

WiNG 5 Installation Guide

VX 9000 – Amazon EC2 Cloud

Published: April 2017

Extreme Networks, Inc.
Phone / +1 408.579.2800
Toll-free / +1 888.257.3000

www.extremenetworks.com

© 2017 Extreme Networks, Inc. All rights reserved.

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries. All other names are the property of their respective owners. All other registered trademarks, trademarks, and service marks are property of their respective owners. For additional information on Extreme Networks trademarks, see www.extremenetworks.com/company/legal/trademarks.

Contents

Overview.....	3
Prerequisites.....	3
VX 9000.....	4
Licensing and Elastic IP addresses on EC2	13

Overview

The VX virtualized controller is a software WLAN controller appliance running as a virtual machine (VM) on a variety of Hypervisor and Amazon EC2 cloud infrastructures. The VX 9000 supports virtually any server and commercially available Hypervisors for fast and seamless integration into an existing network infrastructure, without adding new hardware. You can run multiple instances of the VX 9000 on a single server, reducing cost, space and power in the Network Operations Center (NOC). With the ability to run in a private or public cloud, you have the freedom to choose the model that works best for your deployment needs, install on your own servers or lease a server in the public cloud.

Prerequisites

- VX 9000 ISO image version 5.8.4.0 or above
- [Amazon EC2 CLI tools](#) installed.

Note

Only legacy EC2 CLI tools are supported (tested version 1.7.5.1). New AWS CLI Tools will not work.

After unpacking the CLI Tools following environmental variables needs to be added into the system:

User Variables:

- **%EC2_HOME%** - point to the location of `\ec2-api-tools\ec2-api-tools-1.7.5.1` folder.
- **%JAVA_HOME%** - point to the location of java jre folder, for example `C:\Program Files (x86)\Java\jre1.8.0_101`
- **Path** - add `%EC2_HOME%\bin` and `%JAVA_HOME%`
- Local supported hypervisor available for initial image porting (VMWare ESXi, Citrix XenServer or Microsoft Hyper-V).


VX 9000

1. Create virtual machine on supported hypervisor (recommend minimum 4 GB RAM, 8 GB HDD).


Note

If installing for running Guest Registration or NSight database assign disk space at this stage. Use Thin Provisioning in ESXi. Currently EC2 Elastic Storage is not supported.

2. Start the VX 9000 installation process. Hit Enter when prompted to start the process.
3. When prompted to enable support for Amazon (EC2) HVM mode type yes and hit Enter.

