



6.2 TREND MICRO™ Deep Discovery Inspector

AWS Deployment Guide

Breakthrough Protection Against APTs and Targeted Attacks



Endpoint Security



Network Security



Protected Cloud



Trend Micro Incorporated reserves the right to make changes to this document and to the product described herein without notice. Before installing and using the product, review the readme files, release notes, and/or the latest version of the applicable documentation, which are available from the Trend Micro website at:

<http://docs.trendmicro.com>

Trend Micro, the Trend Micro t-ball logo, Deep Discovery Inspector, Apex Central, and Trend Micro Control Manager are trademarks or registered trademarks of Trend Micro Incorporated. All other product or company names may be trademarks or registered trademarks of their owners.

Copyright © 2022. Trend Micro Incorporated. All rights reserved.

Document Part No.: APEM69459/211129

Release Date: April 2022

Protected by U.S. Patent No.: 8595840; 8925074; 7707635; 8505094

This documentation introduces the main features of the product and/or provides installation instructions for a production environment. Read through the documentation before installing or using the product.

Detailed information about how to use specific features within the product may be available at the Trend Micro Online Help Center and/or the Trend Micro Knowledge Base.

Trend Micro always seeks to improve its documentation. If you have questions, comments, or suggestions about this or any Trend Micro document, please contact us at docs@trendmicro.com.

Evaluate this documentation on the following site:

<https://www.trendmicro.com/download/documentation/rating.asp>

Privacy and Personal Data Collection Disclosure

Certain features available in Trend Micro products collect and send feedback regarding product usage and detection information to Trend Micro. Some of this data is considered personal in certain jurisdictions and under certain regulations. If you do not want Trend Micro to collect personal data, you must ensure that you disable the related features.

The following link outlines the types of data that Deep Discovery Inspector collects and provides detailed instructions on how to disable the specific features that feedback the information.

<https://success.trendmicro.com/data-collection-disclosure>

Data collected by Trend Micro is subject to the conditions stated in the Trend Micro Privacy Notice:

<https://www.trendmicro.com/privacy>

Table of Contents

Chapter 1: About Deployment on AWS

Specialized Knowledge	1-2
AWS Account	1-2
Cost and Licenses	1-2

Chapter 2: Deployment Planning

Planning the Deployment	2-2
Architecture	2-3
System Requirements	2-3
Deployment Options	2-5
Considerations	2-7
Items to Prepare	2-8

Chapter 3: Deployment

Deployment Overview	3-2
Launching a Virtual Appliance	3-2
Configuring the Description for Network Interfaces	3-12
Deploying a Virtual Appliance as a Traffic Mirror Target	3-14
Deploying a Virtual Appliance Behind an NLB	3-22

Chapter 4: Deployment Testing and Troubleshooting

Checkpoints	4-2
Testing the Deployment	4-7
Troubleshooting the Deployment	4-7

Frequently Asked Questions	4-9
What are the changes on the Deep Discovery Inspector virtual appliance on AWS?	4-9
Does the Deep Discovery Inspector virtual appliance support AWS EC2 auto scaling?	4-14
Does Deep Discovery Inspector support creating an Amazon Machine Image (AMI) from an EC2 instance of the Deep Discovery Inspector virtual appliance?	4-14
Does Deep Discovery Inspector support creating an Elastic Block Store (EBS) snapshot from an EC2 instance of the Deep Discovery Inspector virtual appliance?	4-15
Does Deep Discovery Inspector Support AWS Backup Service?	4-16
What are the IAM policies needed to deploy Deep Discovery Inspector's virtual appliance on AWS?	4-16

Index

Index	IN-1
-------------	------

Preface

Preface

Learn more about the following topics:

- *Documentation on page iv*
- *Audience on page v*
- *Document Conventions on page v*

Documentation

The documentation set for Deep Discovery Inspector includes the following:

TABLE 1. Product Documentation

DOCUMENT	DESCRIPTION
Administrator's Guide	The Administrator's Guide contains detailed instructions on how to configure and manage Deep Discovery Inspector, and explanations on Deep Discovery Inspector concepts and features.
AWS Deployment Guide	The AWS Deployment Guide contains information about requirements and procedures for planning deployment, deploying, and troubleshooting Deep Discovery Inspector deployment on AWS.
Inline (LAN bypass) Network Interface Card Installation Guide	The Inline (LAN bypass) Network Interface Card Installation Guide contains information about requirements and procedures for installing an additional bypass network interface card on supported Deep Discovery Inspector appliances.
Installation and Deployment Guide	The Installation and Deployment Guide contains information about requirements and procedures for planning deployment, installing Deep Discovery Inspector, and using the Preconfiguration Console to set initial configurations and perform system tasks.
Syslog Content Mapping Guide	The Syslog Content Mapping Guide provides information about log management standards and syntaxes for implementing syslog events in Deep Discovery Inspector.
Quick Start Card	The Quick Start Card provides user-friendly instructions on connecting Deep Discovery Inspector to your network and on performing the initial configuration.
Readme	The Readme contains late-breaking product information that is not found in the online or printed documentation. Topics include a description of new features, known issues, and product release history.

DOCUMENT	DESCRIPTION
Online Help	Web-based documentation that is accessible from the Deep Discovery Inspector management console. The Online Help contains explanations of Deep Discovery Inspector components and features, as well as procedures needed to configure Deep Discovery Inspector.
Support Portal	The Support Portal is an online database of problem-solving and troubleshooting information. It provides the latest information about known product issues. To access the Support Portal, go to the following website: https://success.trendmicro.com

View and download product documentation from the Trend Micro Online Help Center:

<https://docs.trendmicro.com/en-us/home.aspx>

Audience

The Deep Discovery Inspector documentation is written for IT administrators and security analysts. The documentation assumes that the reader has an in-depth knowledge of networking and information security, including the following topics:





- Network topologies
- Database management
- Antivirus and content security protection

The documentation does not assume the reader has any knowledge of sandbox environments or threat event correlation.

Document Conventions

The documentation uses the following conventions:

TABLE 2. Document Conventions

CONVENTION	DESCRIPTION
UPPER CASE	Acronyms, abbreviations, and names of certain commands and keys on the keyboard
Bold	Menus and menu commands, command buttons, tabs, and options
<i>Italics</i>	References to other documents
Monospace	Sample command lines, program code, web URLs, file names, and program output
Navigation > Path	The navigation path to reach a particular screen For example, File > Save means, click File and then click Save on the interface
 Note	Configuration notes
 Tip	Recommendations or suggestions
 Important	Information regarding required or default configuration settings and product limitations
 WARNING!	Critical actions and configuration options

Chapter 1

About Deployment on AWS

This guide provides additional information that enables you to evolve from an on-premises Deep Discovery Inspector appliance to a Deep Discovery Inspector appliance on AWS. For more details about the Deep Discovery Inspector features and functions, see the *Deep Discovery Inspector Administrator's Guide* on <https://docs.trendmicro.com/en-us/enterprise/deep-discovery-inspector.aspx>.

Specialized Knowledge

This guide assumes familiarity with networking basics. This guide also requires a moderate level of familiarity with AWS. If you are new to AWS, visit the *Getting Started Resource Center* (<https://aws.amazon.com/getting-started/>) and *AWS Training and Certification* (<https://aws.amazon.com/training/>). These sites provide materials for learning how to design, deploy, and operate your infrastructure and applications on AWS.

AWS Account

If you do not already have an AWS account, create one at <https://aws.amazon.com> by following the on-screen instructions. Part of the sign-up process involves receiving a phone call and entering a PIN using the phone keypad.

AWS automatically signs up your account for all AWS services. You are charged only for the services you use.

Cost and Licenses

In order to access and use the AMI version of the Deep Discovery Inspector virtual appliance, you must already have and continually maintain an active AWS Account on the AWS Marketplace and you are responsible for purchasing and maintaining through such AWS Account, your use of the Amazon Web Service platform/infrastructure that is required for your deployment of a Deep Discovery Inspector virtual appliance.

The Deep Discovery Inspector virtual appliance is offered as an AMI in the AWS Marketplace. Access to the AMI can be obtained by searching the AWS Marketplace console.

Chapter 2

Deployment Planning

Planning the Deployment

The following steps provide an overview for planning the deployment of Deep Discovery Inspector virtual appliances in an AWS environment.

