Deployment Guide

VeloCloud SD-WAN AliCloud 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: AliCloud Virtual Edge Deployment Guide..... 1

Chapter 2: AliCloud Virtual Edge Deployment Overview.....2

Chapter 3: Topology A - Virtual Edge Deployment on AliCloud VPC..... 3

Chapter 4: Topology B - Virtual Edge Deployment on AliCloud Single-Arm Topology.....5

Chapter 5: Create a Virtual Private Cloud..... 7

Chapter 6: Create a VSwitch.....9

Chapter 7: Create a Security Group.....11

Chapter 8: Add Security Group Rules..... 13

Chapter 9: Create Custom Route Tables and Associate VSwitches.....15

Chapter 10: Provision an Edge on the Orchestrator.....17

10.1 Create a Virtual Edge Instance on the ECS Console.....22

Chapter 11: Create an Elastic Network Interface.....26

Chapter 12: Create Elastic IP and Assign it to Public Interface of the Edge.....28

Chapter 13: Bind an ENI to an Edge instance..... 30

Chapter 14: Create a LAN Instance.....32

Chapter 15: Add a Custom Route Table Entry.....36

Chapter 16: Create a Jump Host Instance..... 37

Chapter 17: SSH Login to Edge using EIP.....41

Chapter 18: SSH to Private IP of the Edge from Jump Host.....42

Chapter 19: Activate the Edge Against the Orchestrator.....43

Appendix A: References.....44

A.1 Related Documents.....44

Chapter 1

AliCloud Virtual Edge Deployment Guide

This document provides instructions for AliCloud Virtual Edge deployment.

Chapter 2

AliCloud Virtual Edge Deployment Overview

More customers are moving workload to Public Cloud infrastructure and expect to extend VeloCloud SD-WAN™ from remote sites to public cloud to guarantee Service Level Agreements (SLAs). There are multiple options offered by Arista, leveraging distributed VeloCloud Gateways to establish IPSec towards public cloud private network or deploy virtual edge directly in AliCloud.

This document illustrates the high-level workflow of the following two topologies to deploy a virtual VeloCloud Edge (vVCE) on AliCloud:

- Topology A- Virtual Edge Deployment on AliCloud VPC Topology
- Topology B- Virtual Edge Deployment on AliCloud Single-Arm Topology

Prerequisites

- Ensure you have an Alibaba Cloud account and login information for the Alibaba Cloud Console.
- Ensure you have the VeloCloud Edge Cloud Orchestrator host name and admin account to login.

Limitations

When deploying Data Plane Development Kit (DPDK) application such as Virtual Edge on Elastic Compute Service (ECS) instances, it is recommended to choose an ECS instance types in the g5ne instance family for Edge software versions 5.4 or lower. For more information about the g5ne instance family, see Overview of Instance Families.

Deploying a DPDK application like Virtual Edge with software version 5.4 or lower on instances from the sixth-generation instance families (such as g6, c6, or r6) or later may result in exceptions during runtime. For more information on ECS instances, see ECS Instance FAQs.

Chapter 3

Topology A - Virtual Edge Deployment on AliCloud VPC

Describes the Virtual Edge deployment on the AliCloud Virtual Private Cloud (VPC) with three vSwitches, each for a subnet connected to the Edge as shown in the following topology diagram.

Figure 3-1: Topology A

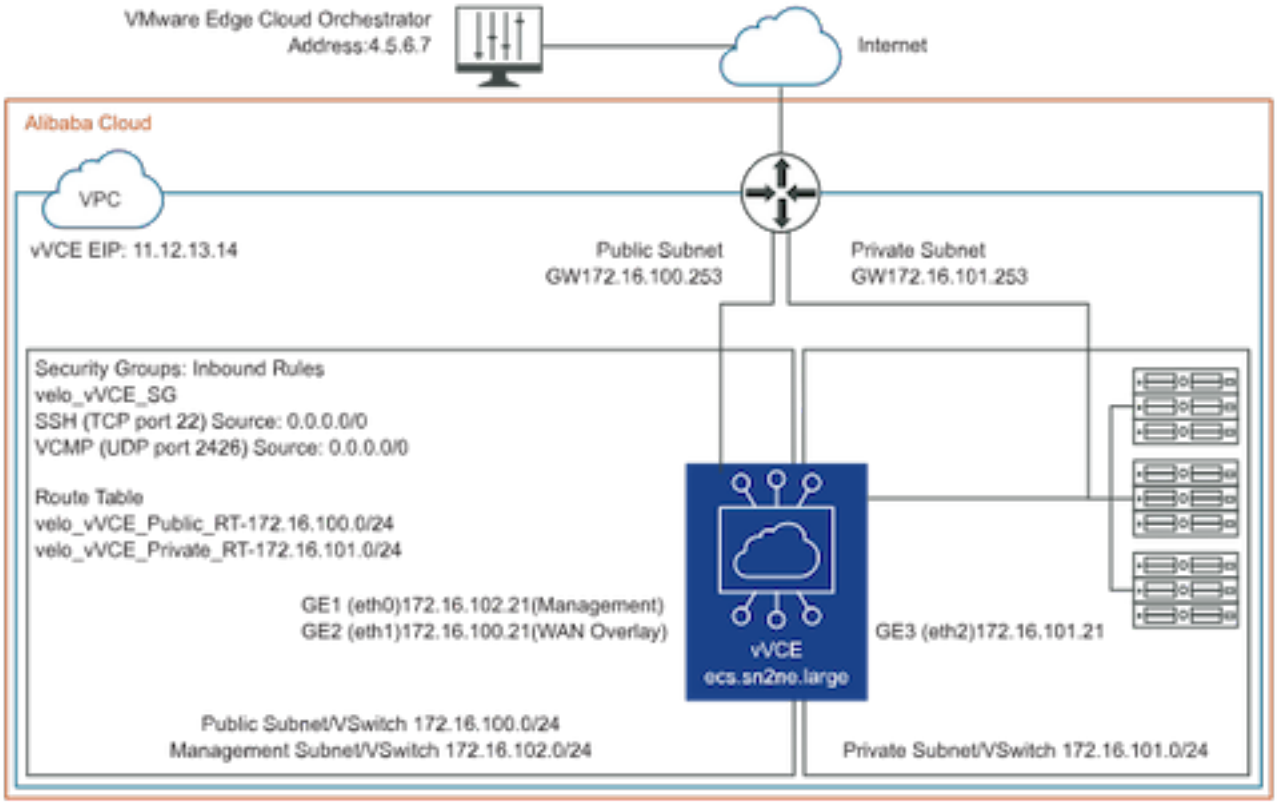

High-Level Workflow

To deploy a Virtual Edge on Alibaba Cloud ECS, perform the following steps:

1. Create a Virtual Private Cloud (VPC). For steps, see Create a VPC.
2. Create three VSwitches, each for a subnet connected to the Edge as shown in the topology diagram. For steps, see Create a VSwitch.
   - Management Subnet/VSwitch for console/management access to the Edge through Management Interface GE1.
   - Public Subnet/VSwitch for Internet access from the Edge through WAN-side Interface GE2.
   - Private Subnet/VSwitch for LAN-side device access through LAN-side Interface GE3.
3. Create a Security Group (velo_vVCE_SG) and add inbound rules. For steps, see Create a Security Group.

4. Create two custom (secondary) route tables (Velo_vVCE_Public_RT and Velo_vVCE_Private_RT) and associate it with the respective VSwitches (Public and Private). For steps, see Create Custom Route Tables and Associate VSwitches.
5. Provision a Edge on the Orchestrator as follows:
   a. Create a **Virtual Edge**.
   b. Change the GE2 interface to **Routed** from **Switched**.
   c. Deactivate **WAN Overlay** for GE3 interface and **NAT Direct Traffic**, which is the next hop for devices connected to Private Subnets (LAN devices).
   d. Add JH IP in firewall SSH access list. For more information, see Provision an Edge on the VCO.
6. Create and launch a virtual Edge (vVCE) instance with Management Interface (GE1). For steps, see Create a vVCE Instance on the ECS Console.
7. Create two Elastic Network Interfaces (ENIs) - one Private LAN-side interface (GE3) and another Public WAN-side interface (GE2). For steps, see Create an Elastic Network Interface.
8. Create an Elastic IP and assign it to the Public Interface (GE2) of the Edge. For steps, see Create Elastic IP and Assign it to Public Interface of the Edge.
9. Bind the Public (GE2) and Private (GE3) interfaces to the Edge instance (vVCE) and then restart the Edge instance to make sure the interfaces connect to the Edge. For steps, see Bind an ENI to an Edge instance.

   The Edge instance activates against the Orchestrator and the Edge establishes the VCMP tunnel to the Gateway.
10. (Optional) Within the VPC, if you want to access your Edge from a Private subnet, not over the Internet, then you have to create a Jump Host (JH) instance (Linux instance) with one interface in Public subnet for Internet connectivity with EIP and the other interface in Management subnet to access the Edge. For steps, see Create a Jump Host Instance.
    a. Create a Jump Host.
    b. Create an EIP and bind it to the Jump Host Instance.

       Note: VCAdmin users can access the Edge over Management subnet interface from JH.

    c. Login to the virtual Edge (vVCE) from Jump Host.
    d. Activate the Edge Against the Orchestrator from Shell.

       Note: After the Edge activation starts, if you want to SSH to the Edge from a Private subnet then you must add the JH IP in the firewall SSH access list.
11. Create a LAN instance with the Primary interface connected to Private subnet. For steps, see Create a LAN Instance.
    a. In the Private routing table (Velo_vVCE_Private_RT), create a new route entry that points to GE3 interface of edge for default route. For steps, see Add a Custom Route Table Entry.
12. Verify the virtual Edge became active in the Orchestrator.

Chapter 4

Topology B - Virtual Edge Deployment on AliCloud Single-Arm Topology

Discusses the Virtual Edge deployment on the AliCloud Virtual Private Cloud (VPC) with three VSwitches, each for a subnet connected to the Edge as shown in the following Single-Arm topology diagram.

