



DART AI

Quick Start Guide For Amazon Web Services

CIRRIES TECHNOLOGIES, INC. SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS MADE HEREIN NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL.

This information is subject to change without notice and is furnished under a license agreement or nondisclosure agreement and may be used or copied only in accordance with that agreement.

© Cirries Technologies, Inc. 2024

Changes to This Document

Revision	Date	Change Summary
V1.0	April 2022	Corresponds to DART Software version 6.0.0
V2.0	January 2024	Template changed

Contacting Cirries Technologies

Cirries Customer Support portal can be accessed at: <https://cirries.com/support/>

New customers will find the default login and password in a file provided when the Cirries DART is installed. The customer should change the login and password as soon as possible.

The Cirries technical support phone line is 469-342-9122.

When contacting Customer Support, the following information helps diagnose and solve problems:

- EC2 instance type, AWS Region, and AWS Availability Zone
- License number, organization's name
- The text of any error messages
- Supporting screen images, logs, and error files as appropriate
- A detailed description of the problem

Table of Contents

1	Introduction	6
1.1	Advanced API Support	6
2	Requirements and Restrictions	7
2.1	Operating Environments	7
2.1.1	AWS Requirements	7
3	DART Deployment Summary.....	9
3.1	Launching Cirries templates	9
3.2	Using the Create Stack Wizard.....	9
3.3	Access Cirries DART Package from AWS Marketplace.....	13
3.4	Accessing Cirries DART Packet Sensor (1GBPS & 5GBPS&10GBPS) from AWS Marketplace.....	13
3.5	AWS Multi-size Provisioning.....	14
3.6	DART Login	14
3.7	DART PacketSensor Settings.....	16
4	AWS VPC application traffic mirroring to the Packet Sensor setup.....	17
4.1	Getting Started with VPC Traffic Mirroring.....	17

List of Tables

Table 1 Specifications for Cirries DART Instances.....	8
Table 2 Configuration Settings.....	15

List of Figures

Figure 3-1 CloudFormation Template to Create Stack with	10
Figure 3-2 Filling Out Parameters of the CloudFormation Template	10
Figure 3-3 Stack Events Page	12
Figure 3-4 New Instances Created in EC2 Management Console.....	12
Figure 3-5 AWS Multi-Size Provisioning.....	14
Figure 3-6 DART Login	15

1 Introduction

The Cirries Smart Data solution from Cirries Technologies comprises the DART Packet Sensor and the DART management center (DART). DART Packet Sensor performs packet recording, flow generation, and packet forwarding, while DART provides centralized control and monitoring of the Packet Sensor instances. The Cirries Smart Data solution delivers comprehensive visibility into network usage and performance, traffic packets, traffic flows, network anomalies, and security attacks.

1.1 Advanced API Support

Cirries provides Alert APIs and network analytics data query APIs for DART. If your monitoring needs could benefit from using these APIs, please contact Technical Support for documents describing them.

2 Requirements and Restrictions

The requirements necessary for DART to run successfully include:

- Licenses to operate Cirries DART and DART Packet Sensor are provided with the AWS Marketplace installation.
- An operating environment that meets, at a minimum, the specifications identified in Section 2.1 of this document.
- Synchronization of the DART system date and time with the AWS NTP server at this URL:
- The latest version of the browser, Google Chrome, is preferred. Microsoft Edge, Opera, and Firefox will operate and are supported but may encounter issues.

2.1 Operating Environments

DART software runs in standard Red Hat Enterprise Linux (RHEL) operating environments supported by AWS virtual machines. Features include:

- Browser-based GUI dashboards
- A set of standard screens that are easily configured and expanded
- Customizable data query templates and dashboards
- Comprehensive reporting, visualization, and analysis of element and network operations

2.1.1 AWS Requirements

The software must be installed with the instance type in the table below. It is recommended to use dedicated EBS-optimized instances with provisioned IOPS SSD or throughput-optimized HDD and enhanced networking for best performance.

For a Cirries DART Packet Sensor supporting 1Gbps traffic, select the instance specs as follows:

- CPU - 2 vCPUs or more
- Storage - 50 GB or more
- Memory - 8 GB or Higher

For a Cirries DART Packet Sensor supporting 5 Gbps traffic, select the instance specs as follows:

- CPU - 4 vCPUs or more
- Storage - 50 GB or more
- Memory - 8 GB or Higher

For Cirries DART, select the instance with the specs as follows:

Table 1 Specifications for Cirries DART Instances

Traffic Rate	vCPU	RAM	Total Storage for 3 Days	Total Storage for a Month
1 Gbps	4 vCPU	16 GB RAM	50 GB or more	500 GB
5 Gbps	4 vCPU	16 GB RAM	100 GB or more	1 TB
10 Gbps	8 vCPU	32 GB RAM	200 GB or more	2 TB

