vAPV Installation Guide for Amazon Web Services (AWS)

Copyright Statement

Warning: Modifications made to the Array Networks unit, unless expressly approved by Array Networks, Inc., could void the user’s authority to operate the equipment.

Contacting Array Networks

Please use the following information to contact us at Array Networks:

- **Website:**

https://www.arraynetworks.com/

- **Telephone:**

Phone: (408)240-8700

Toll Free: 1-866-692-7729 (1-866-MY-ARRAY)

Support: 1-877-992-7729 (1-877-99-ARRAY)

Fax: (408)240-8754

Telephone access to Array Networks, Inc. is available Monday through Friday, 9 A.M. to 5 P.M. PST.

- **E-mail:**

info@arraynetworks.com

- **Address:**

1371 McCarthy Boulevard

Milpitas, California 95035, USA

Revision History

Date	Description
September 22, 2015	Initial official version.
July 28, 2016	Added the information about Pre-license.

Table of Contents

1 Introduction

Array Networks vAPV is an easy-to-use, flexible, secure, high performance/capacity virtual application delivery controller. Amazon Web Services (AWS) is a leading cloud-computing platform that helps enterprises move their business from the physical network infrastructure to the cloud.

vAPV for AWS is a virtual appliance integrated with the AWS cloud environment, providing almost all of the same features as physical APV appliances. vAPV offers a comprehensive feature set including simple-to-use content routing, L3-L7 server load balancing, IPv4/IPv6 dual stack, application security and SSL offloading (using software SSL) for flexible application delivery solutions. vAPV for AWS enables simple and rapid provisioning and on-demand access to computing resources with minimum management effort, and helps achieve up to 99.999% application availability, 5x application acceleration and multi-layer application security.

Array vAPV is available as an Amazon Machine Image (AMI) in the AWS marketplace and can be deployed as an Amazon Elastic Cloud Compute (EC2) instance. With this support, AWS customers can leverage Array vAPV load balancing and other valuable features to better meet their business needs.

1.1 How Array vAPV Works on AWS

AWS provides different types of Web services, such as Amazon Virtual Private Cloud (VPC) and EC2.

Amazon VPC provisions a private, isolated section of the Amazon Web. Services (AWS) cloud where you can launch Amazon AWS resources in a virtual network that you define. With Amazon VPC, you can define a virtual network topology that closely resembles a traditional network that you might operate in your own data center.

Amazon EC2 is a Web service that provides resizable compute capacity in the cloud. Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications.

In deployment, Array vAPV and real services are launched as EC2 instances within an Amazon VPC, as shown in the following figure.

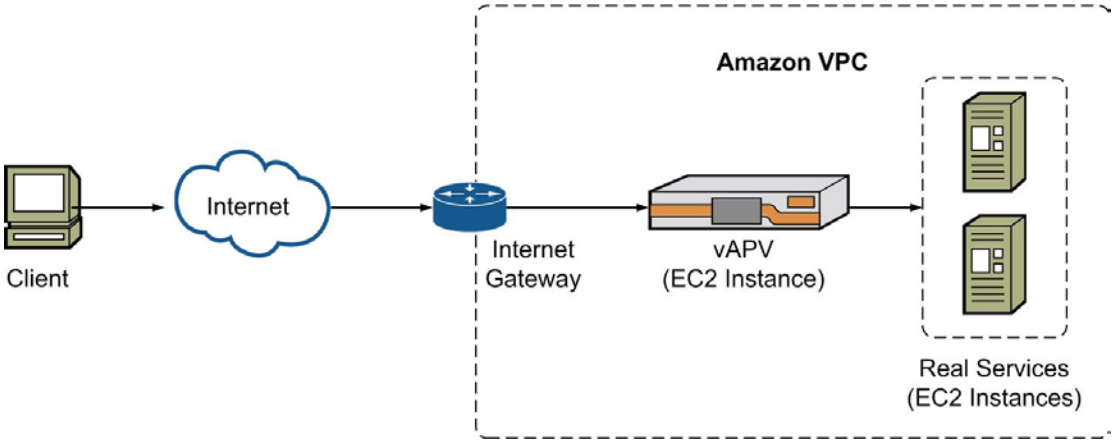

Figure 1–1 Typical Deployment Scenario of vAPV for AWS

Usually, an APV appliance needs to use multiple interfaces, for example one interface for management, one interface for receiving client requests and one interface for connecting real services.

Within the Amazon VPC, the vAPV instance has a default elastic network interface (ENI) that is assigned a primary private IP address. Additional ENIs can then be created and attached to the vAPV instance in the Amazon VPC. Each ENI can have multiple private IP addresses. The total number of supported ENIs and private IP addresses per instance depends on the instance type (refer to 1.3 Supported ENIs). For each private IP address, you can associate a public elastic IP address (EIP) to make the instance reachable from the Internet. You can also configure your Amazon EC2 instance to be assigned a public IP address at launch from Amazon's pool of public IP addresses.

For more information, please refer to http://aws.amazon.com/documentation.

