloud (VPC) networks (MGMT VPC network, Public VPC network, and Private VPC network), each for a subnet connected to the Edge (n1-standard-4) as shown in the topology diagram.
      - Mgmt Subnet for console/management access to the Edge through Management Interface GE1.
      - Public Subnet for Internet access from the Edge through WAN-side Interface GE2.
      - Private Subnet for LAN-side device access through LAN-side Interface GE3.
   b. Modify the Deployment Manager (DM) template. The following is a sample YAML DM template. You can use this template, but ensure to make necessary changes for your environment. The YAML DM template will need the following items modified to fit the intended environment:
      - Project Name
      - Region and Zone
      - VPC-Names and Subnets

- Orchestrator IP or FQDN
- Activation Code (format: xxxx-xxxx-xxxx-xxxx)
- Orchestrator Ignore Cert Errors: true or false

```
# "Arista SD-WAN by VeloCloud GCP Deployment Manager Template (34220201029)" # gcloud
 deployment-manager deployments create velocloud-vce --config gcp_dm.yaml # gcloud
 deployment-manager deployments delete velocloud-vce resources: - type: compute.v1.instance
 name: dm-gcp-vce-01 properties: zone: us-west1-a machineType: https://www.googleapis.com/
compute/v1/projects/gcp-nsx-sdwan/zones/us-west1-a/machineTypes/n1-standard-4 canIpForward:
 true disks: - deviceName: boot type: PERSISTENT boot: true autoDelete: true initializePar
ams: sourceImage: https://www.googleapis.com/compute/v1/projects/vmware-sdwan-public/global/
images/vce-342-102-r342-20200610-ga-3f5ad3b9e2 networkInterfaces: - network: https://www.g
oogleapis.com/compute/v1/projects/gcp-nsx-sdwan/global/networks/velo-mgmt-vpc subnetwork:
 projects/gcp-nsx-sdwan/regions/us-west1/subnetworks/velo-mgmt-sn - network: https://www.g
oogleapis.com/compute/v1/projects/gcp-nsx-sdwan/global/networks/velo-public-vpc subnetwork:
 projects/gcp-nsx-sdwan/regions/us-west1/subnetworks/public-sn accessConfigs: - name:
 External NAT type: ONE_TO_ONE_NAT - network: https://www.googleapis.com/compute/v1/p
rojects/gcp-nsx-sdwan/global/networks/velo-private-vpc subnetwork: projects/gcp-nsx-sdwan/
regions/us-west1/subnetworks/velo-private-sn metadata: items: - key: user-data value: |
 #cloud-config velocloud: vce: vco: vco58-usvi1.velocloud.net activation_code: YPTF-PN33-THT
X-28V5 vco_ignore_cert_errors: false
```

For information about GCLOUD CLI, see https://cloud.google.com/sdk/gcloud/.

4. Verify if the virtual edge is activated in the Orchestrator.

   Once the instance is running in GCP and all information provided was correct, the virtual edge will reach out to the Orchestrator with the activation key, activate and perform software update if needed (and reboot if upgraded). Typical deployment time is between 3 to 4 minutes.

7.1 Enable Deployment Manager

- GCP account and login information.
- Familiarity with GCP Deployment Manager supported resource types. For more information, see https://cloud.google.com/deployment-manager/docs/configuration/supported-resource-types.

Deployment Manager is an infrastructure deployment service that automates the creation and management of Google Cloud resources. Deployment Manager uses the underlying APIs of each Google Cloud service to deploy your resources.

The Google Cloud Deployment Manager V2 API provides services for configuring, deploying, and viewing Google Cloud services and APIs via templates which specify deployments of Cloud resources. To activate the Cloud Deployment Manager V2 API and create credentials, perform the following steps:

1. Log on to the GCP Console.
2. Go to **APIs & Services** > **Dashboard**. The **APIs & Services** page appears.
3. Select **Enable APIS AND SERVICES**.
4. Use the **Search** textbox to find the Deployment Manager API.

