

Elastic Stack Deployment

The files in this repository were used to configure the network depicted below.



These files have been tested and used to generate a live ELK deployment on Azure. They can be used to either recreate the entire deployment pictured above. Alternatively, select portions of the yaml file may be used to install only certain pieces of it, such as Filebeat.

- [Ansible Config](#)
- [Ansible Hosts](#)
- [Install DVWA](#)
- [Install Elk](#)
- [Install FileBeat](#)
- [FileBeat Config](#)
- [Install MetricBeat](#)
- [MetricBeat Config](#)

Download the *ansible.cfg* from <https://ansible.com> and edit to your config or copy [Ansible Configuration](#) to `/etc/ansible`

- For *ansible.cfg* edit:

```
cd /etc/ansible
nano ansible.cfg
change 'remote_user = admin'
```

Assign username and SSH Public Key for Web1, Web2, Web3, and Elk in Azure Portal.

- Web1, Web2, Web3, and Elk | Reset Password | Reset SSH Public Key

```
username: admin
```

```
SSH Key: copy id_rsa.pub from ansible docker container in ~/.ssh/id_rsa.pub
```

Generate SSH Key

```
cd ~/.ssh  
ssh-keygen (Press Enter for all)  
cat id_rsa.pub
```

This document contains the following details:

- Description of the Topology
- Access Policies
- ELK Configuration
 - Beats in Use
 - Machines Being Monitored
- How to Use the Ansible Build

Description of the Topology

The main purpose of this network is to expose a load-balanced and monitored instance of DVWA, the D*mn Vulnerable Web Application.

Load balancing ensures that the application will be highly available, in addition to restricting access to the network.

- Load balancers protect the availability of applications. Should one server go down or receive too many requests the load balancer redirects traffic to other servers with the exact same data.
- Using a jumpbox has many advantages to it. Some of the main advantages are security, automation, and access control. A jumpbox allows the main machines to not be directly accessible, requiring first to connect to the jumpbox before making a connection to the machines. This makes it harder for attackers because there is only one heavily guarded entry point. This also allows you to control who can and can not connect to the machines.

Another advantage is that it allows you to automate the other machines. Instead of updating and installing applications on each machine individually you can use software such as Ansible to install the same thing on multiple machines very easily.

Integrating an ELK server allows users to easily monitor the vulnerable VMs for changes to the data and system logs.

- Filebeat monitors the log files and locations specified, collects log events, and then forwards them to Elasticsearch.
- Metricbeat takes metrics and statistics that it has collected and sends them to Elasticsearch

The configuration details of each machine may be found below.

Name	Function	IP Address (Private \ Public)	Operating System
Jump Box	Gateway	10.0.0.5 \ 104.211.18.77	Linux
Web 1	DWVA	10.0.0.6	Linux
Web 2	DVWA	10.0.0.7	Linux
Web 3	DVWA	10.0.0.10	Linux
Elk	Elk Stack	10.1.0.5 \ 52.159.72.249	Linux
LoadBalancer	Balance Work Load	\ 20.120.87.97	Linux

Access Policies

The machines on the internal network are not exposed to the public Internet.

Only the jumpbox machine can accept connections from the Internet. Access to this machine is only allowed from the following IP addresses:

- MyPublicIP:5601

Machines within the network can only be accessed by Ansible.

- Elk is accessible by

JumpBox | 10.0.0.5:22

My PC | MyPublicIP:5601

A summary of the access policies in place can be found in the table below.

Name	Publicly Accessible	Allowed IP Addresses
Jump Box	Yes	MyPublicIP
Web 1	no	10.0.0.5
Web 2	no	10.0.0.5
Web 3	no	10.0.0.5
Elk	YES	MyPublicIP:5601
LoadBalancer	YES	MyPublicIP

Elk Configuration

Ansible was used to automate configuration of the ELK machine. No configuration was performed manually, which is advantageous because *Ansible allows you to configure multiple machines at a time with the exact same configuration. Ansible allows for quicker installation and updates. Ansible also prevents misconfiguration of one machine by mistyping.*

The playbook implements the following tasks:

- Increase System Memory

```
- name: Set vm.max_map_count to 262144
  systemctl:
    name: vm.max_map_count
    value: '262144'
    state: present
    reload: yes
```

Install Services

```
`docker.io`  
`python3-pip`  
`docker` through pip
```

Launch container with exposed ports

```
`5601:5601`  
`9200:9200`  
`5044:5044`
```

The following screenshot displays the result of running `docker ps` after successfully configuring the ELK instance.

[TODO: Update the path with the name of your screenshot of docker ps output](#)

Target Machines & Beats

This ELK server is configured to monitor the following machines:

- **Web1:** 10.0.0.6
- **Web2:** 10.0.0.7
- **Web3:** 10.0.0.10

We have installed the following Beats on these machines:

- FileBeat
- MetricBeat

These Beats allow us to collect the following information from each machine:

- *FileBeat* is used to log events
- *MetricBeat* is used to get metrics and system statistics

Using the Playbook

In order to use the playbook, you will need to have an Ansible control node already configured. Assuming you have such a control node provisioned:

SSH into the control node and follow the steps below:

- Copy the [Install Elk](#) file to `~/.etc/ansible`.
 - Update the `install-elk.yml` file to change which machines you want to run it on.
 - Run the playbook, and navigate to *Elk Public IP* to check that the installation worked as expected.
 - Repeat the process for each playbook you want to install.
-

Revision #1

Created 2025-11-25 18:21:24 UTC by David Rizzo

Updated 2025-11-25 18:21:43 UTC by David Rizzo