Security Onion Documentation

3.1.0

Security Onion Solutions, LLC

Table of contents

1. About

1.1 Security Onion

Security Onion is a free and open platform built by defenders for defenders. It includes network visibility, host visibility, intrusion detection honeypots, log management, and case management. Security Onion has been downloaded over 2 million times and is being used by security teams around the world to monitor and defend their enterprises. Our easy-to-use Setup wizard allows you to build a distributed grid for your enterprise in minutes!

1.2 Security Onion Solutions, LLC

Doug Burks started Security Onion as a free and open project in 2008 and then founded Security Onion Solutions, LLC in 2014.

Important

Security Onion Solutions, LLC is the only official provider of hardware appliances, training, and professional services for Security Onion.

For more information about these products and services, please see our company site at https://securityonionsolutions.com.

1.3 Documentation

Warning

Documentation is always a work in progress and some documentation may be missing or incorrect. Please let us know if you notice any issues.

1.3.1 License

This documentation is licensed under CC BY 4.0. You can read more about this license at https://creativecommons.org/licenses/by/4.0/.

1.3.2 Formats

This documentation is published online at https://securityonion.net/docs. If you are viewing an offline version of this documentation but have Internet access, you might want to switch to the online version at https://securityonion.net/docs to see the latest version.

This documentation is also available in PDF format at https://securityonion.net/docs/securityonion-docs.pdf.

Many folks have asked for a printed version of our documentation. Whether you work on airgapped networks or simply want a portable reference that doesn't require an Internet connection or batteries, this is what you've been asking for. Thanks to Richard Bejtlich for writing the inspiring foreword! Proceeds go to the Rural Technology Fund! You can purchase your copy at https://securityonion.net/book.

1.3.3 Authors

Security Onion Solutions is the primary author and maintainer of this documentation. Some content has been contributed by members of our community. Thanks to all the folks who have contributed to this documentation over the years!

1.3.4 Contributing

We welcome your contributions to our documentation! We will review any suggestions and apply them if appropriate.

If you are accessing the online version of the documentation and notice that a particular page has incorrect information, you can submit corrections by clicking the `Edit on GitHub` button in the upper-right corner of each page. Once you have made your corrections, you will need to submit your pull request (PR) to the `dev` branch.

To submit a new page, you can submit a pull request (PR) to the `3/dev` branch of the `docs` repo at https://github.com/Security-Onion-Solutions/docs.

Pages are written in Markdown format and you can find several Markdown guides on the Internet including https://www.markdownguide.org/basic-syntax/.

1.3.5 Naming Convention

New documentation pages should use the following naming convention:

- all lowercase
- `.md` file extension
- ideally, the name of the page should be one simple word (for example: `suricata.md`)
- if necessary, the name of the page can be hyphenated (for example: network-visibility.md)

2. Introduction

Security Onion is a free and open platform built by defenders for defenders. It includes network visibility, host visibility, intrusion detection honeypots, log management, and case management.

For network visibility, we offer signature based detection via Suricata, rich protocol metadata and file extraction using either Zeek or Suricata, full packet capture using Suricata, and file analysis. For host visibility, we offer the Elastic Agent which provides data collection, live queries via Osquery, and centralized management using Elastic Fleet. Intrusion detection honeypots based on OpenCanary can be added to your deployment for even more enterprise visibility. All of these logs flow into Elasticsearch and we've built our own user interfaces for Alerts, Dashboards, threat hunting, case management, and grid management.

> **Note**
>
> Check out our Introduction to Security Onion video at https://securityonion.com/demo!

In the diagram below, we see Security Onion in a traditional enterprise network with a firewall, workstations, and servers. You can use Security Onion to monitor north/south traffic to detect an adversary entering an environment, establishing command-and-control (C2), or perhaps data exfiltration. You'll probably also want to monitor east/west traffic to detect lateral movement. As more and more of our network traffic becomes encrypted, it's important to fill in those blind spots with additional visibility in the form of endpoint telemetry. Security Onion can consume logs from your servers and workstations so that you can then hunt across all of your network and host logs at the same time.

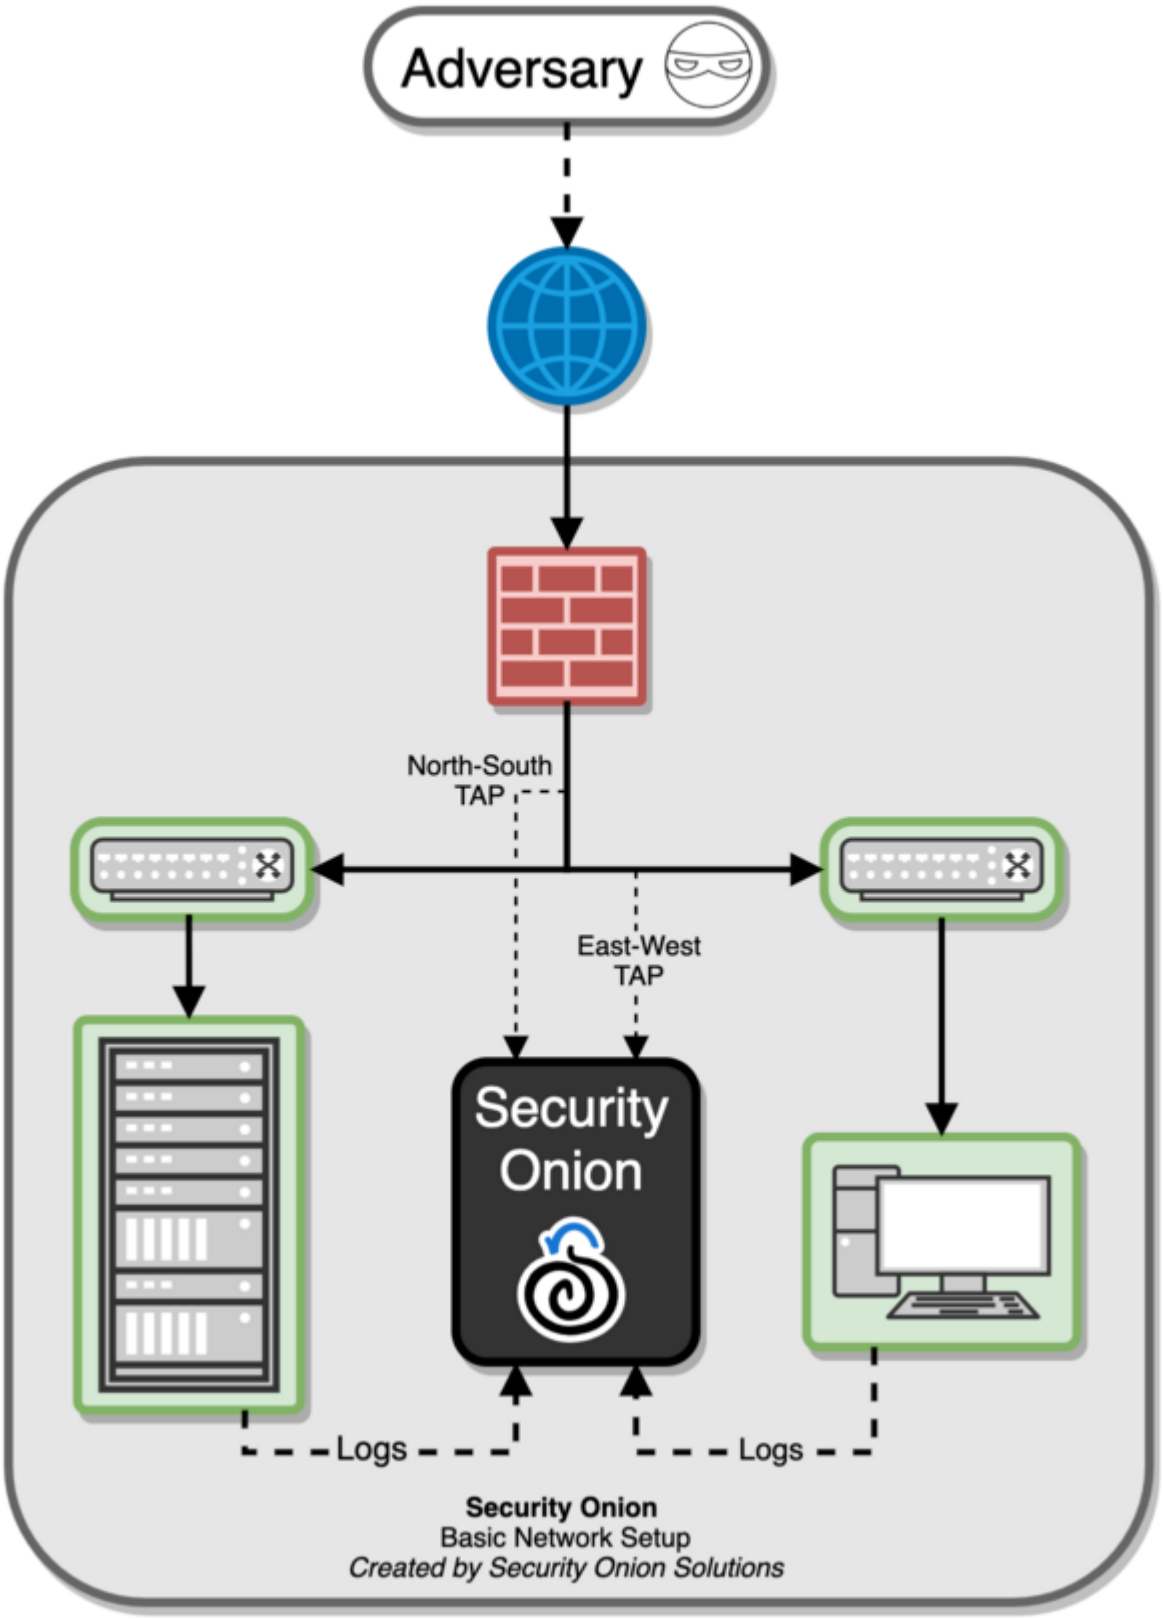
Adversary
North-South
TAP
East-West
TAP
Security
Onion
Logs
Logs
Security Onion
Basic Network Setup
Created by Security Onion Solutions

2.1 Network Visibility

From a network visibility standpoint, Security Onion seamlessly weaves together intrusion detection, network metadata, full packet capture, file analysis, and intrusion detection honeypots.

2.1.1 Intrusion Detection

Security Onion generates NIDS (Network Intrusion Detection System) alerts by monitoring your network traffic and looking for specific fingerprints and identifiers that match known malicious, anomalous, or otherwise suspicious traffic. This is signature-based detection so you might say that it's similar to antivirus signatures for the network, but it's a bit deeper and more flexible than that. NIDS alerts are generated by Suricata.

2.1.2 Network Metadata

Unlike signature-based intrusion detection that looks for specific needles in the haystack of data, network metadata provides you with logs of connections and standard protocols like DNS, HTTP, FTP, SMTP, SSH, and SSL. This provides a real depth and visibility into the context of data and events on your network. Security Onion provides network metadata using your choice of either Zeek or Suricata.

2.1.3 Full Packet Capture

Full packet capture is like a video camera for your network, but better because not only can it tell us who came and went, but also exactly where they went and what they brought or took with them (exploit payloads, phishing emails, file exfiltration). It's a crime scene recorder that can tell us a lot about the victim and the white chalk outline of a compromised host on the ground. There is certainly valuable evidence to be found on the victim's body, but evidence at the host can be destroyed or manipulated; the camera doesn't lie, is hard to deceive, and can capture a bullet in transit. Full packet capture is written to disk using Suricata.

2.1.4 File Analysis

As Zeek and Suricata are monitoring your network traffic, they can extract files transferred across the network. Strelka can then analyze those files and provide additional metadata.

2.1.5 Intrusion Detection Honeypot (IDH)

We also have an IDH node that allows you to build a node that mimics services. Connections to these services automatically generate alerts.

2.2 Host Visibility

In addition to network visibility, Security Onion provides endpoint visibility via the Elastic Agent which provides data collection, live queries via Osquery, and centralized management using Elastic Fleet.

For devices like firewalls and routers that don't support the installation of agents, Security Onion can consume standard syslog.

2.3 Analysis Tools

With all of the data sources mentioned above, there is an incredible amount of data available at your fingertips. Fortunately, Security Onion tightly integrates the following tools to help make sense of this data.

2.3.1 Security Onion Console

Security Onion Console is the first thing you see when you log into Security Onion. It includes our Alerts interface which allows you to see all of your NIDS alerts from Suricata.

