

HPE Aruba Networking EdgeConnect SD-WAN Virtual (EC-V) in Microsoft Azure

Deployment Guide

Contents

- 1. Overview 7
- 2. Why deploy an EC-V gateway in Azure? 7
 - 2.1. Structure of this guide 7
 - 2.2. Decide whether to deploy EC-V into transit VNet or within vWAN hub (Managed NVA) 9
- 3. Deploy EC-V into a Transit Virtual Network (VNet) 11
 - 3.1. Deploy EC-V into a transit VNet from Orchestrator 11
 - 3.1.1. Prerequisites 11
 - 3.1.2. Configure the Azure subscription 12
 - 3.1.3. Add Azure subscription details on Orchestrator 17
 - 3.1.4. Deploy EC-V 17
 - 3.2. Deploy EC-V into a transit VNet from Azure Portal 25
 - 3.2.1. Prerequisites 25
 - 3.2.2. Create an EC-V gateway in the Azure Portal 26
 - 3.3. Deploy EC-V into a transit VNet in an Azure Extended Zone from Azure Portal 41
 - 3.3.1. Things to keep in mind when deploying EC-V to an Extended Zone 41
- 4. Integrate EC-V in transit VNet with Azure Standard Internal Load Balancer (ILB) 43
 - 4.1. EC-V with ILB architecture 43
 - 4.1.1. Horizontal scaling of EC-Vs with ILB 43
 - 4.1.2. ILB limitations 43
 - 4.1.3. Detection of EC-V failures 43
 - 4.1.4. Azure ILB pricing 43
 - 4.1.5. Topology 44
 - 4.1.6. Decide whether to deploy ILB from Orchestrator or manually from Azure Portal 44
 - 4.2. Integrate EC-V with ILB using Orchestrator 45
 - 4.2.1. Prerequisites 45
 - 4.2.2. Add the Azure Subscription to Orchestrator 45
 - 4.2.3. Create ILB 47
 - 4.2.4. Configure LAN interface labels 49
 - 4.2.5. Configure Azure resources in Orchestrator 50
 - 4.2.6. Associate EC-V gateways with ILB 51
 - 4.2.7. Verify connectivity 52
 - 4.2.8. Prevent the static route from being advertised to the SD-WAN fabric 52
 - 4.2.9. Verify health probe responses using the Flows page 52
 - 4.2.10. (Optional) Verify health probe responses using tcpdump 53
 - 4.2.11. Create VNet peering between spoke (workloads) VNet and transit VNet 53
 - 4.2.12. Create an IP SLA rule on EC-V to monitor internet connectivity 53
 - 4.3. Integrate EC-V with ILB manually from Azure Portal 53
 - 4.3.1. Prerequisites 53
 - 4.3.2. Create ILB 53
 - 4.3.3. Create a static route on the EC-V to respond to health probe traffic 57

- 4.3.4. Prevent the static route from being advertised to the SD-WAN fabric 58
- 4.3.5. Verify health probe responses using the Flows page 58
- 4.3.6. (Optional) Verify health probe responses using tcpdump 58
- 4.3.7. Create an IP SLA rule on the EC-V to monitor internet connectivity..... 59
- 4.4. Create a cluster profile to enable flow redirection on the EC-V..... 59
- 4.5. Create VNet peering between spoke (workloads) VNet and transit VNet..... 60
- 5. Integrate EC-V in transit VNet with Azure Route Server (ARS) 61
 - 5.1. EC-V with ARS architecture..... 61
 - 5.1.1. ARS limitations 61
 - 5.1.2. Topology of a single-region EC-V deployment with ARS..... 61
 - 5.1.3. Topology of a multi-region EC-V deployment with ARS..... 63
 - 5.1.4. Horizontal scaling of EC-Vs with ARS..... 64
 - 5.1.5. Detection of EC-V failures 64
 - 5.1.6. Decide whether to deploy ARS from Orchestrator or manually from Azure Portal 64
 - 5.2. Integrate with ARS from Orchestrator 65
 - 5.2.1. Prerequisites..... 65
 - 5.2.2. Add Azure subscription to Orchestrator 65
 - 5.2.3. Create ARS 67
 - 5.2.4. Configure LAN interface labels..... 68
 - 5.2.5. Configure Azure resources in Orchestrator..... 70
 - 5.2.6. Associate EC-V gateways with ARS 71
 - 5.2.7. Verify connectivity..... 71
 - 5.2.8. Create VNet peering between spoke (workloads) VNet and transit VNet..... 72
 - 5.2.9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric.. 72
 - 5.3. Integrate with ARS manually from Azure Portal 72
 - 5.3.1. Prerequisites..... 72
 - 5.3.2. Create a subnet for the ARS 72
 - 5.3.3. Deploy ARS 73
 - 5.3.4. Add EC-V gateways as BGP peers on the ARS 73
 - 5.3.5. Create two static routes for enabling BGP 74
 - 5.3.6. Add ARS as a BGP peer on the EC-V gateways 75
 - 5.3.7. Create VNet peering between spoke (workloads) VNet and transit VNet..... 76
 - 5.3.8. Verify connectivity..... 76
 - 5.3.9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric.. 76
- 6. Transit VNet + Azure Virtual WAN (vWAN) hub design 77
 - 6.1. EC-V with Azure vWAN hub architecture 77
 - 6.1.1. Topology of a single-region EC-V deployment with vWAN hub..... 77
 - 6.1.2. Topology of a multi-region EC-V deployment with vWAN hub..... 78
 - 6.1.3. Horizontal scaling of EC-Vs with vWAN hub..... 79
 - 6.1.4. Detection of EC-V failures 80
 - 6.2. Integrate with vWAN hub using Orchestrator 80
 - 6.2.1. Prerequisites..... 80

- 6.2.2. Add the Azure subscription to Orchestrator 80
- 6.2.3. Verify vWAN hub region 82
- 6.2.4. Configure LAN interface labels..... 82
- 6.2.5. Configure Azure resources in Orchestrator..... 84
- 6.2.6. Associate EC-V gateways with vWAN hub 85
- 6.2.7. Verify connectivity..... 85
- 6.2.8. Create a virtual network connection between spoke (workloads) VNet and vWAN hub 86
- 6.2.9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric.. 86
- 6.3. Integrate with vWAN hub manually from Azure Portal 86
 - 6.3.1. Prerequisites..... 86
 - 6.3.2. Create a vWAN hub..... 86
 - 6.3.3. Create a virtual network connection 87
 - 6.3.4. Add EC-V gateways as BGP peers on the vWAN hub 87
 - 6.3.5. Create two static routes for enabling BGP 88
 - 6.3.6. Add vWAN hub as a BGP peer on the EC-V gateways 88
 - 6.3.7. Verify connectivity..... 89
 - 6.3.8. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric.. 90
- 7. Deploy EC-V within Azure vWAN hub (Managed NVA) 91
 - 7.1. Deploy EC-V within Azure vWAN hub (Managed NVA) from Orchestrator..... 91
 - 7.1.1. Prerequisites..... 91
 - 7.1.2. Configure the Azure subscription 91
 - 7.1.3. Add Azure subscription details on Orchestrator..... 94
 - 7.1.4. Deploy Managed NVA..... 94
 - 7.1.5. Configure LAN interface labels..... 98
 - 7.1.6. Associate EC-V gateways with vWAN hub 100
 - 7.1.7. Verify connectivity..... 100
 - 7.1.8. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric 101
 - 7.2. Deploy EC-V within Azure vWAN hub (Managed NVA) from Azure Portal..... 101
 - 7.2.1. Prerequisites..... 101
 - 7.2.2. Create a vWAN hub..... 101
 - 7.2.3. Create a custom role and assign a user-assigned managed identity 102
 - 7.2.4. Deploy EC-V gateways within the vWAN hub 103
 - 7.2.5. Correlate EC-Vs 107
 - 7.2.6. Add EC-Vs to the SD-WAN fabric 109
 - 7.2.7. Assign IP addresses..... 111
 - 7.2.8. Shut down MGMT0..... 112
 - 7.2.9. Create two static routes for enabling BGP 113
 - 7.2.10. Create static routes for Azure health probes..... 114
 - 7.2.11. Apply templates and overlays 115
 - 7.2.12. Add vWAN hub as a BGP peer on the EC-V gateways 115
 - 7.2.13. Verify connectivity..... 116
 - 7.2.14. Where to find configuration items 116

- 7.2.15. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric 117
- 8. Redirect traffic from spoke (workloads) VNets to EC-V via ILB, ARS, or vWAN hub 118
 - 8.1. Create a virtual network peering session (for ILB and ARS designs) 118
 - 8.2. Create a virtual network connection (for vWAN hub-related designs)..... 119
 - 8.3. Create static routes on Azure Portal to forward outbound traffic from a workloads VNet to ILB 120
 - 8.3.1. Create an Azure route table 121
 - 8.3.2. Associate the workloads subnet to the route table..... 121
 - 8.3.3. Create a static route to forward traffic to the ILB..... 121
 - 8.4. Create static routes on EC-V to advertise Azure subnets to remote EdgeConnect devices 122
- 9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric 124
- 10. Create a cluster profile to enable flow redirection on the EC-V..... 125
 - 10.1. Create a cluster profile..... 126
 - 10.2. Add EC-V gateways to the cluster 126
 - 10.3. Apply a cluster profile to EC-Vs 126

Document Revision History

Project name: HPE Aruba Networking EdgeConnect SD-WAN Virtual (EC-V) in Microsoft Azure

Document status: Revision A

| Document version | Date | Section and text revised |
|------------------|------------------|-----------------------------|
| Revision A | January 14, 2026 | Original document revision. |

1. Overview

An HPE Aruba Networking EdgeConnect SD-WAN Virtual gateway (EC-V) can be deployed as a virtual machine (VM) in Azure. An EC-V offers the same benefits as an EdgeConnect SD-WAN hardware device. This deployment guide provides detailed instructions for deploying an EC-V in Azure and adding it into your HPE Aruba Networking EdgeConnect Orchestrator (in this guide, referred to as *Orchestrator*).

This deployment guide is intended for network engineers, cloud architects, and IT administrators responsible for designing and implementing secure, scalable hybrid cloud connectivity using the EdgeConnect SD-WAN platform. This guide covers supported deployment models for EC-V in both Transit Virtual Network (VNet) and Managed Network Virtual Appliance (NVA) designs. A transit VNet is a “hub-style” Azure VNet that acts as a central point for routing and connectivity between multiple spoke VNets, on-premises networks, and other cloud environments. This guide also provides detailed instructions on integrating EC-V with Azure networking services such as Standard Internal Load Balancer (ILB), Azure Route Server (ARS), and Azure Virtual WAN (vWAN) hub.

Whether deploying via Orchestrator or the Azure Portal, this guide outlines all necessary prerequisites, configuration steps, and best practices to ensure successful deployment and integration of EC-V gateways within an Azure environment.

2. Why deploy an EC-V gateway in Azure?

An EC-V gateway in Azure provides a comprehensive set of capabilities, such as routing, firewall, WAN optimization, application visibility and control, and many other SD-WAN capabilities, all within a single, centrally managed system. Deploying an EC-V gateway in Azure has many benefits to establish and enhance WAN connectivity for various use cases, including:

- Branch offices to Azure
- Data centers to Azure
- One Azure region to another over the Azure backbone or over the Internet
- Other public or private cloud environments to Azure

Your EC-V instance can be deployed from the Orchestrator or the Azure Portal. In general, HPE recommends that you deploy EC-V gateways in public cloud environments from Orchestrator, as it greatly simplifies the deployment. However, some organizations require manual integration of Azure Network services with EC-V, as it allows integration with existing workflows. To support such organizations, this guide also outlines the steps to manually deploy and configure EC-V with ILB, ARS, vWAN hub, and Managed NVA.

The EC-V only supports the bring-your-own-license (BYOL) model. It can be deployed as a single VM or in a high availability (HA) configuration. A single VM deployment is suitable for testing or proof-of-concept environments. An HA deployment is recommended for production environments. This guide covers only HA deployments. When deploying EC-V in an HA configuration, it is highly recommended to deploy each EC-V in a unique availability zone (AZ). Deploying two or more EC-Vs in separate AZs ensures continuous operation even if one AZ experiences a failure. This setup also improves resilience and redundancy, reducing downtime and increasing reliability for mission-critical workloads.

EC-V is available on the Azure Public (Commercial) Cloud Marketplace, Azure Government Cloud Marketplace, and Azure China Marketplace. Also, EC-V can be deployed manually from the Commercial Marketplace into an Azure Extended Zone environment. This is explained in Section [3.2](#).

2.1. Structure of this guide

This guide is organized into nine sections:

- Section [3](#) contains step-by-step instructions for deploying EC-V appliances into a transit VNet.
- Section [4](#), Section [5](#), and Section [6](#) illustrate how to integrate EC-V gateways deployed in a transit VNet with Azure ILB, ARS, and vWAN hub, respectively. Each section illustrates how to connect EC-V gateways to your selected Azure network service, either using Orchestrator or manually from Azure Portal. Figure 1 shows the logical topology of EC-V gateways that are deployed in a transit VNet with ILB, ARS, and vWAN hub. Typically, in a single region in Azure, you would pick only one design for your SD-WAN deployment. Each design has advantages and disadvantages, which are discussed in their respective section in the guide.

— Section 7 contains step-by-step instructions for deploying Managed NVA within a vWAN hub from Orchestrator or manually from Azure Portal and establishing BGP with the virtual hub’s internal route service. Figure 2 shows the logical topology of EC-V gateways deployed within a vWAN hub. As mentioned, this design is known as the Managed NVA design. The HPE team has collaborated with the Azure vWAN team to support this deployment for users who need a turnkey deployment of EC-V gateways within a vWAN hub.

Important

When you use the Managed NVA deployment, you do not deploy VMs or manage IP addresses, NICs, or route tables directly from Azure Portal. The compute instances used in the Managed NVA deployment are deployed within the vWAN’s hub VNet by Azure into a Virtual Machine Scale Set (VMSS). Section 2.2 discusses the pros and cons of each design. Before deploying your EC-V gateways, it is highly recommended that you read each deployment model and select the most appropriate one for your environment.

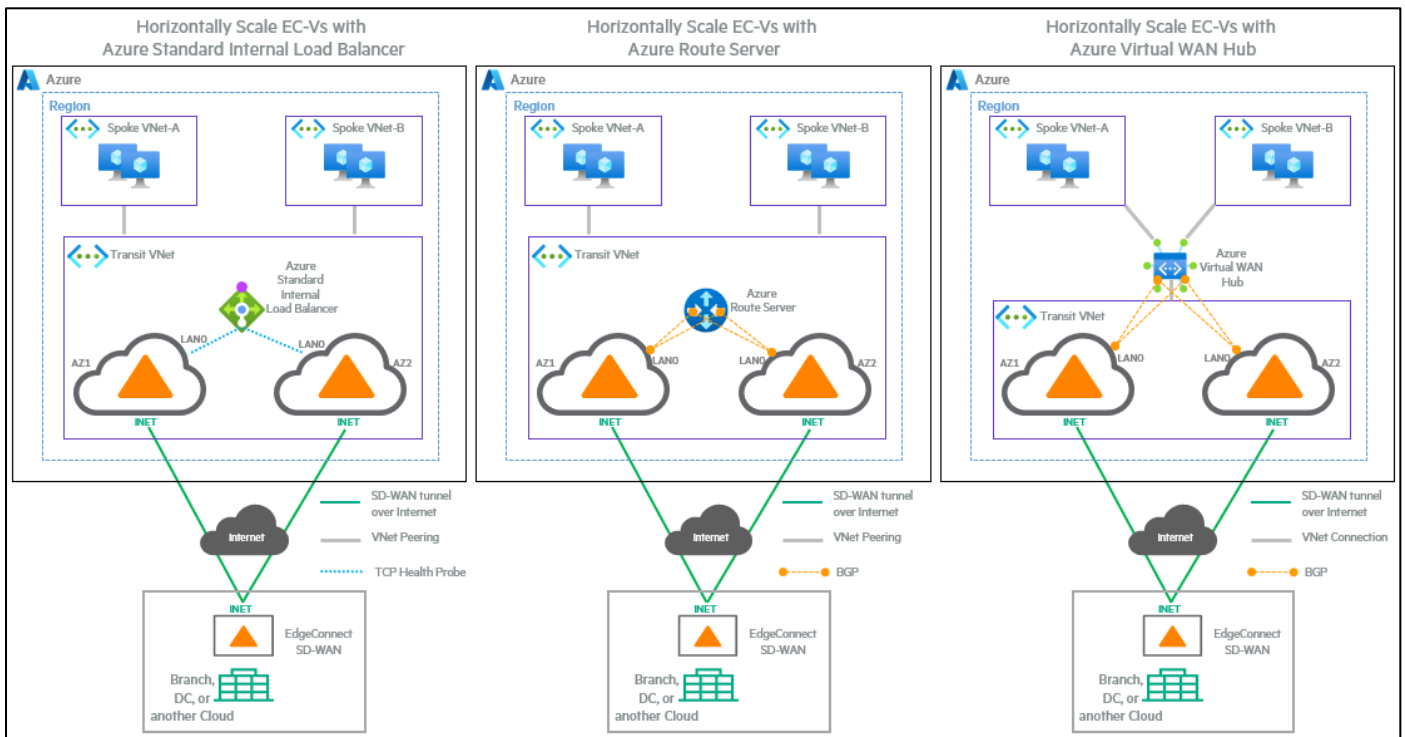


Figure 1. Deploying EC-V in a Transit Virtual Network with ILB, ARS, or vWAN hub.

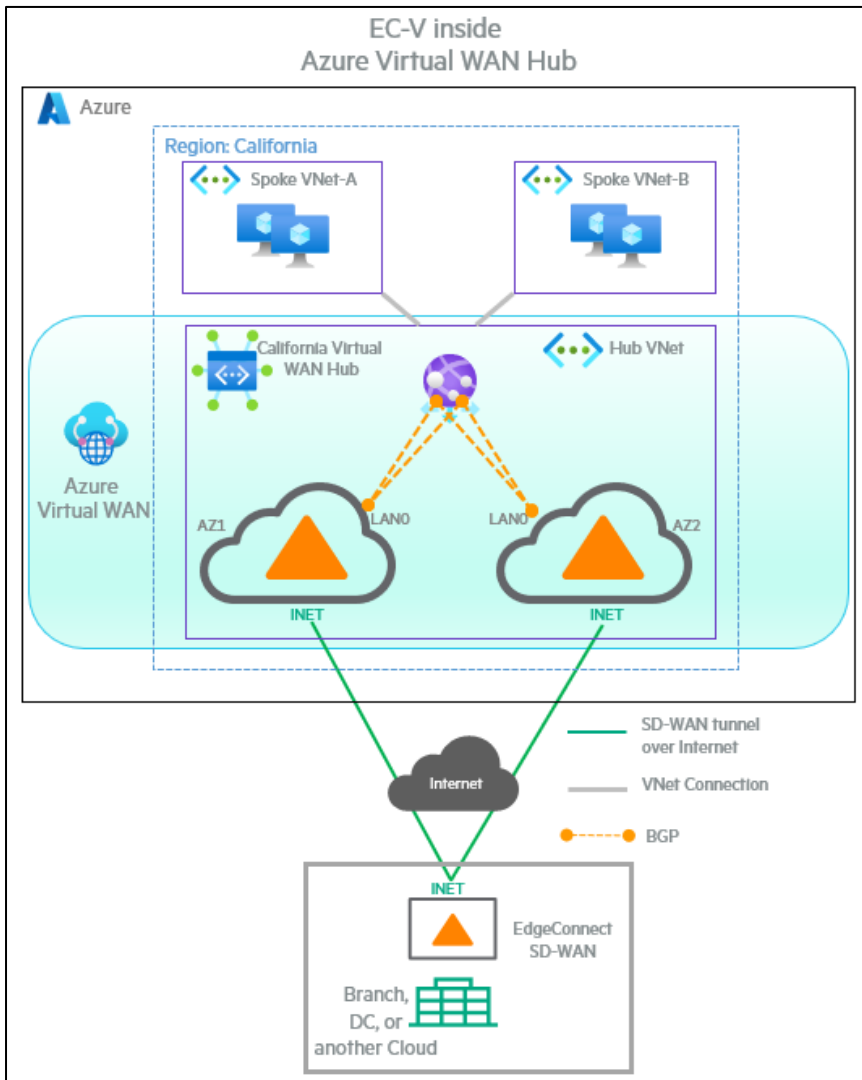


Figure 2. Deploying EC-V within a vWAN hub (Managed NVA design).

2.2. Decide whether to deploy EC-V into transit VNet or within vWAN hub (Managed NVA)

When deciding between deploying an EC-V gateway within a vWAN hub (Managed NVA) or within a transit VNet, you must consider several factors.

The Managed NVA comes with several limitations:

1. Users cannot add or delete network interfaces on the EC-V. Each EC-V has only one WAN interface and one LAN interface. There is no dedicated management (MGMT0) interface available on the EC-V.
2. Users do not have access to modify the default network security group (NSG) assigned to the EC-V. The firewall rules on the default WAN0 and LAN0 NSGs specified by HPE cannot be edited by users.
3. SSH access is disabled on EC-V. There is no Serial Console access to the VM from the Azure Portal. The only way to access the CLI of the EC-V is from the Orchestrator's web CLI.
4. By default, each deployment consists of two EC-Vs. A single EC-V deployment within a vWAN hub is unsupported. If you require three EC-Vs to increase the bandwidth available for your Azure-based service and workloads, you must initiate a new deployment. This new deployment creates two EC-Vs, resulting in a total of four EC-Vs after the deployment.

Deploying an EC-V inside a transit VNet offers more flexibility and control of the EC-V deployment. Advantages of using EC-V in a transit VNet include:

1. The ability to create custom route tables to control traffic forwarding for WAN0 and LAN0 subnets of the EC-V.
2. Full control of the NSGs assigned to each network interface of the EC-V.

3. The ability to add or delete network interfaces on the EC-V. This is important for users who want to establish SD-WAN tunnels over both internet and ExpressRoute simultaneously.
4. Access to the EC-V's CLI via Azure Portal's Serial Console.
5. SSH access is allowed after opening the inbound TCP port 22 on the NSG.

Now that you understand the pros and cons of deploying EC-V within a transit VNet or Managed NVA, the following sections go deeper into each design, examining the distinct configurations and possible challenges of each deployment option to help you make an informed decision.

3. Deploy EC-V into a Transit Virtual Network (VNet)

Starting from Orchestrator version 9.1.1, the Orchestrator supports deploying one or more EC-Vs into an Azure VNet. Deploying EC-Vs from Orchestrator offers several advantages over manual deployment from the Azure Portal, such as:

- **Simplified deployment:** Orchestrator automates the deployment process, reducing manual steps and the potential for errors. It handles the creation of necessary Azure resources such as virtual networks, subnets, network interfaces, public IPs, NSGs, and rules within NSGs. It also enables IP forwarding, assignment of MAC addresses, and configuration of IP addresses on the EC-V, which streamlines the deployment process significantly; manual deployment requires the configuration of each of these components individually. In Orchestrator version 9.5.2 and later, the Orchestrator also allows deployment of an EC-V into an existing VNet.
- **Consistent configuration:** Orchestrator ensures consistent configuration across all EC-V deployments. This is crucial for maintaining standardization and reducing troubleshooting efforts. Manual deployments are more prone to configuration discrepancies.
- **Automated scaling:** Orchestrator allows for easy horizontal scaling of EC-Vs. You can quickly add or remove EC-V instances as needed, directly from the Orchestrator interface. When deploying multiple EC-Vs within a single region, the Orchestrator prompts you to select the availability zone that it should deploy each EC-V.
- **Deployment into multiple Azure subscriptions:** Orchestrator allows users to enter details of multiple Azure subscriptions, making selection for EC-V deployment simple. Using this ability, you can deploy EC-Vs into a test subscription for testing and a production subscription for production from a single Orchestrator.
- **Simplified removal:** Orchestrator simplifies the removal of EC-V deployments. With a single click, it can terminate a deployment and delete all associated resources it created while deploying the EC-V. Manual removal requires deleting each resource individually, which can be time-consuming.

Section [3.1](#) describes how to deploy EC-V from Orchestrator. Section [3.2](#) describes how to deploy EC-V manually from the Azure Portal.

3.1. Deploy EC-V into a transit VNet from Orchestrator

Deploying a new EC-V into a transit VNet from the Orchestrator involves the following steps. Note that Steps 3.1.2 and 3.1.3 are one-time tasks:

- [3.1.2. Configure the Azure subscription](#)
 - [3.1.2.1. Accept Azure Marketplace image terms](#)
 - [3.1.2.2. Create a new app registration](#)
 - [3.1.2.3. Create a new resource group](#)
 - [3.1.2.4. Create a custom role](#)
 - [3.1.2.5. Assign the custom role to the resource group](#)
- [3.1.3. Add Azure subscription details on Orchestrator](#)

3.1.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- An Azure subscription
- HPE Aruba Networking EdgeConnect Orchestrator
- HPE Aruba Networking EdgeConnect licenses
- Check Azure service limits.
 - Before starting the deployment, it is highly recommended to check your Azure subscription quotas in the selected region for compute cores, virtual networks, and subnets to ensure that you have enough available resources to avoid deployment failures.
 - **Number of EC-Vs:** The number of EC-Vs you plan to deploy directly impacts the required number of subnets (if using an existing VNet) and compute cores.
 - **VM SKUs:** The chosen VM SKU determines the required compute cores.

- **Region:** Azure quotas are region-specific. Ensure that you have the necessary quotas in the region you select for deployment.
- **Existing VNet vs. new VNet:** If using an existing VNet, ensure the availability of the required subnets. If you allow the Orchestrator to create a new VNet for the deployment, ensure that the CIDR block you enter does not overlap with an existing CIDR block in your network.
- To learn how to manage your Azure service limits, go to <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits>.

Notes

- When an Orchestrator deploys an EC-V, usually the EdgeConnect OS (ECOS) version of that EC-V is earlier than the version you want. Upon discovering the EC-V on the Orchestrator, you can upgrade it before or after you add it to the Orchestrator's appliance tree. The upgrade takes only a few minutes.
- As of December 2025, Orchestrator does not support deploying EC-V (using the Cloud Hubs in Azure feature) into Azure Government Cloud, Azure Extended Zones, or Azure China. Please follow the instructions in Section [3.2](#) to manually deploy EC-V into these three environments.

3.1.2. Configure the Azure subscription

Before deploying an EC-V via Orchestrator, you must configure your Azure subscription by accepting Marketplace terms for EdgeConnect to enable programmatic deployment, creating an app registration (service principal), resource group, and custom role, and assigning that role to the resource group.

3.1.2.1. Accept Azure Marketplace image terms

Accepting Azure Marketplace's image terms for EdgeConnect is required for the Orchestrator to automatically deploy an EdgeConnect image from the Azure Marketplace. You only need to do this once per Azure subscription.

1. Log in to the Azure Portal.
2. Under Azure services, click **+ Create a resource**.
3. On the Create a resource page, enter **edgeconnect**, and then click **Silver Peak Unity EdgeConnect**.
4. On the Plan drop-down menu, select the latest EdgeConnect version, and then click **Get started**.
5. On the Configure Programmatic Deployment page, select **Enable** next to the subscription ID that you want to use to deploy the EdgeConnect VMs.

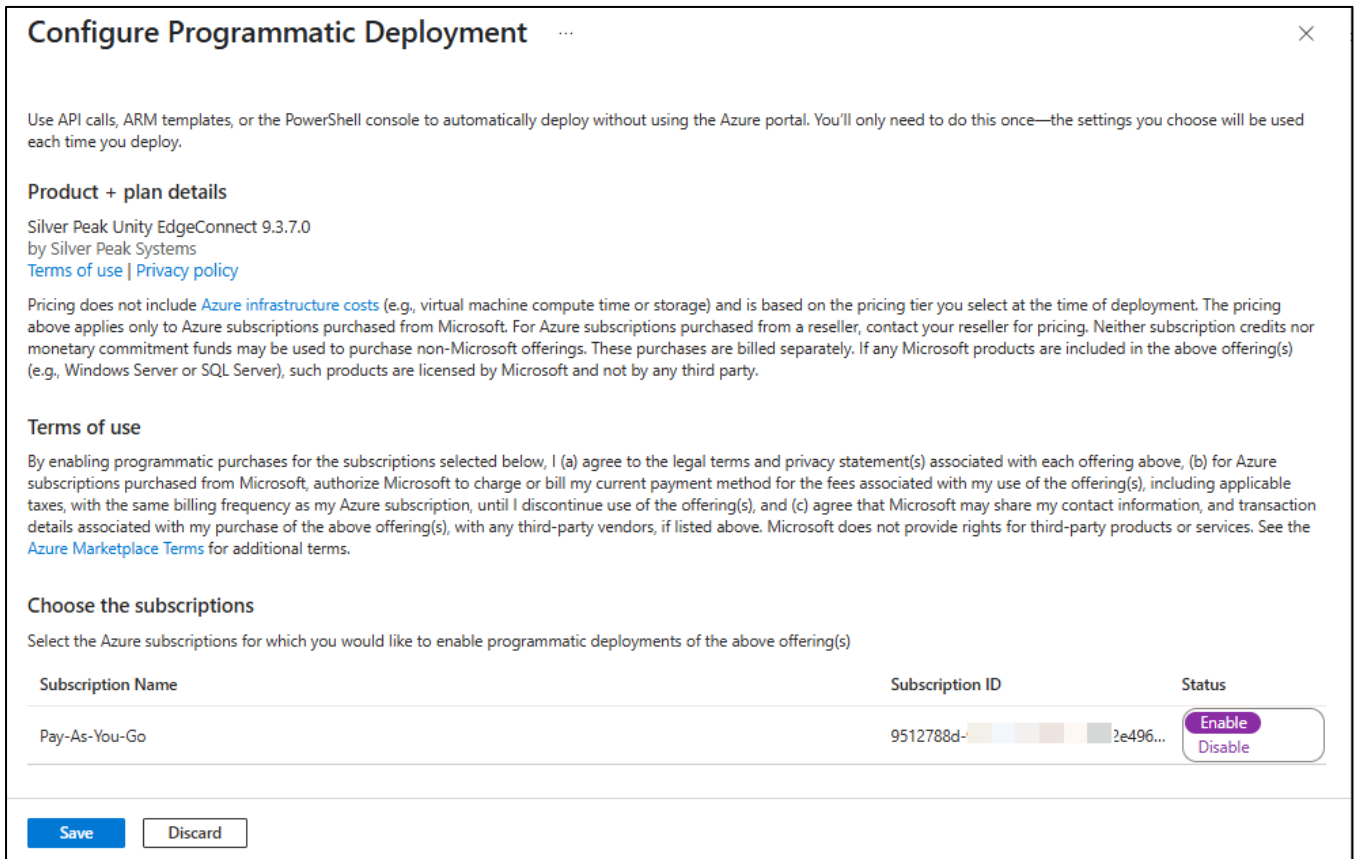


Figure 3. Enabling your subscription on the Configure the Programmatic Deployment page.

6. Click **Save**.

A message at the top of the page notifies you when configuration updates are complete.

3.1.2.2. Create a new app registration

Azure app registrations are used to create service principals, which are identities used by applications or services to access Azure resources securely and with restricted permissions. The Orchestrator uses the app registration to authenticate itself with Azure. An app registration allows you to assign least-privilege access using custom roles. It avoids sharing credentials and exposing elevated privileges unnecessarily.

To create a new app registration for the Orchestrator:

1. Log in to the Azure Portal.
2. In the main search menu, enter **app registrations**, and then click **App registrations**.
3. Click **+ New registration**.
4. On the Register an application page, in the Name field, enter a user-facing display name for the application.
5. Under Supported account types, select **Accounts in this organizational directory only (Default Directory only - single tenant)**.
6. *(Optional)* Enter a redirect URI.
7. Click **Register**.

Note

Note the application (client) ID and directory (tenant) ID. You need these IDs when you add the subscription details on the Orchestrator.

8. Under Manage, click **Certificates & secrets**.
9. Click **New client secret**.
10. Enter a Description and Expiration Date.
11. Click **Add**.

A new client secret is created.

12. Copy the text in the Value column.

ⓘ Important

This text can only be viewed immediately after creation. Be sure to save the secret before leaving the page.

13. On the main search menu bar, enter **subscription**, and then press **Enter**.

14. Copy the subscription ID.

You have successfully registered your application and gathered the details required for adding the Azure subscription details to the Orchestrator.

3.1.2.3. Create a new resource group

Creating a new resource group on the Azure Portal is considered a best practice. This ensures that the SD-WAN Orchestrator only has access to that resource group to deploy EC-Vs. However, it is possible to deploy one or more EC-Vs into an existing resource group that contains other Azure resources.

To create a new resource group:

1. On the main search menu, enter **resource group**, and then click **Resource groups**.
2. Click **+ Create**.
3. On the Create a resource group page, select the subscription that you want to use to create the resource group.
4. Enter a name for the resource group, and then select a region.
5. Click **Review + create**.
6. Click **Create**.

3.1.2.4. Create a custom role

To create custom roles, you must have owner or user access administrator permissions. There are multiple ways to create a custom role. The following steps create a custom role from within the resource group that you created.

1. Select the resource group you created in [Create a new resource group](#), and then click **Access control (IAM)**.
2. Click **Add**, and then click **Add custom role**.

The Custom Roles editor appears (the Basics tab is displayed).

3. In the Custom role name field, enter a name for the custom role. The name must be unique for the Azure AD directory. The name can include letters, numbers, spaces, and special characters.
4. *(Optional)* In the Description field, enter a description for the custom role. The description displays in the tool tip for the custom role.
5. Accept the default value for the Baseline permissions, and then click the **JSON** tab.
6. Click **Edit**.
7. Go to https://arubanetworking.hpe.com/techdocs/SilverPeak/files/cloud-ecv/cloud_ecv_json.htm and click **Permissions required to deploy Cloud Hubs in Azure in a Transit Virtual Network (VNet)**.
8. Copy and paste the list of Azure permissions within the square brackets next to actions (line 10), as shown in the following figures.

Home > ArubaEdgeConnectDeployment >

Create a custom role ...

Got feedback?

If you change tabs and have unsaved edits, those changes will be discarded.

Basics Permissions Assignable scopes **JSON** Review + create

Here is your custom role in JSON format. [Learn more](#)

Download Discard changes Save

```
1 {
2   "properties": {
3     "roleName": "",
4     "description": "",
5     "assignableScopes": [
6       "/subscriptions/9512788d-90b2-42d7-aa25-f172e4969c82/resourceGroups/ArubaEc
7     ],
8     "permissions": [
9       {
10      "actions": [],
11      "notActions": [],
12      "dataActions": [],
13      "notDataActions": []
14    }
15  ]
16 }
17 }
```

Figure 4. Finding the actions field in the Custom a custom role editor.

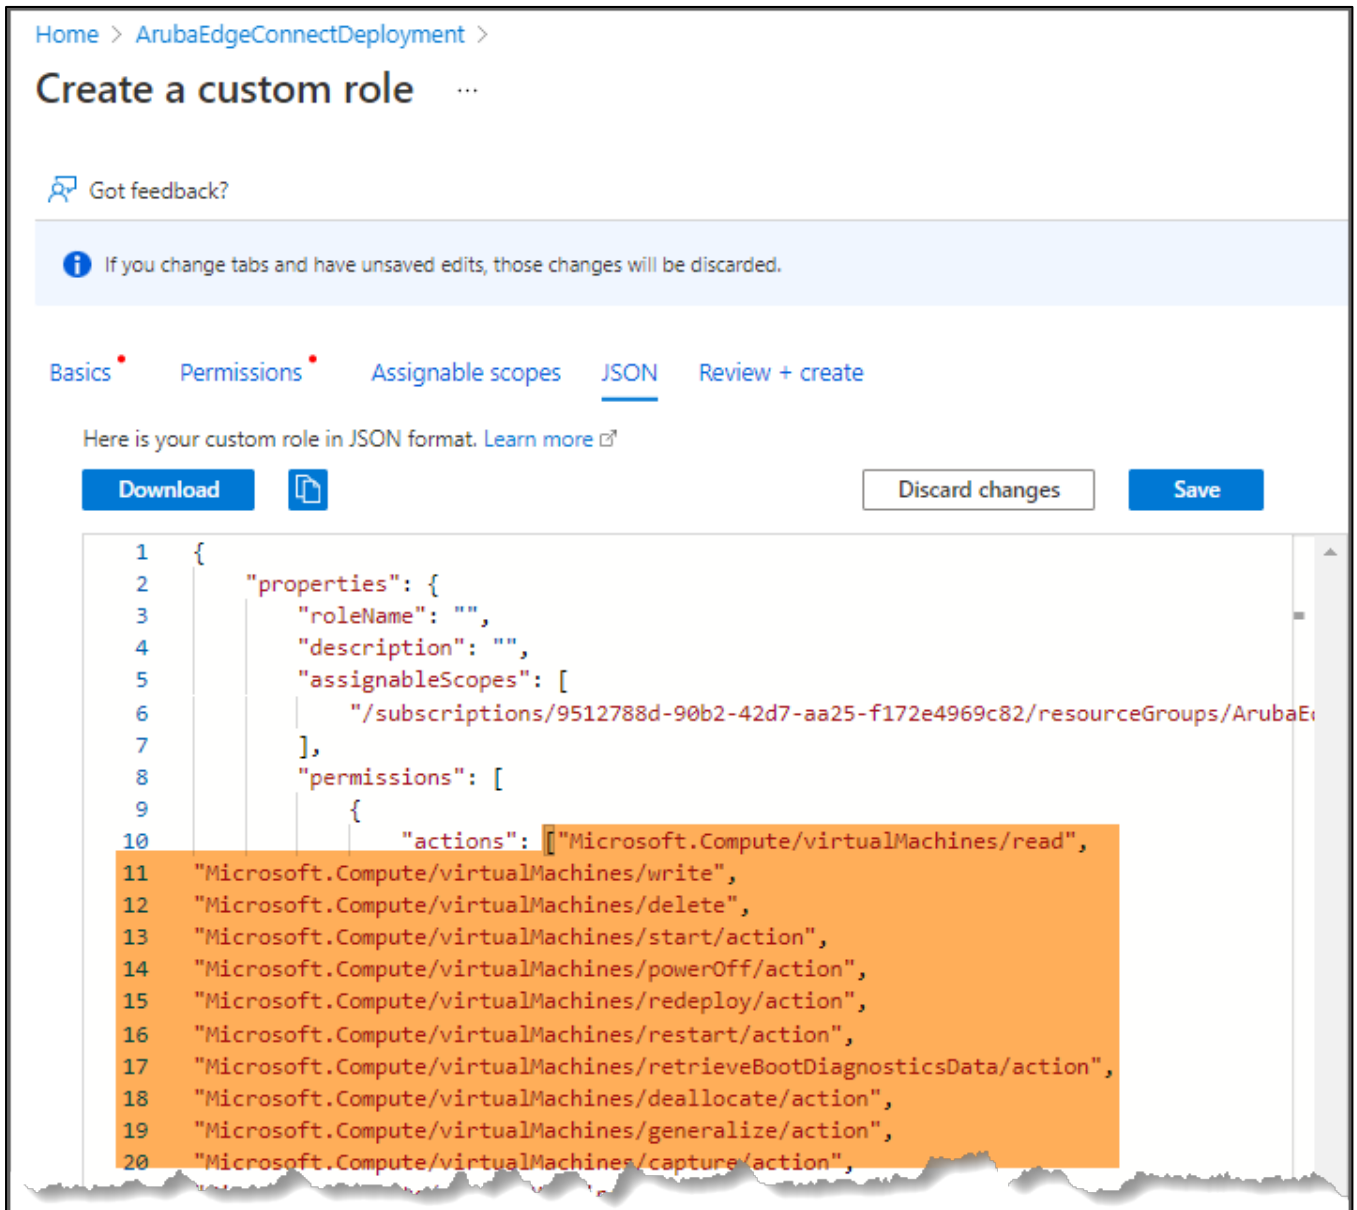


Figure 5. Pasting the Azure permissions into the actions field of the Create a custom role editor.

9. Click **Save**.
10. Click the **Assignable scopes** tab, and then verify that the resource group you created is added as an assignable scope and Type is set to the resource group.
11. Click the **Permissions** tab, and then verify that the permissions, descriptions, and permission types you added are listed.
12. Click **Review + create**.
13. Click **Create**. A message displays to confirm that you have successfully created your custom role.

3.1.2.5. Assign the custom role to the resource group

Assigning the custom role to the resource group ensures that the Orchestrator has the exact set of permissions needed to deploy only the resources within that resource group, in line with the principle of least privilege. However, if you want to deploy EdgeConnect gateways into multiple resource groups, you can assign the custom role at the Azure subscription level. This allows the Orchestrator to list multiple resource groups and lets you select the one you want for deployment. In this example, the custom role is assigned only to the specific resource group.

1. Navigate to the resource group you created, and then click **Access control (IAM)**.

Note

If you just completed the previous task of creating a custom role, the Access control (IAM) page is already open.

2. Click **Add**, and then click **Add role assignment**.

The Role assignment page appears.

3. On the Role tab, enter the name of your custom role.

Note

If the role you created is not displayed, refresh the page.

4. Select the custom role, and then click **Next**.

The Members tab appears.

5. Ensure that **User, group, or service principle** is selected, and then click **+ Select members**.

The Select members page appears.

6. Enter the name of your app registration (service principle), select your app, and then click **Select**.

Your app is added under Members.

7. Click **Review + assign**.

8. Click **Review + assign** again.

You have successfully assigned your custom role to the resource group.

3.1.3. Add Azure subscription details on Orchestrator

Add your Azure subscription details such as subscription ID, directory (tenant) ID, application (client) ID, and client secret value to Orchestrator. You only need to do this once per Azure subscription unless you are a large enterprise with multiple teams that want to use different resource groups to segregate deployments.

To add the Azure subscription to Orchestrator:

1. Log in to Orchestrator.
2. Navigate to **Configuration > IaaS > Deploy Cloud Hubs in Azure**.
3. Click **Azure Subscriptions**.
4. Click **Add Azure Subscription**.
5. Enter a name, subscription ID, directory (tenant) ID, application (client) ID, and client secret for the Azure subscription.

Notes

- You copied the directory (tenant) ID, application (client) ID, and client secret in [3.1.2.2 Create a new app registration](#).
- If you copy and paste the subscription ID, Azure might add a blank space to the beginning of the subscription ID. Be sure to remove all spaces from your subscription ID.

6. Click **Save**.

Orchestrator validates the subscription information.

Note

If you have multiple Azure subscriptions, you can load the credentials for each subscription.

3.1.4. Deploy EC-V

To deploy EC-V into a transit VNet from Orchestrator

1. Log in to the Orchestrator.
2. Navigate to **Configuration > Cloud Services > IaaS > Cloud Hubs in Azure**.
3. Click **Deploy Cloud Hubs in Transit VNet**.

The Azure Deployment Configuration dialog box appears.

- To instantiate EC-V instances, enter the deployment details in the Orchestrator's Azure Deployment Configuration dialog box, as noted in the table below.

