



AWS Deployment Guide

for Version 11.2



Copyright © 1994-2019 Dell Inc. or its subsidiaries. All Rights Reserved.

Contact Information

RSA Link at <https://community.rsa.com> contains a knowledgebase that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

For a list of RSA trademarks, go to www.emc.com/legal/emc-corporation-trademarks.htm#rsa.

License Agreement

This software and the associated documentation are proprietary and confidential to Dell, are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by Dell.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on RSA Link. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

April 2019

Contents

AWS Deployment Overview	5
AWS Environment Recommendations	5
Abbreviations and Other Terminology Used in this Guide	5
AWS Deployment Scenarios	8
Full NetWitness Platform Stack VPC Visibility	8
Hybrid Deployment - Decoder and Log Decoder	9
Hybrid Deployment - Decoder, Log Decoder, and Concentrator	10
Prerequisites	11
Supported Services	11
AWS Deployment	12
Rules	12
Checklist	12
Establish AWS Environment	13
Find NetWitness Platform AMIs	13
Launch an Instance and Configure a Host	13
Partition Recommendations	18
Admin Server, ESA Primary, ESA Secondary and Malware Analysis	18
Log Collector	19
Decoder	19
Log Decoder	21
Concentrator	22
Archiver	24
Endpoint Hybrid or Endpoint Log Hybrid	25
Other Partition Required	25
Installation Tasks	26
Configure Hosts (Instances) in NetWitness Platform	40
Configure Packet Capture	40
Integrate Gigamon GigaVUE with the Network Decoder	40
Integrate Ixia with the Packet Decoder	42
Integrate f5® BIG-IP with the Network Decoder	48
AWS Instance Configuration Recommendations	50
Archiver	51
Broker	52
Concentrator - Log Stream	53
Packet Stream Solutions	54

Concentrator - Gigamon Solution	54
Concentrator - f5 BIG-IP Solution	54
Decoder - Gigamon Solution	55
Decoder - f5 BIG-IP Solution	55
ESA and Context Hub on Mongo Database	57
Log Collector (Syslog, Netflow, and File Collection Protocols)	58
Log Decoder	59
NetWitness Server, Reporting Engine, Respond and Health & Wellness	60
NetWitness Endpoint Hybrid	61
UEBA	61

AWS Deployment Overview

Before you can deploy RSA NetWitness® Platform in the Amazon Web Services (AWS), you need to:

- Understand the requirements of your enterprise.
- Know the scope of a NetWitness Platform deployment.

When you are ready to begin deployment:

- Make sure that you have a NetWitness Platform "Throughput" license.
- For packet capture in AWS, you can purchase either of the following Third-Party solutions. If you engage one of these third-parties, they will assign an account representative and a professional services engineer who will work closely with RSA Support.
 - Gigamon® GigVUE
 - f5BIG-IP

AWS Environment Recommendations

AWS instances have the same functionality as the NetWitness Platform hardware hosts. RSA recommends that you perform the following tasks when you set up your AWS environment.

- Based on the resource requirements of the different components, follow best practices to use the system and dedicated storage Elastic Block Store (EBS) Volumes appropriately.
- Make sure that compute capacity provides a write speed of 10% greater than the required sustained capture and ingest rate for the deployment.
- Build Concentrator directory for index database on the Provisioned IOPS SSD.

Abbreviations and Other Terminology Used in this Guide

Abbreviation	Description
AMI	Amazon Machine Image
AWS	Amazon Web Services
BYOL	Bring your own licensing
CPU	Central Processing Unit

Abbreviation	Description
Dedicated Instance	AWS dedicated instances run in a VPC on hardware that is dedicated to a single customer. Dedicated instances are physically isolated at the host hardware level from instances that belong to other AWS accounts. Dedicated instances may share hardware with other instances from the same AWS account that are not dedicated instances. For more information on dedicated instances, see the AWS "Amazon EC2 Dedicated Instance" documentation (https://aws.amazon.com/ec2/purchasing-options/dedicated-instances/).
Elastic Block Store (EBS) Optimization	An Amazon EBS–optimized instance uses an optimized configuration stack and provides additional, dedicated capacity for Amazon EBS I/O. This optimization provides the best performance for your EBS volumes by minimizing contention between Amazon EBS I/O and other traffic from your instance. For more information on EBS-optimized instances, see the AWS "Amazon EBS–Optimized Instances" documentation (http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html).
EBS Volume	EBS volume is a highly available and reliable storage volume that you can attach to any running instance that is in the same availability zone. For more information on EBS Volumes, see the AWS "Amazon EBS Volumes" documentation (http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumes.html).
EC2 instance	Virtual server in AWS Elastic Compute Cloud (EC2) for running applications on the AWS infrastructure. See also Instance .
Enhanced Networking Enabled	Enhanced networking provides higher bandwidth, higher packet-per-second performance, and consistently lower inter-instance latencies. If your packets-per-second rate appears to have reached its ceiling, you should consider moving to enhanced networking because you have likely reached the upper thresholds of the virtual machine network interface (VIF) driver. For more information on enhanced networking, see the AWS "How do I enable and configure enhanced networking on my EC2 instances" documentation (https://aws.amazon.com/premiumsupport/knowledge-center/enable-configure-enhanced-networking/).
EPS	Events Per Second
GB	Gigabyte. 1 GB = 1,000,000,000 bytes
Gb	Gigabit. 1 Gb = 1,000,000,000 bits.
Gbps	Gigabits per second or billions of bits per second. It measures bandwidth on a digital data transmission medium such as optical fiber.
GHz	GigaHertz 1 GHz = 1,000,000,000 Hz
HDD	Hard Disk Drive
Instance	A virtual host in the AWS (that is, virtual machine or server in the AWS infrastructure on which you run services or applications). See also EC2 Instance .

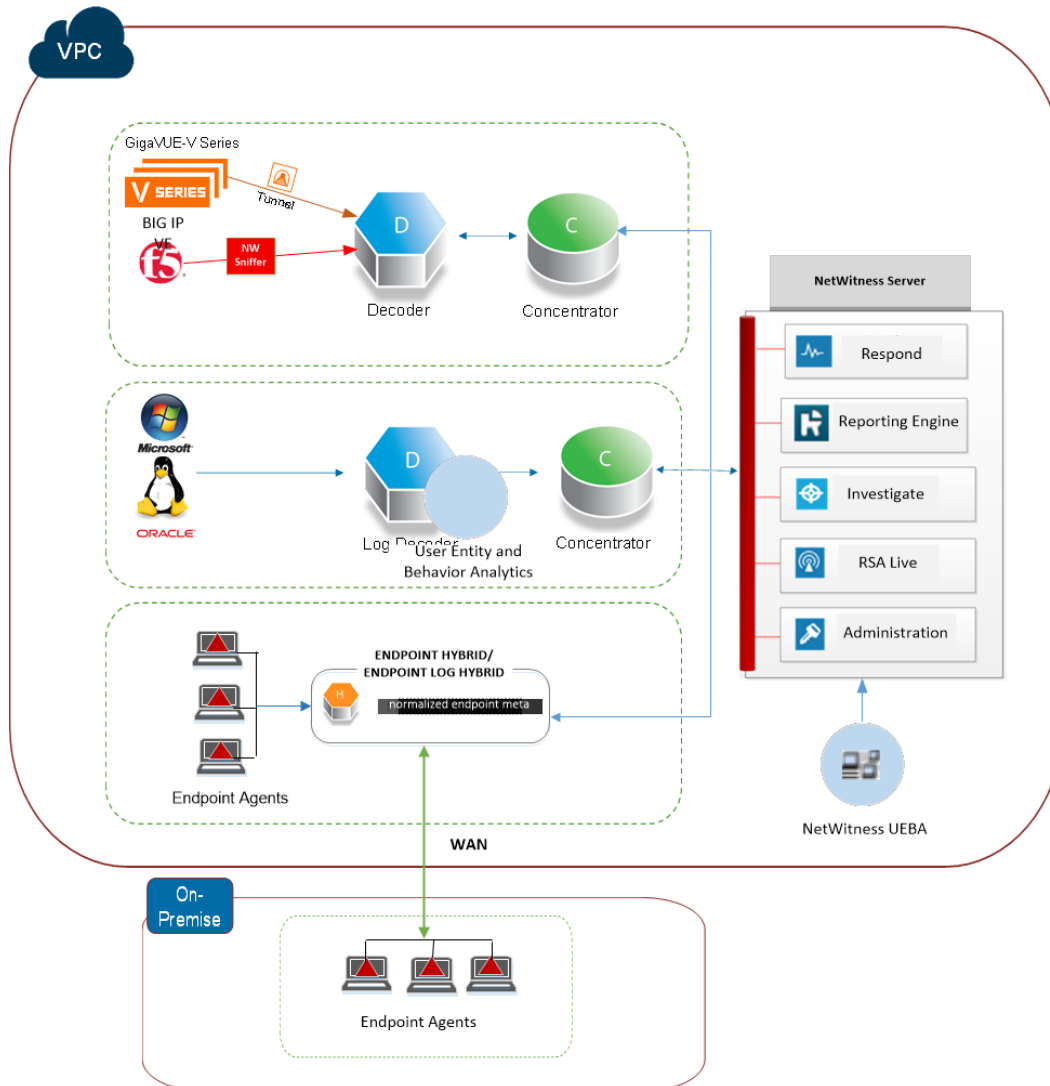
Abbreviation	Description
Instance Type	Specifies the required CPU and RAM for an instance. For more information on the instance types, see the AWS "Amazon EC2 Instance Types" documentation (https://aws.amazon.com/ec2/instance-types/).
IOPS	Input/Output Operations Per Second
Mbps	Megabits per second or millions of bits per second. It measures bandwidth on a digital data transmission medium such as optical fiber.
On-Premise	On-premise hosts are installed and run on computers on the premises (in the building) of the organization using the hosts, rather than in the AWS.
PPS	Packets Per Second
RAM	Random Access Memory (also known as memory)
Security Group	Set of firewall rules. For a comprehensive list of the ports you must set up for all NetWitness Platform components. For more information, see the "Network Architecture and Ports" documentation on RSA Link (https://community.rsa.com/docs/).
SSD	Solid-State Drive
Tag	Meaningful identifier for AWS instance.
Tap Vendor	Network Tapping Vendor
vCPU	Virtual Central Processing Unit (also known as a virtual processor)
VM	Virtual Machine
VPC	Virtual Public Cloud
vRAM	Virtual Random Access Memory (also known as virtual memory)

AWS Deployment Scenarios

The following diagrams illustrate some common AWS deployment scenarios.

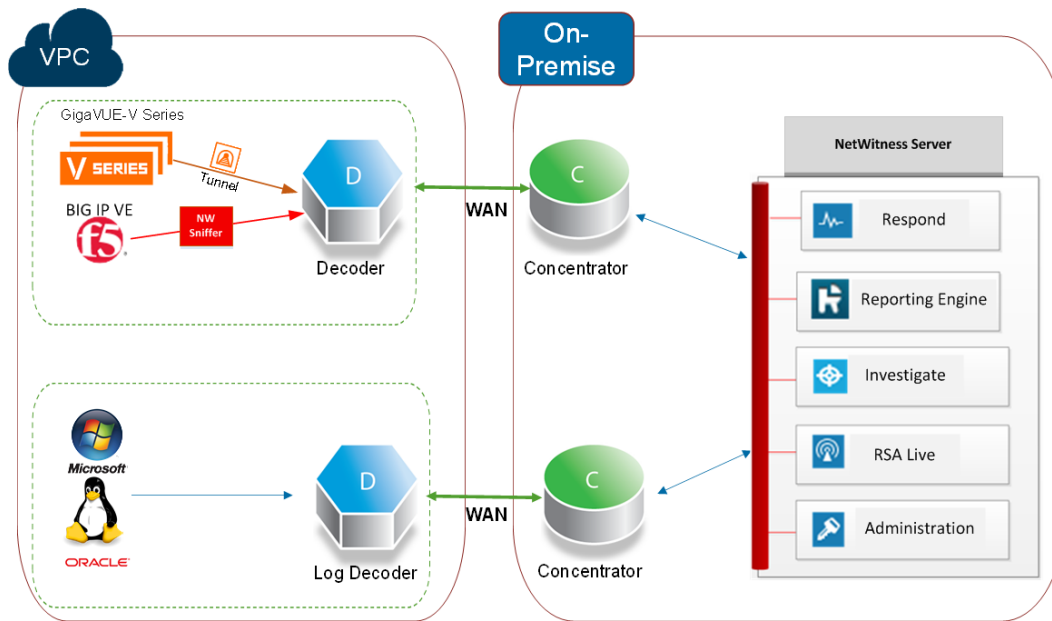
Full NetWitness Platform Stack VPC Visibility

This diagram shows all NetWitness Platform components (full stack) deployed in AWS.



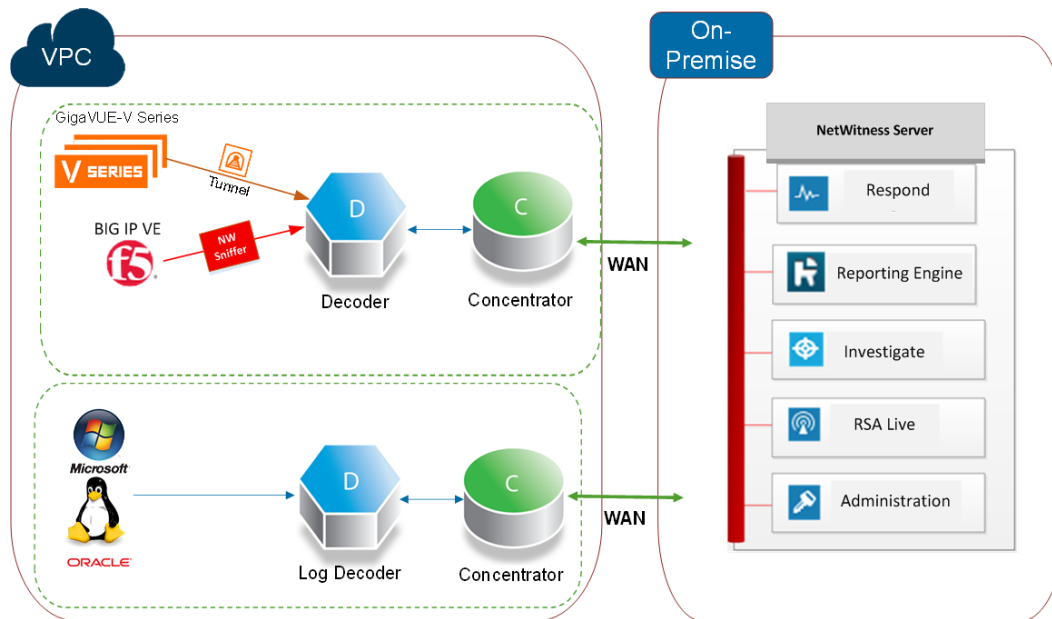
Hybrid Deployment - Decoder and Log Decoder

This diagram shows the Decoder, and Log Decoder deployed in AWS with all other NetWitness Platform components deployed on your premises.