Figure 4-1: Topology B

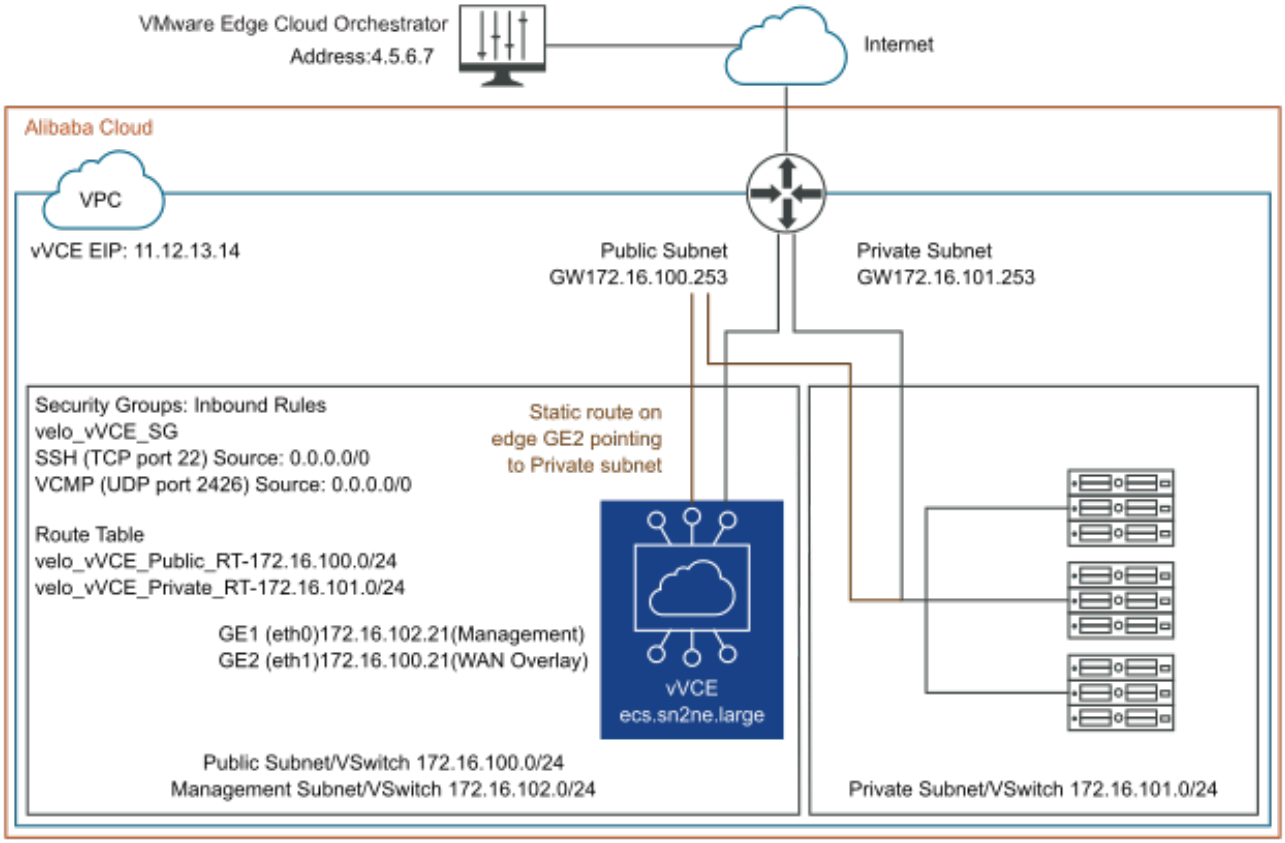

High-Level Workflow

To deploy a Virtual Edge on Alibaba Cloud ECS, perform the following steps:

1. Create a Virtual Private Cloud (VPC). For steps, see Create a VPC.
2. Create three VSwitches, each for a subnet connected to the Edge as shown in the topology diagram. For steps, see Create a VSwitch.
   - Management Subnet/VSwitch for console/management access to the Edge through Management Interface GE1.
   - Public Subnet/VSwitch for Internet access from the Edge through WAN-side Interface GE2.

- Private Subnet/VSwitch for LAN-side device access through LAN-side Interface GE3.

3. Create a Security Group (velo_vVCE_SG) and add inbound rules. For steps, see Create a Security Group.
4. Provision a Edge on the Orchestrator as follows:
   a. Create an edge of type **Virtual Edge**.
   b. Change GE2 interface to **Routed** from **Switched**.
   c. Add a static route on the Edge that points to the Private Subnet/VSwitch.
   d. Add JH IP in firewall SSH access list. For more information, see Provision an Edge on the VCO.
5. Create and launch a virtual Edge (vVCE) instance with Management Interface (GE1). For steps, see Create a vVCE Instance on the ECS Console.
6. Create a public Elastic Network Interface (GE2) on the WAN side. For steps, see Create an Elastic Network Interface.
7. Create an Elastic IP and assign it to the Public Interface (GE2) of the Edge. For steps, see Create Elastic IP and Assign it to Public Interface of the Edge.
8. Bind the Public (GE2) interface to the Edge instance (vVCE) and then restart the Edge instance to make sure the interface connects to the Edge. For steps, see Bind an ENI to an Edge instance.

   The Edge instance activates against the Orchestrator and the Edge establishes the VCMP tunnel to the Gateway.
9. Log into the Edge using the EIP and verify Edge activation.
10. Create a LAN instance with the Primary interface connected to the Management subnet. For steps, see Create a LAN Instance.
    a. In the Primary routing table, create a new route entry that points to GE2 interface of edge for default route. For steps, see Add a Custom Route Table Entry.
11. Verify the virtual Edge (vVCE) activity in the Orchestrator.

Chapter 5

Create a Virtual Private Cloud

A Virtual Private Cloud (VPC) is a virtual private network in which you can deploy your cloud resources. The Cloud resources cannot be directly deployed in a VPC. They must be deployed in a VSwitch (subnet) of the VPC.

Ensure you have an AliCloud account and login information.

1. Log on to the VPC console.
2. Select the region to create a VPC. The VPC must be in the same region as the cloud resources that you want to deploy.
3. In the left navigation pane, select **VPCs**.
4. On the **VPC** page, select **Create VPC**.

   Figure 5-1: Creating the VPC

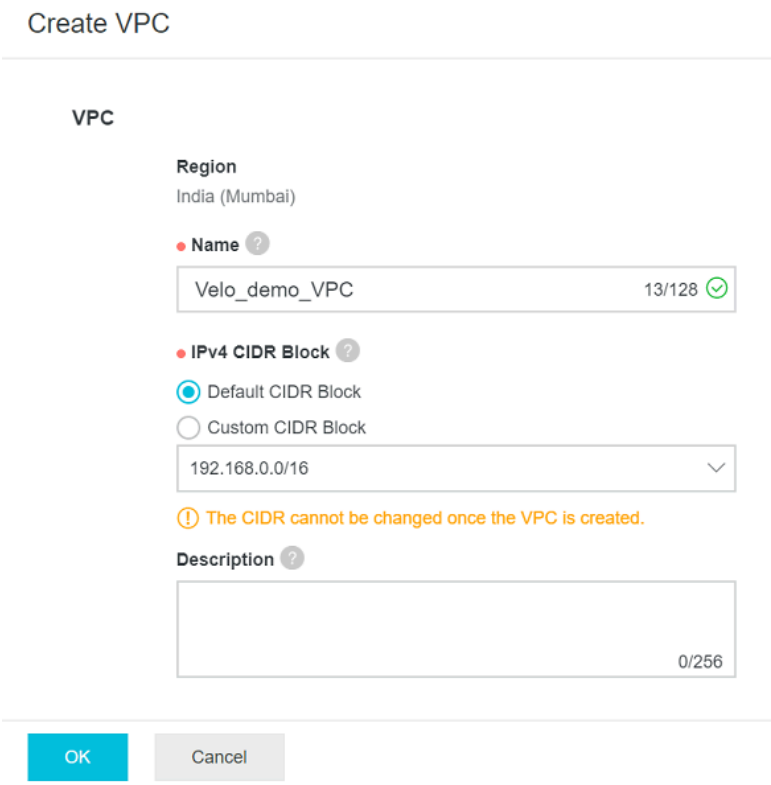

5. On the **Create VPC** page, set the following parameters, and then select **OK**.
   a. In the **Name** text box, enter the name for the VPC.
   b. Under **IPv4 CIDR Block**, select one of the following options:
      - **Default CIDR Block**: Select 192.168.0.0/16, 172.16.0.0/12, or 10.0.0.0/8.

- **Custom CIDR Block**: Select 192.168.0.0/16, 172.16.0.0/12, 10.0.0.0/8, or their subnets. The CIDR block mask must be 8 to 24 bits in length.

During VPC creation, you can create one or more VSwitches by selecting the **+ Add** button. For more information about creating a VSwitch, see Create a VSwitch.

> **Note:** After creating the VPC, you cannot change the IPv4 CIDR block.

Create a VSwitch

A VSwitch is a basic network device of a VPC and used to connect different cloud product instances. After creating a VPC, you can further segment your virtual private network to one or more subnets by creating VSwitches. The VSwitches within a VPC interconnect. Therefore, you can deploy different applications in the different VSwitches of different zones to improve the service availability.

Ensure that you have previously created a VPC.

1. Log on to the VPC console.
2. On the left-side navigation pane, select **VSwitches**.
3. Select the region of the VPC where you want to create a VSwitch.
4. Select **Create VSwitch**. The **Create VSwitch**page appears.

Figure 6-1: Creating a VSwitch

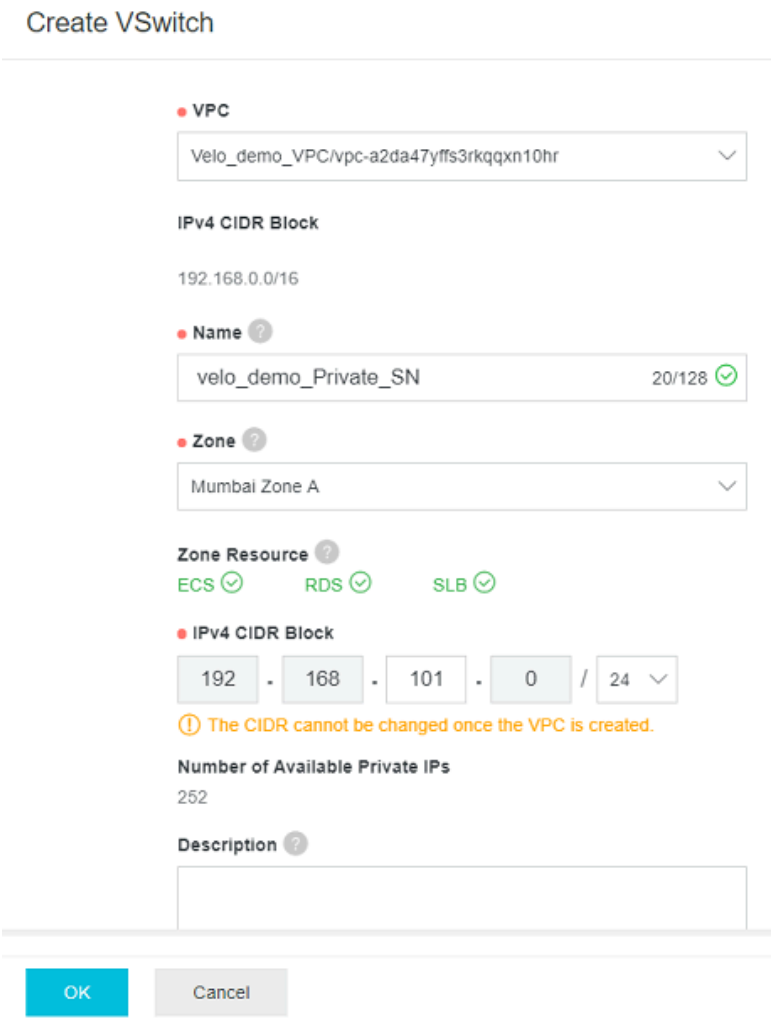

5. On the **Create VSwitch** page, set the following parameters, and then select **OK**.
   a. From the **VPC** menu, select a VPC to add the VSwitch.
   b. In **Name**, enter the name for the VPC.
   c. From the **Zone** menu, select the zone to add the VSwitch.
   d. In the **IPv4 CIDR Block**text box, enter the IPv4 CIDR block of the VSwitch.

   The IPv4 CIDR block of the VSwitch can be the same as the VPC to which the VSwitch belongs or be a subset of the VPC CIDR block.

   For example, if the CIDR block of the VPC is 192.168.0.0/16, the CIDR block of the VSwitch in the VPC can be 192.168.0.0/16, or any CIDR block between 192.168.0.0/17 and 192.168.0.0/29.

   > **Note:** If the CIDR block of the VSwitch is the same as that of the VPC, you can only create one VSwitch.

   The subnet mask of the VSwitch CIDR block can be 16 to 29 bits. It means that the VSwitch can provide 8 to 65,536 IP addresses.

   The first IP address and the last three IP addresses in the VSwitch CIDR block are reserved. For example, if the VSwitch CIDR block is 192.168.1.0/24, the IP addresses 192.168.1.0, 192.168.1.253, 192.168.1.254, and 192.168.1.255 are reserved.

   If the VSwitch needs to communicate with VSwitches of other VPCs or on-premises data centers, make sure that the CIDR blocks involved do not conflict with each other.