Procedure

1. Review the architecture.
For details, see [Architecture on page 2-3](#).
 2. Review the system requirements.
For details, see [System Requirements on page 2-3](#).
 3. Choose a deployment option to integrate with Amazon VPC Traffic Mirroring.
For details, see [Deployment Options on page 2-5](#).
 4. Prepare items before deploying Deep Discovery Inspector.
For details, see [Items to Prepare on page 2-8](#).
 5. Deploy the Deep Discovery Inspector virtual appliance.
For details, see [Deployment on page 3-1](#).
 6. Access the Deep Discovery Inspector virtual appliance management console.
For details, see the [Deep Discovery Inspector Administrator's Guide](#).
-

Architecture

The Deep Discovery Inspector virtual appliance supports deployment on an AWS EC2 environment and can scan as well as analyze mirrored packets from an AWS VPC traffic mirror.

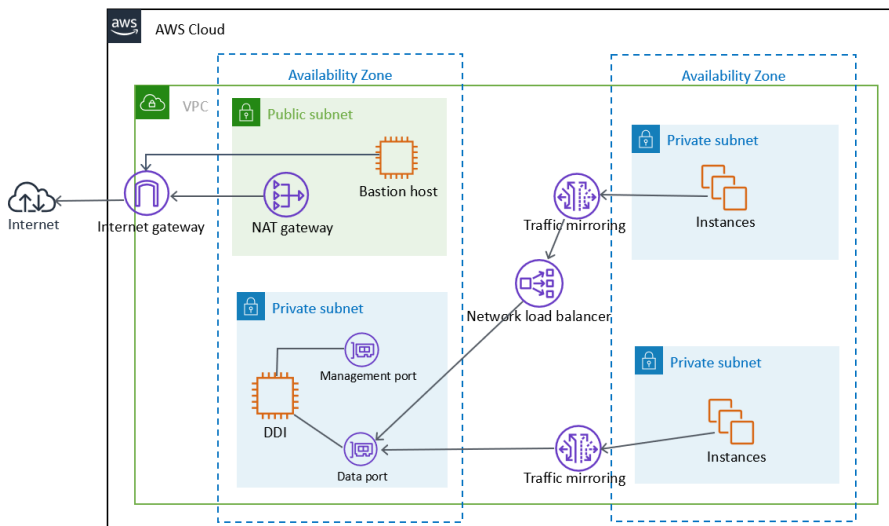


FIGURE 2-1. Deployment Architecture

System Requirements

Trend Micro recommends the following minimum specifications based on your licensed model's throughput.



Note

When using a Deep Discovery Inspector virtual appliance on AWS with Virtual Analyzer, only external Virtual Analyzers and Sandbox as a Service are supported.

TABLE 2-1. System Requirements

THROUGHPUT (MBPS)	AWS VCPU	AWS MEMORY (GiB)	AWS STORAGE (GiB)	AWS ENI (ELASTIC NETWORK INTERFACES)	RECOMMENDED AWS EC2 INSTANCE TYPE
250	8	32	500	2	<ul style="list-style-type: none">• t3.2xlarge• t3a.2xlarge• m5.2xlarge• m5a.2xlarge
500	8	32	500	2	<ul style="list-style-type: none">• t3.2xlarge• t3a.2xlarge• m5.2xlarge• m5a.2xlarge
1000	16	64	1000	2	<ul style="list-style-type: none">• m5.4xlarge• m5a.4xlarge

**Note**

T3 and T3a instances launch as unlimited mode by default. For more details about using unlimited mode or standard mode on the instance types, see <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances.html>.

For details about AWS EC2 instance types, see <https://aws.amazon.com/ec2/instance-types/>.

You can use non-recommended instance types as long as the instance type meets the minimum system requirements.

Deployment Options

By integrating with the Amazon VPC Traffic Mirroring feature, the Deep Discovery Inspector virtual appliance can provide a network security solution via two deployment options:

- **Option 1: Deploy the Deep Discovery Inspector virtual appliance as a traffic mirror target**

Network traffic is mirrored from an ENI (Elastic Network Interfaces) mirror source to a data port of the Deep Discovery Inspector virtual appliance. This option depends on the settings of traffic mirror filter as shown in the figure below.

**Note**

If the Deep Discovery Inspector virtual appliance is attached to more than 1 data port, you can set each data port as traffic mirror target.

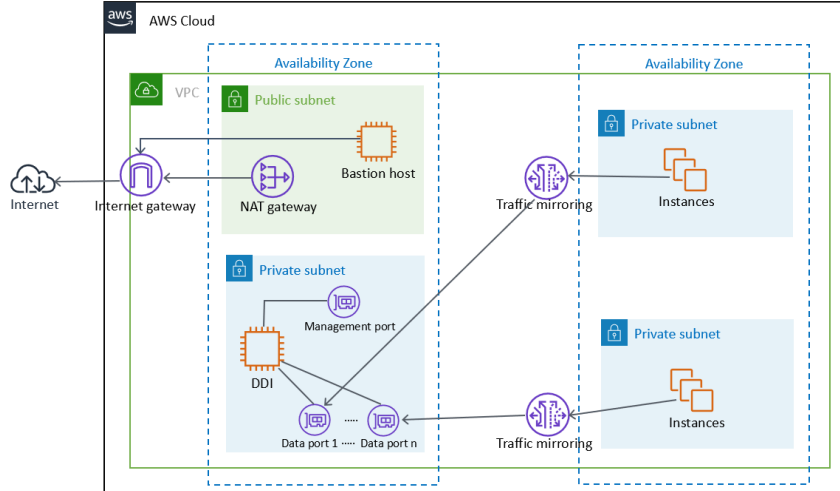


FIGURE 2-2. Option 1: Deploy the Deep Discovery Inspector virtual appliance as a traffic mirror target

- **Option 2: Deploy the Deep Discovery Inspector virtual appliance behind the NLB**

Deploy the Deep Discovery Inspector virtual appliance in the target group behind the NLB (Network Load Balancer). Network traffic is mirrored to the NLB and the NLB forwards traffic to health instances belonging to the target group as shown in the figure below.

 **Note**

The NLB only forwards the mirrored traffic to data port 1 of the Deep Discovery Inspector virtual appliance.




FIGURE 2-3. Option 2: Deploy the Deep Discovery Inspector virtual appliance behind the NLB

Considerations

The quota limitation enforced by AWS traffic mirrors has the following limitations for the deployment options:

- Maximum number of mirror sources per a non-dedicated instance type as target: 10
- Maximum number of mirror sources per a dedicated instance type as target: 100



Note

For more details about the limitation, see <https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-considerations.html>.

You are not limited to a particular deployment option. If you deploy a Deep Discovery Inspector virtual appliance as a traffic mirror target for early validation and later change to deploy a Deep Discovery Inspector virtual

appliance behind an NLB, then it is unnecessary to re-launch a new Deep Discovery Inspector virtual appliance after changing. In addition, advanced deployments can incorporate both deployment options at the same time in the VPC environment.

Items to Prepare

- **Deep Discovery Inspector AMI**

AMI of the Deep Discovery Inspector virtual appliance from the AWS Marketplace

- **Deep Discovery Inspector Activation Code**

Activation Code for the Deep Discovery Inspector virtual appliance

- **AWS VPC and subnet**

A VPC configured with public and private subnets, according to AWS best practices, to provide you with your own virtual network on AWS.



For details about creating a VPC and subnet, see <https://docs.aws.amazon.com/vpc/latest/userguide/working-with-vpcs.html>.

Public subnets and:

- Managed NAT gateways to allow outbound internet access for the Deep Discovery Inspector virtual appliance in the private subnets.



For details about creating a NAT gateway, see <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>.

Private subnets and:

- Management port and Data port of the Deep Discovery Inspector virtual appliance which can be in the same subnet or different subnets in your VPC.
- **AWS VPC Traffic Mirror**

Traffic Mirroring is an AWS VPC feature that you can use to copy network traffic from an elastic network interface (ENI) of Amazon EC2 instances. The security and monitoring appliances can be deployed as individual instances, or as a fleet of instances behind a Network Load Balancer (NLB) with a UDP listener.

**Note**

For details, see <https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-how-it-works.html>.

- One or more instances that create some network connections. The instances act as the traffic mirror sources.

**Important**

There is a limit on the size of the mirrored packet, and packets larger than 8947 bytes are always truncated. Ensure that your traffic mirror source's MTU size is set to equal or less than 8947 bytes. To check and set MTU on your AWS EC2 instance which you want to set as traffic mirror source, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/network_mtu.html#set_mtu and https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/network_mtu.html#set_mtu_windows.