1.2 Supported Instance Types

The vAPV AMI can be launched as an EC2 instance of the types shown in the following table.

Table 1–1 Supported Instance Types

Instance Family	Instance Type
General purpose	m4.large
	m4.xlarge
	m4.2xlarge
	m4.4xlarge

When an EC2 instance is launched, the specified instance type determines the hardware of the host computer used for your instance and offers different compute, memory, and storage capabilities. For details on compute, memory, and storage capabilities of each instance type, please refer to http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html.

1.3 Supported ENIs

Each EC2 instance type supports a different number of ENIs and a different number of private IP addresses per ENI. The following table lists the number of ENIs and private IP addresses per ENI supported by the EC2 instance types that the vAPV AMI supports.

Table 1–2 Supported ENIs

Instance Type	Number of ENIs	Number of Private IP Addresses per ENI
m4.large	2	10
m4.xlarge	4	15
m4.2xlarge	4	15
m4.4xlarge	8	30

1.4 Usage Limitations and Guidelines

- In the AWS cloud service architecture, the vAPV instance should be deployed in an AWS VPC to use multiple subnets (interfaces). It is recommended to use three subnets for one vAPV instance: one for management connection, one for external connection and one for internal connection. You should configure the route table correctly to make sure that the management and external subnets are publicly accessible.
- vAPV for AWS supports the following features:
  - WebWall
  - L4SLB (L4 Server Load Balancing)
  - L7SLB (L7 Server Load Balancing)
  - Caching
  - SSL Acceleration (software SSL only)
  - tProxy
  - SWCompression (Software Compression)
  - LLB (Link Load Balancing)
  - GSLB (Global Server Load Balancing)
  - QoS
  - MultiLang (Multi-language)
  - DynRoute (Dynamic Route)
  - IPv6
  - WAF

- vAPV for AWS supports both the BYOL (Bring Your Own License) model and the pre-license model. For the pre-license model, you need to purchase the AWS vAPV image with the vAPV license loaded beforehand. Users do not need to load the license again. For the BYOL model, please refer to section 2.7 Loading the vAPV License, to learn how to load the vAPV license.

2 Deployment

This section describes the process of deploying the vAPV EC2 instance within the Amazon VPC.

2.1 Creating the Amazon VPC

The deployment of the vAPV on AWS is similar to deployment in a physical network infrastructure. Before deploying the vAPV, you need to use the Amazon VPC to create a private isolated network on the AWS.

When creating the Amazon VPC, it is recommended to add at least three /24 subnets, one for the instance's management interface (management subnet), one for the instance's interface to receive client requests (external subnet) and one for the instance's interface to communicate with real services (internal subnet). Please configure routing tables to ensure that the management subnet and external subnet are accessible from the Internet.

For more information about creating and configuring the VPC, please refer to http://aws.amazon.com/cn/documentation/vpc/.

2.2 Launching the vAPV EC2 Instance

To launch the vAPV EC2 instance on AWS, please perform the following steps:

1. Log into AWS (http://aws.amazon.com) with a valid credential and switch to the AWS EC2 management console, as shown in the following figure:

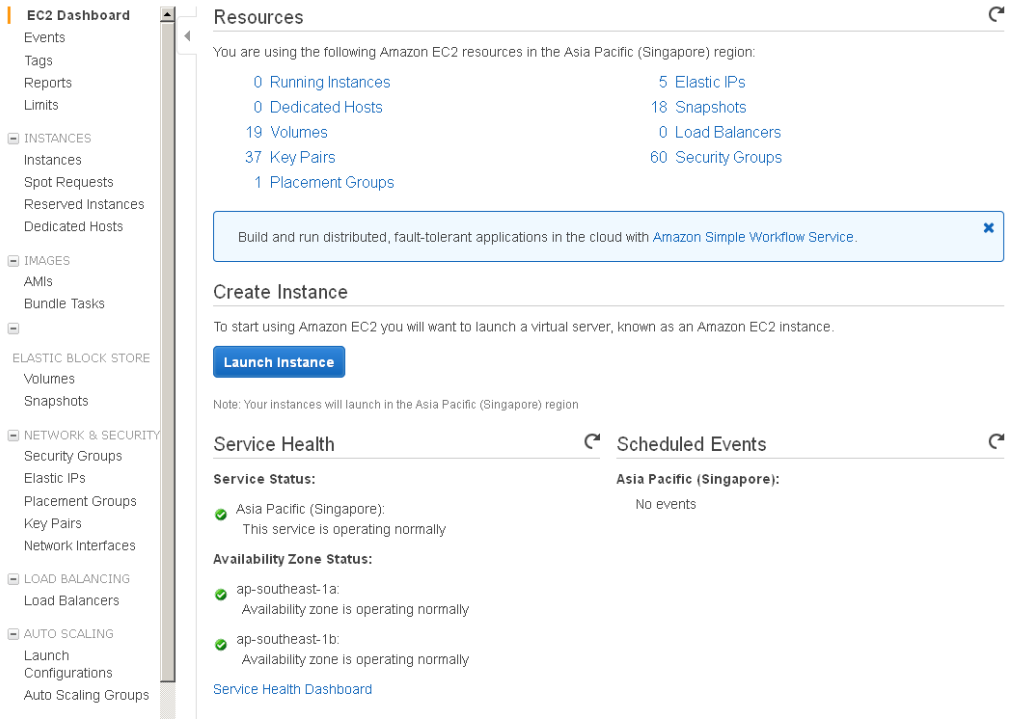

Figure 2–1 EC2 Management Console

2. In the **Create Instance** area of the **EC2 Dashboard** page, click the **Launch Instance** button, as shown in the preceding figure.