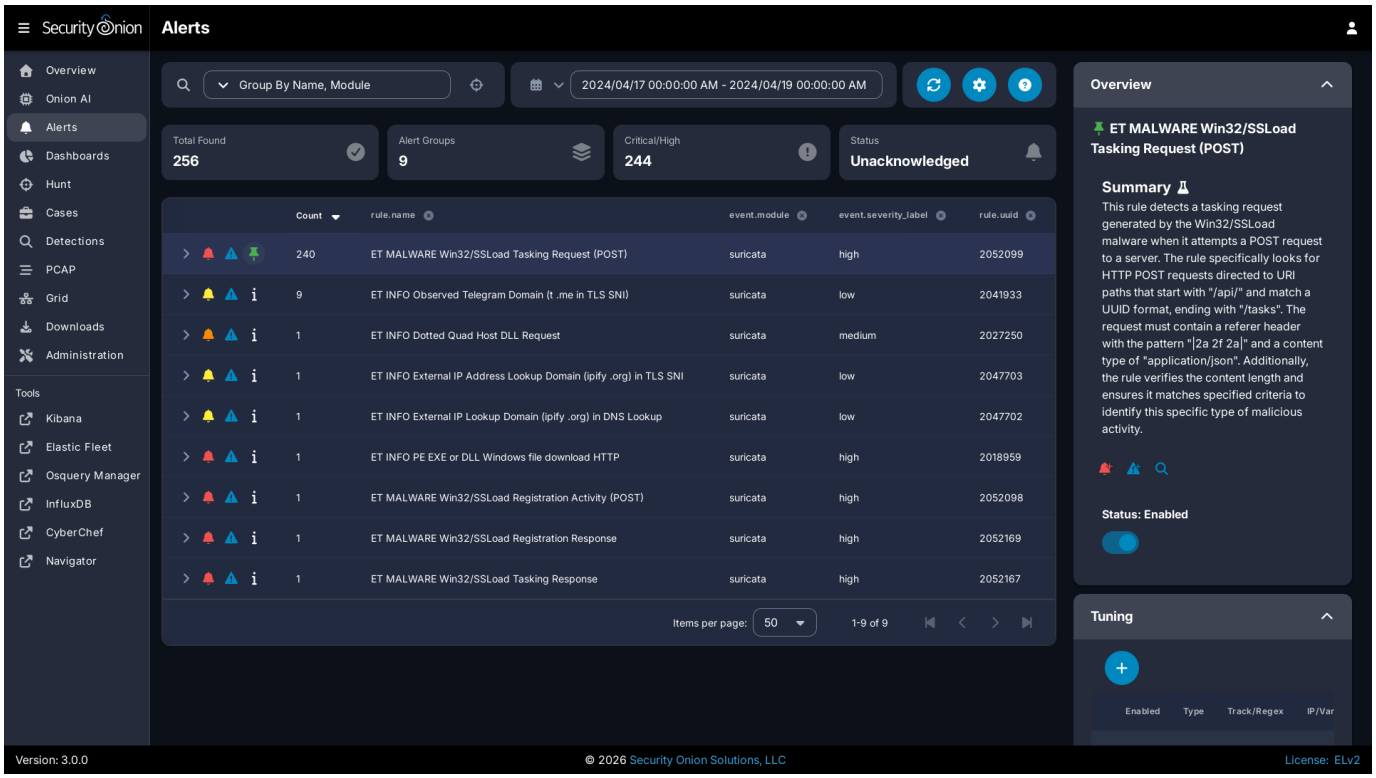

Security Onion Console also includes our Dashboards interface which gives you a nice overview of not only your NIDS alerts but also network metadata logs from Zeek or Suricata and any other logs that you may be collecting.

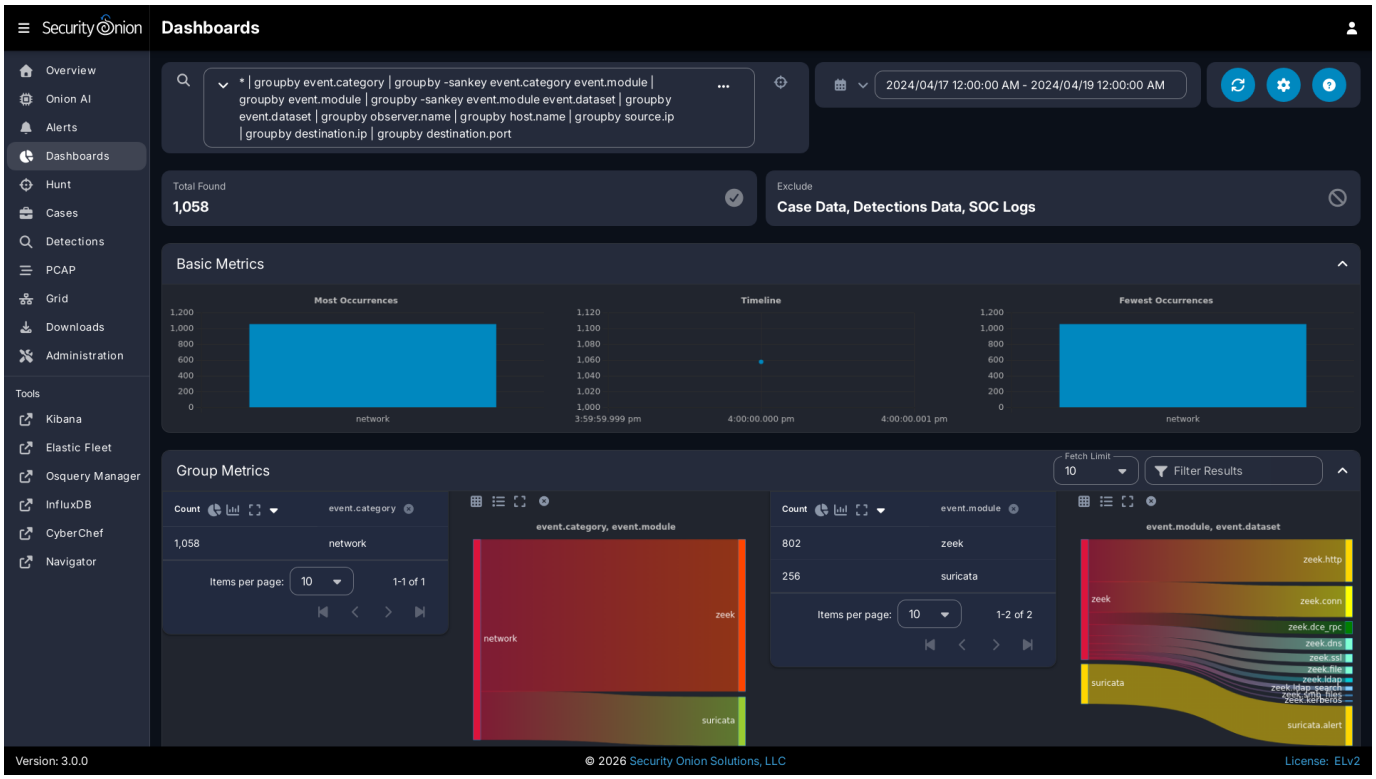

Hunt is similar to Dashboards but its default queries are more focused on threat hunting.

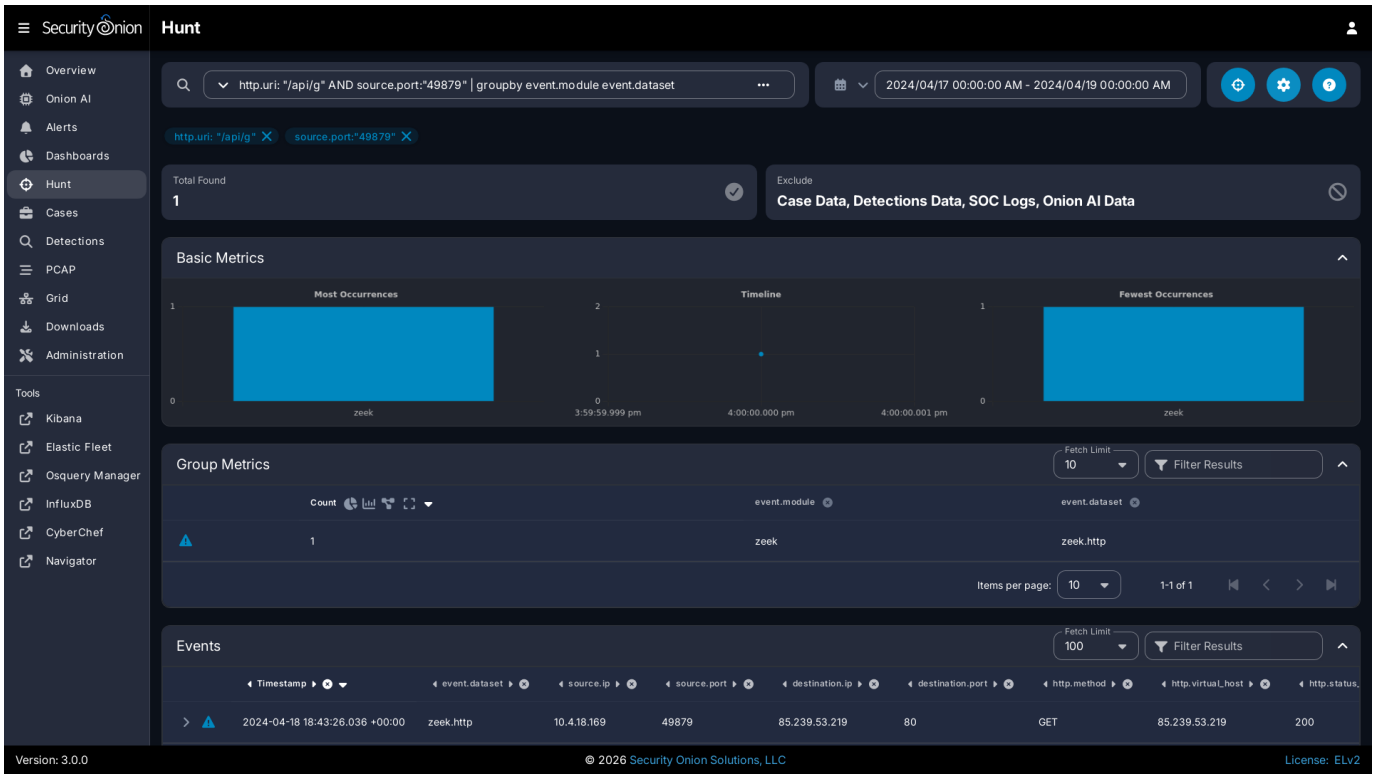

Cases is the case management interface. As you are working in Alerts, Dashboards, or Hunt, you may find alerts or logs that are interesting enough to send to Cases and create a case. Other analysts can collaborate with you as you work to close that case.

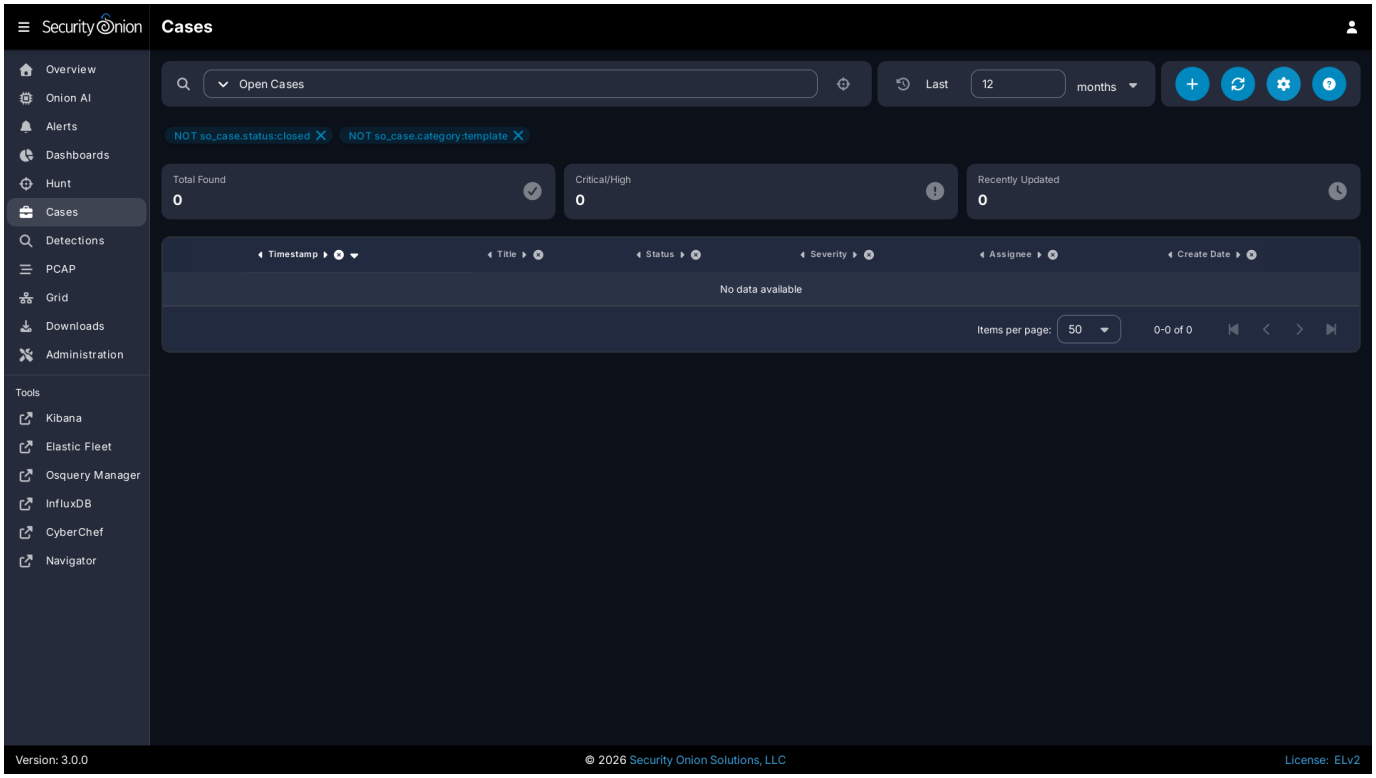

Security Onion Console also includes an interface for full packet capture (PCAP) retrieval.