Hybrid Deployment - Decoder, Log Decoder, and Concentrator

This diagram shows the Decoder, Log Decoder, and the Concentrator deployed in AWS with all other NetWitness Platform components deployed on your premises.



In the diagrams, the:

- **GigaVUE Series** (Gigamon® Solution) is an agent-based solution that uses **Tunneling** (implemented by the NetWitness Platform administrator) to facilitate packet data capture in AWS.
- **BIG-IP** (f5® Solution) is a load balancing solution that uses a Network Decoder acting as a sniffer (customized by the NetWitness Platform administrator) to facilitate packet capture in AWS.
- **Decoder** collects packet data. The **Decoder** captures, parses, and reconstructs all network traffic from Layers 2 – 7.
- **Log Decoder** collects logs. The **Log Decoder** collects log events from hundreds of devices and event sources.
- **Concentrator** indexes metadata extracted from network or log data and makes it available for enterprise-wide querying and real-time analytics while facilitating reporting and alerting.
- **Endpoint Hybrid** - collects endpoint data. The Endpoint Hybrid comprises of Endpoint Server, Log Decoder, and Concentrator. For more information, see *NetWitness Endpoint Insights Configuration Guide*.
- NetWitness Server hosts **Respond, Reporting, Investigate, RSA Live Content Management, Administration, Endpoint Hybrid/Log Hybrid** and other aspects of the user interface.
- **User Entity and Behavior Analytics (UEBA)** provides comprehensive user and entity behavioral analytics to better detect, investigate, and respond to advanced internal attacks and identity-based anomalies.

Prerequisites

You need the following before you begin the integration process:

- Access to AWS console
- Network rout-able (and proper AWS Security Groups) for the containers to transfer data to the NetWitness Platform Decoder.

Supported Services

RSA provides the following NetWitness Platform services.

- NetWitness Server
- Admin Server
- Archiver
- Broker
- Concentrator
- Config Server
- Event Stream Analysis
- Investigate Server
- Orchestration Server
- Reporting Engine
- Respond Server
- Security Server
- Log Decoder
- Decoder
- Remote Log Collector
- Endpoint Server
- User and Entity Behavior Analytics (UEBA)

AWS Deployment

This topic contains the rules and high-level tasks you must follow to deploy RSA NetWitness® Platform components in the AWS.

Rules

You must adhere to the following rules when deploying NetWitness Platform in AWS.

- SSH to the NetWitness Platform instance at least once after deployment to initialize the system.
- Before you enable the out-of-the-box (OOTB) dashboards, set the default data source in the Reporting Engine configuration screen.
- If you reboot the Network Decoder instance, the tunnel is not retained. Create the tunnel on Network Decoder again and restart the decoder service.
- Always use private IP addresses when you provision AWS NetWitness Platform instances.

Note: If you assign a public IP to the NetWitness Server Host, update the `/etc/nginx/conf.d/nginx.conf` configuration file as follows:

```
location /nwrpmrepo
{
alias /var/lib/netwitness/common/repo;
index index.html index.htm;
allow <Subnet-Gateway>/Subnet mask ;
#example
# allow 10.0.0.1/25;
deny all;
autoindex on;
}
```

Checklist

Step	Description
1	Establish AWS Environment
2	Find NetWitness Platform AMIs
3	Launch an Instance and Configure a Host
4	Configure Hosts (Instances) in NetWitness Platform
5	Configure Packet Capture

Establish AWS Environment

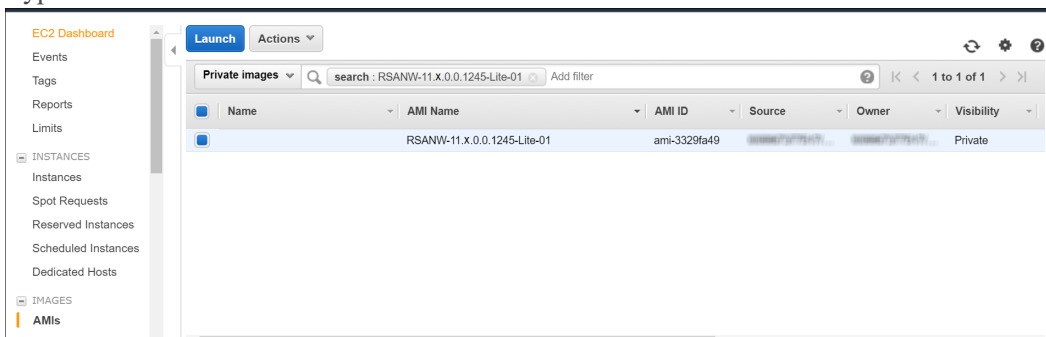
1. Make sure that you have an AWS environment with the capacity to meet or exceed the NetWitness Platform performance guidelines described in [AWS Instance Configuration Recommendations](#).
2. Go to [Find NetWitness Platform AMIs](#).

Find NetWitness Platform AMIs

Search for NW- AMI files within the Public/Shared/Community repository. Use "RSANW" for a key word to search for the AMI files.

Note: For additional instructions, see the AWS **Finding Shared AMIs** documentation (<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/usingsharedamis-finding.html>).

1. Open the Amazon EC2 console (New Subscriber Account) at <https://console.aws.amazon.com/ec2/>.
2. In the Navigation pane, choose AMIs.
3. In the first filter, choose Public images.
4. Type "RSANW" in the search field to find the NetWitness Platform AMIs.



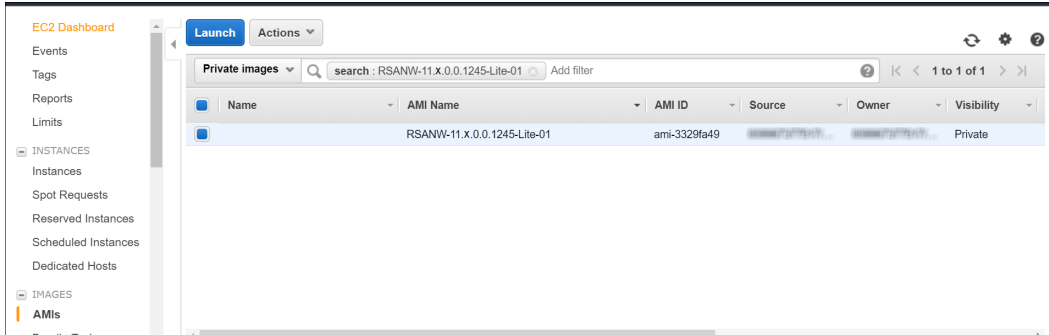
Note: Contact RSA Customer Support (<https://community.rsa.com/docs/DOC-1294>) to obtain access to the **RSANW-11.2.0.0.1245-Full-01**.

5. Go to [Launch an Instance and Configure a Host](#).

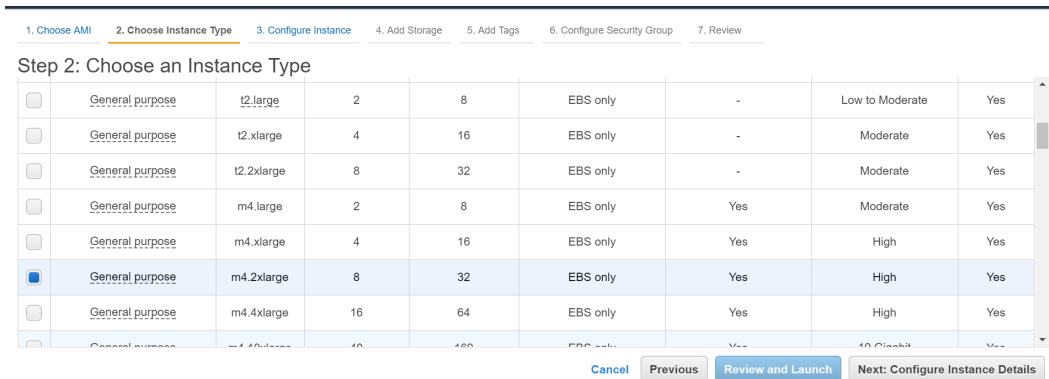
Launch an Instance and Configure a Host

Note: For additional instructions Refer to the AWS "Launching an Instance" documentation (<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/launching-instance.html>).

1. Select an instance from the grid (for example, **RSA-NW-Concentrator-11.2.0.0-01**) and click **Launch**.

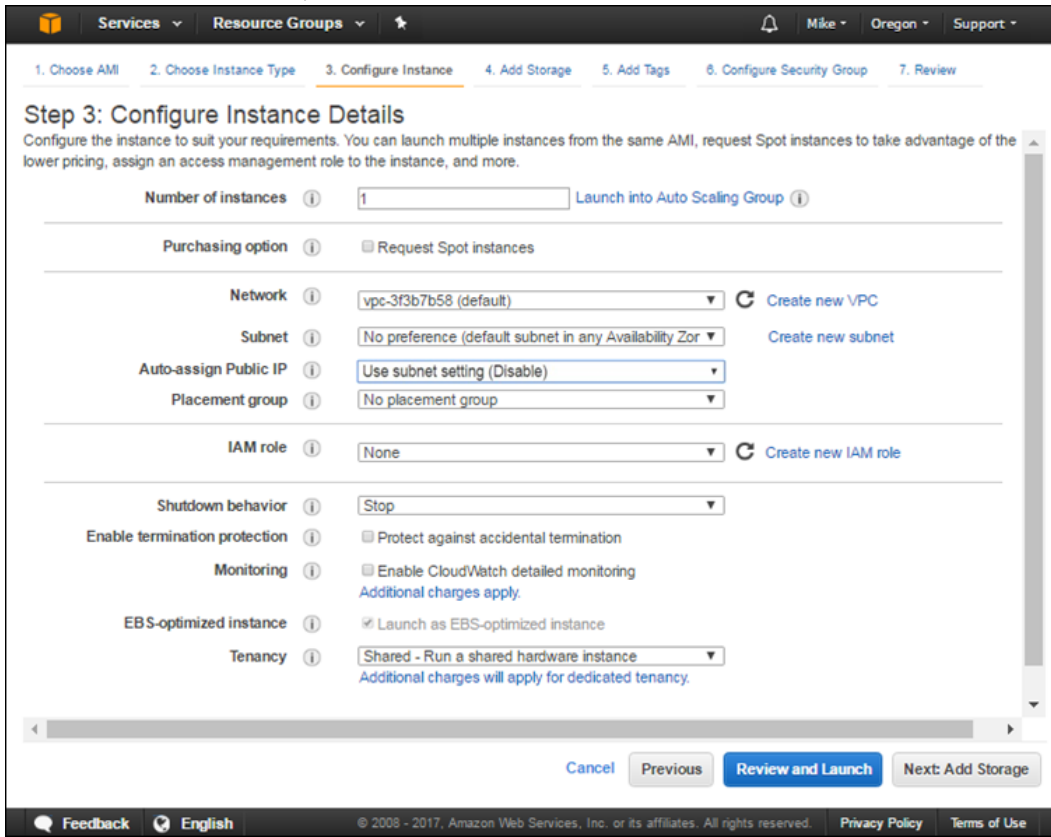


- Choose the RAM and CPUs by selecting instance type. Refer to [AWS Instance Configuration Recommendations](#) for guidelines on how to configure the EC2 Instance based on the requirements of the NetWitness Platform component (that is, service) for which you are launching an instance. The following example has the **m4.2xlarge** instance type selected with **8 CPUs** and **32 GB** of RAM.

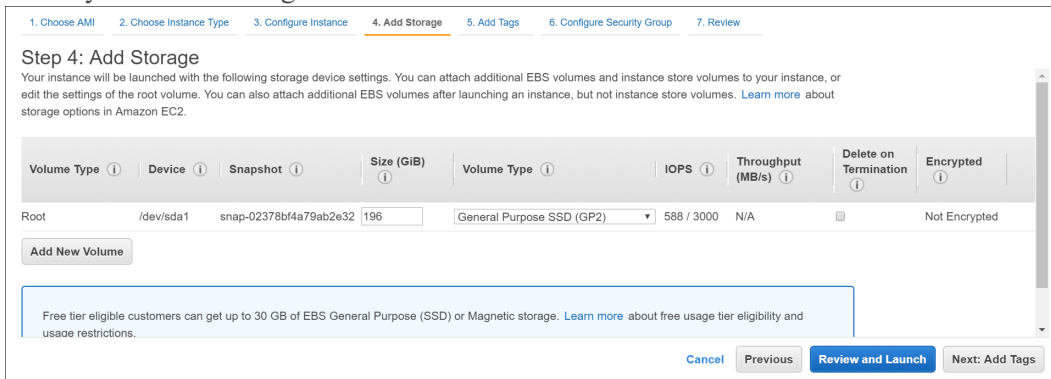


- Click **Next: Configure Instance Details** at the bottom right of the **Step 2: Choose an Instance Type** page. The **Step 3. Configure Instance Details** page is displayed.

For NetWitness Platform, the subnet and VPC are defaulted to the values in the following example.



4. Click **Next: Add Storage** at the bottom right of the **Step 3: Configure Instance Details** page. The **Step 4. Add Storage** page is displayed. Refer to [AWS Instance Configuration Recommendations](#) for guidelines on how to configure storage based on based on the requirements of the NetWitness Platform component (that is, service) for which you are launching an instance.



5. Click **Next: Add Tags** at the bottom right of the **Step 4: Add Storage** page. The **Step 5. Add Tags** page is displayed. Enter the name of your Instance.
6. Click **Next: Configure Security Group** at the bottom right of the **Step 5: Add Tags** page. The **Step 6. Configure Security Group** page is displayed.

- a. Select the Create a new security group option.
- b. Create a rule that opens all the firewall for the NetWitness Platform component.
You must configure the security group correctly to configure the instance (host) from the NetWitness Platform) User Interface and SSH to it.