3. In the **Step 1: Choose an Amazon Machine Image (AMI)** page, click the **AWS Marketplace** tab, enter "vAPV" in the search box and click the **Select** button to select the "Array Networks vAPV - BYOL" image, as shown in the following figure.

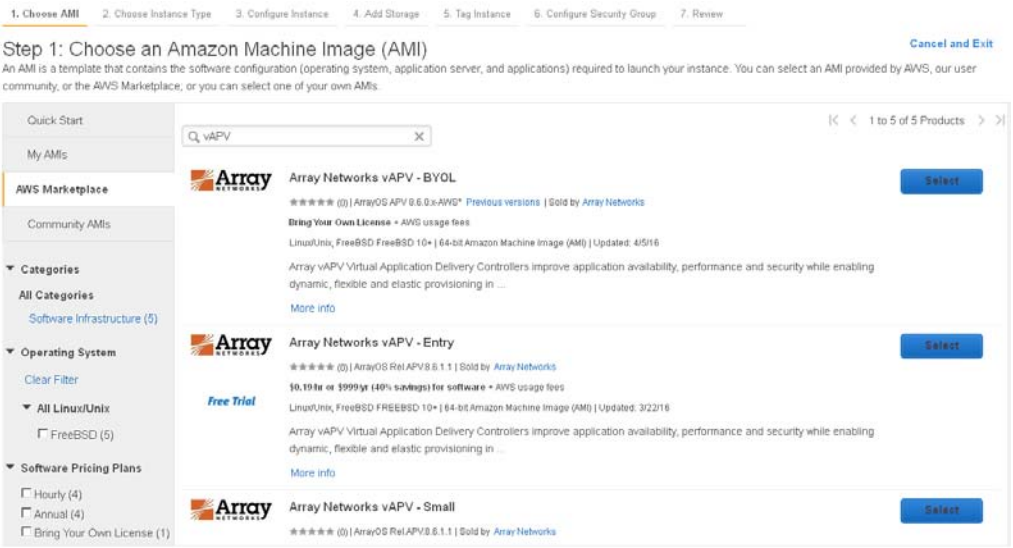

Figure 2–2 Select the vAPV AMI

4. In the **Step 2: Choose an Instance Type** page, select one of the instance types supported by the vAPV AMI, such as "m4.large", and click the **Next: Configure Instance Details** button, as shown in the following figure.

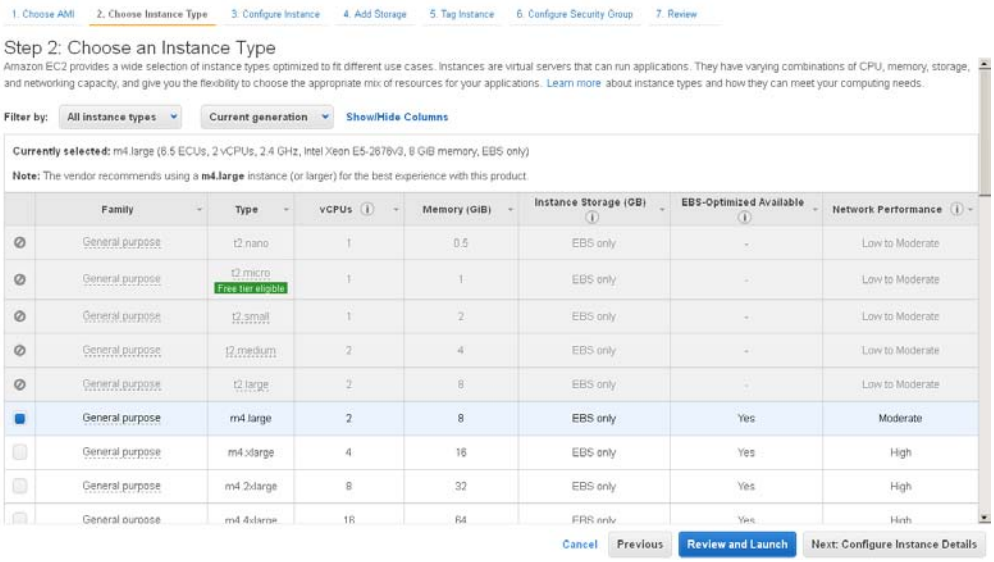

Figure 2–3 Select the Instance Type

5. In the **Step 3: Configure Instance Details** page, set the **Network** parameter to an existing VPC and set the **Subnet** parameter to the management subnet of the VPC. In the **Network** Interface area, click the **Add Device** button to add another interface to the instance and assign the internal subnet to this new interface. Then click the **Review and Launch** button, as shown in the following figure.

