cPacket

cPacket Cloud Suite

Quick Start Guide for AWS cVu-V Packet Broker

Version 25.3.0

Revision History

Document Version	Date	Notes
1	13Aug2024	• The original release of this document.
2	03Dec2024	• Document title/naming modifications.
3	15Apr2025	• Title update • Information about user data validation
4	23Jul2025	• No content changes
5	5Nov2025	• No content changes
6	17Nov2025	• Provide UI instructions for setup
7	29Dec2025	• Updated cloud-init user data information • Updated inbound UDP port • Updated login and license steps • Added Create VXLAN Port procedure • Added installation steps using the AWS Marketplace for the cStor-V AMI • Added installation steps using the cPacket cStor-V shared AMI • Miscellaneous updates

Table of Contents

Introduction

In this guide you will learn how to launch a cPacket cVu-V virtual appliance in your Amazon Web Services (AWS) environment to replicate packets from a Virtual Private Cloud (VPC) with Traffic Mirroring. We recommend using this guide to set up a basic cVu-V deployment in accounts that are primarily used for testing and evaluation. cPacket Solutions Engineering will work with you to set up cPacket solutions at scale with scripting and automation when you are ready to deploy the solutions more broadly in AWS.

Getting started

This AWS VPC deployment is of a cPacket cVu-V which is part of a network visibility solution where traffic from network sources is mirrored and processed through cVu-V (a virtual visibility node) that distributes packets to two downstream components: a security tool for threat detection and a cStor-V for packet capture and storage. For network analytics and visualization, cClear-V then reads the packets from cStor-V. The architecture enables comprehensive, non-intrusive network monitoring by creating copies of live traffic and routing them to specialized analysis tools, providing organizations with security insights, forensic capabilities, and performance visibility within their cloud environment.