- Only instances powered by the AWS Nitro system can be traffic mirror sources. For details, see <https://aws.amazon.com/blogs/aws/new-vpc-traffic-mirroring/>.
- (Optional) A Network Load Balancer, with the settings configured properly:
 - Target group

- Traffic mirror, with the settings configured properly:
 - Traffic mirror filter
 - Traffic mirror target
 - Traffic mirror session

**Note**

For details about creating a traffic mirror target and filter, and then using those resources to create a session, see <https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-getting-started.html>.

- **AWS EC2 Security Group**

INBOUND/ OUTBOUND RULE	TYPE	PROTOCOL	PORT	SOURCE	DESCRIPTION
Inbound	HTTPS	TCP	443	CIDR that can reach your instance	For accessing the Deep Discovery Inspector virtual appliance management console
Inbound	SSH	TCP	22	CIDR that can reach your instance	For accessing the Deep Discovery Inspector virtual appliance pre-configuration console

INBOUND/ OUTBOUND RULE	TYPE	PROTOCOL	PORT	SOURCE	DESCRIPTION
Inbound	Custom UDP	UDP	4789	CIDR of your mirror source or the NLB	For VXLAN traffic required by the AWS traffic mirror
Inbound	Custom TCP	TCP	14789	CIDR of NLB	(Optional) Implemented by the Deep Discovery Inspector virtual appliance for answering the NLB health check.



Note

Outbound Rules in the default security group should allow all traffic. The Deep Discovery Inspector virtual appliance works well with the default outbound rules. The following exceptions may apply:

- For some organizations, whose policies may need more specific protocols and port numbers, see *Chapter 2: About Your System* in the *Deep Discovery Inspector Installation and Deployment Guide*.
- For some organizations, whose infrastructure may need an outbound proxy with domains allowed to access the internet, see https://docs.trendmicro.com/all/ent/ddi/v5.7/en-us/ddi_5.7_olh/access_trend_service.html for detailed addresses.

Chapter 3

Deployment

Deployment Overview

The following is an overview of the steps required to deploy a Deep Discovery Inspector virtual appliance and a VPC traffic mirror in your AWS environment.


1. Launch a Deep Discovery Inspector virtual appliance.
For details, see [Launching a Virtual Appliance on page 3-2](#).
2. (Optional) Configure the description for the virtual appliance network interfaces.
For details, see [Configuring the Description for Network Interfaces on page 3-12](#).
3. Choose one of the following options to deploy the AWS VPC traffic mirror.
 - Deploy a virtual appliance as a traffic mirror target
For details, see [Deploying a Virtual Appliance as a Traffic Mirror Target on page 3-14](#).
 - Deploy a virtual appliance behind an NLB
For details, see [Deploying a Virtual Appliance Behind an NLB on page 3-22](#).

Launching a Virtual Appliance


Procedure

1. Initiate the instance launch.
 - a. Open the **Amazon EC2** console at <https://console.aws.amazon.com/ec2/>.
 - b. In the navigation bar at the top of the screen, select a Region for the instance that meets your needs.

- c. From the **Amazon EC2** console dashboard, select **Launch instance**.



2. Choose the AMI for Deep Discovery Inspector.
 - a. On the **Choose an Amazon Machine Image (AMI)** screen, select **AWS Marketplace** in the left pane.
 - b. In the search box, search for **Trend Micro Deep Discovery Inspector**.



- c. After the search results appear, click **Select** for **Trend Micro Deep Discovery Inspector <version>**.

3. Choose an Instance Type.

- a. On the **Choose an Instance Type** screen, choose an instance type that meets the minimum specifications based on your licensed model's throughput.

For details, see [System Requirements on page 2-3](#).

- b. Choose **Next: Configure Instance Details** to configure your instance further.


Step 2: Choose an Instance Type

Instance Type	vCPUs	Memory (GiB)	Storage	EBS	Network	Price
General purpose r5n.24xlarge	96	768	EBS only	Yes	100 Gigabit	Yes
General purpose r5n.24xlarge	96	768	4 x 900 (SSD)	Yes	100 Gigabit	Yes
General purpose m5.large	2	8	EBS only	Yes	Up to 10 Gigabit	Yes
General purpose m5.xlarge	4	16	EBS only	Yes	Up to 10 Gigabit	Yes
General purpose m5.2xlarge	8	32	EBS only	Yes	Up to 10 Gigabit	Yes
General purpose m5.4xlarge	16	64	EBS only	Yes	Up to 10 Gigabit	Yes
General purpose m5.8xlarge	32	128	EBS only	Yes	10 Gigabit	Yes
General purpose m5.12xlarge	48	192	EBS only	Yes	10 Gigabit	Yes
General purpose m5.16xlarge	64	256	EBS only	Yes	20 Gigabit	Yes
General purpose m5.24xlarge	96	384	EBS only	Yes	25 Gigabit	Yes
General purpose m5.metal	96	384	EBS only	Yes	25 Gigabit	Yes
General purpose m4.large	2	8	EBS only	Yes	Moderate	Yes

Buttons: Cancel, Previous, Review and Launch, Next: Configure Instance Details

4. Configure the Instance Details.

- a. On the **Configure Instance Details** screen, change the following settings.
 - **Network:** Select the VPC.
 - **Subnet:** Select the subnet into which to launch your instance. Select a subnet that is planned for the data port subnet.
 - **Auto-assign Public IP:** Select **Disable**. Trend Micro recommends that you deploy the Deep Discovery Inspector virtual appliance behind an AWS NAT gateway.



- **Network interfaces:** Add a secondary network interface for the Deep Discovery Inspector virtual appliance instance by choosing **Add Device**.



Important

The management port for Deep Discovery Inspector on-premises is fixed at the first NIC port (eth0 in Deep Discovery Inspector). In order to adapt into the AWS environment, the Deep Discovery Inspector virtual appliance has swapped port enumeration for the management port to port 1 (eth1) and data port to port 0 (eth0).

- Device eth0:
 - **Subnet:** The subnet has been configured in a previous step.
 - **Primary IP:** Type a private IPv4 address from the range of your subnet, or leave **Auto-assign** to let AWS choose a private IPv4 address for you.
- Device eth1:
 - **Subnet:** Select a subnet that is planned for the management port subnet.
 - **Primary IP:** Type a private IPv4 address from the range of your subnet, or leave **Auto-assign** to let AWS choose a private IPv4 address for you.

- **IPv6 IPs: (Optional)** Click **Add IP** and type an IPv6 address from the range of the subnet, or leave **Auto-assign** to let AWS choose an IPv6 address for you.

- Click **Next: Add Storage** to specify the root volume size of your instance
5. Add Storage.
- Specify the following settings on the **Add Storage** screen.
 - **Size:** The storage size should meet the minimum specifications based on your licensed model's throughput.


For details, see [System Requirements on page 2-3](#).



Note

To enlarge the storage size, specify the storage size of the **Volume Type: Root**. The Deep Discovery Inspector virtual appliance only partitions the storage when the **Volume Type** is **Root**. The extra storage will not be used.

- **Volume Type:** Use the default value, **General Purpose SSD (gp2)**.




b. Click **Next: Add Tags** to add some custom tags.

6. Add Tags.

a. On the **Add Tags** screen, specify tags by providing the key and value combinations.

For example, for **Key** type `Name` and for **Value** type `vDDI-demo`.

b. Click **Next: Configure Security Group**.



7. Configure Security Group.

a. On the **Configure Security Group** screen, use a security group to define firewall rules for the Deep Discovery Inspector virtual appliance instance.

- To use existing security group, select **Select an existing security group**, and select your security group.
- To create a new security group, select **Create a new security group**.

- b. Verify that your selected security group contains the following rules:

TABLE 3-1. Inbound Rules

TYPE	PROTOCOL	PORT RANGE	SOURCE	REASON
SSH	TCP	22	CIDR that can reach your instance	For accessing Deep Discovery Inspector virtual appliance Pre-Configuration console
HTTPS	TCP	443	CIDR that can reach your instance	For accessing Deep Discovery Inspector virtual appliance management console
Custom UDP	UDP	4789	CIDR of your mirror source or the NLB	For VXLAN traffic required by AWS traffic mirror
Custom TCP	TCP	14789	CIDR of NLB	Implemented by the Deep Discovery Inspector virtual appliance for answering NLB health check

**Note**

Outbound Rules: Rules in default security group allow all traffic. The Deep Discovery Inspector virtual appliance works well with default outbound rules. The following exceptions may occur:

- For some organizations, whose policies may need more specific protocols and port numbers, see *Chapter 2: About Your System > Ports Used by the Appliance* in the *Deep Discovery Inspector Installation and Deployment Guide*.
- For some organizations, whose infrastructures may need an outbound proxy with domains allowed to access the internet, see https://docs.trendmicro.com/all/ent/ddi/v5.7/en-us/ddi_5.7_olh/access_trend_service.html for detailed addresses.