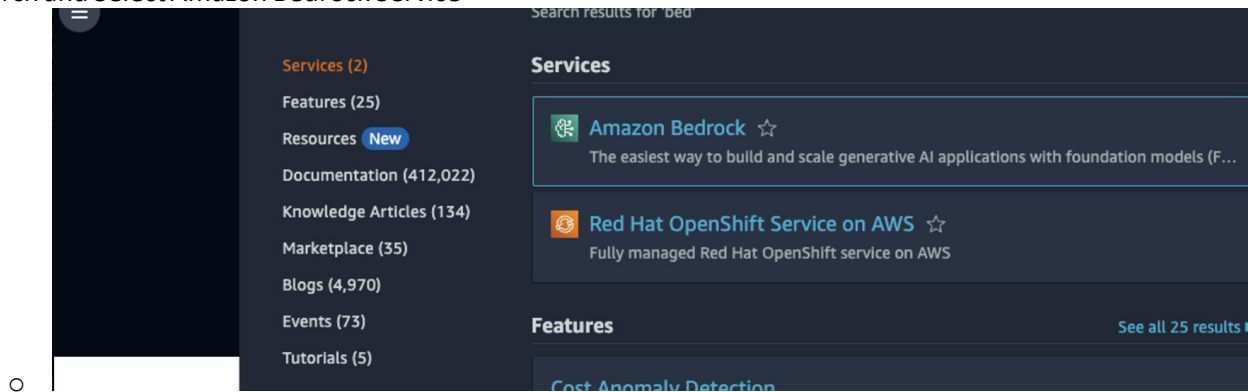
Chapter

3

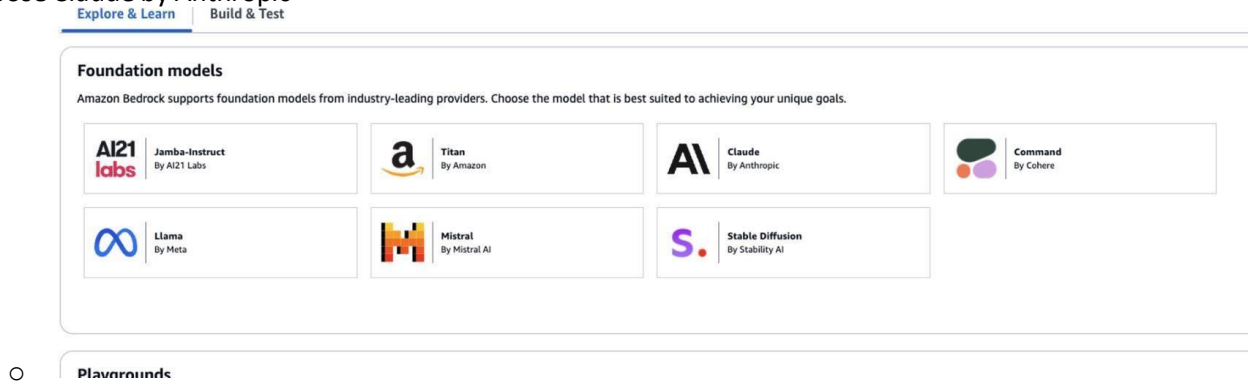
3 DART AI Deployment Summary

Before Deploying Dart AI, user has to manually enable Amazon Bedrock in the AWS account which includes below steps:

- Login to AWS account
- Search and Select Amazon Bedrock Service



- Choose Claude by Anthropic



- Select Model Claude 3.5 Sonnet

The screenshot displays the Amazon Bedrock console interface. On the left, a navigation sidebar lists various sections: Getting started (Overview, Examples, Providers), Foundation models (Base models, Custom models, Imported models), Playgrounds (Chat, Text, Image), Safeguards (Guardrails, Watermark detection), Builder tools (Knowledge bases, Agents, Prompt management, Prompt flows), and Assessment & deployment (Model Evaluation, Provisioned Throughput). Below these are links for Model access, Bedrock Studio, and Settings. The main content area is titled 'Amazon Bedrock' and features 'Anthropic examples' such as 'Create a Side-Scrolling Game' and 'Build a Website'. Below this, the 'Claude models' section highlights 'Claude 3.5 Sonnet' as the latest model. It provides details on supported use cases (customer support, coding, data analysis, visual processing), model attributes (code generation, text generation, complex reasoning and analysis), and model ID (anthropic.claude-3-5-sonnet-20240620-v1:0). Other models like Claude 3 Sonnet, Claude 3 Haiku, and Claude v2.1 are also listed.

- If the Model is not available it can be enable by requesting Model access

Deploying DART and Packet Sensors consists of the following steps:

- Install the DART solution components using the AWS CloudFormation Template
- Obtain the necessary licenses and licensing information
- To connect to the deployed solution over the public internet, assign a public IP address to DART.
- Ensure that DART Packet Sensors are communicating correctly with DART.

3.2 Launching Cirries templates

This section describes how to deploy the Cirries Smart Data solution using the CloudFormation templates and AMIs available on the Cirries site on the AWS Marketplace.

1. Search the Amazon Marketplace for Cirries.
2. Click the Continue to Subscribe button for the Cirries Smart Data Solution for AWS.
3. Accept the Terms and Conditions for the Cirries Smart Data Solution for AWS.
4. Click the Continue to Configuration button.
5. Select the type of deployment you want to perform by choosing from the Fulfillment Options/CloudFormation template.