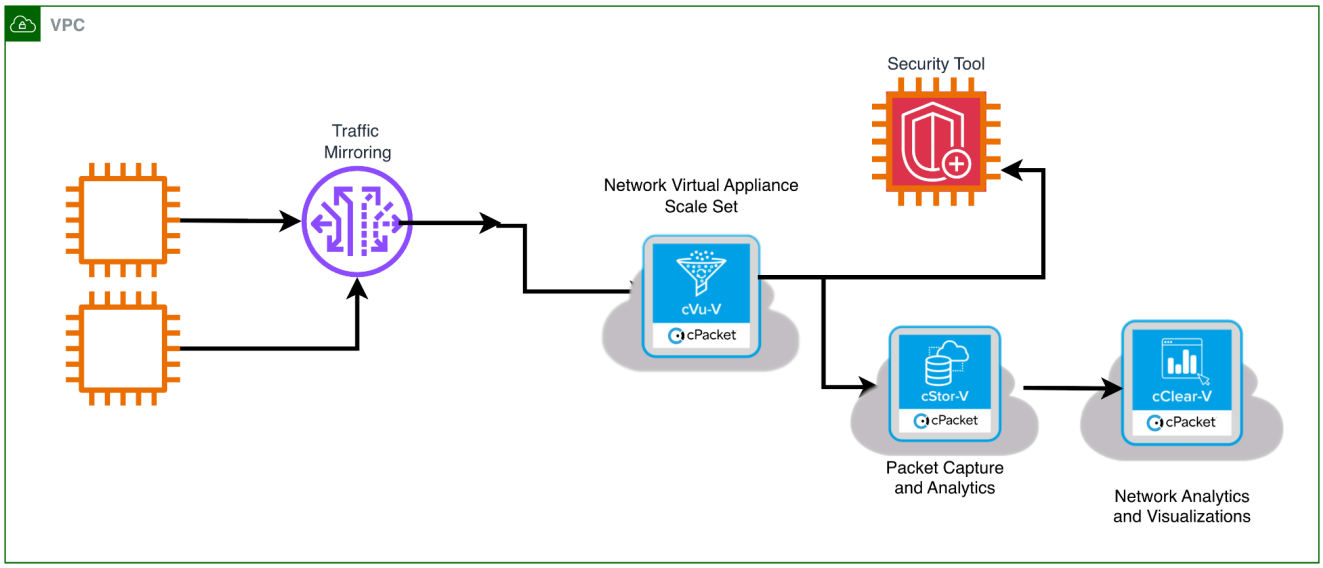

Before you begin

cPacket virtual appliances can be accessed through the AWS Marketplace or via shared Amazon Machine Images (AMI). If you cannot access the AWS Marketplace, you must provide your AWS account ID and Region to your cPacket representative, who will share the latest AMI images for the virtual appliances to your account ID in the selected regions. The Marketplace option (see "Installation Steps using AWS Marketplace cVu-V AMI") provides the simplest installation method for users with Marketplace access.

We recommend that you install cClear-V and optionally cStor-V before installing cVu-V. See the ***Quick Start Guide for AWS cClear-V Control Center*** for detailed information about installing cClear-V and the ***Quick Start Guide for AWS cStor-V Packet Capture*** for detailed information about installing cStor-V.

The following table lists all the requirements necessary to begin installation in AWS.

Requirement	Detail
AWS User ID	You will need a user ID in an AWS account, permissions for the user are listed below.
AWS Account and Region	If you cannot access the AWS Marketplace, an AWS account and region must be provided to cPacket so the virtual appliance images (Amazon Machine Image - AMI's) can be shared. Account numbers are 12 digits in length and example regions names are: us-east-1, us-west-2, …
cPacket License Key	cPacket will provide you with a license key used to activate the cVu-V appliance.
AWS Organization Tagging Policies	Your organization may have requirements for tagging resources created in the cloud. Common tag requirements are: Name, Owner, and CreatedBy. If these tags are mandated by your organization, creating a device without them will fail the organizational policy. Remediation is to simply add the required tags at resource creation time.
Identity and Access Management (IAM) User/Permissions/Role	In the account used for installation, the user needs to have permissions granted to setup the cVu-V virtual appliance. See IAM Policy to install cVu-V for the minimum permissions. Your organization may already have roles defined that grant these permissions through their IAM policies, if not there are existing AWS managed policies that can be granted to the user or the group the user is a member of. They are: - AmazonEC2FullAccess - AmazonEC2RoleforSSM - AWSMarketplaceManageSubscriptions (for marketplace access)
Network Bandwidth	The network bandwidth to be captured will determine the size of the instance for the cVu-V appliance. Recommendation on instance sizing and the number of storage volumes are contained later in this guide and indexed by network bandwidth of Gbps you would like to capture.

<table>
<tr><td>Virtual Private Cloud (VPC)</td><td>You will need a VPC to install the cVu-V into. We recommend working with your organization's AWS cloud support team and requesting a role that allows creating a VPC.</td></tr>
<tr><td>Security Groups/Policies</td><td>The following ports will be opened on the cVu-V for inbound and outbound traffic.

Inbound:
<table><tr><td>TCP 22</td><td>SSH traffic</td></tr><tr><td>TCP 443</td><td>Encrypted HTTP traffic</td></tr><tr><td>UDP 4789</td><td>UDP port to receive VXLAN encapsulated traffic</td></tr></table>
Outbound ports:
- All TCP/UDP ports</td></tr>
<tr><td>SSH Key Pair</td><td>A SSH public/private key pair to control access to the virtual appliances and hosts. If you don’t currently have an SSH Key Pair you can create one during installation.</td></tr>
<tr><td>AWS Cloud Shell Access</td><td>The account and user you are using must have permission to use the AWS Console.</td></tr>
</table>

Create a service VPC and other resources (Optional)

If you do not already have a service VPC with cClear-V and/or cStor-V, use the following procedure to create a service VPC plus the additional VPC resources that you need to run cVu-V, such as subnets, route tables, internet gateways, and NAT gateways.