- c. Click **Review and Launch**.
8. Review Instance Launch and select key pair.
 - a. On the **Review Instance Launch** screen, check the details of your instance, and make any necessary changes by choosing the appropriate **Edit** link.
 - b. Click **Launch**.
 - c. In the **Select an existing key pair or create a new key pair** dialog box, select **Proceed without a key pair**.
 - d. To launch your instance, select the acknowledgment check box, then click **Launch Instances**.

Select an existing key pair or create a new key pair ✕

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about [removing existing key pairs from a public AMI](#).

Proceed without a key pair

acknowledge that I will not be able to connect to this instance unless I already know the password built into this AMI.


[Cancel](#) [Launch Instances](#)

9. Wait for the Deep Discovery Inspector virtual appliance to become ready.

**Note**

The Deep Discovery Inspector virtual appliance takes about 15 minutes to become ready.


- a. View the Deep Discovery Inspector installation progress by using the following steps:
 - i. In the left navigation page, click **Instances**.
 - ii. Select the Deep Discovery Inspector virtual appliance instance.
 - iii. Select **Actions** > **Instance Settings** > **Get Instance Screenshot**.




For more details, see <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/launching-instance.html>.

- b. When the Deep Discovery Inspector virtual appliance pre-configuration console appears, then Deep Discovery Inspector is ready.

- b. Copy the value of **Instance ID**.



3. In the navigation pane, select **Network Interfaces** and find the network interfaces of the Deep Discovery Inspector virtual appliance by searching for the instance ID.




4. Select the network interfaces of the Deep Discovery Inspector virtual appliance and then select **Actions > Change Description**.
5. In the **Change Description** dialog box, type a description for the network interface, select **Save** and then perform the following steps:
- Set description of eth0 to **Data port 1**.
 - Set description of eth1 to **Management port**.

**Tip**

To view which interface is eth0 and which interface is eth1, perform the following:

- a. Select the interface.
- b. Click **Actions** > **Manage IP addresses**.
The port label appears.
- c. Click **Cancel** to return to the previous screen.



Deploying a Virtual Appliance as a Traffic Mirror Target


Procedure

1. Configure the traffic mirror filter.

For details, see <https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-filters.html>.

- a. Open the **Amazon VPC** console at <https://console.aws.amazon.com/vpc/>.
- b. In the **Region** selector, select the AWS Region that you used when you created the VPCs.
- c. On the navigation pane, go to **Traffic Mirroring** > **Mirror Filters**.


- d. Select **Create traffic mirror filter**.
- e. For **Name tag**, type a name for the traffic mirror filter.
For example, type `demo-traffic-mirror-filter`.
- f. (Optional) For **Description**, type a description for the traffic mirror filter.
For example, type `demo-traffic-mirror-filter`.
- g. Select **amazon-dns**.



- h. Add inbound rules. Select **Inbound rules > Add > rule**, and then specify the following information about the traffic mirror source inbound traffic:
 - **Rule number**: Type a priority to assign to the rule.
 - **Rule action**: Select the action to take for the packet.
 - **Protocol**: Select the L4 protocol to assign to the rule.
 - (Optional) **Source port range**: Type the source port range.
 - (Optional) **Destination port range**: Type the destination port range.
 - **Source CIDR block**: Type a source CIDR block.
 - **Destination CIDR block**: Type a destination CIDR block.
 - (Optional) **Description**: Type a description for the rule.

The following is an example of the values.


- Rule number: Use the default number
- Rule action: Select **accept**.
- Protocol: Select **All protocols**.
- Source CIDR block: Type **0.0.0.0/0**.
- Destination CIDR block: Type **0.0.0.0/0**.
- Description: Type **mirror all inbound traffic**.




- Add outbound rules. Select **Outbound rules > Add > rule**, and then specify the following information about the traffic mirror source outbound traffic:
 - **Rule number:** Type a priority to assign to the rule.
 - **Rule action:** Select the action to take for the packet.
 - **Protocol:** Select the L4 protocol to assign to the rule.
 - (Optional) **Source port range:** Type the source port range.
 - (Optional) **Destination port range:** Type the destination port range.
 - **Source CIDR block:** Type a source CIDR block.
 - **Destination CIDR block:** Type a destination CIDR block.
 - (Optional) **Description:** Type a description for the rule.

The following is an example of the values.

- Rule number: Use the default number
- Rule action: Select **accept**.
- Protocol: Select **All protocols**.
- Source CIDR block: Type **0.0.0.0/0**.
- Destination CIDR block: Type **0.0.0.0/0**.
- Description: Type **mirror all outbound traffic**.




- Repeat the previous step for each inbound rule and outbound rule that you want to add.
- Click **Create**.



- Configure the traffic mirror target.
 - On the navigation pane, select **Traffic Mirroring** > **Mirror Targets**.
 - Select **Create Traffic Mirror Target**.

- c. For **Name tag**, type a name for the traffic mirror target.
For example, type `demo-traffic-mirror-target`.
- d. (Optional) For **Description**, type a description for the traffic mirror target.
For example, type `demo-traffic-mirror-target`.




The screenshot shows the AWS console interface for creating a traffic mirror target. The page title is "Create traffic mirror target". Under the "Target settings" section, there are two input fields: "Name tag - optional" and "Description - optional". Both fields contain the text "demo-traffic-mirror-target".

- e. For **Target type**, select **Network Interface**.
- f. For **Target**, select the Deep Discovery Inspector virtual appliance's eth0 (the data port that is connected to your subnet) as the traffic mirror target.

**Note**


You can select any other data port that you have attached on the Deep Discovery Inspector virtual appliance, such as eth2, or eth3.

Do not select the eth1 port that is used as the management port for the Deep Discovery Inspector virtual appliance.




The screenshot shows the AWS console interface for choosing a target. The page title is "Choose target". Under the "Target type" section, there is a dropdown menu set to "Network Interface". Below that, there is a "Target" field with a search icon, a text input containing "eni-...", and a clear button.

- g. Click **Create**.



3. Repeat the previous step to create a traffic mirror target for each Deep Discovery Inspector virtual appliance in your AWS environment.
4. Configure the traffic mirror session.
 - a. On the navigation pane, select **Traffic Mirroring** > **Mirror Sessions**.
 - b. Select **Create traffic mirror session**.
 - c. For **Name tag**, type a name for the traffic mirror session.
For example, type `demo-traffic-mirror-session`.
 - d. (Optional) For **Description**, type a description for the traffic mirror session.
For example, type `demo-traffic-mirror-session`.
 - e. For **Mirror source**, select the network interface of the instance that you want to monitor.
 - f. For **Mirror target**, select the traffic mirror target.
For example, select `demo-traffic-mirror-target`.



g. Under **Additional settings**, perform the following:

- For **Session number**, type the session number **1**.

The session number determines the order that the traffic mirror sessions are evaluated in both of the following situations:

- When an interface is used by multiple sessions
- When an interface is used by different traffic mirror targets and traffic mirror filters.

Traffic is only mirrored one time. Use **1** for the highest priority. Valid values are 1-32766.

- (Optional) For **VNI**, type the VXLAN ID to use for the traffic mirror session.

For details, see <https://tools.ietf.org/html/rfc7348>.

If you do not specify a value, AWS assigns a random, unused number.

- (Optional) For **Packet Length**, type the number of bytes in each packet to mirror.

If you do not want to mirror the entire packet, set **Packet Length** to the number of bytes in each packet to mirror. For example, if you set this value to 100, the first 100 bytes after the

VXLAN header that meet the filter criteria are copied to the target.

To mirror the entire packet, do not enter a value in this field.

- For **Filter**, select the traffic mirror filter that determines what traffic gets mirrored.

For example, select **demo-traffic-mirror-filter**.

- (Optional) Under the **Tags** section, add or remove a tag.

The following are example settings.