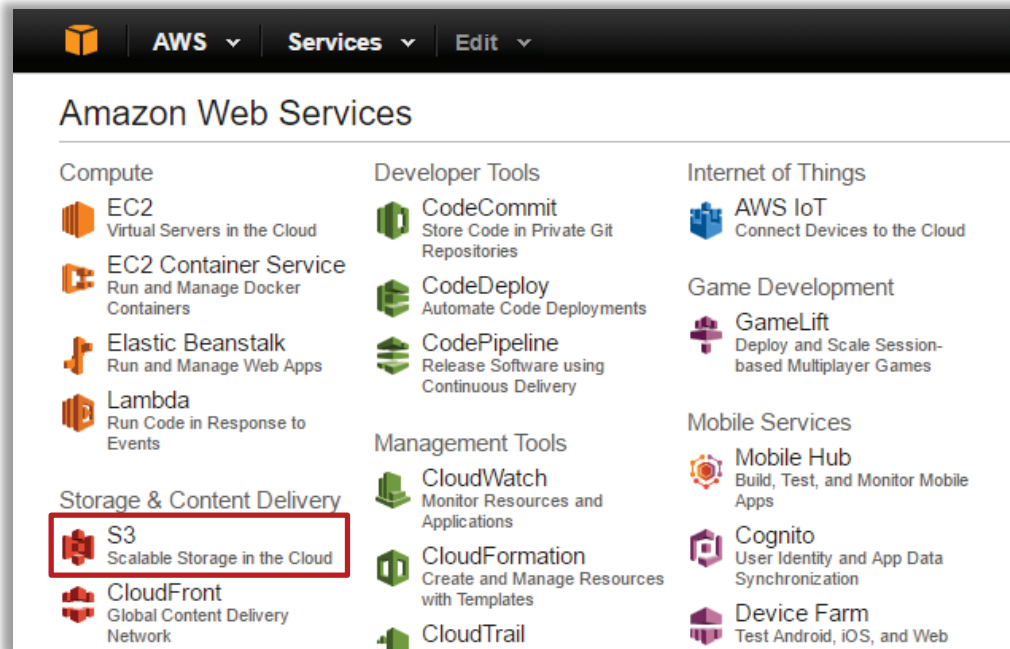


4. Once the install is complete, power off the VM. **Do NOT press Enter at this stage:**




5. Export VM image to local system. In VMWare ESXi use **File -> Export -> Export OVF Template**. Select "Folder of Files" format.
6. After Export in the destination folder locate the .vmdk file with the VM image.

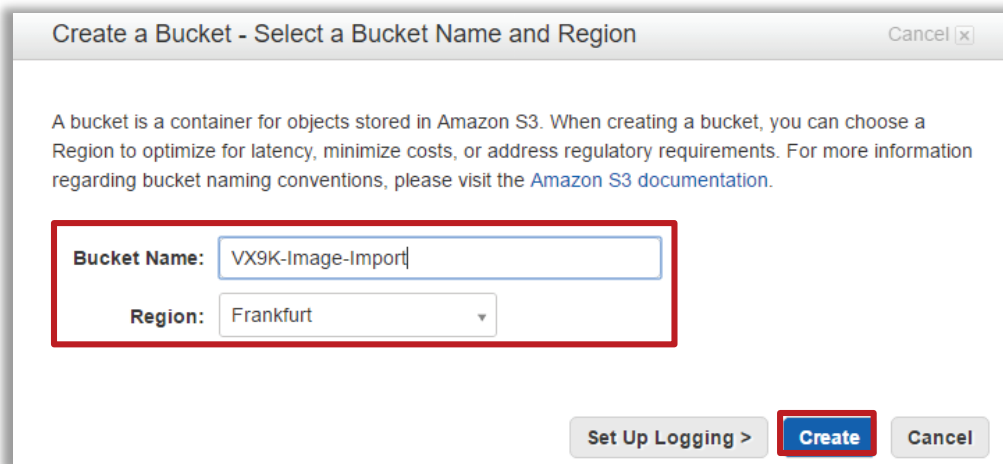
7. Log in into [Amazon AWS Console](#).
8. Go to **S3** and create a data bucket:



9. Select Create Bucket:



10. Specify bucket name and region where you want to run the VX 9000 and select the **Create** button.



11. Open a command shell on your PC and issue the following import commands (assuming Amazon CLI tools are installed):

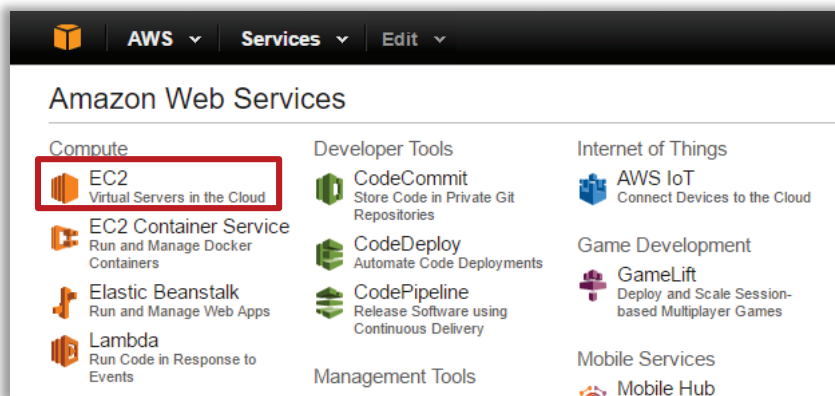
```
C:\ec2-import-volume <image-file-name>.vmdk -f vmdk -b <S3 bucket name> -O <ACCESS KEY> -W <SECURITY KEY> -o <ACCESS KEY> -w <SECURITY KEY> -z <your zone, e.g. eu-central-1a> --region <your region, e.g. eu-central-1>
```