   > **Note:** After creating the VSwitch, you cannot change the IPv4 CIDR block.

- Create a Security Group
- Add Security Group Rules
- Create Custom Route Tables and Associate VSwitches

Chapter 7

Create a Security Group

A security group is a virtual firewall for an ECS instance. This topic describes how to create a security group in the ECS console.

If you want to create a VPC-type security group, ensure that you created a VPC and a VSwitch. For more information, see Create a Virtual Private Cloud and Create a VSwitch.

1. Log on to the ECS console.
2. In the left-side navigation pane, choose > **Network & Security** > **Security Groups**.
3. On the **Security Groups** page, select **Create Security Group**. The **Create Security Group** dialog box appears.

   Figure 7-1: Creating a Security Group

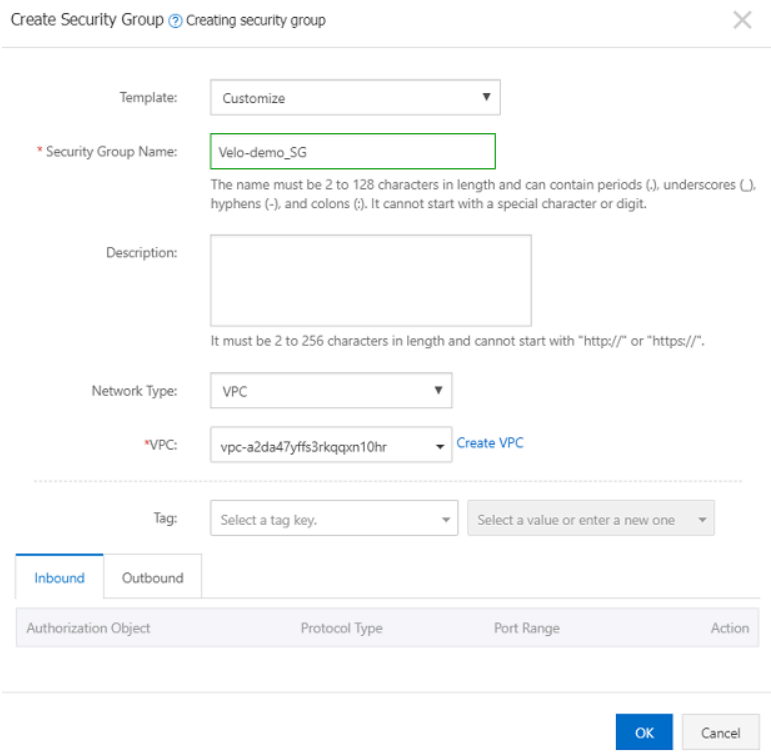

4. In **Create Security Group**, configure the following parameters:
   a. From the **Template** menu, select a suitable template to simplify security group rule configuration. For example, select **Customize**.
   b. In **Security Group Name**, enter a valid name for the security group.
   c. From the **Network Type** menu, select **VPC**.
   d. From the **VPC** menu, select your VPC.

5. Select **OK** and display a message that recommends creating security group rules.

   After creating the security group, a new security group displays in the security group list.

After creating a security group, it is recommended to immediately create security group rules. Otherwise, you may not be able to access the internal network or Internet. For steps, see Add Security Group Rules.

Chapter 8

Add Security Group Rules

You can use security group rules to control the access to public or internal networks of the ECS instances in a security group. To add security group rules, perform the steps on this procedure.

- Ensure that you have created a security group. For more information, see Create a Security Group.
- Ensure that you know which internal or public network requests need to be allowed or denied for your instance.

1. Select **Create Rules Now**. The **Security Group** page appears.
2. Select **Add Security Group Rule**. **Add Security Group Rule** appears.

   Figure 8-1: Adding a Security Group

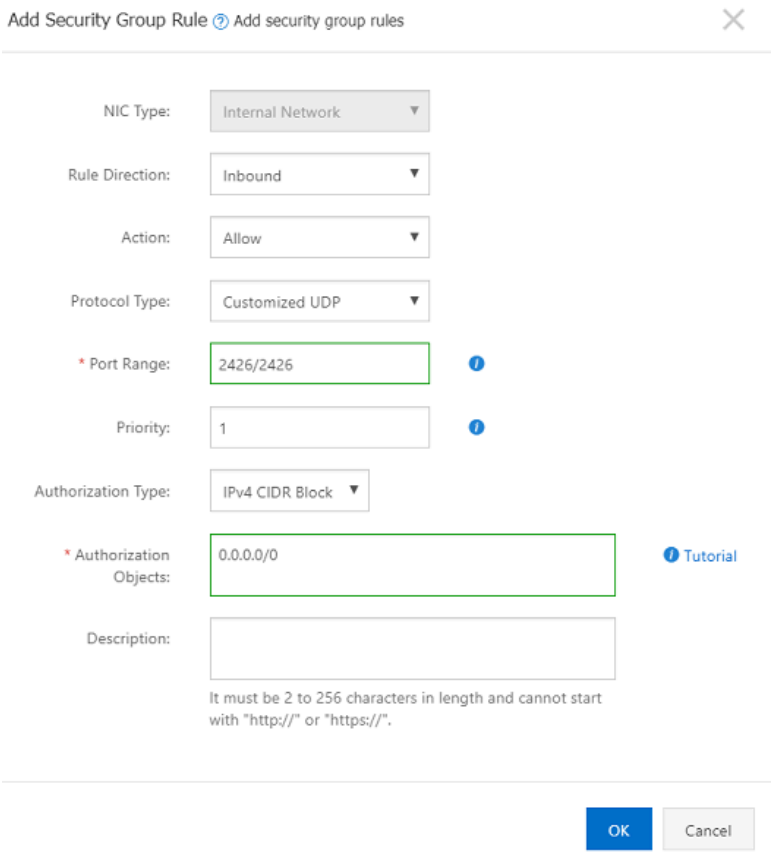

3. From the **Rule Direction** menu, select **Inbound**. By default, all Outbound traffic is allowed.
4. From the **Action** menu, select **Allow**.
5. To allow inbound connectivity to your Edge, select **Protocol Type** and **Port Range**.

   The port range is based on the protocol type. The following are some of the examples:

   - VCMP: UDP port 2426

- SSH: TCP port 22
- SNMP: UDP port 161
- ICMP Request/Reply

6. Select **Authorization Type** and **Authorization Objects**.

   The authorized IP address is based on the authorization type. For example, for IPv4 CIDR block, specifying 0.0.0.0/0 allows or denies all IP addresses, based on the authorization policy.

7. Select **OK**.

   Select the refresh icon to confirm that the addition of the security group rule. Changes to security group rules automatically apply to Elastic Compute Service (ECS) instances in the security group.

Chapter 9

Create Custom Route Tables and Associate VSwitches

A route table consists of a list of route entries in a VPC. The network traffic routes based on the configurations of the route entries in the route table. After creating a VPC, the system automatically creates a default route table and adds system routes to the route table for traffic management.

Ensure that a VPC and VSwitches have been created. For steps, see Create a Virtual Private Cloud and Create a VSwitch.

You cannot create or delete the default route table. However, you can create a custom route table and associate it with a VSwitch to control the routes of the corresponding subnet. To create a custom route table and associate it with a VSwitch, perform the following steps:

1. Log on to the VPC console.
2. In the left-side navigation pane, select **Route Tables**.
3. Select the region of the route table to create.
4. On the **Route Tables** page, select **Create Route Table**. The **Create Route Table** page appears.

   Figure 9-1: Creating a Route Table

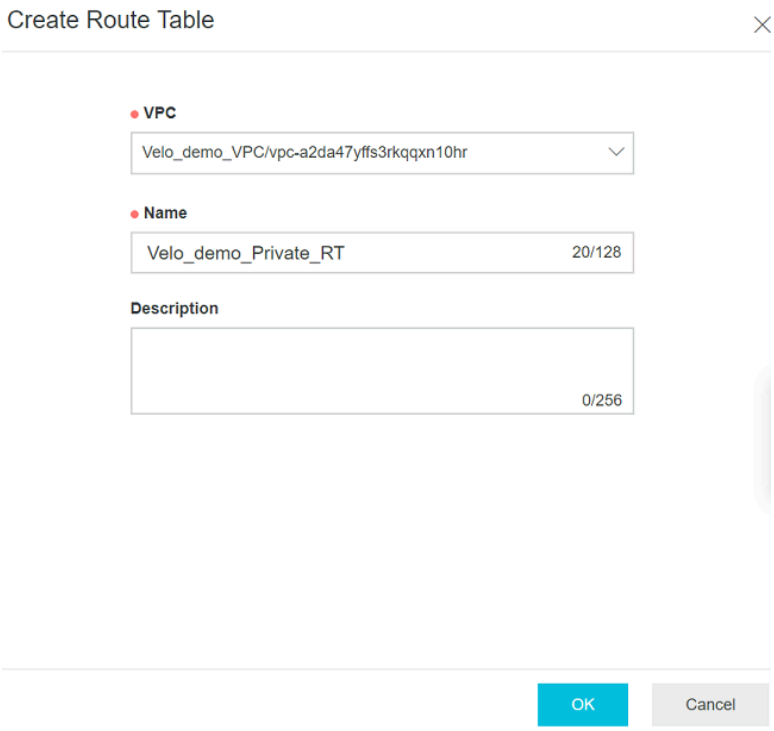

5. Configure the following parameters, and then select **OK**.
   a. From **VPC**, select your VPC.

b. In **Name**, enter the name for the routing table, and create the route table, and now appears in the list of Route Tables.

6. On the **Route Tables**page, select the route table you created and then select the **Associated VSwitches** tab.
7. Select **Associate VSwitch**. The **Associate VSwitch**page appears.

Figure 9-2: Associating a VSwitch

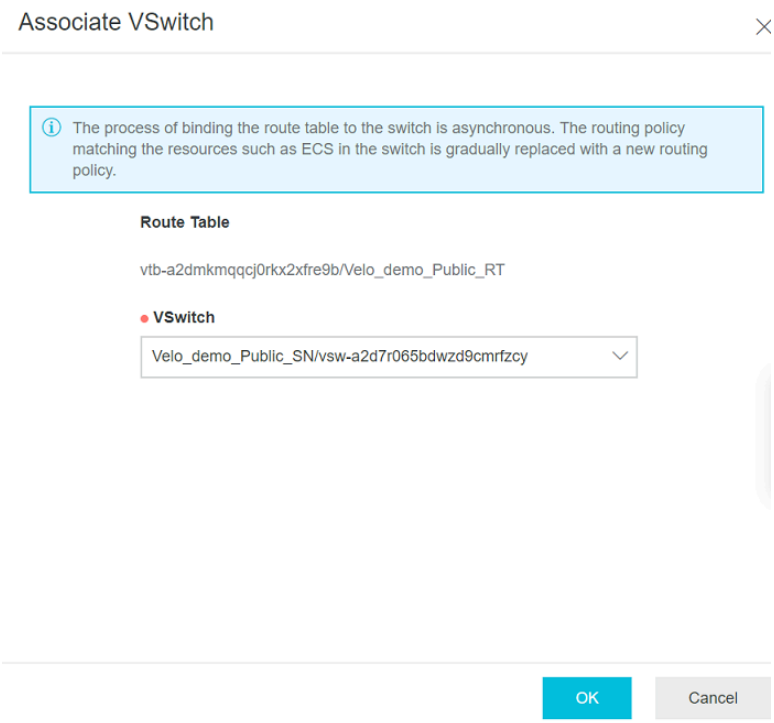

8. From the **VSwitch** menu, select the relevant VSwitch to be associated with the Route Table.
9. Select **Refresh** to confirm that the VSwitch associated with the Routing Table.