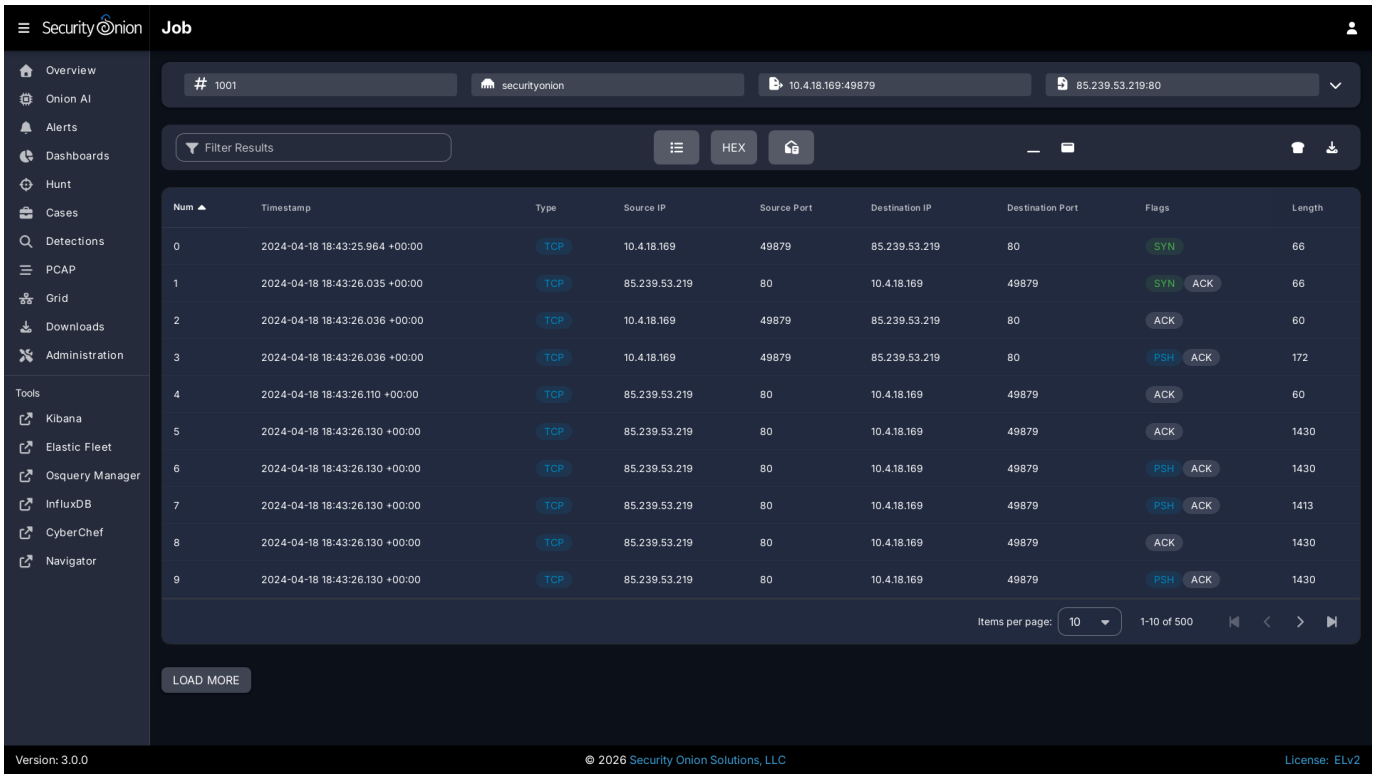

Security Onion Console also includes Detections which makes it quick and easy to tune your NIDS, Sigma, and YARA rules.

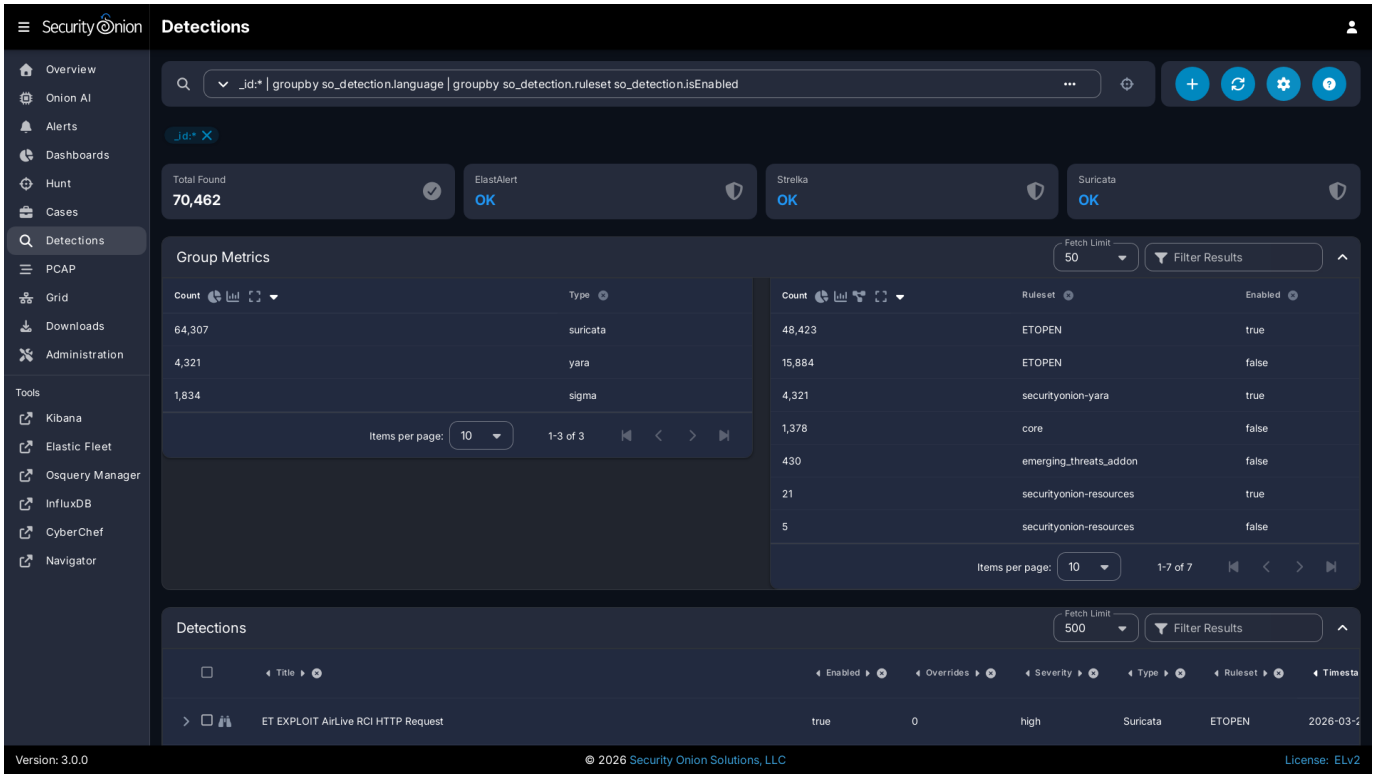

2.3.2 CyberChef

CyberChef allows you to decode, decompress, and analyze artifacts. Alerts, Dashboards, Hunt, and PCAP all allow you to quickly and easily send data to CyberChef for further analysis.

2.4 Workflow

All of these analysis tools work together to provide efficient and comprehensive analysis capabilities. For example, here's one potential workflow:

- Go to the Alerts page and review any unacknowledged alerts.
- Use Detections to tune your rules to increase the signal-to-noise ratio.
- Review Dashboards for anything that looks suspicious.
- Once you've found something that you want to investigate, you might want to pivot to Hunt to expand your search and look for additional logs relating to the source and destination IP addresses.
- If any of those alerts or logs look interesting, you might want to pivot to PCAP to review the full packet capture for the entire stream.
- Depending on what you see in the stream, you might want to send it to CyberChef for further analysis and decoding.
- Escalate alerts and logs to Cases and document any observables. Pivot to Hunt to cast a wider net for those observables.
- If you have the Elastic Agent deployed, then you might want to search for additional host logs or run live queries against your endpoints using Osquery.
- Finally, return to Cases.
- Finally, return to Cases and document the entire investigation and close the case.

2.5 Deployment Scenarios

Analysts around the world are using Security Onion today for many different architectures. The Security Onion Setup wizard allows you to easily configure the best deployment scenario to suit your needs.

2.6 Conclusion

After you install Security Onion, you will have comprehensive network and host visibility for your enterprise. Our analyst tools will enable you to use all of that data to detect intruders more quickly and paint a more complete picture of what they're doing in your environment. Get ready to peel back the layers of your enterprise and make your adversaries cry!

3. License

Security Onion is a free and open platform. Most software included in Security Onion is licensed under open source licenses.

Elastic components and Security Onion components are licensed under the Elastic License 2.0 (ELv2). During installation, you will be prompted to accept the Elastic License:

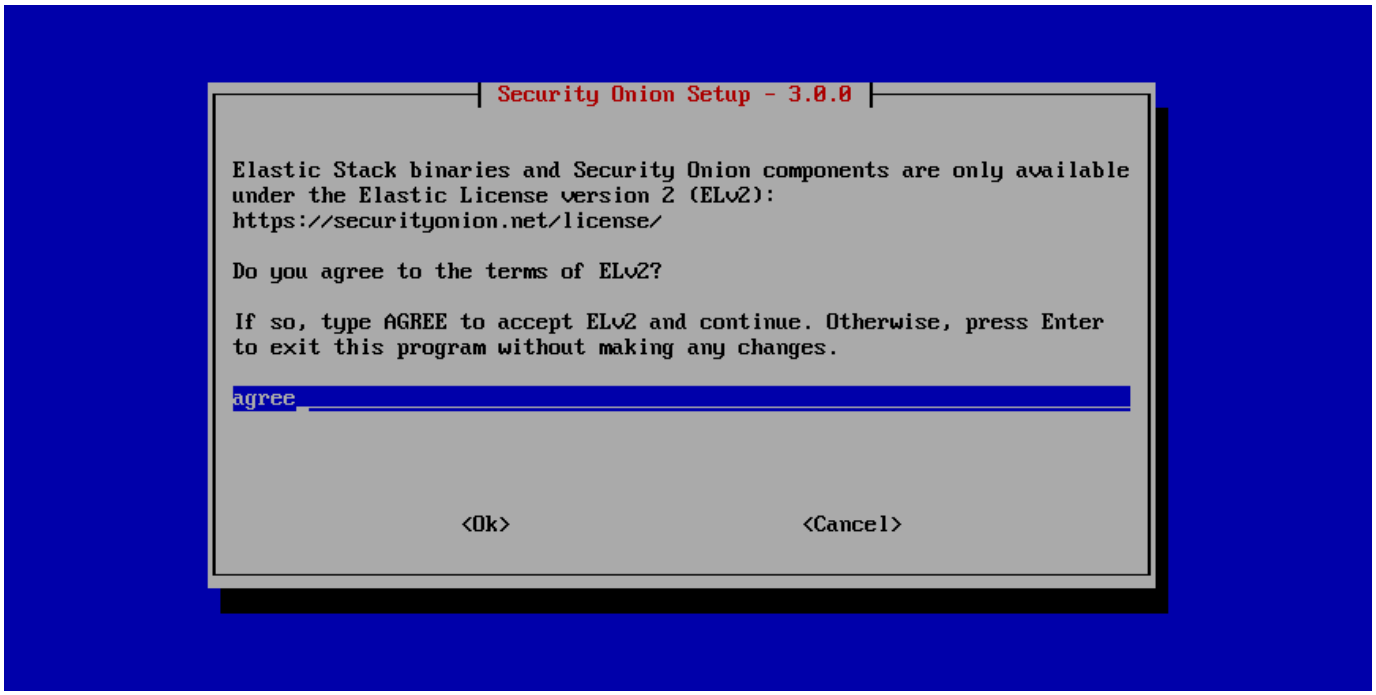

Note

You can find the full text of the Elastic License 2.0 (ELv2) at: https://securityonion.net/license

You can find the Security Onion ELv2 announcement at: https://blog.securityonion.net/2022/08/security-onion-enterprise-features-and.html

4. First Time Users

Welcome, first time users! You're going to be peeling back the layers of your network in just a few minutes!