Note: See the "Network Architecture and Ports" documentation in RSA Link (<https://community.rsa.com/docs/DOC-83050>) for a comprehensive list of the ports you must set up for all NetWitness Platform components..

Step 6: Configure Security Group

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. 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	Protocol	Port Range	Source
SSH	TCP	22	Custom 0.0.0.0/0
Custom TCP Rule	TCP	56005	Custom CIDR, IP or Security Group

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.

Feedback English © 2008 - 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

Note: After you configure a Security Group, you can change it at any time.

7. Click **Review and Launch** at the bottom right of the **Step 6: Configure Security Group** page. The **Step 7. Review Instance Launch** page is displayed.
8. Click **Launch** at the bottom right of the **Step 7. Review Instance Launch** page. The **Select an existing key pair or create a new key pair** dialog is displayed.
9. Choose **Proceed without key pair**.

10. Click **Launch Instance**.

AWS displays the following information as it builds the instance.

Step 6: Configure Security Group
*Select an existing security group

Security	Name	Description
sg-2fb15152	allow-all-traffic	allow-all-traffic
sg-325dfdf1	CentOS 6 (x86_64) - with Updates HVM-1602-AutogenByAWSMP	This security group was generated by AWS Marketplace and is based on recommended settings for CentOS 6 x86_64
<input checked="" type="checkbox"/>	RSA-NW-Concentrator-11.x.0.0-01	launch-wizard-1 created 2016-09-22T10:48:22.210-04:00
sg-81a282f9	default	default VPC security group
sg-8f215af5	Gigamon	launch-wizard-1 created 2016-09-22T15:33:51.373-04:00
sg-8d4602f7	launch-wizard-1	launch-wizard-1 created 2016-09-23T13:26:20.056-04:00
sg-4f32de32	launch-wizard-2	launch-wizard-2 created 2016-10-25T13:30:32.303-04:00
sg-48c0fd34	launch-wizard-3	launch-wizard-3 created 2017-02-22T12:30:46.309-05:00
sg-f8dbe182	SMTP	smtp

Inbound rules for sg-2e631b54 (Selected security groups: sg-2e631b54)

Type	Protocol	Port Range	Source
Custom TCP Rule	TCP	50120	0.0.0.0/0
Custom TCP Rule	TCP	50040	0.0.0.0/0
Custom TCP Rule	TCP	50020	0.0.0.0/0

11. Click **View Instances**.

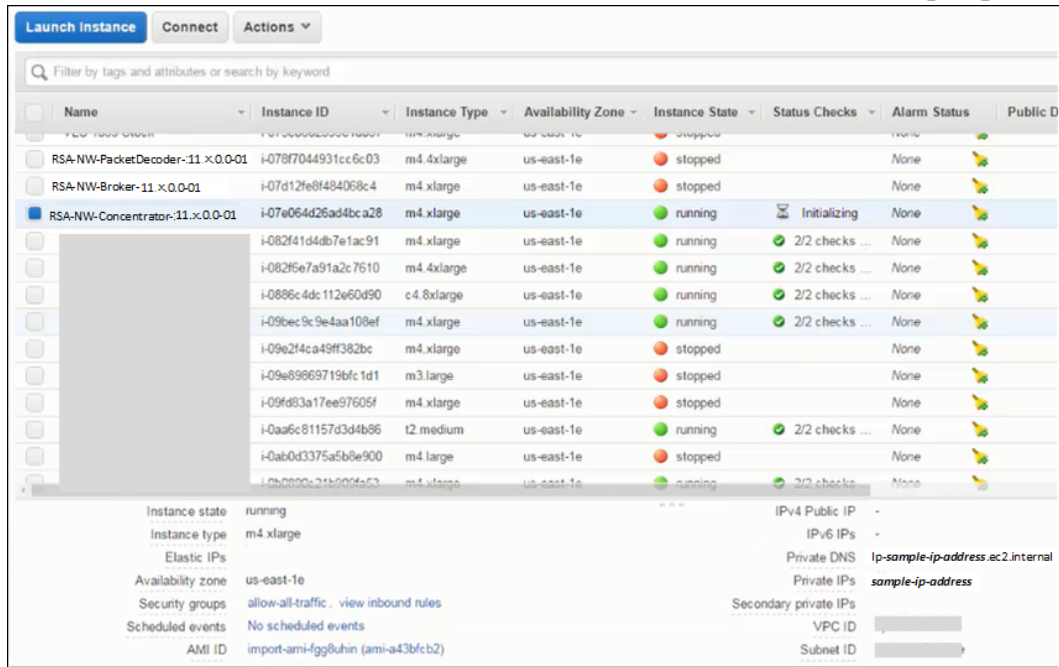
12. Select **Instances** in the left navigation panel to review all instances that AWS is initializing (for example, the **NW-Concentrator**).

Launch Instance Connect Actions

Filter by tags and attributes or search by keyword

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public
RSA-NW-PacketDecoder-11.x.0.0-01	i-078f7044931cc6c03	m4.xlarge	us-east-1e	stopped		None	
RSA-NW-Broker-11.x.0.0-01	i-07d12fe8f484068c4	m4.xlarge	us-east-1e	stopped		None	
RSA-NW-Concentrator-11.x.0.0-01	i-07e064d26ad4bca28	m4.xlarge	us-east-1e	pending	Initializing	None	
RSA-NW-Archiver-11.x.0.0-01	i-082f41d4db7ef1ac91	m4.xlarge	us-east-1e	running	2/2 checks ...	None	

The IP Address for the new **RSA-NW-Concentrator-11.2.0.0-01** host is *sample-ip-address*.



- SSH to newly-created instance using the default NetWitness Platform credentials.
- Create the recommended partitions. For more information, see [Partition Recommendations](#).
- Go to [Configure Hosts \(Instances\) in NetWitness Platform](#).

Partition Recommendations

This topic contains the recommended AWS partition.

Admin Server, ESA Primary, ESA Secondary and Malware Analysis

For an extension of `/var/netwitness/` partition, attach an external volume.

Run `lsblk` to get the physical volume name.

If you attach 2 TB disk, run the following commands:

- `pvcreate <pv_name>` (for example, `pv_name` is `/dev/sdc`)
- `vgextend netwitness_vg00 /dev/sdc`
- `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

RSA recommends the following partition. However, you can change these values based on the retention days.

LVM	Folder	EBS
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	Refer to the EBS Volume (storage) tables.

Log Collector

For an extension of `/var/netwitness/` partition, attach an external volume

Run `lsblk` to get the physical volume name.

If you attach one 500 GB volume, run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvextend -L 600G /dev/netwitness_vg00/nwhome`

RSA recommends the following partition. However, you can change these values based on the retention days.

LVM	Folder	EBS
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	Refer to the EBS Volume (storage) tables.

Decoder

For an extension of `/var/netwitness/` partition, attach an external volume and other external volumes for the Decoder database partitions.

Note: No other partition should reside on this Decoder partition and should be used only for `/var/netwitness/` partition.

Run `lsblk` to get the physical volume name.

If you attach 2 TB disk, run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

Other Partition Required

The following partition should be on the volume group **decodersmall**.

Folder	LVM	Volume Group
<code>/var/netwitness/decoder</code>	<code>decoroot</code>	<code>decodersmall</code>
<code>/var/netwitness/decoder/index</code>	<code>index</code>	<code>decodersmall</code>
<code>/var/netwitness/decoder/metadb</code>	<code>metadb</code>	<code>decodersmall</code>
<code>/var/netwitness/decoder/sessiondb</code>	<code>sessiondb</code>	<code>decodersmall</code>

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md0`
2. `vgcreate -s 32 decodersmall /dev/md0`

3. `lvcreate -L <disk_size> -n <lv_name> decodersmall`
4. `mkfs.xfs /dev/decoderssmall/<lv_name>`
5. Repeat the above steps for all the LVMs mentioned above.

The following partition should be on the volume group **decoder**.

Folder	LVM	Volume Group
<code>/var/netwitness/decoder/packetdb</code>	packetdb	decoder

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md1`
2. `vgcreate -s 32 decoder /dev/md1`
3. `lvcreate -L <disk_size> -n packetdb decoder`
4. `mkfs.xfs /dev/decoder/packetdb`

RSA recommends the following partition. However, you can change these values based on the retention days.

LVM	Folder	EBS
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/decoderssmall/decroot</code>	<code>/var/netwitness/decoder</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/decoderssmall/index</code>	<code>/var/netwitness/decoder/index</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/decoderssmall/metadb</code>	<code>/var/netwitness/decoder/metadb</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/decoderssmall/sessiondb</code>	<code>/var/netwitness/decoder/sessiondb</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/decoder/packetdb</code>	<code>/var/netwitness/decoder/packetdb</code>	Refer to the EBS Volume (storage) tables.

Create each directory and mount the LVM on it in a serial manner, except `/var/netwitness`, which is already created.

After mounting the directory, add the following entries in `/etc/fstab` in the same order:

1. `/dev/decoderssmall/decroot /var/netwitness/decoder xfs noatime,nosuid 1 2`
2. `/dev/decoderssmall/index /var/netwitness/decoder/index xfs noatime,nosuid 1 2`
3. `/dev/decoderssmall/metadb /var/netwitness/decoder/metadb xfs noatime,nosuid 1 2`

4. `/dev/decodersmall/sessiondb /var/netwitness/decoder/sessiondb xfs noatime,nosuid 1 2`
5. `/dev/decoder/packetdb /var/netwitness/decoder/packetdb xfs noatime,nosuid 1 2`

Log Decoder

For an extension of `/var/netwitness/` partition, attach an external volume and other external volumes for the Log Decoder database partitions.

Note: No other partition should reside on this Log Decoder partition and should be used only for `/var/netwitness/` partition.

Run `lsblk` to get the physical volume name.

If you attach 2 TB disk, run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

Other Partition Required

The following partition should be on the volume group **logdecodersmall**.

Folder	LVM	Volume Group
<code>/var/netwitness/logdecoder</code>	<code>decoroot</code>	<code>logdecodersmall</code>
<code>/var/netwitness/logdecoder/index</code>	<code>index</code>	<code>logdecodersmall</code>
<code>/var/netwitness/logdecoder/metadb</code>	<code>metadb</code>	<code>logdecodersmall</code>
<code>/var/netwitness/logdecoder/sessiondb</code>	<code>sessiondb</code>	<code>logdecodersmall</code>

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md0`
2. `vgcreate -s 32 logdecodersmall /dev/md0`
3. `lvcreate -L <disk_size> -n <lvm_name> logdecodersmall`
4. `mkfs.xfs /dev/logdecodersmall/<lvm_name>`
5. Repeat the above steps for all the LVMs mentioned above.

The following partition should be on the volume group **logdecoder**.

Folder	LVM	Volume Group
<code>/var/netwitness/logdecoder/packetdb</code>	<code>packetdb</code>	<code>logdecoder</code>

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md1`
2. `vgcreate -s 32 logdecoder /dev/md1`
3. `lvcreate -L <disk_size> -n packetdb logdecoder`
4. `mkfs.xfs /dev/logdecoder/packetdb`

RSA recommends the following partition. However, you can change these values based on the retention days.

LVM	Folder	EBS
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/logdecoderssmall/decoroot</code>	<code>/var/netwitness/logdecoder</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/logdecoderssmall/index</code>	<code>/var/netwitness/logdecoder/index</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/logdecoderssmall/metadb</code>	<code>/var/netwitness/logdecoder/metadb</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/logdecoderssmall/sessiondb</code>	<code>/var/netwitness/logdecoder/sessiondb</code>	Refer to the EBS Volume (storage) tables.
<code>/dev/logdecoder/packetdb</code>	<code>/var/netwitness/logdecoder/packetdb</code>	Refer to the EBS Volume (storage) tables.

Create each directory and mount the LVM on it in a serial manner, except `/var/netwitness`, which is already created.

After mounting the directory, add the following entries in `/etc/fstab` in the same order:

1. `/dev/logdecoderssmall/decoroot /var/netwitness/logdecoder xfs noatime,nosuid 1 2`
2. `/dev/logdecoderssmall/index /var/netwitness/logdecoder/index xfs noatime,nosuid 1 2`
3. `/dev/logdecoderssmall/metadb /var/netwitness/logdecoder/metadb xfs noatime,nosuid 1 2`
4. `/dev/logdecoderssmall/sessiondb /var/netwitness/logdecoder/sessiondb xfs noatime,nosuid 1 2`
5. `/dev/logdecoder/packetdb /var/netwitness/logdecoder/packetdb xfs noatime,nosuid 1 2`

Concentrator

For an extension of `/var/netwitness/` partition, attach an external disk and other external disks for the Concentrator database partitions.

Note: No other partition should reside on this Concentrator partition and should be used only for /var/netwitness/ partition.

Run `lsblk` to get the physical volume name.

If you attach 2 TB disk, run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

Other Partition Required

The following partition should be on the volume group concentrator.

Folder	LVM	Volume Group
/var/netwitness/concentrator	root	concentrator
/var/netwitness/concentrator /sessiondb	index	concentrator
/var/netwitness/concentrator /metadb	metadb	concentrator

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md0`
2. `vgcreate -s 32 logdecodersmall /dev/md0`
3. `lvcreate -L <disk_size> -n <lvm_name> logdecodersmall`
4. `mkfs.xfs /dev/logdecoderssmall/<lvm_name>`
5. Repeat the above steps all the LVMs mentioned

The following partition should be on volume group index.

Folder	LVM	Volume Group
/var/netwitness/concentrator/index	index	index

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md1`
2. `vgcreate -s 32 lindex /dev/md1`
3. `lvcreate -L <disk_size> -n index index`
4. `mkfs.xfs /dev/index/index`

RSA recommends the following partition. However, you can change these values based on the retention days.

LVM	Folder	EBS
/dev/netwitness_vg00/nwhome	/var/netwitness/	Refer to the EBS Volume (storage) tables.
/dev/concentrator/decoroot	/var/netwitness/concentrator	Refer to the EBS Volume (storage) tables.
/dev/concentrator/metadb	/var/netwitness/concentrator/metadb	Refer to the EBS Volume (storage) tables.
/dev/concentrator/sessiondb	/var/netwitness/concentrator/sessiondb	Refer to the EBS Volume (storage) tables.
/dev/index/index	/var/netwitness/concentrator/index	Refer to the EBS Volume (storage) tables.

Create each directory and mount the LVM on it in a serial manner, except `/var/netwitness`, which is already created.