12. Check the porting status from another command line shell:

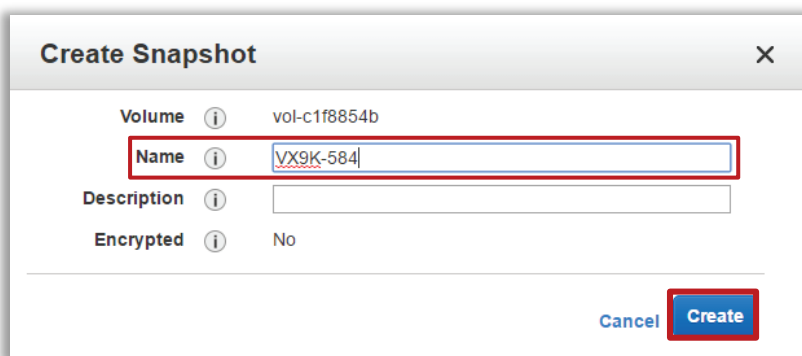
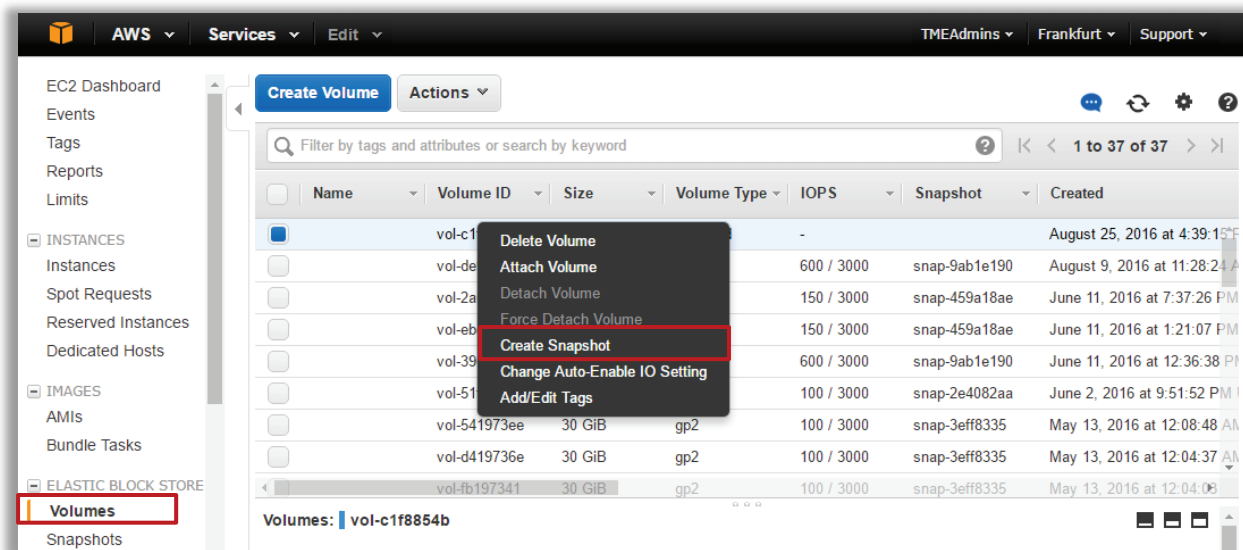
```
C:\ec2-describe-conversion-tasks -region eu-central-1
```

This command will display the volume ID to be created on **S3**, for example “vol-7b596774”.

