CLOUDMON NTM

Getting started guide – AWS

Selvaraj B

Table of Contents

INTRODUCTION

As businesses grow, network infrastructure growth across physical, virtual and public cloud, more often than not are bound to result in complexity and in-efficiency. The very mission critical infrastructure that should help businesses realize the benefits of digital transformation and innovation, often plays spoil-sport because of unknown problems lurking in the network and the resulting performance and availability challenges.

With Veryx Cloudmon NTM, enterprise businesses get 100% network visibility and analytics of all traffic across their mission critical infrastructure –

- *Whether on-premises or private cloud or AWS cloud*
- *Assure the performance of the application in hybrid environments.*
- *Deliver a consistent and high-quality user experience*
- *Enabling better control and realization of the power of digital innovation*
- *At a fraction of the cost of competing solutions.*

Figure 1 illustrates a sample hybrid deployment with a cloud Cloudmon NTM controller operating as a centralized orchestrator. The NTM controller manages

- *Virtual NTM probes in public cloud*
- *Physical NTM probes in on-premises*

Cloudmon NTM probe, minimize public cloud throughput charges.

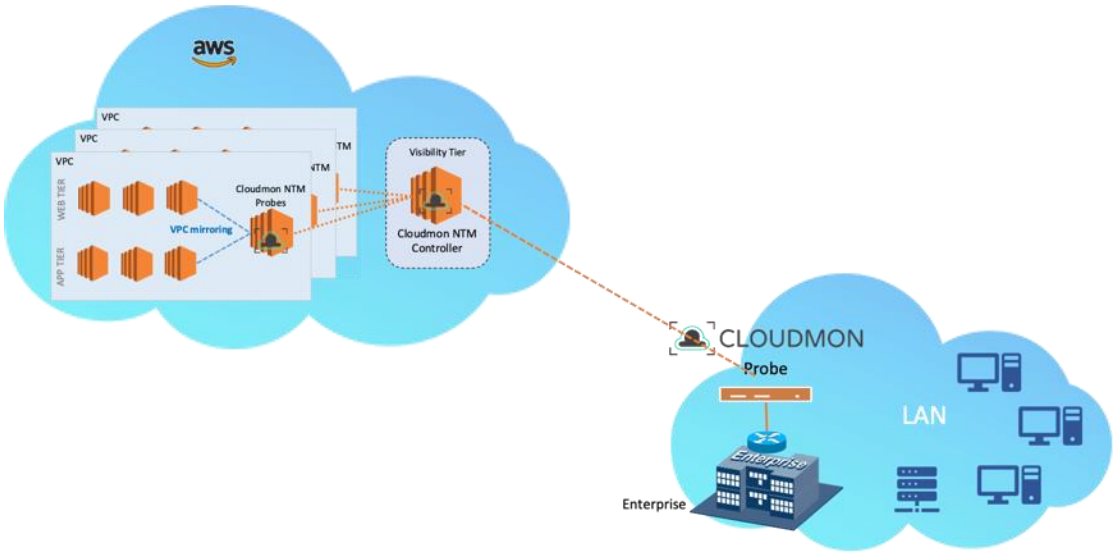

Figure 1 - Sample hybrid deployment

COMPONENTS

NTM Controller

- *Orchestrator metadata from NTM probes and provide intuitive dashboard and reports*
- *Provides data analytics of the orchestrated data*
- *Generates notification / alarms and automated remediation*
- *Provides seamless management of NTM probes, on-premises, and cloud environment*

NTM Probe

- *Deploy as a virtual appliance in AWS EC2 using ready-made AMI. Or deploy as a virtual appliance in private cloud using ready-made VMs. Or deploy as a physical appliance in on-premises using software package*
- *Scalable provisioning depending on Instance Type selected during deployment.*
- *Receives traffic forwarded from multiple virtual taps (vTAPs) or from VPC mirroring for full packet capture & analysis*

SYSTEM REQUIREMENTS – AWS

This document assumes that you are familiar with the networking and configuration of the AWS VPC. In order to provide context for the terms used in this section, here is a brief refresher on the AWS terms (some definitions are taken directly from the AWS glossary) that are referred to in this document:

AWS TERMINOLOGY

Term	Description
EC2	Elastic Compute Cloud A web service that enables you to launch and manage Linux/UNIX and Windows server instances in Amazon's data centers.
Amazon Machine Image (AMI)	An AMI provides the information required to launch an instance, which is a virtual server in the cloud.
Elastic Network Interface (ENI)	An additional network interface that can be attached to an EC2 instance. ENIs can include a primary private IP address, one or more secondary private IP addresses, a public IP address, an elastic IP address (optional), a MAC address, membership in specified security groups, a description, and a source/ destination check flag.
IP address types for EC2 instances	An EC2 instance can have different types of IP addresses. - *Public IP address: An IP address that can be routed across the internet.* - *Private IP address: A IP address in the private IP address range as defined in the RFC 1918. You can choose to manually assign an IP address or to auto assign an IP address within the range in the CIDR block for the subnet in which you launch the EC2 instance.* If you are manually assigning an IP address, Amazon reserves the first four (4) IP addresses and the last one (1) IP address in every subnet for IP networking purposes. - *Elastic IP address (EIP): A static IP address that you have allocated in Amazon EC2 or Amazon VPC and then attached to an instance. Elastic IP addresses are associated with your account, not with a specific instance. They are elastic because you can easily allocate, attach, detach, and free them as your needs change.*

	An instance in a public subnet can have a Private IP address, a Public IP address, and an Elastic IP address (EIP); an instance in a private subnet will have a private IP address and optionally have an EIP.
Instance type	Amazon-defined specifications that stipulate the memory, CPU, storage capacity, and hourly cost for an instance. Some instance types are designed for standard applications, whereas others are designed for CPU-intensive, memory-intensive applications, and so on.
Virtual Private Cloud (VPX)	An elastic network populated by infrastructure, platform, and application services that share common security and interconnection.
Identity and Access Management (IAM) Role	The IAM role defines the API actions and resources the application can use after assuming the role. An IAM role is also required for VM Monitoring.
Subnets	A segment of the IP address range of a VPC to which EC2 instances can be attached. EC2 instances are grouped into subnets based on your security and operational needs. There are two types of subnets: - *Private subnet: The EC2 instances in this subnet cannot be reached from the internet.* - *Public subnet: The internet gateway is attached to the public subnet, and the EC2 instances in this subnet can be reached from the internet.*
Security groups	A security group is attached to an ENI, and it specifies the list of protocols, ports, and IP address ranges that are allowed to establish inbound/outbound connections on the interface.
Key pair	A set of security credentials you use to prove your identity electronically. The key pair consists of a private key and a public key.

Table 1 – AWS terminologies

AWS COMPONENTS & PERMISSIONS

Table 2 summarizes the necessary requirements to deploy Cloudmon NTM in AWS environment.

Component	**Description**
Amazon Web Services Account	You must have an active Amazon Web Services account with access to the EC2 Management Console to deploy in an AWS environment.
Amazon Web Services Permissions	The Amazon Web Services account used to deploy Cloudmon NTM must have appropriate permissions granted. The simplest way to do this is to grant the **AdministratorAccess** policy.

	However, if granting administrator access is not acceptable in your environment, assign the following policies to the account used to deploy Cloudmon NTM components: - *Assign the built-in **AmazonEC2FullAccess** policy.*
Existing AWS VPC	An existing AWS VPC with subnets for both Management and Monitoring.
Route Tables/Security Groups	Appropriate Route Tables and Security Groups for communication between Cloudmon NTM Controller and NTM Probes
Access to Marketplace Images	You must have access to the Veryx Cloudmon NTM AMI images in the AWS Marketplace in the AWS region you are using.
SSH Key Pair	You must have a key pair for SSH access to deployed AMIs. You can create or import the key pair in AWS using these instructions. SSH key pairs created in AWS are: - *Public keys are stored in AWS, are not confidential and are protected at the account level.* - *Private keys are stored by the user and are their responsibility to protect.*

Table 2 – AWS deployment requirements

DEPLOYMENT

CLOUDMON NTM ALL-IN-ONE

This section describes how to deploy the Cloudmon NTM All-in-Once Controller using the AMI available in the Veryx site on the AWS Marketplace

STEP 1

*Log in to the AWS console and select the **EC2 Dashboard***

STEP 2

On the EC2 Dashboard, click **Launch Instance**

STEP 3

*Select the Cloudmon NTM All-in-One Controller. To get the AMI, Search the **Amazon Marketplace** for **Cloudmon NTM***

STEP 4

Click the 

button

STEP 5

Accept the Terms and Conditions

STEP 6

Click Continue and proceed with launch

STEP 7

***Launch** Cloudmon NTM All-in-One on EC2 instance*

1. *Choose the EC2 instance type recommended in the AWS Marketplace. Example: m5.2xlarge or m5.4xlarge based on the model you chosen.*