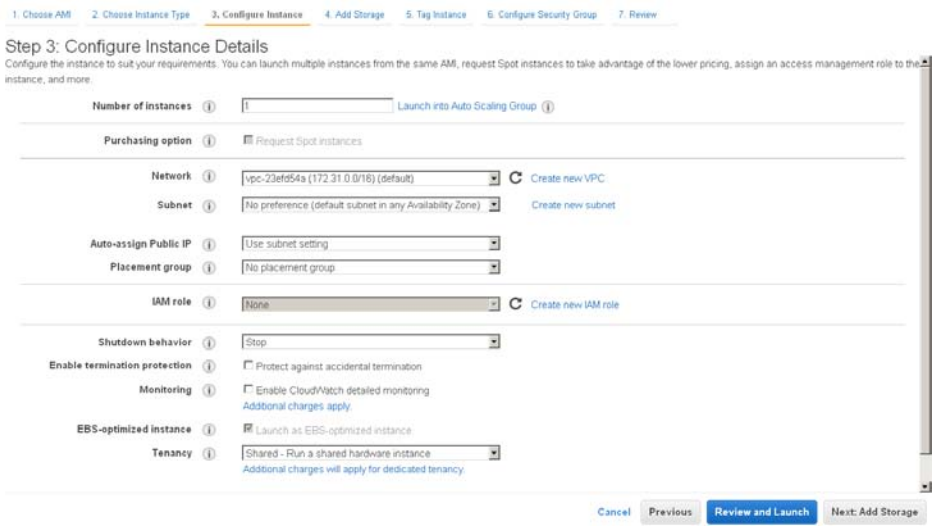

Figure 2–4 Configure Instance Details

6. Click the **6. Configure Security Group** tab, add security group rules to allow access via SSH and WebUI and then click the **Review and Launch** button, as shown in the following figure.

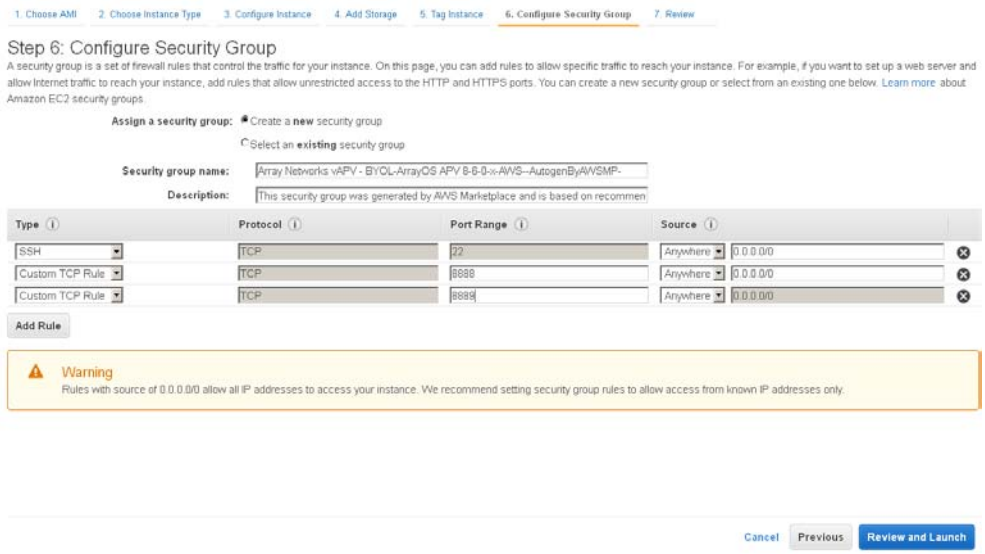

Figure 2–5 Configure Security Group Rules

Note: To allow administrators to access the vAPV instance via SSH and WebUI from the Internet, please add security group rules for port 22 of SSH, for port 8888 of the New

WebUI and for port 8889 of the Legacy WebUI.