5. Select **Cloud Deployment Manager V2 API** and then select **Enable**.

Figure 7-1: Cloud Deployment Manager V2 API

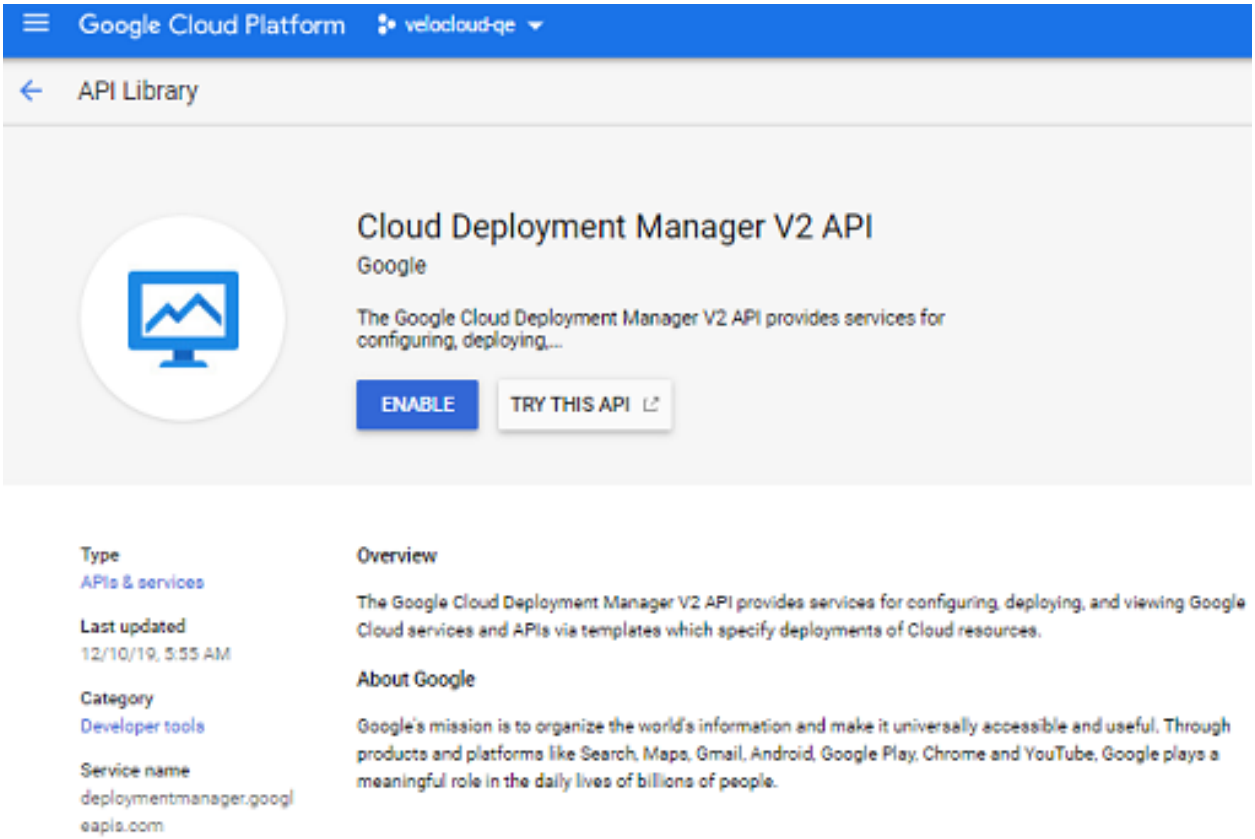

The Cloud Deployment Manager API is activated. To use this API, you must create credentials.