- For **Session number**, type the session number **1**.
- For **VNI**, leave the value empty. AWS will assign a random number.
- For **Packet Length**, leave the value empty. AWS will mirror the entire packet.
- For **Filter**, select **demo-traffic-mirror-filter**.

The screenshot shows the 'Additional settings' section in the AWS console. It includes the following fields and values:

- Session number:** 1 (with a note: 'The order sessions for the same resource are evaluated. Number between 1 and 32768')
- VNI - optional:** (empty, with a note: 'The unique VXLAN network identifier that is included in the encapsulated mirrored packet that is sent to the target. A random unique VNI will be chosen unless specified. Number between 0 and 16777215')
- Packet length - optional:** (empty, with a note: 'The number of bytes in each packet to mirror. eg 255 bytes - the entire packet is default')
- Filter:** demo-traffic-mirror-filter (selected from a dropdown menu)


A 'Create filter' button is located at the bottom right of the settings area.

- h. Click **Create**.



Note

For more details, see *Working with Traffic Mirroring* at <https://docs.aws.amazon.com/vpc/latest/mirroring/working-with-traffic-mirroring.html>.



5. Repeat the previous step to create more traffic mirror sessions when there are multiple sources that you want to monitor.

Deploying a Virtual Appliance Behind an NLB

Procedure

1. Configure a load balancer and a listener.
 - a. Open the **Amazon EC2** console at <https://console.aws.amazon.com/ec2/>.
 - b. On the navigation pane, under **LOAD BALANCING**, select **Load Balancers**.
 - c. Select **Create Load Balancer**.
 - d. For **Network Load Balancer**, select **Create**.
 - e. For **Name**, type a name for your load balancer.
For example, type `demo-nlb`.
 - f. For **Scheme**, select **internal**.
 - g. For **Listeners**, modify protocol to **UDP** and type `4789` for the port to receive mirrored traffic.
 - h. For **Availability Zones**, select the VPC that you used for the Deep Discovery Inspector virtual appliance instance and select the subnet for the data port 1 (known as eth0) subnet.



Note

If you enable multiple **Availability Zones** for your load balancer, ensure each target group has at least one target in each **Availability Zone**. Otherwise, the load balancer will not route traffic to Deep Discovery Inspector. For more details, see <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html#network-load-balancer-components>

- i. For **IPv4 address**, you can select **Assigned from CIDR** to have AWS assign the address or select **Enter IP from CIDR** to specify the address.

The screenshot shows the AWS Management Console interface for configuring a Network Load Balancer. The page title is "Step 1: Configure Load Balancer" and the sub-section is "Basic Configuration".

Basic Configuration:

- Name:** demo-nlb
- Scheme:** internal (selected)

Listeners:

- Load Balancer Protocol:** UDP
- Load Balancer Port:** 4789

Availability Zones:

- VPC:** vpc-192.168.0/22 (vDDI-aws-demo-2)
- Availability Zones:** subnet-vDDI-aws-demo-data
- IPv4 address:** Assigned from CIDR 192.168.0/24
- Private IPv4 address:** Assigned from CIDR 192.168.0/24

Temporary limitation: Choose your Availability Zones and subnets with care. After you create the load balancer, you cannot disable the enabled subnets, but you can enable additional ones.

Buttons at the bottom: Cancel, Next: Configure Security Settings

- j. Click **Next: Configure Security Settings**.

2. Configure the security settings.

- a. No changes are necessary in the **Configure Security Settings** screen.
- b. Click **Next: Configure Routing**.

3. Configure a target group.
 - a. For **Target group**, keep the default, **New target group**.
 - b. For **Name**, type a name for the target group.
For example, type `demo-target-group`.
 - c. For **Target type**, select **Instance**.
 - d. For **Protocol**, select **UDP**.
 - e. For **Port**, type `4789`.
 - f. For **Protocol** under **Health checks**, select **TCP**.
 - g. For **Port** under **Advanced health check settings**, select **override** and type `14789` for the port.
 - h. Leave other settings as default.

The screenshot shows the AWS Management Console interface for configuring a target group. The breadcrumb trail indicates the current step is '3. Configure Routing'. The form is titled 'Step 3: Configure Routing' and includes a note: 'Your load balancer routes requests to the targets in this target group using the protocol and port that you specify, and performs health checks on the targets using these health check settings. Note that each target group can be associated with only one load balancer.'

The form fields are as follows:

- Target group:** New target group
- Name:** demo-target-group
- Target type:** Instance
- Protocol:** UDP
- Port:** 4789
- Health checks:** TCP
- Advanced health check settings:**
 - Port:** override 14789
 - Healthy threshold:** 3
 - Unhealthy threshold:** 3
 - Timeout:** 10 seconds
 - Interval:** 30 seconds

Navigation buttons at the bottom right include 'Cancel', 'Previous', and 'Next: Register Targets'.

- i. Click **Next: Register Targets**.
4. Register targets with the target group.
 - a. For **Instances**, select the Deep Discovery Inspector virtual appliance.

For example, select **demo-ddi**.

- b. Keep the default instance listener port and select **Add to registered**.

Step 4: Register Targets
Register targets with your target group. If you register a target in an enabled Availability Zone, the load balancer starts routing requests to the targets as soon as the registration process completes and the target passes the initial health checks.

Registered targets
To deregister instances, select one or more registered instances and then click Remove.

Instance	Name	Port	State	Security groups	Zone
	vDDI-demo	4789	running	sg-1a2b3c4d	us-east-1a

Instances
To register additional instances, select one or more running instances, specify a port, and then click Add. The default port is the port specified for the target group. If the instance is already registered on the specified port, you must specify a different port.

Add to registered on port 4789

Search instances

Instance	Name	State	Security groups	Zone	Subnet ID	Subnet CIDR
	vDDI-demo	running	sg-1a2b3c4d	us-east-1a	subnet-1a2b3c4d	192.168.0.0/24
	vDDI-aws-demo	running	sg-1a2b3c4d	us-east-1a	subnet-1a2b3c4d	192.168.0.0/24

Cancel Previous Next: Review

- c. Click **Next: Review**.

The **Review** screen appears.

Step 5: Review
Please review the load balancer details before continuing.

Load balancer

- Name: demo-rlb
- Scheme: internal
- Listeners: port 4789 - Protocol UDP
- IP address type: ipv4
- VPC: vpc-1a2b3c4d (vDDI-aws-demo)
- Subnets: subnet-1a2b3c4d (vDDI-aws-demo-data)
- Tags:

Routing

- Target group: New target group
- Target group name: demo-target-group
- Port: 4789
- Target type: instance
- Protocol: UDP
- Health check protocol: TCP
- Health check port: 4789
- Healthy threshold: 3
- Unhealthy threshold: 3
- Interval: 30

Targets

- Instances: vDDI-demo:4789

Cancel Previous Create

5. Create the load balancer.
 - a. On the **Review** screen, click **Create**.
 - b. After the load balancer is created, click **Close**.
 - c. On the navigation pane, under **LOAD BALANCING**, select **Target Groups**.

- d. Select the newly created target group.
For example, select **demo-target-group**.
- e. Select **Targets** and verify that your instances are ready.



Note

If the status of an instance is initial, it's probably because the instance is still in the process of being registered, or it has not passed the minimum number of health checks to be considered healthy. After the status of at least one instance is healthy, you can test your load balancer.

If the Deep Discovery Inspector virtual appliance is launched after the NLB was created, use **Register targets** to add the Deep Discovery Inspector virtual appliance to the NLB target groups. For more details, see <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/target-group-register-targets.html>.

The screenshot shows the AWS Management Console interface for a target group. The left sidebar shows the navigation menu with 'Load Balancing' > 'Target Groups' selected. The main content area shows the configuration for 'demo-target-group'. Under the 'Targets' tab, there is a table of registered targets:


Instance ID	Name	Port	Zone	Status	Status details
i-12345678	vddi-demo	4789	us-east-1b	healthy	

6. Configure the traffic mirror filter.

For details, see <https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-filters.html>.

- a. Open the **Amazon VPC** console at <https://console.aws.amazon.com/vpc/>.
- b. In the **Region** selector, select the AWS Region that you used when you created the VPCs.

- c. On the navigation pane, go to **Traffic Mirroring > Mirror Filters**.
- d. Select **Create traffic mirror filter**.
- e. For **Name tag**, type a name for the traffic mirror filter.
For example, type `demo-traffic-mirror-filter`.
- f. (Optional) For **Description**, type a description for the traffic mirror filter.
For example, type `demo-traffic-mirror-filter`.
- g. (Optional) For **Network services**, select **amazon-dns**.