To create a VPC, subnets, and other VPC resources using the console

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. On the VPC dashboard, choose **Create VPC**.
3. For **Resources to create**, choose **VPC and more**.
4. Keep **Name tag auto-generation** selected to create Name tags for the VPC resources or clear it to provide your own **Name** tags for the VPC resources.
5. For **IPv4 CIDR block**, enter an IPv4 address range for the VPC. A VPC must have an IPv4 address range.

6. Select the **Tenancy** option. Choose the tenancy of the VPC to be `Default`, EC2 instances launched into this VPC will use the tenancy attribute specified when you launch the instance.
7. For **Number of Availability Zones (AZs)**, we recommend that you provision subnets in at least two Availability Zones for a production environment. To choose the AZs for your subnets, expand **Customize AZs**. For this deployment, choose 1 availability zone.
8. To configure your subnets, choose values for **Number of public subnets** and **Number of private subnets**. To choose the IP address ranges for your subnets, expand **Customize subnets CIDR blocks**. For this deployment, choose 1 public subnet and 1 private subnet. The public subnet will be the management subnet through which network operators can reach the cPacket appliances. The private subnet will be the capture subnet where replicated packets will reach cVu-V instances.
9. For **NAT gateways**, for this deployment, choose **1 per AZ**.
10. Select **None** for **VPC Endpoints**.
11. Enable DNS hostnames and DNS resolution.
12. (Optional) To add a tag to your VPC, expand **Additional tags**, choose **Add new tag**, and enter a tag key and a tag value.
13. When you are finished configuring your VPC, choose **Create VPC**.

Installation Steps using AWS Marketplace cVu-V AMI

The following steps describe how to install the cVu-V image from the AWS Marketplace. This is the easiest and simplest way to install a basic cVu-V appliance. Your user must have permission to install subscriptions from the Marketplace for this installation procedure.

1. Sign in to AWS with your username and password.
2. Search the Amazon Marketplace for cPacket, select **Discover Products**, and then select **cPacket cVu-V - Cloud Packet Broker**.
3. Select **Continue to Subscribe**.
4. Accept the Terms and Conditions for the **cPacket cVu-V** and select **Continue to Configuration**.
5. Select the options on the **Configure this software** page to launch the cVu-V.

Field	Value
Fulfillment option	Select **64-bit (x86) Amazon Machine Image (AMI)**.
Software version	Select the software version to deploy, we recommend using the latest release in the Marketplace.
Region	Use the dropdown to specify the region where the software should be deployed.

6. Select **Continue to Launch** to continue.
7. On the **Launch this software** page, select **Launch from EC2 console** for the **Launch method** field.
8. On the **Launch an instance** page, enter or select the following information:

Field	Value
Name and tags	Add any tags you'd like. As a best practice, we recommend using **Name** and **Owner** tags.
Instance type	Leave **c5n.xlarge** as the default.
Key pair (login)	You can use an existing key pair or create a new one.
Network	Select the VPC you want to launch your instance into.
Subnet	Select the subnet you want to use from the dropdown or select **Create a new subnet**.

9. Select **Launch instance** when ready to continue.

A success message confirms the launch of the instance and displays the instance ID, you will need the instance ID in order to login to the cVu-V appliance. Users may need to wait up to 10 minutes after deploying before all services are fully operational. You can now continue to test your installation and setup packet mirroring, proceed to After Installation and Launching.

Installation using the cPacket cVu-V Shared AMI

The following steps describe how to install the cVu-V image shared to your AWS account and region. You only need to follow these steps if your organization and user do not have access to the AWS Marketplace.

1. Sign in to AWS with your username and password.
2. Select **EC2**.
3. In the left navigation panel, under Images, select **AMIs**.
4. Above the table of AMIs, change the **Filter** from **Owned by Me** to **Private Images**.
5. Using the filter box, find the cPacket cVu-V (CVU) AMI and then press ENTER.
6. Select the checkbox next to the cPacket cVu-V AMI and select **Launch Instance** from AMI.
7. Adding any tags required by your organization. The “Name” tag becomes the instance’s name and you can add tags such as Owner or CreatedBy. Make sure you select the

Resource Types to apply the tags to, for cVu-V you can select **Instances**, **Network Interfaces**, and **Volumes**.