6. Select **Credentials** > **CREATE CREDENTIALS** and select one of the following options to create credentials:
   - API key
   - OAuth client ID
   - Service account

- Help me choose

Figure 7-2: Create Credentials

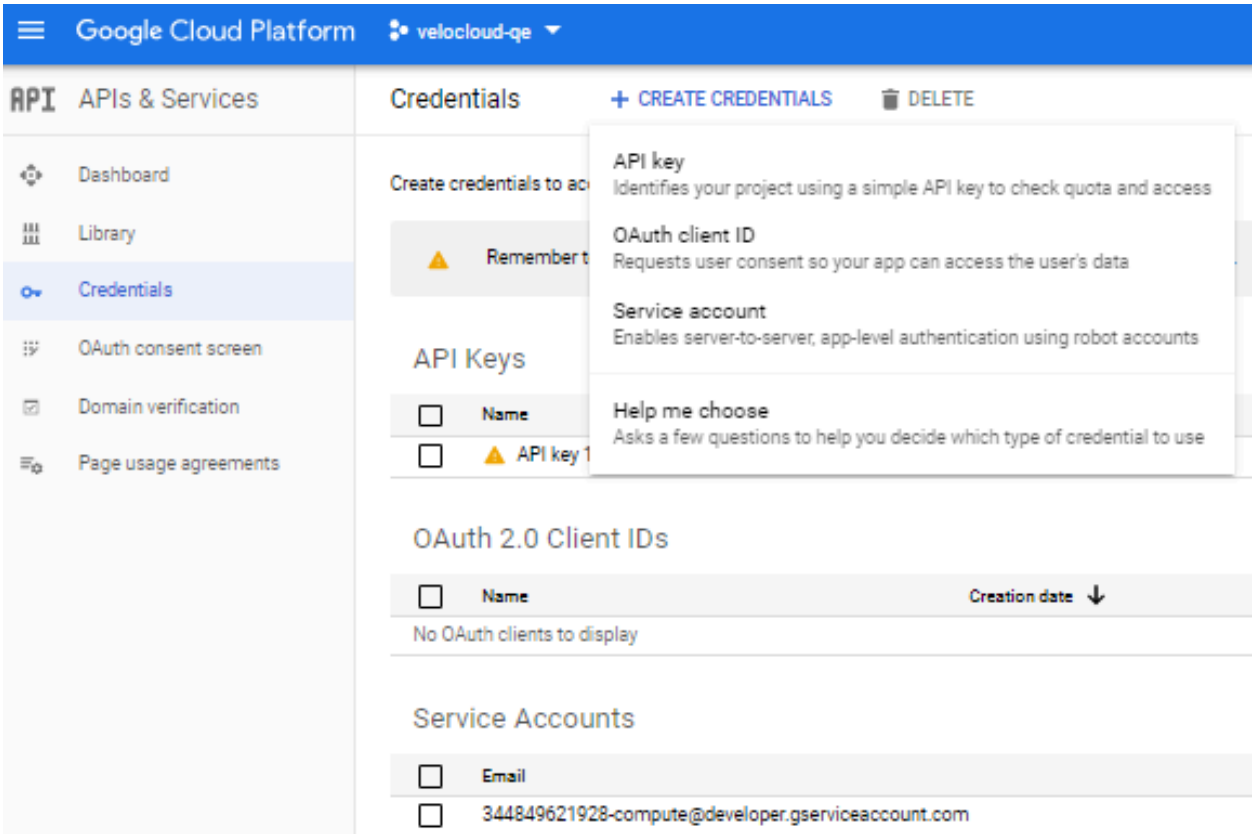

7. Selecting the **API key** will create an API key, which you can use in your application.
8. In the API key created pop-up window, select **RESTRICT KEY**, if you want to restrict your key to prevent unauthorized use in production, or else select **CLOSE**.

   The Deployment Manager and Compute Engine APIs are activated, and you can use the API to deploy your virtual edge resource.

You can deploy Virtual Edge with the Deployment Manager. For complete steps, see Deploy a Virtual Edge with GCP Deployment Manager.

Chapter 8

Verify Edge Activation

This topic describes how to verify the virtual Edge activation in Orchestrator.