- h. Add inbound rules. Select **Inbound rules > Add > rule**, and then specify the following information about the traffic mirror source inbound traffic:
 - **Rule number**: Type a priority to assign to the rule.
 - **Rule action**: Select an action to take for the packet.
 - **Protocol**: Select a L4 protocol to assign to the rule.
 - (Optional) **Source port range**: Type a source port range.
 - (Optional) **Destination port range**: Type a destination port range.
 - **Source CIDR block**: Type a source CIDR block.
 - **Destination CIDR block**: Type a destination CIDR block.

- (Optional) **Description:** Type a description for the rule.

The following is an example of the values.

- **Rule number:** Use the default number
- **Rule action:** Select **accept**
- **Protocol:** Select **All protocols**
- **Source CIDR block:** Type `0.0.0.0/0`.
- **Destination CIDR block:** Type `0.0.0.0/0`.
- **Description:** Type `mirror all inbound traffic`.




- Add outbound rules. Select **Outbound rules > Add > rule**, and then specify the following information about the traffic mirror source outbound traffic:


- **Rule number:** Type a priority to assign to the rule.
- **Rule action:** Select an action to take for the packet.
- **Protocol:** Select a L4 protocol to assign to the rule.
- (Optional) **Source port range:** Type a source port range.
- (Optional) **Destination port range:** Type a destination port range.
- **Source CIDR block:** Type a source CIDR block.
- **Destination CIDR block:** Type a destination CIDR block.
- (Optional) **Description:** Type a description for the rule.

The following is an example of the values.

- **Rule number:** Use the default number
- **Rule action:** Select **accept**
- **Protocol:** Select **All protocols**
- **Source CIDR block:** Type **0.0.0.0/0**.
- **Destination CIDR block:** Type **0.0.0.0/0**.
- **Description:** Type **mirror all outbound traffic**.




- Repeat the previous step for each inbound rule and outbound rule that you want to add.
- Click **Create**.



- Configure the traffic mirror target.
 - Open the **Amazon VPC** console at <https://console.aws.amazon.com/vpc/>.

- b. In the **Region** selector, select the AWS Region that you used when you created the VPCs.
- c. On the navigation pane, go to **Traffic Mirroring > Mirror Targets**.
- d. Select **Create Traffic Mirror Target**.
- e. For **Name tag**, type a name for the traffic mirror target.
For example, type `demo-traffic-mirror-target`.
- f. (Optional) For **Description**, type a description for the traffic mirror target.


For example, type `demo-traffic-mirror-target`.



The screenshot shows the AWS console interface for creating a traffic mirror target. The page title is 'Create traffic mirror target'. Under the 'Target settings' section, there are two input fields: 'Name tag - optional' and 'Description - optional'. Both fields contain the text 'demo-traffic-mirror-target'.

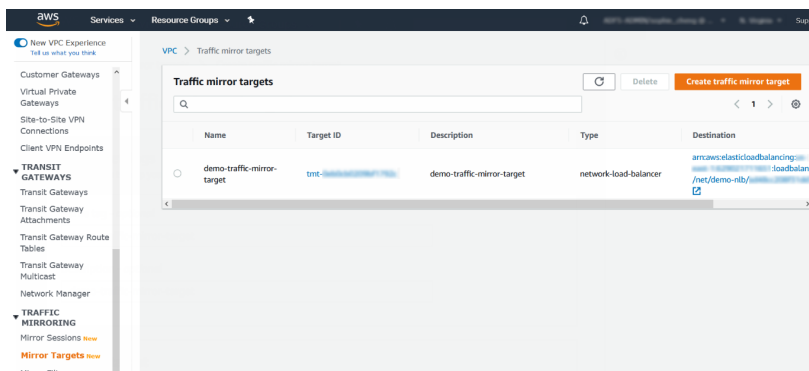
- g. For **Target type**, select **Network Load Balancer**.
- h. For **Target**, select a Network Load Balancer as the traffic mirror target.

For example, select `demo-nlb`.

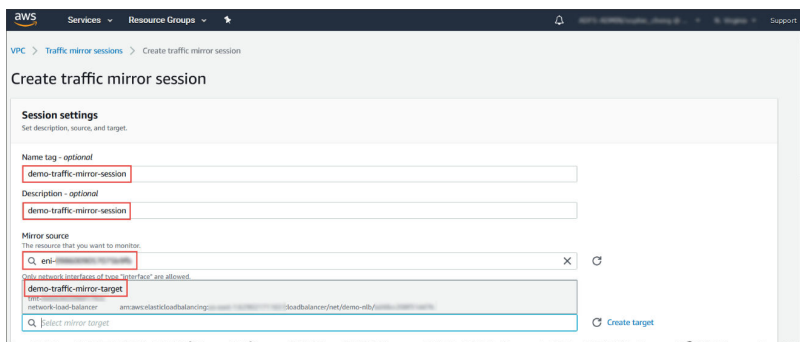


The screenshot shows the AWS console interface for choosing a target. The page title is 'Choose target'. Under the 'Target type' section, a dropdown menu is set to 'Network Load Balancer'. Under the 'Target' section, there is a search bar with the text 'demo-nlb' entered. Below the search bar, a list of search results is displayed, with 'demo-nlb' selected.

- i. Click **Create**.



8. Configure the traffic mirror session.
 - a. On the navigation pane, select **Traffic Mirroring > Mirror Sessions**.
 - b. Select **Create traffic mirror session**.
 - c. For **Name tag**, type a name for the traffic mirror session.
For example, type `demo-traffic-mirror-session`.
 - d. (Optional) For **Description**, type a description for the traffic mirror session.
For example, type `demo-traffic-mirror-session`.
 - e. For **Mirror source**, select the network interface of the instance that you want to monitor.
 - f. For **Mirror target**, select the traffic mirror target.
For example, select **demo-traffic-mirror-target**.



g. Under **Additional settings**, perform the following:

- For **Session number**, type the session number **1**.

The session number determines the order that traffic mirror sessions are evaluated in both of the following situations:

- When an interface is used by multiple sessions.
- When an interface is used by different traffic mirror targets and traffic mirror filters.

Traffic is only mirrored one time. Use **1** for the highest priority. Valid values are 1-32766.

- (Optional) For **VNI**, type the VXLAN ID to use for the traffic mirror session.

For details, see <https://tools.ietf.org/html/rfc7348>.

If you do not specify a value, AWS assigns a random, unused number.

- (Optional) For **Packet Length**, type the number of bytes in each packet to mirror.

If you do not want to mirror the entire packet, set **Packet Length** to the number of bytes in each packet to mirror. For example, if you set this value to 100, the first 100 bytes after the VXLAN header that meet the filter criteria are copied to the target.

To mirror the entire packet, do not enter a value in this field.

- For **Filter**, select the traffic mirror filter that determines what traffic gets mirrored.

For example, select **demo-traffic-mirror-filter**.

- (Optional) Under the **Tags** section, add or remove a tag.

The following are example settings.

- For **Session number**, type the session number **1**.
- For **VNI**, leave the value empty. AWS will assign a random number.
- For **Packet Length**, leave the value empty. AWS will mirror the entire packet.
- For **Filter**, select **demo-traffic-mirror-filter**.

The screenshot shows the 'Additional settings' section in the AWS console. It includes the following fields and values:

- Session number:** 1 (Number between 1 and 32766)
- VNI - optional:** (Empty, Number between 0 and 16777215)
- Packet length - optional:** (Empty, eg 255 bytes - the entire packet is default)
- Filter:** demo-traffic-mirror-filter


A 'Create filter' button is located at the bottom right of the settings area.

- h. Click **Create**.



Note

For more details, see <https://docs.aws.amazon.com/vpc/latest/mirroring/working-with-traffic-mirroring.html>.



9. Repeat the previous step to create more traffic mirror sessions when there are multiple sources that you want to monitor.
-


Chapter 4

Deployment Testing and Troubleshooting


Checkpoints



Pass the following checkpoints to ensure that the deployment is successful.



TABLE 4-1. Checkpoints

#	DESCRIPTION
1	<p>Use an IPv4 address to log in to the management console of the Deep Discovery Inspector virtual appliance.</p> <p>You can find the management IP address on the Amazon EC2 console by following the steps below.</p> <ol style="list-style-type: none"> 1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/. 2. In the navigation pane, select Instances. 3. Select the Deep Discovery Inspector virtual appliance. 4. Select Actions > Networking > Manage IP Addresses. 5. Expand eth1. The Private IP Address is the IP address of the management console for the Deep Discovery Inspector virtual appliance. 
2	Active the Deep Discovery Inspector appliance with the Activation Code.