Cirries Technologies

6. Use the Region dropdown to specify the Availability Zone.
7. Click Continue to Launch to continue.
8. Review the configuration settings in the Launch Page and click Launch when ready to proceed.

3.3 Using the Create Stack Wizard

The above steps will take us to the Create Stack Wizard page with a specified template to create your stack (Figure 3-1 CloudFormation Template to Create Stack with).

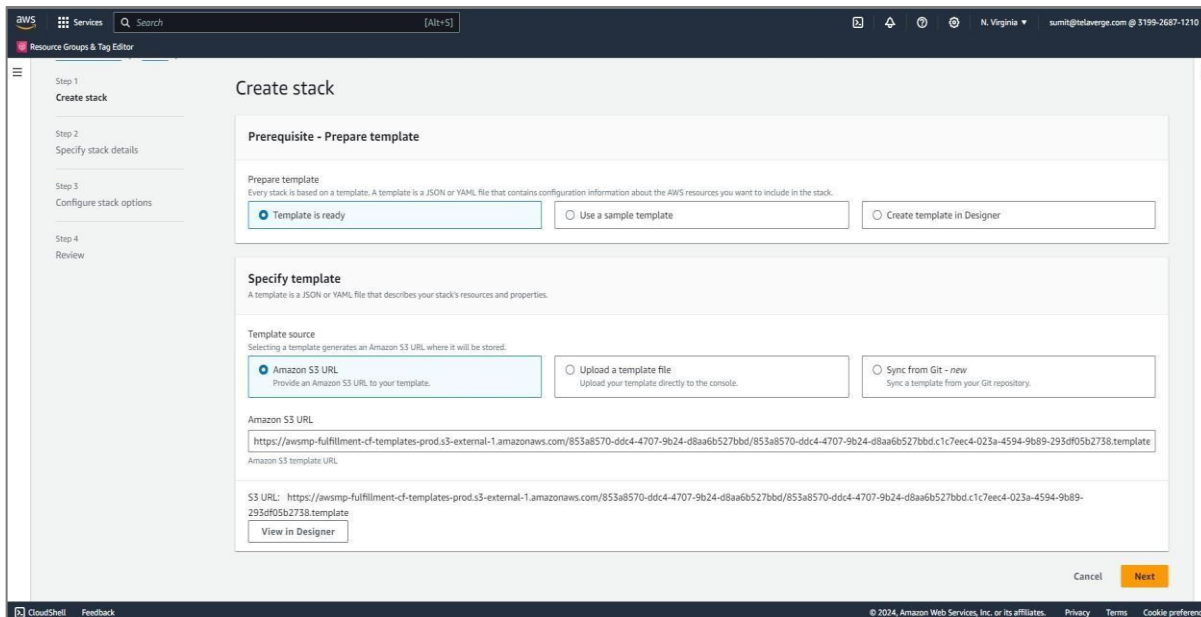


Figure 3-1 CloudFormation Template to Create Stack with

1. Click Next to continue.
2. On the specify stack details page, type a stack name in the Stack name box (Figure 3-2 Filling Out Parameters of the CloudFormation Template).

Note: The stack name is an identifier that helps you find a particular stack from a list of stacks. A stack name can contain only alphanumeric characters (case-sensitive) and hyphens. It must start with an alphabetic character and can't be longer than 128 characters.