Note

If you do not have an Azure subscription configured in Orchestrator, the Azure Deployment Configuration dialog box is blank. Click the **Subscriptions** link to go to the Azure subscription page and enter your Azure credential.

| Field | Description |
|------------------------|--|
| Name | Enter a name for the deployment. This name is used only for identifying the deployment. A deployment consists of one or more EC-Vs that an Orchestrator creates in an Azure virtual network. Only alphabetical letters and hyphens are allowed in the deployment name. The maximum allowed length is 20 characters. |
| Azure account | Select an Azure subscription to use for deploying the EC-V. |
| Resource group | Select an Azure resource group to use for deploying the EC-V. |
| Region | Select an Azure region where you want to deploy the EC-V. |
| Virtual Network | <p>Select Create new virtual network or Existing virtual network.</p> <p>Create new virtual network: If selected, Orchestrator creates a new VNet for the EC-V. Orchestrator creates three subnets (MGMT0, WAN0, and LAN0) for each EC-V you deploy. For example, if you deploy two EC-Vs, Orchestrator creates six subnets. Each subnet is /28 in size.</p> <p>Existing virtual network: If selected, Orchestrator allows you to select an existing VNet and subnets for MGMT0, WAN0, and LAN0 interfaces.</p> <p>Note: If you select Existing virtual network and you deploy multiple EC-Vs using the Horizontally scale setting, the MGMT0, WAN0, and LAN0 interfaces for each EC-V are created using the subnets you select in the Available subnets fields. For example, if you deploy two EC-Vs, the MGMT0 interface for each EC-V is created on the MGMT0 subnet you select in the mgmt0 field, the WAN0 interface for each EC-V is created on the subnet you select in the wan0 field, and the LAN0 interface for each EC-V is created on the subnet you select in the lan0 field. For a two-EC-V (or a multiple EC-V) deployment, you need three subnets. If you select Create new virtual network, Orchestrator creates six new subnets for a two-EC-V deployment. If you select Existing virtual network, the VNet that you select must be created within the same resource group as the EC-V deployment. If your VNet is created in a different resource group, the Orchestrator does not display that VNet under the Existing virtual network option.</p> |
| Virtual Network CIDR | If you select the Create new virtual network option, you need to enter a Virtual Network Classless Inter-Domain Routing (CIDR) block. Orchestrator uses this CIDR block to create a new VNet. The smallest supported CIDR block is /24 and the largest supported CIDR block is /16. Orchestrator creates all Azure resources required for the EC-V deployment within this virtual network. For each EC-V you deploy, Orchestrator creates three subnets that are /28 in size. In other words, if you deploy two EC-Vs, Orchestrator creates six subnets in total. This is true even if both EC-Vs are created in a single Availability Set or Availability Zone. |
| Choose Virtual Network | If you selected Existing virtual network, enter the name of the network in this field. |
| Available subnets | If you selected Existing virtual network, enter a subnet for each network interface. |
| WAN Optimization | After WAN Optimization and an appropriate WAN bandwidth value are selected, Orchestrator displays the appropriate Azure instance types for the deployment on the Instance Type menu. |
| WAN bandwidth | Note: Selecting WAN Optimization does not enable WAN Optimization on the EC-V. It only allows Orchestrator to display appropriate Azure instance types that can support WAN Optimization for the selected WAN bandwidth. To enable WAN Optimization on the EC-V, go to the Deployment page and the Business Intent Overlay (BIO) page after the deployment is complete. |
| Instance type | The WAN bandwidth list displays the current EdgeConnect license tiers. After you select a WAN Bandwidth value, Orchestrator displays the appropriate Azure instance types for the deployment in the Instance type list. |
| Availability option | Based on your selected WAN Optimization and WAN bandwidth values, Orchestrator displays the appropriate instance types. |
| Availability option | Select Availability Set or Availability Zone . Some regions only support Availability Set. HPE Aruba Networking recommends selecting Availability Zone, if available. |

| Field | Description |
|---------------------------|---|
| SSH public key | <p>Generate a public key with an application, such as PuTTYgen, and then input the value here.</p> <p>ⓘ Important: EdgeConnect only supports single-line SSH public keys. Do not use multi-line SSH public keys. Additionally, use an EdDSA (ED25519) key pair, as shown on the image below:</p> <p>Use this:</p> <pre>1 ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIOFDa8CNR5DmJE3EgIOBpYyyL4mlyqarVo/+XBgISULs eddsa-key-20251015</pre> <p>Not this:</p> <pre>1 ---- BEGIN SSH2 PUBLIC KEY ---- 2 Comment: "eddsa-key-20251015" 3 AAAAC3NzaC1lZDI1NTE5AAAAIOFDa8CNR5DmJE3EgIOBpYyyL4mlyqarVo/+XBgI 4 SULs 5 ---- END SSH2 PUBLIC KEY ---- 6</pre> <p>Note: Save the private key file. If you need to log in via SSH to the appliance after it is deployed, you will need this key.</p> |
| Azure tags (Optional) | <p>Any comma-separated tags entered here are applied to all Azure resources that Orchestrator creates while deploying the EC-V. If you do not enter any tags, Orchestrator automatically creates a unique tag for each Azure resource that it creates while deploying the EC-V. This Azure tag is created to identify each resource created by Orchestrator. The tag is formatted as follows:</p> <p>sp-automated-deployment name-instance-index-resource name</p> |
| Comment (Optional) | Enter an optional comment if you want to attach any additional details for the deployment. |
| Advanced settings | Custom VHD: Leave this field blank unless you have an EdgeConnect VHD that you want to use for the deployment. When this field is blank, the Azure Marketplace image is deployed. |
| Horizontally scale | <p>You can deploy multiple EC-Vs by clicking + and selecting the Availability Set or Availability Zone for each EC-V. If the selected region supports multiple availability zones, each availability zone appears on the menu. You can deploy up to five EC-Vs with a CIDR block of /24.</p> <p>If you need to deploy more than five EC-Vs within a single virtual network, select a virtual network CIDR block that is bigger than /24, such as /23 or /22. The maximum number of EC-Vs you can deploy within a single network is 20.</p> |
| Appliance tags (Optional) | Enter an appliance tag. If this field is left blank, Orchestrator automatically assigns an appliance tag for its own configuration purposes. |
| Availability zone | <p>Enter the Azure availability zone for the EC-V.</p> <p>Note: This field only displays if the region supports availability zones.</p> |

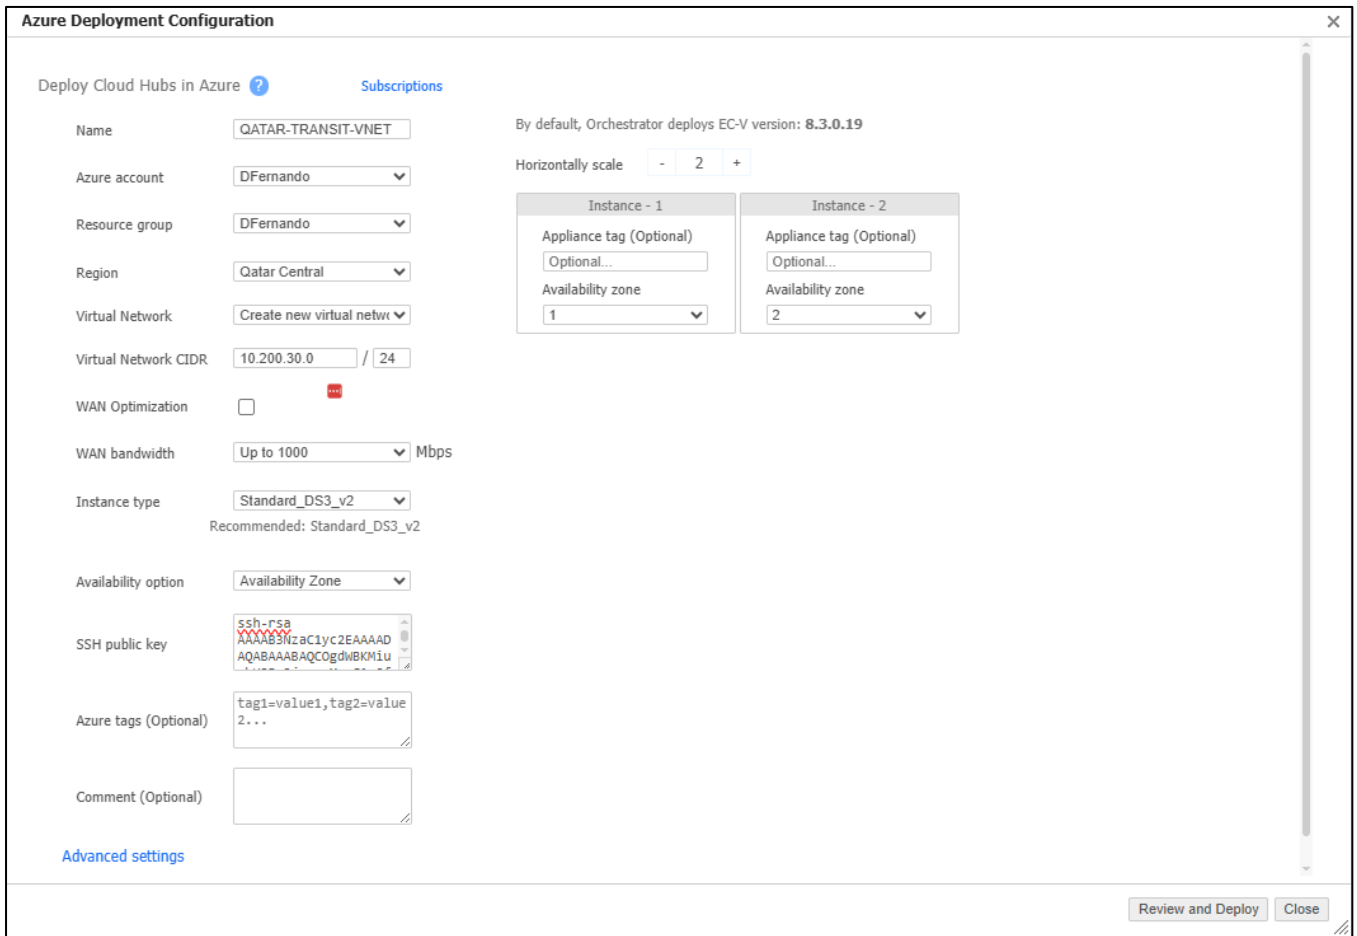


Figure 6. Deploying two EC-Vs into a new transit VNet from Orchestrator.

5. After entering your information as shown in Figure 6, click **Review and Deploy**.
6. Review the configuration summary, and then click **Deploy** to create the EC-V instances.

Note

After a few minutes, click the **refresh** icon next to the Deploy Cloud Hubs in Azure label to check the status.

If your EC-V deployment fails, the status will be shown as *Incomplete* in the Cloud Hubs in Azure table. Click the **info** (i) icon to download the log file. The reason for the failure is typically at the bottom of the log file. To remove (delete) a deployment after a failure, click **Terminate**. All Azure resources created by the Orchestrator are removed. If termination fails, it indicates that an Azure resource is blocking the deletion. To view details of this Azure resource, download the log file and check the last error message. To proceed with the termination, you can manually delete this resource on the Azure Portal.

If your deployment succeeds, the status will be shown as *Deployed*.

7. After the deployment succeeds, the newly deployed EC-Vs appear on the Discovered Appliances tab in the Orchestrator UI. Please allow at least 10 minutes for EC-Vs to be discovered on the Orchestrator. On the Discovered Appliances tab, the Approve button appears after the EC-V is fully configured.

| Serial Number | Appliance | IP Address | Public IP Address... | Location | Tag | Discovered Time | Reachability | Approve | Deny | Software Version... | Model | Account Name |
|---------------|------------|--------------|----------------------|--------------------|-----------|-----------------|--------------|---------|------|---------------------|-------|-----------------------------|
| 001BBC1FEF4E | silverpeak | 10.200.30.4 | 20.21.242.120 | Doha, Ad Dawhah, (| 1_9510522 | 28-May-25 15:36 | Reachable | Approve | Deny | 8.3.0.19_85161 | EC-V | Silver Peak Eng - Dinesh... |
| 001BBC1FEF4F | silverpeak | 10.200.30.52 | 20.21.242.121 | Doha, Ad Dawhah, (| 2_3603390 | 28-May-25 15:36 | Reachable | Approve | Deny | 8.3.0.19_85161 | EC-V | Silver Peak Eng - Dinesh... |

Figure 7. Newly deployed EC-Vs appear on the Discovered Appliances tab.

The following table describes each field on the Cloud Hubs in Azure tab.

| Field | Description |
|-----------------|--|
| Deployment Name | Name given on the deployment configuration page. |

| Field | Description |
|-----------------|--|
| Virtual Network | CIDR block used for deployment. |
| Subscription | Name of the Azure subscription used to deploy the EC-Vs. |
| Instances | <p>Number of EC-V instances in the deployment. To add one or more EC-Vs to the deployment, click +Add. In the New Instance on Azure dialog box, select the availability zone to use and any optional tags to apply to the new instance.</p> <p>Max indicates that the maximum number of instances have been created for this deployment.</p> <p>If the region you selected does not support availability zones, the new Instance in Azure dialog box does not display an Availability Zone menu.</p> |
| Region | Region of the EC-V deployment. |
| Resource Group | Name of the Azure resource group that was used for the EC-V deployment. |
| Status | <p>Status of the deployment. If more information is available, an info (i) icon appears.</p> <p>Note: If the deployment is incomplete, the information dialog contains a link to download the log file and steps to resolve the issue.</p> |
| Terminate | <p>To permanently delete a deployment, click Terminate. This action deletes all resources associated with the EC-Vs, including all Azure resources.</p> <p>Clicking Terminate deletes all EC-Vs in the deployment, as the Orchestrator cannot selectively delete them. The resource group that was used for the deployment is not deleted.</p> |
| Deployment Info | To view deployment and virtual machine details, click the info (i) icon in this column. |
| Resources | To view details about each Azure resource that Orchestrator created during the deployment, click the info (i) icon in this column. |
| Comment | Comments that were added to the deployment when the EC-V was created. To edit the comment, click the edit (pencil) icon. |

Note

When deploying the EC-V from Orchestrator into a transit VNet, inbound SSH traffic to the VM is disabled by default on the MGMT0 interface. Inbound SSH access is not required for the initial configuration of the EC-V gateway. However, if you need to SSH into the EC-V gateway, you must modify the MGMT0 NSG and create a rule to allow inbound SSH traffic from your network.

This completes the deployment of EC-V from Orchestrator into a transit VNet. Your EC-V deployment now matches Figure 8.

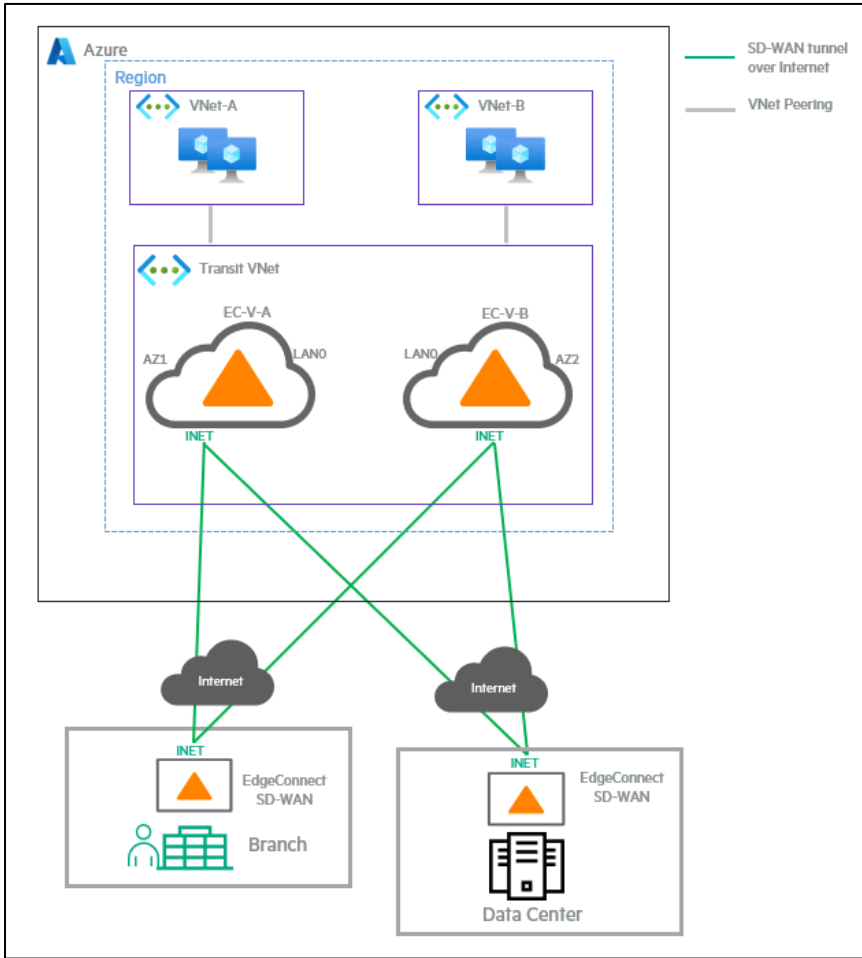


Figure 8. After deploying EC-V into a transit VNet and establishing SD-WAN tunnels.

The next step is to connect the EC-V with your choice of Azure network service and establish LAN-side connectivity. Upcoming sections cover the following EC-V high availability designs:

— [Integrate EC-V in transit VNet with Azure Standard Internal Load Balancer \(ILB\)](#)

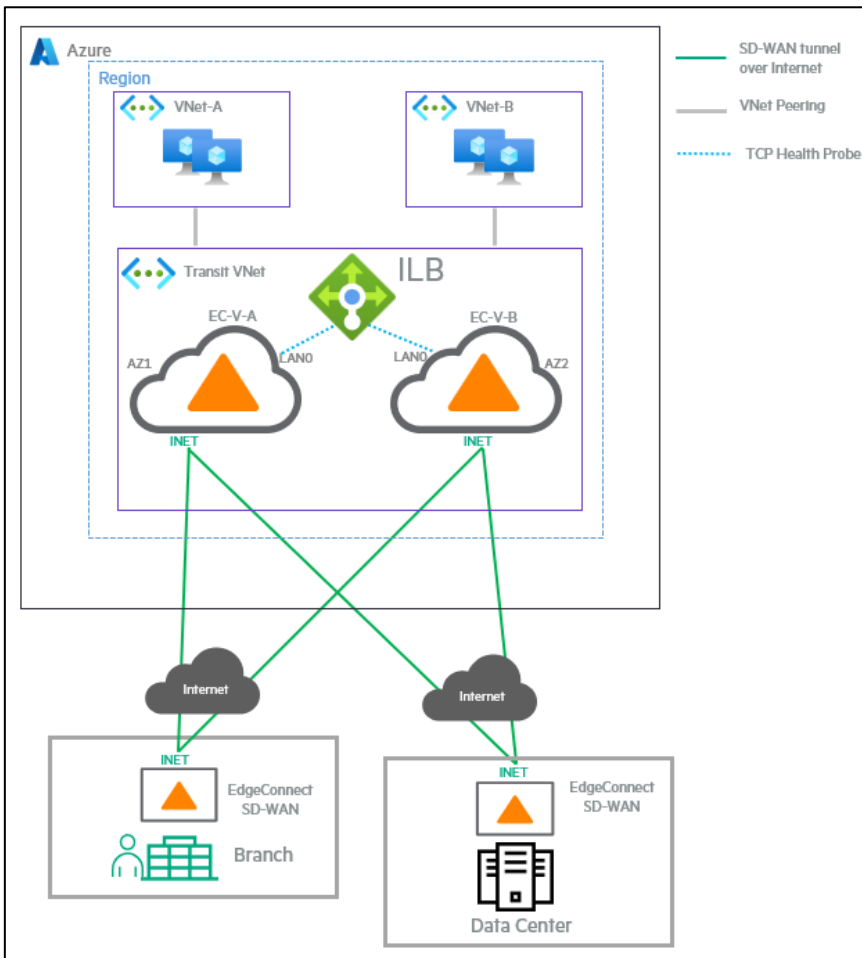


Figure 9. Integrating EC-V in transit VNet with Azure Standard Internal Load Balancer (ILB).

— [Integrate EC-V in transit VNet with Azure Route Server \(ARS\)](#)

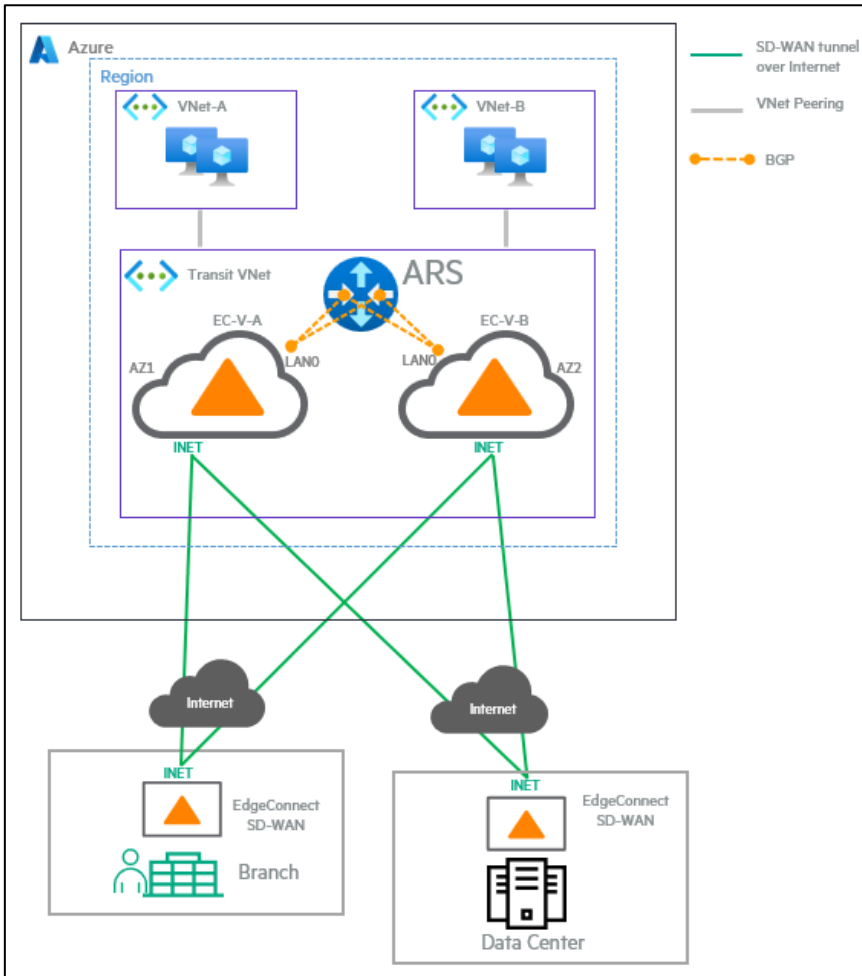


Figure 10. Integrating EC-V in transit VNet with Azure Route Server (ARS).

— [Transit VNet + Azure Virtual WAN \(vWAN\) hub design](#)

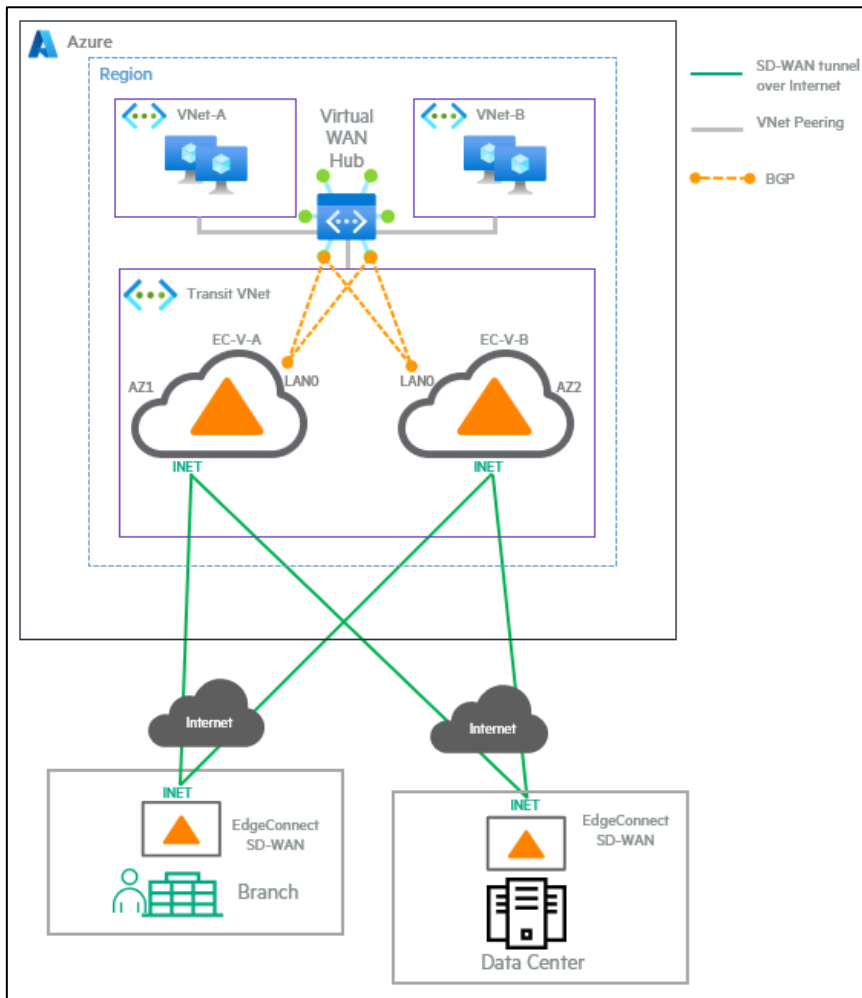


Figure 11. Transit VNet + Azure Virtual WAN Hub design.

Note

If you deployed EC-Vs into a transit VNet from Orchestrator, skip Section 3.2 (which is for users who want to deploy EC-V manually from Azure Portal) and proceed to Section [3.3](#).

3.2. Deploy EC-V into a transit VNet from Azure Portal

The easiest method to deploy an EC-V into a transit VNet is through the Orchestrator. However, some organizations may require alternative deployment methods such as using the Azure Portal, Azure CLI, Azure Resource Manager (ARM), Terraform, or other infrastructure as code (IaC) tools. Furthermore, as of December 2025, Orchestrator cannot deploy EC-V (via the Cloud Hubs in Azure feature) to Azure Government Cloud, Azure Extended Zones, or Azure China. To deploy EC-V in these environments, use the steps provided in this section using the Azure Portal.

3.2.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- A valid Azure Commercial, Azure Government, or Azure China subscription. To deploy EC-V to an Extended Zone, verify that Extended Zone regions are accessible within your Azure subscription.
- HPE Aruba Networking EdgeConnect Orchestrator
- HPE Aruba Networking EdgeConnect licenses
- Check Azure service limits.

- Before starting the deployment, it is highly recommended to check your Azure subscription quotas in the selected region for compute cores, virtual networks, and subnets to ensure that you have enough available resources to avoid deployment failures.
 - **Number of EC-Vs:** The number of EC-Vs you plan to deploy directly impacts the required number of subnets (if using an existing VNet) and compute cores.
 - **VM SKUs:** The chosen VM SKU determines the required compute cores. The recommended VM SKUs can be found at <https://arubanetworking.hpe.com/techdocs/sdwan/docs/sysreq/sys-req-html/ecv-cloud-host-sys-req>.
 - **Region:** Azure quotas are region-specific. Ensure you have the necessary quotas in the region you select for deployment. To deploy EC-V to an Extended Zone, please verify that Extended Zone regions contain the recommended VM SKUs mentioned above.
 - **Existing VNet vs. new VNet:** If using an existing VNet, you are responsible for ensuring the availability of the required subnets. If you create a new VNet for the deployment, you must ensure that the CIDR block you enter does not overlap with an existing CIDR block. Also, you must ensure that the CIDR block you assign to the VNet is sufficient for the subnets required for deploying the EC-V.
- To learn how to manage your Azure service limits, go to <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits>.

Note

When an Orchestrator deploys an EC-V, the EdgeConnect OS (ECOS) version of that EC-V will almost always be earlier than the version you want. Upon deploying and adding the EC-V into the appliance tree, you can upgrade it to the ECOS version you need. The upgrade takes only a few minutes.

3.2.2. Create an EC-V gateway in the Azure Portal

The following tasks configure an existing resource group that meets the requirements for an EC-V deployment and configures a virtual machine as an EC-V gateway.

3.2.2.1. Prepare the resource group

The following instructions prepare the resource group to create the EC-V gateway. Tasks include:

- [3.2.2.1.1 Create a virtual network and subnets](#)
- [3.2.2.1.2 Create network security](#)
- [3.2.2.1.3 Create an SSH private-public key pair](#)
- [3.2.2.1.4 Create data path interfaces](#)

Note

For these steps, you must use an SSH key generator tool to create an SSH private-public key pair. The example in this section uses PuTTYgen to create the keys. When creating a new public/private key pair for EC-V, use the EdDSA (Ed25519) algorithm.

While the following sections provide instructions to create all resources required for the EC-V—including a VNet, subnets, and NSGs—you can use any existing virtual networks (including their subnets and NSGs) and SSH key pair that meet EC-V deployment requirements.

3.2.2.1.1. Create a virtual network and subnets

These steps create a virtual network. If the virtual network you want exists, proceed to step 7 below.

1. Log in to the Azure Portal and navigate to your resource group.
2. On the menu bar, click **+Create**.
3. Enter virtual network in the search field, then select **Virtual Network** from the drop-down menu.
4. On the Virtual Network page, click **Create**.
5. On the Create virtual network page, verify that **Basics** is selected on the menu bar, and then enter the following settings:
 - a. **Subscription:** Depends on your Azure subscription.
 - b. **Resource Group:** Select your resource group.

- c. **Virtual Network Name:** Enter the label for the new virtual network.
- d. **Region:** Select the region where you want to deploy the EC-V. If you are deploying EC-V to an Azure Extended Zone, click **Deploy to an Azure Extended Zone**, and then select the appropriate Extended Zone.
6. Click **Next**.
7. On the Security tab, configure the following settings:
 - a. **Virtual network encryption:** optional
 - b. **Azure Bastion:** optional
 - c. **Azure Firewall:** optional
 - d. **Azure DDoS Network Protection:** optional
8. Click **Next**.
9. On the IP Addresses tab, the virtual network selected for an EC-V deployment requires at least two subnets for WAN0 and LAN0 interfaces. Although optional, it is recommended to create a separate MGMT0 interface and deploy it on a separate subnet from the WAN and LAN subnets. This example implementation uses three subnets.
 - a. Enter your VNet CIDR, and then select the network mask.
 - b. In the address space box in Subnets, select the default subnet.
 - c. Under Subnets, click the **edit** (pencil) icon and enter the MGMT0 subnet information:
 - I. **Subnet purpose:** Default
 - II. **Name:** Enter a name for the subnet, such as *MGMT0_Subnet*.
 - III. **Include an IPv4 address space:** Select the check box.
 - IV. **IPv4 address range:** Enter the MGMT0 subnet CIDR block.
 - V. **Starting address:** Enter the starting IP.
 - VI. **Size:** Select a subnet mask for the MGMT0 CIDR block.
 - d. Private subnet:
 - I. **Enable private subnet (no default outbound access):** Clear this check box. The MGMT0 interface needs to send outbound traffic to the Aruba Cloud Portal.
 - e. Security:
 - I. **NAT gateway:** None
 - II. **Network security group:** None
 - III. **Route table:** None
 - f. **Service Endpoint:** None
 - g. **Subnet Delegation:** None
 - h. Network Policy for Private Endpoints:
 - I. **Private endpoint network policy:** Disabled
 - i. Click **Save**.
10. Click **+ Add a subnet**, and then repeat the steps above to create subnets for WAN0 and LAN0 interfaces. Each vNIC on the EC-V needs to be deployed into a unique subnet.
11. Click **Next**, and then assign tags to your VNet if necessary.
12. Click **Review + create** to begin creating the VNet.
13. When validation passes, click **Create**.

You have successfully created your VNet.

3.2.2.1.2. Create network security groups

Network security groups (NSG) are required when creating the WAN and LAN interfaces on the Azure Portal. The MGMT0 interface's NSG is created automatically when deploying the EC-V. Therefore, it is not created in this procedure.

The following instructions create the WAN0 network security group:

1. Navigate to your resource group, and then click **+Create**.
2. In the search field, enter **network security group**, and then click the **Network security group** tile.
3. Click **Create** to open the Create Network Security Group page.
4. Verify that **Basics** is selected on the menu bar, and then enter the following settings:
 - a. **Subscription:** Depends on your Azure subscription.
 - b. **Resource Group:** Select your resource group.
 - c. **Name:** Enter the label for the new WAN0 NSG.
 - d. **Region:** Select the region where your EC-V will be created.
5. Click **Review + create** at the bottom of the page.
6. After you see the *Validation passed* message, click **Create**.
7. When the Overview page displays the message *Your Deployment is complete*, click **Go to resource**.
The Network Security Group page appears.
8. In the Settings section, click **Inbound security rules**.
9. Click **+Add**.
10. In the Add inbound security rule panel, enter the following settings:
 - a. **Source:** Select **Any**.
 - b. **Source port ranges:** Enter *****.
 - c. **Destination:** Select **Any**.
 - d. **Service:** Select **Custom**.
 - e. **Destination port ranges:** Enter *****.
 - f. **Protocol:** Select **Any**.
 - g. **Action:** Select **Allow**.
 - h. **Priority:** Enter **100**.
 - i. **Name:** Enter a descriptive name.

Note

This example allows all inbound traffic because when the EC-V is deployed, its Firewall Mode is set to *Stateful+SNAT*. Do not allow all inbound traffic on the NSG if you plan to set the EC-V's WAN interface Firewall Mode (on the Deployment page) to *Allow All*.

11. Click **Add**.
12. In the Settings section, click **Outbound security rules**.
13. Click **+Add**.
14. In the Add outbound security rule panel, enter the following settings:
 - a. **Source:** Select **Any**.
 - b. **Source port ranges:** Enter *****.
 - c. **Destination:** Select **Any**.
 - d. **Service:** Select **Custom**.
 - e. **Destination port ranges:** Enter *****.
 - f. **Protocol:** Select **Any**.
 - g. **Action:** Select **Allow**.
 - h. **Priority:** Enter **100**.
 - i. **Name:** Enter a descriptive name.

15. Click **Add**.

The following instructions create the LAN0 network security group:

1. Navigate to your resource group, and then click **+Create**.
2. In the search field, enter **network security group**.
3. Click the **Network security group** tile.
4. Click **Create**.
5. On the Create Network Security Group page, verify that **Basics** is selected on the menu bar, and then enter the following settings:
 - a. **Subscription:** Depends on your Azure subscription.
 - b. **Resource Group:** Select your resource group.
 - c. **Name:** Enter the label for the new LAN0 NSG.
 - d. **Region:** Select the region where your EC-V will be created.
6. Click **Review + create**.
7. After you see the *Validation passed* message, click **Create**.
8. When the Overview page displays the message, *Your Deployment is complete*, click **Go to resource**.
The Network Security Group page appears.
9. In the Settings section, click **Inbound security rules**.
10. Click **+Add**.
11. In the Add inbound security rule panel, enter the following settings:
 - a. **Source:** Select **IP Addresses**.
 - b. **Source IP Addresses/CIDR ranges:** Enter `10.0.0.0/8`.
 - c. **Source port ranges:** Enter `*`.
 - d. **Destination:** Select **Any**.
 - e. **Service:** Select **Custom**.
 - f. **Destination port ranges:** Enter `*`.
 - g. **Protocol:** Select **Any**.
 - h. **Action:** Select **Allow**.
 - i. **Priority:** Enter `100`.
 - j. **Name:** Enter a descriptive name.

Note

Only allow inbound traffic from Azure virtual network CIDR address ranges. For other RFC1918 addresses that need to send traffic to the LAN0 interface, create additional rules.

12. Click **Add**.

13. In the Settings section, click **Outbound security rules**.

14. Click **+Add**.

15. In the Add outbound security rule panel, enter the following settings:
 - a. **Source:** Select **Any**.
 - b. **Source port ranges:** Enter `*`.
 - c. **Destination:** Select **Any**.
 - d. **Service:** Select **Custom**.
 - e. **Destination port ranges:** Enter `*`.
 - f. **Protocol:** Select **Any**.

- g. **Action:** Select **Allow**.
- h. **Priority:** Enter 100.
- i. **Name:** Enter a descriptive name.

Note

It is safe to allow inbound traffic on the LAN interface. It is not assigned a public IP address, resulting in only workloads in your Azure environment sending traffic to the LAN0 interface.

16. Click **Add**.

3.2.2.1.3. Create an SSH private-public key pair

PuTTYGen (or your preferred SSH key generator tool) can be used to create the private-public key pair. Ensure that you create an EdDSA (ED25519)-based key pair. To create a key pair, click **Generate**, follow the tool's on-screen instructions, and then save the private key to an accessible location. Do not share your private key with others.

3.2.2.1.4. Create data path interfaces

When you create a VM in Azure using the Azure Portal, the VM initially contains only one Network Interface Card (NIC). The EC-V requires three NICs; this section explains how to create the WAN0 and LAN0 interfaces, while the MGMT0 interface is created during deployment of the EC-V. The WAN0 and LAN0 interfaces are attached to the EC-V after the EC-V is deployed.

The following steps create the data path interfaces:

1. From the Azure Portal, navigate to your resource group, and then click **+Create**.
2. Use the search field to select **network interface** from the drop-down menu, then click **Create**.
3. In the Create Network Interface menu, enter the following settings:
 - a. **Subscription:** Depends on your Azure subscription.
 - b. **Resource group:** Select your resource group.
 - c. **Name:** Enter a descriptive name for the WAN0 interface.
 - d. **Region:** Select the region where you deployed the EC-V.
 - e. **Virtual Network:** Select the virtual network where you deployed the EC-V.
 - f. **Subnet:** Select the **WAN0** subnet.
 - g. **IP version:** Select **IPv4**.
 - h. **Private IP address assignment:** Select **Dynamic**.
4. Click **Review + create**.
Azure displays the Create network interface page.
5. After you see the *Validation passed* message, click **Create**.
6. When the Overview page displays the message, *Your Deployment is complete*, click **Go to resource**.

Similarly, create the LAN0 network interface:

1. Open the **Resource Group** page to review WAN0 and LAN0 interfaces.
2. Click the **WAN0** network interface.
3. Under Settings, click **Network security group**.
4. Select the WAN0 NSG that you created earlier on the Network security group drop-down menu.
5. Using the steps above, associate the LAN0 NSG to the LAN0 interface.

You have successfully created the WAN0 and LAN0 interfaces and associated the NSGs to them.

3.2.2.2. Create an EC-V gateway

The following procedure utilizes the Azure Portal's Virtual Machine Deployment Wizard to create an EC-V gateway.

1. From the Azure Portal, navigate to your resource group, and then click **+Create**.
2. Use the search field to select **HPE Aruba Networking EdgeConnect SD-WAN** from the drop-down menu.

3. In the Select a software plan drop-down menu, select an EdgeConnect version, then click **Create**. Unless otherwise instructed, select the most recent software version.

The Create Virtual Machine page displays configuration options. The menu bar at the top of the page accesses the multiple pages that configure the virtual machine.

4. Verify that **Basics** is selected on the menu bar, and then enter the following settings:
 - a. **Subscription:** Depends on your Azure account.
 - b. **Resource group:** Select your resource group.
 - c. **Virtual machine name:** Enter a descriptive name for your EC-V.
 - d. **Region:** Select the Azure region to deploy your EC-V. If you are deploying EC-V to an Azure Extended Zone, click **Deploy to an Azure Extended Zone**, and then select the appropriate Extended Zone.
 - e. **Availability options:** The selection depends on whether your deployment is a single EC-V deployment or an HA EC-V deployment.
 - I. **Single EC-V Deployment:** Select **No infrastructure redundancy required**.
 - II. **HA EC-V Deployment:** Select **Availability Zone**, and then select unique availability zones (AZ) for each EC-V. Use a different AZ for each EC-V to ensure that EC-Vs in other AZs continue when one EC-V fails due to an AZ failure. When deploying multiple EC-Vs in a region that does not support AZs, place all EC-Vs in a single availability set. An availability set ensures that deployed EC-Vs are distributed in a cluster across multiple, isolated hardware nodes to protect against Azure hardware failures.
 - III. If you are deploying EC-V to an Azure Extended Zone, select **No infrastructure redundancy required**. Azure Extended Zones do not support availability zones.
 - f. **Security type:** Select **Standard**.
 - g. **Image:** Select the latest EdgeConnect image.
 - h. **Run with Azure Spot discount:** Clear the check box (default value).
 - i. **Size:** The size should match the WAN bandwidth and number of interfaces required on the EC-V. The list of recommended instance types for EC-V is available at <https://arubanetworking.hpe.com/techdocs/sdwan/docs/sysreq/sys-req-html/ecv-cloud-host-sys-req>.
 - j. **Authentication type:** Select SSH public key. The private key associated with the public key entered here is required to log in to the appliance from your SSH client.

ⓘ Important

EC-V does not support password-based authentication in Azure.

- k. **Username:** Enter any valid username except *admin*. The username entered here is used only on Azure to complete the Basics page. The username *admin* is used to log in to the EC-V via SSH or HTTPS.
 - l. **SSH public key source:** Select **Use existing public key**.
-

Note

EdgeConnect does not support the *Generate new key pair* or *Use existing key stored in Azure* options.

- m. **SSH public key:** Enter the single-line SSH public key from the key pair you created earlier.
-

ⓘ Important

Best practice for entering a public key is to open the application that created the key pair, use the private key to regenerate the public key, and then copy it directly from the application into Azure. Saving the private key to a text file and then copying that file into Azure may introduce another text line, causing the appliance to reject the key. As shown in the first image below, use a single-line public key. Do not use a multi-line public key (second image).

Use this:

```
1 ssh-ed25519
   AAAAC3NzaC1lZDI1NTE5AAAAIOFDa8CNR5DmJE3EgIOBpYyyL4m1yqaRVo/+XBgISULs
   eddsa-key-20251015
```

Not this:

```
1 ----- BEGIN SSH2 PUBLIC KEY -----
2 Comment: "eddsa-key-20251015"
3 AAAAC3NzaC1lZDI1NTE5AAAAIOFDa8CNR5DmJE3EgIOBpYyyL4m1yqaRVo/+XBgI
4 SULs
5 ----- END SSH2 PUBLIC KEY -----
6
```

After you deploy the EC-V and give it an admin password, you can log in to the EC-V via SSH with either the admin password or the private key.

- n. Under Public inbound ports, select **Allow selected ports**.
- o. Under Select inbound ports, select **SSH and HTTPS**.
- p. Click **Next : Disks >**.
The Disks page appears.
5. Enter the following settings.
 - a. **OS disk size:** Leave the default value (60 GiB).
 - b. **OS disk type:** Select **Premium SSD**.
 - c. **Delete with VM:** Select the check box.
 - d. **Key management:** Select **Platform-managed key**.
 - e. **Enable Ultra Disk compatibility:** Clear the check box.
 - f. **Encryption type:** Select **(Default) Encryption at-rest with a platform-managed key**.
 - g. Leave all other settings as is, and then click **Next : Networking >**.
The Networking page appears.
6. Enter the following settings:
 - a. **Virtual Network:** Select the virtual network created earlier.
 - b. **Subnet:** Select a subnet to deploy the MGMT0 interface. WAN0 and LAN0 interfaces will be added on the EC-V after the appliance is created and discovered in the Orchestrator.
 - c. **Public IP:** Click **Create New**, and then enter the following in the Create public IP address dialog box:
 - I. **Name:** Enter a descriptive name.
 - II. **SKU:** Select **Standard**.
 - III. **Assignment:** Select **Static**.

Note

Microsoft recommends using Standard SKU public IP addresses. They are zone-redundant by default and work with the Standard SKU Load Balancer.

 - IV. **Routing preference:** Select your preferred routing option.
 - V. **Availability zone:** Select **Zone-redundant**.
 - VI. **NIC network security group:** Select **Advanced**.
- d. **Configure network security group:** It is strongly recommended to tighten security rules to allow incoming traffic from only your network. After approving the EC-V on Orchestrator, you can block inbound access to the MGMT0 interface.

ⓘ Important

DO NOT select the NSGs created earlier in the document. They are created specifically for the WAN0 and LAN0 interfaces. The best practice is to create a separate NSG for the MGMT0 interface.

7. **Delete NIC when VM is deleted:** Select
8. **Accelerated networking:** Select

Note

This parameter is locked (cleared) for instance types that do not support Accelerated Networking.

9. **Load balancing options:** None
10. Click **Next: Management >**.
The Management page appears.
11. Enter the following settings:
 - a. **Enable System assigned managed identity:** Clear (off)
 - b. **Login with Azure AD (Preview):** Clear (off)
 - c. **Enable auto-shutdown:** Clear (off)
 - d. Click **Next: Monitoring >**.
The Monitoring page appears.
12. Enter the following settings:
 - a. **Enable recommended alert rules:** Optional
 - b. **Boot Diagnostics:** Select **Enable with managed storage account (recommended)**.
 - c. **Enable OS guest diagnostics:** Clear (off)
 - d. **Enable application health monitoring:** Clear (off)
 - e. Click **Next: Advanced >**.
The Advanced page appears.

Note

The Custom data section on this page is useful when deploying an EC-V, as it allows you to provide configuration information that runs automatically at first boot. Under Custom Data, by entering your account name, account key, and subnet information as shown below, the EC-V automatically assigns its licenses and sets MAC addresses for interfaces such as MGMT0, WAN0, and LAN0 upon first boot. As a result, the EC-V appears on the Discovered Appliance page within 10 minutes after deployment. Ensure that you enter the same subnet and mask used when creating your subnets earlier. If additional interfaces are required beyond MGMT0, WAN0, and LAN0, you can enter up to 16 subnets in the format shown below. If you do not enter the account name and account key under Custom Data, you must manually license the EC-V. This involves SSHing into the EC-V using your private key, creating a password for the admin user in the CLI, and then logging in to the Appliance Manager Web User Interface to enter the account name and key. To avoid these additional steps, it is recommended to enter the account name and account key under Custom Data. When creating an EC-V gateway, if the MGMT0, WAN0, and LAN0 subnet information was correctly entered under Custom Data, the MAC addresses for WAN0 and LAN0 is automatically assigned when the EC-V powers on and the interfaces are attached. If this information was not provided, you must manually assign the MAC addresses by navigating to Configuration > Interfaces in the EC-V.

These MAC addresses can be found in the Azure Portal under the Properties section of the WAN0 and LAN0 network interface pages.

Enter the following content in the Custom data section. The tag is optional and can be left blank, as illustrated in the example below:

```
{"account": "Your Account Name", "key": "Your Account Key", "tag": "Your tag (optional)", "mgmt0": "Your MGMT0 subnet with mask", "wan0": "Your WAN0 subnet with mask", "lan0": "Your LAN0 subnet with mask"}
```

Example:

```
{"account": "My Account Name", "key": "LQ08wU0B8822RSk3Gqp1ctk500UiTrEL", "tag": "",  
"mgmt0": "10.10.0.0/28", "wan0": "10.10.0.16/28", "lan0": "10.10.0.32/28"}
```

13. Click **Review + create**.
14. After you see the *Validation passed* message, proceed to the next step.
15. Enter a Preferred e-mail address, a Preferred phone number, and then click **Create**.

The Azure Dashboard appears as the EC-V begins to deploy.

Your EC-V is created successfully. After an EC-V gateway is created, it communicates with the HPE Aruba Networking Cloud Portal to coordinate it with your network. Orchestrator then displays the EC-V as a new appliance that is ready for addition to the SD-WAN fabric. If you entered the account name and account key as mentioned above, your EC-V gateway will display under the Discovered Appliances tab on the Orchestrator within 10 minutes. In the next section, you will add the EC-V to the SD-WAN fabric and configure it.

3.2.2.3. Configure the EC-V gateway

In this section, you will perform the following tasks to configure the EC-V:

- [3.2.2.3.1 Add the EC-V to the SD-WAN fabric](#)
- [3.2.2.3.2 Assign a public IP address to WAN0 interface](#)
- [3.2.2.3.3 Enable IP forwarding on LAN0 network interface](#)
- [3.2.2.3.4 Attach WAN0 and LAN0 network interfaces to the EC-V](#)
- [3.2.2.3.5 Configure the EC-V for In-line Router Mode](#)
- [3.2.2.3.6 Enable Accelerated Networking on WAN0 and LAN0 interfaces](#)

3.2.2.3.1. Add the EC-V to the SD-WAN fabric

1. Open the **Discovered Devices** page in the Orchestrator UI.
2. Find your EC-V and click **Approve**.

The Upgrade Appliance screen appears. You can either skip for now or upgrade to the required software version. If deploying multiple EC-Vs, it is better to skip and upgrade all appliances simultaneously after adding them to the Orchestrator tree. If you click **Upgrade & Reboot Appliance**, you must wait for the upgrade process to complete and the Orchestrator to rediscover the EC-V before you can proceed.

The Appliance Wizard appears.

3. Configure the Hostname, Group, Admin Password, and Location.
4. (Optional) If the EC-V will act as a hub for SD-WAN spokes, select the **Hub Site** check box.
5. Enter a name for the appliance in the Appliance field, select a group in which the EC-V will be placed, and then enter a new admin password.
6. Fill in other details, such as location and site name. When deploying multiple EC-Vs in the same Azure region, ensure that you enter the same site name on each EC-V to prevent them from forming tunnels with each other. Orchestrator does not create underlays between EC-Vs with identical site names.
7. Click **Next**.
8. On the Deployment Profile page, DO NOT change the deployment mode from Server to In-line Router Mode yet. You will do this AFTER adding the EC-V to SD-WAN fabric and assigning MAC addresses.
9. Click **Next**.
10. On the Loopback Interfaces page, click **Next**.
11. On the Add Local Routes page, select the **Use SD-WAN Fabric Learned Routes** check box, and then clear the **Automatically include local subnets** check box.
12. On the Add Business Intent Overlays to this Site page, click **Apply**. You can configure the Business Intent Overlays and Template Groups after the EC-V is fully provisioned.
13. After the configuration is applied successfully, click **Close**.

3.2.2.3.2. Assign a public IP address to WAN0 interface

A public IP address is required on the WAN0 interface to establish an SD-WAN tunnel over the internet. The following steps assign a public IP address to the WAN0 interface.

If you have an ExpressRoute circuit and intend to establish an SD-WAN tunnel over it along with your SD-WAN tunnel over the internet, create an additional WAN interface dedicated to all tunnels to and from the ExpressRoute circuit.

1. On the Resource group page, open the WAN0 Network Interface page by clicking its name.
2. In the Settings section, click **IP configurations**.
3. Click **ipv4config**.
4. On the ipv4config dialog box, under Public IP address settings, click **Associate public IP address**.
5. Below the Public IP address field, click **Create a public IP address link**.
 - a. **Name:** Enter a descriptive name for the IP Address.
 - b. **SKU:** Standard
 Microsoft recommends Standard SKU public IP addresses; they are zone-redundant by default and work with the Standard SKU Load Balancer.
 - c. **Assignment:** Select **Static**.
 - d. Click **OK**.

6. Click **Save** to return to the Resource Group page.

3.2.2.3.3. Enable IP forwarding on LAN0 network interface

The following steps enable IP forwarding on the LAN0 network interface:

1. Open the Network Interface page for the LAN0 interface by clicking its name from the Azure Portal Resource group page.
2. In the Settings section, click **IP configurations**.
3. Select the **Enable IP forwarding** check box.
4. Click **Apply**.

The Resource Group page returns.

3.2.2.3.4. Attach WAN0 and LAN0 network interfaces to the EC-V

The EC-V must be powered off before attaching additional interfaces to the appliance. The following steps attach WAN0 and LAN0 network interfaces to the EC-V:

1. From the Azure Portal Resource group page, open the Virtual machine page for the EC-V.
2. Select the EC-V, and then click **Stop** to power it off.
3. If prompted, reserve the public IP address. Do not proceed until the VM has stopped.
4. Under the Networking section, select **Network settings**.
5. On the Network settings page, click **Attach network interface**.
6. Select **WAN0 network interface**, and then click **OK**.
7. Click **Attach network interface**.
8. Select **LAN0 network interface**, and then click **OK**.
9. Return to the Virtual machine page, and then click **Start** to power on the EC-V. Do not proceed until the status is Running.