- Provision an Edge on the Orchestrator
- Create a Virtual Edge Instance on the ECS Console

Chapter 10

Provision an Edge on the Orchestrator

To provision an Edge, perform the following steps:

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

1. Log in to 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 appears.

   Figure 10-1: Provisioning an Edge

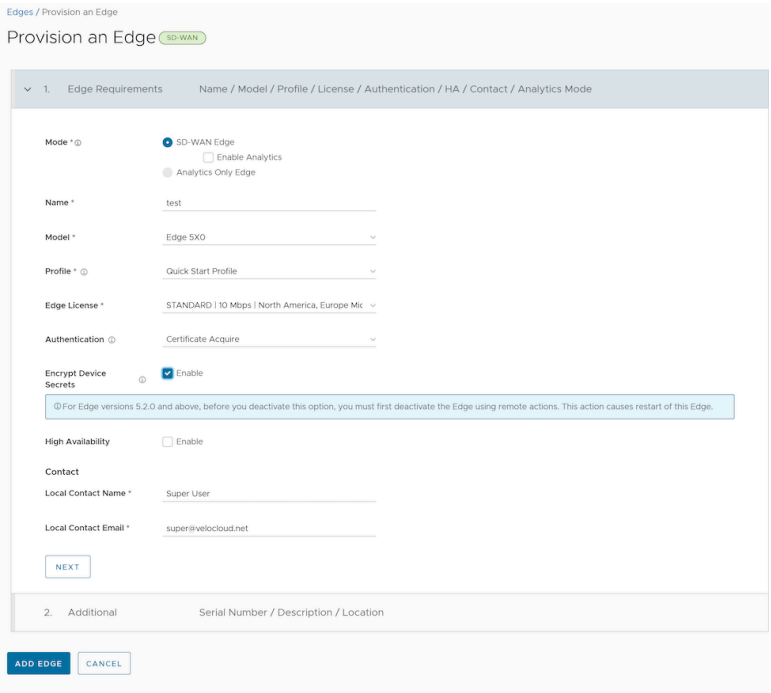

4. You can configure the following options:

Table 1: Option Descriptions

Option	Description
Mode	By default, **VeloCloud Edge** (formerly known as SD-WAN Edge) mode is selected.
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** displays 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	Choose the mode of authentication from the drop-down 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 are issued at the time of Edge activation and renewed 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 are 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:** **Next** activates only when you enter all the required details.

Table 2: 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 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 with 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 **Override Interface** check box. Select the link to the **GE2 Interface** to edit the settings.

Figure 10-2: Creating a Virtual Edge

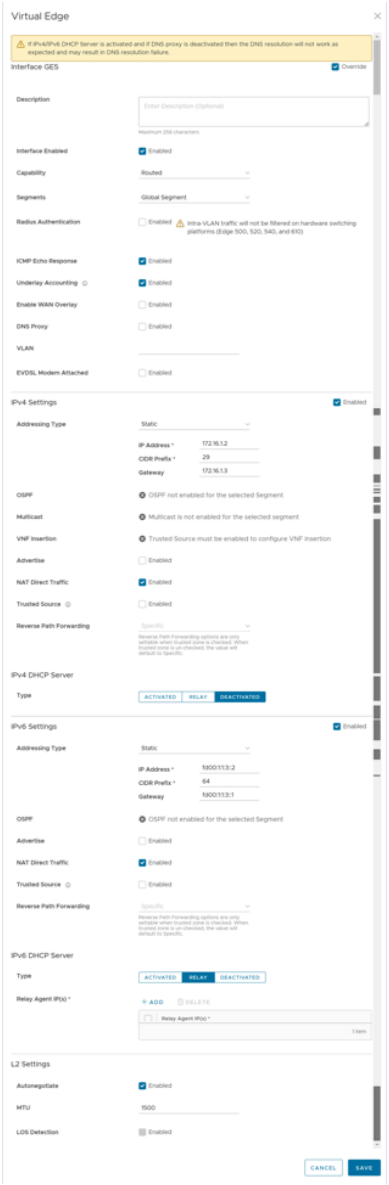

f. Now select the link to the **GE3 Interface** and select **Override Interface** check box.

g. Deactivate **WAN Overlay** and **NAT Direct Traffic**, as this interface will be used LAN-side, and select **Save**. For more information, see the topic *Configure Interface Settings for Edges* in *VeloCloud SD-WAN Administration Guide*.

> **Note:** If you are using an Edge instance with only two interfaces as illustrated in Topology B, then the public interface (GE2) is used for both WAN and LAN connectivity. For the LAN network to point to the GE2 interface, under **Static Route Settings**, configure a static route on the Edge that points to the Private Subnet/VSwitch.

8. Under the **Configure VLAN** area, edit the VLAN settings to update the **Edge LAN IP Address**.

9. (Optional) If using a Jump Host and want to allow SSH access to the Edge from the jump server, make sure to activate the **Support Access** for the Jump Host server IP in the **Firewall** page.
10. Select **Save**.

- Create a Virtual Edge Instance on the ECS Console

10.1 Create a Virtual Edge Instance on the ECS Console

Instances are the core components of Elastic Compute Service (ECS). This topic describes how to create a Pay-As-You-Go Edge instance on the ECS console.

- Ensure you have an AliCloud account and login information.

1. Log on to the ECS console.
2. In the left-side navigation pane, select **Instances & Images** > **Instances**.
3. On the **Instances** page, select **Create Instance**. The **Custom Launch** purchase page appears.
4. Set up **Basic Configurations** by performing the following steps.
   a. Select a billing method. For example, **Pay-As-You-Go**.
   b. From **Region** menu, select a region. The system randomly allocates a zone by default.

      Note: After an instance is created, you cannot change its region or zone.

   c. In the **Instance Type**area, go to **All Generations** > **x86-Architecture** > **General Purpose** and select the **4 vCPU + 16 GiB** memory (ecs. sn2ne.xlarge) instance type. The selected region determines the available instance type family.
   d. In the **Image** area, select **Custom Image** and select an Edge image, for example, *vce-3.3.1-48-R331-20190727-QE*.

e. Select a storage space. By default, a 40 GiB Ultra Cloud Disk is selected.

Figure 10-3: Basic Configuration

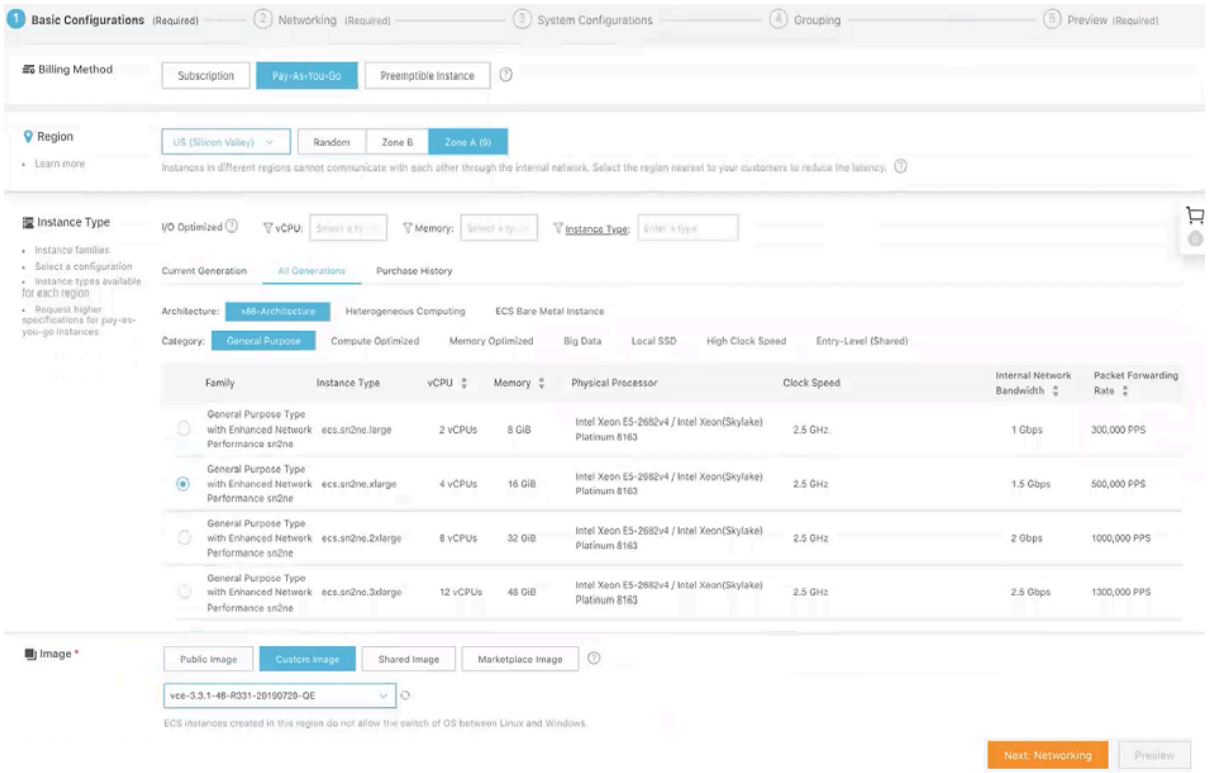

5. Select **Next: Networking** to set up the networking and security group configuration.

a. Select **VPC** as the network type and select the VPC where you are going to deploy your Edge and attach the Console interface of your Edge to MGMT_SN.

b. Set the network billing method.

c. Select your VPC-type security group.

d. Add an Elastic Network Interface (ENI). You can skip this step if the selected instance type does not support ENI.

Figure 10-4: Adding an ENI

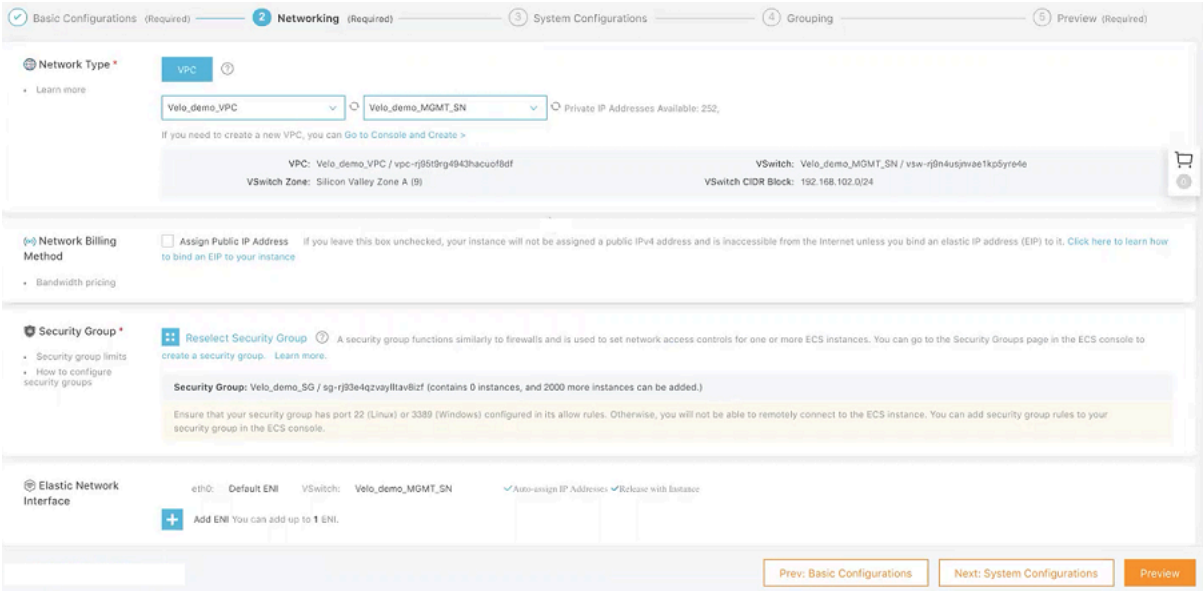

6. Select **Next: System Configurations**.

   a. Configure **Logon Credentials**by selecting one of the following options: **Key Pair**, **Inherit Password From Image**, and **Password**. By default, **Set Later**option is selected.