7. Review the instance information and click the **Launch** button, as shown in the following figure.

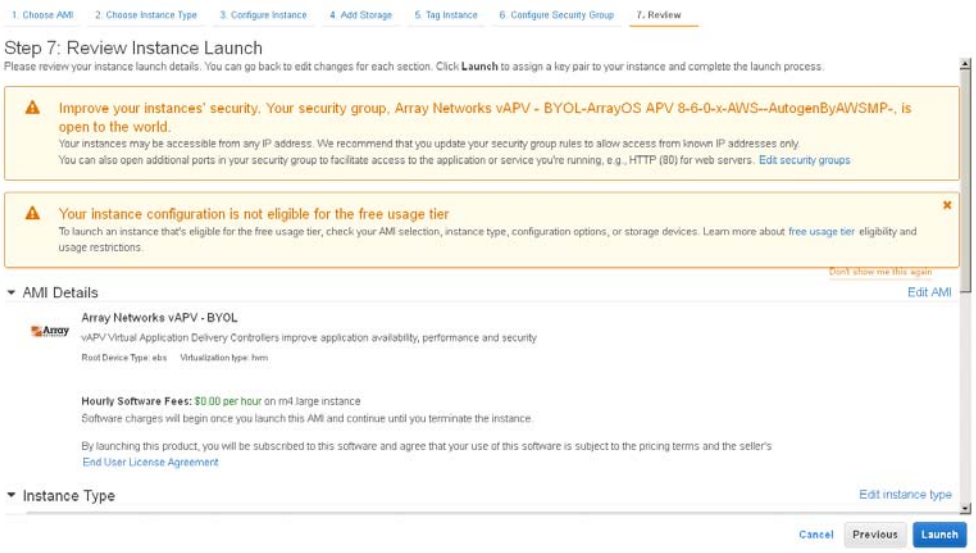

Figure 2–6 Review and Launch the Instance

8. In the prompted dialog box, select **Create a new key pair** from the drop-down list box, specify the **Key pair name** parameter, and click the **Download Key Pair** button to download the key file and then click the **Launch Instances**, as shown in the following figure.

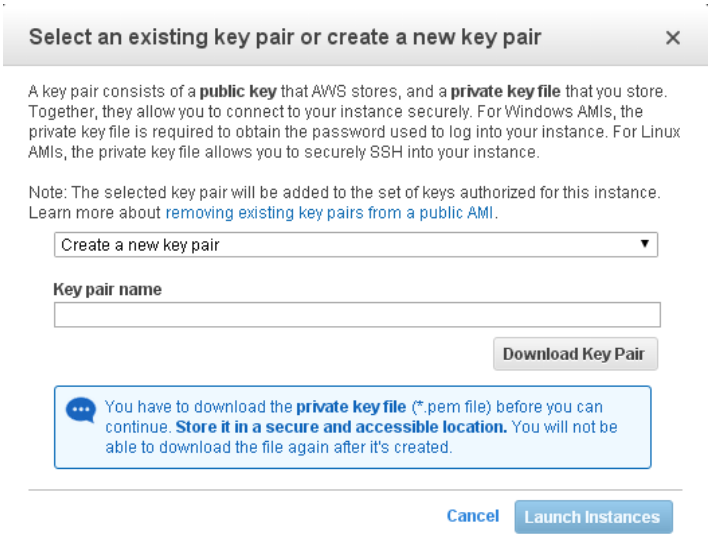

Figure 2–7 Create the New Key Pair

The newly created vAPV instance will be launched successfully as shown the following figure.

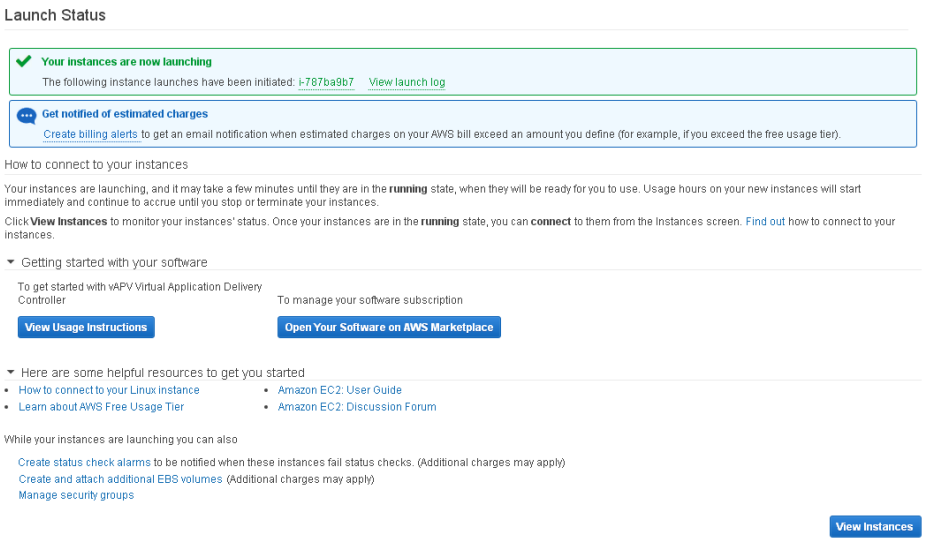

Figure 2–8 vAPV Instance Launched Successfully

2.3 Adding the Other Interfaces to the vAPV Instance

To use more interfaces, you will need to add other interfaces to the vAPV instance after it is launched.

To add a network interface to the vAPV instance:

1. Access the AWS EC2 management console, click the **Network Interfaces** link in the navigation pane, and then click the **Create Network Interface** button, as shown in the following figure.