3.2.2.3.5. Configure the EC-V for In-line Router Mode

After you attach the WAN0 and LAN0 network interfaces to the EC-V, you are ready to configure it in In-line Router Mode. By default, your EC-V comes in Server Mode. Perform the following steps to change the mode from Server to In-line Router Mode:

1. Log in to Orchestrator.
2. Select the EC-V in the appliance tree, and then navigate to **Configuration > Networking > Interfaces**.

3. Ensure that the MGMT0, WAN0, and LAN0 MAC addresses are properly assigned. If the MAC addresses are not properly assigned, assign them to the correct interfaces before you move to the next step.
4. Select the EC-V in the appliance tree, and then navigate to **Configuration > Networking > Deployment**.
5. Configure deployment parameters for the EC-V:
 - a. Click the **edit** (pencil) icon of a deployment row to open the Deployment edit dialog box.
 - b. Click **Router**.
 - c. Under LAN0 IP/Mask, enter the LAN0 interface private IP address and subnet mask.
 - d. Under WAN0 IP/Mask, enter the WAN0 interface private IP address and subnet mask.
 - e. Under LAN0 Next Hop, enter the first IP address of the LAN0 subnet address space. Azure sets the subnet's first IP address as its gateway.
 - f. Under WAN0 Next Hop, enter the first IP address of the WAN0 subnet address space.
 - g. Enter the Total Outbound and Total Inbound bandwidth (Kbps) for the WAN0 interface, and then click **ΣCalc**.
 - h. Set WAN0 Firewall Mode to **Stateful+SNAT**.
 - i. Under WAN Next Hop, click **Not behind NAT**.
 - j. On the NAT Settings dialog box, and select **NAT**. This allows Orchestrator to use the WAN0 public IP address as the tunnel endpoint when establishing underlays to the WAN0 interface.
 - k. Click **Apply**.
 - l. When prompted, click **Apply & Reboot**.
 - m. After the VM reboots, navigate to **Configuration > Networking > Interfaces** and verify that the WAN0 public IP address appears in the table.
 - n. Navigate to **Administration > Software > Upgrade > Upgrade Appliances**, and then upgrade the EC-V appliance to the software version you want. Refer to the release notes for information about available software versions.

3.2.2.3.6. Enable Accelerated Networking on WAN0 and LAN0 interfaces

Accelerated Networking, also known as Single Root I/O Virtualization (SR-IOV), improves networking performance on an EC-V gateway. To attain the maximum throughput on the EC-V gateway, you must ensure that Accelerated Networking is enabled on all data path interfaces.

Perform the following steps to check if Accelerated Networking is enabled on data path interfaces. If it is disabled, perform the Azure CLI command shown below to enable it.

Note

If you deployed the EC-V gateway using Orchestrator, Accelerated Networking is automatically enabled on all data path interfaces.

1. In the Azure Portal, click the **Cloud Shell** icon on the righthand side of the top menu bar to open the Cloud Shell.




Figure 12. Clicking the Cloud Shell icon on the Azure Portal.

2. Verify that Bash is selected on the Cloud Shell drop-down menu, as shown in **Figure 13**.

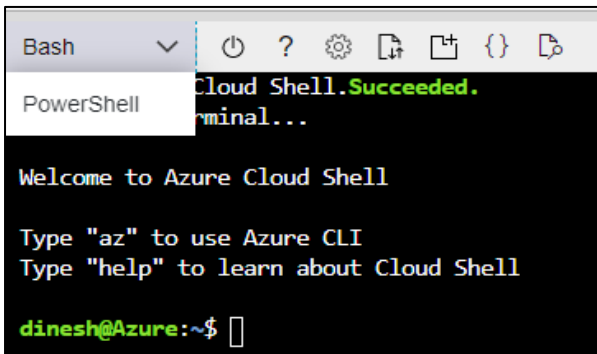


Figure 13. Selecting Bash from the Cloud Shell dropdown menu.

3. Run the following command to display current WAN0 NIC settings:

```
az network nic show --name <WAN0 NIC name> --resource-group <resource group name>
```

Replace *<WAN0 NIC name>* with the name that was assigned to the WAN0 NIC. Replace *<resource group name>* with the name of the resource group.

Figure 14 shows that Accelerated Networking is disabled on the WAN0 NIC.

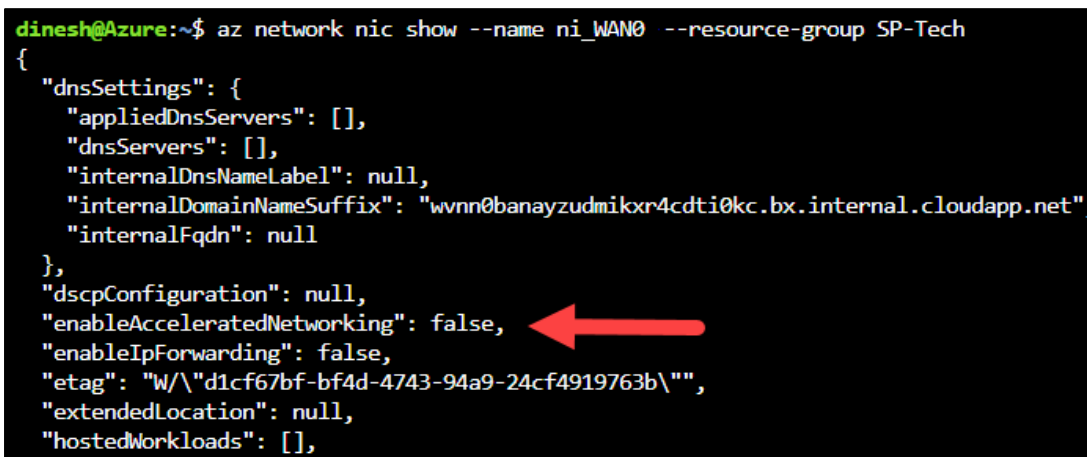


Figure 14. Confirming that Accelerated Networking is disabled on the WAN0 NIC.

4. Run the following command to enable Accelerated Networking on the WAN0 NIC:

```
az network nic update --accelerated-networking true --name <WAN0 NIC name> --resource-group <resource group name>
```

Replace *<WAN0 NIC name>* with the name assigned to the WAN0 NIC. Replace *<resource group name>* with the name of the resource group.

Figure 15 shows that Accelerated Networking is enabled on the WAN0 NIC now.

```
dinesh@Azure:~$ az network nic update --accelerated-networking true --name ni_WAN0 --resource-group SP-Tech
{
  "dnsSettings": {
    "appliedDnsServers": [],
    "dnsServers": [],
    "internalDnsNameLabel": null,
    "internalDomainNameSuffix": "wvnr0banayzudmikr4cdti0kc.bx.internal.cloudapp.net",
    "internalFqdn": null
  },
  "dscpConfiguration": null,
  "enableAcceleratedNetworking": true,
  "enableIpForwarding": false,
  "etag": "W/\c8f8932f-26ce-4fa5-aa0c-fc1c0a6506cd\"",
  "extendedLocation": null,
}
```

Figure 15. Confirming that Accelerated Networking is enabled on the WAN0 NIC.

5. Similarly, enable Accelerated Networking on all other data path interfaces attached to the EC-V.

You have successfully deployed and configured your EC-V gateway. Your EC-V deployment now matches **Figure 16**.

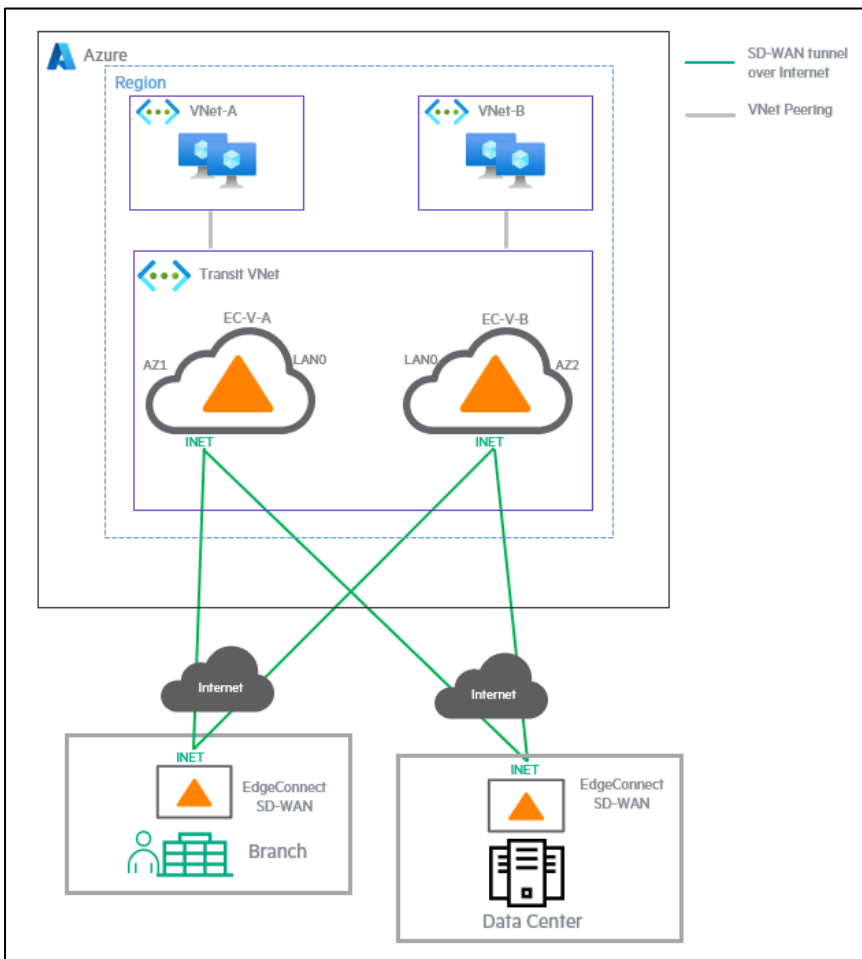


Figure 16. After deploying EC-V into a transit VNet from Azure Portal and establish SD-WAN tunnels.

The next step is to integrate EC-V with your choice of Azure-native service. You have the option to do this using Orchestrator or manually from Azure Portal. Each method is covered in the following sections:

— [Integrate EC-V in transit VNet with Azure Standard Internal Load Balancer \(ILB\)](#)




Figure 17. Integrating EC-V in transit VNet with Azure Standard Internal Load Balancer (ILB) (2).

— [Integrate EC-V in transit VNet with Azure Route Server \(ARS\)](#)




Figure 18. Integrating EC-V in transit VNet with Azure Route Server (ARS) (2).

— [Transit VNet + Azure Virtual WAN \(vWAN\) hub design](#)

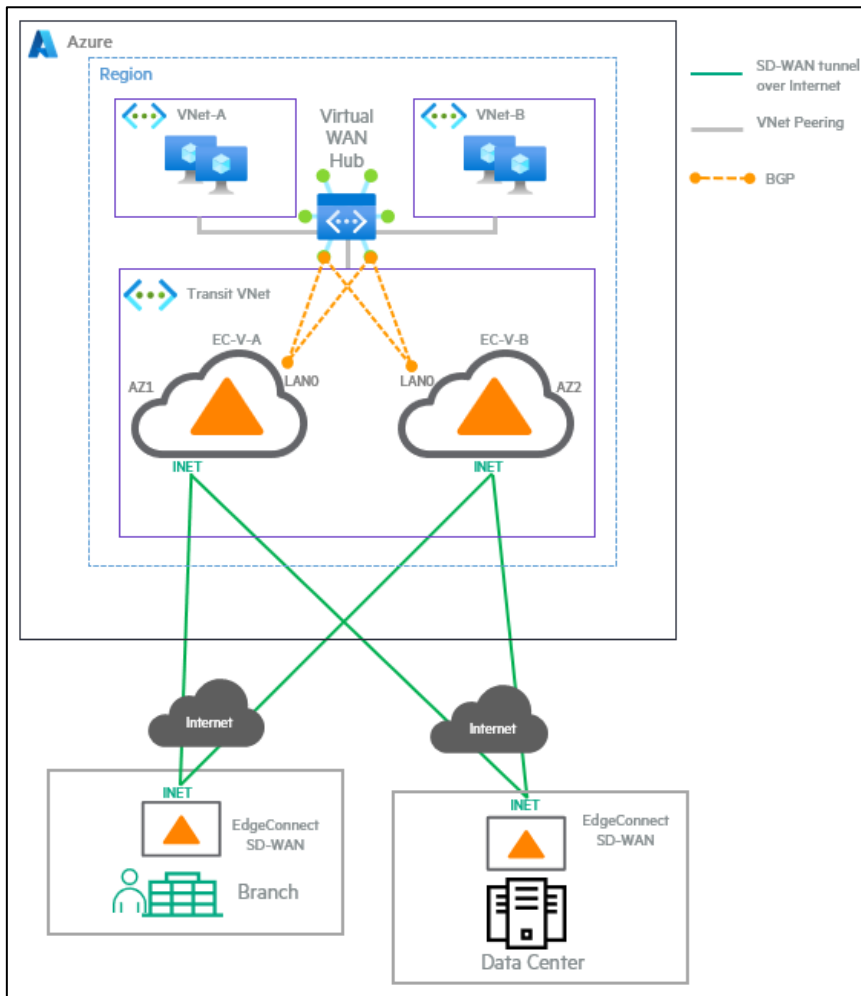


Figure 19. Transit VNet + Azure Virtual WAN Hub design (2).

3.3. Deploy EC-V into a transit VNet in an Azure Extended Zone from Azure Portal

Important

Only use this section if you want to deploy EC-V in an Azure Extended Zone. If you do not intend to deploy EC-V in an Azure Extended Zone, skip to [4. Integrate EC-V in transit VNet with Azure Standard Internal Load Balancer \(ILB\)](#).

EC-V can be deployed as a virtual machine to an Azure Extended Zone, leveraging the same virtualization capabilities available in standard Azure regions. By deploying EC-V in an Azure Extended Zone (such as Los Angeles or Perth), organizations can achieve significantly reduced latency for branch office connectivity to cloud resources compared to traditional deployments in standard Azure regions. This is particularly valuable for latency-sensitive workloads that require real-time responses. Organizations operating in metropolitan areas with Azure Extended Zones can now place EC-V gateways directly in their locality, eliminating geographic distance as a latency factor. This enables enterprises such as those with physical operations in Los Angeles to connect directly to the Azure Los Angeles Extended Zone, avoiding the latency penalty of connecting through distant Azure regions such as West US (San Francisco area) or West US3 (Phoenix). Additionally, this allows organizations to extend their existing SD-WAN fabric directly into Azure Extended Zones, creating low-latency, secure tunnels to Azure workloads while maintaining unified overlay management and Business Intent policies.

3.3.1. Things to keep in mind when deploying EC-V to an Extended Zone

— **Subscription registration requirement:** The most fundamental difference is that Azure Extended Zones require explicit subscription registration before deployment is possible. By default, this capability is not enabled on any subscription.

- **Deployment parameter requirements:** When deploying EC-V to an Extended Zone, you must specify the edge-zone location in your deployment configuration (on Azure Portal).
- **Parent region control plane architecture:** Azure Extended Zones implement a fundamentally different architecture than standard regions. Only the data plane (the EC-V) runs in the Extended Zone, while the control plane remains in the parent region.
- **Restricted virtual machine SKU availability:** Only a subset of VM SKUs is available in Azure Extended Zones. Before you deploy EC-V, ensure that the required VM SKUs are available.
- **Availability zone unavailability:** Azure Extended Zones do not support availability zones. When deploying VMs in Extended Zones, you must select **No infrastructure redundancy required** as the availability option. Standard Azure regions provide availability zones for infrastructure redundancy and fault tolerance. This architectural difference means Extended Zone deployments cannot achieve the same level of distributed resilience within a single location.
- **Network architecture constraints:** Azure Extended Zones support a limited set of networking features compared to standard regions. While Express Route, Standard Load Balancer, VNet peering, and Private Link are supported, other networking services such as ARS and vWAN hub are not supported.

For more information about deploying a VM into an Azure Extended Zone, go to <https://learn.microsoft.com/en-us/azure/extended-zones/deploy-vm-portal>.

4. Integrate EC-V in transit VNet with Azure Standard Internal Load Balancer (ILB)

This section provides detailed instructions for connecting Azure EC-Vs to an ILB.

4.1. EC-V with ILB architecture

This section explains how the ILB distributes traffic across multiple EC-V instances to ensure high availability and scalability, highlights key limitations and design considerations, and describes how EC-V failures are detected and mitigated. It also covers ILB pricing, deployment topology, and the options for deploying the ILB either manually or through the Orchestrator.

For details about Azure load balancers, go to <https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-overview>.

4.1.1. Horizontal scaling of EC-Vs with ILB

Azure ILB distributes outbound (LAN-to-WAN) flows that arrive at the load balancer's frontend private IP address to backend pool instances. Backend pool instances are the VMs to which the load balancer forwards traffic. In this example design, the backend pool instances are the EC-V gateways deployed in Azure. Azure offers two load balancer SKUs: Basic and Standard. Although the Basic Load Balancer is free, it is not recommended for the following reasons:

1. It does not support TCP and UDP flows on all ports simultaneously.
2. It cannot forward traffic to VMs in different availability zones.

Azure ILB operates on Layer 4 of the OSI model. In this example deployment, the ILB is exposed only to Azure workloads. It distributes outbound flows from these workloads to backend pool instances (EC-Vs) based on load-balancing rules and health probes. Azure ILB supports forwarding outbound traffic to multiple healthy EC-V instances, ensuring high availability and increased throughput. If an EC-V in the backend pool fails to respond to health probes, the ILB stops forwarding traffic to that instance. Inbound (WAN-to-LAN) traffic to Azure (from remote sites or other Azure regions) bypasses the ILB and is delivered directly to the destination spoke VNet where the workload resides. Azure guarantees that an ILB serving two or more healthy VM instances is available 99.99% of the time.

4.1.2. ILB limitations

When planning your deployment, be aware of the following ILB limitations:

- The EC-V gateway can only be deployed in active-active failover mode with Azure Standard Internal Load Balancer. Active-standby is not a valid option because Azure load balancer session persistence algorithms do not support sending outbound (Azure to on-premises) traffic from all Azure workloads to an EC-V and failover traffic to another EC-V only when the active EC-V fails.
- Azure ILB does not support BGP. As a result, to advertise Azure virtual network address ranges (or subnets) to the SD-WAN fabric (remote EdgeConnect devices), you must create static routes in your Azure EC-Vs. Similarly, to forward outbound traffic from the Azure workloads (in spoke VNets) to the ILB, you must create static routes (UDRs) in each spoke VNet's route table. These static routes are shown in **Figure 20**.
- To view other limitations on the Azure ILB, go to <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits#standard-load-balancer>.

4.1.3. Detection of EC-V failures

TCP health probes on the ILB monitor the LAN0 interface of an EC-V to detect failures when it stops responding. Additionally, users can create an IP SLA rule on the EdgeConnect to ensure that ILB health probes fail if the EC-V loses internet connectivity. This prevents outbound traffic from Azure being sent to an EC-V that cannot reach the internet, avoiding blackholed traffic. The creation of the IP SLA rule, including the selected monitor and the configuration of the down action, is explained in detail in Section [4.2.12](#) (for automated deployment) and in Section [4.3.7](#) (for manual deployment).

4.1.4. Azure ILB pricing

For pricing details, go to <https://azure.microsoft.com/en-us/pricing/details/load-balancer>.

4.1.5. Topology

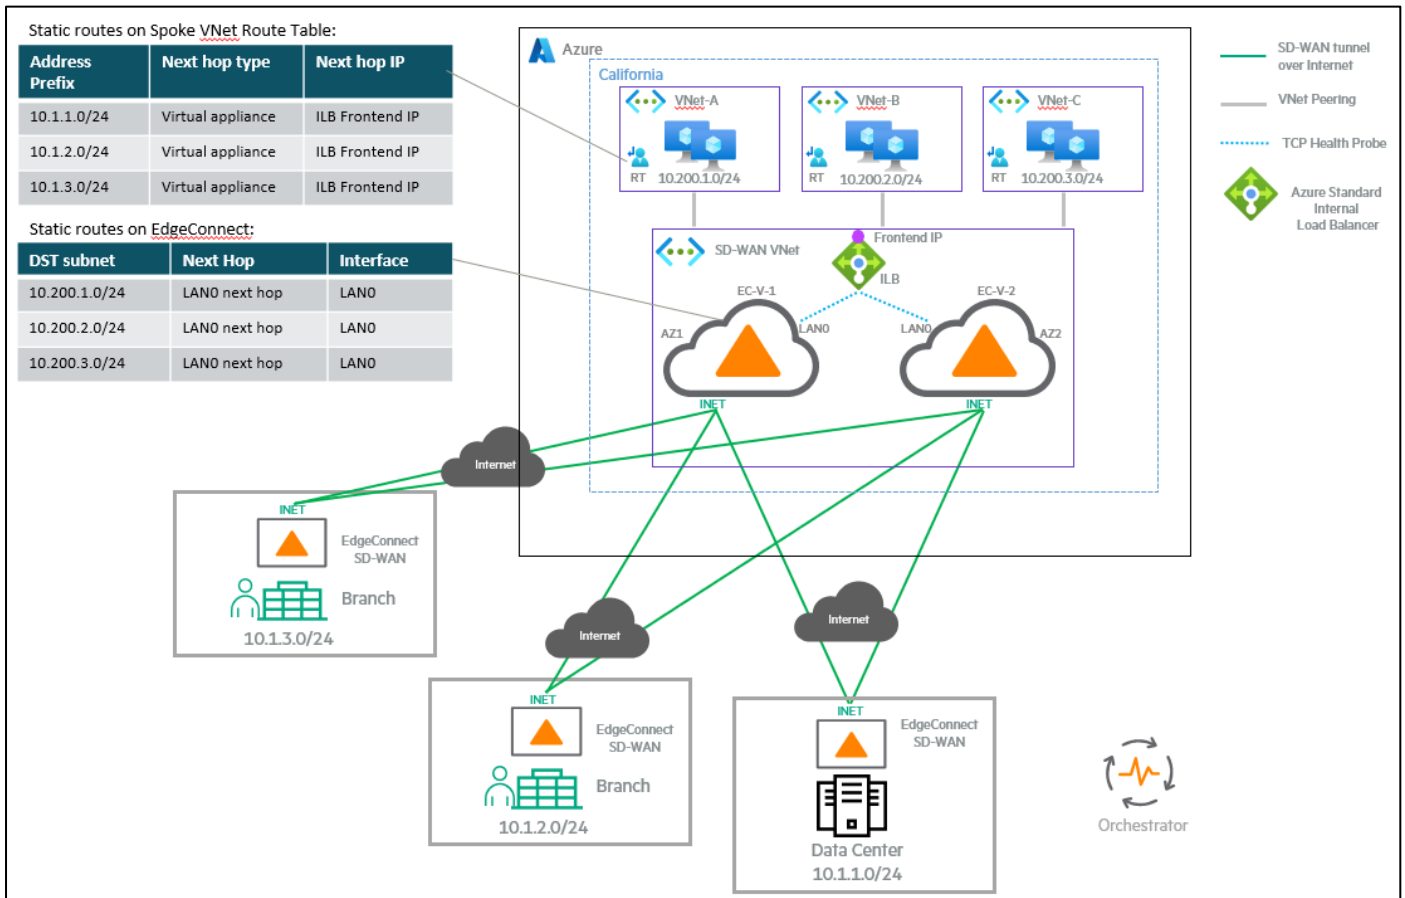


Figure 20. Topology of an EC-V deployment with Azure Standard Internal Load Balancer (ILB).

The diagram above showcases a hybrid cloud architecture leveraging the EC-V and Azure ILB to deliver secure, efficient, and reliable connectivity. It shows how branch offices and a data center can securely and efficiently connect to resources hosted within Azure VNets using EdgeConnect instances.

Below are the components of the architecture:

- **Azure VNets (VNet A, VNet B, VNet C):** Represent isolated networks within Azure, hosting various application workloads.
- **EdgeConnect SD-WAN:** Deployed both on-premises (branches, data center) and as a virtual appliance (EC-V) within Azure. It provides intelligent traffic steering, security, and traffic optimization.
- **Azure ILB:** Distributes traffic across multiple EC-V instances for high availability.
- **Spoke VNet route table:** Manages routing within the Azure VNets, directing traffic to the EC-V via the ILB.
- **EdgeConnect route table:** Controls routes advertised to and from the Azure network.
- **Orchestrator:** The central management platform for EdgeConnect, enabling centralized configuration and monitoring.
- **Internet:** Used as the transport medium for connecting branches and the data center to Azure. Although the diagram only shows tunnel over the internet, EdgeConnect also supports establishing SD-WAN tunnels over ExpressRoute.

4.1.6. Decide whether to deploy ILB from Orchestrator or manually from Azure Portal

After deploying the EC-V in a transit VNet, you can manually deploy the ILB from Azure Portal (or using your preferred Infrastructure-as-Code tool) into the same VNet. Subsequently, you can use the Orchestrator to automate connectivity between the ILB and the Azure EC-Vs. The creation of the ILB and the establishing connectivity from EC-V to ILB are decoupled processes. Therefore, deploying the ILB through the Orchestrator is not required to establish LAN-side connectivity from EC-V to ILB.

However, manually deploying an ILB and establishing connectivity to the Azure EC-Vs from the Azure Portal involve several steps, such as creating a subnet for ILB, deploying the ILB resource, creating a TCP health probe on the ILB, creating a load-balancing rule on the ILB, creating a static route on the EC-V gateway to respond to the health probes from the ILB, and adding the EC-V gateways into the backend pool. All these tasks are automated when you deploy the ILB and set up LAN-side connectivity from the Orchestrator UI.

4.2. Integrate EC-V with ILB using Orchestrator

This section provides steps for deploying an ILB from the Orchestrator and integrating it with Azure EC-V gateways.

4.2.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- **Orchestrator version:** 9.6.0 or later.
- **ECOS version:** 9.3.3 or later.
- Deploy EC-Vs in a transit VNet, as illustrated in Section [3.1](#) or [3.2](#). Ensure that the transit VNet has enough space to create at least a /27 subnet to accommodate the ILB resource.
- **Azure account configuration:** You must have a valid Azure subscription configured in Orchestrator with the necessary permissions to deploy an ILB.
- **EC-V deployment:** The EC-V appliances must be deployed in the same Virtual Network (VNet) where the ILB will be created.

4.2.2. Add the Azure Subscription to Orchestrator

1. In the Orchestrator user interface, navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
2. Click **Subscription**.

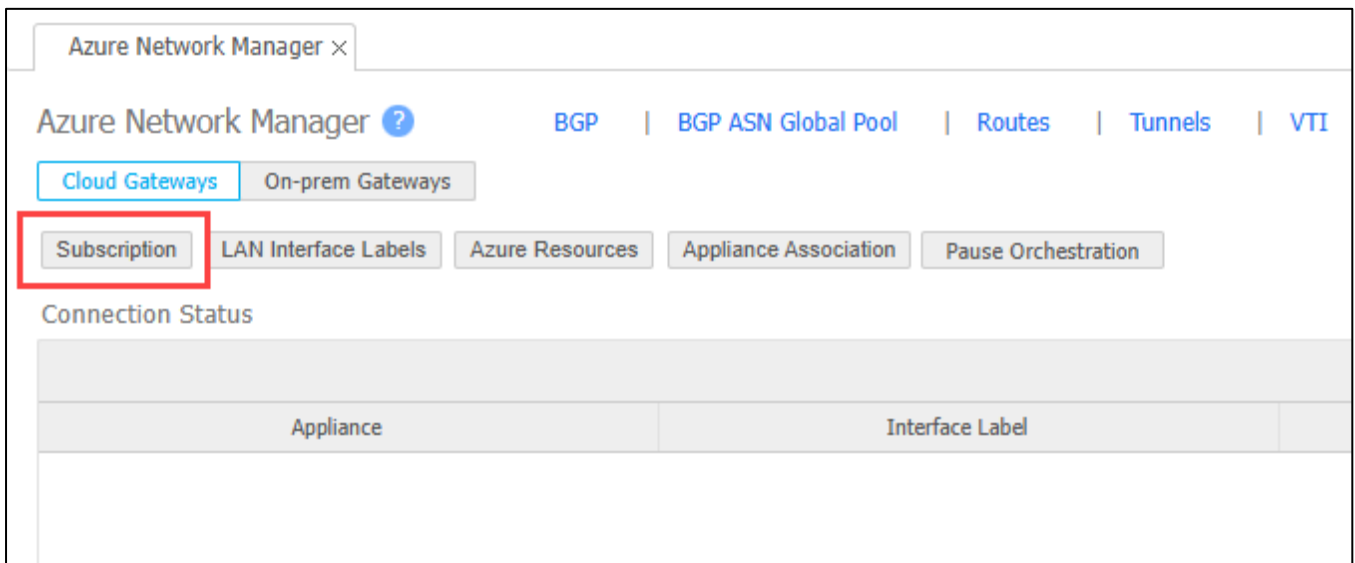


Figure 21. Clicking the Subscription tab in Azure Network Manager.

The Subscription for Azure Cloud Gateways LAN-side Automation dialog box appears.

3. If an Azure subscription has already been added to the Orchestrator (used for automated EC-V gateway deployment), it is possible to add the necessary additional permissions for ILB deployment to its associated custom role. This is required because the permissions that permit the creation of EC-V gateways are not sufficient to deploy an ILB. This is described in section [4.2.2.1 Modify an existing custom role to add the permissions needed for ILB deployment](#). If you do not have an existing Azure subscription in the Orchestrator, go to section [4.2.2.2 Add a new custom role and the permissions needed for ILB deployment](#) and add the permissions required for the ILB deployment.

Note

If any modifications you make (such as creating or deleting subnets) in Azure are not reflected on the Orchestrator after you add the Azure subscription to Orchestrator, click the **refresh** icon, as shown in Figure 22. Clicking the refresh icon retrieves the latest Azure configuration on demand.

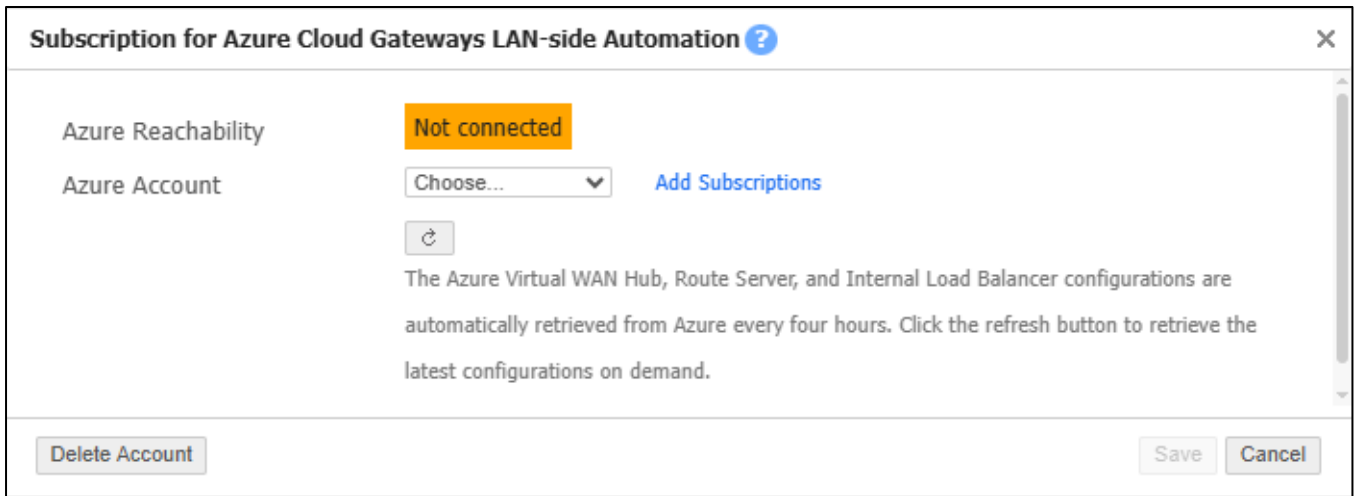


Figure 22. Selecting your Azure subscription in the Subscription for Azure Cloud Gateways LAN-side Automation dialog box of the Orchestrator UI.

4.2.2.1. Modify an existing custom role to add the permissions needed for ILB deployment

If you have added an Azure subscription to deploy EC-V gateways previously, update its custom role to include the extra permissions needed for ILB deployment.

1. Go to the resource group used for the EC-V deployment. This is the same resource group where the existing custom role is defined.
2. Navigate to **Access control (IAM) > Roles**.
3. In the search bar, enter the name of the custom role previously created for Orchestrator.
4. After the custom role is selected, click the three dots on the far-right of the row, and then select **Edit**.

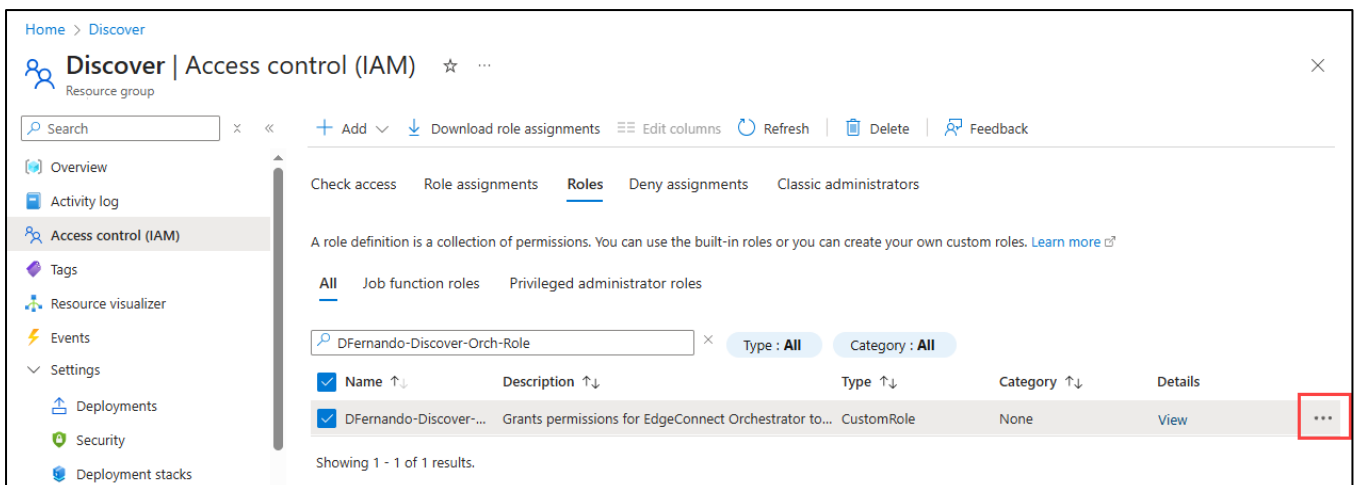


Figure 23. Modifying an existing custom role to add the permissions needed for ILB deployment.

5. On the Update a custom role page, click the **JSON** tab.
6. Click **Edit**.
7. Scroll to the bottom of the Actions section of the JSON editor. After the last permission entry, insert the ILB-related permissions as documented on the [Permissions required to deploy an Azure Standard Internal Load Balancer and establish connectivity to it from EC-Vs deployed in a Transit VNet](#) section of the HPE Aruba Networking EdgeConnect SD-WAN Documentation site. Note that the final permission line does not have a comma. Therefore, be sure to add a comma before adding the new permissions. If you already performed the instructions in this section, skip to Section [4.2.2.2](#).

- After the ILB-related permissions are added, on the Subscription for Azure Cloud Gateways LAN-side Automation dialog box, select the Azure account you added, and then click **Save**. The Azure Reachability label should be displayed as *Connected*.

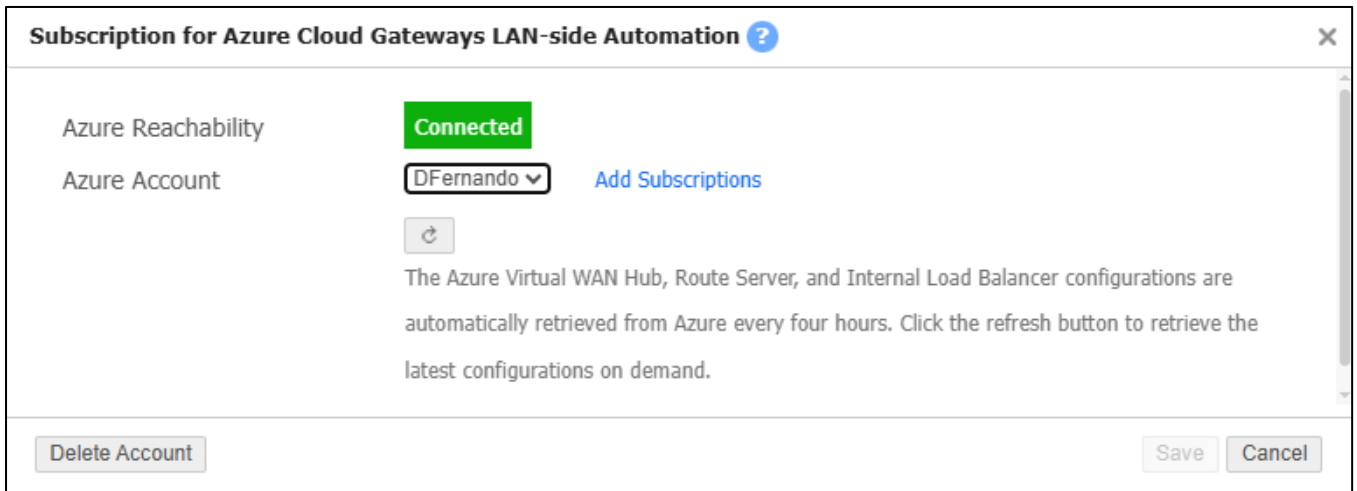


Figure 24. Selecting the Azure account (subscription) required for LAN-side automation.

4.2.2.2. Add a new custom role and the permissions needed for ILB deployment

If EdgeConnect was not deployed from Orchestrator—for example, if it was deployed through the Azure Portal or a tool such as Terraform—but you need to deploy and configure the ILB through Orchestrator, create a new custom role with only the permissions required for ILB deployment.

Creating a new app registration to add the Azure subscription details on Orchestrator is described in Sections [3.1.2.2](#), [3.1.2.3](#), [3.1.2.4](#), [3.1.2.5](#), and [3.1.3](#). The ILB-related permissions you must enter can be found on the [Permissions required to deploy an Azure Standard Internal Load Balancer and establish connectivity to it from EC-Vs deployed in a Transit VNet](#) section of the HPE Aruba Networking EdgeConnect SD-WAN Documentation site.

After the ILB-related permissions are added, on the Subscription for Azure Cloud Gateways LAN-side Automation dialog box, select the Azure account you added, and then click **Save**. The Azure Reachability label should be displayed as *Connected*.

4.2.3. Create ILB

You can deploy a new ILB or associate an existing one with EC-V gateways from Orchestrator. The following steps show how to create a new ILB from Orchestrator.

- From the Azure Network Manager tab, click **Deploy Load Balancer**.
The Azure Load Balancer Deployment Configuration dialog box appears.

Figure 25. Creating an Azure ILB from the Orchestrator UI.

2. Fill in the required details:
 - a. **Name:** Enter a name for the ILB.
 - b. **Azure account:** Select your Azure subscription.
 - c. **Resource group:** Select the resource group added to the Orchestrator.
 - d. **Region:** Select the Azure region that corresponds to the deployment location of EC-V gateways.
 - e. **Virtual network:** The VNet where EC-Vs are deployed. When deploying through Orchestrator, the ILB and EC-Vs must be in the same VNet. Deploying the ILB in a different VNet from the EC-Vs is not supported.
 - f. **New subnet:** If you want Orchestrator to create a new subnet for the ILB, select the check box. When you select the checkbox, two things happen:
 - I. The Orchestrator displays all existing subnets within the VNet you selected.
The Load balancer subnet CIDR field appears.
 - II. You must enter a subnet for the load balancer on this field. The subnet you enter must not overlap with any existing subnets in the virtual network and be /27 or bigger (such as /26 or /25). After the Orchestrator creates the ILB, it will use the fourth IP address of this subnet as the frontend IP address of the load balancer. For instance, if you enter 10.0.0.0/27 as the subnet mask, 10.0.0.4 is assigned as the frontend IP address of the load balancer.
 - g. If the New subnet check box is not selected, you must choose an existing subnet in your VNet for the ILB. This option is available for users whose organizations do not allow a third-party application such as the Orchestrator to create new subnets in a VNet. Ensure that the subnet you select is /27 or bigger (such as /26 or /25) in size. The Orchestrator will not accept a /28 or smaller subnet. This option allows you to create a subnet manually but still use the Orchestrator to deploy the ILB within that subnet. When this option is selected, you must also enter a frontend IP address for the ILB, and the IP address must be available within the selected subnet.
 - h. **TCP health probe port:** Enter a port number that you want to use for the ILB health probe. Port 443 is recommended.
 - i. **Comment:** (*Optional*) Add a comment to identify your ILB deployment.

3. Click **Deploy**.

Orchestrator now deploys the ILB and creates a backend pool, a health probe (on TCP port 443 by default), and a load-balancing rule. The backend pool is empty because you have not associated any EC-V gateways to the ILB yet. In Section 4.2.6, after you associate EC-V gateways to the ILB, the Orchestrator adds the LAN interface IPs of the selected EC-Vs to the ILB's backend address pool.

4.2.4. Configure LAN interface labels

To automate connectivity between the Azure EC-V gateway and the ILB, you must first create a LAN interface label in Orchestrator, assign it to a LAN interface of the Azure EC-Vs that need to connect to the ILB, and drag and drop it into the Primary section of the Establish connectivity using these network interfaces dialog box. These steps are detailed below:

1. In Orchestrator, navigate to **Configuration > Overlay & Security > Interfaces Labels**.
2. Click **New Label**.
3. Select **lan**, and then enter a label name.
4. Click **Save**.

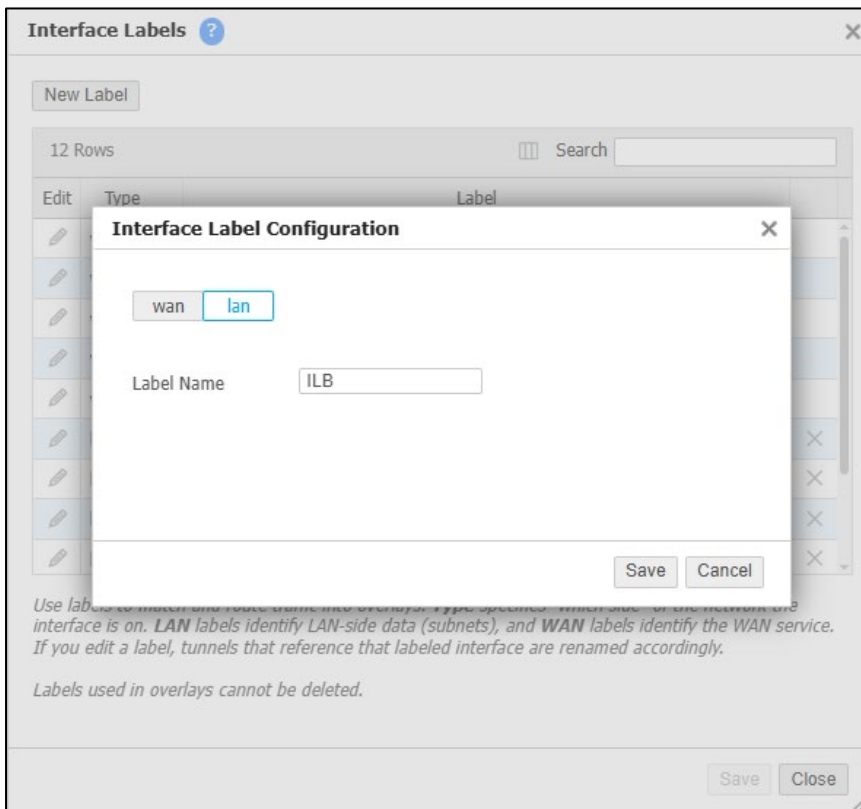


Figure 26. Creating a LAN interface label.

After the label is created, assign it to a LAN interface for each Azure EC-V that needs to receive traffic from the ILB.

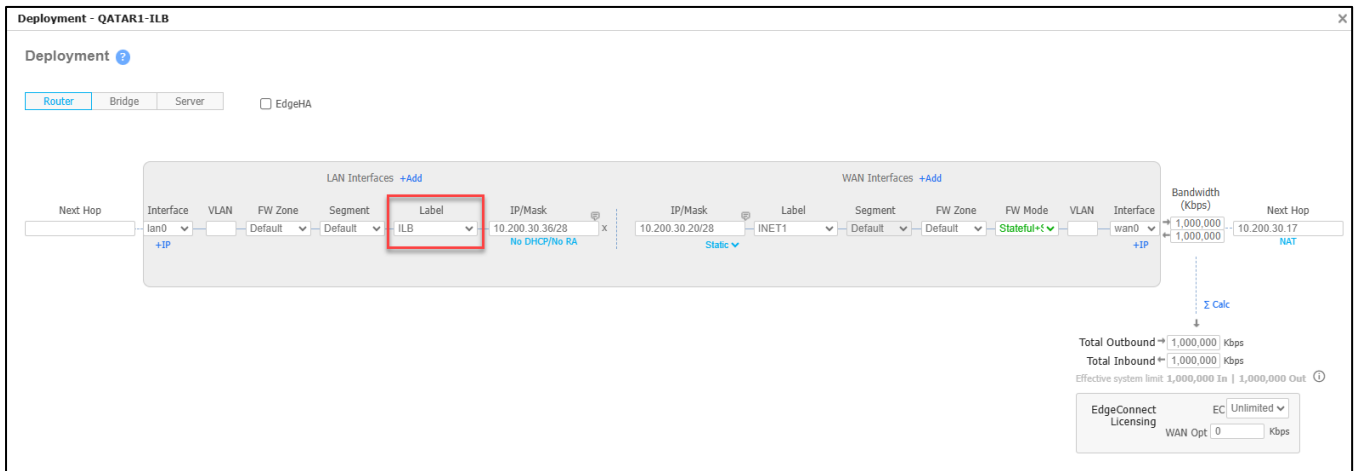


Figure 27. Assigning the label to a LAN interface of an Azure EC-V.

1. Navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
2. Click **Subscription**.
3. Select the Azure account (subscription) you want to use.
4. Click **Save**.
5. Click **LAN Interface Labels**.
6. Drag and drop the LAN-side label into the Primary section of the Establish connectivity using these network interfaces dialog box.

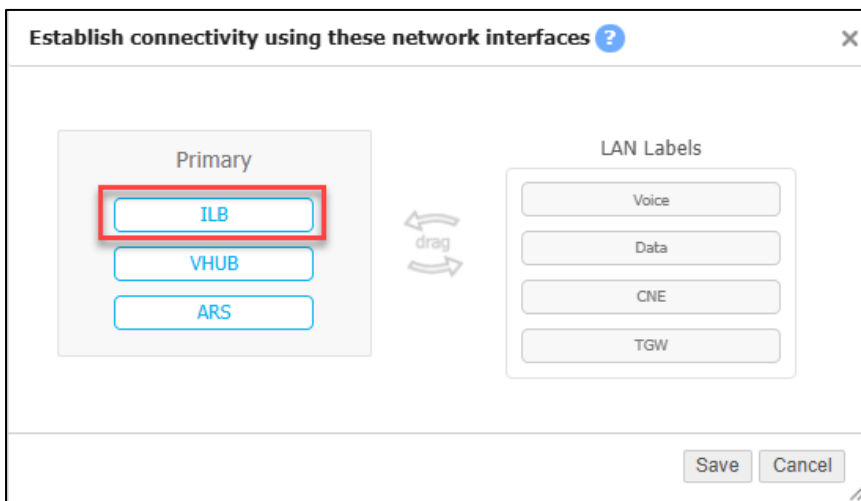


Figure 28. Dragging and dropping the LAN-side ILB label into the Primary section of the Establish connectivity using these network interfaces dialog box.