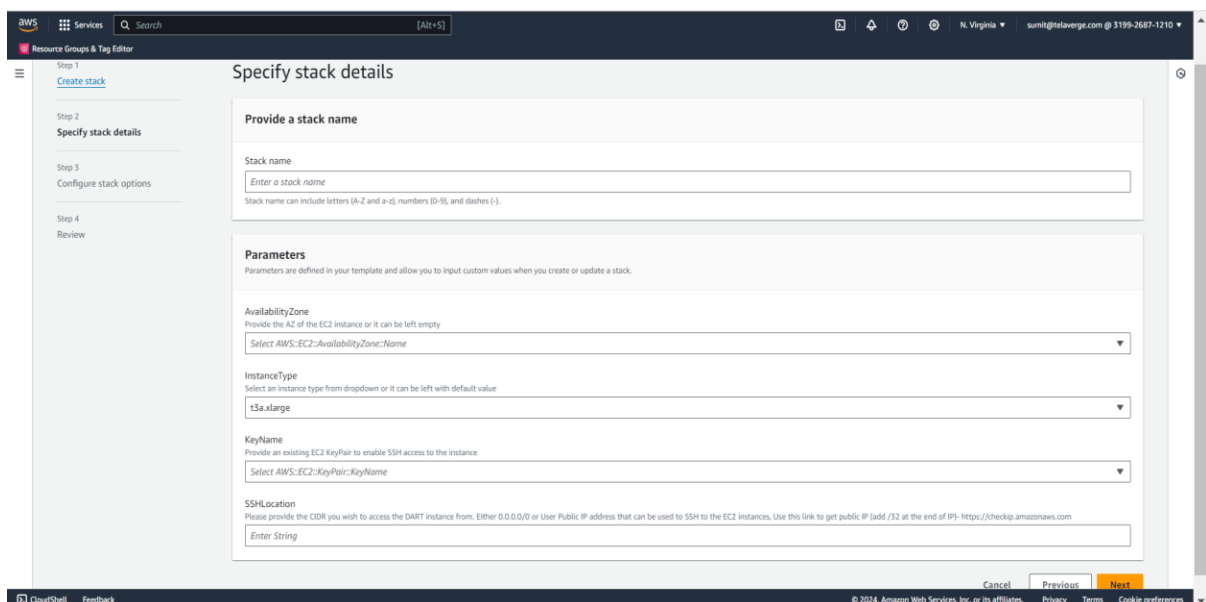


Figure 3-2 Filling Out Parameters of the CloudFormation Template

3. In the Parameters section, specify the parameters defined in the stack template. You can use or change any parameters with default values.
 - AvailabilityZone: Provide the AZ of the EC2 instance from dropdown.
 - InstanceType: Select an instance from the dropdown, or it can be left with a default value.
 - KeyName: Provide an existing EC2 KeyPair to enable SSH access to the instance.
 - EMAIL ID (In the case of PacketSensor Stack): Provide an Email Id for Notification. PacketSensor will use this email id to send notification when traffic reaches more than 1Gbps/5Gbps/10Gbps.

Note: If there is no KeyPair present in your account, please create one and select from the dropdown afterward.

- SSHLocation: Provide the IP address range that can be used to SSH to the EC2 instances. Please provide the CIDR you wish to access the DART instance from. Either 0.0.0.0/0 or a User Public IP address that can be used to SSH to the EC2 instances. Use this link to get public IP (add /32 at the end of IP)- <https://checkip.amazonaws.com>
4. When you are satisfied with the parameter values, choose Next to proceed with setting options for your stack.
 5. After specifying the parameters defined in the template, you can set additional options for your stack. These are all optional and not required. Additional options for your stack include:
 - Tags
 - Permissions
 - Stack Failure Options
 - Stack Policy
 - Rollback Configuration
 - Notification Options
 - Stack Creation Options
 6. To set stack options — on the Configuration Stack Options page of the Create Stack Wizard, you can specify tags and permissions. Use the Advanced Options section to set additional configuration options for your stack.
 7. When you have entered all your stack options, Choose the Next Step to proceed with reviewing your stack.
 8. On the Review page, review the details of your stack. If you need to change any values before launching the stack, choose Edit in the appropriate section to return to the page with the setting you want to change.
 9. Optional: To estimate the cost of your stack, select the Estimate cost link in the Template section. The AWS Pricing Calculator displays values from your stack template and launch settings.
 10. After you review the stack creation settings and the estimated cost, choose Create stack to launch your stack.

CloudFormation displays the Events pane of the Stack details page for your new stack. See Figure 3-3 Stack Events Page. You can view your stack's events, data, or resources from here.

11. Once your stack has been successfully created, you can choose the Outputs tab to view your stack's outputs if you have defined any in the template.

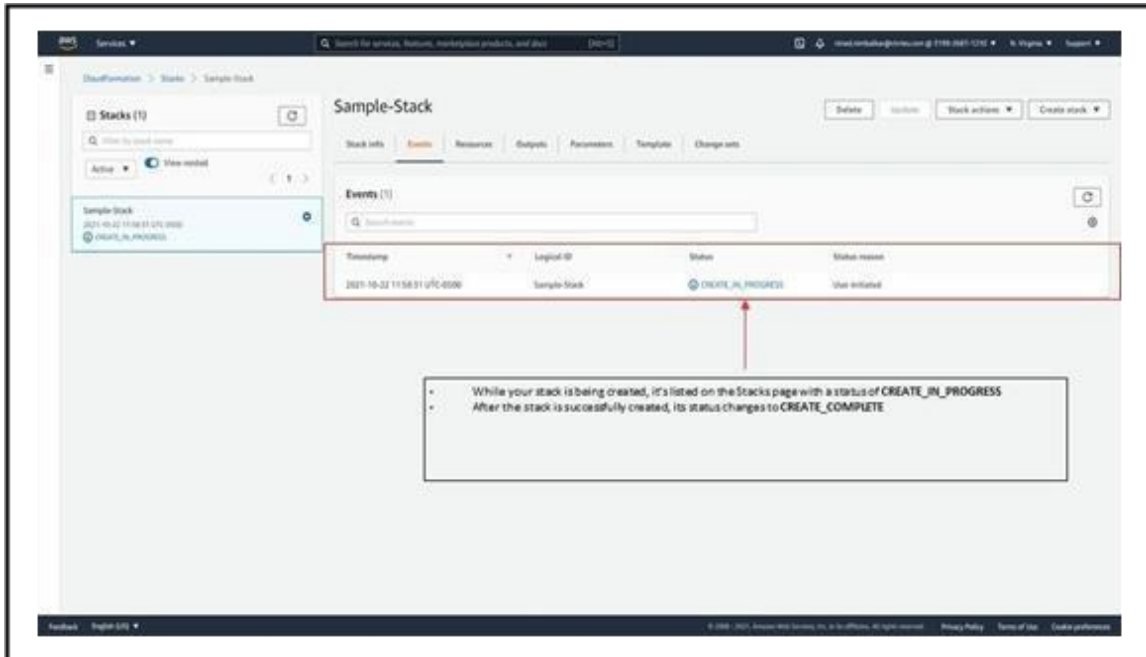


Figure 3-3 Stack Events Page

After a few minutes, you can see the new instance(s) in the EC2 Management Console's Instances list. See Figure 3-4.