#	DESCRIPTION
3	Update the components on the Deep Discovery Inspector appliance.
4	<p>Follow the steps below to perform a network services diagnostic test on the Deep Discovery Inspector appliance and verify that all the tests are successful.</p> <ol style="list-style-type: none"> Go to <code>https://<virtual appliance IP address>/html/troubleshooting.htm</code> and click Network Services Diagnostics. Select one or more enabled services and click Test. When there are no connection issues, the result of all tested services is Connected.



#	DESCRIPTION
5	<p>Verify that the traffic mirror filter contains rules allowing the HTTP protocol in both inbound and outbound traffic.</p>  <p>The screenshot shows the AWS console for a traffic mirror filter named 'demo-traffic-mirror-filter'. The 'Inbound rules' tab is active, showing a table with one rule: Rule number 100, Description 'mirror all inbound traffic', Rule action 'accept', Protocol 'All protocols', Source port range '-', Destination port range '-', Source CIDR block '0.0.0.0/0', and Destination CIDR block '0.0.0.0/0'. The 'Outbound rules' tab is also active, showing a similar rule: Rule number 100, Description 'mirror all outbound traffic', Rule action 'accept', Protocol 'All protocols', Source port range '-', Destination port range '-', Source CIDR block '0.0.0.0/0', and Destination CIDR block '0.0.0.0/0'.</p>
6	<p>If you deploy Deep Discovery Inspector as the traffic mirror target, verify that the mirror target, for example demo-traffic-mirror-target, is configured with destination to the Deep Discovery Inspector virtual appliance.</p>  <p>The screenshot shows the AWS console for traffic mirror targets. A table lists one target: Name 'demo-traffic-mirror-target', Target ID 'tmt-...', Description 'demo-traffic-mirror-target', Type 'network-interface', and Destination 'eni-...'. The 'network-interface' type is highlighted with a red box.</p>

#	DESCRIPTION
7	<p>If you deploy Deep Discovery Inspector behind the NLB, verify that the mirror target, for example demo-traffic-mirror-target, is configured with destination to the NLB.</p> 
8	<p>Verify that the mirror session, for example demo-traffic-mirror-session, is configured properly for the following fields:</p> <ul style="list-style-type: none"> • Source • Target • Session number • Filter
9	<p>If you deploy Deep Discovery Inspector behind the NLB, verify that the status of the registered instance in the target group, for example demo-target-group, is healthy.</p> 

Testing the Deployment

You can perform the following steps to validate the Deep Discovery Inspector virtual appliance deployment:

Procedure

1. Access a test website on your test EC2 instance.

The following example is for a Linux instance.

Your testing EC2 instance must be configured as the traffic mirror source when Deep Discovery Inspector is deployed as a traffic mirror target and when Deep Discovery Inspector is deployed behind and NLB.

In the example below, replace `hxxp` with `http`.

```
~$ curl hxxp://wrs49.winshipway.com/
```

2. Verify the detection on the Deep Discovery Inspector virtual appliance.
 - a. Log in to the management console of the Deep Discovery Inspector virtual appliance.
 - b. Go to **Detections > All Detections**.
 - c. Verify that the website was detected.

The screenshot shows the 'All Detections' page in the management console. It features a search bar at the top, a 'Detection severity: High only' filter, and a table of detections. The table has columns for Details, Status, Timestamp, Source Host, Destination Host, Interested Host, Threat Description, Out., Protocol, Detection Severity, Attack Phase, and Hostile Object. A single detection is visible with a red status icon, a timestamp of 1/7/2020 8:00, source host 'ec2-...', interested host '1...', threat description 'Dangerous URL in Web Reputation Services database - HTTP (Request)', protocol 'HTTP', detection severity 'Medium', attack phase 'Point of Entry', and hostile object 'URL: http://wrs49.winshipway.com'.


Details	Status	Timestamp	Source Host	Destination Host	Interested Host	Threat Description	Out.	Protocol	Detection Severity	Attack Phase	Hostile Object
	!	1/7/2020 8:00	ec2-...		1...	Dangerous URL in Web Reputation Services database - HTTP (Request)		HTTP	Medium	Point of Entry	URL: http://wrs49.winshipway.com

Troubleshooting the Deployment

The following are several tips for troubleshooting packet reception issues on Amazon EC2.

- Use the Deep Discovery Inspector virtual appliance Network Traffic Dump

On the Deep Discovery Inspector virtual appliance, go to **Troubleshooting > Network Traffic Dump** and capture packets to check the data port's reception.



- Verify Network ACLs settings

For details, see <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-network-acls.html>.

- Verify Security Group settings

For details, see https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html. For the traffic mirror target, the traffic mirror target requires the allowance of **VXLAN traffic (UDP port 4789)** from the traffic mirror source in the security groups that are associated with the traffic mirror target.

**Note**

If you are using deploying Deep Discovery Inspector behind an NLB, you may need to allow **custom traffic (TCP port 14789)** to the Deep Discovery Inspector virtual appliance in the security groups that are associated with the Deep Discovery Inspector virtual appliance.

Frequently Asked Questions

- *What are the changes on the Deep Discovery Inspector virtual appliance on AWS? on page 4-9*
- *Does the Deep Discovery Inspector virtual appliance support AWS EC2 auto scaling? on page 4-14*
- *Does Deep Discovery Inspector support creating an Amazon Machine Image (AMI) from an EC2 instance of the Deep Discovery Inspector virtual appliance? on page 4-14*
- *Does Deep Discovery Inspector support creating an Elastic Block Store (EBS) snapshot from an EC2 instance of the Deep Discovery Inspector virtual appliance? on page 4-15*
- *Does Deep Discovery Inspector Support AWS Backup Service? on page 4-16*
- *What are the IAM policies needed to deploy Deep Discovery Inspector's virtual appliance on AWS? on page 4-16*

What are the changes on the Deep Discovery Inspector virtual appliance on AWS?

In order to adapt into the AWS environment, the Deep Discovery Inspector virtual appliance has some minor changes. These changes do not impact any major features and are described in the following list.

- Swapping port enumeration for management port

The management port for Deep Discovery Inspector on-premises is fixed at the first NIC port (known as eth0). This change provides consistent information on Amazon EC2 console.

The Deep Discovery Inspector virtual appliance swapped port enumeration for the management port to port 1 (known as eth1) and the data port to port 0 (known as eth0).

Network Interface ⓘ

Check VLAN tags of each stream to differentiate connections [Show advanced settings](#)

Interface	Function	MAC Address	EC2 Instance Network Interface ⓘ	Status
Management Port	Management	...	eth1	...
Port 1	Data	...	eth0	...
Port 2	Data	...	eth2	...
Port 3	Data	...	eth3	...
Port 4	Data	...	eth4	...

- IPv4 address for management port only supports DHCP

Management ports configured as IPv4 only support DHCP. To modify the IPv4 address that is assigned, use the Amazon EC2 console.

Dashboard | Detections | Reports | Administration | Help

You are here: Administration > System Settings > Network

System Settings

- Network
- Network Interface
- Proxy
- SMTP
- SNMP
- HTTPS Certificate
- Time
- Session Timeout

Network

Appliance Identity

Host name or FQDN:

Use host name instead of IP address as the identity of this Deep Discovery Inspector

Management Port

IPv4 Type:

IPv4 address:

IPv4 subnet mask: 255.255.255.0

IPv4 gateway:

IPv4 DNS server 1:

IPv4 DNS server 2:

Enable IPv6 address


To modify the IPv4 address that is assigned, perform the following steps on the Amazon EC2 console.

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, select **Instances** and select the Deep Discovery Inspector virtual appliance.
3. Go to **Actions > Networking > Detach Network Interface**.

4. In the drop-down list, select **eth1** and click **Detach**.
5. In the navigation pane, select **Network interfaces**.

You can create a network interface (For details, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#create_eni) or find the IPv4 address that you want to attach to the management port of the Deep Discovery Inspector virtual appliance.

6. Select the network interface that you created or found in the previous step, and then click **Attach**.
7. Select the instance ID of the Deep Discovery Inspector virtual appliance, and then click **Attach**.