7. Click **Save**.

4.2.5. Configure Azure resources in Orchestrator

This step identifies the resource group, region, and VNet of the Azure EC-V gateways that participate in the automation.

1. On the Azure Network Manager tab, click **Azure Resources**.
2. Click **Azure Gateways**.

The Configure SD-WAN VNet Resources for cloud-deployed EdgeConnect appliances dialog box appears.

3. Click **Add**.

The Configure Azure Resources dialog box appears.

- a. **Rule Name:** Enter a descriptive name to help identify the EC-V gateways and their VNet.

- b. **Appliances:** Select the EC-V gateways that should be associated with the ILB. You can either enter the appliance names or click **Use Tree Selection** to browse and select them.
 - c. **Resource Group:** From the drop-down menu, select the Azure resource group where the EC-V gateways are deployed.
 - d. **Region:** Select the Azure region that corresponds to the deployment location of the EC-V gateways.
 - e. **Virtual Network:** Choose the VNet where the EC-V gateways reside. This ensures that the appliances are correctly associated with the ILB within the specified VNet.
4. To apply the configuration, click **Save**.

Figure 29. Selecting the resource group, region, and VNet of the Azure EC-V gateways that participate in the ILB automation.

4.2.6. Associate EC-V gateways with ILB

This step prompts you to select the EC-Vs and the ILB that need to establish connectivity.

1. In the appliance tree in Orchestrator, select each EC-V that you want to establish connectivity with the ILB.
2. On the Azure Network Manager tab, click **Appliance Association**.
3. Click **Azure Standard Internal Load Balancer**. You will see all ILBs the Orchestrator has access to. If you do not see the ILB that you want to associate your EC-Vs with, check the permissions assigned to the Orchestrator's custom role on Azure Portal.
4. From the lefthand panel, choose the ILB you want to associate.
5. Click **Save**.

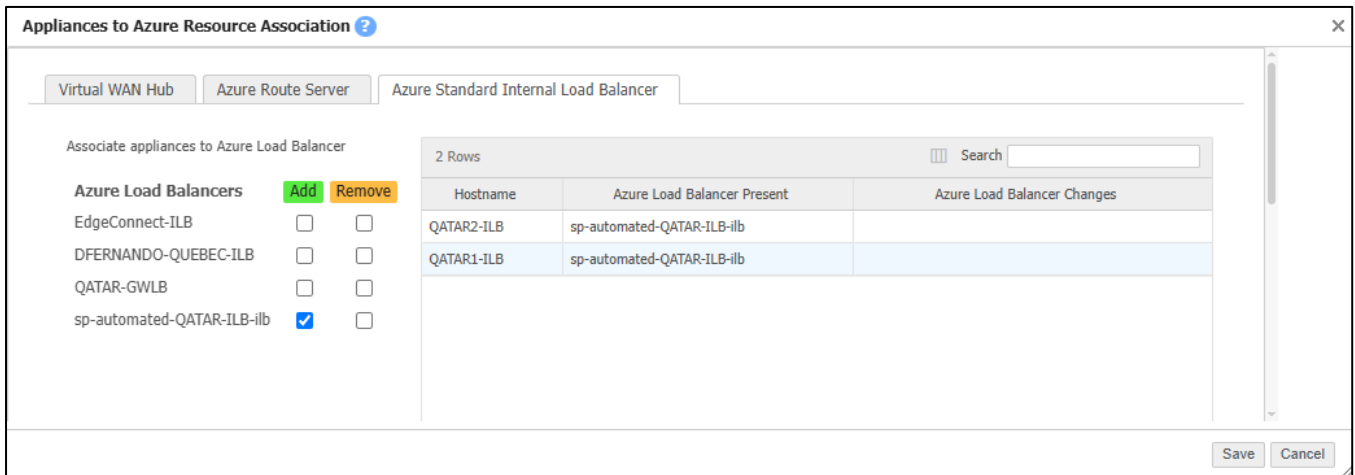


Figure 30. Associating EC-Vs to the ILB.

4.2.7. Verify connectivity

After the deployment or association is complete, you should verify the status.

1. **Check audit logs:** In Orchestrator, navigate to **Orchestrator > Orchestrator Server > Tools > Audit Logs**, and filter for **AzureLanConfigurationManager**. This shows progress and any issues with ILB integration.
2. **Connection status:** In Orchestrator, navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**. The connection status section displays the details of the association and its status, confirming that the EC-Vs are part of the ILB's backend pool.
3. **Check in Azure Portal:** Log in to Azure Portal to verify that the EC-V network interfaces have been added to the backend pool of the specified ILB.

This process ensures that your EC-V HA pair is correctly integrated with Azure ILB, providing a highly available solution for your network traffic.

4.2.8. Prevent the static route from being advertised to the SD-WAN fabric

When Orchestrator automated connectivity between EC-V and ILB, it created a static route on each EC-V to respond to the health probes from the ILB. You do not want the EC-V to advertise this route to remote EdgeConnect devices via Subnet Sharing. The following steps prevent this from happening.

1. On the Routes tab (Configuration > Networking > Routing > Routes), click the **edit** (pencil) icon located to the right of the Redistribute routes to SD-WAN Fabric drop-down menu.
The SD-WAN Fabric Route Distribution Maps dialog box appears.
2. To open the Add Rule dialog box, click **Add rule**, and then enter the following values:
 - a. **Select Match Criteria – Source Protocol:** Local/Static
 - b. **Select Match Criteria – Prefix:** Select and enter 168.63.129.16/32.
 - c. **Set Actions – Permit:** Clear
3. Click **Add**.
4. Click **Apply**.
5. Click **Apply**.
6. Starting from the Routes tab again, repeat the steps in this section to prevent the static route from being advertised to the SD-WAN fabric from EC-V-B.

4.2.9. Verify health probe responses using the Flows page

Perform the following tasks to verify if each EC-V is receiving health probes from the ILB:

1. Navigate to **Monitoring > Bandwidth > Flows > Active and Recent Flows**.
2. Click the **refresh** icon until the Outbound Tunnel column of a new flow indicates a passthrough tunnel with a LAN0 next hop. This takes a few seconds.

This confirms that the EC-V responds to health probes on the LAN interface instead of forwarding probes through its WAN interface, allowing the load balancer to forward traffic to the EC-V when probes indicate the EC-V is healthy.

4.2.10. (Optional) Verify health probe responses using tcpdump

You can further verify health probe responses by running a quick packet capture. Complete the following steps to run a packet capture on the Web CLI.

1. In Orchestrator, right-click an EC-V in the appliance tree, and then select **CLI Session**.

A CLI session opens on a new tab.

2. Enter the following commands to capture the traffic received on the LAN0 interface:

```
enable
```

```
tcpdump -i lan0 port 443 -nn
```

4.2.11. Create VNet peering between spoke (workloads) VNet and transit VNet

Now that you have automated connectivity between the EC-V and ILB, you need to establish a VNet peering connection between the spoke VNets where your workloads reside and the transit VNet. VNet peering allows workloads in spoke VNets to send and receive traffic from the EC-V gateways in the transit VNet. Because this is a common topic for both ILB and ARS designs, it is explained in Section 8. To create the VNet peering, follow the instructions in [8.1. Create a virtual network peering session \(for ILB and ARS designs\)](#).

4.2.12. Create an IP SLA rule on EC-V to monitor internet connectivity

It is important to create an IP SLA rule on EC-Vs to ensure that ILB health probes fail if the EC-V loses internet connectivity. This prevents outbound Azure traffic from being sent to an EC-V that cannot reach the internet, avoiding blackholed traffic.

4.2.12.1. Configure an IP SLA for external reachability

- Create an IP SLA rule to send ICMP echo requests (ping) to 8.8.8.8 via the WAN interface (wan0).

- The IP SLA continuously monitors outbound internet connectivity.

4.2.12.2. Define IP SLA down action

- Set the SLA Down Action to disable the passthrough data interface (pass-through_<LAN_IF_Label>_lan0) if the configured SLA target (8.8.8.8) becomes unreachable.

- Disabling pass-through_<LAN_IF_Label>_lan0 blocks ILB health probe traffic destined for the EC-V, which in turn causes the ILB to mark the instance as unhealthy and remove it from the backend pool.

4.2.12.3. Validation

- Confirm that ILB health probe traffic is reaching lan0.

- Simulate a WAN outage or block 8.8.8.8 to verify that the passthrough interface goes down and the ECV is removed from the backend pool.

- Restore connectivity and confirm automatic re-inclusion of the ECV in the pool.

4.3. Integrate EC-V with ILB manually from Azure Portal

Follow the instructions in this section to manually integrate EC-V gateways with ILB from Azure Portal.

4.3.1. Prerequisites

Before you begin, ensure the following prerequisites are met:

- Deploy EC-Vs in a transit VNet as illustrated in Section [3.1](#) or [3.2](#). Ensure that the transit VNet has enough space to create at least a /27 subnet to accommodate the ILB resource.

- **Azure account configuration:** You must have a valid Azure subscription configured in Orchestrator with the necessary permissions to deploy an ILB.

- **EC-V deployment:** The EC-V appliances must be deployed in the same VNet where the ILB is created.

4.3.2. Create ILB

In the following section, you will create an Azure ILB and configure it by entering details such as frontend IP configuration, backend pools, and load-balancing rules.

4.3.2.1. Create an Azure Standard ILB

1. From the Azure Portal, navigate to your resource group, and then click **+Add**.
2. Use the search field to select **Load Balancer** from the drop-down menu, and then click **Create**.
3. On the Load Balancer page, verify that **Basics** is selected on the menu bar, and then enter the following settings:
 - a. Project Details
 - I. **Subscription:** Depends on your Azure subscription.
 - II. **Resource Group:** Select your resource group.
 - b. Instance Details
 - I. **Name:** Enter the label for the new Load Balancer.
 - II. **Region:** Select the region where the EC-V is deployed.
 - III. **SKU:** Select **Standard (Distribute traffic to backend resources)**.
 - IV. **Type:** Select **Internal**.
 - V. Tier: Select **Regional**.

Note:

This selection is locked when Type is set to *Internal*.

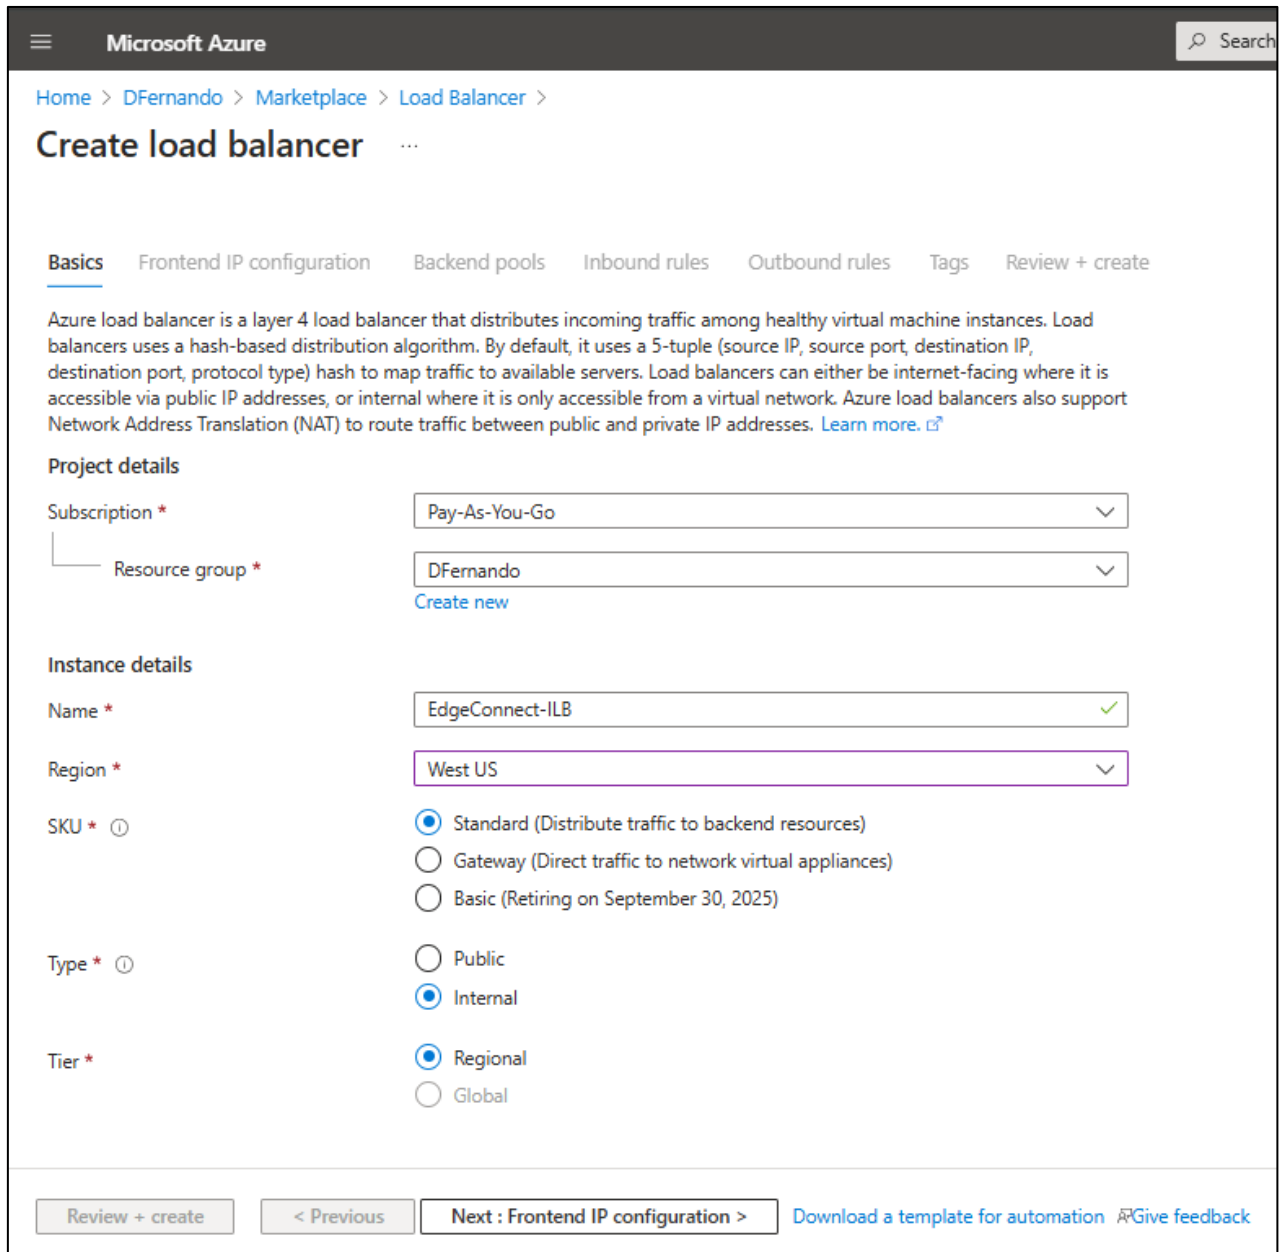


Figure 31. Configuring basic Internal Load Balancer (ILB) configuration settings.

4. Click **Next: Frontend IP configuration >**.
5. Click **Add a frontend IP configuration**, and then enter the following settings:
 - a. **Name:** Enter a name for the Frontend IP (example: *EdgeConnect-ILB-Frontend*).
 - b. **IP version:** IPv4
 - c. **Virtual network:** Select the virtual network where the EC-Vs are deployed.
 - d. **Subnet:** Select the subnet you created for the load balancer. If you have not created a subnet for the load balancer, create one.
 - e. **Assignment:** Select **Dynamic**.
 - f. **Availability zone:** Select **Zone-redundant**.

Note:

This parameter is not available in regions that do not support availability zones.

6. Click **Save**.
7. Click **Next : Backend pools >**.

8. On the Backend pools tab, click **+ Add a backend pool**, and then enter the following settings:
 - a. **Name:** Enter a backend pool name (example: *EC-V_Backend_Pool*).
 - b. **IP Backend Pool Configuration:** Select **NIC**.
 - c. Under IP configurations, click **+ Add**.
 - d. On the Add IP configurations to backend pool page, select the EC-Vs you want to include in the backend pool. These EC-Vs will receive health probes from the Azure ILB.
 - e. Click **Add**.
 - f. Click **Save**.
The two EC-Vs appear on the Backend Pool page.
9. Click **Next: Inbound rules >**.
10. Click **+ Add a load balancing rule**, and then enter the following settings:
 - a. **Name:** Enter a name for the load-balancing rule.
 - b. **IP Version:** Select **IPv4**.
 - c. **Frontend IP address:** Select the load balancer's frontend private IP address.
 - d. **Backend pool:** Select the backend pool that comprises the EC-Vs.
 - e. **High availability ports:** Select. This enables load balancing on all ports for TCP and UDP protocols.
 - f. **Health probe:** Click **Create new**.
 - g. **Name:** Enter a name for the health probe.
 - h. **Protocol:** Select **TCP**.

ⓘ Important

DO NOT use HTTPS.

- i. **Port:** Enter **443**.
- j. **Interval:** Set to **5** seconds. This is the shortest supported interval.
- k. **Session persistence:** Select a session persistence policy. Persistence policies specify traffic handling processes from a client by the same virtual machine in the backend pool for the duration of a session. Options include:
 - I. **None:** Successive requests from the same client may be handled by any VM in the backend pool.
 - II. **Client IP:** Successive requests from the same client IP address are handled by the same backend pool VM.
 - III. **Client IP and protocol:** Successive requests from the same client IP address and protocol combination are handled by the same backend pool VM.
 - IV. **Guide example:** Select **Client IP and protocol**.
- l. **Idle timeouts (minutes):** This setting decides to keep a TCP or HTTP connection open without relying on clients to send keep-alive messages.
Guide example: Enter **4**.
- m. **TCP Reset:** Enabled
- n. **Floating IP:** Disabled
11. Click **Save**.
 - a. **Inbound NAT rule:** Not required
 - b. **Outbound rules:** Not required
 - c. **Tags:** Optional
12. Click **Next : Review + create**.

The Create load balancer page appears. After you see the *Validation passed* message, proceed to the next step.

13. At the bottom of the page, click **Create**.

The Deployment is Underway page appears.

14. When the page displays *Your Deployment is complete*, click **Go to Resource**.

The Load Balancer page for the new load balancer appears.

15. In the Settings menu on the left-side column, click **Frontend IP configuration**.

The Frontend IP Configuration page appears. This page lists the load balancer's frontend IP. When creating a static route later, use this private IP as the next hop for outbound traffic from Azure workloads.

4.3.2.2. Verify the IP address of the ILB

After creating the load balancing rule, you can view the incoming health probe traffic on the EC-V.

1. In Orchestrator, select the EC-Vs in the appliance tree, and then navigate to **Monitoring > Bandwidth > Flows > Active and Recent Flows**.
2. In the IP/Subnet textbox and the TCP port on the Port textbox, enter **168.63.129.16**. This filters traffic sent from the Azure ILB to the EC-V. Azure uses this virtual IP address to facilitate health probes from the Azure ILB to determine the health state of VMs (EC-Vs).
3. Select the **Include Built-in** check box.
4. Click **Apply**. Note the health probes that are received from the ILB.

This panel indicates the health probes are going passthrough on the EC-V's WAN0 interface, as shown in the Outbound Tunnel column in Figure 32.

Section 4.3.3 creates a static route on the EC-V that replies to the health probes from the ILB instead of forwarding them on its WAN interface. This allows the ILB to detect the health of the EC-V.

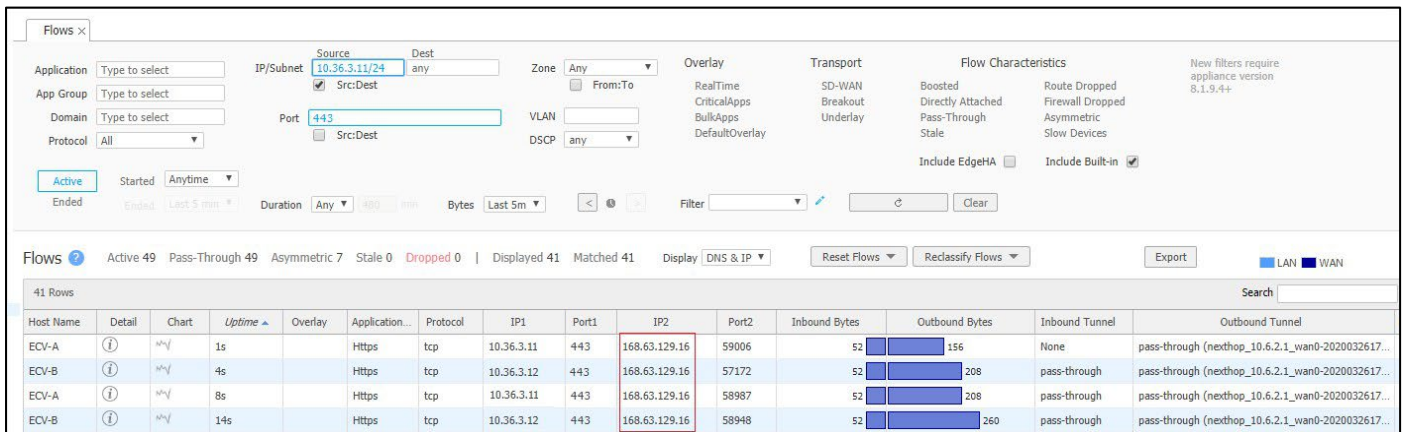


Figure 32. Verifying ILB health probe flows.

4.3.3. Create a static route on the EC-V to respond to health probe traffic

Complete the following steps to create a static route that responds to health probes from the ILB.

1. In Orchestrator, select **ECV-A** in the appliance tree, and then navigate to **Configuration > Networking > Routing > Routes**.
2. Click an **edit** (pencil) icon in the Routes table, and then enter the following settings:
 - a. **Use shared subnet information:** Select the check box.
 - b. **Automatically advertise local LAN subnets:** Clear the check box.
 - c. **Automatically advertise local WAN subnets:** Clear the check box.
 - d. **Metric for automatically added subnets:** Enter **50**.
 - e. **Redistribute routes to SD-WAN Fabric:** Select an available route map.
 - f. **Guide Example:** default_rtmap_to_subsh
 - g. **Include BGP Local ASN to routes sent to SD-WAN Fabric:** Clear the check box.
 - h. **Tag BGP communities to routes:** Clear the check box.

3. Click the **edit** (pencil) icon to the right of the Redistribute routes to SD-WAN Fabric data field.
The SD-WAN Fabric Route Distribution Maps dialog box appears.
4. Verify that the dialog box includes a rule with the following settings:
 - a. **Match Criteria:** Source Protocol Local/Static
 - b. **Permit:** Allow
This ensures that the static routes are advertised to the SD-WAN fabric.
5. If necessary, modify or add a rule, and then click **Apply**. If no modification is required, click **Cancel**.
6. Click **Add Route**, and then enter the following settings:
 - a. **Subnet/Mask:** Enter the load balancer IP address, as displayed on the Flows page.
 - b. **Next Hop:** Enter the LAN0 next hop address.
 - c. **Interface:** Enter **lan0**.
 - d. **Metric:** Enter **50**.
 - e. **Tag:** Select **ALL**.
 - f. **Comments:** Optional
7. Click **Add**.
8. From the Routes tab, create the same static route on EC-V-B.

4.3.4. Prevent the static route from being advertised to the SD-WAN fabric

In the previous step, you created a static route on each EC-V to respond to the health probes from the ILB. You do not want the EC-V to advertise this route to remote EdgeConnect devices. The following steps prevent this from happening:

1. On the Routes tab, click the **edit** (pencil) icon located to the right of the Redistribute routes to SD-WAN Fabric drop-down menu.
The SD-WAN Fabric Route Distribution Maps dialog box appears.
2. Click **Add rule**, and then enter the following values:
 - a. **Select Match Criteria – Source Protocol:** Local/Static
 - b. **Select Match Criteria – Prefix:** Select the check box, and then enter **168.63.129.16/32**.
 - c. **Set Actions – Permit:** Clear the check box.
3. Click **Add**, click **Apply**, and then click **Apply** again.
4. Repeat the steps in this section to prevent the static route from being advertised to the SD-WAN fabric from EC-V-B.

4.3.5. Verify health probe responses using the Flows page

1. Navigate to **Monitoring > Bandwidth > Flows > Active and Recent Flows**.
2. Click the **refresh** icon until the Outbound Tunnel column of a new flow indicates a passthrough tunnel with a LAN0 next hop. This takes a few seconds.

This confirms that the EC-V responds to health probes on the LAN interface instead of forwarding the probes through its WAN interface. The load balancer can now forward traffic to the EC-V when probes indicate that the EC-V is healthy.

4.3.6. (Optional) Verify health probe responses using tcpdump

You can further verify health probe responses by running a quick packet capture. Complete the following steps to run a packet capture on the Web CLI:

1. In Orchestrator, right-click an EC-V in the appliance tree, and then select **CLI Session**.
A CLI session opens on a new tab.

2. Enter the following commands to capture the traffic received on the LAN0 interface:

```
enable
tcpdump -i lan0 port 443 -nn
```

4.3.7. Create an IP SLA rule on the EC-V to monitor internet connectivity

It is important to create an IP SLA rule on EC-Vs to ensure that ILB health probes fail if the EC-V loses internet connectivity. This prevents outbound traffic from Azure being sent to an EC-V that cannot reach the internet, avoiding blackholed traffic.

4.3.7.1. Configure an IP SLA for external reachability

- Create an IP SLA rule to send ICMP echo requests (ping) to 8.8.8.8 via the WAN interface (wan0).
- The IP SLA continuously monitors outbound internet connectivity.

4.3.7.2. Define IP SLA down action

- Set the SLA Down Action to disable the passthrough data interface (pass-through_<LAN_IF_Label>_lan0) if the configured SLA target (8.8.8.8) becomes unreachable.
- Disabling pass-through_<LAN_IF_Label>_lan0 blocks ILB health probe traffic destined for the ECV, which in turn causes the ILB to mark the instance as unhealthy and remove it from the backend pool.

4.3.7.3. Validation

- Confirm that ILB health probe traffic is reaching lan0.
- Simulate a WAN outage or block 8.8.8.8 to verify that the passthrough interface goes down and the ECV is removed from the backend pool.
- Restore connectivity and confirm automatic re-inclusion of the ECV in the pool.

4.4. Create a cluster profile to enable flow redirection on the EC-V

As mentioned, EC-V can only be deployed in active-active failover mode with Azure ILB. Active-active failover mode results in flow asymmetry. Generally, flow asymmetry is not an issue for most applications. However, it can break traffic when firewalls are deployed on the LAN side of the EdgeConnect gateways, as stateful firewalls may drop asymmetric flows due to mismatched session states. Additionally, if you need the WAN Optimization (Boost) capability enabled on your Azure EC-V deployed with the ILB, you must enable flow redirection on the Azure EC-Vs. This is because the TCP acceleration capability on the EdgeConnect SD-WAN gateways requires symmetric TCP traffic. Flow Redirection merges the traffic of an asymmetric flow into a single appliance, removing flow asymmetry. Figure 33 depicts the redirection of a traffic flow through ECV-A (AZ1).

Note

Flow redirection only redirects TCP traffic. It does not redirect any non-TCP traffic such as UDP and ICMP.

Starting from ECOS 9.4.1.0 and later, the following limitations apply when you enable flow redirection:

- When segmentation is enabled, flow redirection works only for intra-segment traffic. It does not work for intersegment traffic.
- Flow redirection does not work for IPv6 traffic.
- When Zone-Based Firewall is enabled, flow redirection only works if the participating interface is in the Default zone.

Enable flow redirection on EC-Vs using Cluster Profile. This lets you manage multiple EdgeConnect gateways as a cluster and orchestrate flow redirection within it. Applying the Cluster Profile to EdgeConnect gateways makes each gateway inherit those settings.

Note

Starting with release 9.5, *Site Names* are called *Site/Cluster Names*. The site/cluster name must match precisely for each appliance in the cluster.

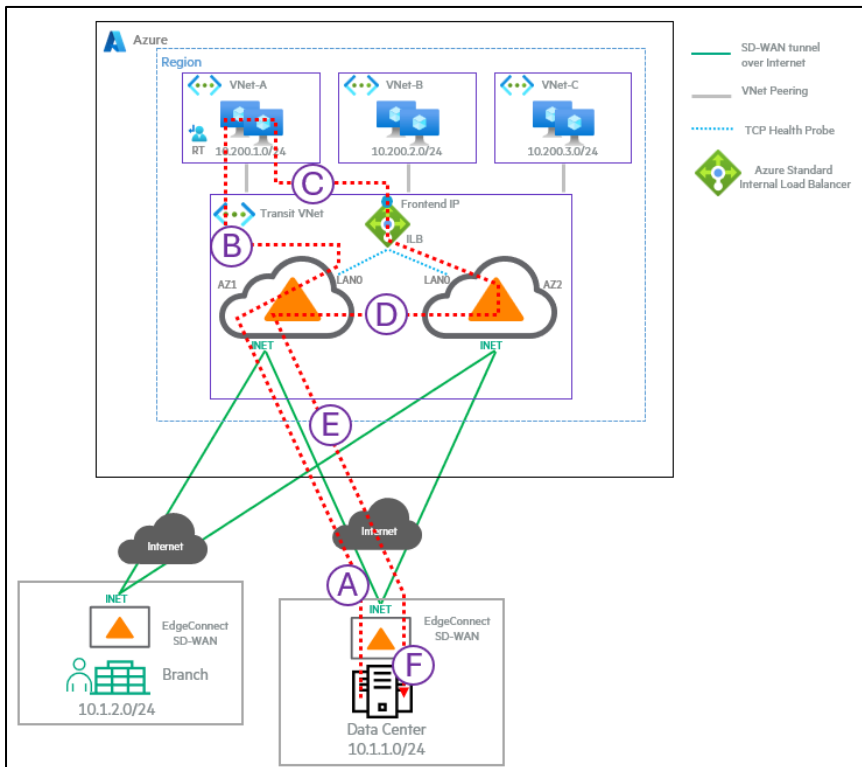


Figure 33. Viewing the traffic path when flow redirection is enabled.

To configure flow redirection on EC-V gateways, you must perform the following tasks, which are explained below:

- [10.1. Create a cluster profile](#)
- [10.2. Add EC-V gateways to the cluster](#)
- [10.3. Apply a cluster profile to EC-Vs](#)

4.5. Create VNet peering between spoke (workloads) VNet and transit VNet

Now that you have established connectivity between the EC-V and ILB, you need to establish a VNet peering connection between the spoke VNETs where your workloads reside and the transit VNET. VNet peering allows workloads in spoke VNETs to send and receive traffic from the EC-V gateways in the transit VNET. Because this is a common topic for both ILB and ARS designs, this is explained in [Section 8](#). To create the VNet peering, follow the instructions in [8.1. Create a virtual network peering session \(for ILB and ARS designs\)](#).

This concludes the EC-V HA deployment using the ILB.

5. Integrate EC-V in transit VNet with Azure Route Server (ARS)

This section provides detailed instructions for connecting Azure EC-Vs to an ARS.

5.1. EC-V with ARS architecture

ARS is a fully managed service that simplifies routing between a network virtual appliance such as an EC-V gateway and Azure virtual networks. It enables direct exchange of routing information through BGP between an EC-V and the Azure Software Defined Network (SDN). This eliminates the need to manually configure or maintain route tables and static routes, simplifying network management.

To learn more about ARS, go to <https://learn.microsoft.com/en-us/azure/route-server/overview>.

Integration between EC-V and ARS provides the following key benefits:

1. Dynamic route exchange between on-premises networks and Azure virtual networks via the Azure EC-Vs.
2. Elimination of manual route table updates when network changes occur, such as when you add a new spoke VNet, delete a VNet, add a new SD-WAN site, remove an existing SD-WAN site, and more.
3. Support for active-active or active-passive EC-V failover configurations.

5.1.1. ARS limitations

When planning your deployment, be aware of the following ARS limitations:

1. Maximum of 8 BGP peers from an EC-V (or EC-Vs) per route server deployment.
2. Each BGP peer from an EC-V can advertise up to 1,000 routes to an ARS.
3. Support for up to 4,000 VMs in the virtual network (including peered virtual networks).
4. The maximum number of peered VNets that a single ARS instance can dynamically exchange routes with is 500.
5. Support for up to 10,000 total on-premises and Azure VNet prefixes.

The same limits are mentioned at <https://learn.microsoft.com/en-us/azure/route-server/overview#route-server-limits>.

5.1.2. Topology of a single-region EC-V deployment with ARS

Figure 34 illustrates an SD-WAN deployment in Azure, demonstrating connectivity between on-premises SD-WAN sites to workloads in spoke VNets via EC-Vs in a transit VNet and ARS.

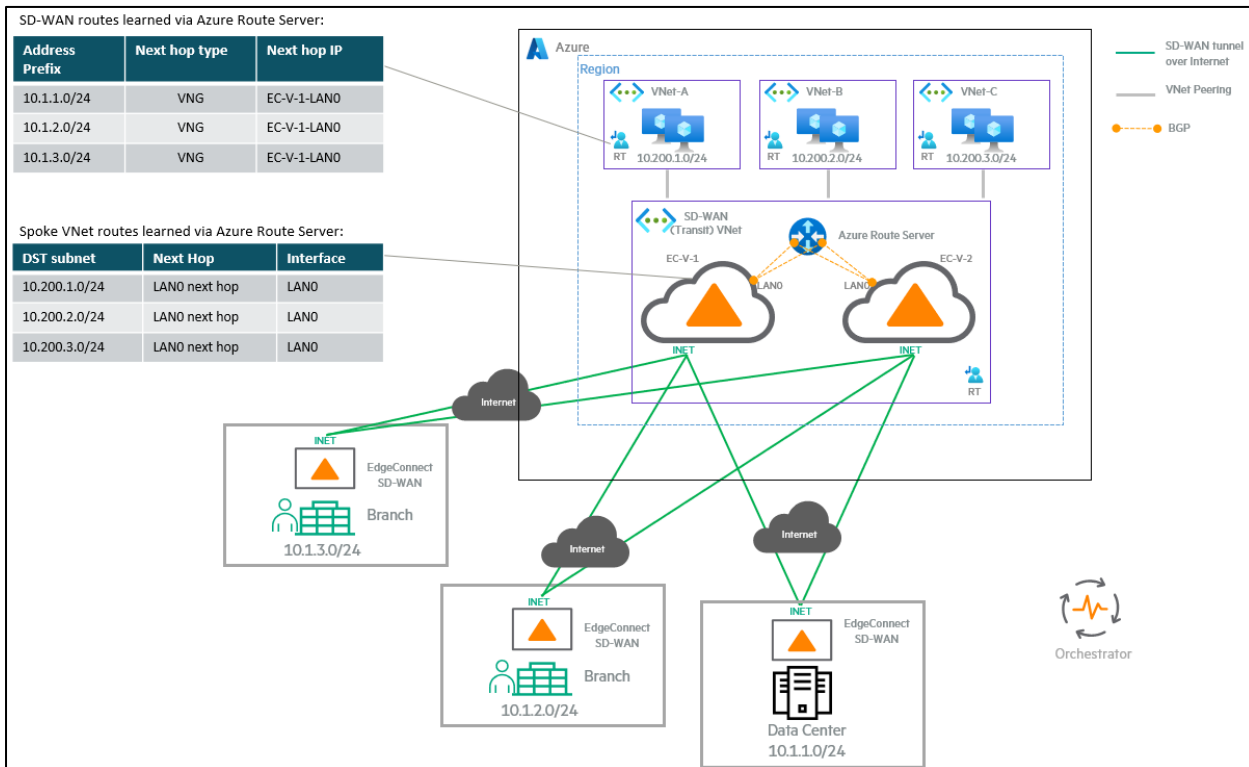


Figure 34. A single-region EC-V deployment with Azure Route Server.

Figure 34 also illustrates the following items:

1. Azure-based resources:
 - a. Three spoke VNets (VNet-A, VNet-B, VNet-C) with VNet CIDR blocks of 10.200.1.0/24, 10.200.2.0/24, and 10.200.3.0/24.
 - b. An SD-WAN transit VNet that hosts two EC-V gateways (EC-V-1 and EC-V-2). Each EC-V is deployed in a unique availability zone.
 - c. ARS is used to exchange routes dynamically between EC-V gateways and Azure VNets.
2. Route advertisements to spoke VNets:
 - a. Yellow lines represent BGP sessions between the EC-V gateways and ARS. The LAN interface of an EC-V gateway establishes two BGP sessions with the ARS. This dual-session approach prevents single points of failure of an ARS and maintains consistent connectivity.
 - b. Gray lines represent VNet peering connections between the SD-WAN Transit VNet and the spoke VNets.
 - c. The Azure EC-V gateways advertise on-prem networks (10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24) to spoke VNets via ARS. The on-prem EdgeConnect devices advertise these network prefixes to Azure EC-Vs via subnet sharing.
 - d. The next hop type for these on-prem network prefixes on the Azure workloads will appear as *Virtual Network Gateway* (VNG).
3. Route advertisements from spoke VNets:
 - a. The EC-V gateways learn spoke VNet CIDR blocks (10.200.1.0/24, 10.200.2.0/24, and 10.200.3.0/24) dynamically via ARS.
 - b. These routes are then readvertised to the on-prem SD-WAN devices via subnet sharing.
4. On-premises SD-WAN connectivity:
 - a. Branch sites and the data center connect to Azure EC-V gateways over the internet. In environments with an Azure ExpressRoute connection between your on-premises data center and Azure, it is possible to establish an SD-WAN tunnel over the ExpressRoute link by configuring an additional WAN interface on both your on-prem EdgeConnect device and the Azure EC-V gateways. By associating your ExpressRoute gateway with the SD-WAN transit VNet, the EC-V gateways can establish an SD-WAN tunnel to the on-prem data center's EdgeConnect device using private IP addresses.

- b. Green lines represent SD-WAN tunnels over the internet between on-prem SD-WAN devices and Azure EC-V gateways.
- 5. Key takeaways:
 - a. SD-WAN tunnels are used for secure communication between on-prem locations and Azure workloads.
 - b. ARS facilitates dynamic route learning between SD-WAN instances and Azure VNets.
 - c. The SD-WAN transit VNet acts as the central routing hub for Azure-hosted workloads and on-prem locations.
 - d. BGP is used for route exchange between the SD-WAN appliances and ARS.

This setup ensures optimized routing, security, and performance for hybrid cloud deployments.

5.1.3. Topology of a multi-region EC-V deployment with ARS

Figure 35 is similar to the previous figure, but introduces a multi-region SD-WAN architecture with Azure-based inter-region connectivity.

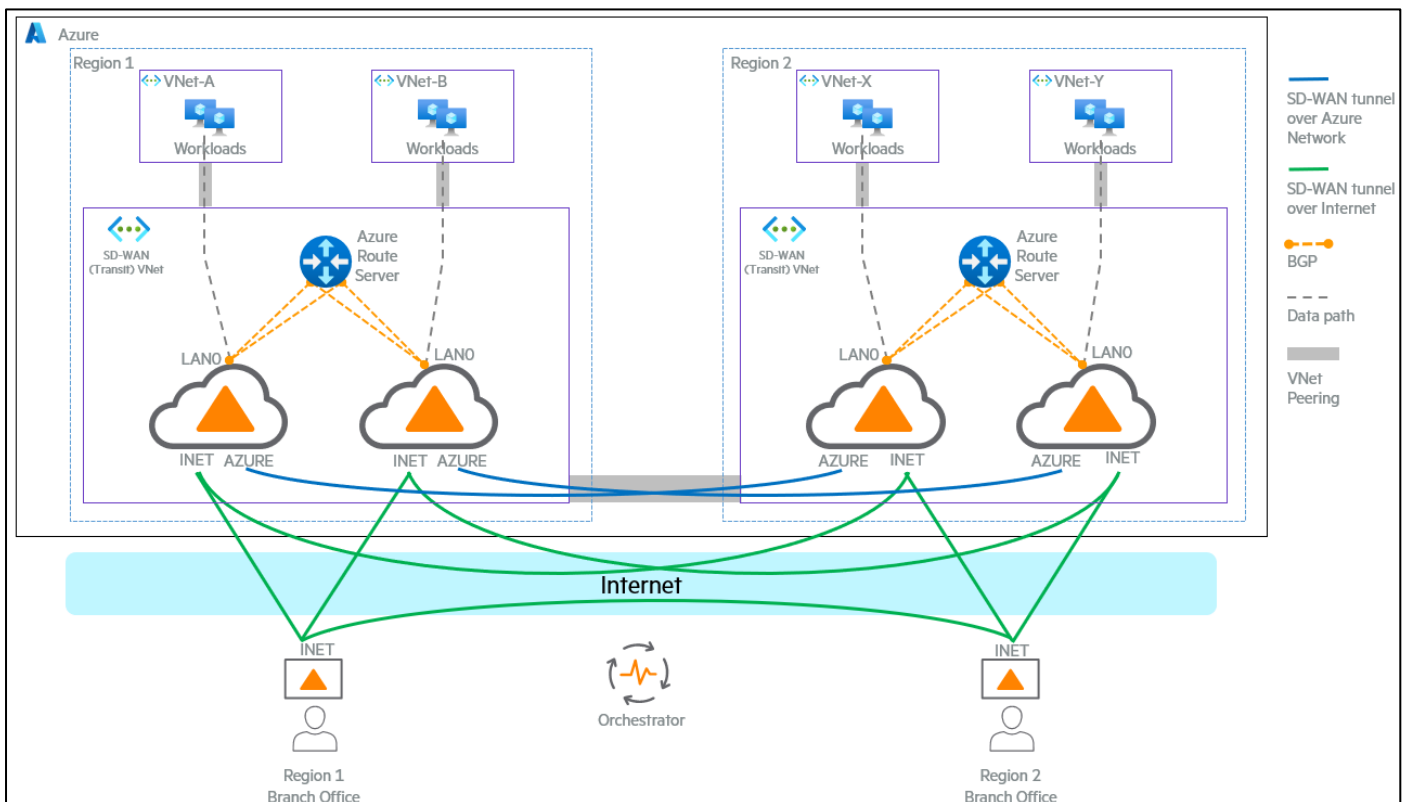


Figure 35. A multi-region EC-V deployment with ARS.

Figure 35 also illustrates the following items:

1. Multi-region deployment:
 - a. Instead of a single Azure region, this design spans two Azure regions.
 - b. Each region has its own SD-WAN transit VNet, EC-V instances, and ARS.
2. Azure network-based SD-WAN tunnels:
 - a. In addition to internet-based SD-WAN tunnels, this topology includes SD-WAN tunnels over the Azure backbone (shown in blue).
 - b. An inter-region VNet peering connection is established between the SD-WAN VNet in Region 1 and the SD-WAN VNet in Region 2.
 - c. These tunnels improve inter-region traffic performance by keeping SD-WAN traffic within Microsoft's private network instead of traversing the public internet.
3. Inter-region VNet connectivity:
 - a. Workloads in VNet-A/B (Region 1) and VNet-X/Y (Region 2) are interconnected through their respective EC-V appliances and ARS.

- b. This setup enables optimized routing across Azure regions while maintaining SD-WAN control over traffic flows.
4. Branch-to-region connectivity:
 - a. Branch offices in Region 1 and Region 2 are connected via SD-WAN tunnels to both Azure regions.
 - b. This provides regional redundancy, allowing traffic to be routed dynamically based on availability and performance.
5. Key takeaways:
 - a. Multi-region SD-WAN deployment for resilience and scalability.
 - b. SD-WAN tunnels over Azure's private network to optimize inter-region traffic.
 - c. Workload connectivity across Azure regions via SD-WAN transit VNets.
 - d. Redundant paths for branch offices to improve resiliency and performance.

This architecture is designed for organizations that require global SD-WAN connectivity with Azure-hosted workloads while optimizing for latency, availability, and cost efficiency.

5.1.4. Horizontal scaling of EC-Vs with ARS

ARS enables horizontal scaling for EC-V gateways by leveraging Equal Cost Multi-Path (ECMP) routing within the Azure SDN stack. When multiple EC-Vs are deployed and each establishes a BGP peering with ARS, the EC-Vs can advertise identical routes (for the same prefixes) with equally preferred BGP attributes. In this configuration, ARS injects multiple forwarding entries for these destinations into the virtual network's route table, each pointing to a different EC-V as the next hop. The Azure SDN then distributes flows across all EC-Vs with ECMP, providing both scalability (traffic is load balanced) and high availability (failover if one EC-V fails).

ECMP is automatically triggered by ARS when two or more EC-Vs are BGP-peered with ARS in the same VNet and each EC-V advertises the same routes with identical BGP attributes (AS Path, MED, Local Preference, and the like). When these conditions are met, ARS programs multiple routes for the destination prefix into the subnet's effective route table, with each EC-V as an equal-cost next hop.

If you prefer to configure an active-standby deployment instead of an active-active configuration, you must modify the AS Path Prepend values advertised by each EC-V to ARS. By assigning a longer AS Path Prepend on the standby EC-V, its routes become less preferred by the ARS, ensuring that traffic is primarily forwarded through the designated active EC-V. This approach maintains high availability while preventing load balancing across multiple EC-Vs.

5.1.5. Detection of EC-V failures

In an ARS deployment, the health of EC-V gateways is monitored using the BGP session's status between ARS and each EC-V. ARS continuously exchanges BGP keepalive messages (TCP port 179) with the EC-V LAN0 interfaces to maintain peering. If an EC-V gateway experiences a failure, the BGP session between ARS and that EC-V drops, alerting ARS to withdraw any routes previously learned from that EC-V. This dynamic withdrawal immediately removes the failed EC-V's prefixes from route distribution across all connected VNets, thereby preventing the blackholing of traffic through the inoperative EC-V. The best practice dictates that each EC-V establishes BGP sessions to both ARS IP endpoints. This ensures continued route propagation and high availability even when one BGP endpoint on the ARS becomes unreachable.

5.1.6. Decide whether to deploy ARS from Orchestrator or manually from Azure Portal

After following the instructions in Section 3 to deploy EC-V gateways in Azure, you are ready to deploy the ARS. Decide whether to deploy it and establish BGP using Orchestrator or manually from the Azure Portal and establish BGP by hand. Starting Orchestrator version 9.6.0, you can deploy an ARS within a transit VNet directly from Orchestrator and establish BGP with EC-Vs; this process is explained in [5.2. Integrate with ARS from Orchestrator](#). Alternatively, if you want to deploy ARS manually from Azure Portal and establish BGP with it, refer to [5.3. Integrate with ARS manually from Azure Portal](#). Manually deploying ARS and configuring BGP from the Azure Portal involves several steps, such as creating a subnet for ARS, deploying the ARS resource, adding EC-V gateways as BGP peers on the Azure Portal, creating two static routes on the EdgeConnect to enable BGP, and enabling BGP on the EdgeConnect. If you deploy the ARS and establish BGP from Orchestrator UI, all these tasks are automated.

Note

Creating the ARS and establishing BGP are independent tasks. You do not need to create ARS using Orchestrator to set up BGP using Orchestrator. In other words, you can create ARS using Azure Portal or another tool and then use Orchestrator to establish BGP from EC-V to the ARS.

5.2. Integrate with ARS from Orchestrator

This section provides steps for deploying a High Availability (HA) pair of EC-V gateways in Azure and integrating them with ARS using Orchestrator.