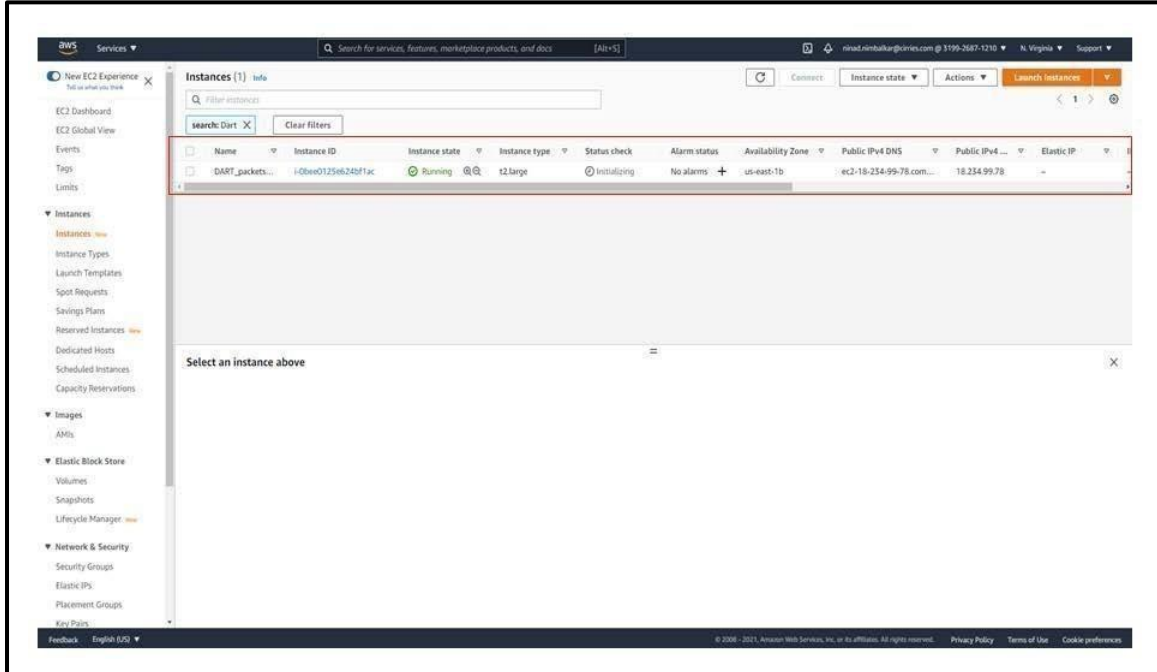


Figure 3-4 New Instances Created in EC2 Management Console

3.4 Access Cirries DART Package from AWS Marketplace:

Security Notice:

This template grants permissions to modify IAM roles and policies. We recommend deploying it in a separate AWS account for security best practices.

WARNING: <product name> uses some broadly permissive IAM policies to set up a variety of resources in the account, as well as cross-account read-only access for debugging via the AWS console. For this reason, please only install this template into a new, empty AWS account.

This product creates administrative IAM roles to enable creation of additional IAM roles to follow least privilege principles for each instance in the cluster so that it is able to orchestrate your HPC environment. We recommend utilizing separate AWS accounts for different groups of users to ensure separation.

AWS Resources required by Cirries DART template:

Reason: This template contains Identity and Access Management (IAM) resources that provide entities access to make changes to your AWS account. These policies grant specific actions (Action) on various AWS resources (Resource) for different services such as EC2, DynamoDB, etc., allowing or denying certain operations based on conditions or resource constraints. This template defines permissions related to EC2 instances, DynamoDB tables, and specific actions like termination based on instance types.

- Identity and Access Management (IAM) Role.
- Virtual Private Cloud (VPC).
- Ec2
- Cloudwatch
- Lambda
- Amazon Bedrock

To access Cirries DART, perform the following:

1. Select the instance (Dart_Instance) created on the successful launch of DART templated, which has deployed Cirries DART in a separate instance.
2. Connect to EC2 instance in SSH as ec2-user
user Eg: `ssh ec2-user@<public_ip_of_Dart_Instance>`
 - It will download required python modules for DART AI and ask for user input to select s3 bucket
 - Select S3 bucket for DART AI to read VPC flow logs
 - Give bucket number with separated comma (eg: 2,11,13)
3. Open the publicip.txt file to get the DART server IP address. For example, Cat publicip.txt
4. Configure the DART server IP address in the DART Packet Sensor: 50.17.179.25

Note: This public IP address is the input for the DART Packet Sensor Collector-IP address of the Dart instance

3.5 Accessing Cirries DART Packet Sensor (1GBPS, 5GBPS & 10GBPS) from AWS Marketplace:

AWS Resources required by Cirries DART Packet Sensor template:

Security Notice:

This template grants permissions to modify IAM roles and policies. We recommend deploying it in a separate AWS account for security best practices.