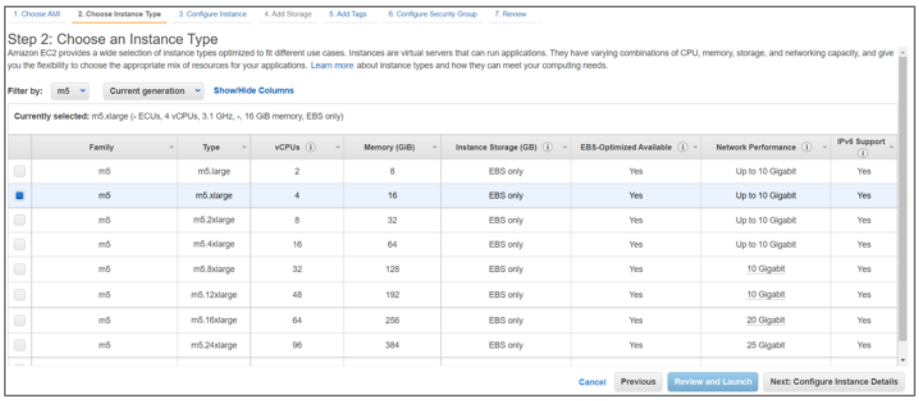

Figure 2 – Selecting EC2 instance type

2. *Select the VPC*
3. *Select the public subnet to which the **management** interface will attach.*
4. *Select **Automatically assign a public IP address.** This allows you to obtain a publicly accessible IP address for the management interface of the Cloudmon NTM.*
    - *You can later attach an Elastic IP address to the management interface; unlike the public IP address that is disassociated from the instance when the instance is terminated, the Elastic IP address provides persistence.*
5. *Add another network interface so that you can mirrored traffic can be attached to it (eth0 and eth1).*
    - *Expand the Network Interfaces section and click **Add Device** to add another network interface. Make sure that your VPC has more than one subnet so that you can add additional ENIs at launch.*

Figure 3 – Adding additional ENI

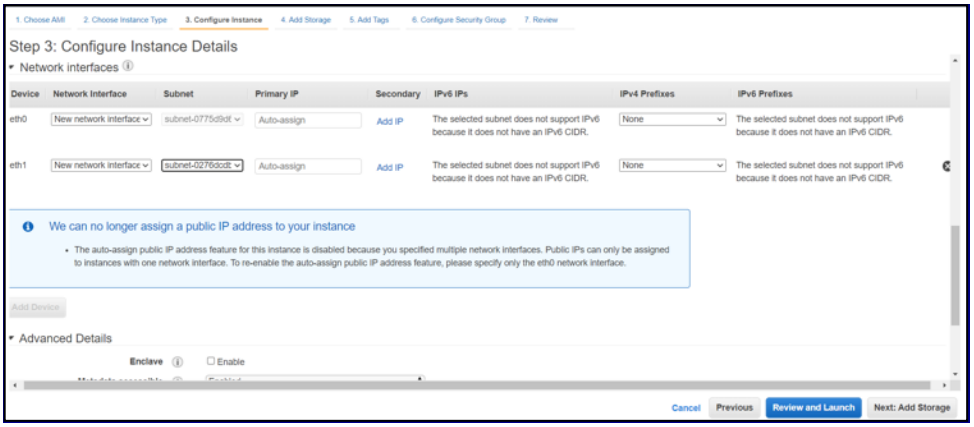

Figure 4 – Configuring new ENI

6. *Click next to add the storage. The storage size is based on the NTM controller model you have chosen. We recommend a minimum storage size of 250GB and can be configured up to 2TB. Use the default storage volume type SSD (gp2).*
7. *Tagging (optional), add one or more tags to create your own metadata to identify and group.*
8. *Select an existing **Security Group** or create a new one. This security group is for restricting access to the management interface. At a minimum enable HTTP, HTTPS. and SSH access for the management interface.*
    - *Ensure that the monitored interface security group allows traffic using **GRE or UDP port 4789**.This ensures Cloudmon NTM to capture AWS mirrored packets.*
9. *Select **Review and Launch**. Review that your selections are accurate and click **Launch.***
10. *Select an existing key pair or create a new one and acknowledge the key disclaimer.*
    - *Download and save the private key to a safe location; the file extension is ".pem". You cannot regenerate this key, if lost. It takes 1-3 minutes to launch the Clooudmon NTM. You can view the progress on the EC2 Dashboard. When the process completes, displays on the Instances page of the EC2 Dashboard.*

STEP 8

The Cloudmon NTM All-in-One controller is Up. The controller will can be accessed via configured public IP via HTTP or HTTPS

AWS TRAFFIC MIRRORING

In addition to forwarding packets from vTAP Agents to Cloudmon NTM All-in-Once controller or NTM probe, AWS provides additional tools that help Cloudmon NTM provide visibility on cloud-based traffic:

- *Amazon VPC traffic mirroring allows you to acquire packet data from multiple application workloads in an Amazon VPC and mirror it to a Cloudmon NTM All-in-One controller or NTM probe instance's monitor port.*

CONFIGURING AWS VPC TRAFFIC MIRRORING

AWS VPC Traffic Mirroring lets you send packets from a mirror source to a destination – a Cloudmon NTM All-in-One controller or NTM probe instance's monitoring interface for our purposes. Traffic Mirroring sessions consist of the following main components:

- *Mirror source. This is where the traffic will be mirrored from.*
- *Mirror destination. This is where mirrored traffic will be sent – a Cloudmon NTM All-in-One controller or NTM probe instance's monitoring port in our case.*
- *Optional filter - This lets you limit which traffic is mirrored to just the packets of interest.*

Refer Amazon AWS documentation for more information.