5.2.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- **Orchestrator version:** 9.6.0 or later
- **ECOS version:** 9.3.3 or later
- Deploy EC-Vs in a transit VNet, as illustrated in Section [3.1](#) or [3.2](#). Ensure that the transit VNet has enough space for a /26 (or bigger) subnet to accommodate the ARS subnet requirement. Microsoft requires a /26 or bigger subnet for the ARS subnet.
- **Orchestrator BGP ASN pool:** The BGP ASN pool must be enabled and the BGP ASN range must be specified in Orchestrator’s configuration. This is required for assigning an ASN to appliances.

Note

When several EC-V gateways share an identical site name, the automatic BGP ASN assignment feature allocates the same ASN to all EC-Vs associated with that site name.

- **Azure account configuration:** You must have a valid Azure subscription configured in Orchestrator with the necessary permissions to deploy an ARS.
- **EC-V deployment:** The EC-V appliances must be deployed in the same Virtual Network (VNet) where the ARS is created.

5.2.2. Add Azure subscription to Orchestrator

The following steps explain how to add an Azure subscription to Orchestrator:

1. In Orchestrator, navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
2. Click **Subscription**.

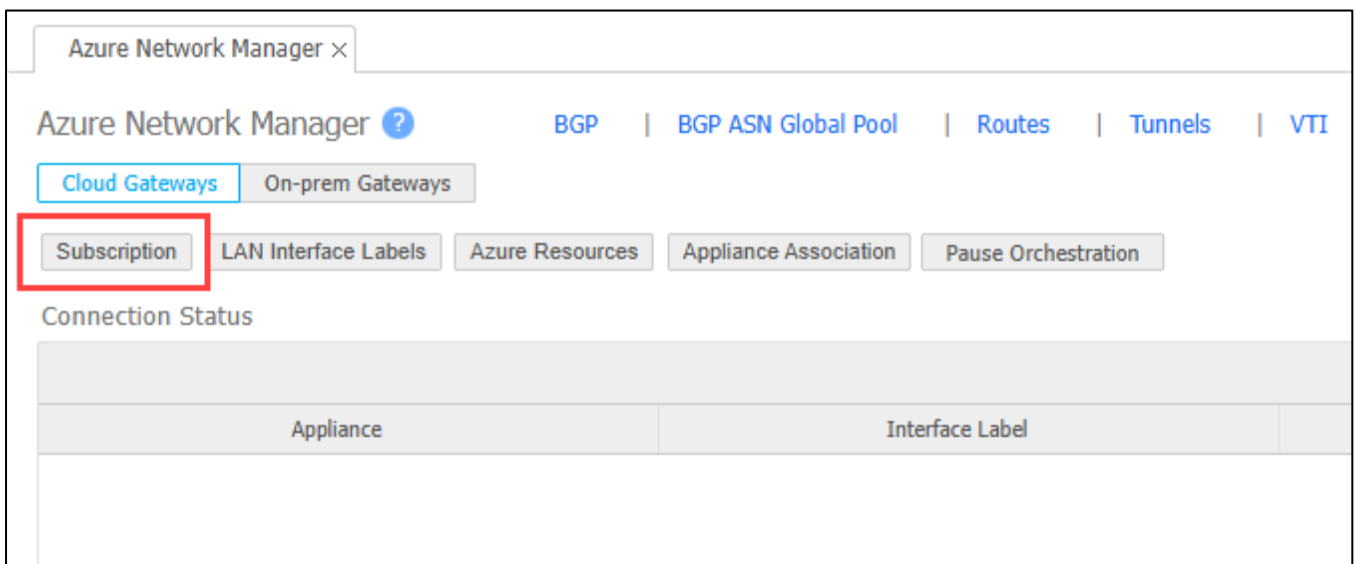


Figure 36. Clicking the Subscription tab in Azure Network Manager (2).

The Subscription for Azure Cloud Gateways LAN-side Automation dialog box appears.

3. If an Azure subscription has already been added to the Orchestrator (used for automated EC-V gateway deployment), it is possible to add the required permissions for the ARS deployment to its associated custom role. This is required because the EC-V gateway deployment does not include permissions needed to automate BGP connectivity to an ARS; this is described in [5.2.2.1. Modify an existing custom role to add the permissions needed for ARS deployment](#). If you do not have an existing Azure subscription in the Orchestrator, go to [5.2.2.2. Add a new custom role and the permissions needed for ARS deployment](#) and add the permissions needed for ARS deployment.

Note

If any modifications you make in Azure (such as creating or deleting subnets) are not reflected on the Orchestrator after you add the Azure subscription to Orchestrator, click the **refresh** icon, as shown in Figure 37. Clicking the refresh icon retrieves the latest Azure configuration on demand.

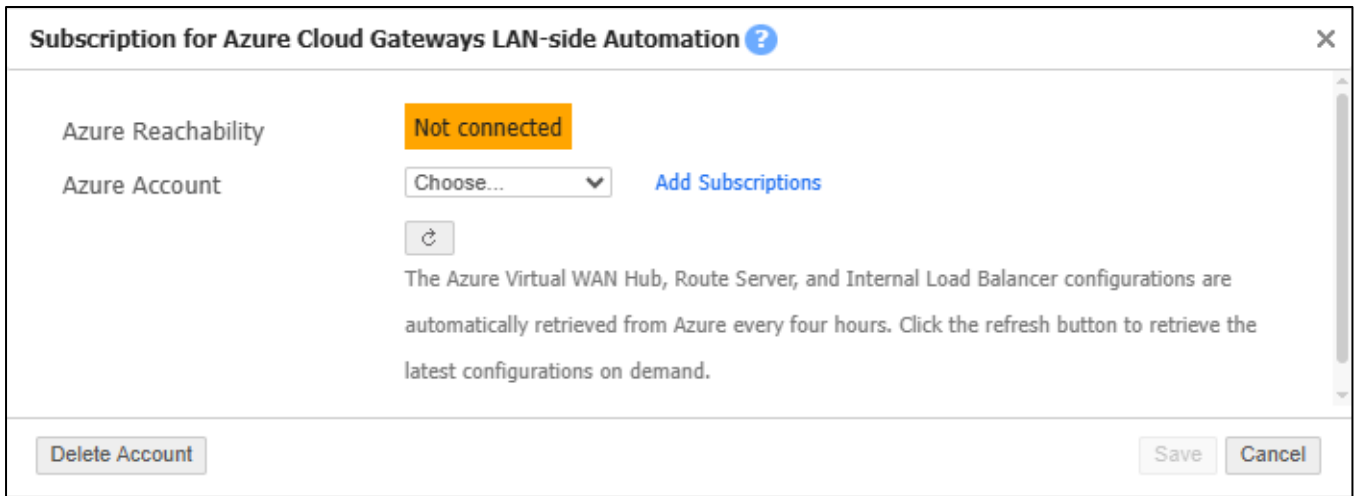


Figure 37. Selecting your Azure subscription in the Subscription for Azure Cloud Gateways LAN-side Automation dialog box of the Orchestrator UI (2).

5.2.2.1. Modify an existing custom role to add the permissions needed for ARS deployment

If you have added an Azure subscription to deploy EC-V gateways previously, update its custom role to include the extra permissions needed for ARS deployment:

1. Go to the resource group used for the EC-V deployment. This is the same resource group where the existing custom role is defined.
2. Navigate to **Access control (IAM) > Roles**.
3. In the search bar, enter the name of the custom role previously created for Orchestrator.
4. After the custom role is selected, click the three dots on the far right of the row, and select then **Edit**.

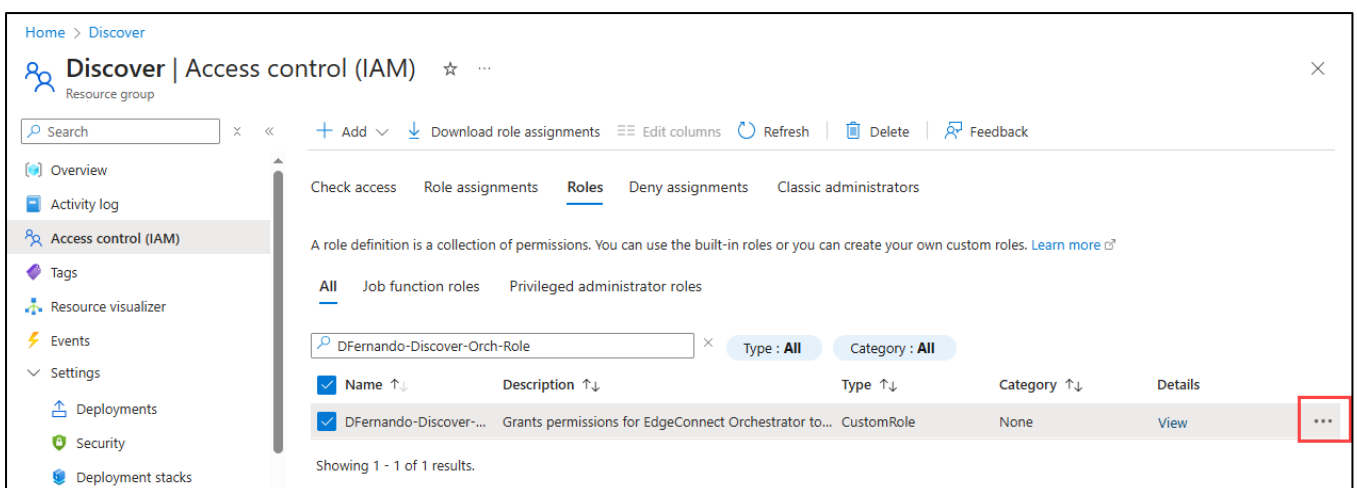


Figure 38. Modifying an existing custom role to add the permissions needed for ARS deployment.

5. On the Update a custom role page, click the **JSON** tab.
6. Click **Edit**.
7. Scroll to the bottom of the Actions section in the JSON editor and, after the last permission entry, insert the ARS-related permissions as documented on the [Permissions required to deploy an Azure Route Server and establish BGP connectivity from EC-Vs deployed in a Transit VNet](#) section of the HPE Aruba Networking EdgeConnect SD-WAN Documentation site. Note that the final permission line does not have a comma. Therefore, be sure to add a comma before adding the new permissions.

5.2.2.2. Add a new custom role and the permissions needed for ARS deployment

If EdgeConnect was not deployed from Orchestrator but you need to deploy and configure the ARS through Orchestrator, create a new custom role with only the permissions required for ARS deployment. Because the EC-V has already been deployed via the Azure Portal or Infrastructure-as-Code tools, only ARS-related permissions need to be assigned to this custom role.

Creating a new app registration to add the Azure subscription details on Orchestrator is described in Sections [3.1.2.2](#), [3.1.2.3](#), [3.1.2.4](#), [3.1.2.5](#), and [3.1.3](#). The ARS-related permissions you must enter can be found on the [Permissions required to deploy an Azure Route Server and establish BGP connectivity from EC-Vs deployed in a Transit VNet](#) section of the HPE Aruba Networking EdgeConnect SD-WAN Documentation site.

5.2.3. Create ARS

You can create a new ARS or associate an existing one with EC-V gateways from Orchestrator. The following steps explain how to create a new ARS from Orchestrator within a transit VNet:

1. In Orchestrator, navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
2. Under the Route Servers header, click **Deploy Route Server**.
3. On the Azure Route Server Deployment Configuration dialog box, enter the following settings:
 - a. **Name:** Enter a descriptive name for the ARS.
 - b. **Azure account:** Select the Azure subscription for deployment.
 - c. **Resource group:** Select the resource group where the ARS is created.
 - d. **Region:** Select the Azure region for the EC-V gateways.
 - e. **Virtual network:** Select the transit VNet where the EC-Vs are deployed.

Note

For Orchestrator to establish BGP between the EC-V and the ARS, they both must reside within the same VNet.

- f. **Route server subnet:** If your transit VNet already has a RouteServerSubnet of /26 or larger, it will be auto-selected. This is useful for users deploying the subnet themselves and using Orchestrator only for BGP setup. If not, you will be asked to provide a CIDR block so Orchestrator can create the subnet for ARS deployment.
 - g. **Comment (Optional):** Add comments for tracking purposes.
4. Click **Deploy**.

Note

Route Server creation can take up to 30 minutes. You can monitor progress in the audit logs in Orchestrator or the activity logs in the Azure Portal.

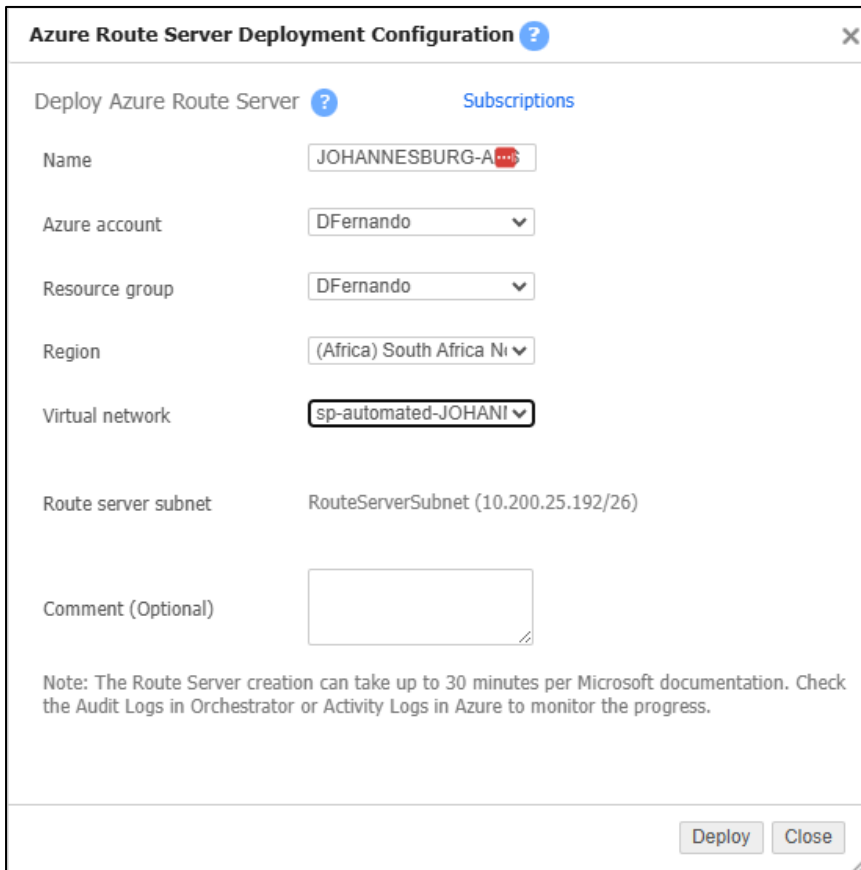


Figure 39. Creating an ARS from Orchestrator.

5.2.3.1. Post-deployment verification

- Confirm that the ARS resource appears in the Azure resource group.
- Check its status and assigned IP addresses.
- In Orchestrator, on the Azure Network Manager tab, verify that the ARS is listed by clicking **Appliances Association**, and then clicking **Azure Route Server**.

5.2.4. Configure LAN interface labels

Now that you have deployed the ARS, to automate connectivity between the Azure EC-V gateway and ARS, you must create a LAN interface label in Orchestrator, assign it to a LAN interface on the EC-Vs that need to connect to the ARS, and drag and drop it into the Primary section of the Establish connectivity using these network interfaces dialog box. These steps are detailed below:

1. In Orchestrator, navigate to **Configuration > Overlays & Security > Interface Labels**.
2. Click **New Label**.
3. Select **lan**, and then enter a label name.
4. Click **Save**.

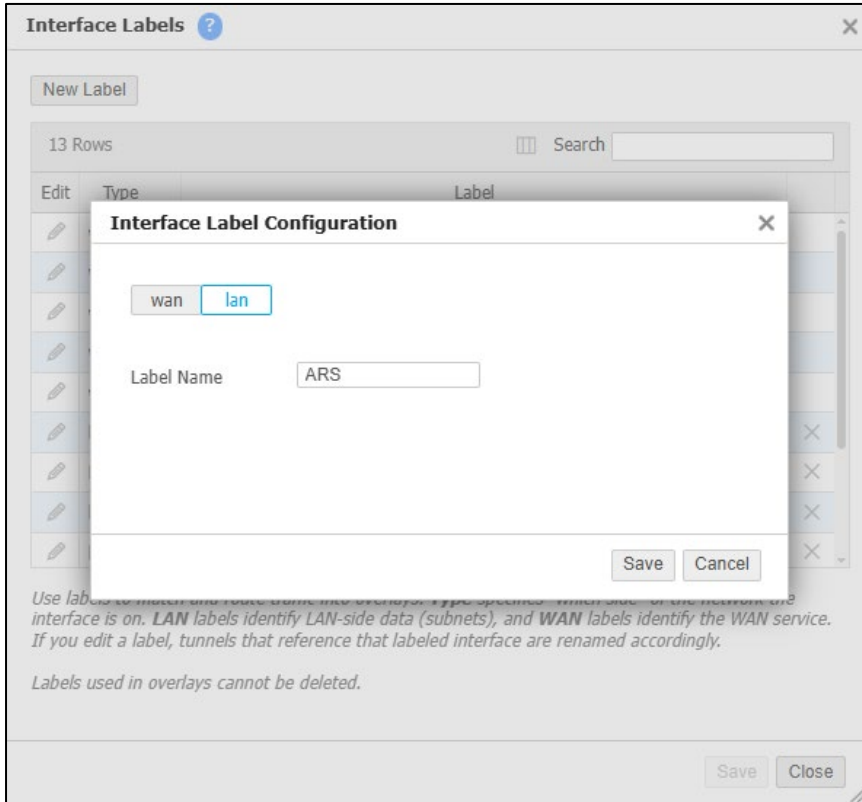


Figure 40. Creating a LAN interface label (2).

After the label is created, assign it to a LAN interface on each Azure EC-V that needs to establish BGP with ARS.

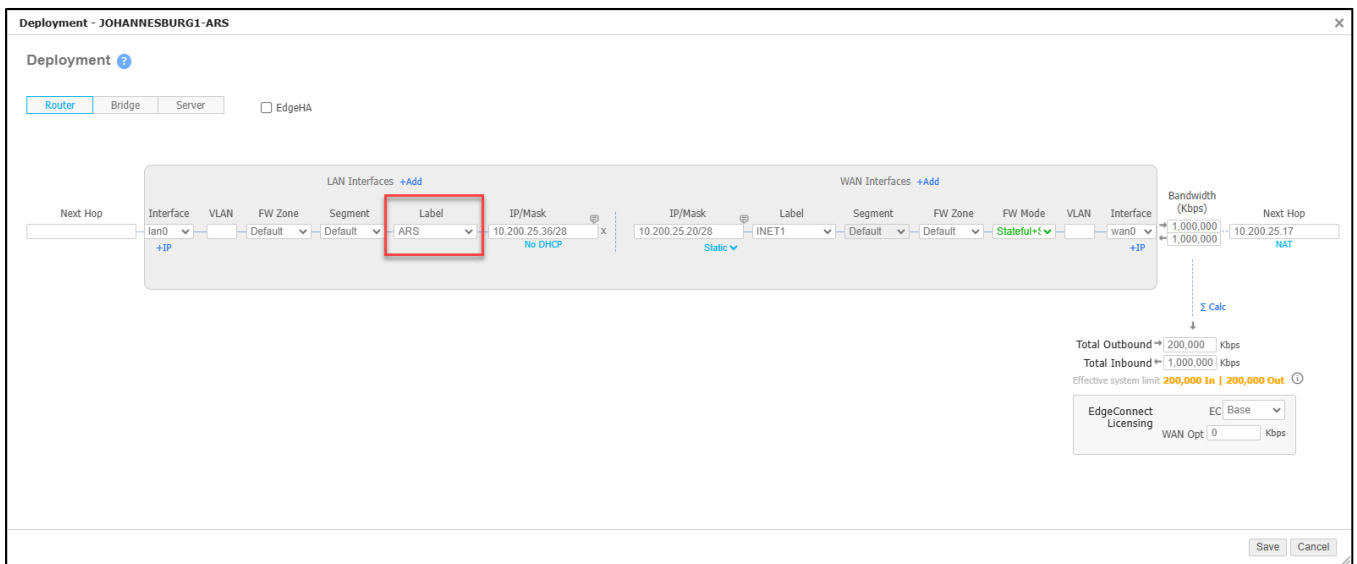


Figure 41. Assigning the label to a LAN interface of an Azure EC-V (2).

1. Navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
2. Click **Subscription**.
3. Select the Azure account (subscription) you want to use.
4. Click **Save**.
5. Click **LAN Interface Labels**.
6. On the Establish connectivity using these network interfaces dialog box, drag and drop the LAN-side label into the Primary section.

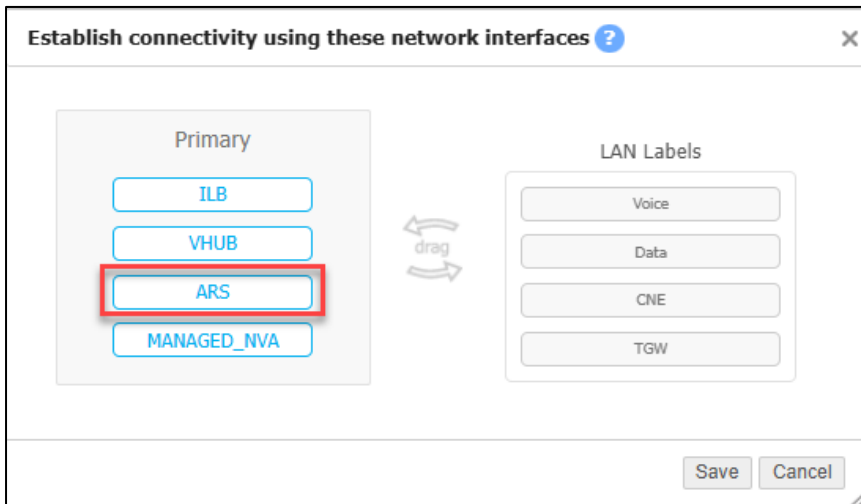


Figure 42. Dragging and dropping the LAN-side ARS label into the Primary section of the Establish connectivity using these network interfaces dialog box.

7. Click **Save**.

5.2.5. Configure Azure resources in Orchestrator

This step identifies the resource group, region, and VNet of the Azure EC-V gateways that participate in the automation:

1. On the Azure Network Manager tab, click **Azure Resources**.
2. Click **Azure Gateways**.
3. On the Configure SD-WAN VNet Resources for cloud-deployed EdgeConnect appliances dialog box, click **Add**.
4. On the Configure Azure Resources dialog box, enter the following settings:
 - a. **Rule Name:** Enter a descriptive name to identify the EC-V gateways and their VNet.
 - b. **Appliances:** Select the EC-V gateways that should be associated with the ILB. You can either enter the appliance names or click **Use Tree Selection** to browse and select them.
 - c. **Resource Group:** From the drop-down list, select the Azure resource group where the EC-V gateways are deployed.
 - d. **Region:** Select the Azure region that corresponds to the deployment location of the EC-V gateways.
 - e. **Virtual Network:** Choose the VNet where the EC-V gateways reside. This ensures that the appliances are correctly associated with the ILB within the specified VNet.
5. When all fields are completed, click **Save** to apply the configuration.

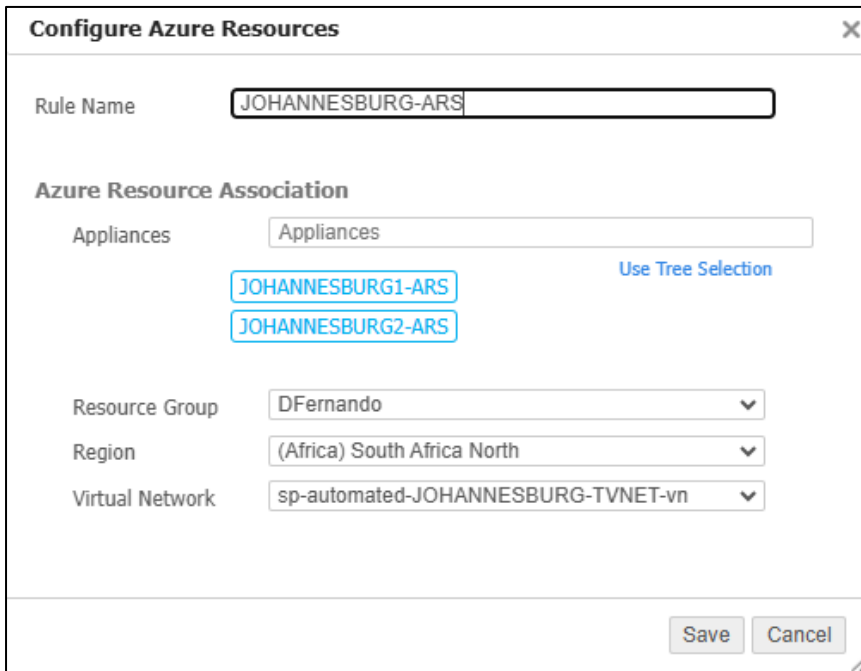


Figure 43. Selecting the resource group, region, and VNet of the Azure EC-V gateways that participate in the ARS automation.

5.2.6. Associate EC-V gateways with ARS

This step prompts you to select the EC-Vs and the ARS that need to establish BGP sessions.

Note

In Orchestrator version 9.6.0, if you select multiple EC-Vs on the Orchestrator appliance tree and associate them with an ARS, only one EC-V gateway's BGP sessions comes up. This is fixed in Orchestrator version 9.6.1 and later.

1. In the appliance tree in Orchestrator, select each EC-V that you want to establish BGP with the ARS.
2. On the Azure Network Manager tab, click **Appliance Association**.
3. Click **Azure Route Server**. You will see all Azure Route Servers to which the Orchestrator has access. If you do not see the ARS you want to associate your EC-Vs with, check the permissions assigned to the Orchestrator's custom role on Azure Portal.
4. From the left-side panel, select the ARS you want to associate.
5. Click **Save**.

5.2.7. Verify connectivity

After you click Save, the Orchestrator begins automating connectivity and establishing BGP with the ARS. To check the progress of the automation, perform the following steps:

1. In Orchestrator, navigate to **Orchestrator > Orchestrator Server > Tools > Audit Logs**.
2. Enter **AzureLanConfigurationManager** to filter the logs. This shows the progress of automation.
3. After a few minutes, navigate back to the Azure Network Manager tab in Orchestrator. The connection status section should display the details of the association and its status.
4. Log in to your Azure Portal and:
 - a. Check the activity logs.
 - b. Verify the status of the BGP sessions.
 - c. Confirm route exchange by checking EdgeConnect route tables and Azure effective routes.

| Appliance | Interface Label | Azure Service | Region | Connection Status |
|-------------------|-----------------|---|------------------|------------------------|
| JOHANNESBURG1-ARS | ARS LAN | sp-automated-JOHANNESBURG-ARS-route-server Route Server | southafricanorth | Up |
| JOHANNESBURG2-ARS | ARS LAN | sp-automated-JOHANNESBURG-ARS-route-server Route Server | southafricanorth | Up |
| PUNE1-VHUB | VHUB LAN | CentralIndia-Hub Vhub | centralindia | Up |
| PUNE2-VHUB | VHUB LAN | CentralIndia-Hub Vhub | centralindia | Up |
| QATAR1-ILB | ILB LAN | sp-automated-QATAR-ILB-ib Load Balancer | qatarcentral | Load Balancer Insights |
| QATAR2-ILB | ILB LAN | sp-automated-QATAR-ILB-ib Load Balancer | qatarcentral | Load Balancer Insights |

Figure 44. Successfully associating EC-Vs with ARS, as indicated by the *Up* connection status for the Johannesburg EC-Vs.

| Edit | Appliance | Segment | Peer IP | Address Family | Route Target | Local Interface | Peer ASN | Peer State | Soft Reset | Established Time | Peer Details |
|------|-------------------|---------|---------------|----------------|--------------|-----------------|----------|-------------|------------|------------------|--------------|
| | JOHANNESBURG1-ARS | Default | 10.200.25.196 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 5h 31m 2s | |
| | JOHANNESBURG1-ARS | Default | 10.200.25.197 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 5h 31m 2s | |
| | JOHANNESBURG2-ARS | Default | 10.200.25.196 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 5h 31m 2s | |
| | JOHANNESBURG2-ARS | Default | 10.200.25.197 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 5h 31m 2s | |

Figure 45. Verifying BGP sessions of both EC-Vs.

You have successfully associated your EC-Vs with the ARS and established BGP.

5.2.8. Create VNet peering between spoke (workloads) VNet and transit VNet

After establishing BGP sessions, you need to establish a VNet peering connection between the spoke VNets where your workloads reside and the transit VNet. The VNet peering allows workloads in spoke VNets to send and receive traffic from the EC-V gateways in the transit VNet. Since this is a common topic for both ILB and ARS designs, this is explained in Section 8. To create the VNet peering, follow the instructions in [8.1. Create a virtual network peering session \(for ILB and ARS designs\)](#).

After creating the VNet peering, perform the following steps to verify if the destination routes are being properly learned by the workloads in spoke VNets:

1. Navigate to a vNIC of one of the workloads in your spoke VNet that is peered with the transit VNet.
2. Under Help, select **Effective routes**.
3. Verify that routes advertised by your EC-V appliances are visible. Confirm that the route's Next Hop Type is *Virtual network gateway*.
4. On the Routes page of the Azure EC-Vs, ensure that the routes are being learned and advertised to and from ARS.

5.2.9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric

To avoid the two static routes that were created on each EC-V (to establish BGP sessions) from being advertised to other SD-WAN devices via subnet sharing, see [9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric](#).

5.3. Integrate with ARS manually from Azure Portal

This section explains how to manually integrate EC-V gateways with ARS from the Azure Portal.

5.3.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- Deploy EC-Vs in an SD-WAN transit VNet, as illustrated in Section [3.1](#) or [3.2](#). Ensure that the SD-WAN transit VNet has enough space for a /26 (or bigger) subnet to accommodate the ARS subnet requirement. Microsoft requires a /26 or bigger subnet for the ARS subnet
- EC-V and ARS must be in the same VNet. Make sure that all EC-Vs for a region are in one VNet, which is also used for ARS deployment, as outlined below.

5.3.2. Create a subnet for the ARS

After the EC-Vs are deployed, perform the following steps to create a subnet in the SD-WAN transit VNet for ARS:

1. From the Azure Portal, navigate to your resource group, and then click the name of the SD-WAN transit VNet.
2. From the left-side navigation pane, click **Subnets**.
3. Click **+ Subnet**.
4. On the Add subnet page, enter the following settings:
 - a. **Subnet purpose:** Select **Route Server**.
 - b. **Include an IPv4 address space:** Select the check box.
 - c. **IPv4 address:** Ensure that your VNet CIDR block is selected.
 - d. **Subnet address range:** Select the starting address of ARS's subnet.
 - e. **Size:** Select **/26 (64 addresses)** or bigger.
5. Leave all other settings as is, and then click **Add**.

You have successfully created a subnet for the ARS.

5.3.3. Deploy ARS

Now that you have created the RouteServerSubnet, perform the following tasks to deploy the ARS:

1. From the Azure Portal, navigate to your resource group.
2. Click **+ Create**.
3. Search for **route server**, and then select **Route Server**.
4. Enter the following settings:
 - a. **Subscription:** Select your subscription.
 - b. **Resource group:** Select your resource group.
 - c. **Name:** Enter a name for your Route Server (for example, *EC-V-RouteServer*).
 - d. **Region:** Select the same region as your SD-WAN transit VNet.
 - e. **Routing Preference:** Select **ASPath**.
 - f. **Virtual network:** Select your SD-WAN transit VNet.
 - g. **Subnet:** Because you already created a subnet for the Route Server, the **RouteServerSubnet** should be automatically selected.
 - h. **Public IP address:** Create a new standard public IP address.

Note

Azure's underlying SDN and management platform require this public IP address to communicate with ARS. Because ARS is part of the customer's private network, compliance restrictions prevent Azure's platform from directly accessing or managing it through private endpoints. Access to ARS's public endpoints is secured using certificate-based authentication, and Azure performs regular security audits on these endpoints. Consequently, they do not pose a security risk to your virtual network.

- i. **Public IP address name:** Enter a name for the ARS's public IP address.
 - j. **Assignment:** Static
5. Click **Review + create**.
6. After validation passes, click **Create**.
7. Wait for the deployment to complete. (This can take up to 30 minutes.)
8. After deployment is complete, navigate to your route server's Overview page and note the ASN (65515) and Route Server IP addresses. You need these when configuring BGP on your EC-V appliances.

You have successfully created the ARS.

5.3.4. Add EC-V gateways as BGP peers on the ARS

1. In the Azure Portal, select the route server you created.
2. From the left-side navigation pane, click **Peers**.

3. Click **+ Add**.
4. Enter the following information about the EC-V gateway:
 - a. **Name:** Enter a name for the peering between the ARS and the EC-V gateway.
 - b. **ASN:** Enter the ASN of the EC-V. Enter any BGP ASN between 64512 and 65534, excluding the following ASNs:
 - I. ASNs reserved by Azure:
 1. Public ASNs: 8074, 8075, 12076
 2. Private ASNs: 65515, 65517, 65518, 65519, 65520
 - II. ASNs reserved by IANA:
 1. 23456, 64496-64511, 65535-65551

Note

ARS only supports 16-bit ASNs. Do not enter a 32-bit ASN. The ASN you enter here is entered on the EC-V gateway.

- c. **IPv4 Address:** Enter the IP address of the EC-V gateway's LAN0 interface. This is the interface that the route server communicates with to establish BGP.
 5. Click **Add** to add the first EC-V gateway as a BGP peer.
 6. Repeat steps 2–5 to add the second EC-V as a BGP peer. You must enter the same ASN on both BGP peers.
- You have successfully added EC-V gateways as BGP peers on the ARS.

5.3.5. Create two static routes for enabling BGP

Because the ARS contains two endpoints, you must create two static routes on each EC-V gateway to establish BGP with the ARS. To create the static routes, perform the following steps:

1. From the Azure dashboard Overview page, find the Azure Route Server BGP endpoints. The two BGP endpoint IP addresses appear in the Essentials section as *Route Server IP addresses*.
2. In Orchestrator, select the first EC-V in the appliance tree.
3. Navigate to **Configuration > Networking > Routing > Routes**, and then select **Local/Static**.
4. Click the **edit** (pencil) icon on the Default segment.
5. On the Routes page for the Default segment, enter the following settings:
 - a. **Automatically advertise local LAN subnets:** Clear the check box.
 - b. **Automatically advertise local WAN subnets:** Clear the check box.
 - c. **Metric for automatically added subnets:** Enter **50**.
 - d. **Redistribute routes to SD-WAN Fabric:** Select an available route map.
 - e. **Filter routes from SD-WAN fabric with matching local ASN:** Select the check box.
 - f. **Include BGP Local ASN to routes sent to SD-WAN Fabric:** Select the check box.
 - g. **Tag BGP communities to routes:** Clear the check box.
 - h. Click **Apply**.
6. Click **Add Route**.
7. Enter the following settings:
 - a. **Subnet/Mask:** Enter the first ARS BGP endpoint (/32) IP address.
 - b. **Next Hop:** Enter the EC-V's LAN0 interface's next hop IP address. This is the first IP address of the LAN0 subnet. (If you are establishing BGP sessions from other interfaces, such as the LAN1 interface, make sure to enter the first IP address of that interface's subnet.)
 - c. **Interface:** LAN0
 - d. **Zone:** Default

- e. **Metric:** Leave the default value (**50**).
 - f. **Tag:** ANY
 - g. **Comments:** Optional
8. To add the first static route, click **OK**.
 9. Repeat steps 5–7 and create a static route for the second ARS BGP endpoint.
 10. Repeat steps 2–7 on the second EC-V, creating two static routes on that EC-V as well.

5.3.6. Add ARS as a BGP peer on the EC-V gateways

Now that the static routes are created and BGP peers are created on the ARS, perform the following steps to enable BGP on the EC-V gateways:

1. In Orchestrator, select the first EC-V in the appliance tree.
2. Navigate to **Configuration > Networking > Routing > BGP**.
3. Click the **edit** (pencil) icon on the Default segment row..
4. Click the **Enable BGP** toggle.
5. Enter the following BGP settings:
 - a. **Autonomous system number:** Enter the EC-V's ASN that you entered in [5.3.4. Add EC-V gateways as BGP peers on the ARS](#).
 - b. **Route Target:** Leave empty.
 - c. **Router ID:** The Router ID is an IPv4 address by which the remote peer can identify this EC-V gateway for purposes of BGP.
 - d. **Graceful restart:** Select the check box.
 - e. **AS path propagate:** Select the check box.
 - f. **Log BGP update messages:** Select the check box.
 - g. **Max route updates per peer:** 10
 - h. **Detection interval:** 15 minutes
 - i. Click **Add Peer**, and then enter the following settings:
 - I. **Peer IP:** Primary Azure BGP Peer (same address as the /32 Static Route added in the previous step)
 - II. **Peer Adjacency:** Multi-Hop
 - III. **EVPN Peer:** Clear the check box.
 - IV. **Local Interface:** lan0
 - V. **Peer ASN:** 65515
 - VI. **Override ASN:** Select the check box.
 - VII. **Peer Type:** Branch
 - VIII. **Admin Status:** Up
 - IX. **Soft Reconfiguration:** Select the check box.
 - X. **Next-Hop Self:** Select the check box.
 - XI. **Inbound route map:** Select the appropriate inbound route map.
 - XII. **Outbound route map:** Select the appropriate outbound route map.
 - XIII. **BFD:** Clear the check box.
 - XIV. **Keep Alive Timer:** 5
 - XV. **Hold Timer:** 15
 - XVI. **Enable MD5 Password:** Clear the check box.
 - XVII. To create the first BGP session, click **Add**.
6. Repeat steps 2–5 to create the second BGP session.

7. Repeat the same steps on the second EC-V.
8. Verify BGP configuration by checking the BGP neighbor status in the Orchestrator.

5.3.7. Create VNet peering between spoke (workloads) VNet and transit VNet

After establishing BGP sessions, you must establish a VNet peering connection between the spoke VNets where your workloads reside and the transit VNet. VNet peering allows workloads in spoke VNets to send and receive traffic from the EC-V gateways in the transit VNet. Because this is a common topic for both ILB and ARS designs, it is explained in Section 8. To create the VNet peering, follow the instructions in [8.1. Create a virtual network peering session \(for ILB and ARS designs\)](#).

5.3.8. Verify connectivity

After adding the VNet peering session, perform the following tasks to verify that the routes are being exchanged between the SD-WAN fabric and the Azure workloads:

1. In the Azure Portal, navigate to your route server.
2. Under Settings, select **Peers**.
3. Confirm that the Provisioning State of your EC-V peers is *Provisioned*.
4. Navigate to a vNIC of one of the workloads in your spoke VNet that is peered with the transit VNet.
5. Under Help, select **Effective routes**.
6. Verify that the routes advertised by your EC-V appliances are visible. Confirm that the route's Next Hop Type is *Virtual network gateway*.
7. On the Routes page of the Azure EC-Vs, ensure that the routes are being learned and advertised to and from ARS.

You have successfully established BGP between the EC-Vs and ARS.

5.3.9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric

To avoid the two static routes that were created on each EC-V (to establish BGP sessions) from being advertised to other SD-WAN devices via subnet sharing, see [9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric](#).

6. Transit VNet + Azure Virtual WAN (vWAN) hub design

This section provides detailed instructions for connecting Azure EC-Vs deployed in a transit VNet to a vWAN hub.

6.1. EC-V with Azure vWAN hub architecture

An Azure vWAN hub is a Microsoft-managed service that is part of the Azure vWAN service. It provides a unified way to connect, secure, and manage branch offices, remote users, and workloads across Azure regions. By leveraging a global transit network architecture, vWAN enables connectivity between on-premises environments, Azure VNets, and remote users.

6.1.1. Topology of a single-region EC-V deployment with vWAN hub

Figure 46 illustrates an SD-WAN deployment in Azure, demonstrating connectivity between on-premises SD-WAN sites to workloads in spoke VNets via EC-Vs in a transit VNet and vWAN hub.

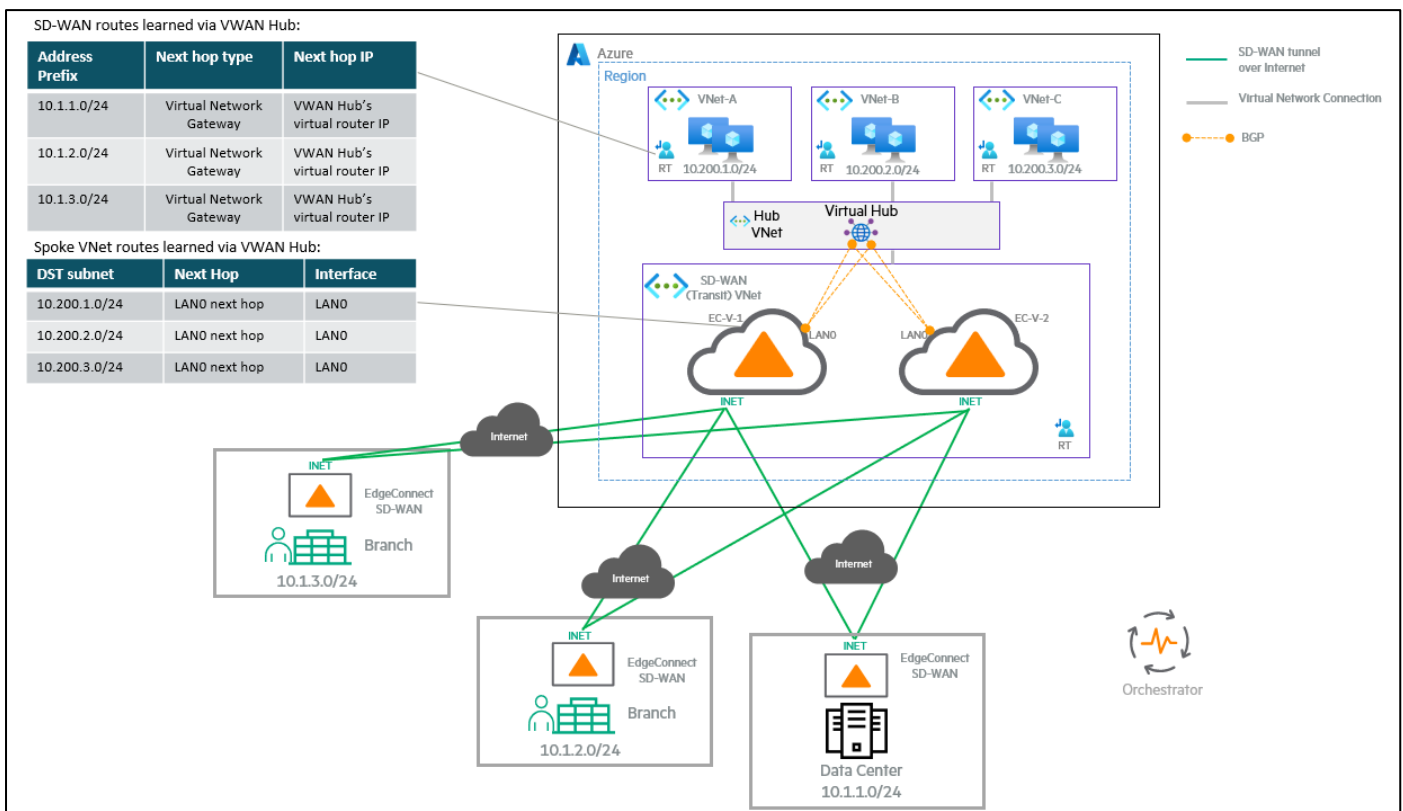


Figure 46. A single-region EC-V deployment with vWAN hub.

The above diagram illustrates the following items:

1. Azure-based resources:
 - a. Three spoke VNets (VNet-A, VNet-B, VNet-C) with VNet CIDR blocks of 10.200.1.0/24, 10.200.2.0/24, and 10.200.3.0/24.
 - b. An SD-WAN transit VNet that hosts two EC-V gateways (EC-V-1 and EC-V-2). Each EC-V is deployed in a unique availability zone.
 - c. Azure vWAN hub is used to exchange routes dynamically between EC-V gateways and Azure VNets.
2. Route advertisements to spoke VNets:
 - a. Yellow lines represent BGP sessions between the EC-V gateways and vWAN hub. The LAN interface of an EC-V gateway establishes two BGP sessions with the vWAN hub. This dual-session approach prevents single points of failure of a vWAN hub and maintains consistent connectivity.
 - b. Gray lines represent SD virtual network connections between the VNets and the vWAN hub.

- c. The Azure EC-V gateways advertise on-prem networks (10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24) to spoke VNets via vWAN hub. The on-prem EdgeConnect devices advertise these network prefixes to Azure EC-Vs via subnet sharing.
 - d. The next hop type for these on-prem network prefixes on the Azure workloads will show as *Virtual Network Gateway* (VNG).
3. Route advertisements from spoke VNets:
 - a. The EC-V gateways learn spoke VNet CIDR blocks (10.200.1.0/24, 10.200.2.0/24, and 10.200.3.0/24) dynamically via vWAN hub.
 - b. These routes are then readvertised to the on-prem SD-WAN devices via subnet sharing.
 4. On-premises SD-WAN connectivity:
 - a. Branch sites and the data center connect to Azure EC-V gateways over the internet. In environments with an Azure ExpressRoute connection between your on-premises data center and Azure, it is possible to establish an SD-WAN tunnel over the ExpressRoute link by configuring an additional WAN interface on both your on-premises EdgeConnect device and the Azure EC-V gateways. By associating your ExpressRoute gateway with the SD-WAN transit VNet (or vWAN hub), the EC-V gateways can establish an SD-WAN tunnel to the on-premises data center's EdgeConnect device using private IP addresses.
 - b. Green lines represent SD-WAN tunnels over the internet between on-prem SD-WAN devices and Azure EC-V gateways.
 5. Key takeaways:
 - a. SD-WAN tunnels are used for secure communication between on-prem locations and Azure workloads.
 - b. vWAN hub facilitates dynamic route learning between SD-WAN instances and Azure VNets.
 - c. The SD-WAN transit VNet acts as the central routing hub for Azure-hosted workloads and on-prem locations.
 - d. BGP is used for route exchange between the SD-WAN appliances and vWAN hub.

This setup ensures optimized routing, security, and performance for hybrid cloud deployments.

6.1.2. Topology of a multi-region EC-V deployment with vWAN hub

Figure 47 is similar to Figure 46, but it introduces a multi-region SD-WAN architecture with Azure-based inter-region connectivity.

Note

While Azure EC-V gateways support establishing SD-WAN tunnels between regions within the Azure network, the diagram focuses solely on tunnels established over the Azure internet gateway.