1. Log into the **Orchestrator**.
2. Go to **Monitor** > **Edges**.

   This screen displays the details of the Edges, like the status, links, Gateways, and other information. You can verify whether your Virtual Edge is activated successfully.

   Figure 8-1: Verify Edge Activation

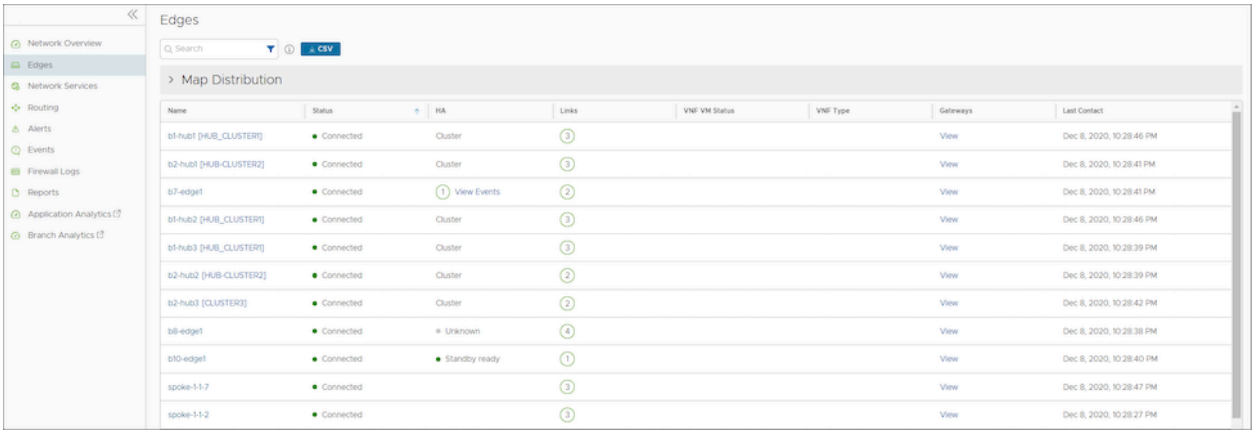

Appendix A

References

A.1 Related Documents

The following documentation is available for ***Arista VeloCloud SD-WAN***:

- *Arista VeloCloud SD-WAN Operator Guide*
- *Arista VeloCloud SD-WAN Administration Guide*
- *Arista VeloCloud SD-WAN Gateway Monitoring Guide*
- *Arista VeloCloud SD-WAN Orchestrator Deployment and Monitoring Guide*
- *Arista VeloCloud SD-WAN Partner Guide*
- *Arista VeloCloud SASE Global Settings Guide*
- *Arista VeloCloud SD-WAN Troubleshooting Guide*
- *Arista VeloCloud SD-WAN Design Guide for Enhanced Firewall Services*
- *Arista VeloCloud SD-WAN 6.4 API*
- *Arista VeloCloud Portal API 6.4*
- *Arista AliCloud Virtual Edge Deployment Guide*
- *Arista AWD Virtual Edge Deployment Guide*
- *Arista Azure Virtual Edge Deployment Guide*
- *Arista Google Cloud Platform Virtual Edge Deployment Guide*
- *Arista VeloCloud SASE and IBM QRadar SIEM Integration Guide*
- *Arista VeloCloud SD-WAN and Cloud on AWS Deployment Guide*
- *Arista VeloCloud SD-WAN and Forcepoint SSE Integration Guide*
- *Arista VeloCloud SD-WAN and Google Network Connectivity Center Integration Guide*
- *Arista VeloCloud SD-WAN and Microsoft Route Server Integration Guide*
- *Arista VeloCloud SD-WAN and Netskope SSE Integration Guide*
- *Arista VeloCloud SD-WAN Azure Private Multi-Access Edge Compute Deployment Guide*
- *Arista VeloCloud SD-WAN License Management Guide*