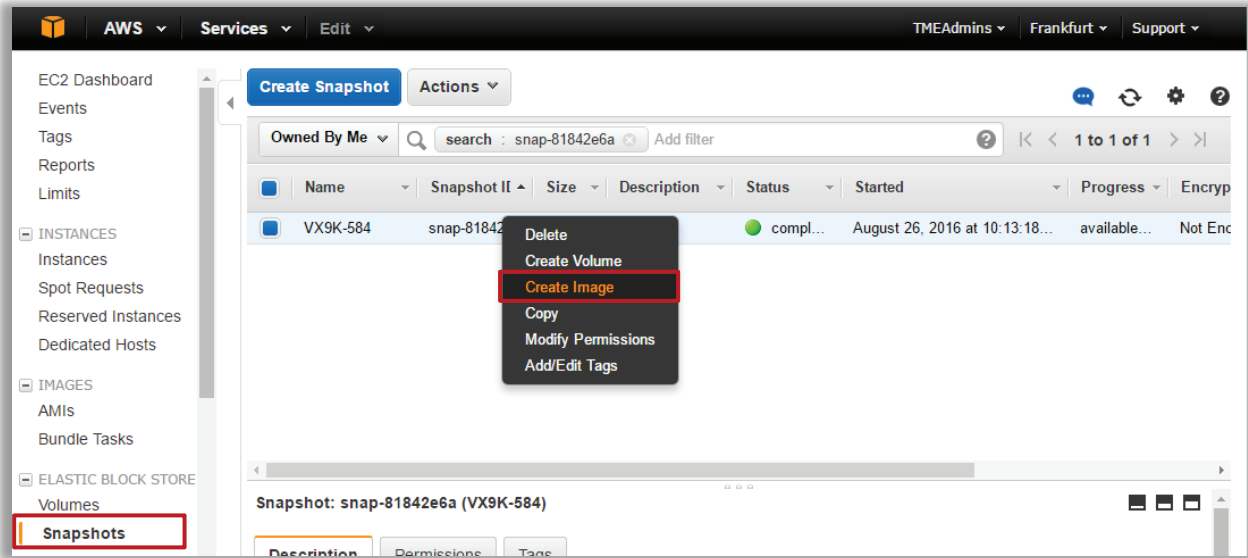
- On AWS navigate to **EC2**.



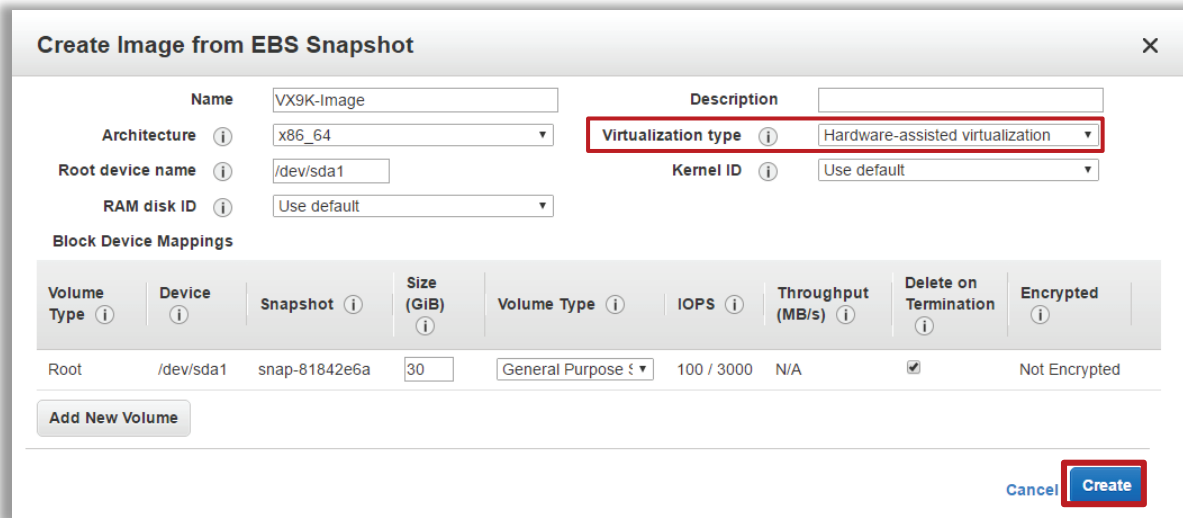
- Go to **Volumes**. Highlight imported volume and select **Create Snapshot**:




- You will be redirected to the snapshots screen. From that screen highlight newly created snapshot and select **Create Image**:



- Give a name to this new AMI Image, set Virtualization type to Hardware-assisted virtualization, provision IOPS for the disk performance as required:



- Once AMI image is created highlight it and select **Launch**.