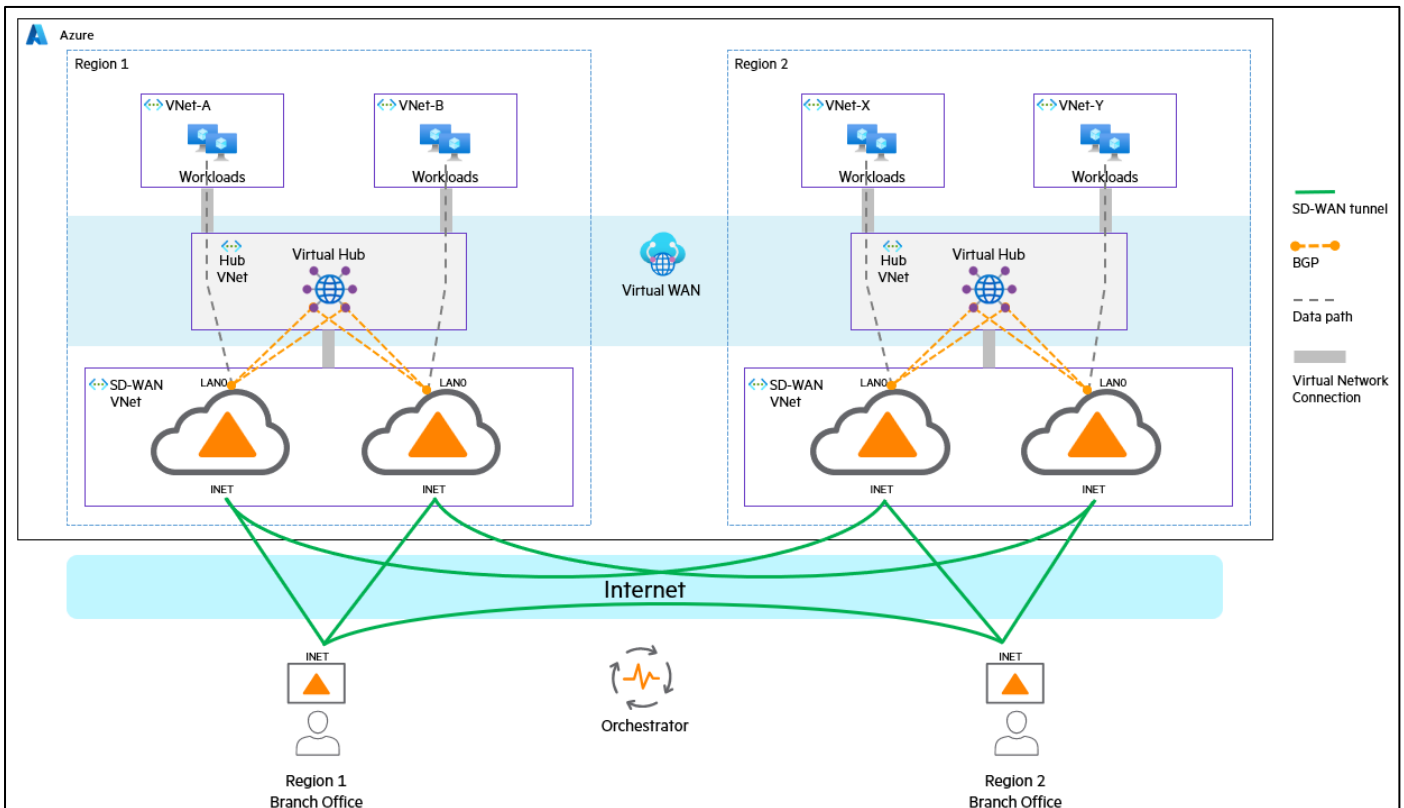


Figure 47. A multi-region EC-V deployment with vWAN hub.

The above diagram illustrates the following items:

1. Multi-region deployment:
 - a. Instead of a single Azure region, this design spans two Azure regions.
 - b. Each region has its own SD-WAN transit VNet, EC-V instances, and vWAN hub.
2. Azure network-based SD-WAN tunnels:
 - a. As the two vWAN hubs are created within the same vWAN resource, an automatic, inter-region virtual network connection is established between Region 1 and Region 2.
3. Inter-region VNet connectivity:
 - a. Workloads in VNet-A/B (Region 1) and VNet-X/Y (Region 2) can communicate with each other over the Azure network.
4. Branch-to-region connectivity:
 - a. Branch offices in Region 1 and Region 2 connect via SD-WAN tunnels to both Azure regions.
 - b. This provides regional redundancy, allowing traffic to be routed dynamically based on availability and performance.
5. Key takeaways:
 - a. Multi-region SD-WAN deployment for resilience and scalability.
 - b. Workload connectivity across Azure regions via SD-WAN transit VNets.
 - c. Redundant paths for branch offices to improve resiliency and performance.

This architecture is designed for organizations that require global SD-WAN connectivity with Azure-hosted workloads while optimizing for latency and availability.

6.1.3. Horizontal scaling of EC-Vs with vWAN hub

vWAN hub enables horizontal scaling for EC-V gateways by leveraging Equal Cost Multi-Path (ECMP) routing within the Azure SDN stack. When multiple EC-Vs are deployed and each establishes a BGP peering with vWAN hub, the EC-Vs can advertise identical routes (for the same prefixes) with equally preferred BGP attributes. In this configuration, vWAN hub injects multiple forwarding entries for these destinations into the VNet's route table, each

pointing to a different EC-V as the next hop. The Azure SDN then distributes flows across all EC-Vs with ECMP, providing both scalability (traffic is load balanced) and high availability (failover if one EC-V fails).

ECMP is automatically triggered by vWAN hub when two or more EC-Vs are BGP-peered with vWAN hub and each EC-V advertises the same routes with identical BGP attributes (AS Path, MED, Local Preference, and the like). When these conditions are met, vWAN hub programs multiple routes for the destination prefix into the subnet's effective route table, with each EC-V as an equal-cost next hop.

If you prefer to configure an active-standby deployment instead of an active-active configuration, you must modify the AS Path Prepend values advertised by each EC-V to vWAN hub. By assigning a longer AS Path Prepend on the standby EC-V, its routes become less preferred by the vWAN hub, ensuring that traffic is primarily forwarded through the designated active EC-V. This approach maintains high availability while preventing load balancing across multiple EC-Vs.

6.1.4. Detection of EC-V failures

In a vWAN hub deployment, the health of EC-V gateways is monitored using the BGP session's status between vWAN hub and each EC-V. vWAN hub continuously exchanges BGP keepalive messages (TCP port 179) with the EC-V LAN0 interfaces to maintain peering. If an EC-V gateway experiences a failure, the BGP session between vWAN hub and that EC-V drops, alerting vWAN hub to withdraw any routes previously learned from that EC-V. This dynamic withdrawal immediately removes the failed EC-V's prefixes from route distribution across all connected VNets, thereby preventing the blackholing of traffic through the inoperative EC-V. The best practice dictates that each EC-V establishes BGP sessions to both vWAN hub private IP endpoints. This ensures continued route propagation and high availability even when one BGP endpoint on the vWAN hub becomes unreachable.

6.2. Integrate with vWAN hub using Orchestrator

Follow the instructions in this section to establish BGP between EC-V gateways and vWAN hub using Orchestrator.

6.2.1. Prerequisites

Before you begin, ensure the following prerequisites are met:

- **Orchestrator version:** 9.6.0 or later.
- **ECOS version:** 9.3.3 or later.
- Deploy EC-Vs in a transit VNet, as illustrated in Section [3.1](#) or [3.2](#).
- **Orchestrator BGP ASN pool:** The BGP ASN pool must be enabled and the BGP ASN range must be specified in Orchestrator's configuration. This is required for assigning an ASN to appliances. Note that when several EC-V gateways share an identical site name, the automatic BGP ASN assignment feature allocates the same ASN to all EC-Vs associated with that site name.
- **Azure account configuration:** You must have a valid Azure subscription configured in Orchestrator with the necessary permissions. This is outlined in [6.2.2. Add the Azure subscription to Orchestrator](#).

6.2.2. Add the Azure subscription to Orchestrator

1. In Orchestrator, navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
2. Click **Subscription**.

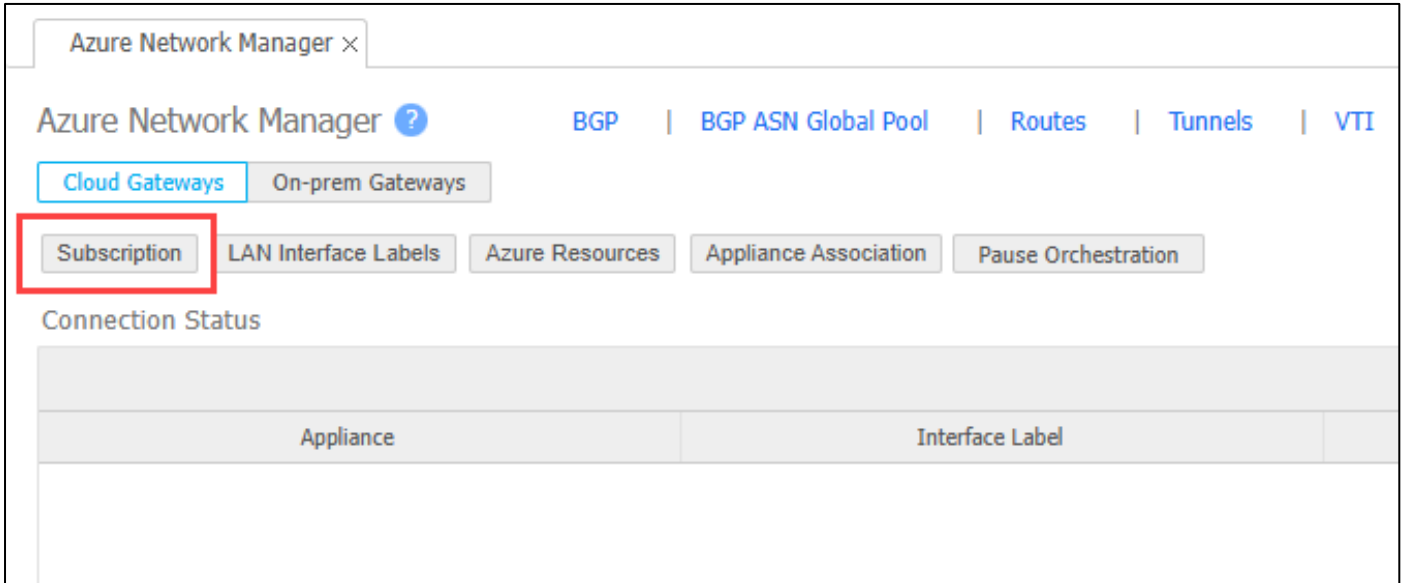


Figure 48. Clicking the Subscription tab in Azure Network Manager (3).

The Subscription for Azure Cloud Gateways LAN-side Automation dialog box appears.

3. If an Azure subscription has already been added to the Orchestrator (used for automated EC-V gateway deployment), it is possible to add the required permissions for the vWAN hub BGP automation to its associated custom role. This is required because the EC-V gateway deployment does not include permissions needed to automate BGP connectivity to a vWAN hub. This is described in [6.2.2.1. Modify an existing custom role to add the permissions needed for vWAN hub deployment](#). If you do not have an existing Azure subscription in the Orchestrator, go to [6.2.2.2. Add a new custom role and the permissions needed for vWAN hub deployment](#).

Note

If any modifications you make (such as creating or deleting subnets) in Azure are not reflected on the Orchestrator after you add the Azure subscription to Orchestrator, click the **refresh** icon, as shown in Figure 49. Clicking the refresh icon retrieves the latest Azure configuration on demand.

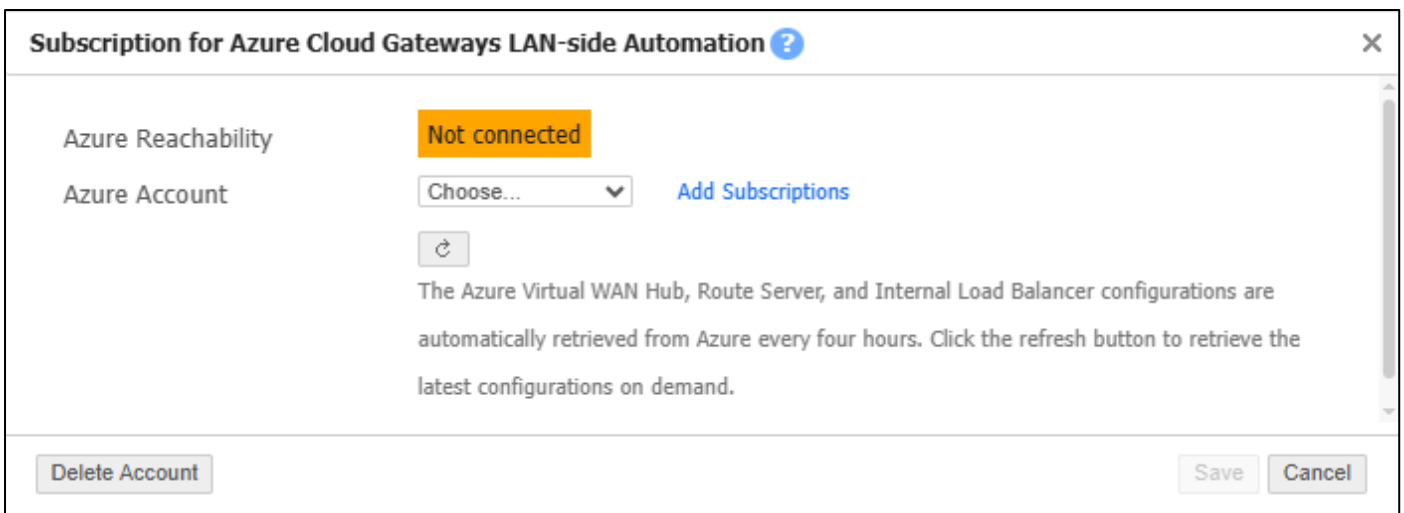


Figure 49. Selecting your Azure subscription in the Subscription for Azure Cloud Gateways LAN-side Automation dialog box of the Orchestrator UI (3).

6.2.2.1. Modify an existing custom role to add the permissions needed for vWAN hub deployment

If you have added an Azure subscription to deploy EC-V gateways previously, update its custom role to include the extra permissions needed to associate the EC-V gateways to the vWAN hub.

1. Go to the resource group used for the EC-V deployment. This is the same resource group where the existing custom role is defined.
2. Navigate to **Access control (IAM) > Roles**.

3. In the search bar, enter the name of the custom role previously created for Orchestrator.
4. With the custom role selected, click the three dots on the far right of the row, and then select **Edit**.

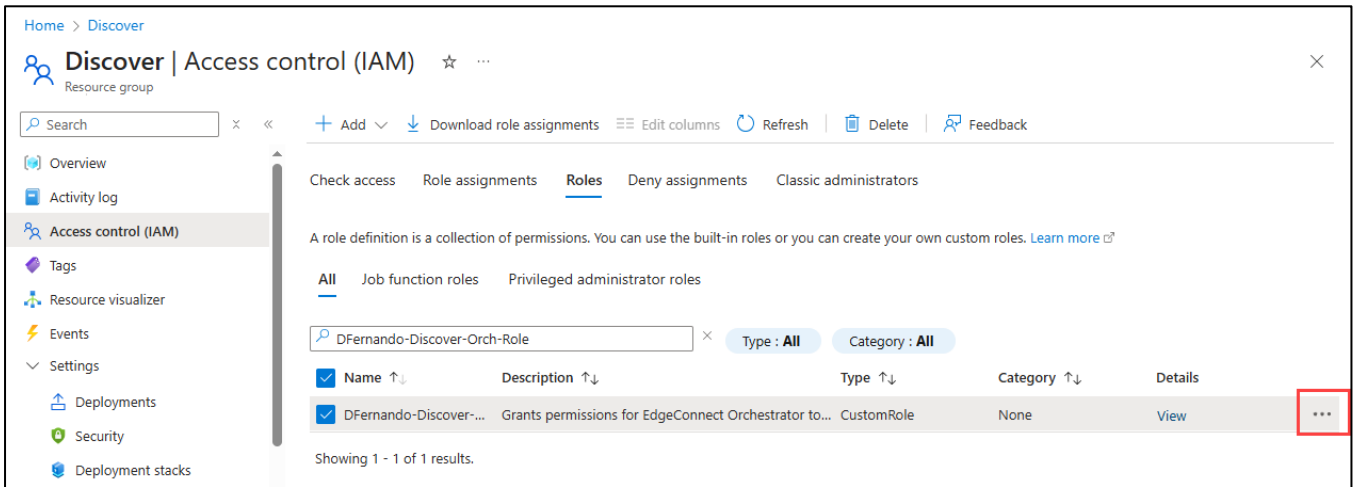


Figure 50. Modifying an existing custom role to add the permissions needed for ILB deployment (2).

5. On the Update a custom role page, click the **JSON** tab.
6. Click **Edit**.
7. Scroll to the bottom of the Actions section in the JSON editor and, after the last permission entry, insert the vWAN hub-related permissions as documented on the [Permissions required to establish BGP connectivity with an Azure Virtual WAN Hub from EC-Vs deployed in a Transit VNet](#) section of the HPE Aruba Networking EdgeConnect SD-WAN Documentation site. Note that the final permission line does not have a comma. Be sure to add a comma before adding the new permissions.

6.2.2.2. Add a new custom role and the permissions needed for vWAN hub deployment

If EdgeConnect was not deployed from Orchestrator previously, but you need to associate a vWAN hub to EC-V gateways through Orchestrator, create a new custom role with only the permissions required to associate the EC-V to the vWAN hub.

Creating a new app registration to add the Azure subscription details on Orchestrator is described in Sections [3.1.2.2](#), [3.1.2.3](#), [3.1.2.4](#), [3.1.2.5](#), and [3.1.3](#). The vWAN-related permissions you must enter to establish BGP can be found on the [Permissions required to establish BGP connectivity with an Azure Virtual WAN Hub from EC-Vs deployed in a Transit VNet](#) section of the HPE Aruba Networking EdgeConnect SD-WAN Documentation site.

6.2.3. Verify vWAN hub region

Unlike ILB and ARS, you cannot create a new vWAN hub directly from Orchestrator. Make sure you have created the vWAN hub from Azure Portal or with your preferred IaaS tool, and that it is in the same region as the EC-V gateways.

6.2.4. Configure LAN interface labels

After you verify the vWAN hub's region, to automate BGP connectivity between the EC-V gateway and vWAN hub, you must create a LAN interface label in Orchestrator, assign it to a LAN interface of the EC-Vs, and drag and drop it into the Primary section of the Establish connectivity using these network interfaces dialog box. These steps are detailed below:

1. In Orchestrator, navigate to **Configuration > Overlay & Security > Interfaces Labels**.
2. Click **New Label**.
3. Select **lan**, and then enter a label name.
4. Click **Save**.

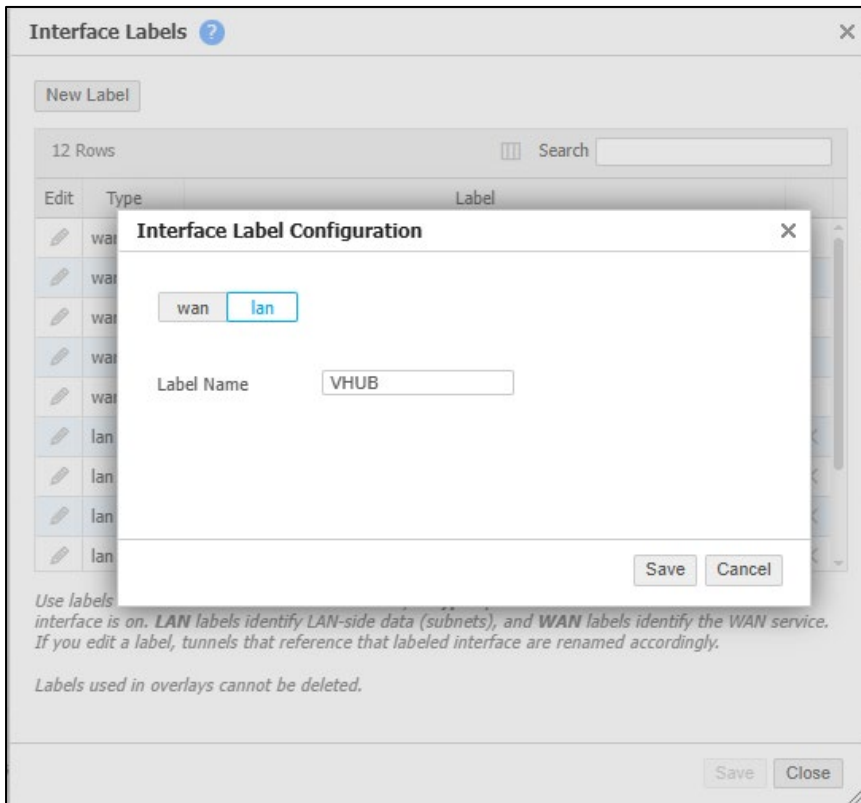


Figure 51. Creating a LAN interface label (3).

- After the label is created, assign it to a LAN interface of each Azure EC-V that needs to establish BGP with vWAN hub.

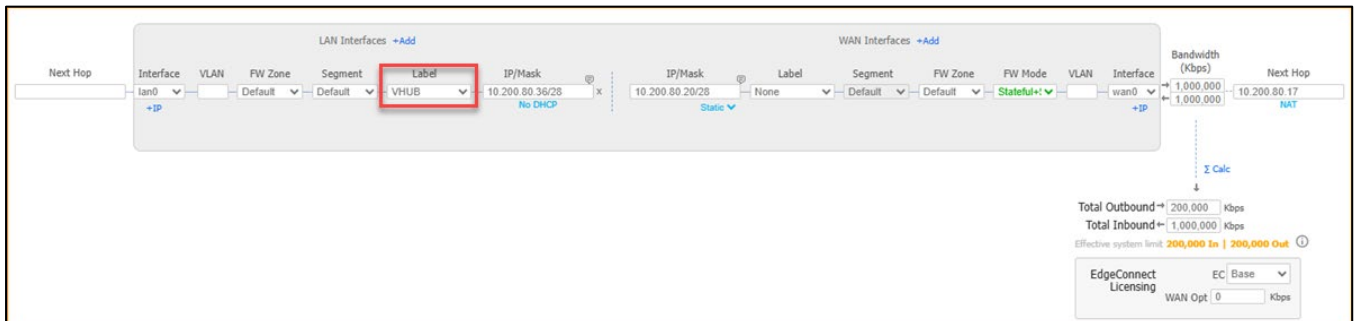


Figure 52. Assigning the label to a LAN interface of an Azure EC-V (3).

- Navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
- Click **Subscription**.
- Select the Azure account (subscription) you want to use.
- Click **Save**.
- Click **LAN Interface Labels**.
- Drag and drop the LAN-side label into the Primary section of the Establish connectivity using these network interfaces dialog box.

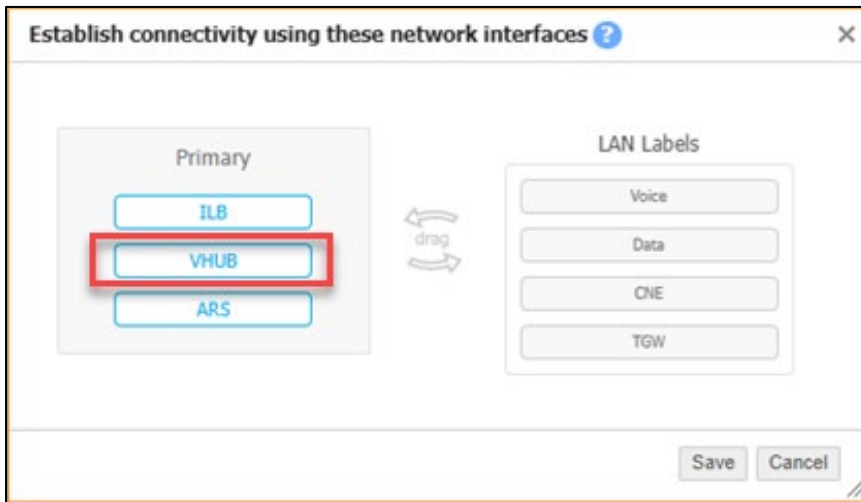


Figure 53. Dragging and dropping the LAN-side VHUB label into the Primary section of the Establish connectivity using these network interfaces dialog box.

12. Click **Save**.

6.2.5. Configure Azure resources in Orchestrator

This step identifies the resource group, region, and VNet of the Azure EC-V gateways that participate in the automation.

1. On the Azure Network Manager tab, click **Azure Resources**.
2. Click **Azure Gateways**.
3. On the Configure SD-WAN VNet Resources for cloud-deployed EdgeConnect appliances dialog box, click **Add**.
4. On the Configure Azure Resources dialog box, enter the following settings:
 - a. **Rule Name:** Enter a descriptive name to identify the EC-V gateways and their VNet.
 - b. **Appliances:** Select the EC-V gateways that should establish BGP with the vWAN hub. You can either enter the appliance names or click **Use Tree Selection** to browse and select them.
 - c. **Resource Group:** From the drop-down list, select the Azure resource group where the EC-V gateways are deployed.
 - d. **Region:** Select the Azure region that corresponds to the deployment location of the EC-V gateways.
 - e. **Virtual Network:** Choose the VNet where the EC-V gateways reside. This ensures that the appliances are correctly associated to the vWAN hub.
5. When all fields are completed, click **Save** to apply the configuration.

Figure 54. Selecting the resource group, region, and VNet of the Azure EC-V gateways that participate in the vWAN hub automation.

6.2.6. Associate EC-V gateways with vWAN hub

This step prompts you to select the EC-Vs and the vWAN hub that need to establish BGP sessions.

Note

In Orchestrator version 9.6.0, if you select multiple EC-Vs on the Orchestrator appliance tree and associate them with a vWAN hub, only one EC-V gateway's BGP sessions come up. This is fixed in Orchestrator version 9.6.1 and later.

1. In the appliance tree in Orchestrator, select each EC-V that you want to establish BGP with the vWAN hub.
2. On the Azure Network Manager tab, click **Appliance Association**.
3. Click **Virtual WAN Hub**. You will see all vWAN hubs to which the Orchestrator has access. If you do not see the vWAN hub that you want to associate your EC-Vs with, check the permissions assigned to the Orchestrator's custom role on Azure Portal.
4. From the left-side panel, select the vWAN hub you want to associate.
5. Click **Save**.

6.2.7. Verify connectivity

After you click Save, the Orchestrator begins automating connectivity and establishing BGP with the vWAN hub. To check the progress of the automation, perform the following tasks:

5. **Check audit logs:** In Orchestrator, navigate to **Orchestrator > Orchestrator Server > Tools > Audit Logs**, and then enter **AzureLanConfigurationManager** to filter the logs. This progress of automation will be displayed.
1. **Connection status:** After a few minutes, navigate back to **Configuration > Cloud Services > Microsoft Azure Network Manager**. The connection status section displays the details of the association and its status.
2. **Check in Azure Portal:** Log in to your Azure Portal and check activity logs. Also, verify the status of the BGP sessions. Confirm route exchange by checking EdgeConnect route tables and Azure effective routes.

| Appliance | Interface Label | Azure Service | Region | Connection Status |
|-------------------|-----------------|---|------------------|------------------------|
| JOHANNESBURG1-ARS | ARS LAN | sp-automated-JOHANNESBURG-ARS-route-server Route Server | southafricanorth | Up |
| JOHANNESBURG2-ARS | ARS LAN | sp-automated-JOHANNESBURG-ARS-route-server Route Server | southafricanorth | Up |
| PUNE1-VHUB | VHUB LAN | CentralIndia-Hub Vhub | centralindia | Up |
| PUNE2-VHUB | VHUB LAN | CentralIndia-Hub Vhub | centralindia | Up |
| QATAR1-ILB | ILB LAN | sp-automated-QATAR-ILB-ilb Load Balancer | qatarcentral | Load Balancer Insights |
| QATAR2-ILB | ILB LAN | sp-automated-QATAR-ILB-ilb Load Balancer | qatarcentral | Load Balancer Insights |

Figure 55. Successfully associating EC-Vs with vWAN hub, as indicated by the *Up* connection status for the Pune EC-Vs.

| Edit | Appliance | Segment | Peer IP | Address Family | Route Target... | Local Interface | Peer ASN | Peer State | Soft Reset | Established Time | Peer Details |
|------|------------|---------|---------------|----------------|-----------------|-----------------|----------|-------------|------------|------------------|--------------|
| | PUNE2-VHUB | Default | 192.168.50.68 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 2d 28m 19s | |
| | PUNE2-VHUB | Default | 192.168.50.69 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 2d 28m 19s | |
| | PUNE1-VHUB | Default | 192.168.50.68 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 2d 27m 18s | |
| | PUNE1-VHUB | Default | 192.168.50.69 | IPv4 Unicast | N/A | Ian0 | 65515 | Established | | 2d 27m 21s | |

Figure 56. Verifying BGP sessions of both EC-Vs (2).

You have successfully associated your EC-Vs with the vWAN hub and established BGP.

6.2.8. Create a virtual network connection between spoke (workloads) VNet and vWAN hub

After establishing BGP sessions, you must establish a virtual network connection between the spoke VNets where your workloads reside and the vWAN hub. The virtual network connection allows workloads in spoke VNets to send and receive traffic from the EC-V gateways in the transit VNet. Because this is a common topic for both vWAN hub (transit VNet and Managed NVA) designs, this is explained in Section 8. To create the VNet peering, follow the instructions in [8.2. Create a virtual network connection \(for vWAN hub-related designs\)](#). After it is done, perform the following steps to verify if the destination routes are being properly learned by the workloads in spoke VNets:

1. Navigate to a vNIC of one of the workloads in your spoke VNet that is connected to the vWAN hub.
2. Under Help, select **Effective routes**.
3. Verify that routes advertised by your EC-V appliances are visible. Confirm that the route's Next Hop Type is *Virtual network gateway*.
4. On the Routes page of the Azure EC-Vs, ensure that the routes are being learned and advertised to and from ARS.

6.2.9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric

To avoid the two static routes that were created on each EC-V (to establish BGP sessions) from being advertised to other SD-WAN devices via subnet sharing, see [9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric](#).

6.3. Integrate with vWAN hub manually from Azure Portal

Follow the instructions in this section to manually integrate EC-V gateways with vWAN hub from the Azure Portal.

6.3.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- Deploy EC-Vs in an SD-WAN transit VNet, as illustrated in Section [3.1](#) or [3.2](#).

6.3.2. Create a vWAN hub

To create a vWAN hub from the Azure Portal:

1. In the Azure Portal, navigate to **Virtual WANs**, and then click **+ Create**.
2. On the Create WAN page, verify that **Basics** is selected on the menu bar, and then enter the following settings:
 - a. **Subscription:** Select your subscription.

- b. **Resource group:** Create new or use existing.
 - c. **Resource group location:** Choose a location. You can select any region.
 - d. **Name:** Give the vWAN a name.
 - e. **Type:** Select **Standard**. *Basic* is not sufficient for BGP and hub routing.
 - f. Click **Review + create**.
3. Open the vWAN, and then click **Hubs**.
 4. Click **+ New Hub**, and then enter the following settings:
 - a. **Region:** Select the same region as the EC-Vs.
 - b. **Name:** Enter a name for the hub.
 - c. **Hub private address space:** Specify a hub private address space. The minimum address space is /24. This range should not overlap with your VNets or on-premises networks.
 - d. **Virtual hub capacity:** Select your required virtual hub capacity.
 - e. **Hub routing preference:** Select **AS Path**. This ensures that the vWAN hub prefers routes with the shortest BGP AS-Path length irrespective of the source of the route advertisements.
 5. Click **Review + create**, and then click **Create**.

Note

A site-to-site gateway was not created when the vWAN hub was created. This is because the EC-V establishes BGP directly with the vWAN hub's virtual router. This avoids the need for a VPN gateway within the vWAN hub.

6. Wait for the new hub to finish provisioning. When it is ready, the Routing status will appear as *Provisioned* on the hub's Overview page.

6.3.3. Create a virtual network connection

After the vWAN hub is provisioned, you must connect your transit VNet (where the EC-V gateways reside) to the vWAN hub. This is explained in [8.2. Create a virtual network connection \(for vWAN hub-related designs\)](#). Similarly, connect your spoke VNets (where your workloads that need to forward traffic to the EC-V gateways are deployed) by creating a virtual network connection to the vWAN hub.

6.3.4. Add EC-V gateways as BGP peers on the vWAN hub

With the transit VNet attached, you are ready to set up BGP peering in the vWAN hub so it can exchange routes with the EdgeConnect appliances.

1. In the Azure Portal, select the route server you created.
2. From the left-side navigation pane, click **BGP Peers**.
3. Click **+ Add**.
4. Enter the following information about the EdgeConnect appliance:
 - a. **Name:** Enter a name (for example, *EC-V-Gateway-1*).
 - b. **ASN:** The BGP ASN of the EC-V gateway. Enter any BGP ASN between 64512 and 65534, excluding the following ASNs:
 - I. ASNs reserved by Azure:
 1. Public ASNs: 8074, 8075, 12076
 2. Private ASNs: 65515, 65517, 65518, 65519, 65520
 - II. ASNs reserved by IANA:
 1. 23456, 64496-64511, 65535-65551

Note

vWAN hub only supports 16-bit ASNs. Do not enter a 32-bit ASN. The ASN you enter here is entered on the EC-V gateway when enabling BGP.

- c. **IPv4 Address:** Enter the IP address of the EC-V gateway's LAN0 interface in the transit VNet. Azure requires peering with an interface IP; you cannot use a loopback IP.
- d. **Virtual network connection:** Select the transit VNet connection you created in the previous step.
5. Click **Add** to add the first EC-V gateway as a BGP peer.
6. Repeat steps 2–5 to add the second EC-V as a BGP peer. You must enter the same ASN on both BGP peers. After adding the peers, Azure displays two BGP endpoint IP addresses for the hub's router (for high availability).
7. You have successfully added EC-V gateways as BGP peers on the vWAN hub.

Note

Azure provides two IP addresses for the vWAN hub. You must configure both as BGP neighbors on each EC-V gateway. vWAN hub uses active-active BGP, so peering with both hub router IPs is required for full redundancy.

6.3.5. Create two static routes for enabling BGP

The vWAN hub contains two BGP endpoints. This means you must create two static routes on each EC-V gateway to establish BGP with the vWAN hub. Perform the following tasks to create the static routes on EC-Vs:

1. From the Azure dashboard, find the vWAN hub endpoint IPs.
2. In Orchestrator, select the first EC-V in the appliance tree.
3. Navigate to **Configuration > Networking > Routing > Routes**, and then select **Local/Static**.
4. Click the **edit** (pencil) icon on the Default segment.
5. On the Routes page for the Default segment, enter the following settings:
 - a. **Automatically advertise local LAN subnets:** Clear the check box.
 - b. **Automatically advertise local WAN subnets:** Clear the check box.
 - c. **Metric for automatically added subnets:** Enter **50**.
 - d. **Redistribute routes to SD-WAN Fabric:** Select an available route map.
 - e. **Filter routes from SD-WAN fabric with matching local ASN:** Select the check box.
 - f. **Include BGP Local ASN to routes sent to SD-WAN Fabric:** Select the check box.
 - g. **Tag BGP communities to routes:** Clear the check box.
 - h. Click **Apply**.
6. Click **Add Route**.
7. Enter the following settings:
 - a. **Subnet/Mask:** Enter the first vWAN hub BGP endpoint (/32) IP address.
 - b. **Next Hop:** Enter the EC-V's LAN0 interface's next hop IP address. This is the first IP address of the LAN0 subnet. (If you are establishing BGP sessions from other interfaces, such as the LAN1 interface, make sure to enter the first IP address of that interface's subnet.)
 - c. **Interface:** LAN0
 - d. **Zone:** Default
 - e. **Metric:** Leave the default value (**50**).
 - f. **Tag:** ANY
 - g. **Comments:** Optional
8. To add the first static route, click **OK**.
9. Repeat steps 5–7 and create a static route for the second vWAN hub BGP endpoint.
10. Repeat steps 2–7 on the second EC-V, creating two static routes on that EC-V as well.

6.3.6. Add vWAN hub as a BGP peer on the EC-V gateways

Now that the static routes are created and BGP peers are created on the vWAN hub, perform the following steps to enable BGP on the EC-V gateways:

1. In Orchestrator, select the first EC-V in the appliance tree.
2. Navigate to **Configuration > Networking > Routing > BGP**.
3. Click the **edit** (pencil) icon on the Default segment row.
4. Click the **Enable BGP** toggle.
5. Enter the following BGP settings:
 - a. **Autonomous system number:** Enter the EC-V's ASN that you entered in [6.3.4. Add EC-V gateways as BGP peers on the vWAN hub](#).
 - b. **Route Target:** Leave empty.
 - c. **Router ID:** The Router ID is an IPv4 address by which the remote peer can identify this EC-V gateway for purposes of BGP.
 - d. **Graceful restart:** Select the check box.
 - e. **AS path propagate:** Select the check box.
 - f. **Log BGP update messages:** Select the check box.
 - g. **Max route updates per peer:** 10
 - h. **Detection interval:** 15 minutes
 - i. Click **Add Peer**, and then enter the following settings:
 - I. **Peer IP:** Primary Azure BGP Peer (same address as the /32 Static Route added in the previous step)
 - II. **Peer Adjacency:** Multi-Hop
 - III. **EVPN Peer:** Clear the check box.
 - IV. **Local Interface:** lan0
 - V. **Peer ASN:** 65515
 - VI. **Override ASN:** Select the check box.
 - VII. **Peer Type:** Branch
 - VIII. **Admin Status:** Up
 - IX. **Soft Reconfiguration:** Select the check box.
 - X. **Next-Hop Self:** Select the check box.
 - XI. **Inbound route map:** Select the appropriate inbound route map.
 - XII. **Outbound route map:** Select the appropriate outbound route map.
 - XIII. **BFD:** Clear the check box.
 - XIV. **Keep Alive Timer:** 5
 - XV. **Hold Timer:** 15
 - XVI. **Enable MD5 Password:** Clear the check box.
 - XVII. To create the first BGP session, click **Add**.
6. Repeat steps 2–5 to create the second BGP session.
7. Repeat the same steps on the second EC-V.
8. Verify BGP configuration by checking the BGP neighbor status in the Orchestrator.

6.3.7. Verify connectivity

After enabling BGP on the EC-V gateways, verify that the BGP sessions come up on both sides.

On Orchestrator's BGP tab (Configuration > Networking > Routing > BGP), check the Peer State for the Azure hub neighbors on each EC-V. A state of *Established* indicates that the BGP adjacency is up and routes can be exchanged. To see information such as number of routes learned/advertised via that peer, click the **peer details** icon.

In the Azure Portal, the vWAN hub's BGP Peers page should show the EdgeConnect peers with a status of *Connected/Established*.

After BGP is established, the EC-V gateways dynamically advertise their routes to the vWAN hub, and the vWAN hub advertises routes present on its route table to the EC-V gateways. Verify that the EC-Vs are receiving Azure routes by checking the routing table on each EC-V (in Orchestrator's Routes table) to see BGP-learned routes from Azure. Likewise, Azure's route tables (in the default route table of the vWAN hub or Effective Routes for the VNet connection) should now include routes for the networks advertised by the EC-V gateways. This confirms end-to-end dynamic routing is in place.

You have successfully established BGP between your EC-Vs and vWAN hub.

6.3.8. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric

To avoid the two static routes that were created on each EC-V (to establish BGP sessions) from being advertised to other SD-WAN devices via subnet sharing, see [9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric.](#)

7. Deploy EC-V within Azure vWAN hub (Managed NVA)

This section describes how to deploy EC-V gateways directly within an Azure vWAN hub and establish BGP. This design is also referred to as *Managed NVA* or *NVA in vWAN hub*. To learn more about this design, see [2.2. Decide whether to deploy EC-V into transit VNet or within vWAN hub \(Managed NVA\)](#).

7.1. Deploy EC-V within Azure vWAN hub (Managed NVA) from Orchestrator

7.1.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- **Orchestrator version:** 9.6.0 or later.
- **ECOS version:** 9.3.3 or later.
- **Azure account configuration:** You must have a valid Azure subscription configured in Orchestrator with the necessary permissions. This is outlined in [7.1.2. Configure the Azure subscription](#).

Notes

- When an Orchestrator deploys a Managed NVA, usually the EdgeConnect OS (ECOS) version of that EC-V is earlier than the version you want. Upon discovering the EC-V on the Orchestrator, you can upgrade it before or after you add it to the Orchestrator's appliance tree. The upgrade takes only a few minutes.
- As of November 2025, Orchestrator does not support deploying Managed NVA (using the Cloud Hubs in Azure feature) into Azure Government Cloud or Azure China. To manually deploy EC-V into these two environments, follow the instructions in [7.2. Deploy EC-V within Azure vWAN hub \(Managed NVA\) from Azure Portal](#).

7.1.2. Configure the Azure subscription

Before deploying Managed NVA via Orchestrator, you must configure your Azure subscription by creating an app registration (service principal), a resource group, and a custom role, and assigning that role to the resource group.

7.1.2.1. Create a new app registration

Azure app registrations are used to create service principals, which are identities used by applications or services to access Azure resources securely and with restricted permissions. The Orchestrator uses the app registration to authenticate itself with Azure. An app registration allows you to assign least-privilege access using custom roles. It avoids sharing credentials or exposing elevated privileges unnecessarily.

To create a new App registration for the Orchestrator:

1. Log in to the Azure Portal.
2. In the main search menu, enter **app registrations**, and then click **App registrations**.
3. Click **+ New registration**.
4. On the Register an application page, in the Name field, enter a user-facing display name for the application.
5. Under Supported account types, select **Accounts in this organizational directory only (Default Directory only - single tenant)**.
6. *(Optional)* Enter a redirect URI.
7. Click **Register**. Note the Application (client) ID and Directory (tenant) ID. You will need these IDs when you add the subscription details on the Orchestrator.
8. Under Manage, click **Certificates & secrets**.
9. Click **New client secret**.
10. Enter a Description and Expiration Date.
11. Click **Add**.
A new client secret is created.
12. Copy the text in the Value column.

Note

This text can only be viewed immediately after creation. Be sure to save the secret before leaving the page.

13. On the main search menu bar, enter **subscription**, and then press **Enter**.

14. Copy the subscription ID.

You have successfully registered your application and gathered the details required for adding the Azure subscription details to the Orchestrator.

7.1.2.2. Create a new resource group

Creating a new resource group on the Azure Portal is considered the best practice. This ensures that the SD-WAN Orchestrator only has access to that resource group to deploy EC-Vs. However, it is possible to deploy one or more EC-Vs into an existing resource group that contains other Azure resources.

To create a new resource group:

1. From the main search menu in Azure Portal, enter **resource group**, and then click **Resource groups**.
2. Click **+ Create**.
3. On the Create a resource group page, select the subscription that you want to use to create the resource group.
4. Enter a name for the resource group, and then select a region.
5. Click **Review + create**.
6. Click **Create**.

7.1.2.3. Create a custom role

To create custom roles, you must have Owner or User Access Administrator permissions. There are multiple ways to create a custom role. The following steps create a custom role from within the resource group you created.

1. Select the resource group you created in the previous section, and then click **Access control (IAM)**.
2. Click **Add**, and then click **Add custom role**.
The Custom Roles editor appears (the Basic tab is displayed).
3. In the Custom role name field, enter a name for the custom role. The name must be unique for the Azure AD directory. The name can include letters, numbers, spaces, and special characters.
4. *(Optional)* In the Description field, enter a description for the custom role. The description displays in the tool tip for the custom role.
5. Accept the default value for the Baseline permissions, and then click the **JSON** tab.
6. Click **Edit**.
7. Go to https://arubanetworking.hpe.com/techdocs/SilverPeak/files/cloud-ecv/cloud_ecv_json.htm, and then click the **Permissions required to deploy EC-Vs inside an Azure Virtual WAN Hub-Managed NVA and establish BGP connectivity**.
8. Copy the list of Azure permissions, and then paste the list within the square brackets under Actions (line 10), as shown in Figure 57.

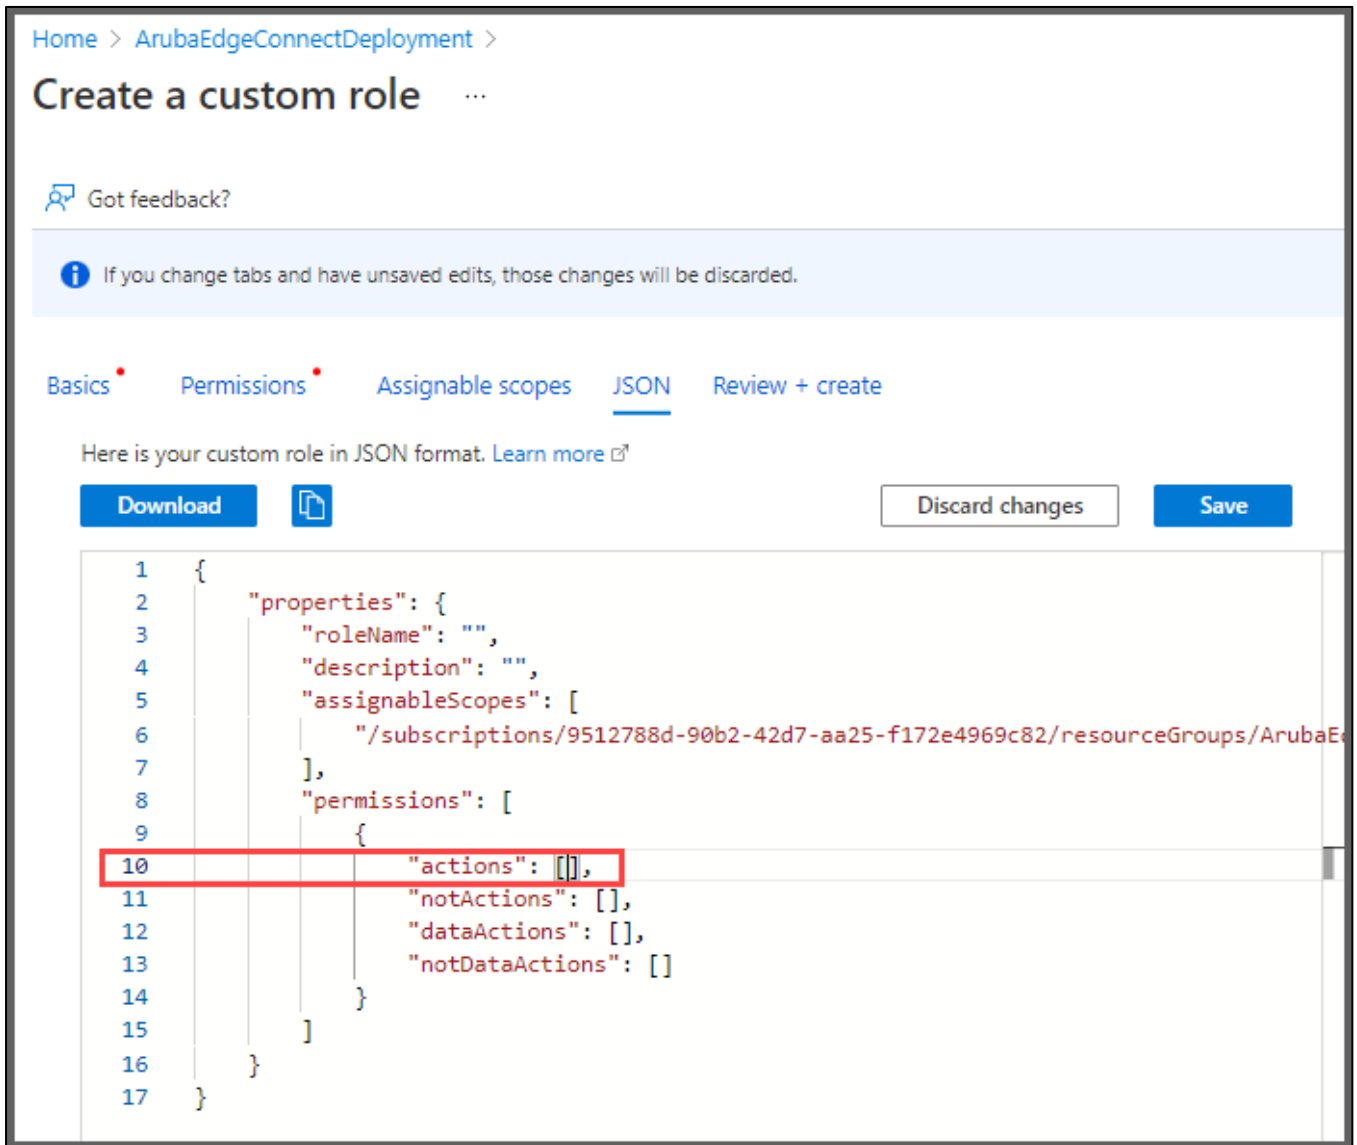


Figure 57. Finding the actions field in the Custom Roles editor (2).