   b. In the **Instance Name** text box, enter a unique name for your Edge instance.

c. Under the **Advance** area, you can provide the `cloud-init` user data for your edge in the following sample format for activation purpose. According to your Orchestrator set up, you must change the orchestrator name and activation code.

```
#cloud-config velocloud: vce: vco: 1.211.224.11 activation_code: 12XX-ABC1-6DD3-3EFG
 vco_ignore_cert_errors: true
```

Figure 10-5: System Configuration

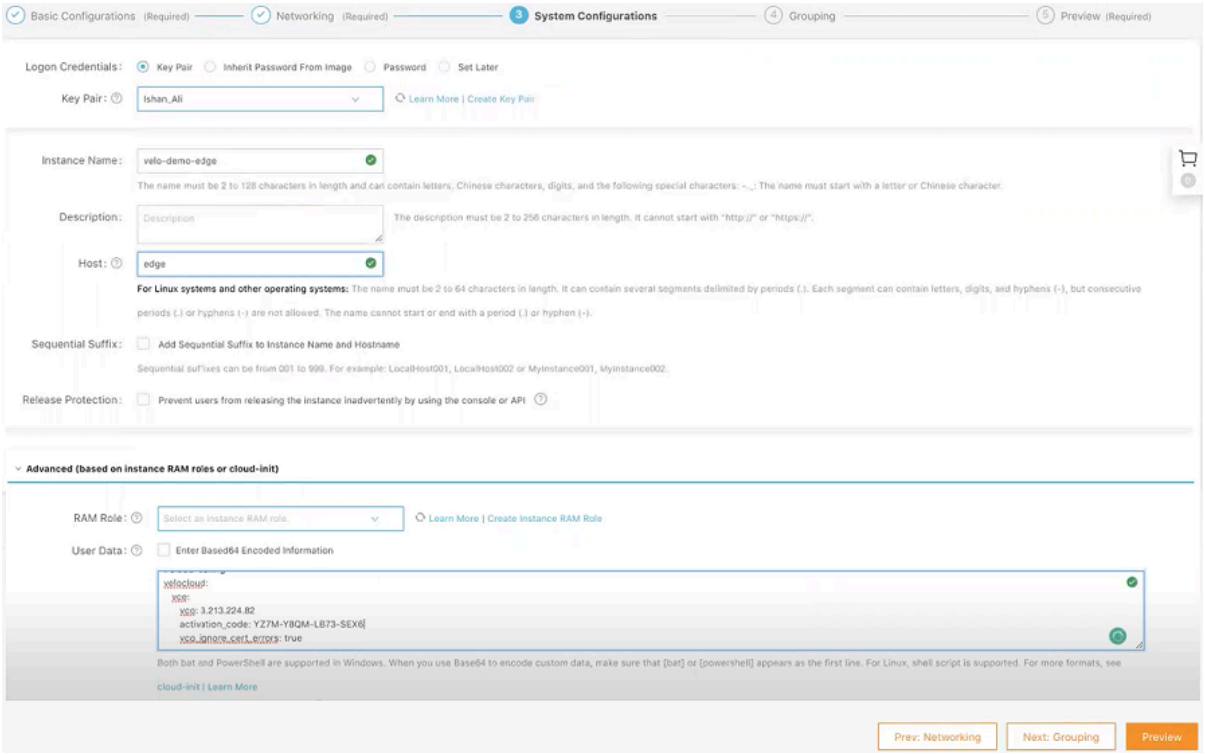

7. Select **Next: Grouping** and set the options as needed.
8. Select **Next: Preview** and confirm the selected configuration. You can also select the edit icon to modify the configurations.
9. Read and confirm **Terms of Service**, and then select **Create Instance**.
10. Select **Console** to return to the ECS console. Select the refresh button to check if the Edge instance is created. If the newly created Edge instance is in a **Running** status, then the Edge is created successfully.

- Create an Elastic Network Interface

Chapter 11

Create an Elastic Network Interface

An Elastic Network Interface (ENI) is a virtual network interface that can be attached to an ECS instance in a VPC. This topic describes how to create an ENI in the ECS console.

- A VPC and a VSwitch are created. For steps, see Create a Virtual Private Cloud.
- A security group is created in the same VPC. For steps, see Create a Security Group.

1. Log on to the ECS console.
2. In the left-side navigation pane, select **Network & Security** to display **Network Interfaces**.
3. Select **Create ENI** to display **Create ENI** page.

   Figure 11-1: Creating ENI

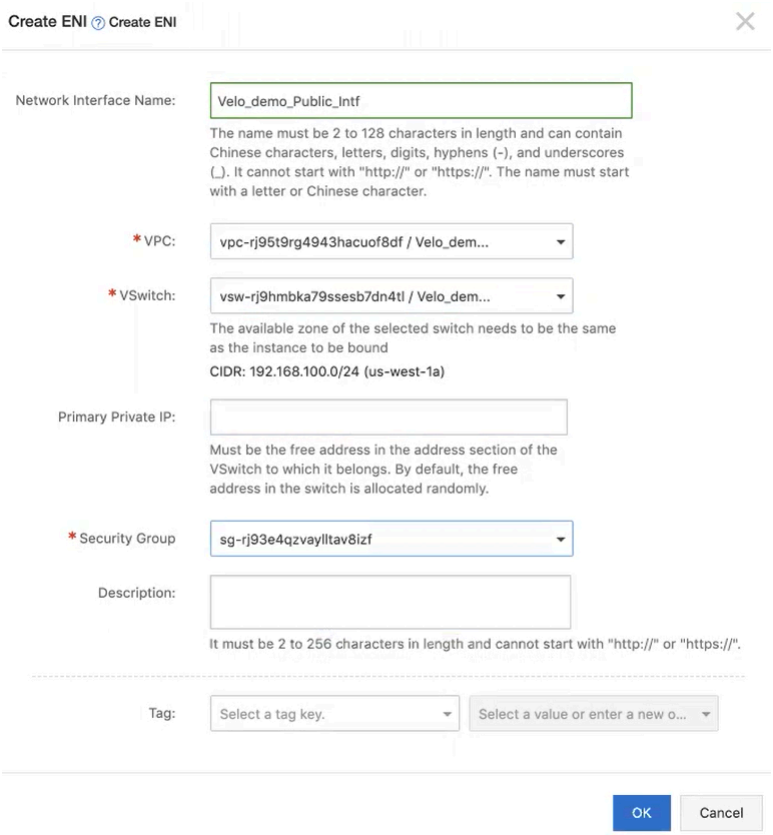

4. Provide the following configuration details:
   a. In **Network Interface Name**, enter a unique name for the ENI.
   b. From the **VPC** menu, select the same VPC associated with the instance to which you want the ENI to be bound. When you attach an ENI to an instance, they must be in the same VPC.

> **Note:** After creating an ENI, you cannot change the VPC.

c. From the **VSwitch** menu, select the VSwitch to bind to the ENI. When you attach an ENI to an instance, they must be in the same zone, but they do not have to be in the same VSwitch.

> **Note:** After creating an ENI, you cannot change the VSwitch.

d. From the **Security Group** menu, select your security group in the selected VPC.

5. Select **OK**.
6. On the **Network Interfaces** page, select **Refresh** to view the newly created ENI instance.

- After creating an ENI, you can attach an ENI to an Edge instance. For steps, see Bind an ENI to an Edge instance.

Chapter 12

Create Elastic IP and Assign it to Public Interface of the Edge

Elastic IP Addresses (EIPs) are public IP address resources that you can purchase and hold independently. You can create an EIP or reinstate a released EIP through the console. This topic describes how to create an EIP and bind it to the secondary (public) interface.

- Ensure you have an AliCloud account and login information.
- Ensure you have an Elastic Network Interface (ENI) to assign the Elastic IP. For steps to create ENI, see Create an Elastic Network Interface.

1. Log on to the VPC console.
2. In the left-side navigation pane, select **Elastic IP Addresses**. The **Elastic IP Addresses** page appears.
3. Select **Create EIP**. The **Elastic IP** page appears.
4. Configure EIP.
   a. Select the region of the EIP to be created. Ensure that the EIP and the cloud instance to be associated with the EIP belong to the same region.
   b. Set the maximum bandwidth for the EIP to be created, depending on the requirement.
   c. Select the number of the EIPs that you want to create with the same configurations. The default value for **Quantity** is 1.
5. Select **Buy Now**. The **Confirm Order**page appears.
6. Select **Elastic IP Agreement of Service** check box and select **Activate**. On the **Elastic IP Addresses** page, select **Refresh** to view the newly created EIP instance.

7. To associate the EIP instance to a secondary (public) Elastic Network Interface (ENI) of the Edge, select **Bind** under the **Actions** column. The **Bind Elastic IP Address** page appears.

   Figure 12-1: Binding an Elastic IP Address

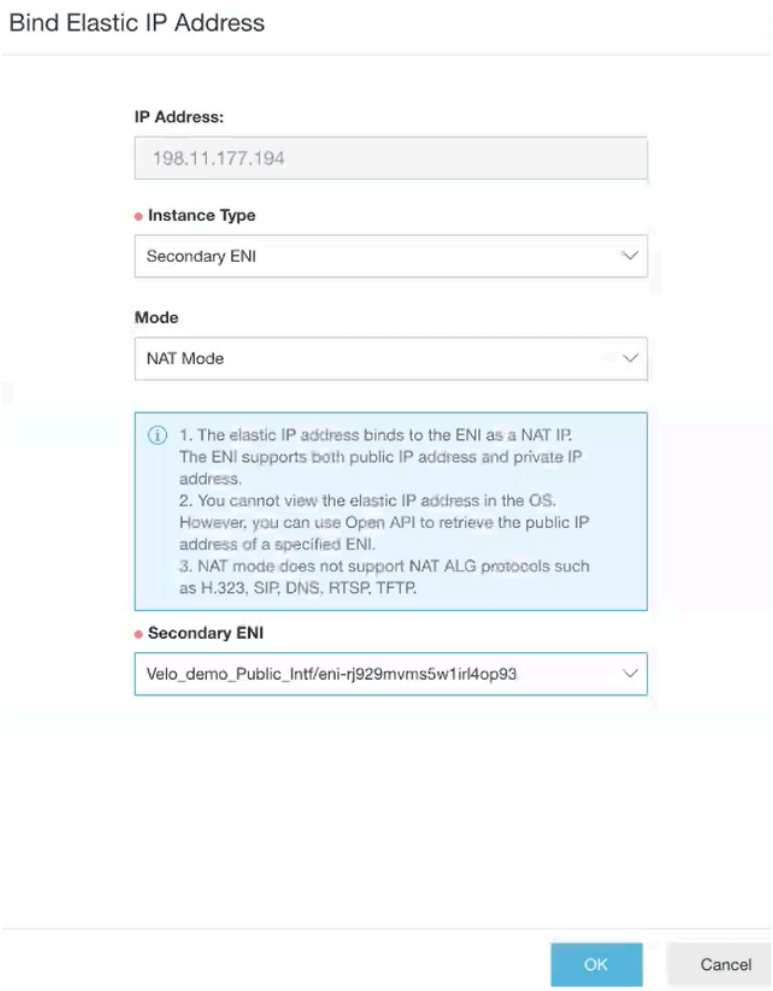

8. From the **Instance Type** menu, select **secondary ENI**.
9. From the **Secondary ENI** menu, select the interface to which you want to bind the EIP.
10. Select **OK**.