19. Select Instance type – minimum supported instance type for HVM mode is m3.large.

The screenshot shows the AWS Management Console interface for creating an EC2 instance. The current step is 'Step 2: Choose an Instance Type'. A table lists various instance types with their specifications. The 'm4.xlarge' instance type is selected, indicated by a blue radio button and a red border around its row. At the bottom right, the 'Next: Configure Instance Details' button is also highlighted with a red border.

Instance Type	General Purpose	Memory (GiB)	Storage (GiB)	Network (Gbps)	Root Volume	Performance
t2.small	1	2	EBS only	-	Low to Moderate	
t2.medium	2	4	EBS only	-	Low to Moderate	
t2.large	2	8	EBS only	-	Low to Moderate	
m4.large	2	8	EBS only	Yes	Moderate	
<input checked="" type="radio"/> m4.xlarge	4	16	EBS only	Yes	High	
m4.2xlarge	8	32	EBS only	Yes	High	
m4.4xlarge	16	64	EBS only	Yes	High	
m4.10xlarge	40	160	EBS only	Yes	10 Gigabit	
m3.medium	1	3.75	1 x 4 (SSD)	-	Moderate	

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

20. During network configuration it is recommended to assign a specific internal IP address to the VX instance. Select desired subnet in the availability zone and then specify desired IP address under Network Interfaces section.

Step 3: Configure Instance Details
 Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances [Launch into Auto Scaling Group](#)

Purchasing option Request Spot instances

Network [Create new VPC](#)

Subnet [Create new subnet](#)
 4085 IP Addresses available

Auto-assign Public IP

Placement group

IAM role [Create new IAM role](#)

Shutdown behavior

Enable termination protection Protect against accidental termination

Monitoring Enable CloudWatch detailed monitoring
 Additional charges apply.

EBS-optimized instance Launch as EBS-optimized instance

Tenancy
 Additional charges will apply for dedicated tenancy.

Network interfaces

Device	Network Interface	Subnet	Primary IP	Secondary IP addresses
eth0	<input type="text" value="New network interface"/>	<input type="text" value="subnet-e0d30289"/>	<input type="text" value="172.31.16.85"/>	Add IP

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Storage](#)

22. On the Storage configuration you may specify desired IOPS (if instance will be running NSight or Captive Portal), otherwise leave default values:

Step 4: Add Storage
 Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encrypted
Root	/dev/sda1	snap-81842e6a	30	General Purpose	100 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypted

Add New Volume

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Cancel Previous **Review and Launch** **Next: Tag Instance**

23. Optionally add tags to the instance as required.

Step 5: Tag Instance
 A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. [Learn more](#) about tagging your Amazon EC2 resources.

Key (127 characters maximum)	Value (255 characters maximum)
Name	#VX-WING5

Create Tag (Up to 50 tags maximum)

Cancel Previous **Review and Launch** **Next: Configure Security Group**

24. Under **Security Group** configuration open necessary ports for management, i.e. SSH and HTTPS, optionally SNMP. Open other ports based on the requirements (for example MINT 24576 UDP to allow AP adoption)

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Tag Instance 6. Configure Security Group 7. Review

Step 6: Configure Security Group

You can create a new security group or select from an existing one below. [Learn more](#) about Amazon EC2 security groups.

Assign a security group: Create a **new** security group
 Select an **existing** security group

Security group name:

Description:

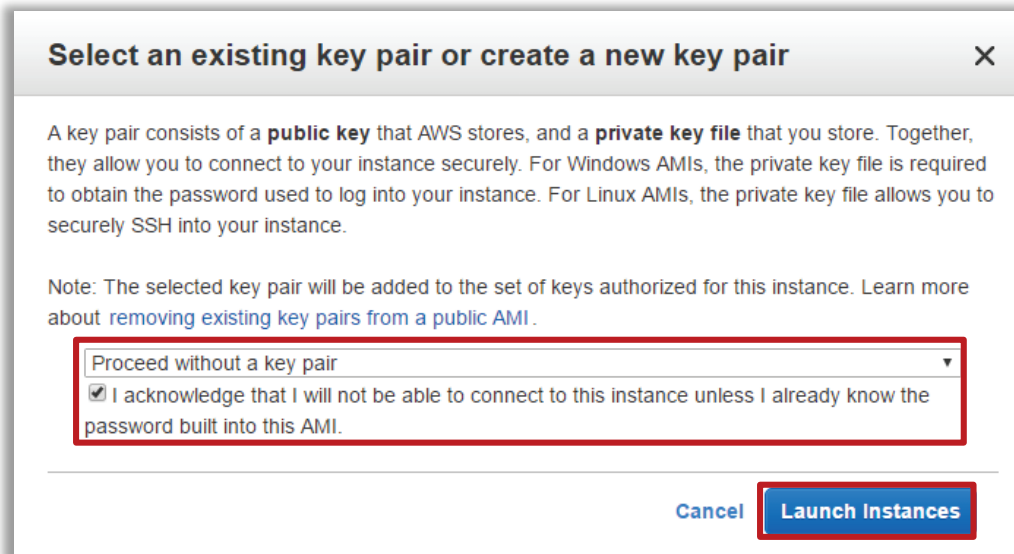
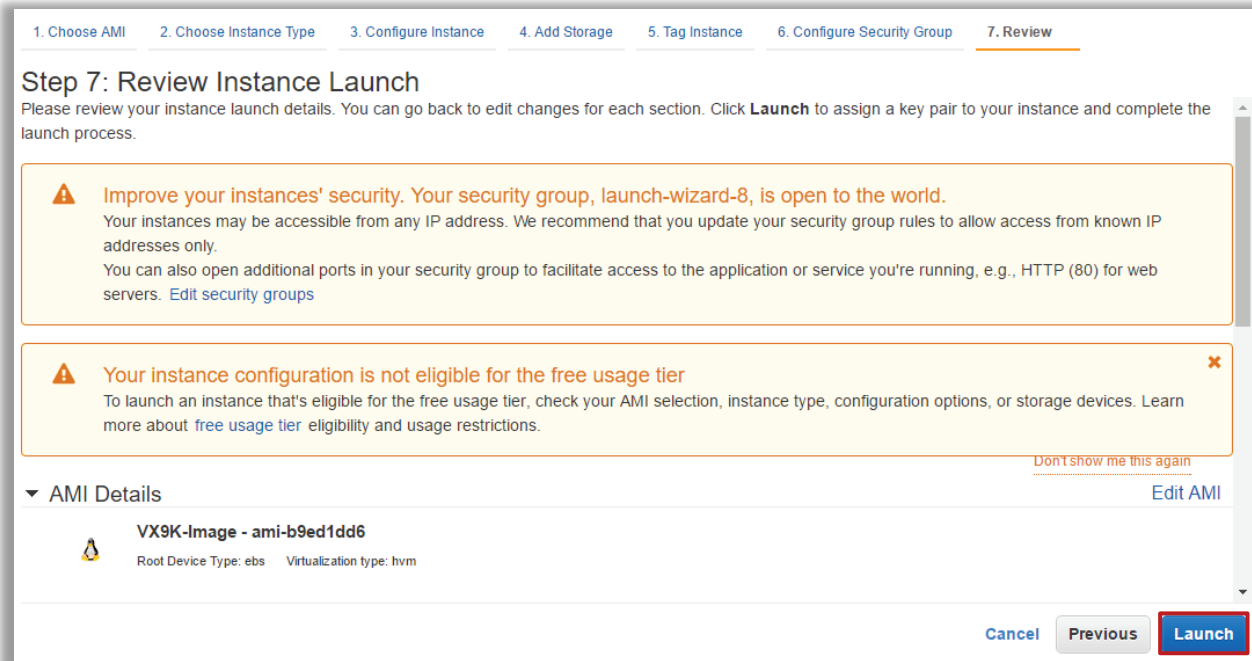
Type ⁱ	Protocol ⁱ	Port Range ⁱ	Source ⁱ
SSH	TCP	22	Anywhere 0.0.0.0/0
HTTPS	TCP	443	Anywhere 0.0.0.0/0
Custom UDP Rule	UDP	24576	Anywhere 0.0.0.0/0

Add Rule

Warning
 Rules with source of 0.0.0.0/0 allow all IP addresses to access your instance. We recommend setting security group rules to allow access from known IP addresses only.