8. For **Instance type**, choose **m5n.xlarge**.
9. Select or create a SSH Key pair. This key is used to access the instance remotely over SSH. If you create a new one it will download the .PEM file to your **Downloads** folder.
10. Select the **Edit** button on the **Network Settings** to edit VPC and network policies for the cVu-V appliance.
11. Select the VPC to run from the VPC drop down menu.
12. Make sure the **Auto-assign public IP** is enabled, this will allow access to the web UI on the cVu-V from the Internet.
13. You can select an existing security group used by your organization or create one when you launch this cVu-V. The following security rules and protocols are needed for the cVu-V, you can add additional protocols by selecting **Add security group rule**.

Type	Port	Source Type
SSH	22	Select My IP for restricted or Anywhere for unrestricted access
HTTPS	443	Select My IP for restricted or Anywhere for unrestricted access
UDP	4789	Select 0.0.0.0/0 – VXLAN port to receive mirrored traffic **

** Note: In production deployments this CIDR should be limited to subnets that will be mirroring traffic.

14. From the **Shutdown** behavior drop-down list, select **Stop**.
15. Select **enable** in the **Termination protection** drop down list.
16. Select **Launch Instance**.

After Installation and Launching

Please allow up to 10 minutes for the cVu-V to become accessible.

Log In and License

1. Open a web browser and navigate to the cVu-V's IP address: https://<cVu-V Management IP>.
2. The password will be the EC2 Instance ID from the AWS Console. To find this, navigate to the EC2 console, select your cVu-V instance, and locate the Instance ID in the

instance details (it will be in the format *i-0123456789abcdef0*). Use this Instance ID as the password to log in.

Edit cVu-V Mode

The cVu-V mode indicates the mode in which cVu-V is operating. cVu-V must be configured in **endpoint** mode to handle mirroring packets to downstream tools in an out-of-band fashion.

1. Navigate to the **Admin->Settings** section of the Web UI.
2. In the **General** section, click the **Edit** (pencil) icon in the upper right corner.
3. The **Edit General Settings** dialog will appear, allowing you to modify the **cVu-V mode**.
4. Select **endpoint**.

Create VXLAN Port

A VXLAN port is required to create the overlay network connection between the cVu-V and the cStor-V.

1. Navigate to the **Admin->Settings** section of the Web UI.
2. In the **VXLAN Ports (packet mirror output)** section, click **Add new VXLAN**.
3. In the **Add new VXLAN** dialog, configure the following:

Parameter	Value	Configuration Details
VXLAN ID	1337	This ID must match the corresponding ID on the downstream tool.
Remote IP	<IP Address of cStor-V>	The IP address of the cStor-V appliance. This should be set to the cStor internal IP.

4. Click **Save** or **Activate** to finalize the VXLAN port creation.

cVu-V requires you to have a valid license to replicate traffic to downstream tools. You should connect this cVu-V to an existing cClear-V with an active cVu-V license. The cClear-V must have a network path to the cVu-V.

Traffic Mirroring

Once you have completed the installation and activated the license, your cVu-V is ready to receive mirrored traffic from EC2 Instances. The AWS guide for Traffic Mirroring can be found here (AWS Traffic Mirroring). This guide will walk you through the necessary steps to receive mirrored traffic from an EC2 instance.

There are three resources to create for Traffic Mirroring:

1. **Traffic Mirror Target** – This is the Gateway Load Balancer endpoint.
2. **Traffic Mirror Filter** – Rules to determine which traffic to mirror from the source(s).
3. **Traffic Mirror Session** – Session that identifies the mirrored source and target.

Create a traffic mirror target

To set up a Traffic Mirror Target you will need the identifier for the Gateway Load Balancer endpoint.