- Bind an ENI to an Edge instance

Chapter 13

Bind an ENI to an Edge instance

To bind secondary Elastic Network Interfaces (ENIs) to an Edge instance, perform the steps on this procedure.

- Ensure that you have created an ENI. For steps. see Create an Elastic Network Interface.

1. Log on to the ECS console.
2. In the left-side navigation pane, select **Network & Security** > **ENI**.
3. On the **Network Interfaces** page, select an ENI and select **Bind to Instance** under the **Actions** column. The **Bind to Instance**pop-up window appears.

   Figure 13-1: Binding an ENI to an Instance

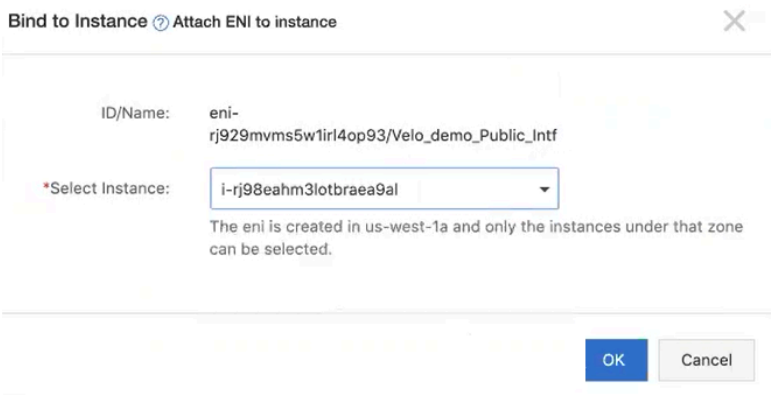

4. From **Select Instance** menu, select the Edge instance to bind the ENI.
5. Select **OK**.

   The selected ENI binds to the Edge instance.
6. Restart your Edge instance for the newly associated interfaces to take effect, by performing the following steps:

   a. In the left-side navigation pane, select **Instances**.

   b. On the **Instances** page, go to your Edge instance, and select **More** under the **Actions** column.

   c. Go to **Instance Status** > **Restart**.

d. Select **OK**.

Figure 13-2: Instances

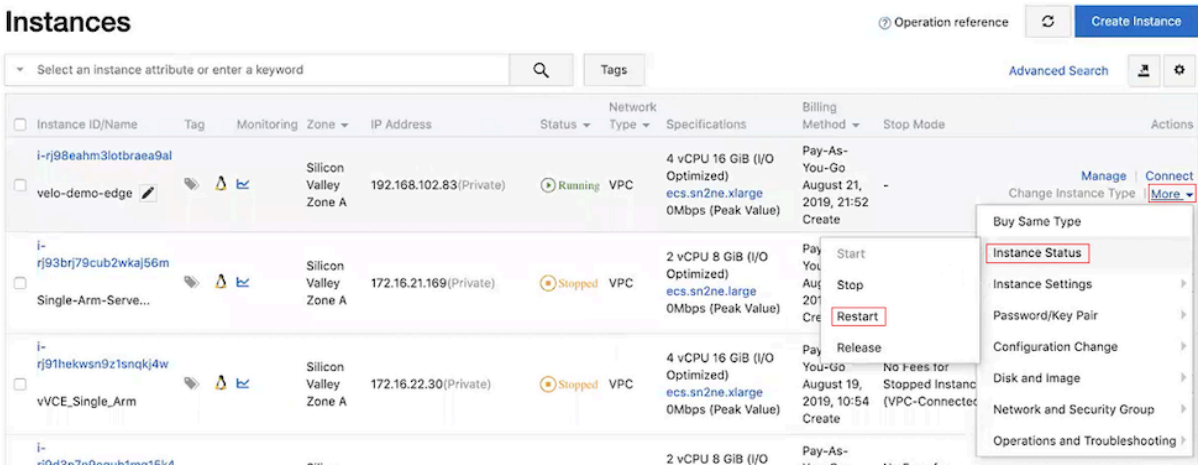

Chapter 14

Create a LAN Instance

Describes how to create a LAN (Linux) instance on the ECS console.

Ensure you have an AliCloud account and login information.

1. Log on to the ECS console.
2. In the left-side navigation pane, select **Instances & Images** > **Instances**.
3. On the **Instances** page, select **Create Instance**. The **Custom Launch** purchase page appears.
4. Set up **Basic Configurations** by performing the following steps.
   a. Select a billing method. For example, **Pay-As-You-Go**.
   b. From the **Region** drop-down menu, select a region. The system randomly allocates a zone by default.

      > **Note:** After creating an instance, you cannot change the region or zone.

   c. In the **Instance Type**area, go to **All Generations** > **x86-Architecture** > **General Purpose** and select an instance type. The available instance type family is determined by the region you selected.
   d. In the **Image** area, select **Public Image** and select an Ubuntu image, for example, 16.04.64.bit.

e. Select a storage space. By default, a 40 GiB Ultra Cloud Disk is selected.

Figure 14-1: Configuring the Image

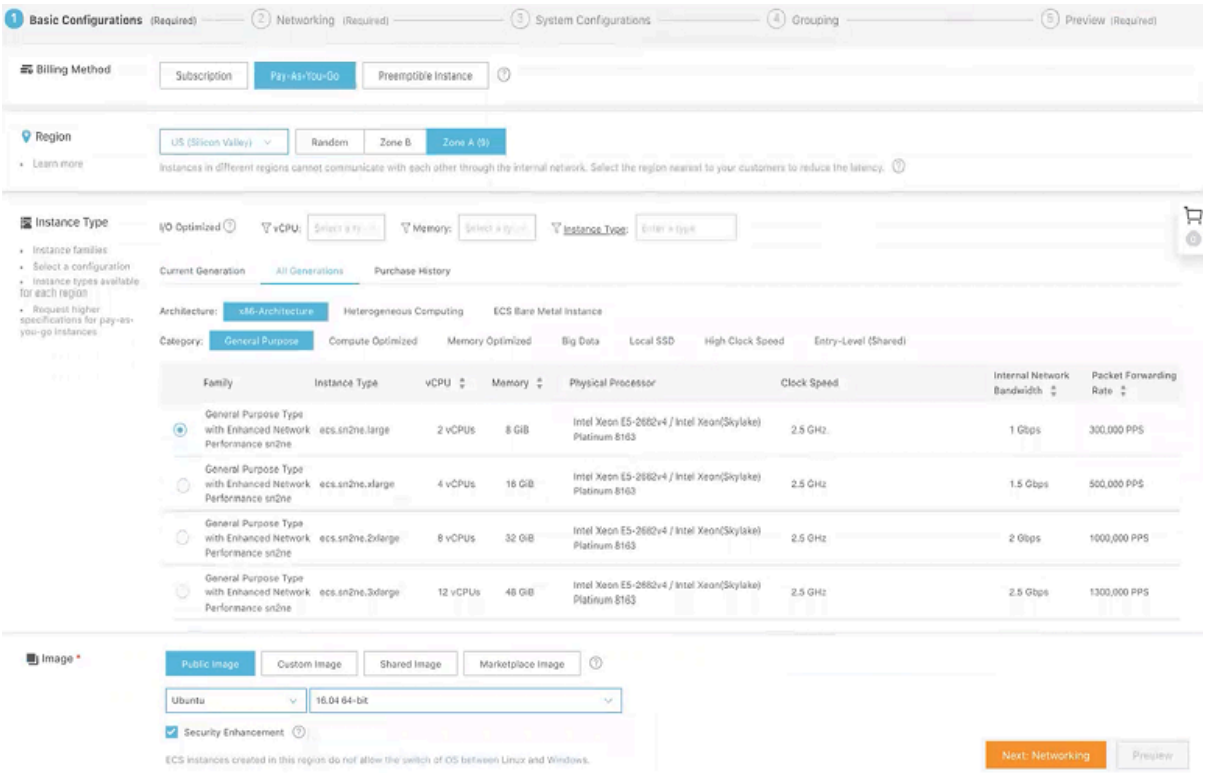

5. Select **Next: Networking** to set up the networking and security group configuration.

a. Select **VPC** as the network type and select the VPC where you are going to deploy your LAN device and the Subnet/VSwitch to which you want the primary interface to be connected.

b. Set the network billing method.

c. Select your VPC-type security group.

d. Add an Elastic Network Interface (ENI). You can skip this step if the selected instance type does not support ENI.

Figure 14-2: Adding an ENI

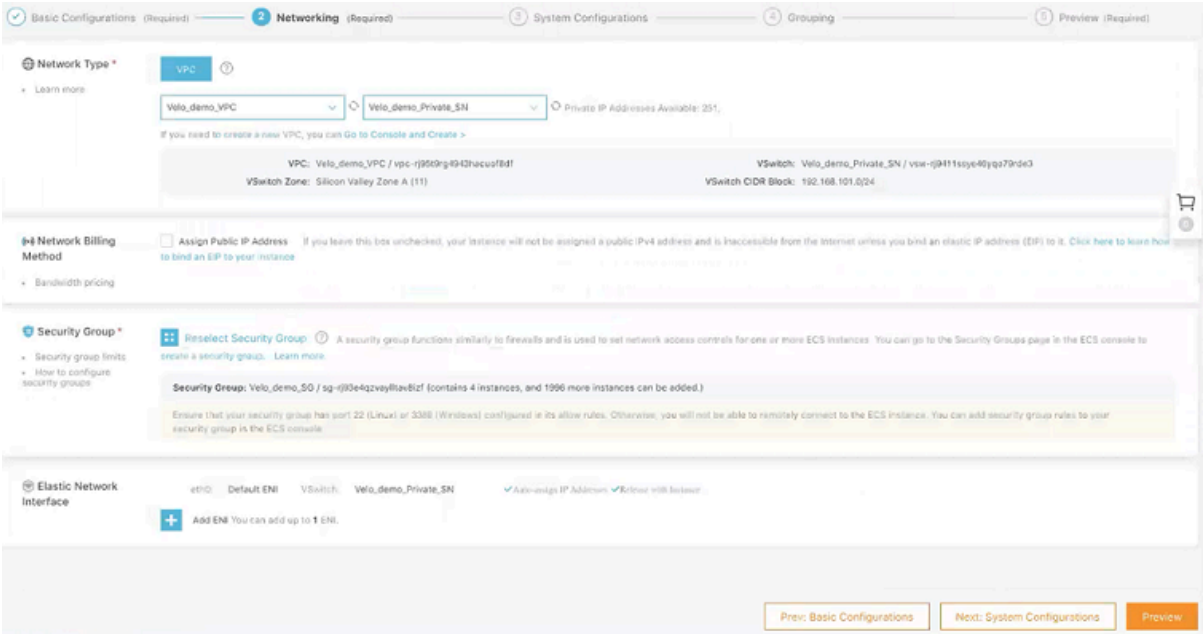

6. Select **Next: System Configurations**.

   a. Configure **Logon Credentials**by selecting one of the following options: **Key Pair**, **Inherit Password From Image**, and **Password**. By default, **Set Later** option is selected.

   b. In the **Instance Name** text box, enter a unique name for the LAN instance.