After mounting the directory, add the following entries in `/etc/fstab` in the same order:

1. `/dev/concentrator/root /var/netwitness/concentrator xfs noatime,nosuid 1 2`
2. `/dev/concentrator/sessiondb /var/netwitness/concentrator/sessiondb xfs noatime,nosuid 1 2`
3. `/dev/concentrator/metadb /var/netwitness/concentrator/metadb xfs noatime,nosuid 1 2 2`
4. `/dev/index/index /var/netwitness/concentrator/index xfs noatime,nosuid 1 2`

Archiver

For an extension of `/var/netwitness/` partition, attach an external volume and other external disks for the Archiver database partitions.

Note: No other partition should reside on this Archiver partition and should be used only for `/var/netwitness/partition`.

Run `lsblk` to get the physical volume name.

If you attach 2 TB disk, run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvextend -L 1.9T /dev/netwitness_vg00/nwhome`

Other Partition Required

The following partition should be on the volume group archiver.

Folder	LVM	Volume Group
/var/netwitness/archiver	archiver	archiver

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md0`
2. `vgcreate -s 32 archiver /dev/md0`
3. `lvcreate -L <disk_size> -n archiver archiver`
4. `mkfs.xfs /dev/archiver/archiver`

RSA recommends the following partition. However, you can change these values based on the retention days.

LVM	Folder
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>
<code>/dev/archiver/archiver</code>	<code>/var/netwitness/archiver</code>

Create each directory and mount the LVM on it in a serial manner, except `/var/netwitness`, which is already created.

After mounting the directory, add the following entries in `/etc/fstab` in the same order:

1. `/dev/archiver/archiver /var/netwitness/archiver xfs noatime,nosuid 1 2`

Endpoint Hybrid or Endpoint Log Hybrid

For an extension of `/var/netwitness/` partition, attach an additional volume, and make sure that no other partition resides on this Endpoint Hybrid or Endpoint Log Hybrid. Also, attach

other additional volumes for the endpoint database partitions

Run `lsblk` to get the physical volume name.

If you attach 1 TB disk, run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvextend -L 1T /dev/netwitness_vg00/nwhome`
4. `xfs_growfs /dev/netwitness_vg00/nwhome`

Other Partition Required

The following partition should be on the volume group endpoint and should be in a single RAID 0 array.

Folder	LVM	Volume Group
<code>/var/netwitness/mongo</code>	hybrid-mongo	endpoint
<code>/var/netwitness/concentrator</code>	concentrator-concroot	endpoint
<code>/var/netwitness/concentrator/index</code>	hybrid-concindex	endpoint
<code>/var/netwitness/logdecoder</code>	hybrid-ldecroot	endpoint

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md0`
2. `vgcreate -s 32 endpoint /dev/md0`
3. `lvcreate -L <disk_size> -n <lv_name> endpoint`
4. `mkfs.xfs /dev/ endpoint /<lv_name>`
5. Repeat the above steps for all the LVMs mentioned.

RSA recommends the following partition. However, you can change these values based on the retention days.

LVM	Folder	EBS
/dev/netwitness_vg00/nwhome	/var/netwitness/	Refer to the EBS Volume (storage) tables.
/dev/endpoint/hybridmongo	/var/netwitness/mongo	Refer to the EBS Volume (storage) tables.
/dev/endpoint/concentratorconcroot	/var/netwitness/concentrator	Refer to the EBS Volume (storage) tables.
/dev/endpoint/hybridconcinde	/var/netwitness/concentrator/index	Refer to the EBS Volume (storage) tables.
/dev/endpoint/hybridldecroot	/var/netwitness/logdecoder	Refer to the EBS Volume (storage) tables.

Installation Tasks

Task 1 - Install 11.2.0.0 on the NetWitness Server (NW Server) Host

Note: You can perform this task for RSANW-11.2.0.0.1245-Full-01 instance.

1. Run the `nwsetup-tui` command to set up the host.

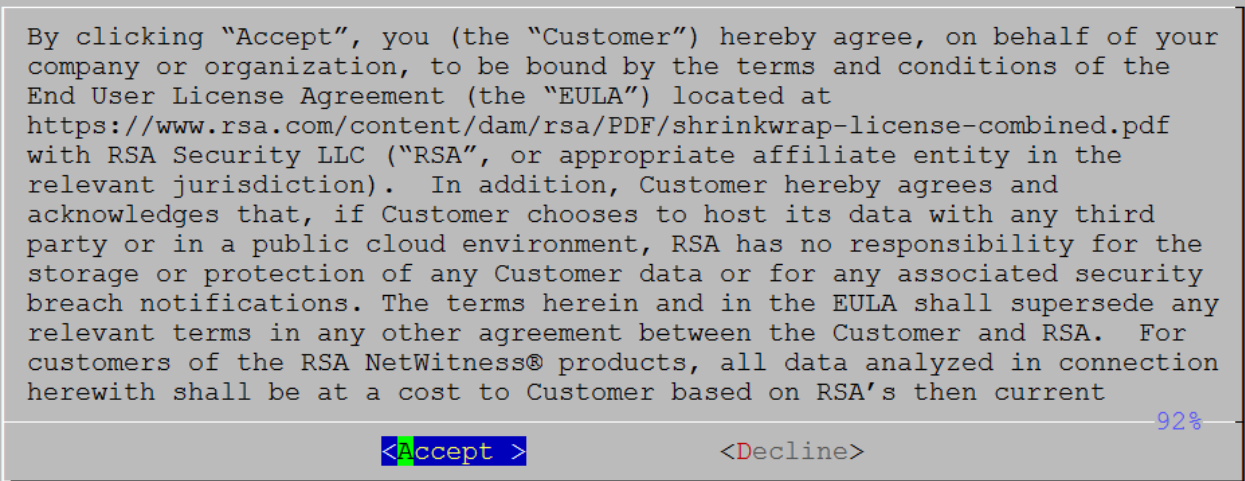
This initiates the Setup program and the EULA is displayed.

Note: 1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, use Tab key to move to and from commands (such as `<Yes>`, `<No>`, `<OK>`, and `<Cancel>`). Press **Enter** to register your command response and move to the next prompt.

2.) The Setup program adopts the color scheme of the desktop or console you use access the host.

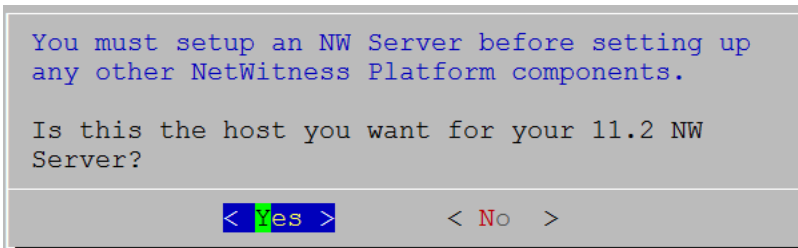
3.) If you specify DNS servers during Setup program (`nwsetup-tui`) execution, they **MUST** be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` to proceed. Any misconfigured DNS servers cause the Setup to fail. If you need to reach DNS server after setup that unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see the "Post Installation Tasks" topic in the *Physical Host Installation Guide*.

If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Platform Update Repository** prompt in step 12 (the DNS servers are not defined so the system cannot access the external repo).



2. Tab to **Accept** and press **Enter**.

The **Is this the host you want for your 11.2 NW Server** prompt is displayed.

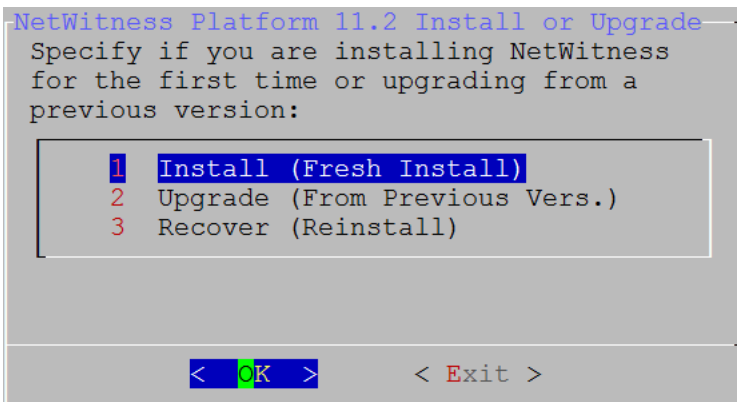


3. Tab to **Yes** and press **Enter**.

Choose **No** if you already installed 11.2 on the NW Server.

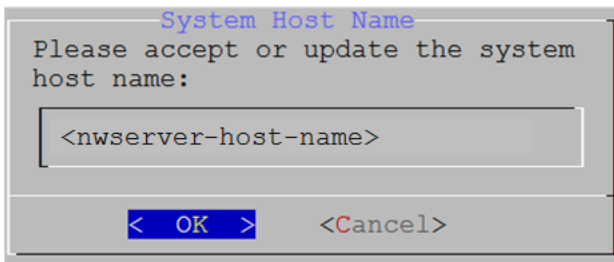
Caution: If you choose the wrong host for the NW Server and complete the Setup, you must restart the Setup Program (step 2) and complete all the subsequent steps to correct this error.

The **Install** or **Upgrade** prompt is displayed. (**Recover** does not apply to the installation. It is for 11.2 Disaster Recovery.).



4. Press **Enter**. **Install (Fresh Install)** is selected by default.

The **Host Name** prompt is displayed.



Caution: If you include "." in a host name, the host name must also include a valid domain name.

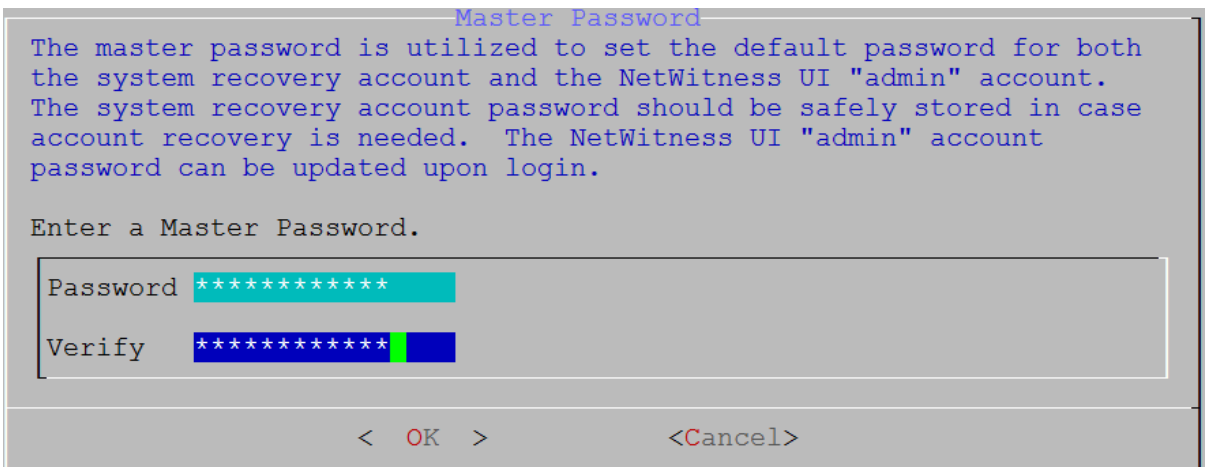
5. Press **Enter** if you want to keep this name. If you do not edit the host name, tab to **OK**, and press **Enter** to change it.

The **Master Password** prompt is displayed.

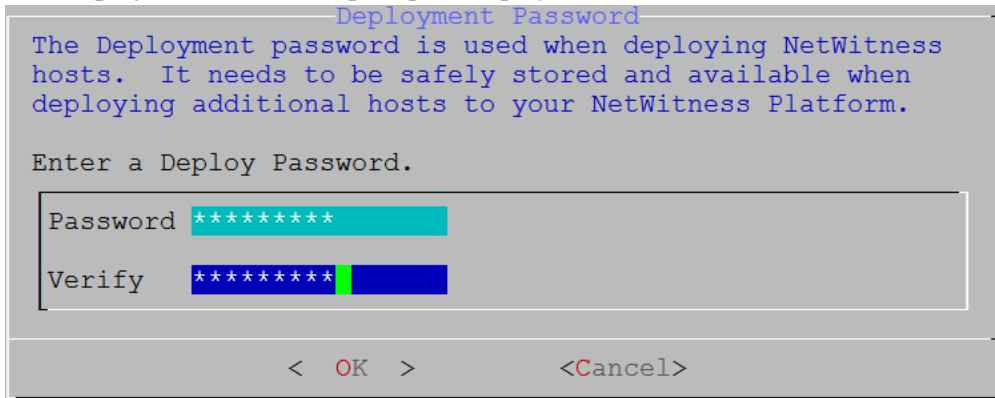
The following list of characters are supported for Master Password and Deployment Password:

- Symbols : ! @ # % ^ + ,
- Numbers : 0-9
- Lowercase Characters : a-z
- Uppercase Characters : A-Z

No ambiguous characters are supported for Master Password and Deployment Password (for example: space { } [] () / \ ' " ` ~ ; : . < > -).



- Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**. The **Deployment Password** prompt is displayed.



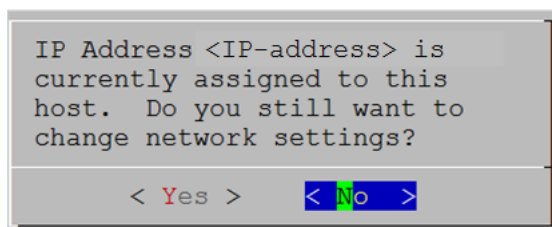
```
Deployment Password
The Deployment password is used when deploying NetWitness
hosts. It needs to be safely stored and available when
deploying additional hosts to your NetWitness Platform.

Enter a Deploy Password.

Password *****
Verify *****

< OK >      <Cancel>
```

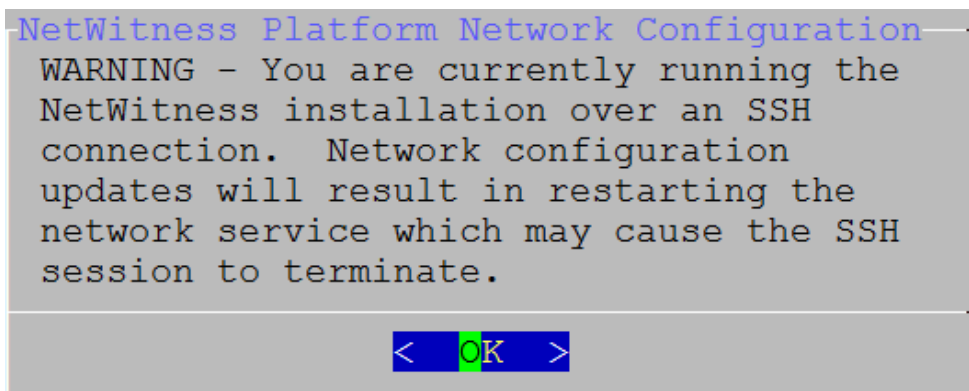
- Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**. If:
 - The Setup program finds a valid IP address for this host, the following prompt is displayed.



```
IP Address <IP-address> is
currently assigned to this
host. Do you still want to
change network settings?

< Yes >      < No >
```

- Press **Enter** if you want to use this IP and avoid changing your network settings. Tab to **Yes** and press **Enter** if you want to change the IP configuration found on the host.
- If you are using an SSH connection, the following warning is displayed. Press **Enter** to close warning prompt.