1. In the AWS Management Console, in the top menu, select **Services**.
2. Select **VPC**.
3. In the left pane, under Traffic Mirroring, select **Mirror Targets**.
4. Select **Create Traffic Mirror Target**.
5. In the **Name** tag field, type a descriptive name for the target.
6. In the **Description** field, type a description for the target.
7. From the **Target** type drop-down list, select **Gateway Load Balancer endpoint**.
8. From the **Target** search box, enter the cVu-V gateway load balancer endpoint ID recorded after its creation earlier.
9. Select **Create**.

Note:	The Traffic Mirror Target ID that is created, it will look like tmt-01421338b23ede911 with unique digits for your ID.

Create a traffic mirror filter

You must create a filter to allow or restrict traffic sent from the mirroring sources to cVu-V. For the purposes of this guide we will mirror all inbound and outbound traffic from EC2 instances.

1. In the AWS Management Console, in the left pane under Traffic Mirroring, select **Mirror Filters**.
2. Select **Create traffic mirror filter**.
3. In the **Name tag** field, type a name for the filter.
4. In the **Description** field, type a description for the filter.
5. Under **Network services**, select the **amazon-dns** checkbox.
6. In the **Inbound rules** section, select **Add rule**.
7. Configure an inbound rule:
    1. In the **Number** field, type a number for the rule, such as 100.
    2. From the **Rule action** drop-down list, select **accept**.
    3. From the **Protocol** drop-down list, select **All protocols**.
    4. In the **Source CIDR block** field, type 0.0.0.0/0.
    5. In the **Destination CIDR block** field, type 0.0.0.0/0.
    6. In the **Description** field, type a description for the rule.
8. In the **Outbound rules** section, select **Add rule**.
9. Configure an outbound rule:
    1. In the **Number** field, type a number for the rule, such as 100.
    2. From the **Rule action** drop-down list, select **accept**.
    3. From the **Protocol** drop-down list, select **All protocols**.
    4. In the **Source CIDR block** field, type 0.0.0.0/0.
    5. In the **Destination CIDR block** field, type 0.0.0.0/0.

6. In the **Description** field, type a description for the rule.

10. Select **Create**.

For VPC and subnets that contain application instances that will be mirrored, we recommend the following filtering rules to help avoid mirroring duplicate frames from peer EC2 instances that are in a single VPC to the cVu-V.

- All outbound traffic is mirrored to the cVu-V, whether the traffic is sent from one peer device to another on the subnet or if the traffic is sent to a device outside of the subnet.
- Inbound traffic is only mirrored to the cVu-V when the traffic is from an external device. For example, this rule ensures that an app server request is not mirrored twice: once from the sending app server and once from the database that received the request.
- Rule numbers determine the order in which the filters are applied. Rules with lower numbers, such as 100, are applied first.

Important: These filters should only be applied when mirroring all the instances in a CIDR block including app servers.

Create a traffic mirror session

You must create a session for each AWS resource that you want to monitor.

Note: To prevent mirror packets from being truncated, set the traffic mirror source interface MTU value to 54 bytes less than the traffic mirror target MTU value for IPv4 and 74 bytes less than the traffic mirror target MTU value for IPv6. For more information about configuring the network MTU value, see the following AWS documentation: Network Maximum Transmission Unit (MTU) for Your EC2 Instance.

1. In the AWS Management Console, in the left pane, under Traffic Mirroring, select **Mirror Sessions**.
2. Select **Create traffic mirror session**.
3. In the **Name** tag field type a descriptive name for the session.
4. In the **Description** field type a description for the session.