WARNING: <product name> uses some broadly permissive IAM policies to set up a variety of resources in the account, as well as cross-account read-only access for debugging via the AWS console. For this reason, please only install this template into a new, empty AWS account.

This product creates administrative IAM roles to enable creation of additional IAM roles to follow least privilege principles for each instance in the cluster so that it is able to orchestrate your HPC environment. We recommend utilizing separate AWS accounts for different groups of users to ensure separation.

Reason: This template contains Identity and Access Management (IAM) resources that provide entities access to make changes to your AWS account. These policies grant specific actions (Action) on various AWS resources (Resource) for different services such as EC2, DynamoDB, etc., allowing or denying certain operations based on conditions or resource constraints. This template defines permissions related to EC2 instances, DynamoDB tables, and specific actions like termination based on instance types.

- Identity and Access Management (IAM) Role.
- Virtual Private Cloud (VPC).
- Ec2

- Cloudwatch
- Lambda

To access Cirries DART Packet Sensor, perform the following:

1. Select the instance (PacketSensor_Instance[1G/5G/10G]) created on the successful launch of DART PacketSensor1Gbps or PacketSensor5Gbps template, which has deployed Cirries DART PacketSensor in a separate instance.
2. Connect to the EC2 instance in SSH as a ec2-user user
Eg: `ssh ec2-user@<public_ip_of_PacketSensor_instance>`
3. Enter the IP address obtained from the DART instance. This should be the IP address of the DART instance that you will deploy. For example: 50.17.179.25
4. After this, it will take some time (2-3mins)for configuration for PacketSensor to Dart. Please don't do anything in between.

3.6 AWS Multi-size Provisioning

Cirries offers two (2) AMIs for Packet Sensor, one used for 1 Gbps traffic and a second used for traffic up to 5 Gbps. All the deployed Packet Sensors can be directed to a single DART application. After each Packet Sensor deployment, the customer must direct the Packet Sensor to the DART instance. When PacketSensors are in different VPC in that case we need to connect two VPC's via AWS peering connection.

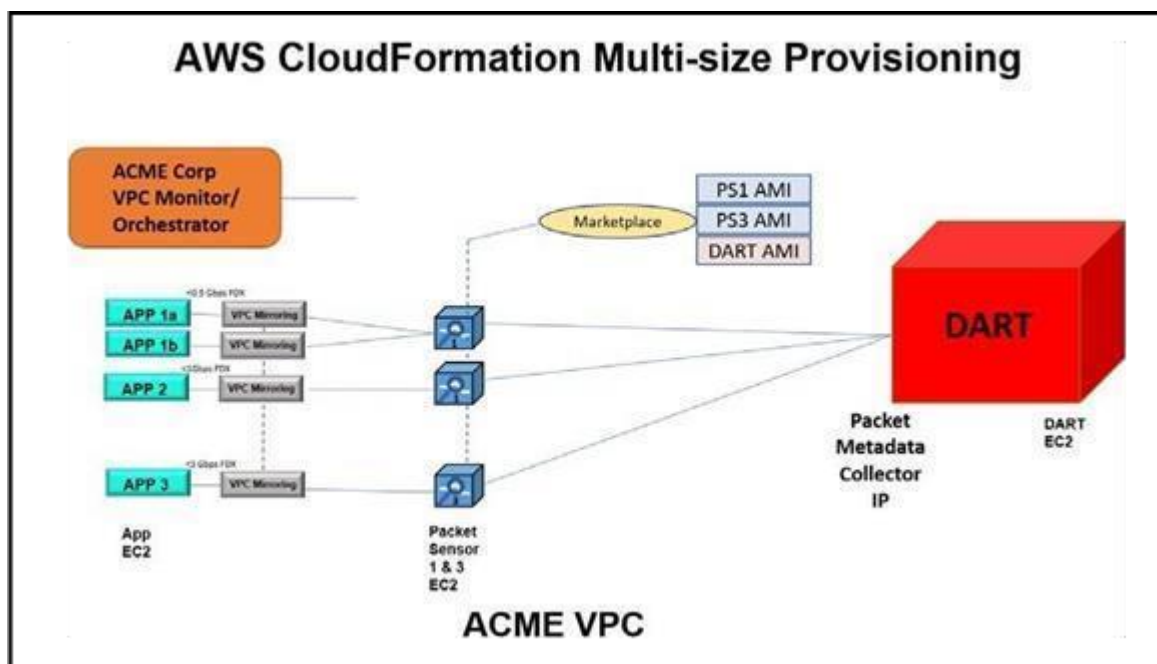


Figure 3-5 AWS Multi-Size Provisioning

3.7 DART Login

Once DART has been initiated, go to `http://<public or privateIP>:3020` to log into the application using the Login screen as illustrated in Figure 3-6 DART Login.