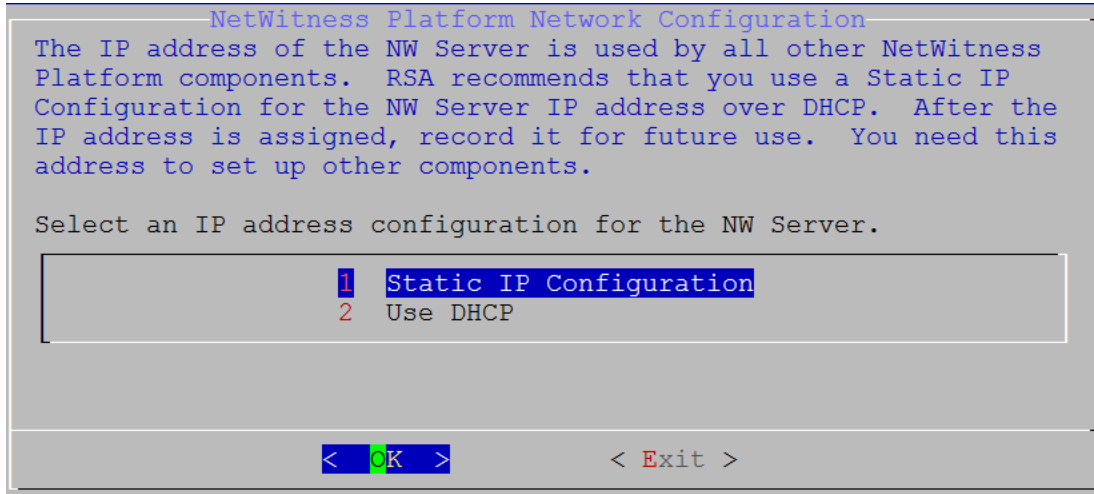


```
NetWitness Platform Network Configuration
WARNING - You are currently running the
NetWitness installation over an SSH
connection. Network configuration
updates will result in restarting the
network service which may cause the SSH
session to terminate.

< OK >
```

Note: If you connect directly from the host console, the following warning will not be displayed.

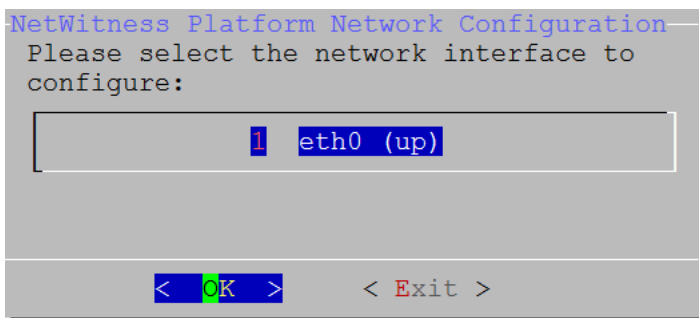
- If the Setup Program found an IP configuration and you chose to use it, the Update Repository prompt is displayed. Go to step 10 to and complete the installation.
- If the Setup Program did not find an IP configuration or if you chose to change the existing IP configuration, the Network Configuration prompt is displayed.



8. Tab to **OK** and press **Enter** to use **Static IP**.

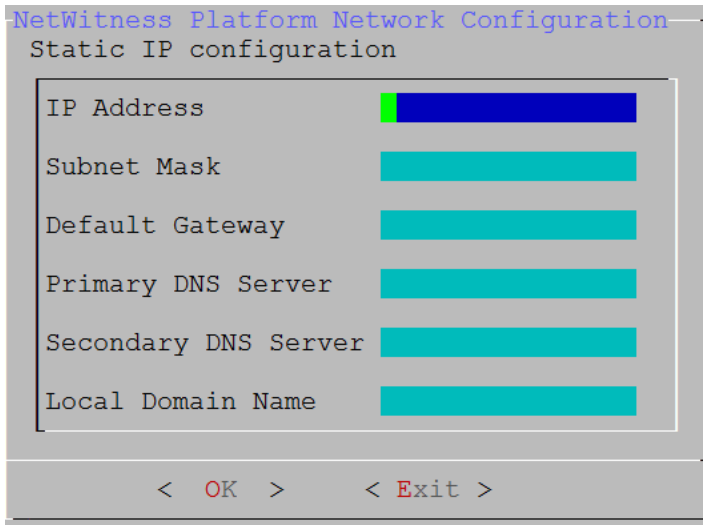
If you want to use **DHCP**, down arrow to 2 Use DHCP and press **Enter**.

The **Network Configuration** prompt is displayed.



9. Down arrow to the network interface you want, tab to **OK**, and press **Enter**. If you do not want to continue, tab to **Exit**.

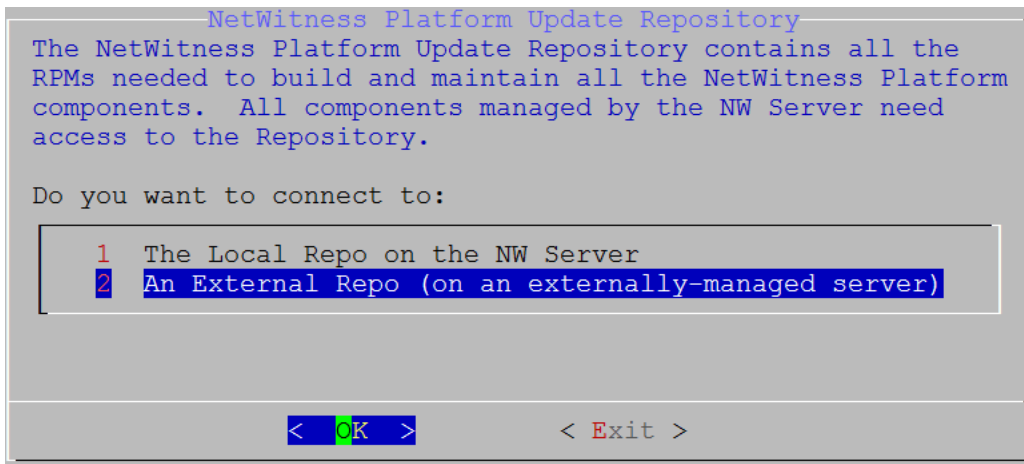
The **Static IP Configuration** prompt is displayed.



10. Type the configuration values (using the down arrow to move from field to field), tab to **OK**, and press **Enter**.
 If you do not complete all the required fields, an **All fields are required** error message is displayed (**Primary DNS Server**, **Secondary DNS Server**, and **Local Domain Name** fields are not required.)
 If you use the wrong syntax or character length for any of the fields, an **Invalid field-name** error message is displayed.

Caution: If you select DNS Server, make sure that the DNS Server is correct and the host can access it before proceeding with the install.

The **Update Repository** prompt is displayed.



11. Apply the standard firewall configuration, press **Enter**.
 - Disable the standard configuration, tab to **Yes** and press **Enter**.
 The Disable firewall prompt is displayed.

```

Disable Firewall
Do you need to apply custom
firewall rules to this host?
("No" enforces the standard
NetWitness firewall rule set to
the host)

< Yes > < No >

```

The disable firewall configuration confirmation prompt is displayed.

```

Warning: you chose to disable the default NetWitness
firewall configuration which means you must set up
firewall rules manually.

Select "Yes" to confirm that you will set up firewall
rules manually.

< Yes > < No >

```

- Tab to **Yes** and press **Enter** to confirm (press **Enter** to use standard firewall configuration).
12. Press **Enter** to install 11.2 on the NW Server.

The **Start Install** prompt is displayed.

```

Start Install/Upgrade
All the required information has been gathered.

Select "1 Install Now" to start the installation
on this host.

1 Install Now
2 Restart

< OK > < Exit >

```

When **Installation complete** is displayed, you have installed the 11.2 NW Server on this host.

Note: Ignore the hash code errors similar to the errors shown in the following screen shot that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.


```
ValueError: error:3207A06D:lib(50):B_HASH_init:cr new
Checksum type 'md5' disabled
(skipped due to only_if)
* file[/etc/yum.repos.d/CentOS-Base.repo] action delete (up to date)
* ruby_block[yum-cache-reload-CentOS-Base] action nothing (skipped due to action :nothing)
(up to date)
* yum_repository[Remove CentOS-CR repository] action delete
* execute[yum clean all CentOS-CR] action runERROR:root:code for hash md5 was not found.
Traceback (most recent call last):
File "/usr/lib64/python2.7/hashlib.py", line 129, in <module>
    globals()[__func_name] = __get_hash(__func_name)
File "/usr/lib64/python2.7/hashlib.py", line 98, in __get_openssl_constructor
    f(usedforsecurity=False)
```

Task 2 - Install 11.2 on Other Component Hosts

Note: You can perform this task for RSANW-11.2 1245-Lite-01 instance.

1. Run the `nwsetup-tui` command to set up the host.

This initiates the Setup program and the EULA is displayed.

Note: 1.) When you navigate through the Setup program prompts, use the down and up arrows to move among fields, use Tab key to move to and from commands (such as `<Yes>`, `<No>`, `<OK>`, and `<Cancel>`). Press **Enter** to register your command response and move to the next prompt.
 2.) The Setup program adopts the color scheme of the desktop or console you use access the host.
 3.) If you specify DNS servers during Setup program (`nwsetup-tui`) execution, they **MUST** be valid (valid in this context means valid during setup) and accessible for the `nwsetup-tui` to proceed. Any misconfigured DNS servers cause the Setup to fail. If you need to reach DNS server after setup that unreachable during setup, (for example, to relocate a host after setup that would have a different set of DNS Servers), see the "Post Installation Tasks" topic in the *Physical Host Installation Guide*.
 If you do not specify DNS Servers during setup (`nwsetup-tui`), you must select **1 The Local Repo (on the NW Server)** in the **NetWitness Platform Update Repository** prompt in step 12 (the DNS servers are not defined so the system cannot access the external repo).

```
By clicking "Accept", you (the "Customer") hereby agree, on behalf of your
company or organization, to be bound by the terms and conditions of the
End User License Agreement (the "EULA") located at
https://www.rsa.com/content/dam/rsa/PDF/shrinkwrap-license-combined.pdf
with RSA Security LLC ("RSA", or appropriate affiliate entity in the
relevant jurisdiction). In addition, Customer hereby agrees and
acknowledges that, if Customer chooses to host its data with any third
party or in a public cloud environment, RSA has no responsibility for the
storage or protection of any Customer data or for any associated security
breach notifications. The terms herein and in the EULA shall supersede any
relevant terms in any other agreement between the Customer and RSA. For
customers of the RSA NetWitness® products, all data analyzed in connection
herewith shall be at a cost to Customer based on RSA's then current
```

92%

`<Accept >``<Decline>`

2. Tab to **Accept** and press **Enter**.

The **Is this the host you want for your 11.2 NW Server** prompt is displayed.


```
You must setup an NW Server before setting up
any other NetWitness Platform components.
```

```
Is this the host you want for your 11.2 NW
Server?
```

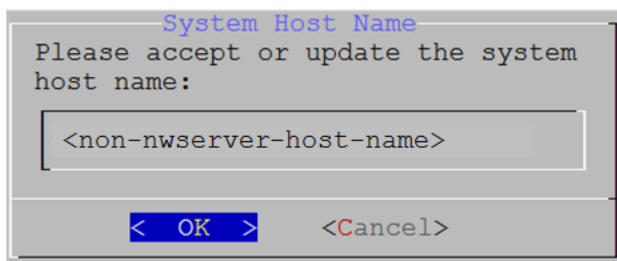
`< Yes >``< No >`

Caution: If you choose the wrong host for the NW Server and complete the Setup, you must restart the Setup Program (step 2) and complete all the subsequent steps to correct this error.

3. Press **Enter**(No).



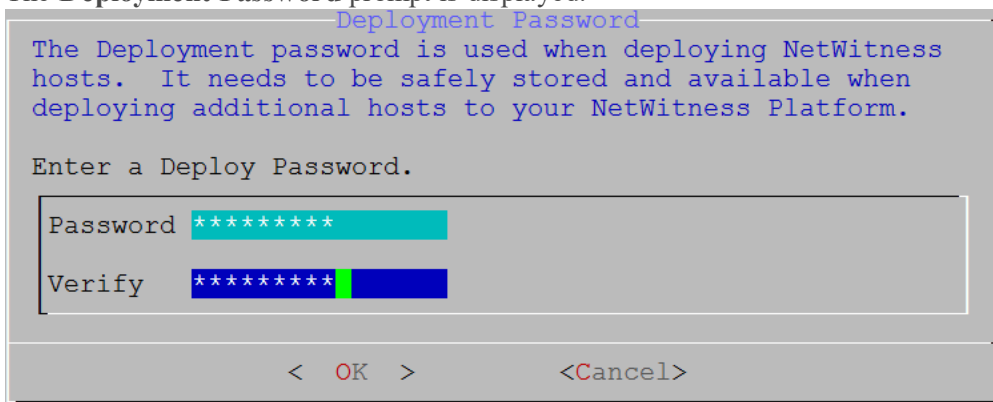
4. Press **Enter**. **Install (Fresh Install)** is selected by default. The **Host Name** prompt is displayed.



Caution: If you include "." in a host name, the host name must also include a valid domain name.

5. If want to keep this name, press **Enter**. If you want to change this name, edit it, tab to **OK**, and press **Enter**.