9. Click **Save**.
10. Click the **Assignable scopes** tab, and then verify that the resource group you created is added as an assignable scope and Type is set to the resource group.
11. Click the **Permissions** tab, and then verify that the permissions, descriptions, and permission types you added are listed.
12. Click **Review + create**.
13. Click **Create**.

A message displays to confirm that you have successfully created your custom role.

7.1.2.4. Assign the custom role to the resource group

Assigning the custom role to the resource group ensures that the Orchestrator has the exact set of permissions needed to deploy only the resources within that resource group, in line with the principle of least privilege. However, if you want to deploy EdgeConnect gateways into multiple resource groups, you can assign the custom role at the Azure subscription level. This allows the Orchestrator to list multiple resource groups and lets you select the one you want for deployment. This example assigns the custom role only to the specific resource group.

1. Navigate to the resource group you created, and then click **Access control (IAM)**.

Note

If you just completed the previous task of creating a custom role, the Access control (IAM) page is already open.

2. Click **Add**, and then click **Add role assignment**.
3. On the Role tab, enter the name of your custom role.

Note

If the role you created is not displayed, refresh the page.

4. Select the custom role, and then click **Next**.
5. On the Members tab, ensure that **User, group, or service principle** is selected, and then click **+ Select members**.
6. On the Select members tab, enter the name of your App registration (Service Principle), select your app, and then click **Select**.
Your app is added under Members.
7. Click **Review + assign**.
8. Click **Review + assign** again.

You have successfully assigned your custom role to the resource group.

7.1.3. Add Azure subscription details on Orchestrator

This section explains how to add your Azure subscription details—such as Subscription ID, Directory (tenant) ID, Application (client) ID, and Client secret value—to Orchestrator. You only need to do this once per Azure subscription unless you are a large enterprise with multiple teams that want to use different resource groups to segregate deployments.

To add the Azure subscription to Orchestrator:

1. Log in to Orchestrator, and then navigate to **Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in Azure**.
2. Click **Azure Subscriptions**, and then click **Add Azure Subscription**.
3. Enter a name, subscription ID, directory (tenant) ID, application (client) ID, and client secret for the Azure subscription.

Notes

- You copied the directory (tenant) ID, application (client) ID, and client secret in [7.1.2.1 Create a new app registration](#).
- If you copy and paste the subscription ID, Azure might add a blank space to the beginning of the subscription ID. Be sure to remove all spaces from your subscription ID.

4. Click **Save**.
Orchestrator validates the subscription information.

Note

If you have multiple Azure subscriptions, you can load the credentials for each subscription.

7.1.4. Deploy Managed NVA

To deploy Managed NVA from Orchestrator:

1. Log in to the Orchestrator, and then navigate to **Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in Azure**.
2. Click **Deploy Cloud Hubs in Virtual WAN Hub**.
3. Enter the deployment details shown below.

Note

If you do not have an Azure subscription configured in Orchestrator, the Azure Deployment Configuration dialog box is blank. Click the **Subscriptions** link to go to the Azure subscription page and enter your Azure credential.

| Field | Description |
|--------------------|--|
| Name | Enter a name for the deployment. This name is used only for identifying the deployment. A deployment consists of two EC-Vs created within a vWAN Hub. Only alphabetical letters and hyphens are allowed in the deployment name. The maximum allowed length is 20 characters. |
| Azure account | Select an Azure subscription to use for deploying the Managed NVA. |
| Resource group | Select an Azure resource group to use for deploying the Managed NVA. |
| Region | Select an Azure region where you want to deploy the Managed NVA. |
| Virtual WAN hub | Select Create new virtual network or Existing virtual network . Select your vWAN Hub. If no vWAN Hubs are shown, the region you selected does not have any vWAN Hubs created, or you have not entered the required permissions to your custom role. |
| Scale units | The scale unit determines the resource allocation on the Managed NVA. The higher the scale unit, the greater the amount of traffic that can be handled. |
| EdgeConnect ASN | <p>Enter an ASN between 64512 and 65534, excluding the following ASNs:</p> <ul style="list-style-type: none"> i) ASNs reserved by Azure: <ul style="list-style-type: none"> (1) Public ASNs: 8074, 8075, 12076 (2) Private ASNs: 65515, 65517, 65518, 65519, 65520 ii) ASNs reserved by IANA: <ul style="list-style-type: none"> (1) 23456, 64496-64511, 65535-65551 <p>Note: vWAN Hub only supports 16-bit ASNs. As a result, do not enter a 32-bit ASN. The ASN you enter here will be assigned to both EC-Vs.</p> |
| SSH public key | <p>Generate a public key with an application, such as PuTTYgen, and then input the value here.</p> <p>Important: EdgeConnect only supports single-line SSH public keys. Do not use multi-line SSH public keys. Additionally, use an EdDSA (ED25519) key pair as shown on the image below:</p> <p>Use this:</p> <pre>1 ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIOFDa8CNR5DmJE3EgIOBpYyyL4mlyqaRVo/+XBgISULs eddsa-key-20251015</pre> <p>Not this:</p> <pre>1 ---- BEGIN SSH2 PUBLIC KEY ---- 2 Comment: "eddsa-key-20251015" 3 AAAAC3NzaC1lZDI1NTE5AAAAIOFDa8CNR5DmJE3EgIOBpYyyL4mlyqaRVo/+XBgI 4 SULs 5 ---- END SSH2 PUBLIC KEY ---- 6</pre> <p>Note: Save the private key file. If you need to log in via SSH to the appliance after it is deployed, you will need this key.</p> |
| Comment (Optional) | Enter an optional comment if you want to attach any additional details for the deployment. |

Figure 58. Deploying Managed NVA from Orchestrator.

4. After entering your information as shown in Figure 58, click **Review and Deploy**.
5. Review the configuration summary, and then click **Deploy** to create the EC-V instances. The deployment begins.

Note

To check the status after a few minutes, click the **refresh** icon next to the Deploy Cloud Hubs in Azure button.

6. If your EC-V deployment fails, the status appears as *Incomplete* in the Cloud Hubs in Azure table. To download the log file, click the **info** (i) icon. The reason for the failure is typically at the bottom of the log file.
 - a. To remove (delete) a deployment after a failure, click **Terminate**. All Azure resources created by the Orchestrator will be removed. If termination fails, an Azure resource is blocking the deletion.
 - b. To view details of this Azure resource, download the log file and then check the last error message. To proceed with the termination, you can manually delete this resource on the Azure Portal.

If your deployment succeeds, the status appears as *Deployed*.

7. After the deployment succeeds, the newly deployed EC-Vs appear on the Discovered Appliances tab on the Orchestrator. Please allow at least 10 minutes for EC-Vs to be discovered on the Orchestrator. On the Discovered Appliances tab, the Approve button appears after the EC-V is fully configured.

The following table describes each field on the Cloud Hubs in Azure tab.

| Field | Description |
|-----------------|---|
| Deployment Name | Name given to the Managed NVA deployment. |
| Subscription | Name of the Azure subscription used. |
| Region | Region of the Managed NVA deployment. |
| Resource Group | Name of the Azure Resource Group used for the Managed NVA deployment. |
| Platform | The azure_nva platform. |
| Virtual Network | CIDR block of the vWAN Hub. |
| Instances | Number of EC-V instances deployed in the vWAN Hub. |

| Field | Description |
|-----------------|---|
| Status | Status of the deployment. |
| Terminate | To permanently delete a deployment, click Terminate . This action deletes both EC-Vs. |
| Deployment Info | Click the info (i) icon in this column to view information such as vWAN hub name, region, and private and public IP addresses assigned to each EC-V. |
| Resources | The NVA resource ID. |
| Comment | Comments that were added to the deployment. |

This completes the Managed NVA deployment from Orchestrator. Your EC-V deployment now matches Figure 59.

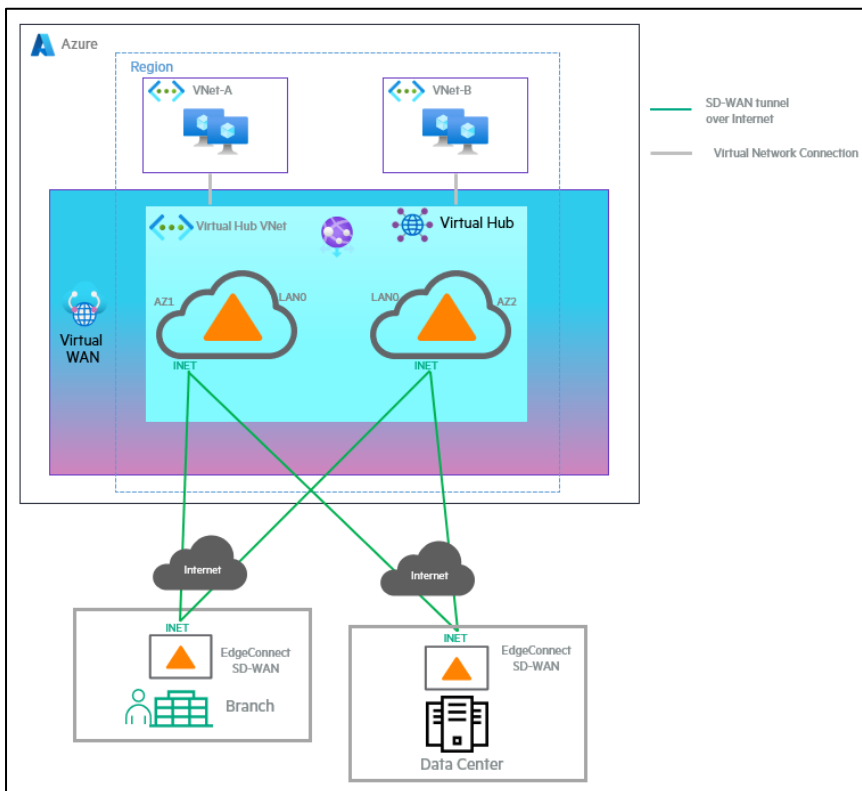


Figure 59. After deploying Managed NVA into a vWAN hub and establishing SD-WAN tunnels.

The next step is to establish BGP with the route service within the virtual hub using Orchestrator’s automation, as shown in Figure 60. (The yellow dotted lines indicate the BGP sessions from the LAN0 interface.)

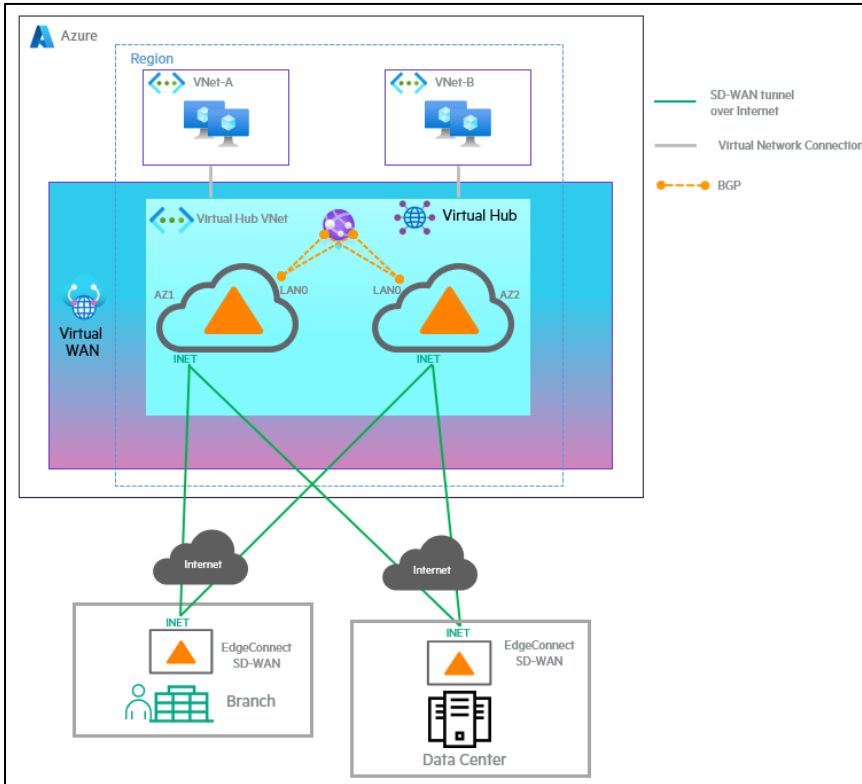


Figure 60. After deploying managed NVA with SD-WAN tunnels and LAN-side BGP sessions established.

7.1.5. Configure LAN interface labels

To begin automating BGP session creation, you must first create a LAN interface label in Orchestrator, assign it to a LAN interface of the EC-Vs, and drag and drop it into the Primary section of the Establish connectivity using these network interfaces dialog box. These steps are detailed below:

1. In Orchestrator, navigate to **Configuration > Overlay & Security > Interfaces Labels**.
2. Click **New Label**.
3. Select **lan**, and then enter a label name.
4. Click **Save**.

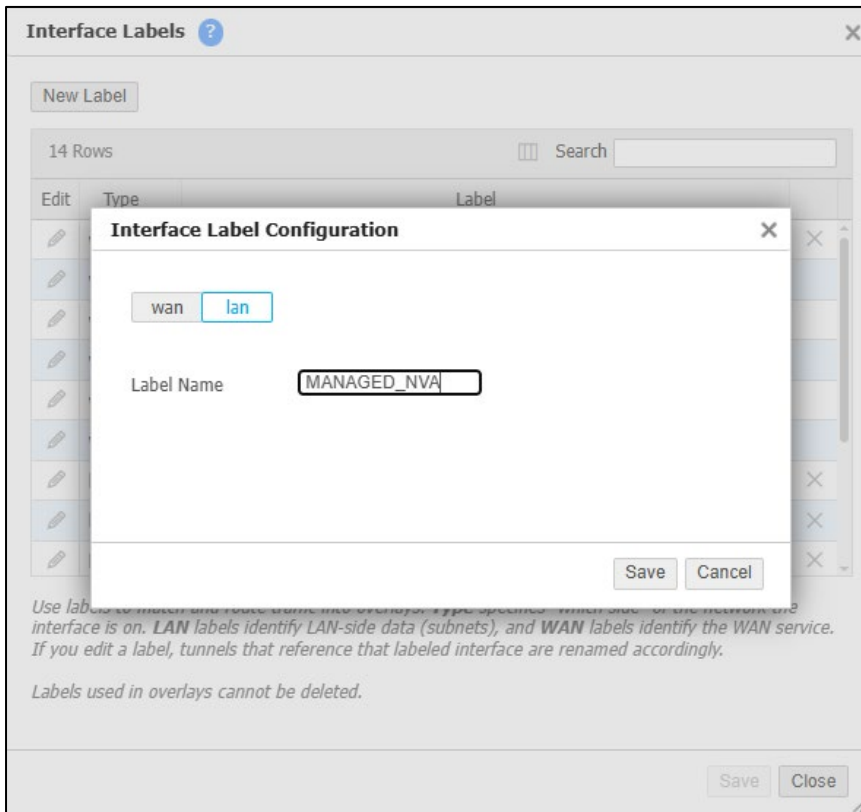


Figure 61. Creating a LAN interface label (4).

5. After the label is created, assign it to a LAN interface of each EC-V that needs to establish BGP with vWAN hub's route service.

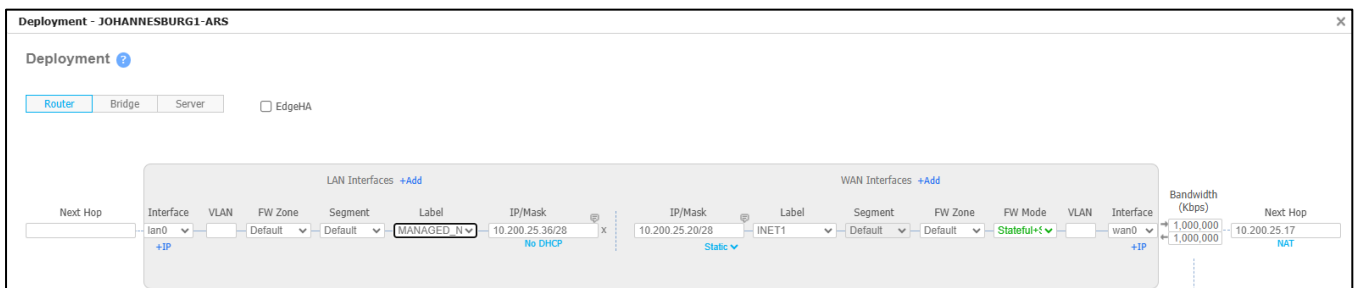


Figure 62. Assigning the label to a LAN interface of an Azure EC-V (4).

6. Navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager**.
7. Click **Subscription**.
8. Select the Azure account (subscription) you want to use.
9. Click **Save**.
10. Click **LAN Interface Labels**.
11. Drag and drop the LAN-side label into the Primary section of the Establish connectivity using these network interfaces dialog box.

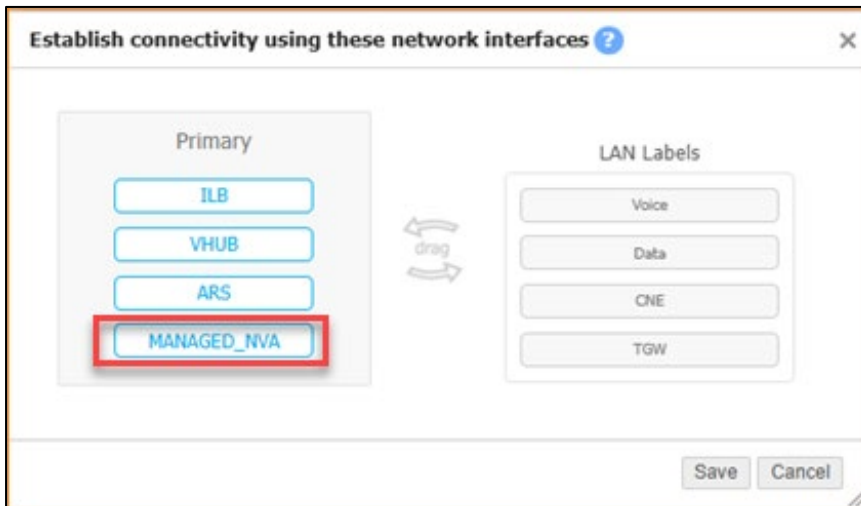


Figure 63. Dragging and dropping the LAN-side MANAGED_NVA label into the Primary section of the Establish connectivity using these network interfaces dialog box.

12. Click **Save**.

7.1.6. Associate EC-V gateways with vWAN hub

This step prompts you to select the EC-Vs and the vWAN hub that need to establish BGP sessions.

Notes

- In a Managed NVA deployment, you do not need to use the Azure Resources dialog box to enter details such as rule name, appliances, resource group, region, and virtual network. Orchestrator does not need this information to establish BGP, as the EC-Vs are already deployed within the vWAN hub's Hub VNet.
- In Orchestrator version 9.6.0, if you select multiple EC-Vs on the Orchestrator tree and associate them with a vWAN hub, only one EC-V gateway's BGP sessions come up. This is fixed in Orchestrator version 9.6.1 and newer.

1. In the appliance tree in Orchestrator, select each EC-V you want to establish BGP with the vWAN hub.
2. On the Azure Network Manager tab, click **Appliance Association**.
3. Click **Virtual WAN Hub**. You can see all vWAN hubs the Orchestrator has access to. If you do not see the vWAN hub that you want to associate your EC-Vs with, check the permissions you have assigned to the Orchestrator's custom role on Azure Portal.
4. From the left-side panel, select the vWAN hub you want to associate.
5. Click **Save**.

7.1.7. Verify connectivity

After you click Save, the Orchestrator begins automating connectivity and establishing BGP with the vWAN hub. To check the progress of the automation:

1. **Check audit logs:** navigate to **Orchestrator > Orchestrator Server > Tools > Audit Logs**, and filter for **AzureLanConfigurationManager**. This shows the progress of automation.
2. **Connection status:** After a few minutes, navigate to **Configuration > Cloud Services > Microsoft Azure Network Manager** in Orchestrator. The connection status section displays the details of the association and its status.
3. **Check in Azure Portal:** Log in to your Azure Portal to check activity logs and verify the status of the BGP sessions. Confirm route exchange by checking EdgeConnect route tables and Azure effective routes.

| Appliance | Interface Label | Azure Service | Region | Connection Status |
|-------------------|-----------------|---|------------------|------------------------|
| JOHANNESBURG1-ARS | ARS LAN | sp-automated-JOHANNESBURG-ARS-route-server Route Server | southafricanorth | Up |
| JOHANNESBURG2-ARS | ARS LAN | sp-automated-JOHANNESBURG-ARS-route-server Route Server | southafricanorth | Up |
| PUNE1-VHUB | VHUB LAN | CentralIndia-Hub VHub | centralindia | Up |
| PUNE2-VHUB | VHUB LAN | CentralIndia-Hub VHub | centralindia | Up |
| QATAR1-ILB | ILB LAN | sp-automated-QATAR-ILB-ilb Load Balancer | qatarcentral | Load Balancer Insights |
| QATAR2-ILB | ILB LAN | sp-automated-QATAR-ILB-ilb Load Balancer | qatarcentral | Load Balancer Insights |

Figure 64. Successful association of EC-Vs with vWAN hub, as indicated by the *Up* connection status for the Pune EC-Vs.

| Edit | Appliance | Segment | Peer IP | Address Family | Route Target... | Local Interface | Peer ASN | Peer State | Soft Reset | Established Time | Peer Details |
|------|------------|---------|---------------|----------------|-----------------|-----------------|----------|-------------|------------|------------------|--------------|
| | PUNE1-VHUB | Default | 192.168.50.68 | IPv4 Unicast | N/A | lan0 | 65515 | Established | | 2d 28m 19s | |
| | PUNE2-VHUB | Default | 192.168.50.69 | IPv4 Unicast | N/A | lan0 | 65515 | Established | | 2d 28m 19s | |
| | PUNE1-VHUB | Default | 192.168.50.68 | IPv4 Unicast | N/A | lan0 | 65515 | Established | | 2d 27m 18s | |
| | PUNE1-VHUB | Default | 192.168.50.69 | IPv4 Unicast | N/A | lan0 | 65515 | Established | | 2d 27m 21s | |

Figure 65. Verifying BGP sessions of both EC-Vs (3).

You have successfully associated your EC-Vs with the vWAN hub and established BGP.

7.1.8. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric

To avoid the two static routes that were created on each EC-V (to establish BGP sessions) from being advertised to other SD-WAN devices via subnet sharing, see [9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric.](#)

7.2. Deploy EC-V within Azure vWAN hub (Managed NVA) from Azure Portal

Follow the instructions in this section to manually integrate EC-V gateways with vWAN hub from the Azure Portal.

7.2.1. Prerequisites

Before you begin, ensure that the following prerequisites are met:

- An Azure subscription
- HPE Aruba Networking EdgeConnect Orchestrator
- HPE Aruba Networking EdgeConnect licenses

7.2.2. Create a vWAN hub

To create a vWAN hub from the Azure Portal:

1. In the Azure Portal, click **Virtual WANs**, and then click **+ Create**.
2. On the Create WAN page, enter the following information on the Basics page:
 - a. **Subscription:** Select your subscription.
 - b. **Resource group:** Create new or use existing.
 - c. **Resource group location:** Choose a location. You can select any region.
 - d. **Name:** Give the vWAN a name.
Type: Select **Standard**. (Basic is not sufficient for BGP and hub routing.)
 - e. Click **Review + create**.
3. After creating the vWAN, open it and then click **Hubs**.
4. Click **+ New Hub**, and then enter the following information:
 - a. **Region:** Select the same region as the EC-Vs.

- b. **Name:** Enter a name for the hub.
 - c. **Hub private address space:** Specify a hub private address space. The minimum address space is /24. This range should not overlap with your VNets or on-premises networks.
 - d. **Virtual hub capacity:** Select your required Virtual Hub capacity.
 - e. **Hub routing preference:** Select **AS Path**. When you select AS Path, the vWAN hub prefers routes with the shortest BGP-AS Path length irrespective of the source of the route advertisements.
5. Click **Review + create**, and then click **Create**.

Note

A site-to-site gateway was not created when creating the vWAN hub. This is because the EC-V establishes BGP directly with the vWAN hub's route service. This avoids the need for a VPN gateway within the vWAN hub.

Wait for the new hub to finish provisioning. When it is ready, the Routing status appears as *Provisioned* on the hub's Overview page.

7.2.3. Create a custom role and assign a user-assigned managed identity

Deploying EdgeConnect SD-WAN instances within a Virtual WAN Hub requires the following permissions:

- "Microsoft.Resources/subscriptions/providers/read",
- "Microsoft.Resources/subscriptions/resourcegroups/resources/read",
- "Microsoft.Resources/subscriptions/resourceGroups/read",
- "Microsoft.Network/virtualHubs/read",
- "Microsoft.Network/networkVirtualAppliances/delete",
- "Microsoft.Network/networkVirtualAppliances/read",
- "Microsoft.Network/networkVirtualAppliances/write"

Creating a custom role in Azure with the above permissions and assigning it to a user-assigned managed identity, scoped to the resource group where the vWAN hub and EdgeConnect Managed NVA are deployed, ensures a secure and controlled deployment of SD-WAN instances within a vWAN hub. It is important that the user-assigned managed identity has READ permissions over the vWAN hub you are deploying the Managed NVA into. If not, you will see an error message such as:

```
Caller does not have access on one or more referenced resource(s). Details: The client with object id '<>' does not have authorization to perform action 'Microsoft.Network/virtualHubs/read' over scope '/subscriptions/<>/resourceGroups/<>/providers/Microsoft.Network/virtualHubs/<>'
```

Below are the instructions to create a custom role and associate it with a user-assigned managed identity.

7.2.3.1. Create a custom role

To create a custom role on Azure Portal and assign the above permissions to it, follow the instructions at <https://learn.microsoft.com/en-us/azure/role-based-access-control/custom-roles-portal#start-from-scratch>.

7.2.3.2. Create a user-assigned managed identity

To create a user-assigned managed identity, follow the instructions at <https://learn.microsoft.com/en-us/entra/identity/managed-identities-azure-resources/manage-user-assigned-managed-identities-azure-portal>.

7.2.3.3. Associate the custom role with the user-assigned managed identity

Associate the custom role with the user-assigned managed identity, as shown in Figure 66.

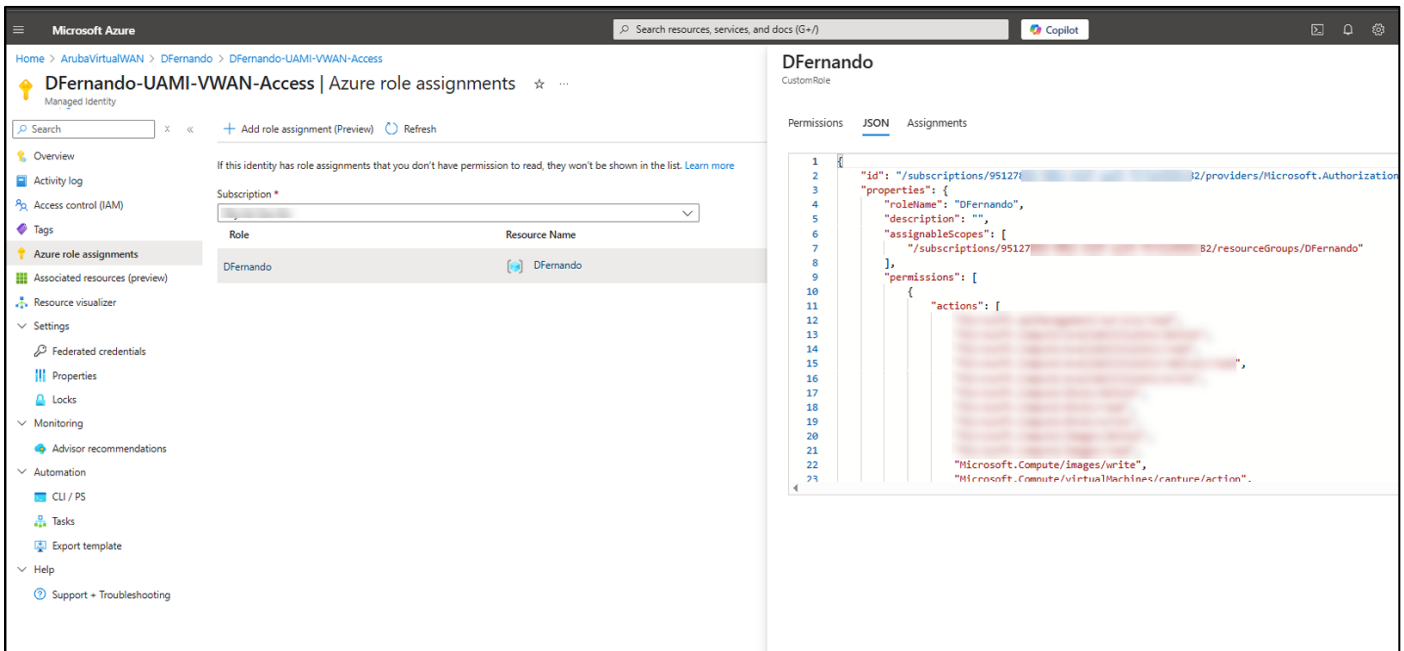


Figure 66. Associating the custom role with the user-assigned managed identity.

7.2.4. Deploy EC-V gateways within the vWAN hub

To deploy EC-V gateways within a vWAN hub:

1. Navigate to your vWAN hub, and then select **Network Virtual Appliance** (under Third party providers).
2. Click **Create network virtual appliance**.
3. Click **Aruba EdgeConnect**.
4. Click **Create**.
5. Select the Azure subscription you want to use, and then click **Create**.
6. On the Create Aruba EdgeConnect Enterprise in Azure Virtual WAN page, enter the following details on the Basics page:
 - a. **Subscription:** Select your subscription.
 - b. **Resource group:** Create new or use existing.
 - c. **Region:** Select the Region where you created the vWAN hub.
 - d. **Application Name:** Enter an application name. The application name can only contain letters and numbers and be between 3 and 32 characters.

ⓘ Important

Do NOT modify the content of the Managed Resource Group field.

- e. Click **Next**.

The screenshot shows the Microsoft Azure portal interface for creating an Aruba EdgeConnect Enterprise in Azure Virtual WAN. The breadcrumb trail is 'Home > Aruba EdgeConnect Enterprise in Azure Virtual WAN >'. The main heading is 'Create Aruba EdgeConnect Enterprise in Azure Virtual WAN'. Below the heading are tabs for 'Basics', 'Aruba EdgeConnect SD-WAN NVA', 'JIT Configuration', and 'Review + create'. The 'Basics' tab is active. Under 'Project details', there is a description: 'Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.' The 'Subscription' dropdown is set to 'Pay-As-You-Go' and the 'Resource group' dropdown is set to 'DFernando'. Under 'Instance details', the 'Region' dropdown is set to 'West Central US'. Under 'Managed Application Details', there is a description: 'Provide a name for your managed application, and its managed resource group. Your application's managed resource group holds all the resources that are required by the managed application which the consumer has limited access to.' The 'Application Name' field contains 'WestCentralUSECV' and the 'Managed Resource Group' field contains 'mrg-aruba_edgeconnect_enterprise_-20251103145708'. At the bottom, there are three buttons: 'Previous', 'Next', and 'Review + create'.

Figure 67. Deploying the Managed NVA Basics tab.

7. Enter the following details on the Aruba EdgeConnect SD-WAN NVA tab:
 - a. **Deployment Name:** This must be a unique name. This is the name that appears on the Network Virtual Appliances page on the Azure Portal. The name can contain letters, numbers, and dashes, and must be between 3 and 32 characters.
 - b. Click **+Add**, and then select the user-assigned managed identity you created earlier.
 - c. Select the Virtual WAN Hub in which to deploy the EC-V gateways.
 - I. If no vWAN hubs are shown, the region you selected does not have any vWAN hubs created yet.
 - II. After you select the vWAN hub, two separate tool tips appear at the bottom of the page. The first tool tip displays the Azure Virtual Hub's BGP endpoint IP addresses and the ASN you must use when establishing BGP from the EC-V to the Azure Virtual Hub. Write down these IP addresses and the ASN. You need this information in Section [7.2.11](#) when creating the static routes.
 - III. The second tool tip displays that you must create two static routes for establishing BGP sessions. This is described in in Section [7.2.11](#).
 - d. **Scale Unit:** The scale unit determines the resource allocation on the Managed NVA. The higher the scale unit, the greater the amount of traffic that can be handled.
 - e. **EdgeConnect SD-WAN ASN:** Enter an ASN between 64512 and 65534, excluding the following ASNs:

- I. ASNs reserved by Azure:
 - 1. Public ASNs: 8074, 8075, 12076
 - 2. Private ASNs: 65515, 65517, 65518, 65519, 65520
- II. ASNs reserved by IANA:
 - 1. 23456, 64496-64511, 65535-65551

Note

vWAN hub only supports 16-bit ASNs. As a result, do not enter a 32-bit ASN.

- f. **EdgeConnect SD-WAN Account Name:** This can be found on the Orchestrator under Cloud Portal.
- g. **EdgeConnect SD-WAN Account Key:** This can be found on the Orchestrator under Cloud Portal.
- h. **EdgeConnect SD-WAN Tag:** (*Optional*) Enter the EdgeConnect tag.
- i. **EdgeConnect SD-WAN single-line SSH key:** Do NOT enter a multi-line SSH key. Multi-line SSH keys typically start with the text “-----BEGIN SSH2 PUBLIC KEY”.
- j. Click **Review + create**.

Microsoft Azure Search resources, services, and docs (G+/) Copilot

Home > Aruba EdgeConnect Enterprise in Azure Virtual WAN >

Create Aruba EdgeConnect Enterprise in Azure Virtual WAN

Basics **Aruba EdgeConnect SD-WAN NVA** JIT Configuration Review + create

Deployment Name * ⓘ wvus-ecv-nva ✓

User assigned managed identity

Add user assigned identities to grant the EdgeConnect SD-WAN managed application to access your Virtual WAN Hub.

+ Add Remove

| Name | ↑↓ Resource group | ↑↓ Subscription |
|---|-------------------|-----------------------------------|
| <input type="checkbox"/> DFernando-UAMI-VWAN-Access | DFernando | 9512788d-90b2-42d7-aa25-f172e4... |

Virtual WAN Hub * ⓘ WestCentralUS-Hub

EdgeConnect SD-WAN Version * ⓘ 9.3.70

Scale Unit * ⓘ 4 Scale Units (4 vCPUs / 14GB RAM)

EdgeConnect SD-WAN ASN * ⓘ 65199 ✓

EdgeConnect SD-WAN Account Name * ⓘ

EdgeConnect SD-WAN Account Key * ⓘ

EdgeConnect SD-WAN Tag ⓘ

EdgeConnect SD-WAN single-line SSH key * ⓘ ssh-ed25519 AAAAB3Nkuyioewytopolkjlqwueuy3h1uahHjkYfqEwcgpo...

i Virtual WAN Hub BGP endpoints: ["192.168.253.68","192.168.253.69"] Virtual WAN Hub BGP ASN: 65515

i Configure two /32 static routes on the EdgeConnect SD-WAN VM and establish two BGP sessions using the information above.

Previous Next **Review + create**

Figure 68. Creating an EdgeConnect Managed NVA on Azure Portal.

8. (Optional) Set the JIT Configuration tab.
9. Click **Next**.
10. Review the settings you entered.
11. Select the **I agree to the terms and conditions above** check box.
12. Click **Create**.

EC-V deployment begins.

7.2.5. Correlate EC-Vs

Within approximately five minutes of the deployment, two EC-Vs appear on your Orchestrator’s Discovered Appliances page (Configuration > Overlays & Security > Discovery > Discovered Appliances).

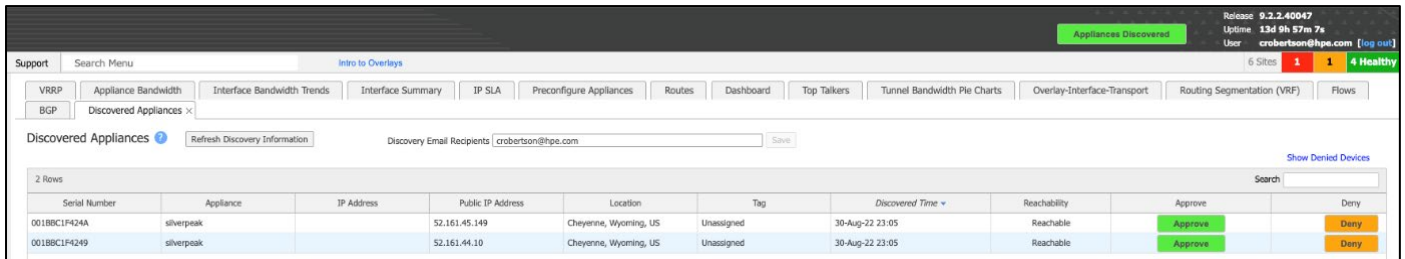


Figure 69. Verifying that EC-Vs appear in the Orchestrator UI.

Note

If you deploy the EC-Vs in a region where Azure supports availability zones, each EC-V is deployed in a unique availability zone. If you deploy the EC-Vs in a region where Azure does not support availability zones, each EC-V is deployed in an availability set.

Azure numbers each EC-V according to the order it deploys them. To find this order, perform the tasks below. The objective is to ensure that order is preserved when you add the EC-Vs to the SD-WAN fabric. You want to approve the first EC-V (that Azure deployed) first, followed by the second EC-V:

1. Go to the vWAN hub in which you deployed the EC-Vs.
2. Under Third party providers, click **Network Virtual Appliance**.

Note

If you encounter a message on the Network Virtual Appliance page that says You haven't created a network virtual appliance for this virtual hub yet, the appliances are likely still initializing. Please wait for 10–15 minutes for the appliances to fully provision.

3. In the Instances info column, click **Click here** for your deployment.

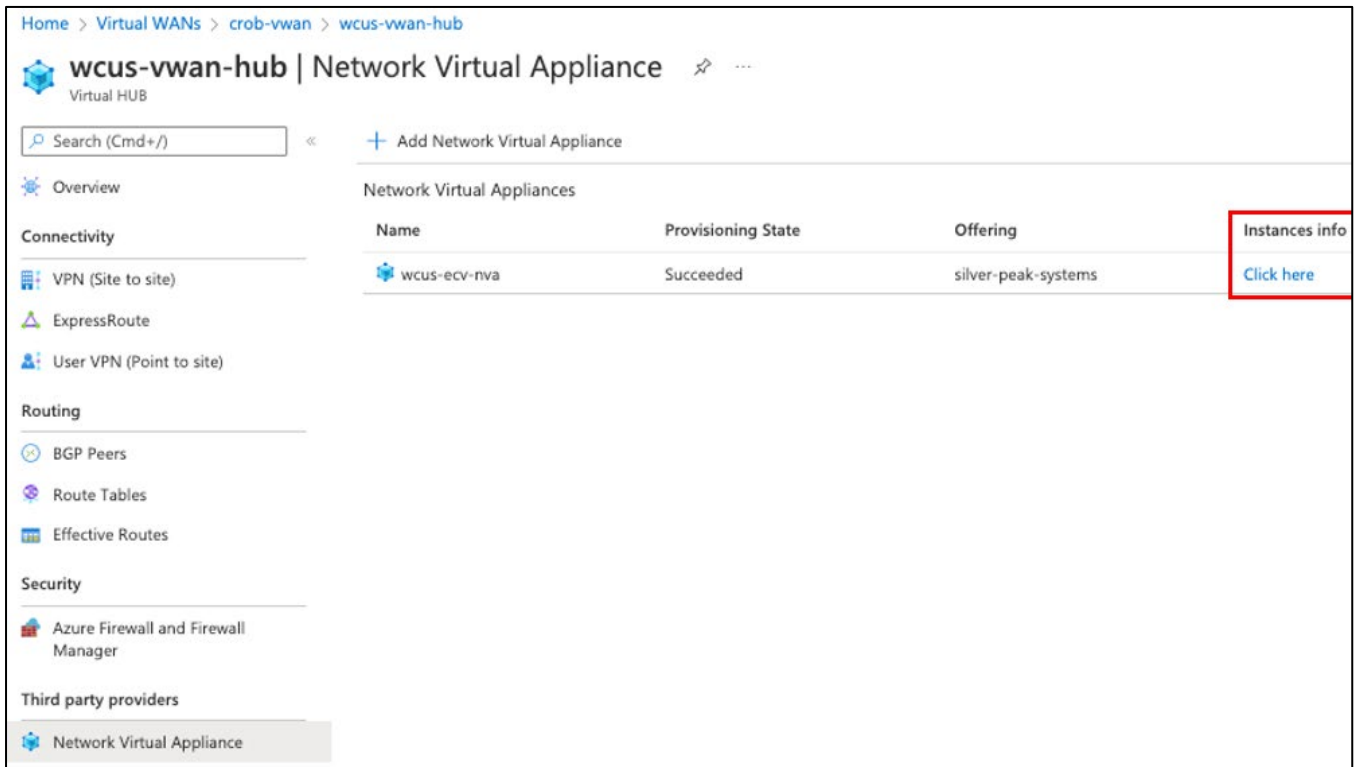


Figure 70. Finding the IP addresses assigned to each instance by clicking the Click here link in the Instances info column.

A window opens that shows the IP addresses of each EC-V in the deployment. Under the NVA instance name, *0* denotes the first EC-V and *1* denotes the second EC-V, as shown in Figure 71.

Notes

- Azure only assigns one public IP address per EC-V. This public IP address is assigned to the WAN0 interface.
- Each subnet is created as a /28 by Azure. This is not modifiable. The next hop IP address of each interface is always the first useable IP of the assigned subnet.

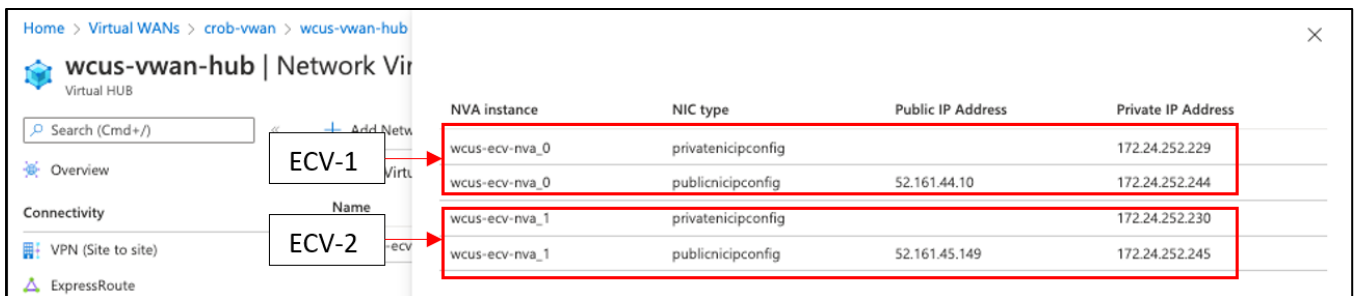


Figure 71. Identifying EC-Vs and their public and private IP addresses.

EC-V-1 details:

- In EC-V-1, the public IP address of the WAN0 interface is 52.161.44.10.
- In EC-V-1, the private IP address of the WAN0 interface is 172.24.252.244.
- The WAN0 next hop IP address is 172.24.252.241.
- In EC-V-1, the private IP address of the LAN0 interface is 172.24.252.229.
- The LAN0 next hop IP address is 172.24.252.225.

EC-V-2 details:

- In EC-V-2, the public IP address of the WAN0 interface is 52.161.45.149.
- In EC-V-2, the private IP address of the WAN0 interface is 172.24.252.245.

- The WAN0 next hop IP address is 172.24.252.241.
- In EC-V-2, the private IP address of the LAN0 interface is 172.24.252.230.
- The LAN0 next hop IP address is 172.24.252.225.

Note

Although Azure assigned the above IP addresses to the EC-Vs at the hypervisor level, these IP addresses are not yet assigned at ECOS level, meaning that the EC-V is still in Server Mode with no IP addresses.

Upcoming sections of this guide detail how to change the deployment mode from Server Mode to In-line Router Mode and assign the WAN0 and LAN0 IP addresses on the EC-V's Deployment page.

Important

In ECOS 9.3.7.0 and later, when you deploy EC-Vs within a vWAN Hub (from Orchestrator or manually from Azure Portal, as explained in this section) the MAC addresses are automatically assigned to WAN0 and LAN0 interfaces when the EC-V is provisioned. This means you no longer need to assign the MAC addresses on the Interfaces page after the EC-V is deployed. Although the MAC addresses are automatically assigned, the EC-V still remains in Server Mode when you deploy it for the first time. This means that you should not attempt to reboot or upgrade the EC-V after it is deployed for the first time until the deployment mode is set to In-line Router Mode and you assign WAN0 and LAN0 IP addresses on the Deployment page.

If you reboot or upgrade the EC-V before the deployment mode is set to In-line Router Mode and assign WAN0 and LAN0 IP addresses, you will lose connectivity to the EC-V and you will not be able to recover the EC-V. Your only option is to redeploy new EC-Vs.

7.2.6. Add EC-Vs to the SD-WAN fabric

Now that you are aware of IP addresses for each EC-V, find the public IP address of EC-V-1. In this example, the IP address is 52.161.44.10.

1. To start adding the EC-V to the SD-WAN fabric, click **Approve**.




Figure 72. Adding the first EC-V to the Orchestrator.

2. While adding the EC-Vs to the SD-WAN fabric, you are prompted to upgrade the EC-V. **DO NOT UPGRADE THE EC-V** at this time. You can upgrade the EC-V after WAN0 and LAN0 IP addresses are assigned to the EC-V.

To start the appliance wizard, click **Skip**.

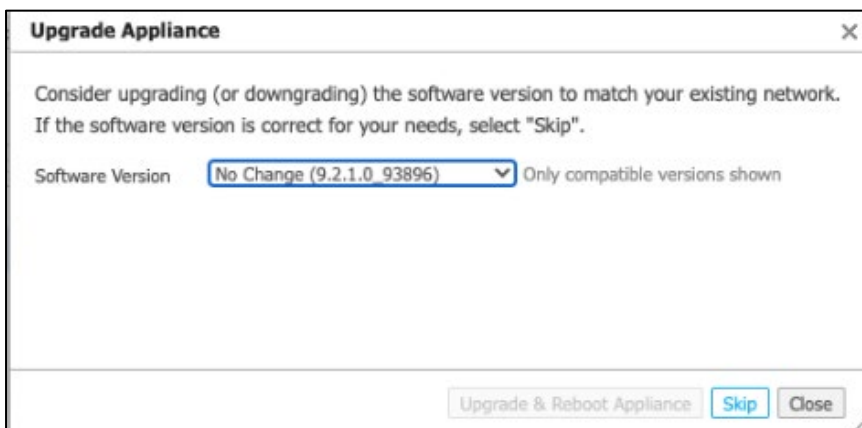


Figure 73. Using the optional Upgrade Appliance dialog box in the Orchestrator UI.

- On the Appliance Setup page, configure the hostname, group, admin password, and location information, and then select the **Hub Site** check box if the NVA will act as a hub for SD-WAN spokes.

Figure 74. Configuring the first page in the Appliance Wizard of the Orchestrator UI.

Note

Ensure that you enter the same site name on both EC-Vs to prevent them from forming tunnels with each other.