Log-in default values are in the "README" file, located in `/opt/cirries/nm-web/README.txt`.

Login credentials are set in the User Management dashboard. The number of users supported is unlimited.



Figure 3-6 DART Login

DART Application Configuration

The configuration file may be edited to change some aspects of the DART web application. The configuration file is located at `/opt/cirries/nm-web/config.js`. Configuration settings include:

Table 2 Configuration Settings

Configuration	Description
<code>config.logging</code>	The logging configuration for the application consists of configuration for the logging files to be created and the process that will clean up the old log files. Some of the fields inside this configuration should not be modified to avoid application instability.
<code>config.logging.dir</code>	The directory where the logging files will be created.
<code>config.logging.filename</code>	Fixed part of the name that the logging files will have. The second part will be generated using the <code>datePattern</code> defined inside the logging configuration.
<code>config.logging.datePattern</code>	The date pattern is to be appended to the second part of the logging filename.
<code>config.http.port</code>	The port to start the dashboard application. Default in port 3020.

config.influx	The configuration for the communication between the application and the database server. Some of the fields inside this configuration should not be modified to avoid application instability.
config.influx.address	Defines the host and port where the database server is running. By default, the values are 'inhost' and 8086 respectively.
config.influx.timeout.get_data	Defines the time that the application server will wait for the database server response. The default value is 90 seconds. It is recommended not to decrease or increase this value higher than one. (2) minute
config.packetpoint	Not Applicable.
config.packetpoint.middleware	Not Applicable
config.packetpoint.port	Not Applicable
config.packetpoint.timeout	Not Applicable
config.packetpoint.timeout.get_data	Not Applicable.
config.packetpoint.timeout.export_pcap	Not Applicable.
config.packetpoint.max_size_cluster	Not Applicable.
config.packetpoint.cron_job_clean_storage	Not Applicable.
config.packetpoint.delete_old_minutes_storage	Not Applicable
config.emdc	Configure the host and port to communicate between the application and the Element Manager server. By default, the values are 'emhost' and 8595 respectively.

Configuration fields not mentioned in the previous table or marked Not Applicable should not be modified; otherwise, application instability may occur.

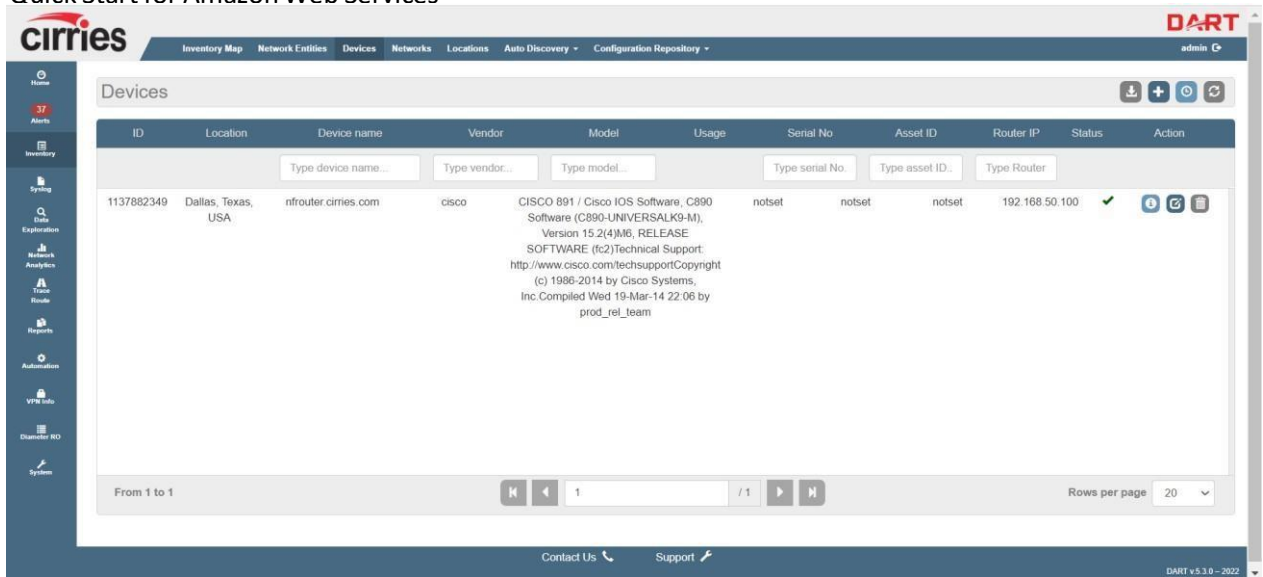
3.8 Traceroute Configuration

In order to use Dart Generative AI feature user need to add device via autodiscovery SNMP scan.