The **Deployment Password** prompt is displayed.



6. Type in the **Password**, down arrow to **Verify**, retype the password, tab to **OK**, and press **Enter**.
- If the Setup program finds a valid IP address for this host, the following prompt is displayed.

```
IP Address <IP-address> is
currently assigned to this
host. Do you still want to
change network settings?

< Yes > < No >
```

- Press **Enter** if you want to use this IP and avoid changing your network settings.
- Tab to **Yes** and press **Enter** if you want to change the IP configuration found on the host. If you are using an SSH connection, the following warning is displayed. Press **Enter** to close warning prompt.

```
NetWitness Platform Network Configuration
WARNING - You are currently running the
NetWitness installation over an SSH
connection. Network configuration
updates will result in restarting the
network service which may cause the SSH
session to terminate.

< OK >
```

The Setup Program found an IP configuration and you chose to use it, the Update Repository prompt is displayed. Go to step 10 to and complete the installation.

The Setup Program did not find an IP configuration or if you chose to change the existing IP configuration, the Network Configuration prompt is displayed.

```
NetWitness Platform Network Configuration
The IP address of the NW Server is used by all other NetWitness
Platform components. RSA recommends that you use a Static IP
Configuration for the NW Server IP address over DHCP. After the
IP address is assigned, record it for future use. You need this
address to set up other components.


Select an IP address configuration for the NW Server.

1 Static IP Configuration
2 Use DHCP

< OK > < Exit >
```

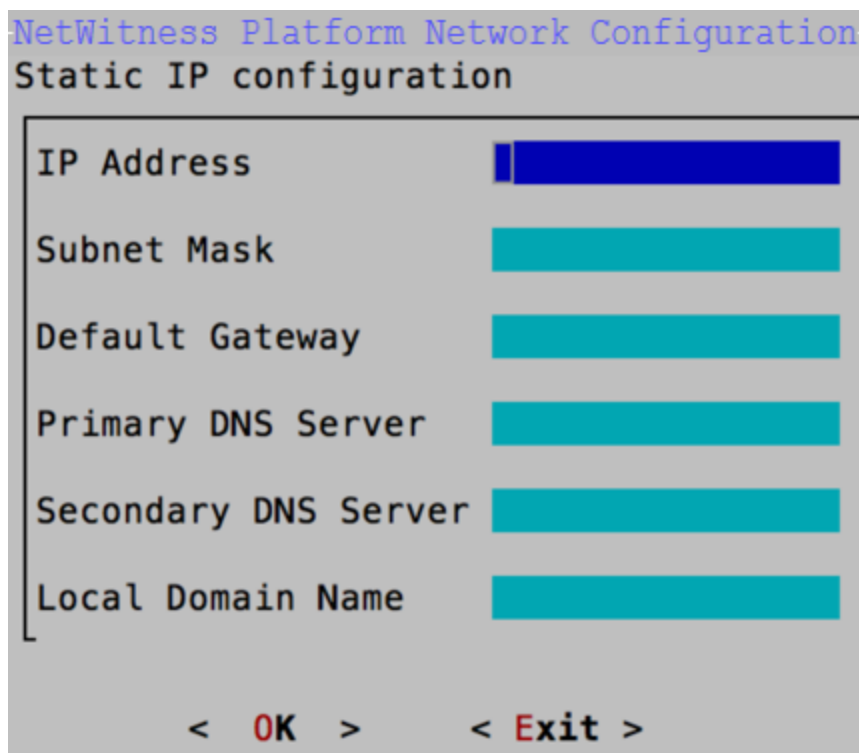
7. Tab to **OK** and press **Enter** to use **Static IP**.
If you want to use **DHCP**, down arrow to 2 Use DHCP and press **Enter**.

The **Network Configuration** prompt is displayed.



8. Down arrow to the network interface you want, tab to **OK**, and press **Enter**. If you do not want to continue, tab to **Exit**.

The **Static IP Configuration** prompt is displayed.



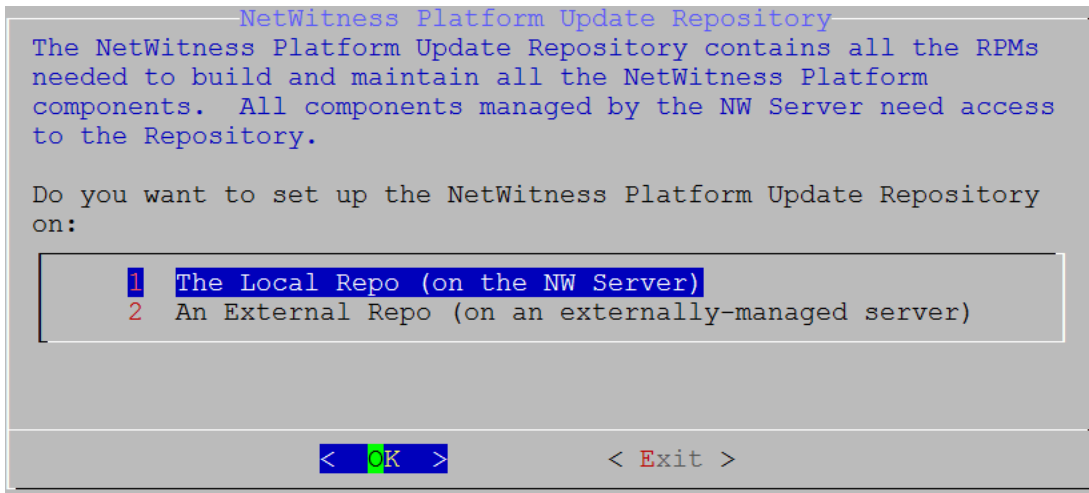
9. Type the configuration values (using the down arrow to move from field to field), tab to **OK**, and press **Enter**.

If you do not complete all the required fields, an **All fields are required** error message is displayed (**Primary DNS Server**, **Secondary DNS Server**, and **Local Domain Name** fields are not required.) If you use the wrong syntax or character length for any of the fields, an **Invalid field-name** error message is displayed.

Caution: If you select DNS Server, make sure that the DNS Server is correct and the host can access it before proceeding with the install.

10. The **Update Repository** prompt is displayed. Press **Enter** to choose the **Local Repo** on the NW

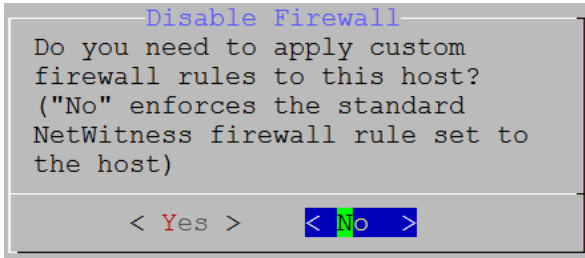
Server.



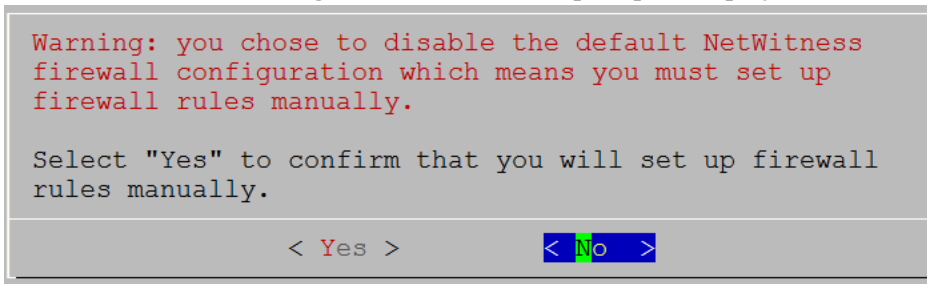
11. To:

- Apply the standard firewall configuration, press **Enter**.
- Disable the standard configuration, tab to **Yes** and press **Enter**.

The Disable firewall prompt is displayed.




The disable firewall configuration confirmation prompt is displayed.



- Tab to **Yes** and press **Enter** to confirm (press **Enter** to use standard firewall configuration).

12. The **Start Install** prompt is displayed.



13. Press **Enter** to install 11.2 on the NW Server.

When **Installation Complete** is displayed, you have installed the 11.2 NW Server on this host.

Note: Ignore the hash code errors similar to the errors shown in the following screen shot that are displayed when you initiate the `nwsetup-tui` command. Yum does not use MD5 for any security operations so they do not affect the system security.

```

ValueError: error:3207A06D:lib(50):B_HASH_init:cr new
Checksum type 'md5' disabled
(skipped due to only_if)
 * file[/etc/yum.repos.d/CentOS-Base.repo] action delete (up to date)
 * ruby_block[yum-cache-reload-CentOS-Base] action nothing (skipped due to action :nothing)
  (up to date)
 * yum_repository[Remove CentOS-CR repository] action delete
 * execute[yum clean all CentOS-CR] action runERROR:root:code for hash md5 was not found.
Traceback (most recent call last):
  File "/usr/lib64/python2.7/hashlib.py", line 129, in <module>
    globals()[__func_name] = __get_hash(__func_name)
  File "/usr/lib64/python2.7/hashlib.py", line 98, in __get_openssl_constructor
    f(usedforsecurity=False)

```

Configure Hosts (Instances) in NetWitness Platform

Configure individual hosts and services as described in RSA NetWitness® Platform *Host and Services Configuration Guide*. This guide also describes the procedures for applying updates and preparing for version upgrades.

Note: After you successfully launch an instance, AWS assigns a default hostname to it. See the "Change the Name and Hostname of a Host" documentation in RSA Link (<https://community.rsa.com>) for instructions on changing a hostname.

Configure Packet Capture

You can integrate any of the following Third-Party solutions with the Network Decoder to capture packets in the AWS cloud:

- [Gigamon® GigaVUE](#)
- [Ixia CloudLens™](#)
- [f5® BIG-IP](#)

Integrate Gigamon GigaVUE with the Network Decoder

There are two main tasks to configure the Gigamon® third-party Tap vendor packet capture solution:

Task 1. Integrate the Gigamon Solution

Gigamon® Visibility Platform on AWS is available through the AWS Marketplace and activated by a BYOL license. A thirty-day free trial is also available.

For more information on the Gigamon® solution refer to the "Gigamon® Visibility Platform for AWS Data Sheet" (<https://www.gigamon.com/sites/default/files/resources/datasheet/ds-gigamon-visibility-platform-for-aws-4095.pdf>).

For deployment details, see the "Gigamon® Visibility Platform for AWS Getting Started Guide" (<https://www.gigamon.com/sites/default/files/resources/deployment-guide/dg-visibility-platform-for-aws-getting-started-guide-4111.pdf>).

After the “Monitoring Session” is deployed within the Gigamon GigaVUE-FM, you can configure the Network Decoder Tunnel.

Task 2. Configure Tunnel on the Network Decoder

1. SSH to the Decoder.

2. Enter the following command strings.

```
$ sudo ip link add tun0 type gretap local any remote <ip_address_of_VSERIES_NODE_TUNNEL_INTERFACE> ttl 255 key 0
```

```
$ sudo ip link set tun0 up mtu <MTU-SIZE>
```

```
$ sudo ifconfig (to verify if the tunnel tun0 is being listed in the list of interfaces)
```

```
$ sudo lsmod | grep gre ( to make sure if the below kernel modules are running:
```

```
ip_gre 18245 0
```

```
ip_tunnel 25216 1)
```

If they are not running then execute the below commands to enable the modules

```
$ sudo modprobe act_mirred
```

```
$ sudo modprobe ip_gre
```

3. Create a firewall rule in the Network Decoder to allow traffic through the tunnel.

a. Open the iptables file.

```
vi /etc/sysconfig/iptables
```

b. Append the line `-A INPUT -p gre -j ACCEPT` before the commit statement

c. Restart iptables by executing the following commands.

```
service iptables restart
```

4. Set the interface in the Network Decoder.

a. Log in to NetWitness Platform, select the `decoder/config` node in Explorer view for the Network Decoder service.

b. Set the `capture.selected = packet_mmap_, tun0`.

Parameter	Value
capture.autostart	off
capture.buffer.size	128 MB
capture.device.params	
capture.selected	packet_mmap_, tun0 (bpf)
export.cache.expire	60
export.packet.enabled	no

5. (Conditional) - If you have multiple tunnels on the Network Decoder.

a. Restart Decoder service after you create the tunnel in Network Decoder.

b. Log in to NetWitness Platform, select the `decoder/config` node in Explorer view for the Network Decoder service, and set the following parameters.

```
capture.device.params = interfaces=tun0,tun1,tun2
```

```
capture.selected = packet_mmap_,All
```

The screenshot shows the PacketDecoder configuration interface. On the left, a tree view shows the 'config' folder expanded. On the right, a table lists configuration parameters. The 'capture.selected' parameter is highlighted with a red box and is set to 'packet_mmap_ALL'. Other parameters include 'capture.buffer.size' (128 MB), 'capture.device.params' (interfaces=tun0,tun1,tun2), and 'capture.autostart' (off).

Parameter	Value
/decoder/config	PacketDecoder (Decoder)
capture.autostart	off
capture.buffer.size	128 MB
capture.device.params	interfaces=tun0,tun1,tun2
capture.selected	packet_mmap_ALL
export.cache.expire	60
export.packet.enabled	no

- Restart decoder service.

```
$ sudo restart nwdecoder
```

The user should be all set to capture the network traffic in Decoder.

Integrate Ixia with the Packet Decoder

You must complete the following tasks to integrate the Packet Decoder with Ixia CloudLens.

[Task 1. Deploy Client Machines](#)

[Task 2. Create CloudLens Project](#)

[Task 3. Install Docker Container on Decoder](#)

[Task 4. Install Docker Container on Clients](#)

[Task 5. Map Packet Decoder to Ixia Clients](#)

[Task 6. Validate CloudLens Packets Arriving at Decoder](#)

[Task 7. Set Interface in Packet Decoder](#)

Task 1. Deploy Client Machines

- Deploy client machines onto which you want to route the traffic to the Packet Decoder. See the Ixia CloudLens documentation (<https://www.ixia.cloud/help/Default.htm>) for specifications needed for supported client machines.
- For Client Machines (as well as Decoder machine) the following ports must be opened on AWS Security Group Inbound Rules; UDP 19993 from all, TCP 22 from Admin IP.

Task 2. Create CloudLens Project

Complete the following steps to create a new project and get your project key.