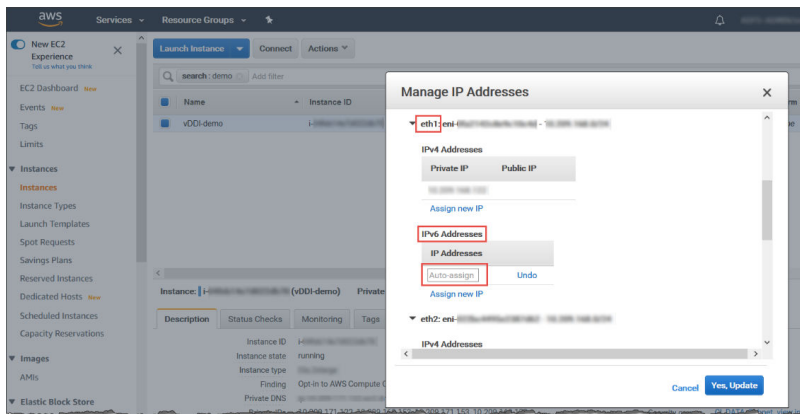


8. **Reboot** the Deep Discovery Inspector virtual appliance.
 9. Verify that the management port (eth1) of the Deep Discovery Inspector virtual appliance is assigned to the new IPv4 address.
- IPv6 address for management port only supports DHCP

On AWS, the IPv6 address is managed on the Amazon EC2 console. The Deep Discovery Inspector virtual appliance on AWS retrieves the IPv6 address automatically when IPv6 is assigned to a network interface on the Amazon EC2 console.

To assign an IPv6 address, perform the following steps.

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, select **Instances**.
3. Select the Deep Discovery Inspector virtual appliance, and then select **Actions** > **Networking** > **Manage IP Addresses**.
4. For **eth1**, under **IPv6 Addresses**, select **Assign new IP**. You can specify an IPv6 address in the subnet range, or leave the **Auto-assign** value to let Amazon choose an IPv6 address for you.




5. Click **Yes, Update**.
6. Log in to the management console of the Deep Discovery Inspector virtual appliance.
7. Go to **Administration** > **System Settings** > **Network**.
8. In **Management Port** section, select **Enable IPv6 address**.
9. Click **Save**.
10. **Reboot** the Deep Discovery Inspector virtual appliance.
11. Go to **Administration** > **System Settings** > **Network** and verify that the Deep Discovery Inspector virtual appliance is assigned an IPv6 address.

The screenshot displays the 'Deep Discovery Inspector' Administration console. The top navigation bar includes 'Dashboard', 'Detections', 'Reports', 'Administration', and 'Help'. The breadcrumb trail indicates the current location: 'Administration > System Settings > Network'. On the left, a 'System Settings' sidebar lists 'Network' as the active section, along with 'Network Interface', 'Proxy', 'SMTP', 'SNMP', 'HTTPS Certificate', 'Time', and 'Session Timeout'. The main content area is titled 'Network' and is divided into two sections: 'Appliance Identity' and 'Management Port'. The 'Appliance Identity' section contains a 'Host name or FQDN' field and a checkbox for 'Use host name instead of IP address as the identity of this Deep Discovery Inspector'. The 'Management Port' section is further divided into 'IPv4' and 'IPv6' configurations. The 'IPv4' section includes fields for 'IP4 Type' (set to 'Dynamic IP address (DHCP)'), 'IP4 address', 'IP4 subnet mask', 'IP4 gateway', 'IP4 DNS server 1', and 'IP4 DNS server 2'. The 'IPv6' section includes a checked checkbox for 'Enable IPv6 address', 'IPv6 Type' (set to 'Dynamic IP address (DHCP)'), 'IPv6 address' (with a port of 2600), 'IPv6 subnet prefix length' (set to 64), 'IPv6 gateway', and 'IPv6 DNS server'. Red boxes highlight the 'Management Port' section header, the 'Enable IPv6 address' checkbox, and the IPv6 configuration fields.

- No support for internal Virtual Analyzer

When launching a Deep Discovery Inspector virtual appliance on AWS, only external Virtual Analyzer and Sandbox as a Service are supported.




Does the Deep Discovery Inspector virtual appliance support AWS EC2 auto scaling?

No. The Deep Discovery Inspector virtual appliance does not support AWS EC2 auto scaling.


Does Deep Discovery Inspector support creating an Amazon Machine Image (AMI) from an EC2 instance of the Deep Discovery Inspector virtual appliance?

No. Deep Discovery Inspector does not support creating an AMI from an EC2 instance of the Deep Discovery Inspector virtual appliance.



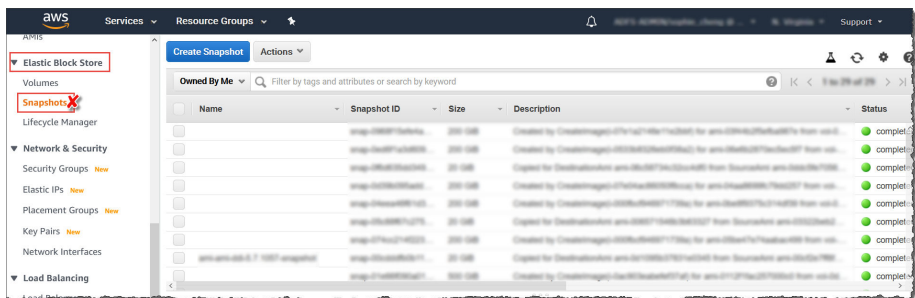
After installation, the Deep Discovery Inspector virtual appliance creates a UUID automatically and this UUID is used everywhere when communicating with Trend Micro global services. Creating a VM clone will corrupt the health status of bounded services.

If the Deep Discovery Inspector virtual appliance detects that the instance ID has changed, there is a warning message on the Deep Discovery Inspector virtual appliance management console.



Does Deep Discovery Inspector support creating an Elastic Block Store (EBS) snapshot from an EC2 instance of the Deep Discovery Inspector virtual appliance?


No. Deep Discovery Inspector does not support creating an EBS snapshot from an EC2 instance of the Deep Discovery Inspector virtual appliance.



After installation, the Deep Discovery Inspector virtual appliance creates a UUID automatically and this UUID is used everywhere when communicating

with Trend Micro global services. Creating a VM clone will corrupt the health status of bounded services.

If the Deep Discovery Inspector virtual appliance detects that the instance ID has changed, there is a warning message on the Deep Discovery Inspector virtual appliance management console.



Does Deep Discovery Inspector Support AWS Backup Service?

Deep Discovery Inspector does not support **AWS Backup** service.

After installation, Deep Discovery Inspector virtual appliance creates a UUID automatically, and this UUID is used everywhere when communicating with Trend Micro global services. Creating a VM clone will corrupt the health status of integrated services.

When Deep Discovery Inspector has detected that the instance ID has changed, a warning message appears in the Deep Discovery Inspector virtual appliance management console.

What are the IAM policies needed to deploy Deep Discovery Inspector's virtual appliance on AWS?

IAM (Identity and Access Management) is an AWS feature you can use to control who is authenticated and authorized to use resources. To deploy Deep Discovery Inspector, ensure your IAM user has the following permissions.

AWS SERVICE	POLICY NAME
EC2 instances	<ul style="list-style-type: none">• AmazonEC2FullAccess• IAMReadOnlyAccess• AllowAssumeCIEC2Deployment• AmazonEC2SpotFleetTaggingRole
EC2 Network & Security	<ul style="list-style-type: none">• AmazonEC2FullAccess• IAMReadOnlyAccess• AllowAssumeCIEC2Deployment• AmazonEC2SpotFleetTaggingRole
EC2 Load Balancing	<ul style="list-style-type: none">• AmazonEC2FullAccess• IAMReadOnlyAccess• AllowAssumeCIEC2Deployment• AmazonEC2SpotFleetTaggingRole
VPC TRAFFIC MIRRORING	<ul style="list-style-type: none">• AmazonEC2FullAccess• IAMReadOnlyAccess• AllowAssumeCIEC2Deployment• AmazonEC2SpotFleetTaggingRole
AWS Marketplace	AWSMarketplaceManageSubscriptions
AWS Compute Optimizer finding	ComputeOptimizerReadOnlyAccess



TREND MICRO INCORPORATED

225 E. John Carpenter Freeway, Suite 1500
Irving, Texas 75062 U.S.A.
Phone: +1 (817) 569-8900, Toll-free: (888) 762-8736
Email: support@trendmicro.com

www.trendmicro.com

Item Code: APEM69459/211129