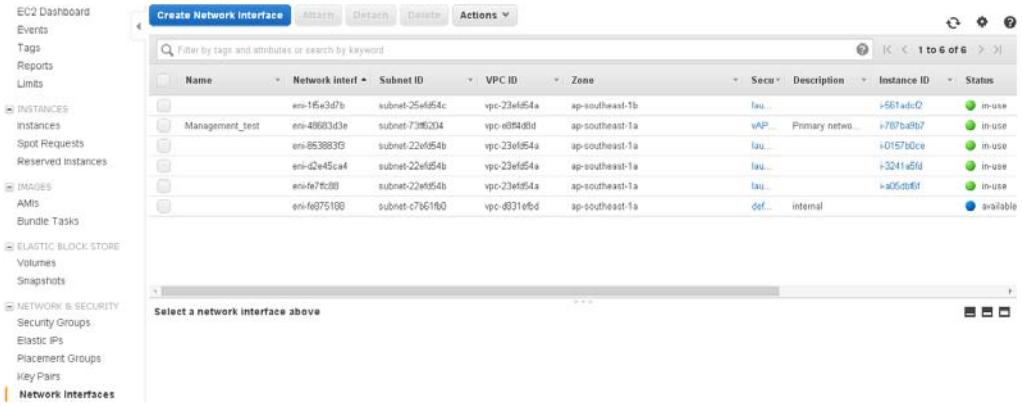

Figure 2–9 Create a Network Interface

2. In the **Create New Interface** dialog box, specify the parameters and click the **Yes, Create** button, as shown in the following figure.

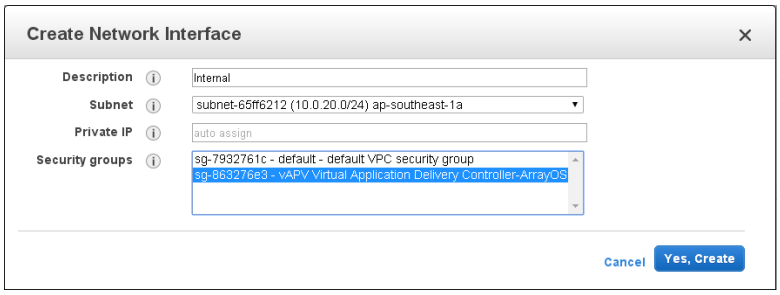

Figure 2–10 Set the Parameters of the Network Interface

3. Select the entry of the newly created network interface and click the **Attach** button. In the **Attach Network Interface** dialog box, specify the **Instance ID** parameter and click the **Attach** button, as shown in the following figure.

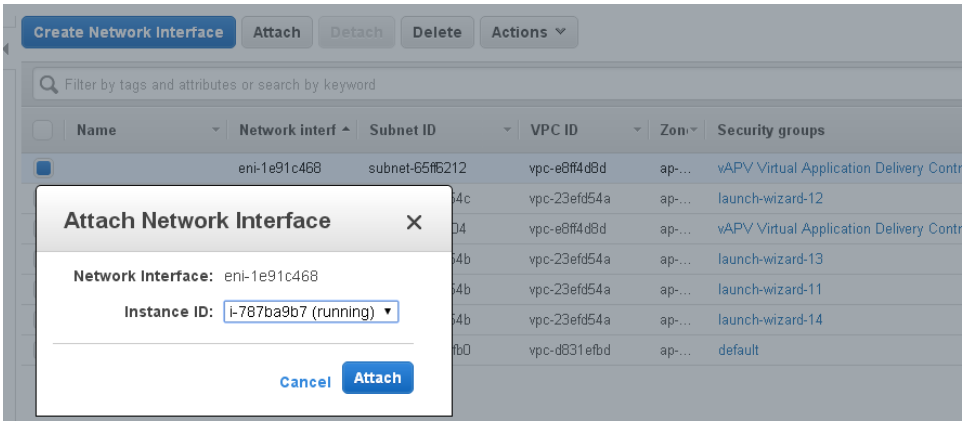

Figure 2–11 Attach the Network Interface to the Instance

2.4 Adding Secondary Private IP Addresses to the External Interface

To provide virtual services, secondary private IP addresses need to be added to the external interface of the vAPV instance.

To add a secondary private IP address to the external interface:

1. Access the AWS EC2 management console, and click the **Network Interfaces** link in the navigation pane. Select the entry of the external interface and click the **Actions** button to select **Manage Private IP Address**, as shown in the following figure.