- Get Cloudlens login credentials and access to a free trial. You can create an Ixia login account at <https://www.ixiacom.com/products/cloudlens-trial-a>.
- Go to the Cloudlens public site (<https://www.ixia.cloud>).
- Click + (add) to create a new project with a name of your choosing (for example, **Netwitness-IxiaIntegration**).

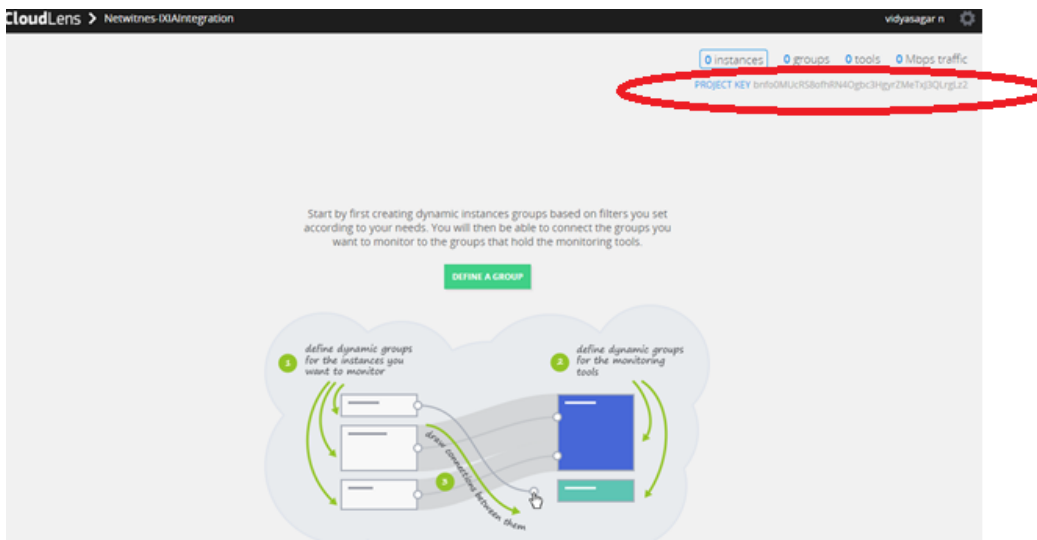
CREATE NEW PROJECT

Project Name
Netwitness-IXIAIntegration

Project Description

OK Cancel

- Click on your newly created project and make note of your Project Key.
You need the key later for the API key configured on the **Host & Tool agents**.



Task 3. Install Docker Container on Decoder

Complete the following steps to install the Docker container onto the Packet Decoder.

- SSH to the Packet Decoder.
- Enter the following commands to complete the install the Docker service on the Decoder.


```
# yum clean all
# yum -y install docker
```
- Enter the following command string to start the Docker service.


```
# service docker start
```

4. Enter the following commands to:

- Access the Ixia repository and obtain the `cloudlens-agent` container.
- Replace the **ProjectKeyFromIxiaProjectPortal** variable, which identifies your project key in Ixia portal, with the Project Key you created in [Task 2. Create CloudLens Project](#).

```
sudo docker run \
--name cloudlens \
-v /:/host \
-v /var/run/docker.sock:/var/run/docker.sock \
-d --restart=always \
--net=host \
--privileged \
ixiacom/cloudlens-agent:latest \
--server agent.ixia.cloud \
--accept_eula y \
--apikey ProjectKeyFromIxiaProjectPortal \
```

Task 4. Install Docker Container on Clients

Complete the follow steps to Y install the Docker Container onto the client machines for which you want to route the traffic to the Packet Decoder.

1. SSH to the AWS Client instance.
2. Enable root access to OS CLI (for example `sudo su -`).
3. Enter the following commands to install Docker.

```
# yum -y install docker
```

Caution: The above example of the installed docker engine is for CentOS7. The instructions may vary slightly for different Linux Distributions. For more information, see the Docker docs at <https://docs.docker/install>.

4. Enter the following commands to start the Docker service.

```
# service docker start
```

5. Enter the following commands to:

- Access the Ixia repository and obtain the **cloudlens-agent** container.
- Replace the variable **ProjectKeyFromIxiaProjectPortal**, which identifies your project key in Ixia portal, with the Project Key you created in the previous section.

```
sudo docker run \
--name cloudlens \
-v /:/host \
-v /var/run/docker.sock:/var/run/docker.sock \
-d --restart=always \
--net=host \
--privileged \
ixiacom/cloudlens-agent:latest \
--server agent.ixia.cloud \
--accept_eula y \
--apikey ProjectKeyFromIxiaProjectPortal \
```

Warning: If you cut and paste commands from a PDF, first paste them into a test editor such as Notepad to confirm the syntax before pasting into the OS CLI. Direct cut and paste between PDF and CLI can contain dashes or other special characters that should not be part of the commands.

Task 5. Map the Packet Decoder to Ixia Clients

Complete the following steps to map the Packet Decoder to the client machines to route the traffic to the Packet Decoder.

1. Go to the Cloudlens public site (<https://www.ixia.cloud>).

2. Double-click on your project to open it.

3. Click the **Define Group** button or the Instances count.

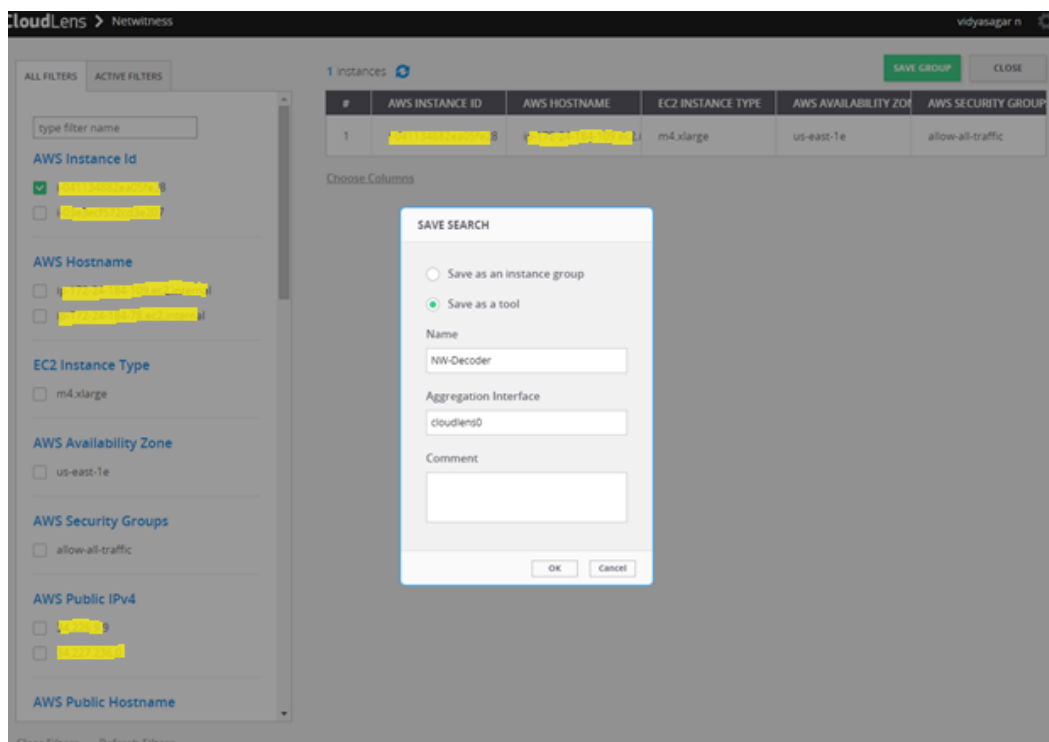
You should see two instances listed, one for your decoder and the other for the client machines.

4. Filter for the decoder instance and click **Save Search**.

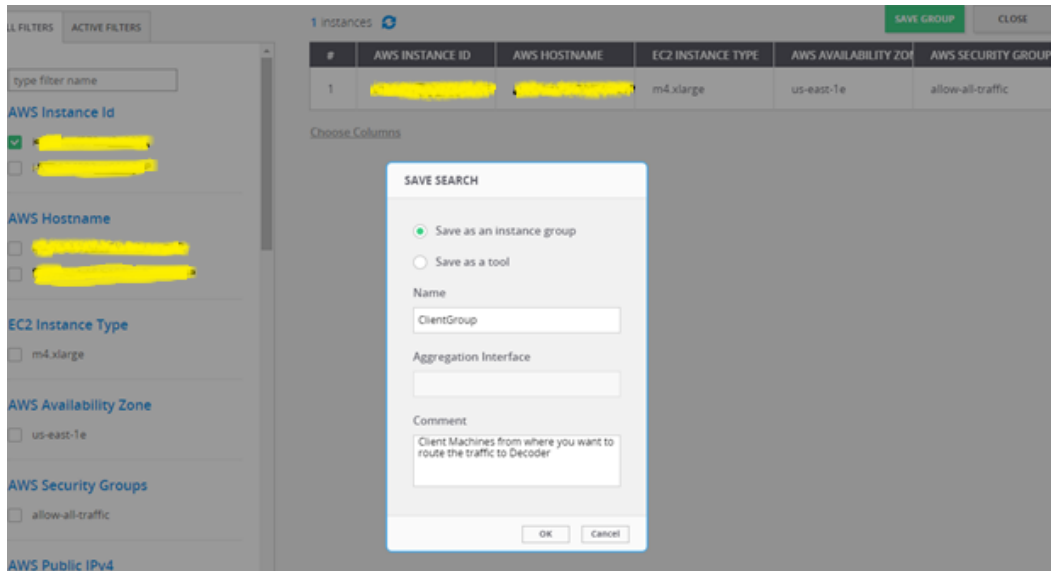
5. Choose **Save as a tool**.

6. Specify a name for the tool, and the **Aggregation Interface**.

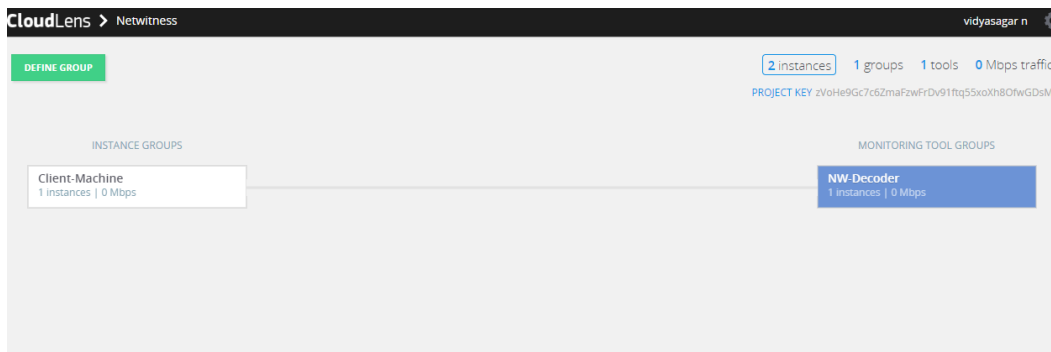
Use a meaningful name for the Aggregation Interface (for example **cloudlens0**). This is a virtual interface that appears in the OS where your Tool is installed. You need to instruct your tool to 'listen' to that interface in a subsequent step.



7. Filter the client host instance from the list, and click **Save Search**.



8. Navigate back to the top-level view of the project.
Your client machine instance and Decoder instance are now displayed.
9. Drag a connection between the your client machine instance and Decoder instance to allow the flow of packets.



Task 6. Validate CloudLens Packets Arriving at Decoder

Complete the following steps to validate that packets are actually arriving at the Packet Decoder.

1. SSH to the Packet Decoder.
2. Enter the following command.

```
ifconfig
```

The new aggregation interface you created is displayed.

```
[root@ip-172-31-101-10 ~]# ifconfig
cloudlens0 Link encap:Ethernet HWaddr 08:00:27:00:00:00
inet6 addr: fe80::208:0:0:0:fe80::0:0:0:0/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:9100 Metric:1
RX packets:6 errors:0 dropped:0 overruns:0 frame:0
TX packets:6 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:468 (468.0 b) TX bytes:468 (468.0 b)
```

3. Generate traffic from the client OS instance CLI (for example, wget http://www.google.com/).

```
root@ip-172-24-164-104-ec2-user]# wget https://172.24.164.104 --no-check-certificate
--2017-06-19 14:33:05-- https://172.24.164.104/
Connecting to 172.24.164.104:443... connected.
WARNING: cannot verify 172.24.164.104's certificate, issued by 欵樁N=Puppet CA: cc4bfb66-8746-4b2f-88ee-3f82862c7069欵?
Unable to locally verify the issuer's authority.
WARNING: certificate common name 欵樁cc4bfb66-8746-4b2f-88ee-3f82862c7069欵? doesn't match requested host name 欵? 172.24.164.104?
HTTP request sent, awaiting response... 302 Found
Location: https://172.24.164.104/login [following]
--2017-06-19 14:33:05-- https://172.24.164.104/login
Reusing existing connection to 172.24.164.104:443.
HTTP request sent, awaiting response... 200 OK
length: unspecified
Saving to: 欵樁index.html.7欵?

index.html.7          [ <=>          ]  2.01K  --.-KB/s  in 0s

2017-06-19 14:33:05 (246 MB/s) - 欵樁index.html.7欵? saved [2062]
```

4. SSH to Packet Decoder to go to your Packet Decoder instance CLI.
5. Enter the following commands to look for suitable results in the tcpdump.



```
tcpdump -i Cloudlens0
```

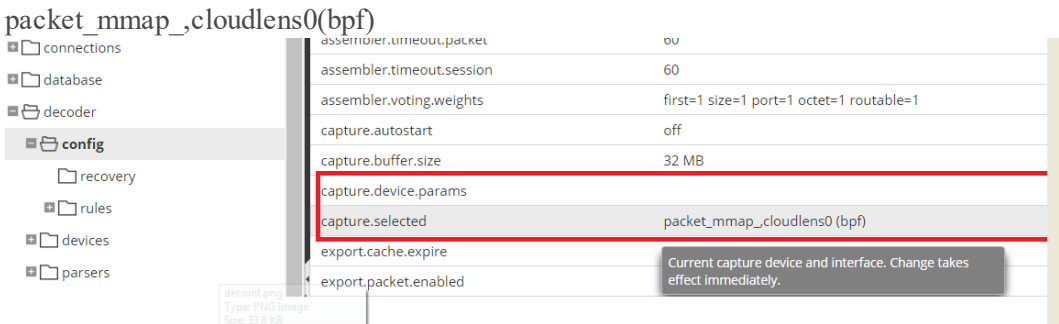
```
4 packets dropped by kernel
root@ip-172-24-164-104-ec2-user]# tcpdump -i cloudlens0
tcpdump: WARNING: cloudlens0: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on cloudlens0, link-type EN10MB (Ethernet), capture size 65535 bytes

14:37:11.408308 IP 175.2.141.156 > ip-172-24-164-104.ec2.internal: ICMP echo request, id 132, seq 32849, length 8
14:37:11.408318 IP ip-172-24-164-104.ec2.internal > 175.2.141.156: ICMP echo reply, id 132, seq 32849, length 8
14:37:11.781923 IP 175.2.141.156 > ip-172-24-164-104.ec2.internal: ICMP 175.2.141.156 protocol 1 unreachable, length 36
```

Task 7. Set the Interface in the Packet Decoder

Complete the following steps in the Packet Decoder to set the interface to use for the Ixia integration.