   Figure 14-3: Instance Name

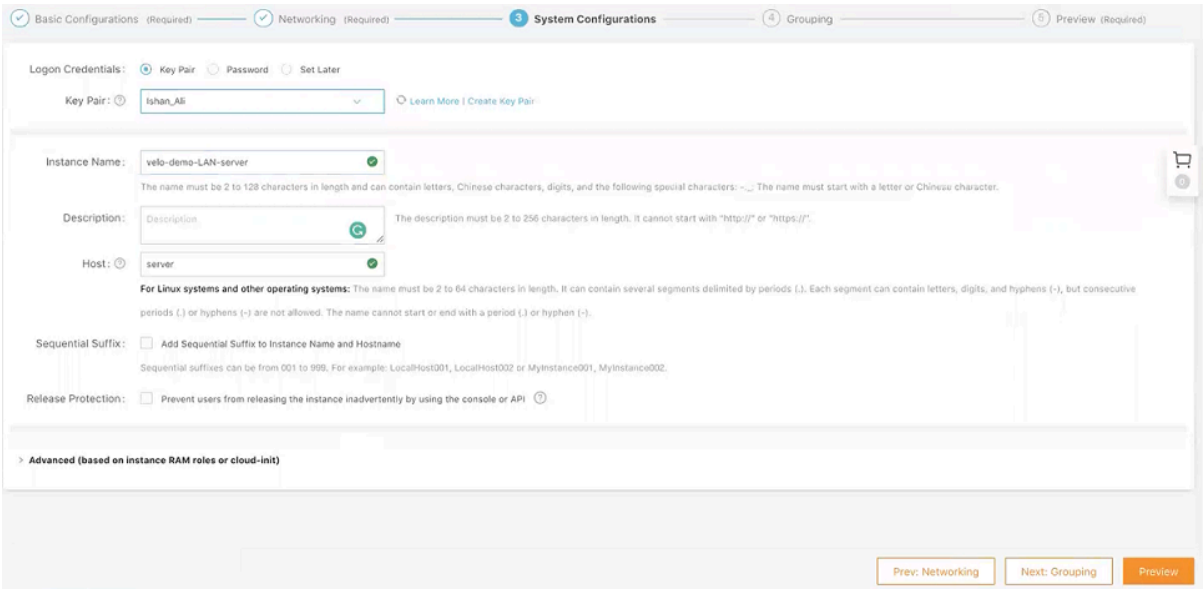

7. Select **Next: Grouping** and set the options as needed.

8. Select **Next: Preview** and confirm the selected configuration. You can also select the edit icon to modify the configurations.
9. Read and confirm **Terms of Service**, and then select **Create Instance**.
10. Select **Console** to return to the ECS console. Select the refresh button to check if the LAN instance is created. If the newly created LAN instance has a **Running** status, then the instance created successfully.

Add a Custom Route Table Entry

Describes how to add a custom route entry in a custom route table.

Ensure you have created a VPC and VSwitches. For steps, see Create a Virtual Private Cloud and Create a VSwitch.

Ensure you have a custom route table associated with a VSwitch. For steps, see Create Custom Route Tables and Associate VSwitches.

1. Log on to the VPC console.
2. In the left-side navigation pane, select **Route Tables**.
3. On the **Route Tables** page, find the target route table, and then select **Manage** in the **Actions** column.
4. In the **Route Table Details** area, select **Route Entry List** tab, and then select **Add Route Entry**. The **Add Route Entry** page appears.

   Figure 15-1: Add Route Entry

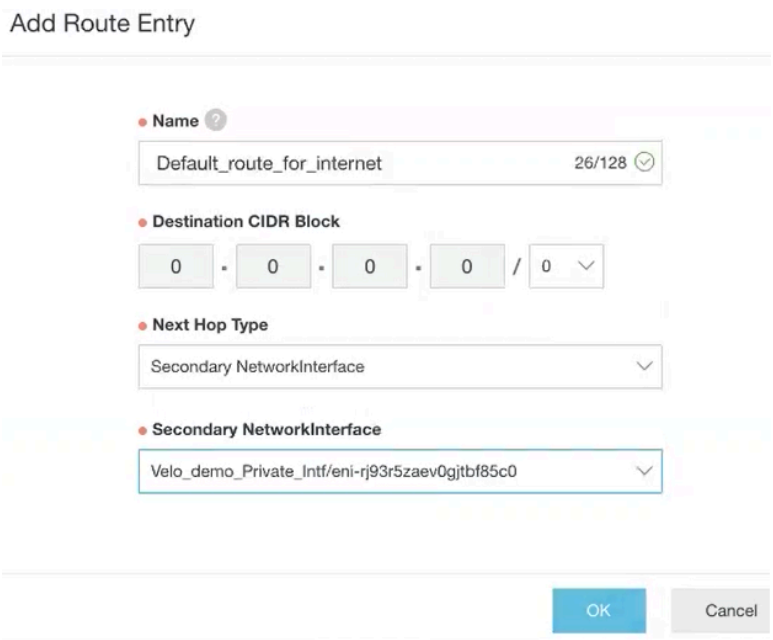

5. In the **Name** text box, enter the unique name for the route entry.
6. In **Destination CIDR Block**, specify the destination CIDR block address.
7. From the **Next Hop Type** drop-down menu, select the next hop interface type, for example, **Secondary Network Interface**, and then select the interface.
8. Select **OK**.

Chapter 16

Create a Jump Host Instance

Creating a Jump Host (JH) instance is an optional step for the Edge deployment. However, to locally manage the virtual Edge, you must deploy a JH and assign an Elastic IP to it. To SSH to an Edge over a private network using a JH, create a JH (Linux instance) in VPC with one interface in the Public subnet (for Internet connectivity with EIP), and the other interface in the management subnet.

Ensure you have an AliCloud account and login information.

To create a JH instance on the ECS console, perform the following steps.

1. Log on to the ECS console.
2. In the left-side navigation pane, select **Instances & Images** > **Instances**.
3. On the **Instances** page, select **Create Instance**. The **Custom Launch** purchase page appears.
4. Set up **Basic Configurations** by performing the following steps.
   a. Select a billing method. For example, **Pay-As-You-Go**.
   b. From the **Region** drop-down menu, select a region. The system randomly allocates a zone by default.

      Note: After creating an instance is created, you cannot change the region or zone.

   c. In the **Instance Type**area, go to **All Generations** > **x86-Architecture** > **General Purpose** and select an instance type. The available instance type family is determined by the region you selected.
   d. In the **Image** area, select **Public Image** and select an Ubuntu image, for example, 16.04.64.bit.

e. Select a storage space. By default, a 40 GiB Ultra Cloud Disk is selected.

Figure 16-1: Configuring the Image

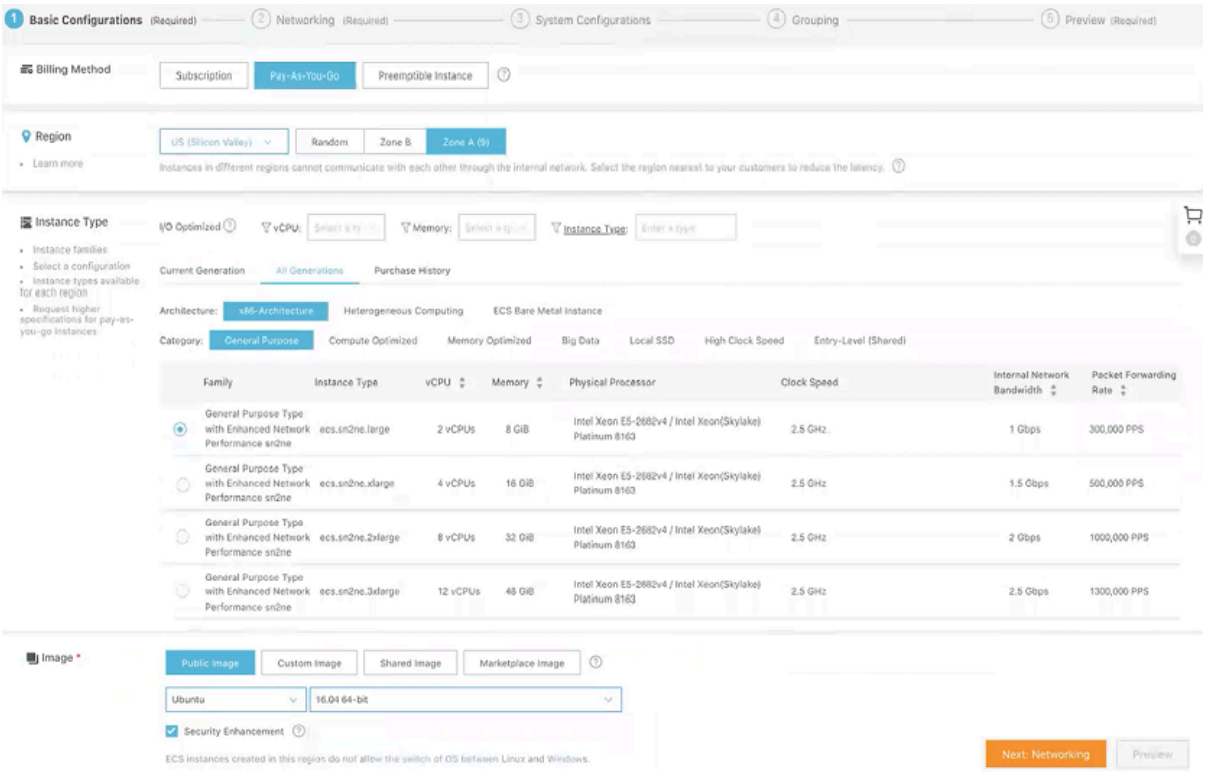

5. Select **Next: Networking** to set up the networking and security group configuration.

   a. Select **VPC** as the network type and select the VPC where you are going to deploy your JH and attach the Console interface of your Edge to MGMT_SN.

   b. Set the network billing method.

   c. Select your VPC-type security group.

d. Add an Elastic Network Interface (ENI). You can skip this step if the selected instance type does not support ENI.

Figure 16-2: Adding an ENI

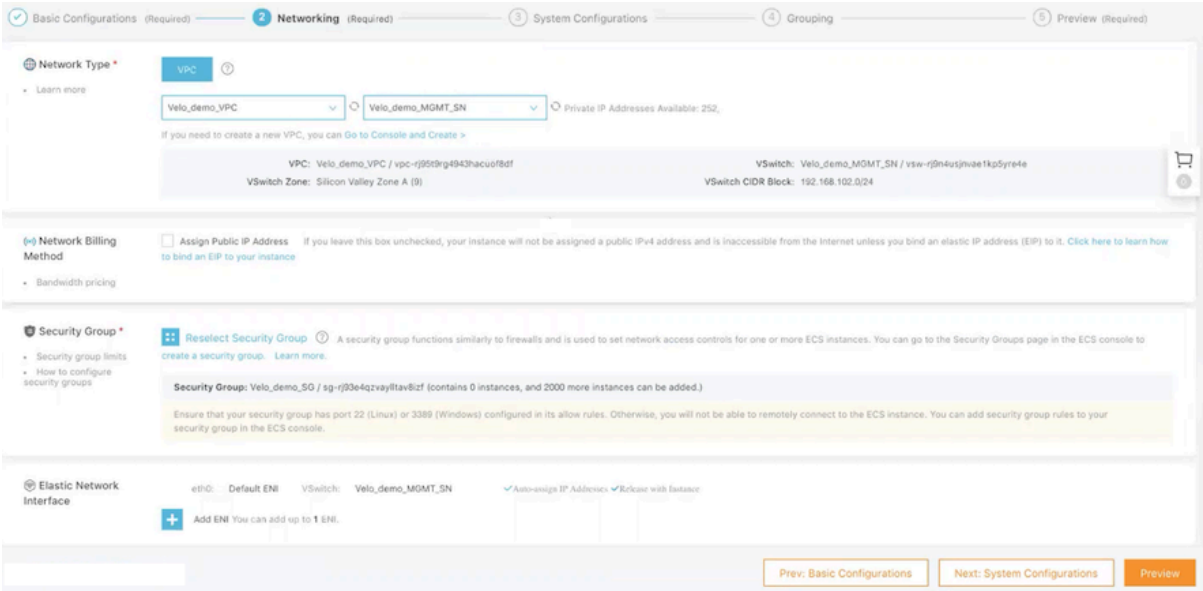

6. Select **Next: System Configurations**.

a. Configure **Logon Credentials**by selecting one of the following options: **Key Pair**, **Inherit Password From Image**, and **Password**. By default, **Set Later** option is selected.

b. In the **Instance Name** text box, enter a unique name for the JH instance.