Traffic Mirroring Prerequisites and Rules

In general:

- *Traffic mirror sources and destinations must either be in the same VPC or in VPCs that are reachable from one another.*
- *The traffic mirror target must have **UDP Port 4789** open to receive traffic.*
- *The traffic mirror source must have a route table entry for the target.*
- *If mirrored traffic is not reaching the destination, check to see if there are any security group or access control list (ACL) rules that are preventing the traffic from arriving.*

Creating a Traffic Mirror Target

1. *Log in to the AWS Management Console and select the **Services > Networking & Content Delivery > VPC** option to launch the VPC Console.*
2. *Locate and select the **Traffic Mirroring > Mirror Targets** option in the navigation panel at the left of the console.*

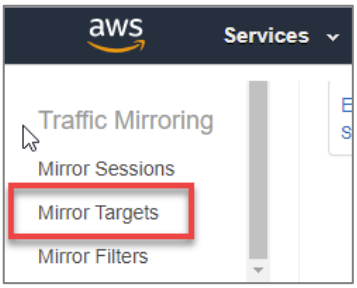

Figure 5 – Creating a Mirror Target

The Traffic mirror targets page appears, listing the existing targets for traffic mirroring.

3. *Click the **Create traffic mirror target** button to create a new target.*
4. *Use the options in the **Create traffic mirror target** page to define the new target*
   - *Target – Cloudmon NTM All-in-One controller or NTP probe's monitoring interface*
   - *Tags – optional*
5. *Click the **Create** button to create the new traffic mirror target and add it to the list of available traffic mirror targets*

Creating a Traffic Mirroring Filter

1. *Log in to the AWS Management Console and select the **Services > Networking & Content Delivery > VPC** option to launch the VPC Console.*
2. *Locate and select the **Traffic Mirroring > Mirror Filters** option in the navigation panel at the left of the console.*

Figure 6 – Creating a Mirror Filter

The **Traffic mirror filters** page appears, listing the existing filters for traffic mirroring.

3. *Click the **Create traffic mirror filter** button to create a new filter.*
4. *Use the options in the **Create traffic mirror filter** page to define the Inbound and Outbound rules for the new filter:*
    - ***Inbound** rules - apply to traffic arriving at whatever mirror source port you apply this filter to. Click the Add rule button and then use the available criteria to define the filter. You can accept or reject traffic based on L4 protocol, source/destination port ranges (optional), and source/destination CIDR blocks (mandatory). Filters are applied based on their priority, as specified by the Number field at the left of each rule's entry in the list.*
    - ***Outbound** rules - apply to traffic sent out of whatever mirror source port you apply this filter to. The same filtering criteria available for Inbound rules are available for Outbound rules.*
5. *Click the Create button to create new traffic mirror filter.*

Creating a Traffic Mirroring Session

Once you've created both a target and a filter, you're ready to establish a traffic mirroring session. Use the following procedure:

1. *Log in to the AWS Management Console and select the **Services > Networking & Content Delivery > VPC** option to launch the VPC Console.*
2. *Locate and select the **Traffic Mirroring > Mirror Sessions** option in the navigation panel at the left of the console.*

Figure 7 – Creating a Mirror Session

The Traffic mirror sessions page appears, listing the existing traffic mirroring sessions.

3. *Click the **Create traffic mirror session** button to create a new session.*
4. *Use the options in the **Create traffic mirror session** page to set up the mirroring session*
   - *Mirror source – select the interface whose traffic you want to mirror and monitor*
   - *Mirror target – select the interface where you want to send the monitored traffic. This will be monitored interface of Cloudmon NTM All-in-One or NTM probe monitored interface in our case.*
   - *VNI – optional VXLAN network identifier*
   - *Packet length – By default, the entire packet will be mirrored. You can optionally slice the packet before sending it to mirror target*
   - *Filters – optional, select any of the filters created.*
5. *Click the **Create** button to create the new traffic mirror session and add it to the list of active sessions.*

Note

Keep in mind that the mirroring session remains active until you cancel it from the Traffic mirror sessions page.