[Cancel](#) [Previous](#) [Review and Launch](#)

26. Launch instance and select “Proceed without a key pair” when prompted. Default username and password admin/admin123 will be used to login to the VX 9000.



Licensing and Elastic IP addresses on EC2


VX 9000 licenses are bound to the **Serial Number** of the VX 9000 instance. Serial Number is automatically generated using a combination of **Base MAC address** (shown in “**show version**”) and **current IP address** of the management VLAN interface (VLAN 1 by default):

```

NSIGHT-PRIMARY#show version
VX9000 version 5.8.2.0-025R
Copyright (c) 2004-2015 Symbol Technologies, Inc. All rights reserved.
Booted from primary

NSIGHT-PRIMARY uptime is 1 days, 04 hours 28 minutes
CPU is Intel(R) Xeon(R) CPU E5-2676 v3 @ 2.40GHz
Base ethernet MAC address is 06-71-B1-5D-77-51
System serial number is 5C4E917EF1158BED
Model number is VX-9000
    
```

Amazon EC2 default behavior is to assign an IPv4 address to each instance via DHCP from a private IP range (RFC1918), and then perform NAT to a dynamically assigned public IPv4 address that will be shown under EC2 Instances Tab, for example:



This automatically assigned public IPv4 address is only reassigned in case the instance is stopped or rebooted from EC2 console (you can safely reboot the VX from WiNG without losing your assigned public IP or licenses).

In a situation where VX instance needs to be stopped for a period of time, it is important to remember that instance will get a different internal IPv4 address upon next boot, **thus all the licenses assigned previously will be lost**. To prevent this from happening it is possible to assign a static DHCP binding to always receive the same internal IPv4 address for the VX instance (this should be done at Instance Launch phase - see step #17of the previous chapter):

Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

- Number of instances:** 1 [Launch into Auto Scaling Group](#)
- Purchasing option:** Request Spot instances
- Network:** vpc-a1be72c8 (172.31.0.0/16) (default) [Create new VPC](#)
- Subnet:** subnet-e0d30289(172.31.16.0/20) | Default in eu-c [Create new subnet](#)
4089 IP Addresses available
- Auto-assign Public IP:** Use subnet setting (Enable)
- IAM role:** None [Create new IAM role](#)
- Shutdown behavior:** Stop
- Enable termination protection:** Protect against accidental termination
- Monitoring:** Enable CloudWatch detailed monitoring
[Additional charges apply.](#)
- Tenancy:** Shared - Run a shared hardware instance
[Additional charges will apply for dedicated tenancy.](#)

Network interfaces

Device	Network Interface	Subnet	Primary IP	Secondary IP addresses
eth0	New network interfac	subnet-e0d3028	172.31.16.30	Add IP

Additionally it is possible to retain the same public IPv4 address assigned in case the instance is stopped for a period of time. Amazon provides Elastic IP feature that accomplishes that goal:

AWS console screenshot showing the Elastic IP management interface. The 'Elastic IPs' menu item is highlighted in the left sidebar. The main area shows a table of allocated Elastic IP addresses with columns for IP address, Allocation ID, Instance, and Private IP Address. The 'Allocate New Address' button is highlighted in red.

Elastic IP	Allocation ID	Instance	Private IP Address
52.28.64.246	eipalloc-321df25b	i-46966187	172.31.20.10
52.29.20.233	eipalloc-c779b8ae	i-03309bbf (#NSIGHT-PRIM...	172.31.0.49
52.29.42.40	eipalloc-d979b8b0	i-9a309b26 (#NSIGHT-SEC...	172.31.2.248
52.29.67.169	eipalloc-c679b8af	i-9c309b20 (#NSIGHT-ARBI...	172.31.5.121
52.29.194.51	eipalloc-eb6cad82		

Associate Address dialog box. The 'Instance' dropdown is selected, showing 'i-a373891e (#VX-9000-TEST) (running)'. The 'Private IP Address' is set to '172.31.16.30* - 52.28.228.7'. A warning message is displayed: 'Warning: If you associate an Elastic IP address with your instance, your current public IP address is released. Learn more about public IP addresses.' The 'Associate' button is highlighted in red.