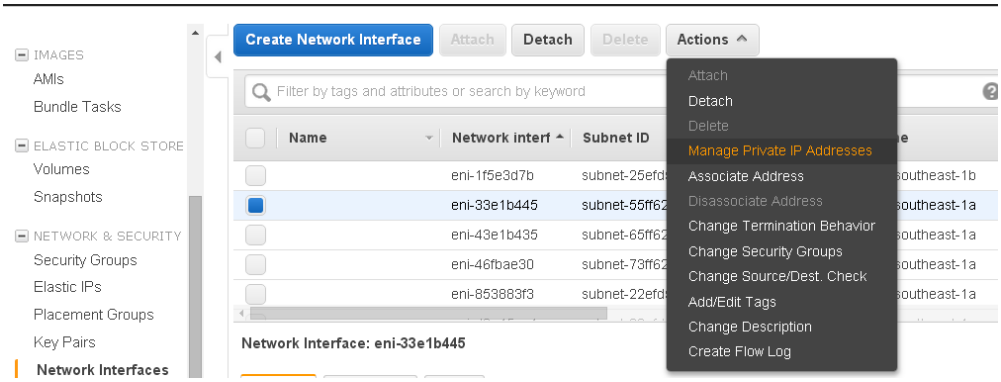

Figure 2–12 Manage Private IP Addresses

2. In the **Manage Private IP Addresses** dialog box, click the **Assign new IP** link to add a new secondary IP address, and then click the **Yes, Update** button, as shown in the following figure.

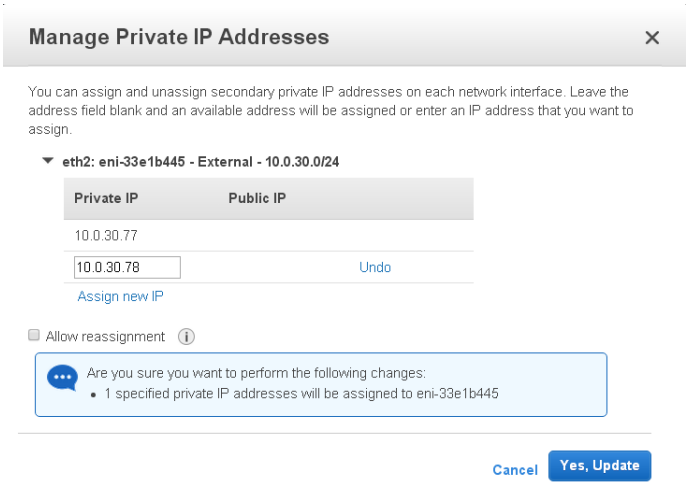

Figure 2–13 Add a New IP Address

When configuring server load balancing (SLB) virtual services, you should use the secondary private IP address as the VIP. To make the secondary private IP address publicly accessible, you need to add an elastic IP and associate it with the secondary private IP address. The next section describes this process.

2.5 Adding Elastic IPs to the Network Interfaces

To make the management interface and the external interface of the vAPV instance publicly accessible, you need to add the elastic IPs and associate them with the private IP addresses of the network interfaces.

To add the elastic IP to a network interface:

1. Access the AWS EC2 management console, click the **Elastic IPs** link in the navigation pane, and then click the **Allocate New Address** button. In the **Allocate New Address** dialog box, click the **Yes, Allocate** button, as shown in the following figure.

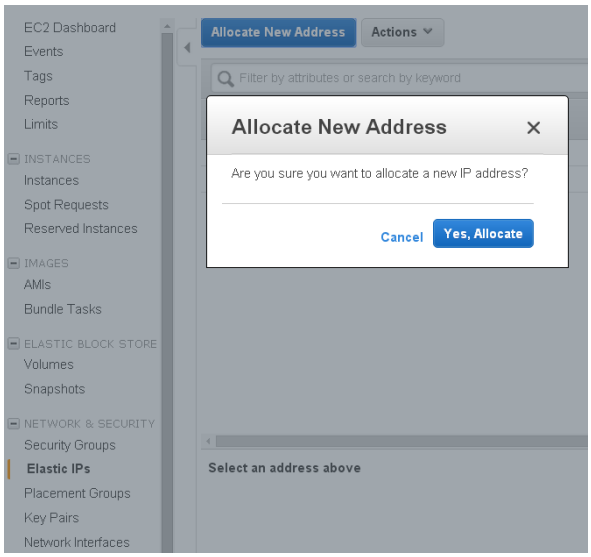

Figure 2–14 Allocate the New Address

2. Select the entry of the newly created elastic IP, click the **Actions** button and select the **Associate Address** item, as shown in the following figure.

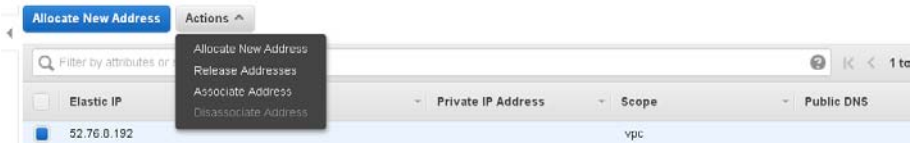

3. In the **Associate Address** screen, specify the parameters **Network Interface** and **Private IP Address** and click the **Associate** button, as shown in the following figure.