Figure 16-3: Adding Logon Credentials

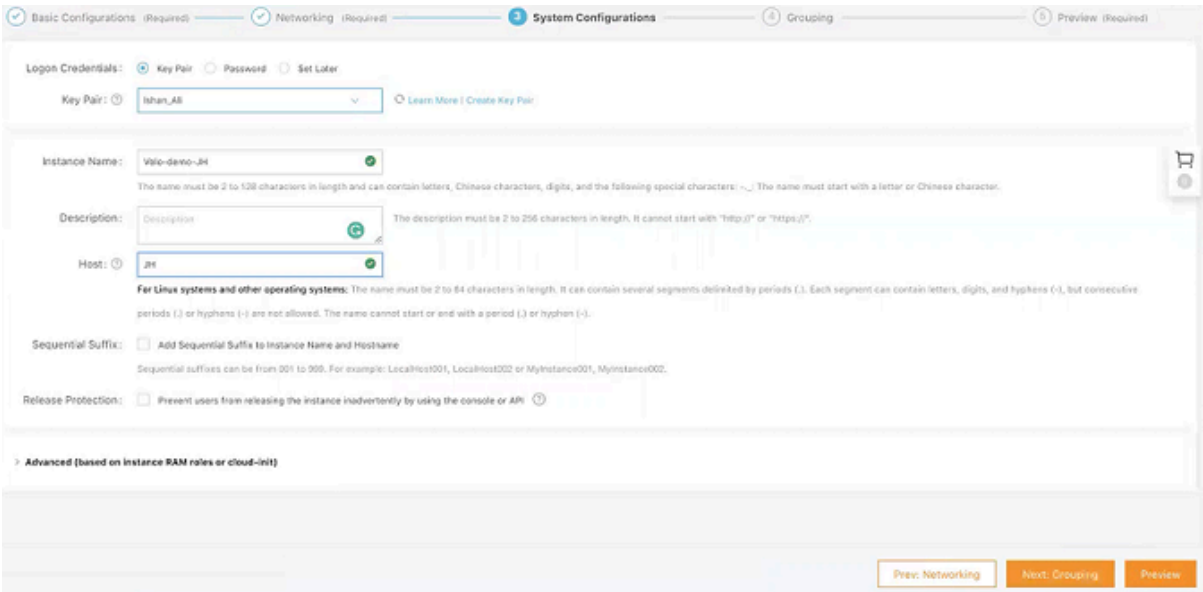

7. Select **Next: Grouping** and set options as needed.

8. Select **Next: Preview** and confirm the selected configuration. You can also select the edit icon to modify the configurations.

9. Read and confirm **Terms of Service**, and then select **Create Instance**.
10. Select **Console** to return to the ECS console. Select the refresh button to check if the JH instance is created. If the newly created JH instance is in a **Running** status, then the instance is created successfully.

SSH Login to Edge using EIP

To use SSH to log into an edge using EIP and verify activation, enter the following command.

```
ssh-i private key vcadmin@ Elastic IP of the edge
```

Where:

- *private key* - Use the private key generated in Create a Virtual Edge Instance on the ECS Console.
- *Elastic IP of the edge* - Use the Elastic IP created in Create Elastic IP and Assign it to Public Interface of the Edge.

Figure 17-1: SSH to an Edge

SSH to Private IP of the Edge from Jump Host

This is an optional step. To SSH to Private IP of the Edge from Jump Host (JH), enter the following command.

```
ssh-i private key vcadmin@ Private IP of the edge on interface eth0
```

The Private IP of interface eth0 can be derived from AliCloud console.

Figure 18-1: SSH to a Private IP Address of the Edge

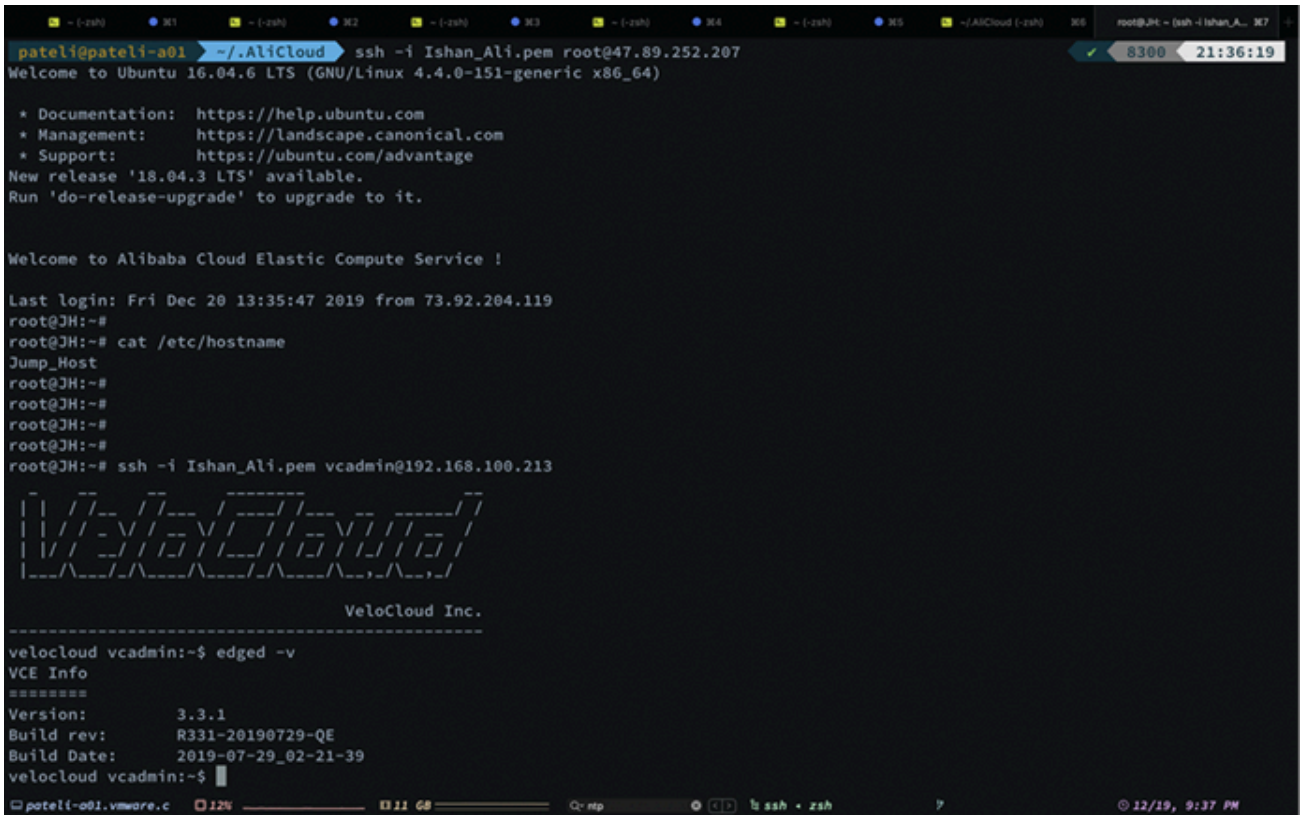

Activate the Edge Against the Orchestrator

Activation occurs on step 6 of Create a Virtual Edge Instance on the ECS Console.

Figure 19-1: Activating the Edge

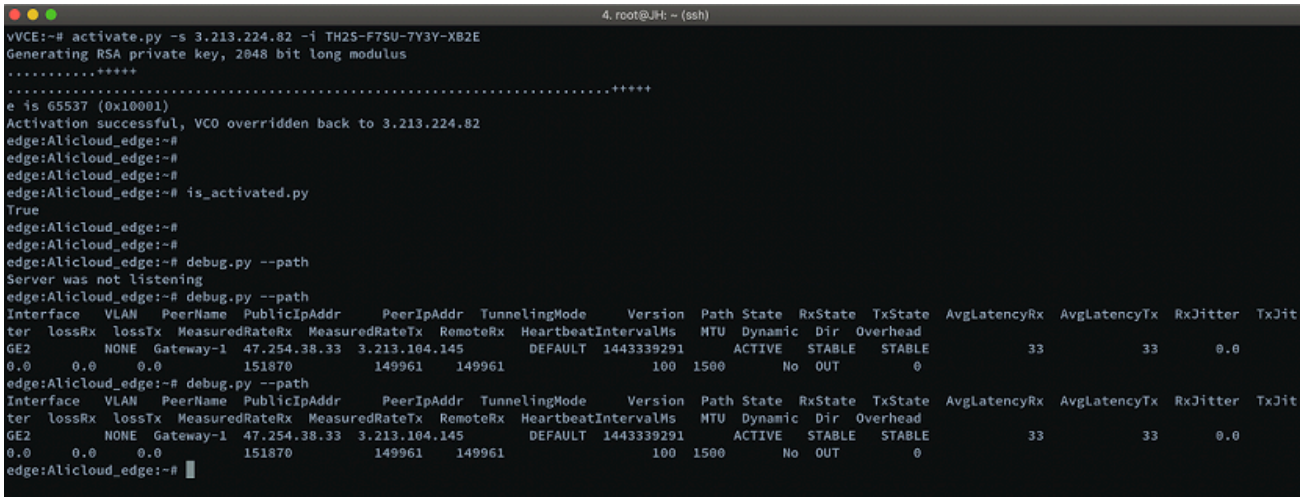

```
4. root@JH: ~ (ssh)
vVCE:~# activate.py -s 3.213.224.82 -i TH2S-F7SU-7Y3Y-XB2E
Generating RSA private key, 2048 bit long modulus
...........+++++
.......................................................................+++++
e is 65537 (0x10001)
Activation successful, VCO overridden back to 3.213.224.82
edge:Alicloud_edge:~#
edge:Alicloud_edge:~#
edge:Alicloud_edge:~#
edge:Alicloud_edge:~# is_activated.py
True
edge:Alicloud_edge:~#
edge:Alicloud_edge:~#
edge:Alicloud_edge:~# debug.py --path
Server was not listening
edge:Alicloud_edge:~# debug.py --path
Interface   VLAN   PeerName  PublicIpAddr     PeerIpAddr  TunnelingMode     Version  Path State  RxState  TxState  AvgLatencyRx  AvgLatencyTx  RxJitter  TxJit
ter  lossRx  lossTx  MeasuredRateRx  MeasuredRateTx  RemoteRx  HeartbeatIntervalMs   MTU  Dynamic  Dir  Overhead
GE2         NONE  Gateway-1  47.254.38.33  3.213.104.145        DEFAULT  1443339291      ACTIVE   STABLE   STABLE            33            33       0.0
0.0     0.0     0.0          151870          149961    149961                   100  1500       No  OUT         0
edge:Alicloud_edge:~# debug.py --path
Interface   VLAN   PeerName  PublicIpAddr     PeerIpAddr  TunnelingMode     Version  Path State  RxState  TxState  AvgLatencyRx  AvgLatencyTx  RxJitter  TxJit
ter  lossRx  lossTx  MeasuredRateRx  MeasuredRateTx  RemoteRx  HeartbeatIntervalMs   MTU  Dynamic  Dir  Overhead
GE2         NONE  Gateway-1  47.254.38.33  3.213.104.145        DEFAULT  1443339291      ACTIVE   STABLE   STABLE            33            33       0.0
0.0     0.0     0.0          151870          149961    149961                   100  1500       No  OUT         0
edge:Alicloud_edge:~#
```

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*