Follow below steps to add device

- Once login go to Inventory -> Devices
- Click on add device button
- Add device manually or using SNMP scan
- To use SNMP scan add IP address, Range, interface, community string, version

Cirries Technologies
DART
Quick Start for Amazon Web Services



- Navigate to System -> PacketSensor settings -> Add relation
- In "Packet Sensor to Router" table select Router IP (host IP's) from dropdown for Packetsensor interface
- In "Trace Route Source table" select Route from (host) dropdown

Router - Interface Relation ✕

Parameter: Packet device
Description: Ingress packet device

New Interface: Add

Packet Sensor to Router			
Interface ID	Interface	Router IP	Actions
2	enp2s0f0	<input type="text" value="10.10.200.2"/> ▾	
3	enp2s0f1	<input type="text" value="10.10.200.3"/> ▾	
1	eno2	<input type="text"/> ▾	

Trace Route Source		
Interface ID	Route From:	Actions
2	<input type="text" value="PC-A3"/> ▾	
3	<input type="text" value="PC-A3"/> ▾	
1	<input type="text" value="PC-A3"/> ▾	

Close

1 AWS VPC application traffic mirroring to the Packet Sensor setup

Traffic Mirroring is an Amazon VPC feature that you can use to copy network traffic from an elastic network interface of Amazon EC2 instances.

The following are the key concepts for Traffic Mirroring:

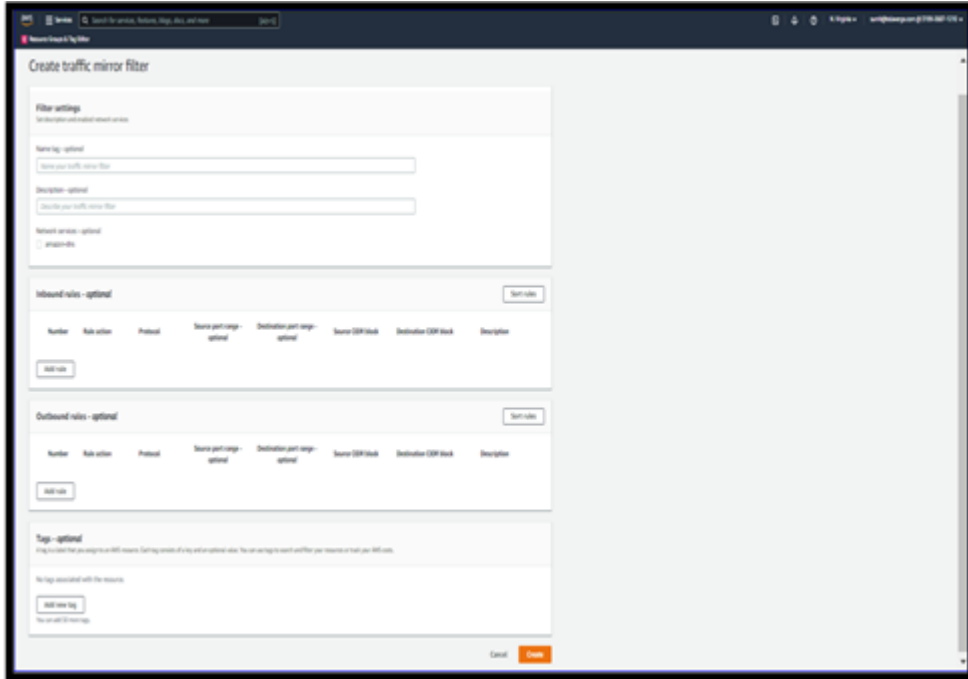
- **Source:** An AWS network resource within a particular VPC that can be used as the traffic source. VPC Traffic Mirroring supports using Elastic Network Interfaces (ENIs) as mirror sources.
- **Target:** The destination for mirrored traffic. This is created by Packet Sensor CFT.
- **Filter:** A set of rules defining the traffic copied in the traffic mirror session.
- **Session:** An entity that describes Traffic mirroring from a source to a target using filters.

1.1 Getting Started with VPC Traffic Mirroring

1. **Create Mirror Source** – You can mirror traffic from any EC2 instance that is powered by the AWS Nitro system (A1, C5, C5d, C5n, I3en, M5, M5a, M5ad, M5d, p3dn.24xlarge, R5, R5a, R5ad, R5d, T3, T3a, and z1d as I write this).

Below are steps (no: 2,3) we need to do under VPC console -> Traffic mirroring.

2. **Create a Mirror Filter** – A specification of the inbound or outbound (with respect to the source) traffic that is to be captured (accepted) or skipped (rejected). The filter can specify a protocol, ranges for the source and destination ports, and CIDR blocks for the source and destination. Rules are numbered and processed in order within the scope of a particular Mirror Session.
 - a. Go to VPC -> Traffic mirroring -> Mirror filters -> Create traffic mirroring



3. Create a Traffic Mirror Session – A connection between a mirror source and target that uses a filter. Sessions are numbered and evaluated in order, and the first match (accept or reject) is used to determine the fate of the packet. A given packet is sent to at most one target.