5. From the **Mirror source** drop-down list, select the source **ENI** to mirror, this is typically attached to the EC2 instance that you want to monitor. This EC2 instance should be in a private subnet of the VPC selected as the GWLBe VPC.
6. From the **Mirror target** drop-down list, select the traffic mirror target ID generated for the Gateway Load Balancer endpoint that was created earlier.
7. In the **Session number** field, type 1.
8. For the VNI field, leave this field empty and the system assigns a random unique VNI.
9. For the **Packet length** field, leave this field empty, this will mirror the entire packet.
10. From the **Filter** drop-down list select the ID for the traffic mirror filter you created in the previous step.
11. Select **Create.**

The traffic from the ENI you selected should now be mirrored to the cVu-V. You can verify operation by running a packet capture in the cVu-V and verifying the source address in the packet capture includes the IP address of the ENI being mirrored.

Verifying Operation

1. In cClear, add your cVu. For detailed instructions, please refer to the *cClear User Guide* which can be accessed from the help menu.

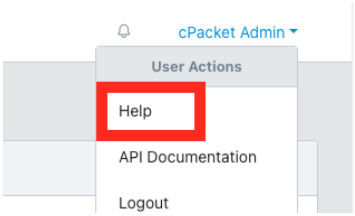

   a. Go to **Configure > cVus**.
   b. Click **Add cVu**.

      For **Auth Type**, select **User Login**.

      i. Enter the cVu name, IP address, and login credentials.
      ii. Click **Save** to add the cVu.
2. Verify that cClear-V is indicating that traffic is flowing to your cVu-V packet capture appliance.

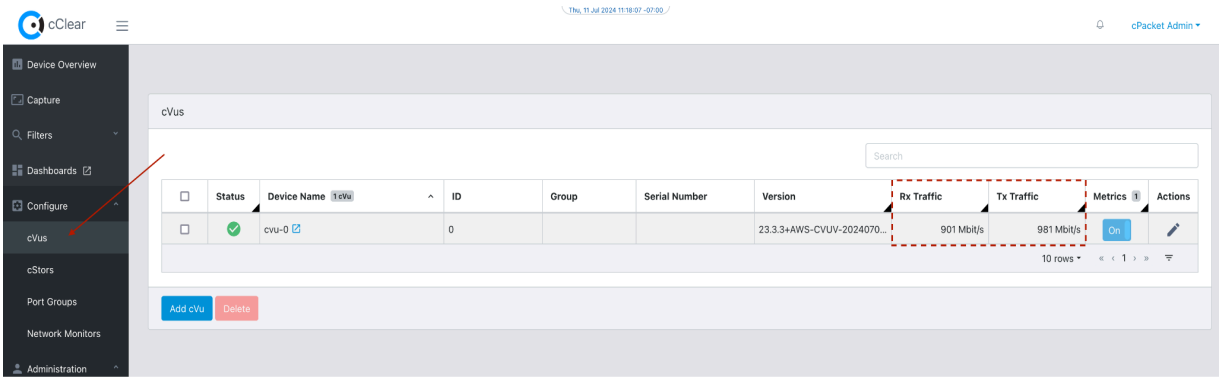

IAM Policy to Install cVu-V

The following policy is used to install and operate the cVu-V appliance. It is a very restricted policy and defines the minimum permissions necessary. To install cVu-V from the AWS marketplace, the user or the role used will need to attach the AWSMarketplaceManageSubscriptions policy.

```
{
   "Version": "2012-10-17",
   "Statement": [
       {
           "Sid": "cStorMinInstall",
           "Effect": "Allow",
           "Action": [
               "iam:ListInstanceProfiles",
               "ec2:*"
           ],
           "Resource": "*"
       },
       {
           "Sid": "cStorMinSSHKey",
           "Effect": "Allow",
           "Action": [
               "iam:GetSSHPublicKey",
               "iam:ListSSHPublicKeys",
               "iam:UpdateSSHPublicKey",
               "iam:UploadSSHPublicKey"
           ],
           "Resource": [
               "arn:aws:iam::*:user/${aws:username}"
           ]
```



```
        },
        {
            "Sid": "cStorMinCloudWatch",
            "Effect": "Allow",
            "Action": [
                "cloudwatch:GetMetricData",
                "cloudwatch:DescribeAlarms"
            ],
            "Resource": "*"
        }
    ]
}
```