1. SSH to the Packet Decoder.
2. Enter the following commands to restart decoder service.
\$ sudo restart nwdecoder
The Packet Decoder is now set to capture network traffic.
3. Log in to NetWitness Platform and click **Administration > Services**.
4. In the Admin Services view, select a Decoder service and click   > **View > Explore**.
5. Expand the **decoder** node and click **config** to view the configuration settings.
6. Set the **capture.selected** parameter to the following value.



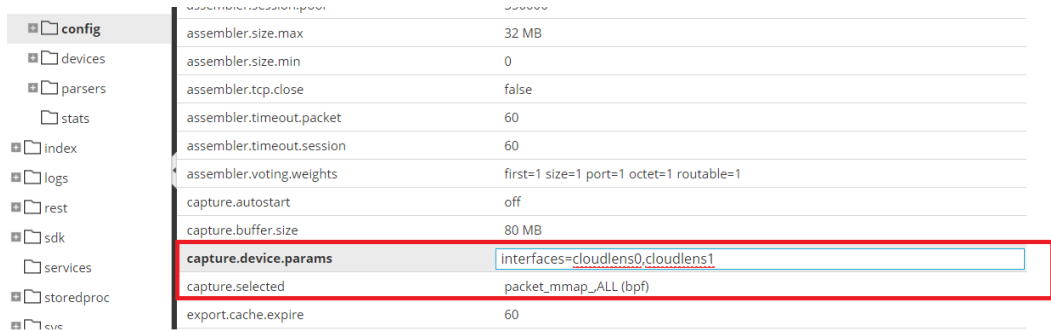
packet_mmap_cloudlens0(bpf)

connections	assembler.timeout.packet	00
database	assembler.timeout.session	60
decoder	assembler.voting.weights	first=1 size=1 port=1 octet=1 routable=1
config	capture.autostart	off
recovery	capture.buffer.size	32 MB
rules	capture.device.params	
devices	capture.selected	packet_mmap_cloudlens0 (bpf)
parsers	export.cache.expire	
	export.packet.enabled	

Current capture device and interface. Change takes effect immediately.

7. (Conditional) - If you have multiple capture interfaces on the Packet Decoder, set the parameters with the following values.

```
capture.device.params --> interfaces=cloudlens0,cloudlens1
capture.selected --> packet_mmap_,All
```



config	assembler.session.timeout	300000
config	assembler.size.max	32 MB
config	assembler.size.min	0
config	assembler.tcp.close	false
config	assembler.timeout.packet	60
config	assembler.timeout.session	60
config	assembler.voting.weights	first=1 size=1 port=1 octet=1 routable=1
config	capture.autostart	off
config	capture.buffer.size	80 MB
config	capture.device.params	interfaces=cloudlens0,cloudlens1
config	capture.selected	packet_mmap_,All (bpf)
config	export.cache.expire	60

8. Restart the Decoder service after you set the **capture.selected** parameter.

Integrate f5® BIG-IP with the Network Decoder

IG-IP Virtual Edition (VE) is an inline virtual server and load balancer. A common use case would be for the f5® box to be a virtual web server that presents a single IP address and host name that manages requests to a pool of web servers in the cloud.

All traffic to RSA NetWitness® Platform flows through the f5® BIG-IP VE virtual server.

The virtual server functions of the BIG-IP clone all traffic to a designated computer by re-writing mac addresses and loading them into a subnet shared with the destination sniffer. This section describes how to set up the Decoder as the sniffer.

f5® BIG-IP VE Deployment Information

f5® BIG-IP VE on AWS is available through the AWS Marketplace and activated by a BYOL license. A thirty-day free trial is also available.

For more information on this solution refer to the f5® BIG-IP DNS Data Sheet (<https://www.f5.com/pdf/products/big-ip-dns-datasheet.pdf>).

Task 1: Set Up a BIG-IP VE Virtual Server Instance

Set up a BIG-IP VE Virtual Server Instance according to the instructions in the "BIG-IP Virtual Edition 12.1.0 and Amazon Web Services: Multi-NIC Manual" (https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-ve-multi-nic-setup-amazon-ec2-12-1-0.html). Complete all the steps through the last steps, "Creating a virtual server."

This virtual server performs packet capture. You may need to create multiple virtual servers to depending on your volume.

As part of creating the virtual server, you must have at least one server in your NetWitness Platform domain to handle the traffic routed by the virtual server (for example, you can create another instance in AWS to host the internal server).

Task 2: Create a Clone Pool

1. Make sure that your Decoder has a network interface on the same subnet as one of the network interfaces on the BIG-IP VE instance.
The clone pool sends packets to the Decoder by rewriting MAC addresses and sending them out a network interface. MAC address rewriting can be used to route packets to another subnet.
2. Set up the clone pool within the BIG-IP VE virtual server according to the instructions in "K13392: Configuring the BIG-IP system to send traffic to an intrusion detection system (11.x - 13.x)" article (<https://support.f5.com/kb/en-us/solutions/public/13000/300/sol13392.html>).
This document explains how to create the clone pool, and how to make an existing virtual server copy traffic to the clone pool. In this case, we will place the Decoder instance in the clone pool.

Guidelines

The following guidelines help you to configure packet capture correctly using BIG-IP VE.

- The Decoder instance must have its own IP address on one of the same subnets as BIG-IP VE. BIG-IP VE uses that IP address to identify the Decoder as being part of the clone pool.
- When adding the Decoder instance to the clone pool, BIG-IP asks for a port number in addition to the IP address. This port number does not matter for the cloned traffic. The Decoder will receive all the cloned traffic, regardless of what port number was used here.
- By default, the AWS subnet shared by the Decoder and BIG-IP VE does not allow the cloned traffic to travel from the BIG-IP VE interface to the Decoder interface. You must disable the **source/dest. check** on both the Decoder and BIG-IP VE network interfaces in AWS.
- The Decoder instance must have a single network interface, eth0, by default. The Decoder captures traffic on this interface, but it may also receive administrative traffic on this interface. RSA recommends using network rules to filter out ssh and nwdecoder traffic from the capture stream. These are ports 22 (ssh) and 50004/56004 (nwdecoder).

Troubleshooting Tips

There are areas to troubleshoot if packets are not being accepted by the Decoder.

- Make sure that the BIG-IP VE is sending the packets out of the correct interface.
The BIG-IP VE instance contains `tcpdump`. Use it to verify the cloned packets are being sent out the expected interface. If they are not, there is a problem in the setup of the clone pool or the virtual server.
- Make sure that the Decoder is receiving packets.
The Decoder has `tcpdump` installed on it. Use it to verify that the Decoder is receiving packets. If the Decoder is not capturing packets, make sure that:
 - The AWS **source/dest. check** is turned off.
 - The Decoder is on the same subnet as the interface the BIG-IP VE is using to clone packets.

AWS Instance Configuration Recommendations

Note: These recommendations can be used as a baseline for 11.2.0.0 and adjusted as needed.

Note: For a description of terms and abbreviations used in this topic, see [Abbreviations and Other Terminology Used in this Guide](#).

This topic contains the minimum AWS instance configuration settings recommended for the RSA NetWitness® Platform virtual stack components.

- EC2 Instance:
 - Minimum instance type - **m4-2xlarge** is the minimum instance type required for any NetWitness Platform component AMI so that it can function.
 - Instance type adjustments you must adjust instance types according to your ingestion rate, content and parsers, dashboard reports, scheduled reports, investigations, and active users.
 - Recommended settings - the recommended settings in the SA component instance tables below were calculated under the following conditions.
 - Ingestion rates of 15,000 EPS and 1.5 Gbps were used.
 - All the components were integrated.
 - The Log stream includes a Log Decoder, Concentrator, and Archiver.
 - The Packet stream includes a Network Decoder and Concentrator.
 - The Endpoint Hybrid stream includes a Endpoint Server, Concentrator and Log Decoder.
 - Respond was receiving alerts from the Reporting Engine and Event Stream Analysis.
 - The background load includes reports, charts, alerts, investigation, and respond.

- EBS Volumes (Storage)

Contact RSA Customer Support (<https://community.rsa.com/docs/DOC-1294>) for assistance on how to increase the number of volumes based on your storage requirements using the RSA Sizing and Scoping Calculator.

Note: The Concentrator index volume must be allocated on Provisioned IOPS SSD.

- Index
- Meta
- Session
- Packet

Archiver

EC2 Instance			
EPS	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
5,000	m4.xlarge No of CPU: 4 Memory: 16 GB	No	Yes
10,000	m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
15,000	m4.4xlarge No of CPU: 16 Memory: 64 GB	No	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
archiver	/dev/sdg	Throughput Optimized HDD	240 MB/s
workbench	/dev/sdh	Throughput Optimized HDD	N/A

Broker

EC2 Instance		
Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
m4.xlarge No of CPU: 4 Memory: 16 GB	No	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
broker	/dev/sdg	General Purpose SSD	N/A

Concentrator - Log Stream

EC2 Instance			
EPS	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
5,000	m4.xlarge No of CPU: 4 Memory: 16 GB	No	Yes
10,000	m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
15,000	m4.4xlarge No of CPU: 16 Memory: 64 GB	No	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/(root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index	/dev/sdg	Provisioned IOPS	10,000
session, metadb	/dev/sdh	Throughput Optimized HDD	240 MB/s

Packet Stream Solutions

Concentrator - Gigamon Solution

EC2 Instance			
Mbps/Gbps	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
500 Mbps	c4.4xlarge No of CPU: 16 Memory: 30 GB	No	Yes
1,000 Mbps	c4.8xlarge No of CPU: 36 Memory: 60 GB	No	Yes
1.5 Gbps	m4.10xlarge No of CPU: 40 Memory: 160 GB	No	Yes

Concentrator - f5 BIG-IP Solution

To be updated when f5 BIG-IP performance testing is complete.

EC2 Instance			
Mbps/Gbps	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
230 Mbps	m4.4xlarge No. of CPU: 16 Memory: 64 GB	No	No

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index	/dev/sdg	Provisioned IOPS	15,000
session, metadb	/dev/sdh	Throughput Optimized HDD	240 MB/s

Decoder - Gigamon Solution

EC2 Instance			
Mbps/Gbps	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
500 Mbps	c4.2xlarge No of CPU: 8 Memory: 15 GB	Yes	Yes
1000 Mbps	c4.4xlarge No of CPU: 16 Memory: 30 GB	Yes	Yes
1.5 Gbps	c4.8xlarge No of CPU: 36 Memory: 60 GB	Yes	Yes

Decoder - f5 BIG-IP Solution

To be updated when f5 BIG-IP performance testing is complete.

EC2 Instance			
Mbps/Gbps	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
230 Mbps	m4.4xlarge No. of CPU: 16 Memory: 64 GB	No	No

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index,session,meta	/dev/sdg	Throughput Optimized HDD	240 MB/s
packet	/dev/sdh	Throughput Optimized HDD	240 MB/s

ESA and Context Hub on Mongo Database

	EC2 Instance		
EPS	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
9,000	m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
18,000	r4.2xlarge No of CPU: 8 Memory: 61 GB	No	Yes
30,000 Aggregation Rate	r4.4xlarge No of CPU: 16 Memory: 122 GB	No	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
apps (/opt/rsa)	/dev/sdg	General Purpose SSD	N/A

Log Collector (Syslog, Netflow, and File Collection Protocols)

EC2 Instance			
EPS	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
30,000 NON SSL	c4.2xlarge No of CPU: 8 Memory: 15 GB	No	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
logcollector	/dev/sdg	General Purpose SSD	N/A

Log Decoder

EC2 Instance			
EPS	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
5,000	c4.2xlarge No of CPU: 8 Memory: 15 GB	Yes	Yes
10,000	c4.4xlarge No of CPU: 16 Memory :30 GB	Yes	Yes
15,000	c4.8xlarge No of CPU: 36 Memory: 60GB	Yes	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index,session,meta	/dev/sdg	Throughput Optimized HDD	240 MB/s
packet	/dev/sdh	Throughput Optimized HDD	240 MB/s

NetWitness Server, Reporting Engine, Respond and Health & Wellness

EC2 Instance		
Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
m4.2xlarge No of CPU: 8 Memory: 32 GB	No	Yes
m4.4xlarge No of CPU: 16 Memory: 64 GB	No	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/ (root)	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
uax,ipdb	/dev/sdg	General Purpose SSD	N/A
redb,rehome	/dev/sdh	General Purpose SSD	N/A

NetWitness Endpoint Hybrid

	EC2 Instance		
Agents	Instance Type	Enhanced Networking Enabled	Tenancy Type - Dedicated - Run a Dedicated Instance
15,000 agents	m4.10xlarge No of CPU: 40 Memory: 160 GB RAM	Yes	Yes

EBS Volumes (Storage)			
Volumes	Device	Volume Type	IOPS/Baseline Throughput
/root	/dev/sda1	General Purpose SSD	N/A
usr,var,opt,home,tmp	/dev/sdf	General Purpose SSD	N/A
index,session,meta (Log Decoder)	/dev/sdg	Throughput Optimized HDD	240 MB/s
packet (Log Decoder)	/dev/sdh	Throughput Optimized HDD	240 MB/s
index (Concentrator)	/dev/sdi	Provisioned IOPS	10,000
session,meta (Concentrator)	/dev/sdj	Throughput Optimized HDD	240 MB/s
mongoDB	/dev/sdl	Throughput Optimized HDD	240 MB/s

UEBA

	UEBA Instance		
CPU	Memory	Read IOPS	Write IOPS
16 or 24GHz	64GB	500MB	500MB