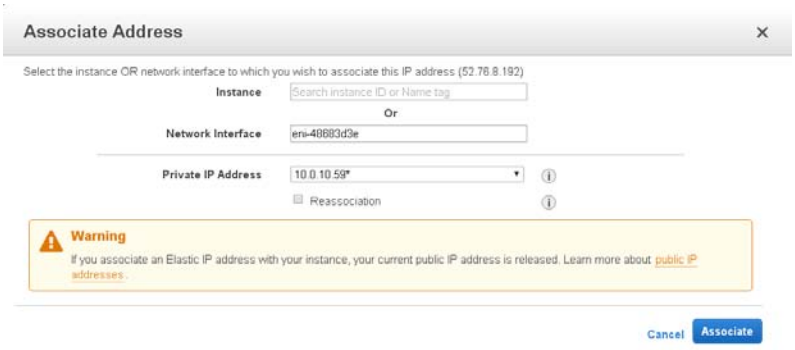

Figure 2–15 Associate the Elastic IP

2.6 Accessing the vAPV Instance

2.6.1 Accessing the vAPV Instance via SSH

You can connect to the vAPV instance via SSH after the status of the newly created vAPV instance becomes “Running”.

To access the vAPV instance via SSH, use the public DNS Name or IP address and the SSH port 22 as the SSH access point.

For the Bring Your Own License model, the downloaded private key file will be required for logging into the vAPV instance. After you successfully log into the vAPV instance, the following screen will be displayed.

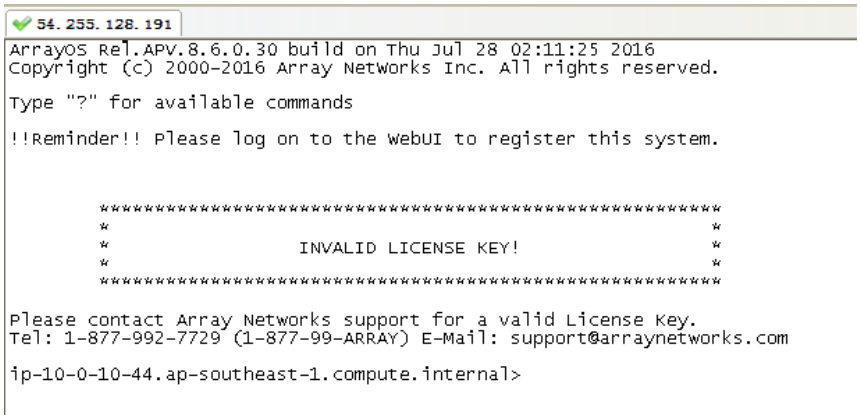

Figure 2–16 Access the vAPV Instance via SSH

Note: For the Bring Your Own License model, the INVALID LICENSE KEY message is displayed because you have not yet entered the license information. That process is described in section 2.7 Loading the vAPV License.

2.6.2 Accessing the vAPV Instance via WebUI

To access the vAPV via WebUI, you will first need to access the vAPV instance via SSH to perform the following configurations in the Config mode:

- Change the password of the default account (array) using the "**passwd user array** *new_password*" command.
- Enable the WebUI using the "**webui on**" command.
- (Optional) Configure the WebUI port using the "**webui port**" command.

After the preceding configurations are completed, you can access the WebUI of the vAPV instance at https://<EIP>:<WebUI_port> using a Web browser. On the login page, enter "array" as the username and the previously configured "new_password" as the password to pass the authentication.

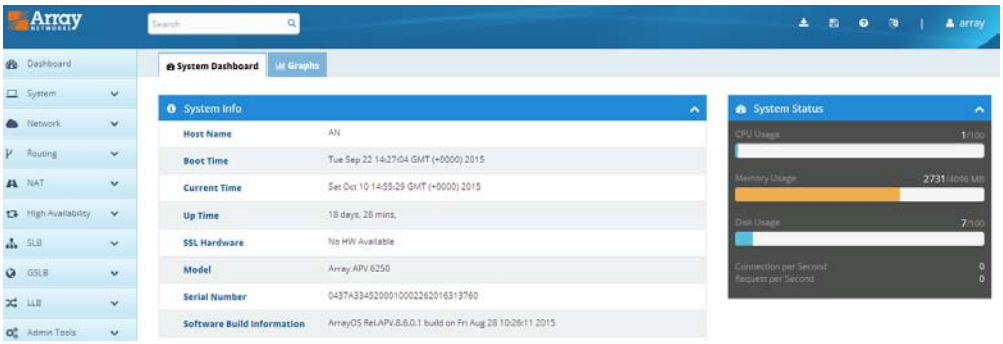

Figure 2–17 Access the vAPV Instance via WebUI

2.7 Loading the vAPV License

If you are using the pre-license model for the vAPV, your license is already loaded and you can skip this step. If you are using the Bring Your Own License model, to purchase a license from Array Networks and load the license to the system, please execute the following steps:

1. Access the vAPV instance via SSH.
2. View the software version, model and serial number of the vAPV by executing the "**show version**" command.
3. Contact Array Networks Customer Support to obtain a valid license key.
4. Execute the "**system license**" command in the Config mode, paste the license key and press "Enter". The license will be successfully loaded.

Alternatively, you can load the license via WebUI. To load the license, select **System > System Management > System License > License Key**.