- Click **Next**, but do not make any changes on this screen. You will do this **AFTER** adding the EC-V to SD-WAN fabric. Do not change the deployment mode from Server Mode to In-line Router Mode yet.

Figure 75. Configuring the Deployment Profile of the Appliance Wizard in the Orchestrator UI.

- Continue through the appliance wizard. When you reach step 5, clear all templates and do not apply overlays. You will apply overlays after assigning the IP addresses.

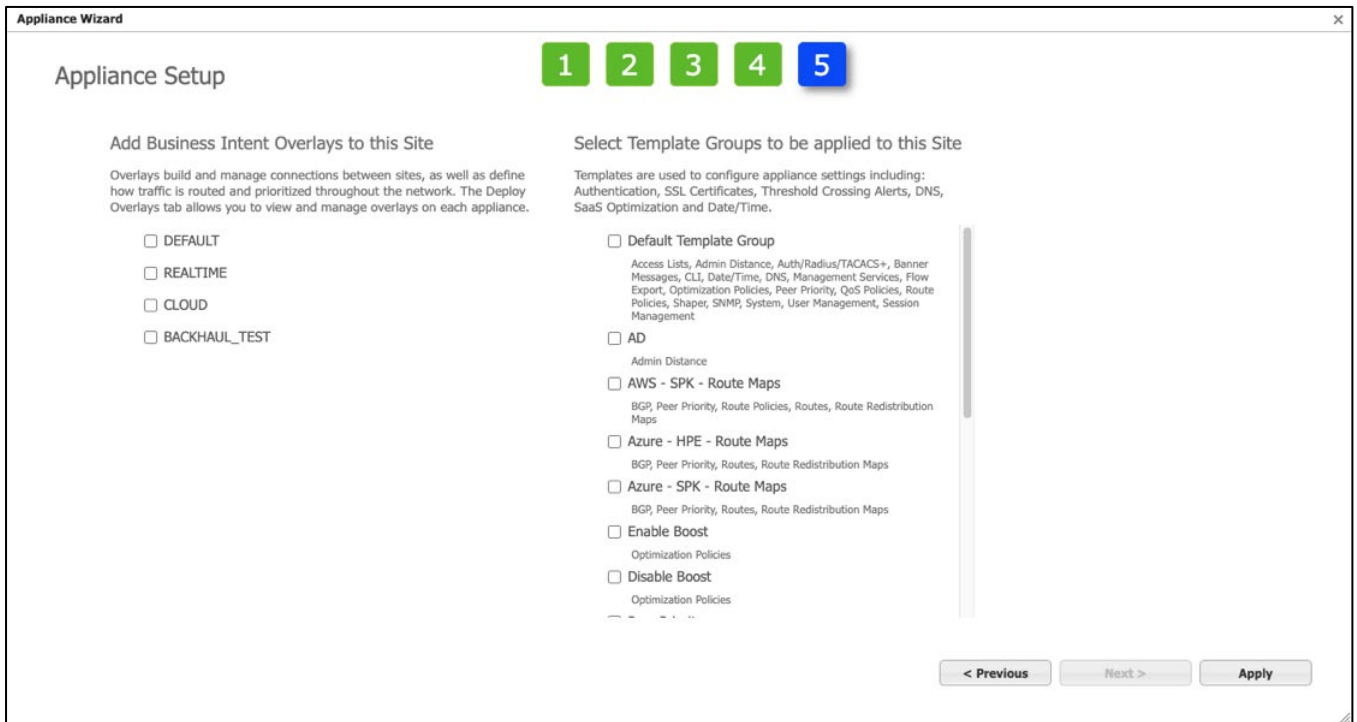


Figure 76. Configuring BIOs and template groups in the Appliance Wizard of the Orchestrator UI.

- Repeat the steps above for the other EC-V. Ensure that you enter the same site name on both EC-Vs to prevent them from forming tunnels with each other.

7.2.7. Assign IP addresses

Perform the following tasks to assign IP addresses on the EC-V:

- On the Orchestrator appliance tree, select **EC-V-1**, and then navigate to **Configuration > Networking > Deployment**.
- Click the **edit** (pencil) icon.
- On the Deployment dialog box, click the **Router** tab.
- Set the FW Mode to **Stateful+SNAT**.
- Assign the WAN0 IP/Mask and LAN0 IP/Mask based on the information you collected in [7.2.5 Correlate EC-Vs](#).
- Enter the WAN0 Next Hop IP address. The WAN0 subnet's first useable IP address is the next hop IP address of the WAN0 subnet.
- Enable **Stateful+SNAT** on the WAN0 interface.
- Set the **WAN0** label.

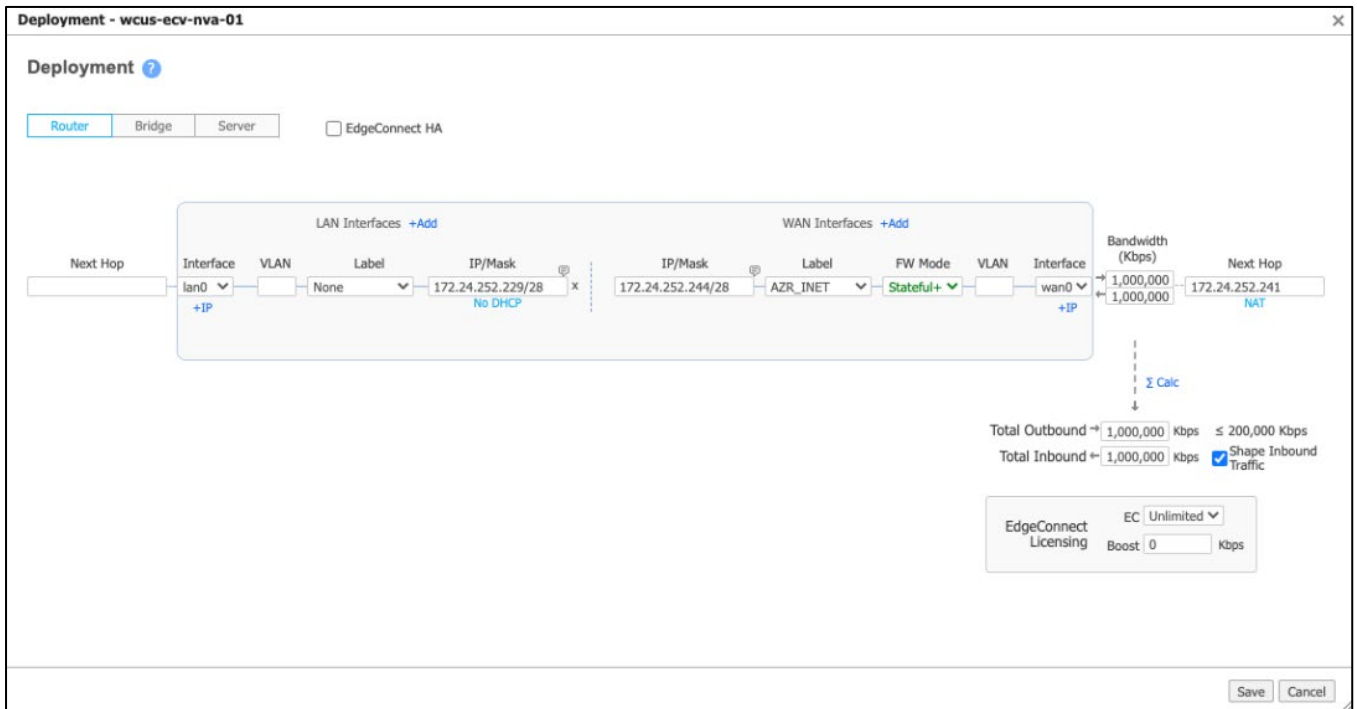


Figure 77. Changing mode from Server Mode to In-line Router and assigning IP addresses to interfaces.

10. (Optional) Enter Boost bandwidth.
11. Click **Save**.
12. On the Reboot Appliance dialog box, click **Apply and Reboot**. The appliance begins rebooting.
13. Repeat the above steps on EC-V-2. Please wait approximately 15 minutes for the appliance to be fully operational again.

7.2.8. Shut down MGMT0

Perform the following tasks to disable the unassigned mgmt0 interface to clear the alarm it is generating:

1. On the Orchestrator appliance tree, select **EC-V-1**, and then navigate to **Configuration > Networking > Interfaces**.
2. Click the **edit** (pencil) icon.




Figure 78. Clicking the edit (pencil) icon on the Interfaces tab of the Orchestrator UI.

3. On the Interfaces dialog box, click **All Interfaces**.

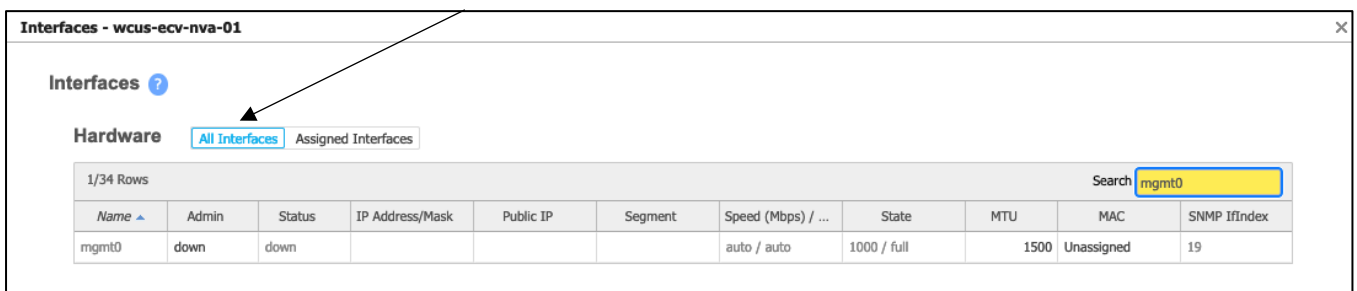


Figure 79. Clicking the All Interfaces tab on the Interfaces page of the Orchestrator UI.

4. Search for the mgmt0 interface.
5. Click the drop-down menu in the Admin column, and then select **down**.

6. Click **Apply**.
7. Repeat these steps for EC-V-2.

7.2.9. Create two static routes for enabling BGP

The vWAN hub's route service contains two BGP endpoints. This means you must create two static routes on each EC-V gateway to establish BGP with the vWAN hub's BGP endpoints. To create the static routes on EC-V:

1. Refer to Figure 68 to find the IP addresses of the vWAN hub's BGP endpoints.
2. On the Orchestrator appliance tree, select **EC-V-1**, and then navigate to **Configuration > Networking > Routing Routes**.
3. Click **Local/Static**.
4. Click the **edit** (pencil) icon on the Default segment.
5. On the Routes page for the Default segment, enter the following settings:
 - a. **Automatically advertise local LAN subnets:** Clear the check box.
 - b. **Automatically advertise local WAN subnets:** Clear the check box.
 - c. **Metric for automatically added subnets:** Enter 50.
 - d. **Redistribute routes to SD-WAN Fabric:** Select an available route map.
 - e. **Filter routes from SD-WAN fabric with matching local ASN:** Select the check box.
 - f. **Include BGP Local ASN to routes sent to SD-WAN Fabric:** Select the check box.
 - g. **Tag BGP communities to routes:** Clear the check box.
 - h. Click **Apply**.
6. Click **Add Route**, and then enter the following settings:
 - a. **Subnet/Mask:** Enter the first vWAN hub BGP endpoint (/32) IP address.
 - b. **Next Hop:** Enter EC-V's LAN0 interface's next hop IP address. This is the first IP address of the LAN0 subnet.

Note

If you are establishing BGP sessions from other interfaces, such as the LAN1 interface, make sure to enter the first IP address of that interface's subnet.

- c. **Interface:** LAN0
- d. **Zone:** *(Optional)* Default
- e. **Metric:** 50 (default)
- f. **Tag:** Select **ANY**.
- g. **Comments:** Optional

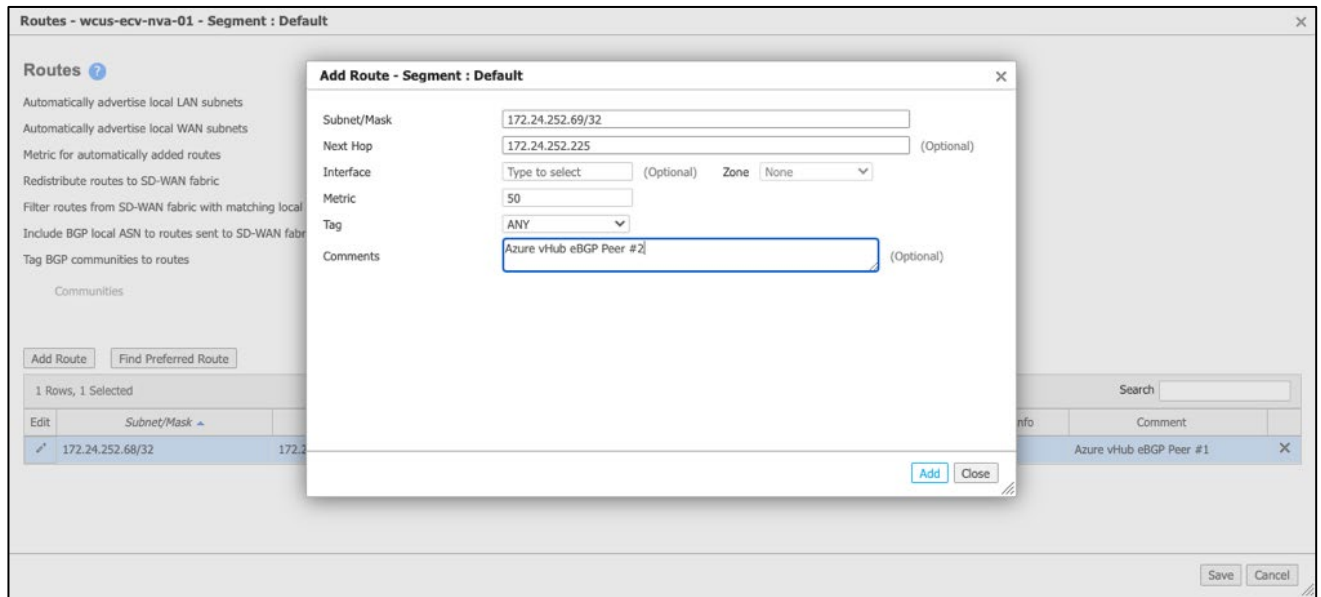


Figure 80. Adding a static route.

7. Click **Add** to add the first static route.
8. Repeat steps 6 and 7 to create a static route for the second vWAN hub BGP endpoint.

| Edit | Appliance | Segment | Subnet/Mask | Next Hop | Interface | Zone | State | Metric | Advertise To Peers... | Type | Additional Info | Comment |
|------|-----------------|---------|------------------|----------------|-----------|----------------|-------|-----------|-----------------------|------|-----------------|---------------------|
| ✓ | wcus-ecv-nva-01 | Default | 172.24.252.69/32 | 172.24.252.225 | lan0 | None (Default) | UP | 50 (AD-1) | N/A | | Tag ANY | Azure vHub eBGP ... |
| ✓ | wcus-ecv-nva-01 | Default | 172.24.252.68/32 | 172.24.252.225 | lan0 | None (Default) | UP | 50 (AD-1) | N/A | | Tag ANY | Azure vHub eBGP ... |

Figure 81. After adding both static routes.

9. Repeat steps 2–7 on the second EC-V to create two static routes on that EC-V.

7.2.10. Create static routes for Azure health probes

When you deploy the Managed NVA, Azure automatically deploys an ILB as part of the hub infrastructure. This ILB monitors the health of each EC-V by sending health probes to the LAN0 interface. To respond to the health probes, you must create a static route on each EC-V:

1. On the Orchestrator appliance tree, select **EC-V-1**, and then navigate to **Configuration > Networking > Routing Routes**.
2. Click **Local/Static**.
3. Click the **edit** (pencil) icon.
4. On the Routes page, enter the following settings:
 - a. **Subnet/Mask:** Enter 168.63.129.16/32.
 - b. **Next Hop:** Enter the LAN0 next hop IP address.
 - c. **Interface:** LAN0
 - d. **Zone:** (Optional) Default
 - e. **Metric:** 50 (default)
 - f. **Tag:** Select **ANY**.
 - g. **Comments:** Optional
5. Repeat these steps for EC-V-2.

| Routes | | | | | | | | | | | | |
|--------|-----------------|---------|------------------|----------------|-----------|------|-------|-----------|--------------------|------|-----------------|-------------------------|
| 3 Rows | | | | | | | | | | | | |
| Edit | Appliance | Segment | Subnet/Mask | Next Hop | Interface | Zone | State | Metric | Advertise To Peers | Type | Additional Info | Comment |
| ✓ | wcus-ecv-nva-01 | Default | 172.24.252.69/32 | 172.24.252.225 | | None | UP | 50 (AD:1) | N/A | | Tag ANY | Azure vHub eBGP Peer #2 |
| ✓ | wcus-ecv-nva-01 | Default | 172.24.252.68/32 | 172.24.252.225 | | None | UP | 50 (AD:1) | N/A | | Tag ANY | Azure vHub eBGP Peer #1 |
| ✓ | wcus-ecv-nva-01 | Default | 168.63.129.16/32 | 172.24.252.225 | lan0 | None | UP | 50 (AD:1) | N/A | | Tag ANY | Azure Health Probes |

Figure 82. After creating the static route to respond to the health probes.

7.2.11. Apply templates and overlays

Now that the appliances are operational, templates and overlays can be applied to these appliances.

Note

If you do not currently have BGP route maps configured as part of your templates, you can configure them now.

1. On the Orchestrator appliance tree, select **EC-V-1**, and then navigate to **Configuration > Templates & Policies > Apply Template Groups**.
2. Select the template groups you want, and then click **Apply**.

Figure 83. Applying template groups on the Apply Template Groups tab of the Orchestrator UI.

3. On the Orchestrator appliance tree, select **EC-V-1**, and then navigate to **Configuration > Overlays & Security > Apply Overlays**.
4. Select the overlays you want, and then click **Apply**.

7.2.12. Add vWAN hub as a BGP peer on the EC-V gateways

Now that the static routes are created and BGP peers are created on the vWAN hub, perform the following steps to enable BGP on the EC-V gateways:

1. On the Orchestrator appliance tree, select **EC-V-1**, and then navigate to **Configuration > Networking > Routing BGP**.
2. Click the **Enable BGP** toggle.
3. Configure the following BGP parameters:
 - a. **Autonomous system number:** Enter the EC-V's ASN that you entered in [6.3.4. Add EC-V gateways as BGP peers on the vWAN hub](#).
 - b. **Route Target:** Leave empty.
 - c. **Router ID:** The Router ID is an IPv4 address by which the remote peer can identify this EC-V gateway for purposes of BGP.
 - d. **Graceful restart:** Select the check box.
 - e. **AS path propagate:** Select the check box.

- f. **Log BGP update messages:** Select the check box.
- g. **Max route updates per peer:** 10
- h. **Detection interval:** 15 minutes
- i. Click **Add Peer**, and then enter the following settings:
 - I. **Peer IP:** Primary Azure BGP Peer (same address as the /32 Static Route added in the previous step)
 - II. **Peer Adjacency:** Multi-Hop
 - III. **EVPN Peer:** Clear the check box.
 - IV. **Local Interface:** lan0
 - V. **Peer ASN:** 65515
 - VI. **Override ASN:** Select the check box.
 - VII. **Peer Type:** Branch
 - VIII. **Admin Status:** Up
 - IX. **Soft Reconfiguration:** Select the check box.
 - X. **Next-Hop Self:** Select the check box.
 - XI. **Inbound route map:** Select the appropriate inbound route map.
 - XII. **Outbound route map:** Select the appropriate outbound route map.
 - XIII. **BFD:** Clear the check box.
 - XIV. **Keep Alive Timer:** 5
 - XV. **Hold Timer:** 15
 - XVI. **Enable MD5 Password:** Clear the check box.
 - XVII. To create the first BGP session, click **Add**.
4. Repeat the steps above to create the second BGP session.
5. Repeat the same steps on the second EC-V.
6. Verify BGP configuration by checking the BGP neighbor status in the Orchestrator.

7.2.13. Verify connectivity

After enabling BGP on the EC-V gateways, verify that the BGP sessions come up on both sides. On Orchestrator's BGP tab, check the Peer State for the Azure hub neighbors on each EC-V. A state of *Established* indicates that the BGP adjacency is up and routes can be exchanged. Click the peer details icon in Orchestrator to see information such as the number of routes learned/advertised via that peer.

After BGP is established, the EC-V gateways dynamically advertise their routes to the vWAN hub, and the vWAN hub advertises routes present on its route table to the EC-V gateways. Verify that the EC-Vs are receiving Azure routes—you can check the routing table on each EC-V (in Orchestrator's Routes table) to see BGP-learned routes from Azure. Likewise, Azure's route tables (in the default route table of the vWAN hub) now include routes for the networks advertised by the EC-V gateways. This confirms that end-to-end dynamic routing is in place.

You have successfully established BGP between your EC-Vs and vWAN hub.

7.2.14. Where to find configuration items

If you do not have the EdgeConnect ASN that was configured in the NVA deployment steps, you can find it on the Azure Portal by navigating to **Subscriptions > Resources**, filtering the Type column to find **Managed Applications**, and then clicking the managed application.

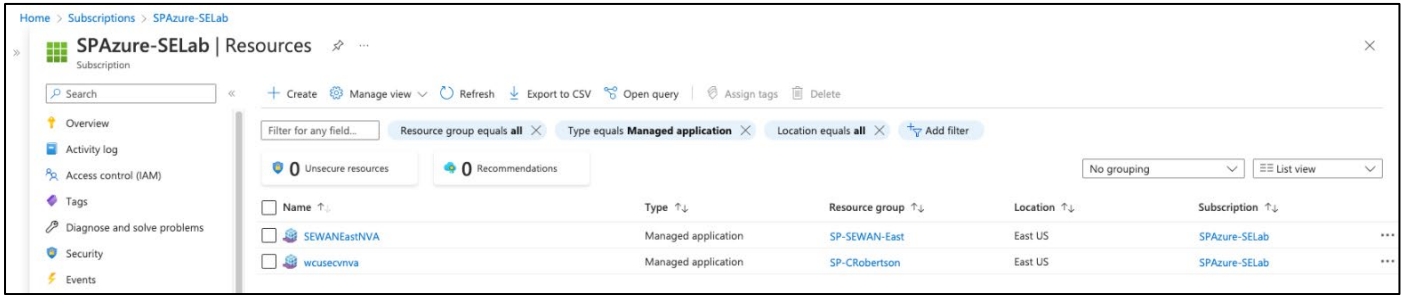


Figure 84. Finding the EdgeConnect ASN on the Azure Portal.

Click **Parameters and Outputs**.



Figure 85. Viewing the Parameters and Outputs sections on the Azure Portal.

If you do not have the Virtual WAN Hub ASN or BGP Peer IPs, you can find them by navigating to **Virtual WANs > Select your Virtual WAN > Hubs > Select your Hub**, and then selecting **BGP Peers** under Routing Select.

The peer IPs and ASN can be found in the upper-right corner of the page.

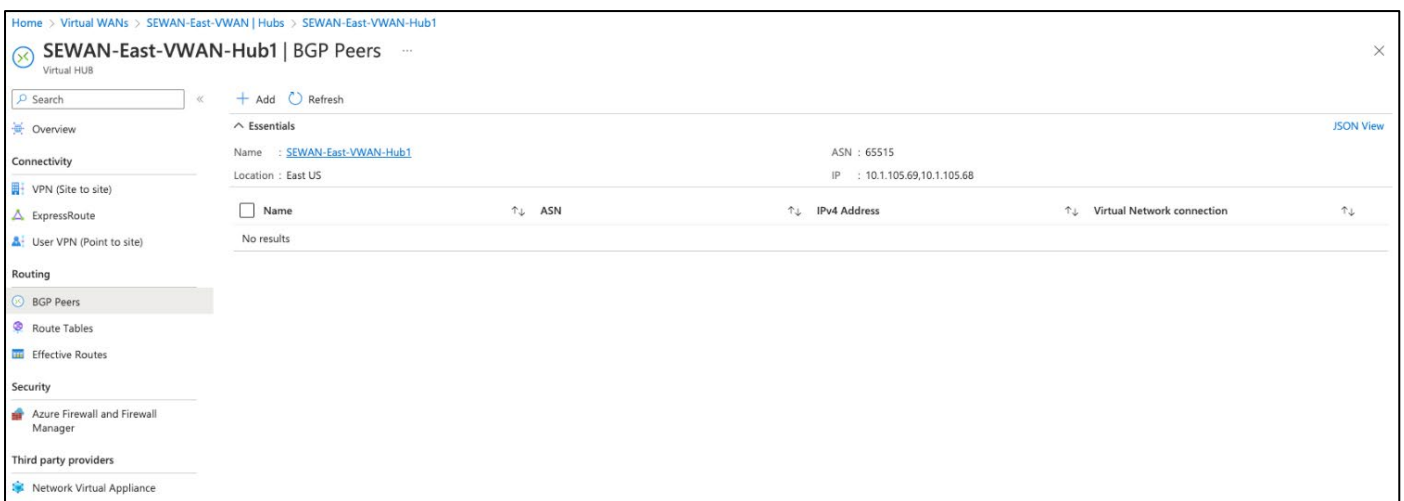


Figure 86. Viewing the BGP Peers page on the Azure Portal.

7.2.15. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric

To avoid the two static routes that were created on each EC-V (to establish BGP sessions) from being advertised to other SD-WAN devices via subnet sharing, see [9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric.](#)

8. Redirect traffic from spoke (workloads) VNets to EC-V via ILB, ARS, or vWAN hub

In ILB and ARS designs, redirecting network traffic from spoke VNets (where user workloads/applications run) to the Azure EC-V gateways (via ILB or ARS), is done using VNet peering sessions.

In transit VNet + vWAN hub and Managed NVA designs, redirecting network traffic from spoke VNets (where user workloads/applications run) to the Azure EC-V gateways (via vWAN hub), is done using virtual network connections.

The following section explains this in more detail:

1. ILB design
 - a. Create a VNet peering to establish connectivity between the spoke VNet (where workloads reside) and the transit VNet (where EC-V and ILB are deployed). This is explained in Section [8.1](#).
 - b. Next, in the spoke VNet, create a route table with a static route that forwards outbound traffic to the ILB's frontend private IP address. This is explained in Section [8.3](#).
 - c. After ILB receives traffic from the spoke VNet, it distributes it to healthy EC-V instances in its backend pool.
2. ARS design
 - a. Create a VNet peering to establish connectivity between the spoke VNet (where workloads reside) and the transit VNet (where EC-V and ARS are deployed). This is explained in Section [8.1](#).
 - b. After you establish BGP between EC-V and ARS, EC-V advertises routes to ARS, which propagates them to spoke VNets. This eliminates the creation of UDRs in spoke VNets.
3. Transit VNet + vWAN hub design
 - a. Create a virtual network connection to establish connectivity between the spoke VNet (where workloads reside) and the vWAN hub. This is explained in Section [8.2](#).
 - b. After you establish BGP between EC-V and vWAN hub, EC-V advertises routes to vWAN hub, which propagates them to spoke VNets. This eliminates the creation of UDRs in spoke VNets.
4. Managed NVA design
 - a. Create a virtual network connection to establish connectivity between the spoke VNet (where workloads reside) and the vWAN hub. This is explained in Section [8.2](#).
 - b. After you establish BGP between EC-V and vWAN hub, EC-V advertises routes to vWAN hub, which propagates them to spoke VNets. This eliminates the creation of UDRs in spoke VNets.

8.1. Create a virtual network peering session (for ILB and ARS designs)

If you are deploying EC-V with ILB or ARS, follow the instructions in this section to create a virtual network peering session.

1. From Azure Portal, navigate to your resource group, and then click the EC-V's VNet.
2. From the left-side Settings menu, click **Peerings**.
3. On the Peerings page, Click **+ Add** to open the Add peering page, and then enter the following settings:
 - a. Remote virtual network summary
 - I. **Peering link name:** Enter a name (example: *Workloads-VNet-to-EC-V-VNet*).
 - II. **Virtual network deployment model:** Select **Resource Manager**.
 - III. **I know my resource ID:** Clear the check box.
 - IV. **Subscription:** Select the subscription that you used to deploy the EC-Vs.
 - V. **Virtual Network:** Select the workload's VNet.
 - b. Remote virtual network peering settings
 - I. **Allow <Workloads VNet> to access <EC-V's VNet>:** Select the check box.
 - II. **Allow <Workloads VNet> to receive forwarded traffic from <EC-V's VNet>:** Select the check box.

- III. **Allow gateway or route server in <Workloads VNet> to forward traffic to <EC-V's VNet>**: Clear the check box.
- IV. **Enable <Workloads VNet> to use <EC-V's VNet>s' remote gateway or route server**: For ILB design, clear the check box. For ARS design, select the check box.
- c. Local virtual network peering summary
 - I. **Peering link name**: EC-V-VNet-to-Workloads-VNet
- d. Local virtual network peering settings
 - I. **Allow <EC-V's VNet> to access <Workloads VNet>**: Select the check box.
 - II. **Allow <EC-V's VNet> to receive forwarded traffic from <Workloads VNet>**: Select the check box.
 - III. **Allow gateway or route server in <EC-V's VNet> to forward traffic to <Workloads VNet>**: Select the check box.
 - IV. **Enable <EC-V's VNet> to use <Workloads VNet>s' remote gateway or route server**: Clear the check box.
4. At the bottom of the page, click **Add**.

This establishes the VNet peering between the transit VNet and spoke VNet.

Proceed to [8.3 Create static routes on Azure Portal to forward outbound traffic from a workloads VNet to ILB](#). This step forwards outbound traffic from Azure to the ILB.

8.2. Create a virtual network connection (for vWAN hub-related designs)

If you are deploying EC-V in a transit VNet with vWAN hub or within a vWAN hub (Managed NVA), follow the instructions in this section to create a virtual network connection:

1. In the Azure Portal, go to your vWAN resource, and then click **Virtual network connections**.
2. On the Virtual network connections page, click **+ Add connection**.
3. On the Add connection page, configure the connection settings. For information about routing settings, see [Microsoft's About virtual hub routing page](#).
 - a. **Connection name**: Name your connection.
 - b. **Hubs**: Select the hub you want to associate with this connection.
 - c. **Subscription**: Verify the subscription.
 - d. **Resource group**: Select the resource group that contains the virtual network to which you want to connect.
 - e. **Virtual network**: Select the virtual network you want to connect to this hub. The virtual network you select cannot have an already existing virtual network gateway.
 - f. **Propagate to none**: By default, this is set to *No*. Selecting *Yes* makes the configuration options for Propagate to Route Tables and Propagate to labels unavailable for configuration.
 - g. **Associate Route Table**: From the drop-down menu, select a route table that you want to associate. Associating a connection to a route table allows the traffic from that connection to be sent to the destination indicated as routes in the route table.
 - h. **Propagate to labels**: Select a label. Labels are a logical group of route tables. Labels provide a mechanism to logically group route tables. This is especially helpful during propagation of routes from connections to multiple route tables. If no label is specified in the list of labels that a virtual network connection is propagating to, then the virtual network connection automatically propagates to the *Default* label.
 - i. **Static routes**: Static routes are not necessary because you established BGP between EC-V and vWAN hub.
 - j. **Bypass Next Hop IP for workloads within this VNet**: Select **No**.
 - k. **Propagate static route**: Select **Yes**.

Add connection ✕

Connection name *

Hubs * ⓘ

Subscription *

Resource group *

Virtual network *

Routing configuration ⓘ

Propagate to none ⓘ
 Yes No

Associate Route Table

Propagate to Route Tables

Propagate to labels ⓘ

Static routes ⓘ

| Route name | Destination prefix | Next hop IP |
|--|--|--|
| <input style="width: 95%;" type="text"/> | <input style="width: 95%;" type="text"/> | <input style="width: 95%;" type="text"/> |

Bypass Next Hop IP for workloads within this VNet ⓘ
 Yes No

Propagate static route ⓘ
 Yes No

Figure 87. Creating a virtual network connection.

4. Click **Create**.

8.3. Create static routes on Azure Portal to forward outbound traffic from a workloads VNet to ILB

Note

This section is applicable only to ILB design.

After you create a VNet peering session between transit VNet and spoke VNet, you need to create static routes on the spoke VNet route table to forward traffic to the ILB. Static routes in Azure are known as user-defined routes (UDRs). Azure automatically routes traffic between Azure subnets, virtual networks, and on-premises networks. To modify

these default routes of a VNet, you should create UDRs in a route table, which overrides the Azure default system routes. In Azure, you create a route table and then associate the route table to the subnets of a VNet. When you do this, the route table's routes are combined with the subnet's default routes. If there are conflicting route assignments, UDRs override the default routes. When you create a UDR in a route table, you can specify the following next hop types:

- Virtual appliance
- Virtual network gateway
- None
- Virtual network
- Internet

This deployment will select **Virtual appliance** as the next hop type and enter the ILB's frontend private IP address. This redirects outbound traffic for a specific destination route to the ILB. To create a UDR in this deployment, create an Azure route table, associate the workloads subnet to the route table, and then create a static route (UDR) to forward traffic to the ILB.

8.3.1. Create an Azure route table

The following steps create a route table for your Azure workload.

1. From the Azure Portal, navigate to your resource group, and then click **+Add**.
2. Use the search field to select **Route Table** from the drop-down menu.
3. On the Route table dialog box, click **Create**.
4. On the Create Route Table page, enter the following settings:
 - a. **Name:** Enter a descriptive name for the route.
 - b. **Subscription:** Your Azure subscription.
 - c. **Resource Group:** Select your resource group.
 - d. **Location:** Select the region where your workloads VNet reside.
 - e. **Virtual Network Gateway Route Propagation:** Select **Disabled**.
5. At the bottom of the Create Route Table dialog box, click **Create**.

8.3.2. Associate the workloads subnet to the route table

Route tables are not associated to VNets. If you have multiple subnets that contain workloads that need to forward traffic to the ILB in your workloads VNet, you must associate the route table you created above to each subnet on the VNet.

The following steps associate a subnet where your workloads are deployed to the route table:

1. Open the Route Table page for the previously created route table (accessible from the Resource Group page).
2. From the left-side Settings menu, click **Subnets**.
3. On the Subnets page, click **+ Associate**.
4. On the Associate subnet panel, enter the following settings:
 - a. **Virtual network deployment model:** Select the deployment's virtual network.
 - b. **Subnet:** Select the virtual network that accesses the workload.
5. Click **OK**.

The specified subnet appears on the Subnets page, indicating that it is associated to the route table.

6. Repeat these steps for each subnet that needs to forward traffic to the ILB.

8.3.3. Create a static route to forward traffic to the ILB

After associating the workload VNet's subnets to the route table, the following steps create a static route on the Azure Portal to forward outbound traffic to the ILB:

1. Open the Route Table page for the previously created route table (accessible from the Resource Group page).

2. From the left-side Settings menu, click **Routes**.
3. On the Routes page, click **+ Add**.
4. On the Add route page, enter the following settings:
 - a. **Route name:** Enter a name for the route.
 - b. **Destination type:** Select **IP Addresses**.
 - c. **Destination IP addresses/CIDR ranges:** The destination IP address range (CIDR notation) to which this applies. Packets match this route when their destination IP address falls in this range. Azure selects a route for traffic based on longest prefix match.
 - d. **Next hop type:** Select **Virtual appliance**.
 - e. **Next hop address:** Enter the ILB's frontend private IP address.
5. Click **OK** to create the route and close the page.
The new route subnet appears on the Routes page.

8.4. Create static routes on EC-V to advertise Azure subnets to remote EdgeConnect devices

Note

This section applies only to ILB design. This topic creates static routes on EC-V to advertise Azure subnets to remote EdgeConnect devices.

Repeat these steps to create static routes on each Azure EC-V in your transit VNet:

1. In Orchestrator, select **ECV-A** in the appliance tree, and then navigate to **Configuration > Networking > Routing > Routes**.
2. Click an **edit** (pencil) icon to open the Routes dialog box, and then enter the following settings:
 - a. **Automatically advertise local LAN subnets:** Clear the check box.
 - b. **Automatically advertise local WAN subnets:** Clear the check box.
 - c. **Metric for automatically added subnets:** Enter **50**.
 - d. **Redistribute routes to SD-WAN Fabric:** Select an available route map.
 - e. **Filter Routes From SD-WAN Fabric With Matching Local ASN:** Select the check box.
 - f. **Include BGP Local ASN to routes sent to SD-WAN Fabric:** Select the check box.
 - g. **Tag BGP communities to routes:** Clear the check box.
 - h. **Use SD-WAN fabric learned routes:** Select the check box.
 - i. **Enable Equal Cost Multi Path (ECMP):** Select the check box.
3. Click the **edit** (pencil) icon next to the Redistribute routes to SD-WAN Fabric data field.
The SD-WAN Fabric Route Distribution Maps dialog box appears.
4. Verify that the dialog box includes a rule with the following settings:
 - a. **Match Criteria:** Source Protocol Local/Static
 - b. **Permit:** Allow
This ensures that the static routes are advertised to the SD-WAN fabric.
5. Modify or add a rule (if necessary), and then click **Apply**. If no modification is required, click **Cancel**.
6. Click **Add route**.
7. On the Add Route dialog box, enter the following settings:
 - a. **Subnet/Mask:** Enter an IP address space to the workload subnet.
 - b. **Next Hop (Optional):** Enter the LAN0 next hop IP address.
 - c. **Interface:** Enter the interface that connects to the Load Balancer (LAN0 interface).

- d. **Metric:** Enter 50.
 - e. **Tag:** Select **FROM_WAN**.
 - f. **Comments:** Optional
8. Click **Add** to return to the Routes dialog box and verify that the new route is listed in the routes table.
 9. Click **Apply** to return to the Routes tab and verify that the new route is listed in the routes table.
 10. If you have more Azure subnets that need to be advertised to the remote SD-WAN devices, repeat these steps. Static routes are required for all Azure subnets that need to be reachable by remote EdgeConnect devices.

Similarly, repeat these steps to create static routes on EC-V-B. Static routes are identical in each EC-V. Figure 88 shows the static routes created in the environment.

| Edit | Appliance | Segm... | Subnet/Mask | Next Hop | Label | Interface | Zone | State | Metric | Type | Additional Info | Comment |
|------|------------|---------|------------------|--------------|-------|-----------|----------------|-------|-----------|------|-----------------|--------------------|
| | QATAR1-ILB | Default | 168.63.129.16/32 | 10.200.30.33 | | Ian0 | None (Default) | UP | 50 (AD-1) | | Tag ANY | ThirdParty_Azur... |
| | QATAR1-ILB | Default | 10.200.1.0/24 | 10.200.30.33 | | Ian0 | None (Default) | UP | 50 (AD-1) | | Tag ANY | |
| | QATAR2-ILB | Default | 168.63.129.16/32 | 10.200.30.81 | | Ian0 | None (Default) | UP | 50 (AD-1) | | Tag ANY | ThirdParty_Azur... |
| | QATAR2-ILB | Default | 10.200.1.0/24 | 10.200.30.81 | | Ian0 | None (Default) | UP | 50 (AD-1) | | Tag ANY | |

Figure 88. Creating static routes on each Azure EC-V to advertise Azure VNets to remote EdgeConnect devices and static routes for responding to health probes.

This concludes traffic redirection tasks for the ILB design.

9. Prevent the static routes created for BGP establishment from being advertised to the SD-WAN fabric

In ARS, transit VNet + vWAN hub, and Managed NVA designs, two static routes were created on each EC-V to reach the Azure eBGP endpoints. These routes should not be advertised to other SD-WAN devices via subnet sharing.

The following steps prevent the static routes from being advertised to the SD-WAN fabric via subnet sharing by creating a rule on the SD-WAN Fabric Route Redistribution Maps table:

1. In Orchestrator, navigate to **Configuration > Networking > Routing > Routes**.
2. Click the **edit** (pencil) icon next to the Redistribute routes to SD-WAN Fabric drop-down menu on the Routes tab. The SD-WAN Fabric Route Redistribution Maps dialog box appears.
3. Click **Add Rule**, and then enter the following settings:
 - a. **Select Match Criteria – Source Protocol:** Local/Static
 - b. **Select Match Criteria – Prefix:** Select and enter the IP address of the ARS's or the vWAN hub's first BGP endpoint.
 - c. **Set Actions – Permit:** Clear the check box.
4. Click **Add**.
5. Click **Add Rule** again and perform step 3 above to add the ARS's or the vWAN hub's second BGP endpoint.
6. Click **Add**.
7. Click **Apply** to close the SD-WAN Fabric Route Redistribution Maps dialog box.
8. Click **Apply** to close the Routes panel and return to the Routes tab.
9. Repeat the steps in this section on the second EC-V.

You have successfully prevented the static routes from being advertised to the SD-WAN fabric on both EC-Vs.

10. Create a cluster profile to enable flow redirection on the EC-V

EC-V can only be deployed in active-active failover mode with Azure ILB. Active-active failover mode results in flow asymmetry. Generally, flow asymmetry is not an issue for most applications. However, it can break traffic when firewalls are deployed on the LAN side of the EdgeConnect gateways, as stateful firewalls could drop asymmetric flows due to mismatched session states. Additionally, if WAN Optimization (Boost) capability should be enabled on your Azure EC-V deployed with the ILB, flow redirection must be enabled on the Azure EC-Vs. This is because the TCP acceleration capability on the EdgeConnect SD-WAN gateways requires symmetric TCP traffic. Flow redirection merges the traffic of an asymmetric flow into a single appliance, thus removing the flow asymmetry. **Figure 89** depicts the redirection of a traffic flow through ECV-A (AZ1).

Note

Flow redirection only redirects TCP traffic. It does not redirect any non-TCP traffic such as UDP and ICMP.

In ECOS 9.4.1.0 and later, the following limitations apply when flow redirection is enabled:

- When segmentation is enabled, flow redirection works only for intra-segment traffic. It does not work for intersegment traffic.
- Flow redirection does not work for IPv6 traffic.
- When Zone-Based Firewall is enabled, flow redirection only works if the interface participating in flow redirection is in the Default zone.
- Enable flow redirection on EC-Vs using cluster profiles. This lets you manage multiple EdgeConnect gateways as a cluster and orchestrate flow redirection within it. Applying a cluster profile to EdgeConnect gateways makes each gateway inherit those settings.

Note

Starting with release 9.5, *Site Names* are called *Site/Cluster Names*. The Site/Cluster Name must match precisely for each appliance in the cluster.

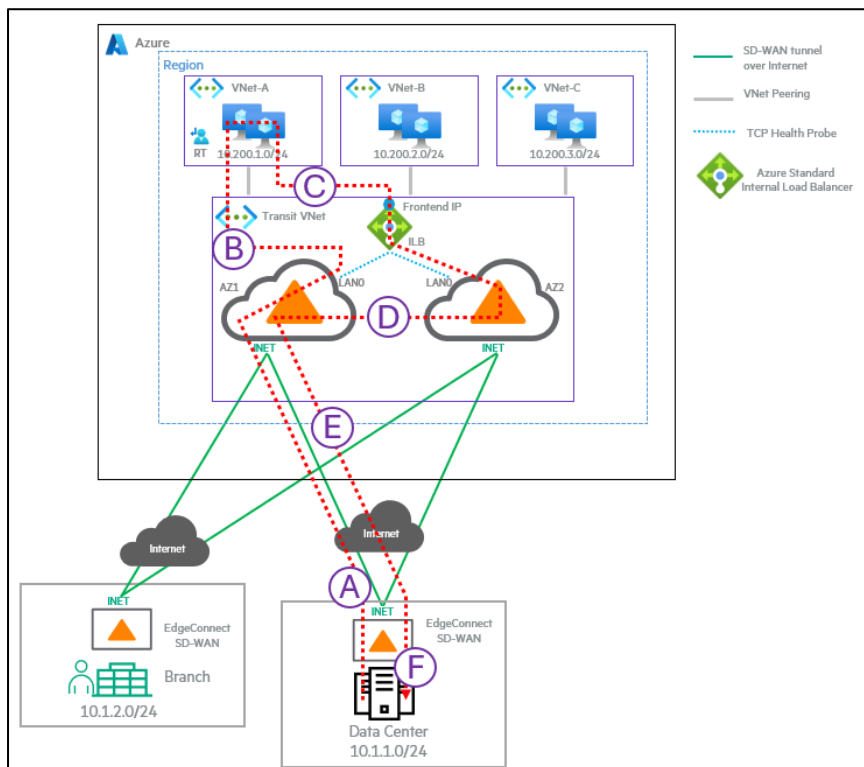


Figure 89. Viewing the traffic path when flow redirection is enabled (2).

To configure flow redirection on EC-V gateways, perform the following tasks:

- [10.1. Create a cluster profile](#)
- [10.2. Add EC-V gateways to the cluster](#)
- [10.3. Apply a cluster profile to EC-Vs](#)

10.1. Create a cluster profile

The following steps create a cluster profile:

1. Log in to Orchestrator, and then navigate to **Configuration > Networking > Cluster Profiles**.
2. Click **+Add**, and then enter the following settings:
 - a. **Name:** Enter a name for the cluster profile.
 - b. **Interface/label:** If you have already created a label for your LAN0 interface, select that label. If not, select **LAN0**.
 - c. **Flow redirection:** Select the check box.
 - d. **Wait time:** Enter **50**.
 - e. **User session sync:** Select the check box.
 - f. **Secure:** Select the check box.
3. Click **Save**.

10.2. Add EC-V gateways to the cluster

To add an EC-V gateway to a cluster, you need to assign the cluster name to the EC-V. There are three ways to do this:

1. Assign a site/cluster name when setting up a new appliance using the Appliance Configuration Wizard (Configuration > Overlays & Security > Discovery > Configuration Wizard).
2. Assign a site/cluster name to a specific appliance by accessing the System Settings for the appliance from either the appliance tree or the System Information tab.
 - a. In the appliance tree, locate an appliance, click the **menu** button, then click **System Information**. Or, navigate to **Administration > Software > Upgrade > System Information**, and then click the **edit** (pencil) icon for an appliance.
The System Information dialog box appears.
 - b. Click **System Settings**.
 - c. To add an appliance to a cluster, in the Site/Cluster name field, enter the name of the cluster.
 - d. To remove an appliance from a cluster, in the Site/Cluster name field, delete the name of the cluster.
 - e. Click **Save**.
3. Assign a site/cluster name to multiple appliances on the Cluster tab or the Tunnels tab.
 - a. From either the Clusters tab (Configuration > Networking > Clusters) or the Tunnels tab (Configuration > Networking > Tunnels > Tunnels), click **Sites/Clusters**.
The Appliance Site/Cluster Info dialog box appears.
 - b. To add an appliance to a cluster, find the appliance in the list, click in the **Site/Cluster** column, and then enter the name of the cluster.
 - c. To remove an appliance from a cluster, find the appliance in the list, click in the **Site/Cluster** column, and then delete the name of the cluster.
 - d. Click **Apply**.

10.3. Apply a cluster profile to EC-Vs

To apply the cluster profile to the EC-Vs:

1. In the appliance tree, select the appliances to which you want to apply or remove a cluster profile.

2. Click **Apply Cluster Profiles**.

The Apply Cluster Profiles dialog box appears.

3. From the Cluster Profile menu, select the profile.
4. To apply the cluster profile to the selected appliances, select the **Add** check box. To remove the cluster profile from the selected appliances, select the **Remove** check box.
5. Click **Apply**.

Visit [HPE.com](https://www.hpe.com)

[Chat now](#)

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Trademark acknowledgments, if needed. All third-party marks are property of their respective owners.

HEWLETT PACKARD ENTERPRISE

[HPE.com](https://www.hpe.com)