First, please note that Security Onion only supports x86-64 architecture (standard Intel or AMD 64-bit processors). If you don't have an x86-64 box available, then one option may be to run Security Onion in the cloud. For more information, please see the Amazon Cloud, Azure Cloud, and Google Cloud sections.

Otherwise, if you have an x86-64 box for your Security Onion IMPORT installation, then check to make sure it meets the MINIMUM hardware requirements of 4GB RAM, 2 CPU cores, and 200GB of storage. If you will be installing Security Onion in a virtual machine, then the VM will need those specs at minimum and the host machine will have higher hardware requirements since it will be running the host operating system and possibly other VMs or apps. For more information about virtualization, please see the VMware, VirtualBox, and Proxmox sections. Once you've verified that you have an appropriate installation target, you can proceed to download our ISO image as shown in the Download section and then install the ISO image as shown in the Installation section.

Once you have Security Onion installed either in the cloud or on-prem, you can configure for IMPORT as shown below (also see the Configuration section).

Once you're comfortable with your IMPORT installation, then you can move on to more advanced installations as shown in the Architecture section.

After booting the ISO image, the boot menu appears:

When prompted, enter your desired username and password:

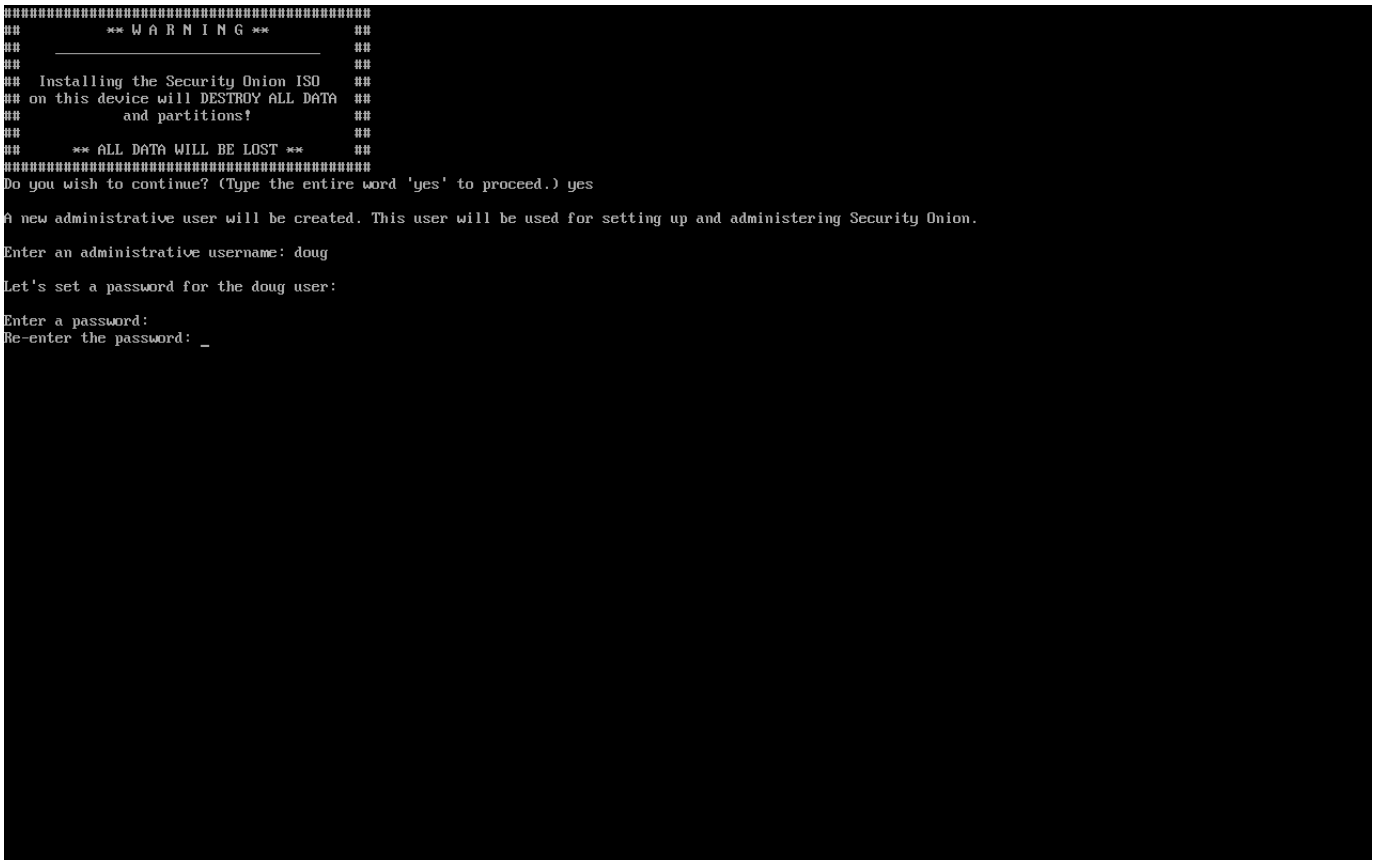

Once installation is complete, you are prompted to reboot:

After rebooting, login using the username and password that you specified and then Setup will start automatically:

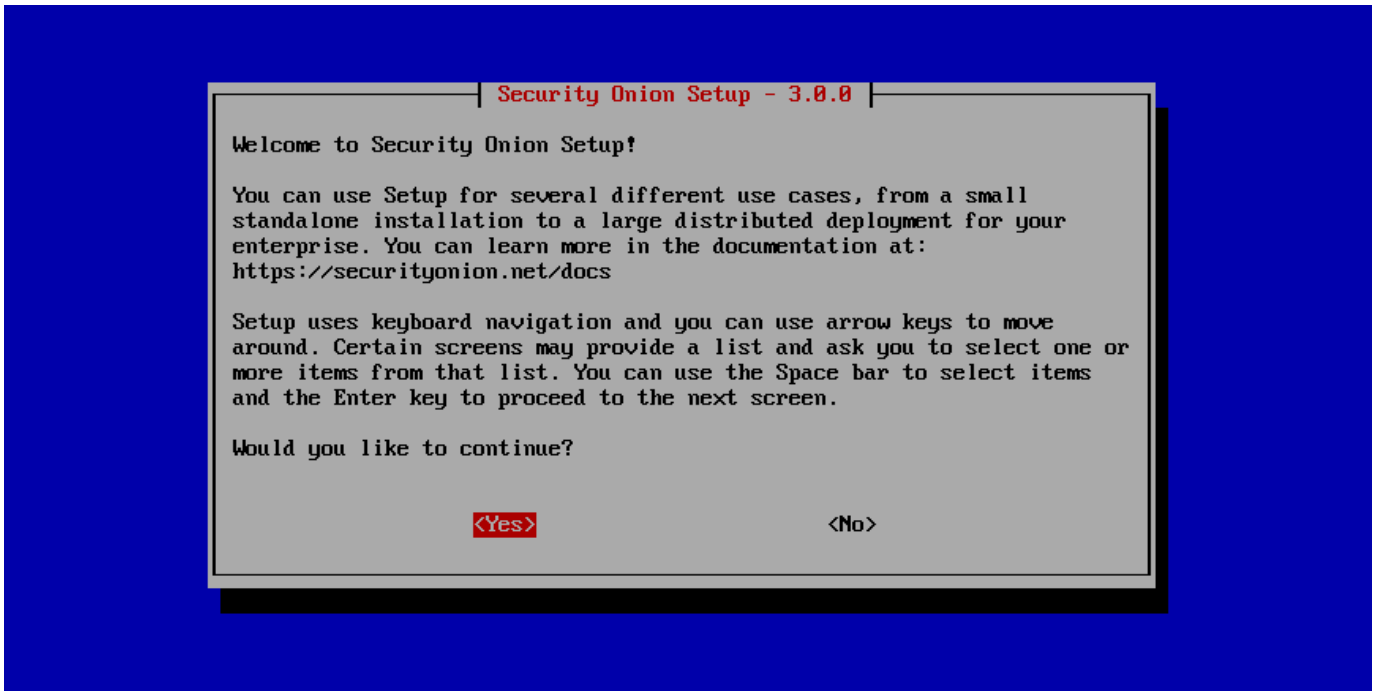

Perform a standard installation:

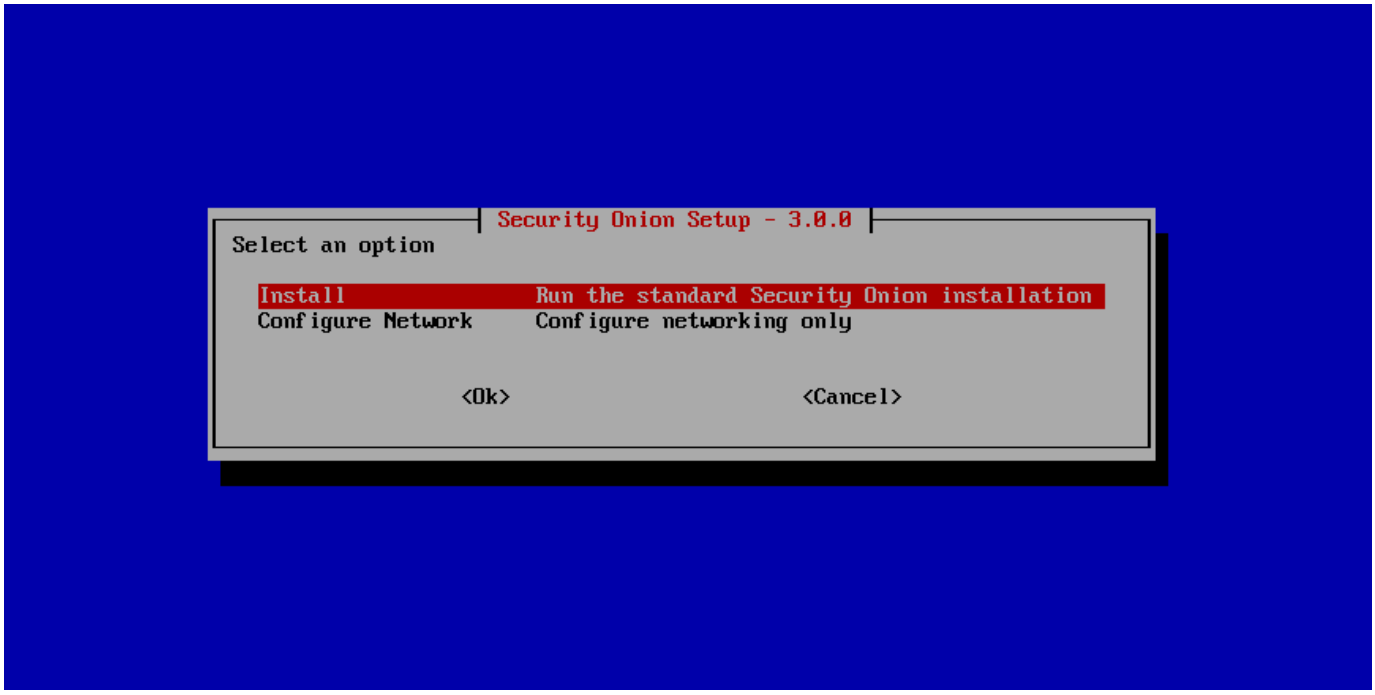

When prompted for installation type, select IMPORT:

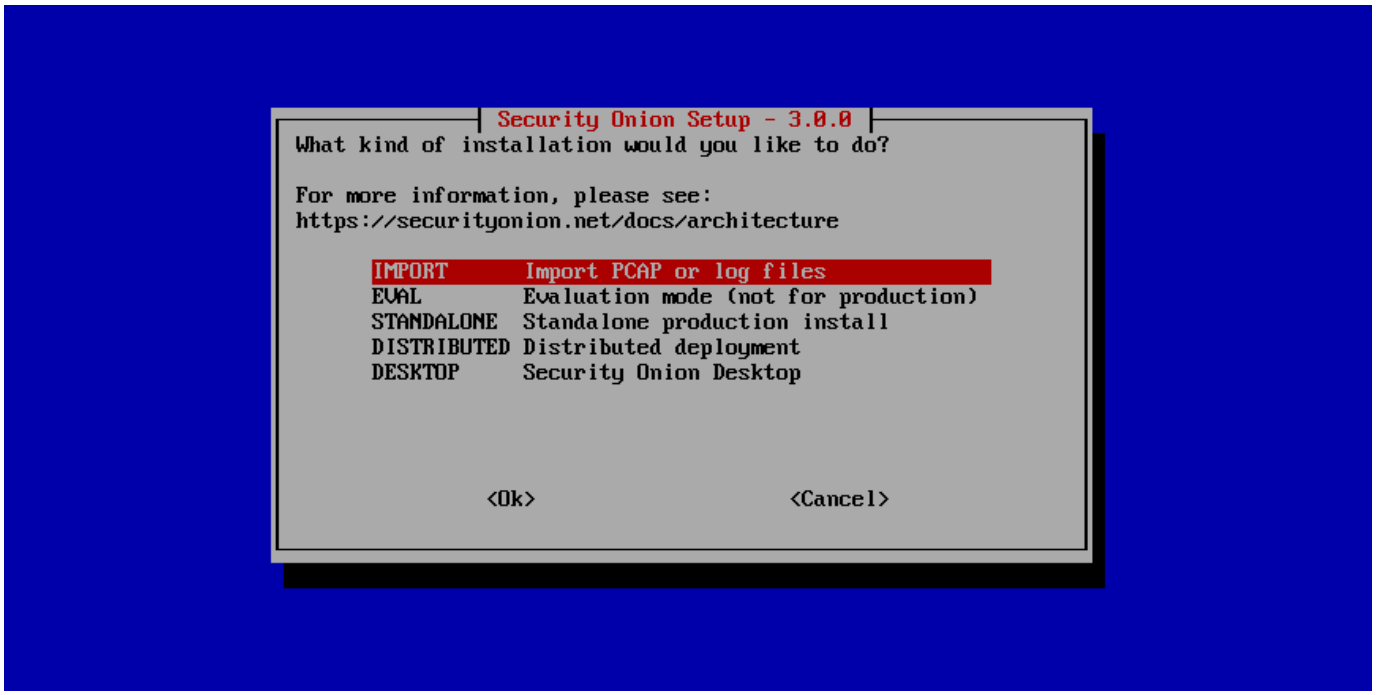

If your Security Onion machine has full Internet access as described in the Firewall section, select Standard. Otherwise, select Airgap:

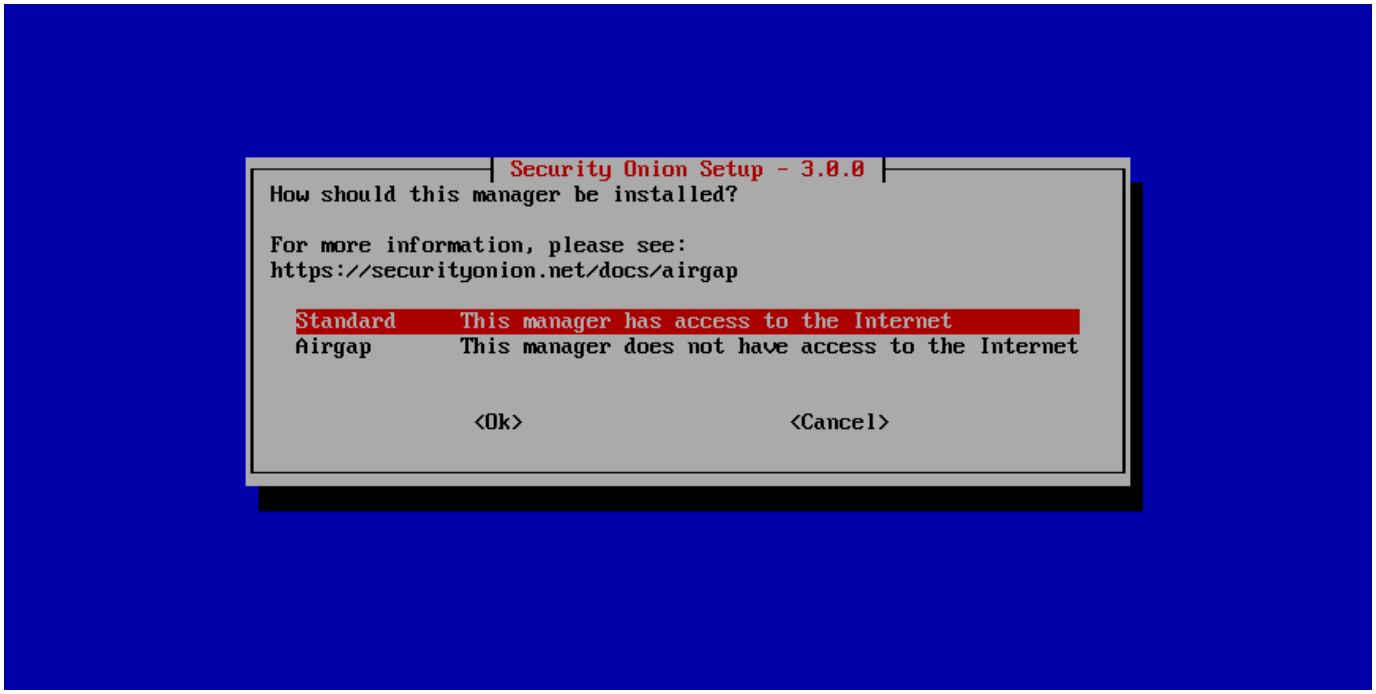

Review the license and agree:

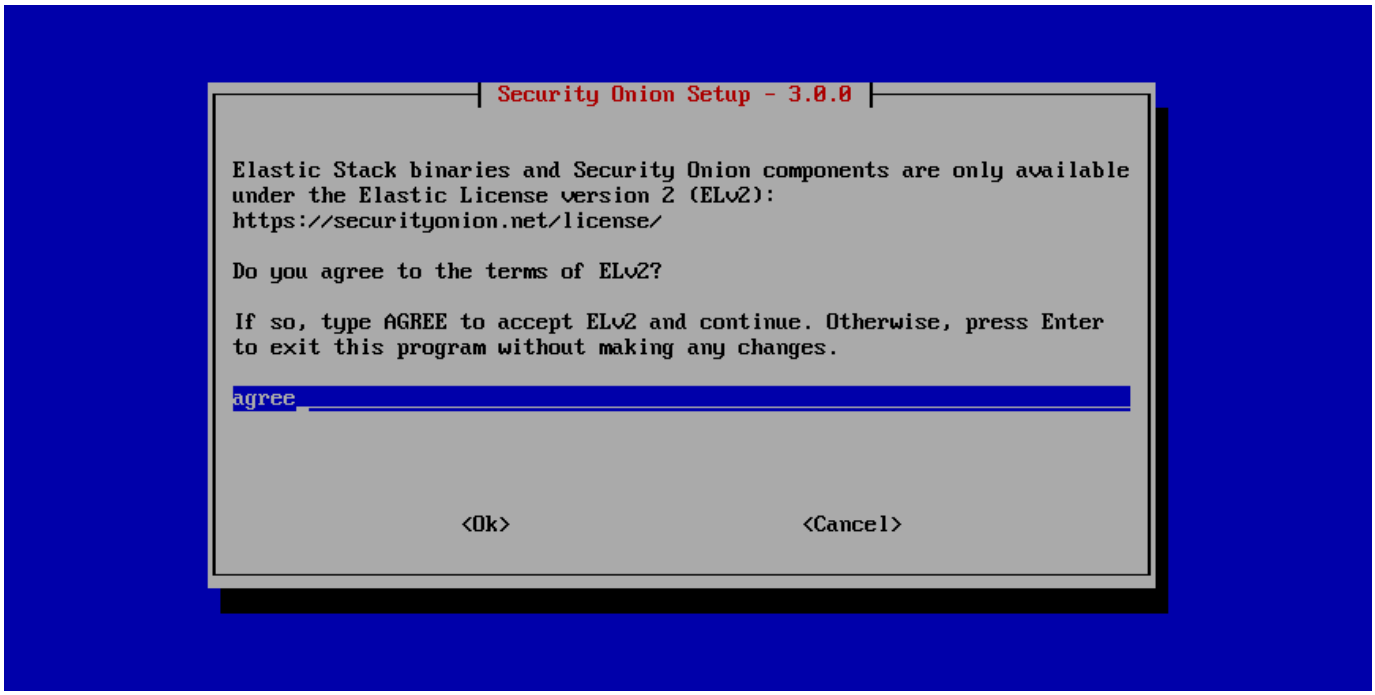

Set the hostname:

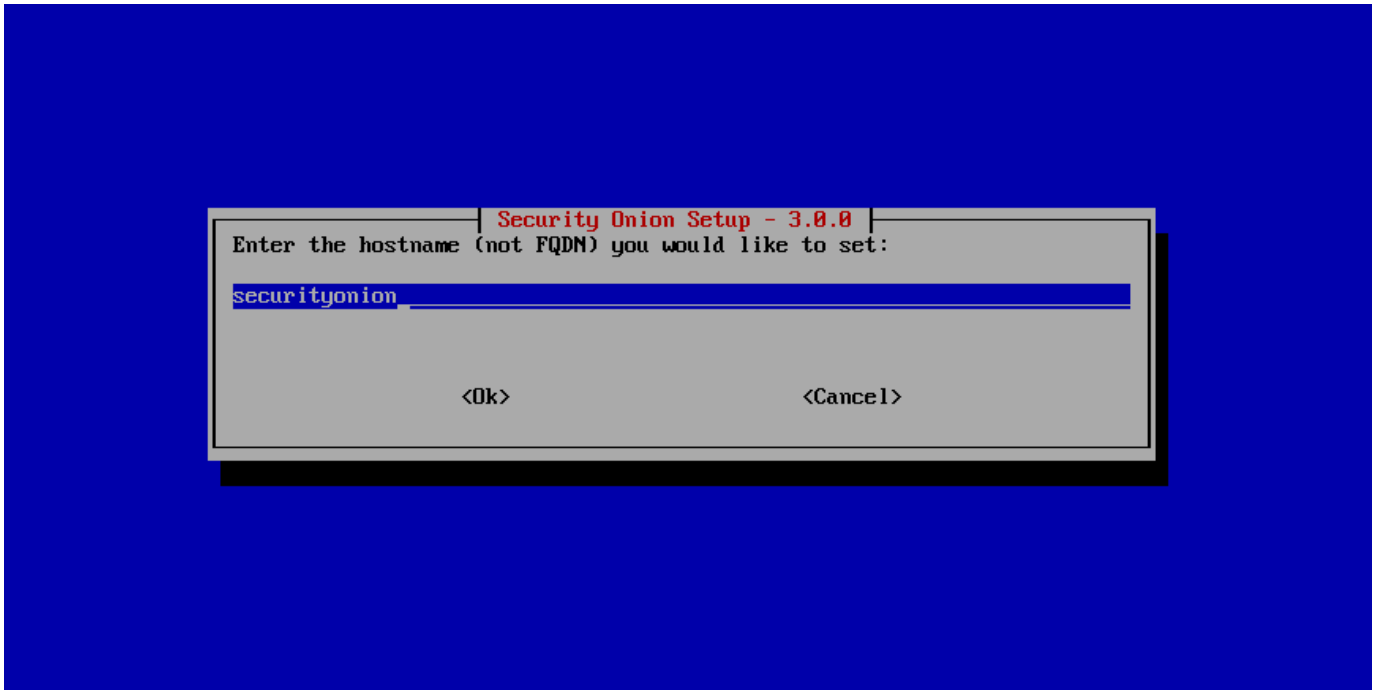

If you use the default hostname of `securityonion`, you will see a warning:

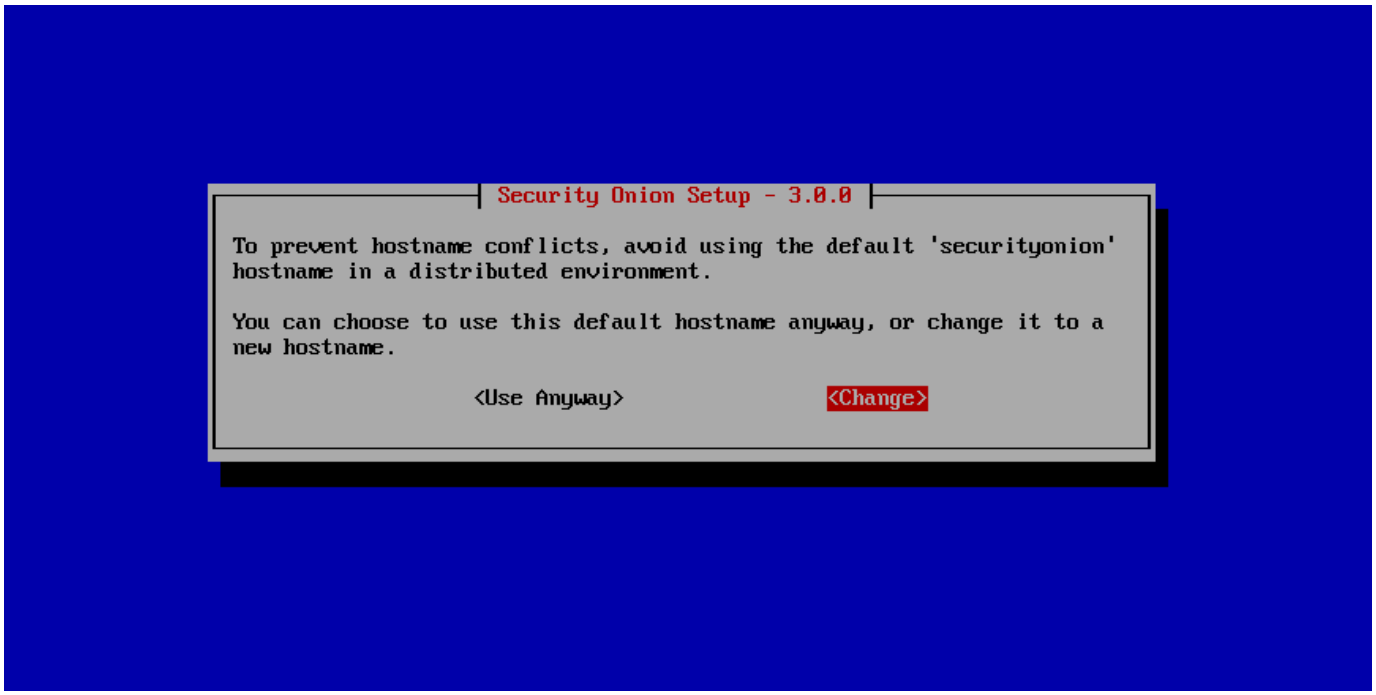

Select your management interface:

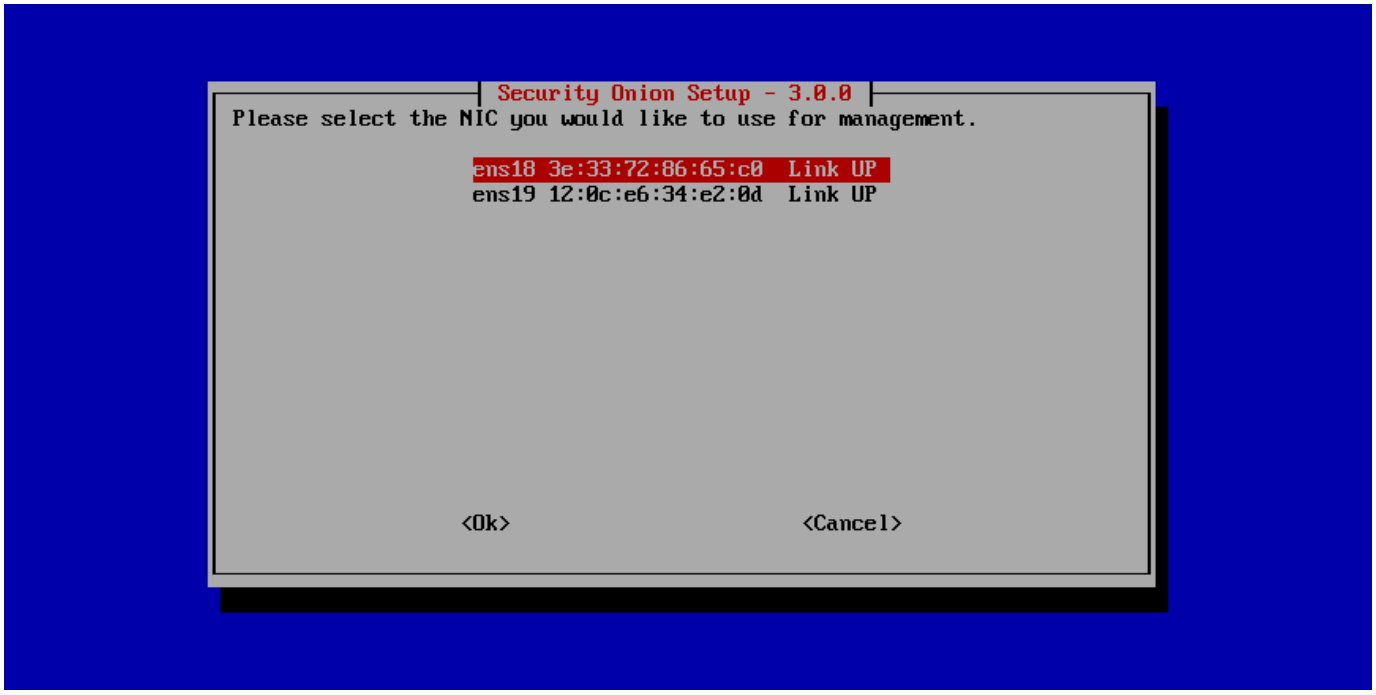

Select static IP addressing (recommended) or DHCP:

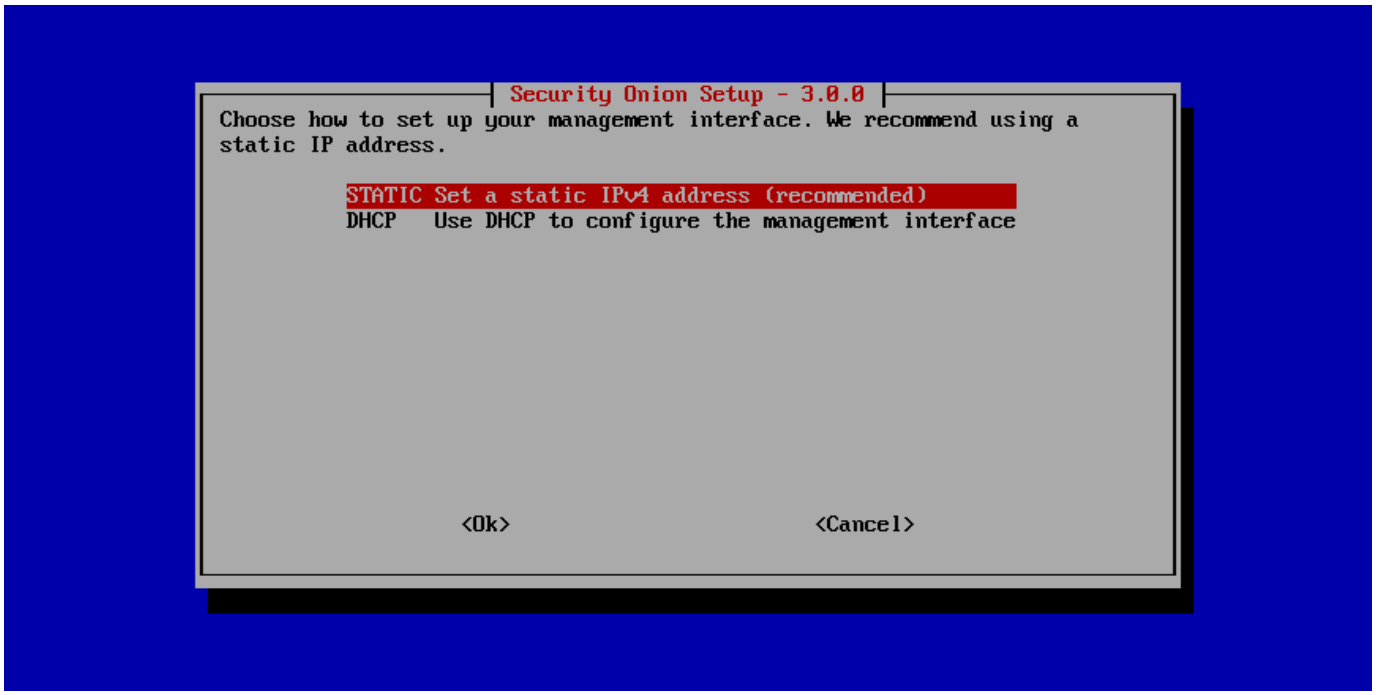

Specify IP address and CIDR mask:

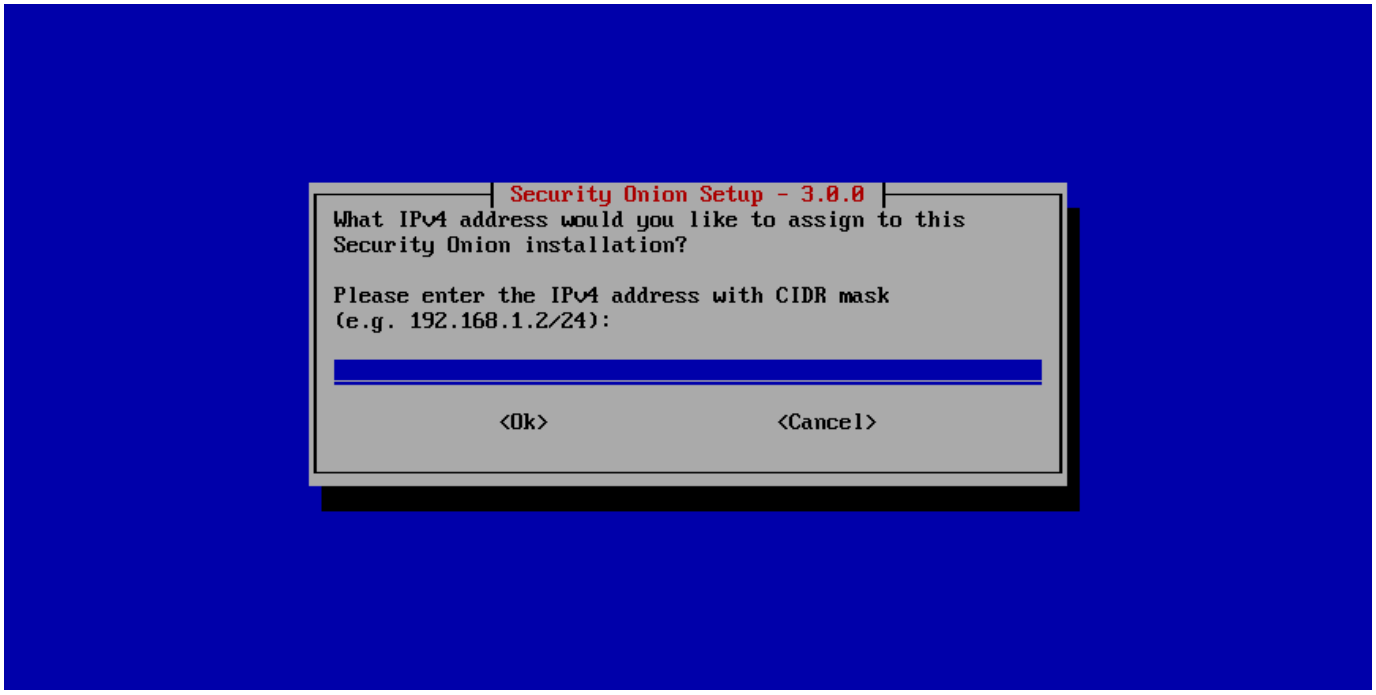

Set gateway address:

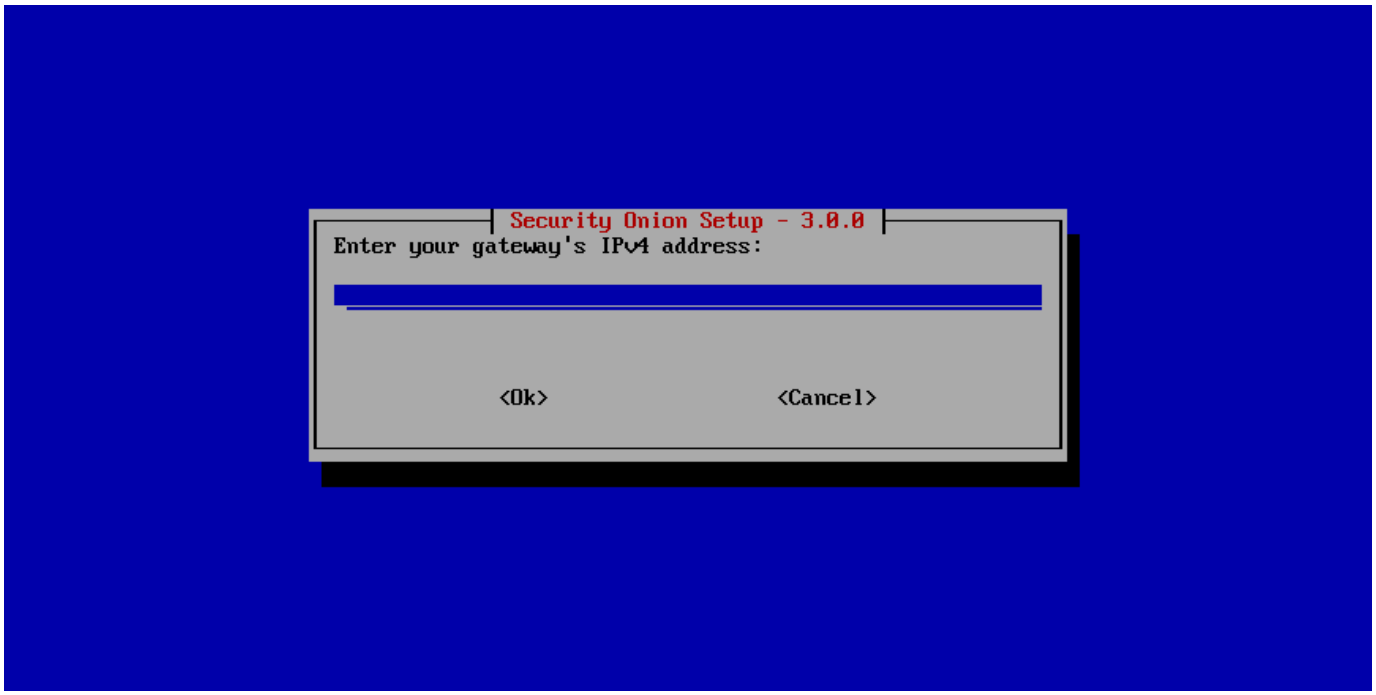

Enter DNS servers:

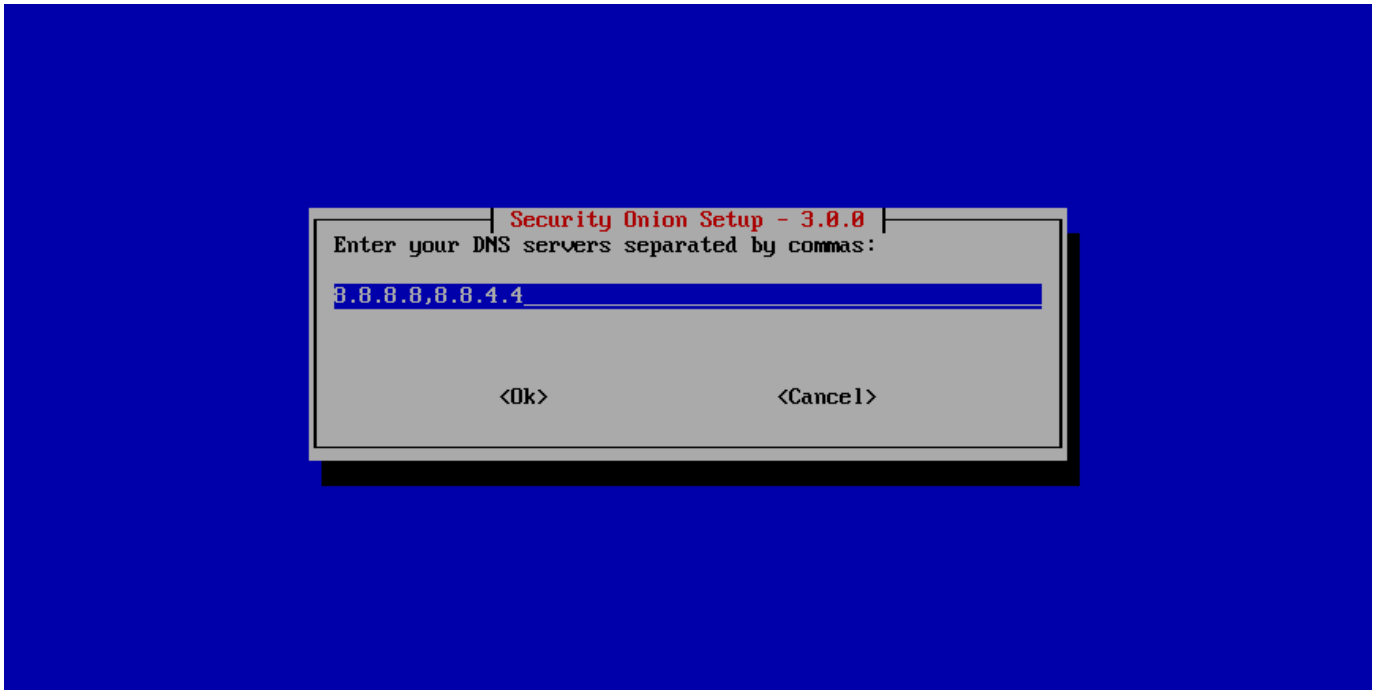

Configure DNS search domain:

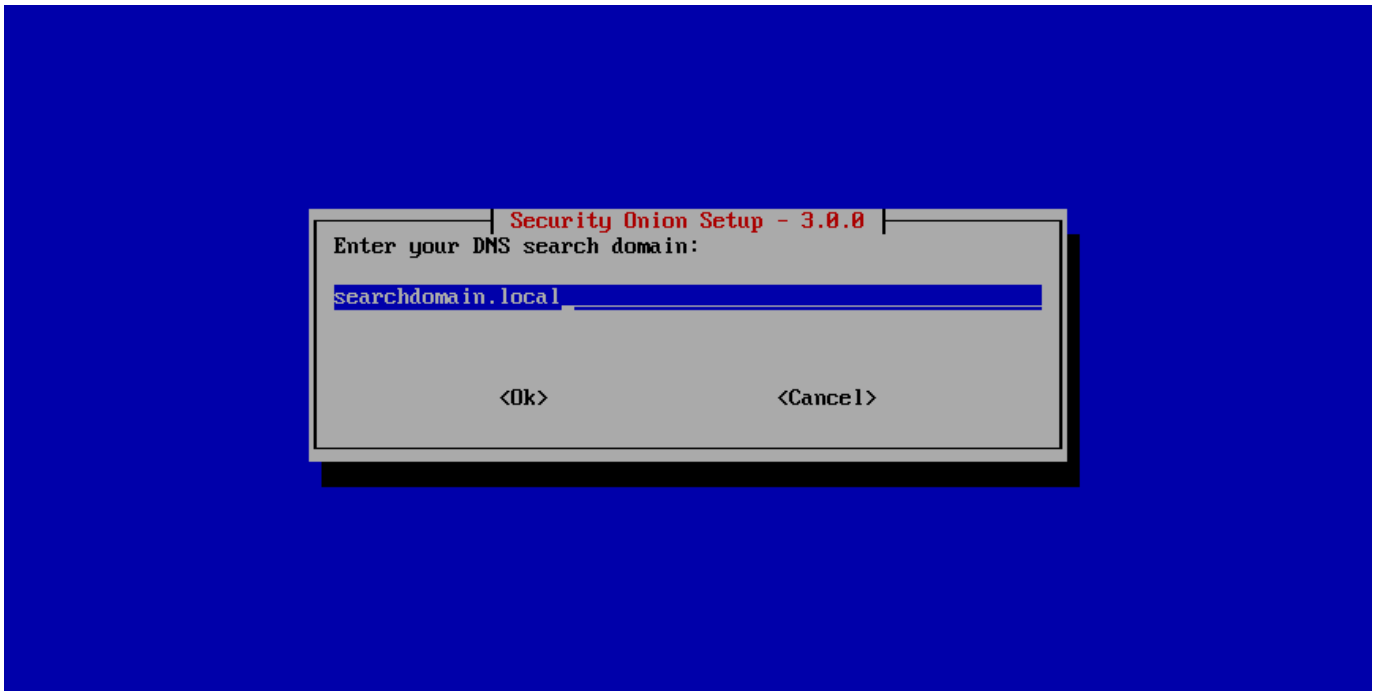

If necessary, you can change the default Docker IP range:

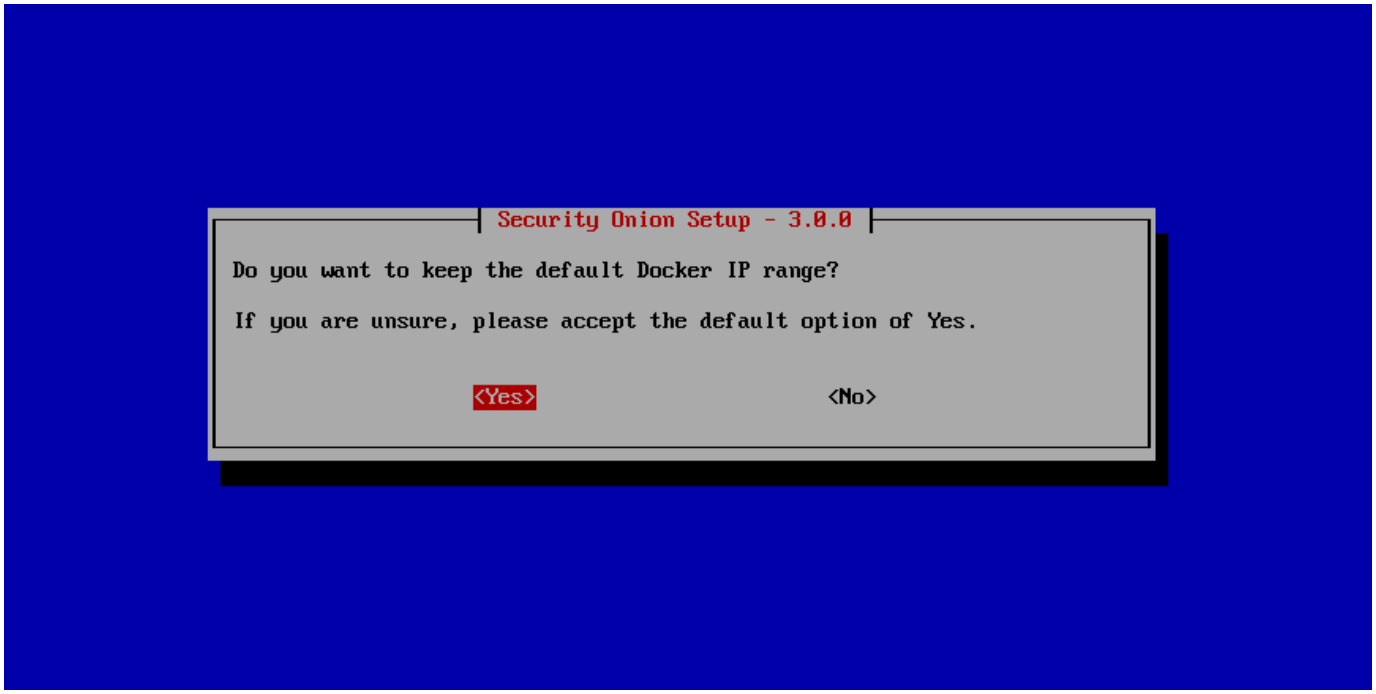

If you are connected to the Internet, select whether it is direct or via proxy:

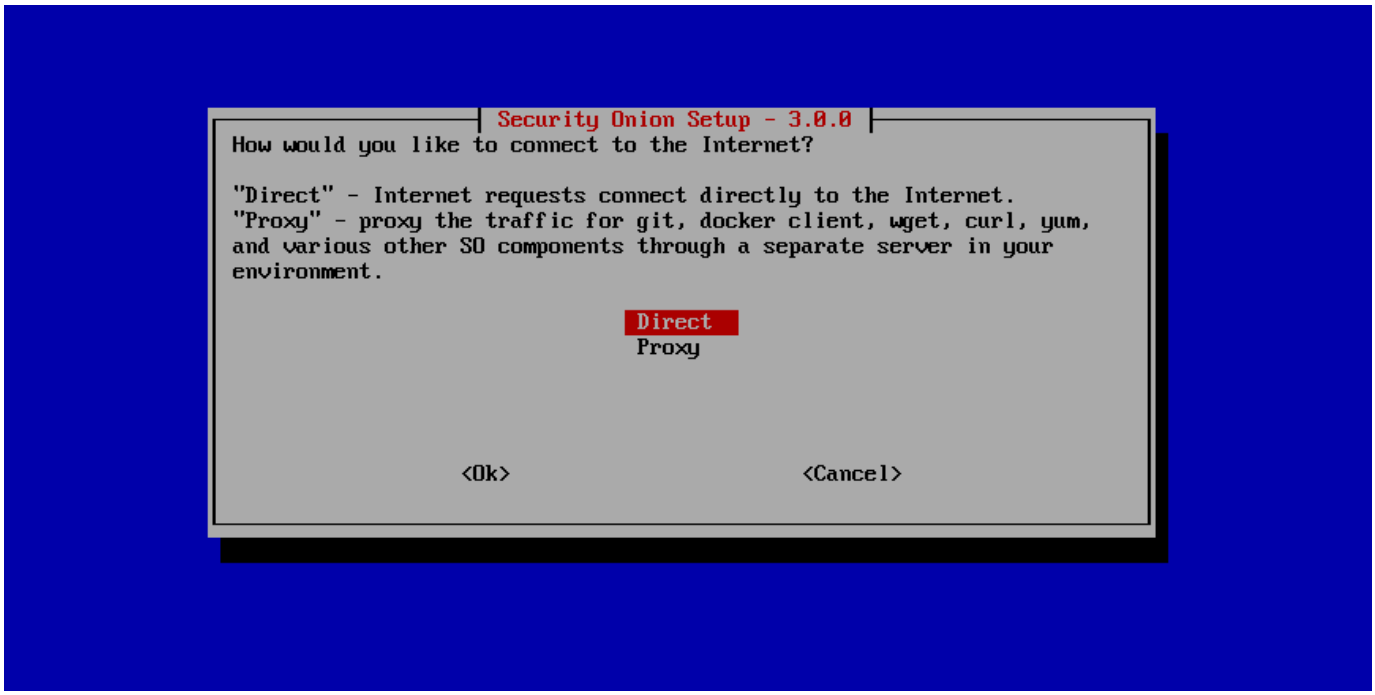

Create username for Security Onion Console:

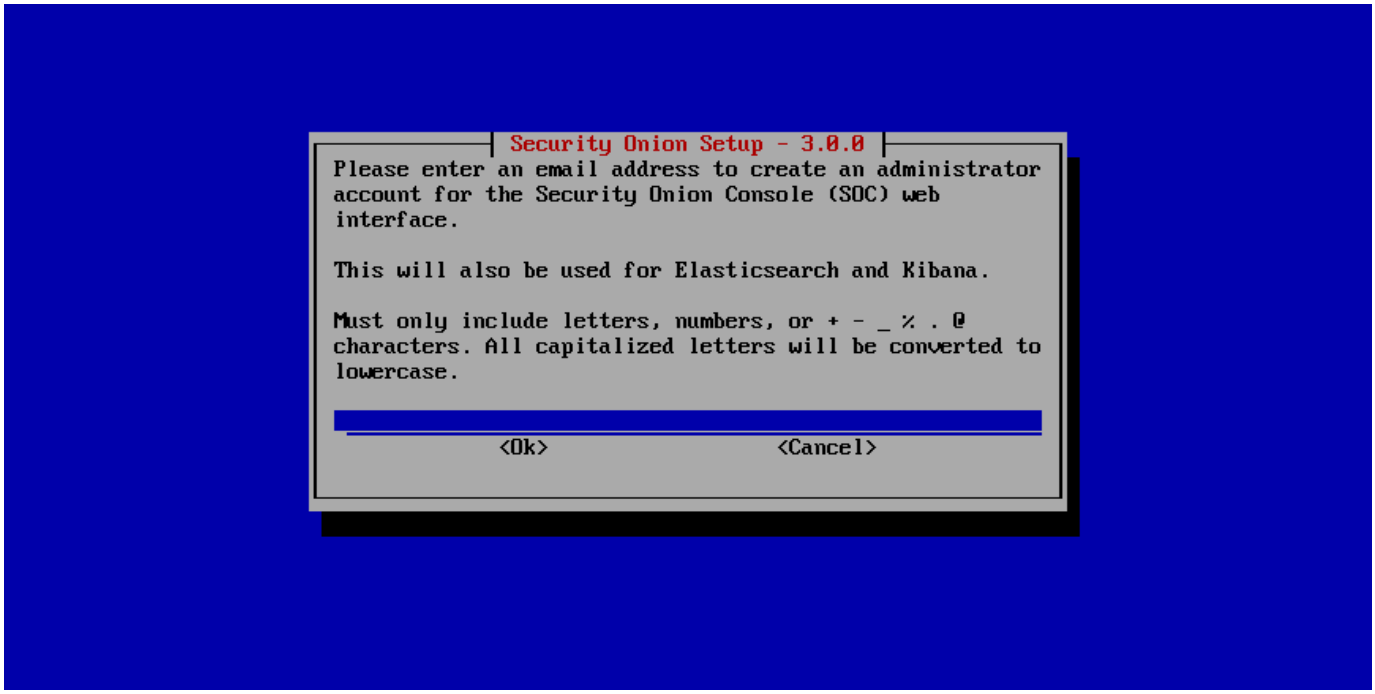

Set password for Security Onion Console:

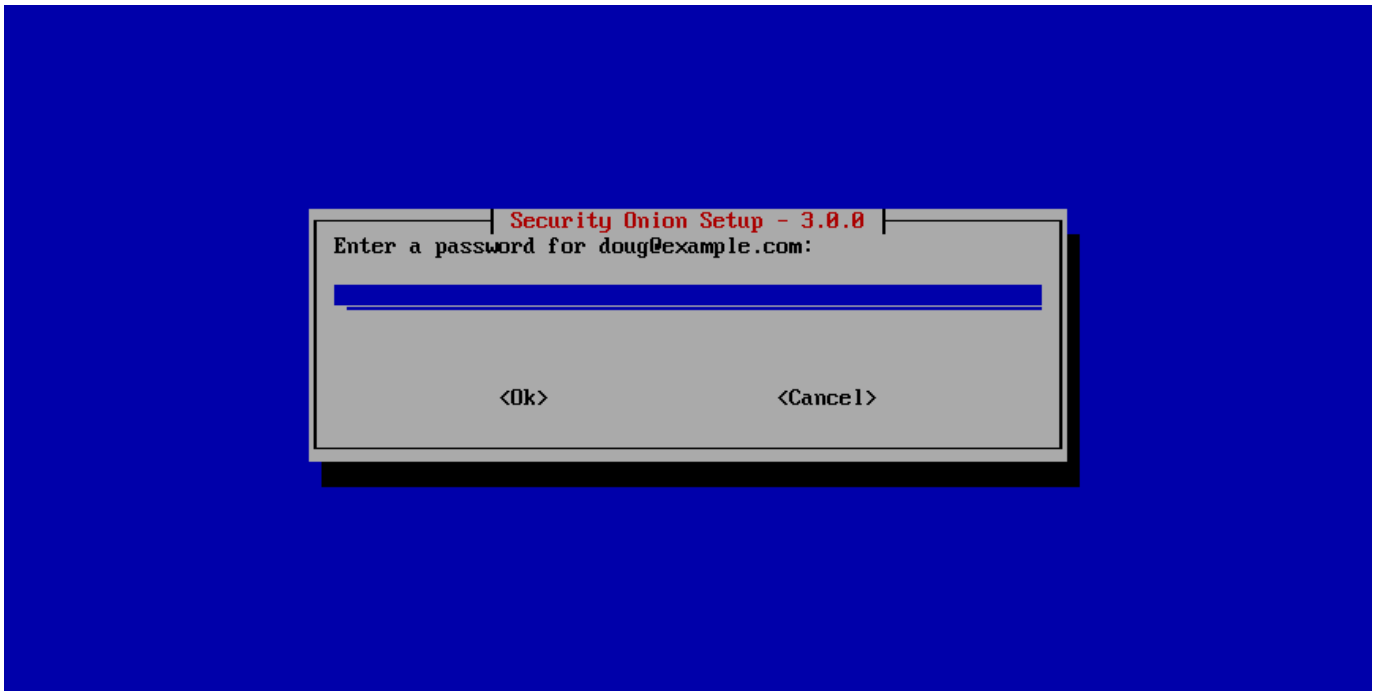

Confirm password for Security Onion Console:

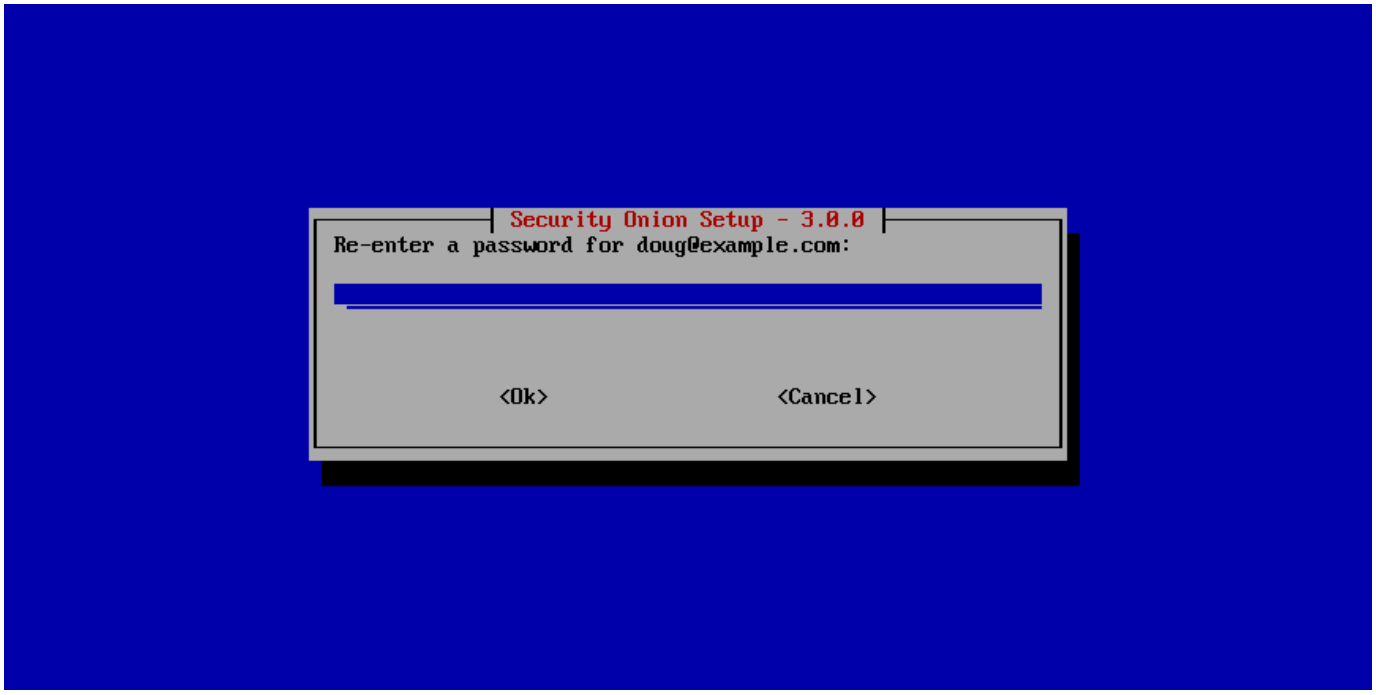

Select how to access Security Onion Console:

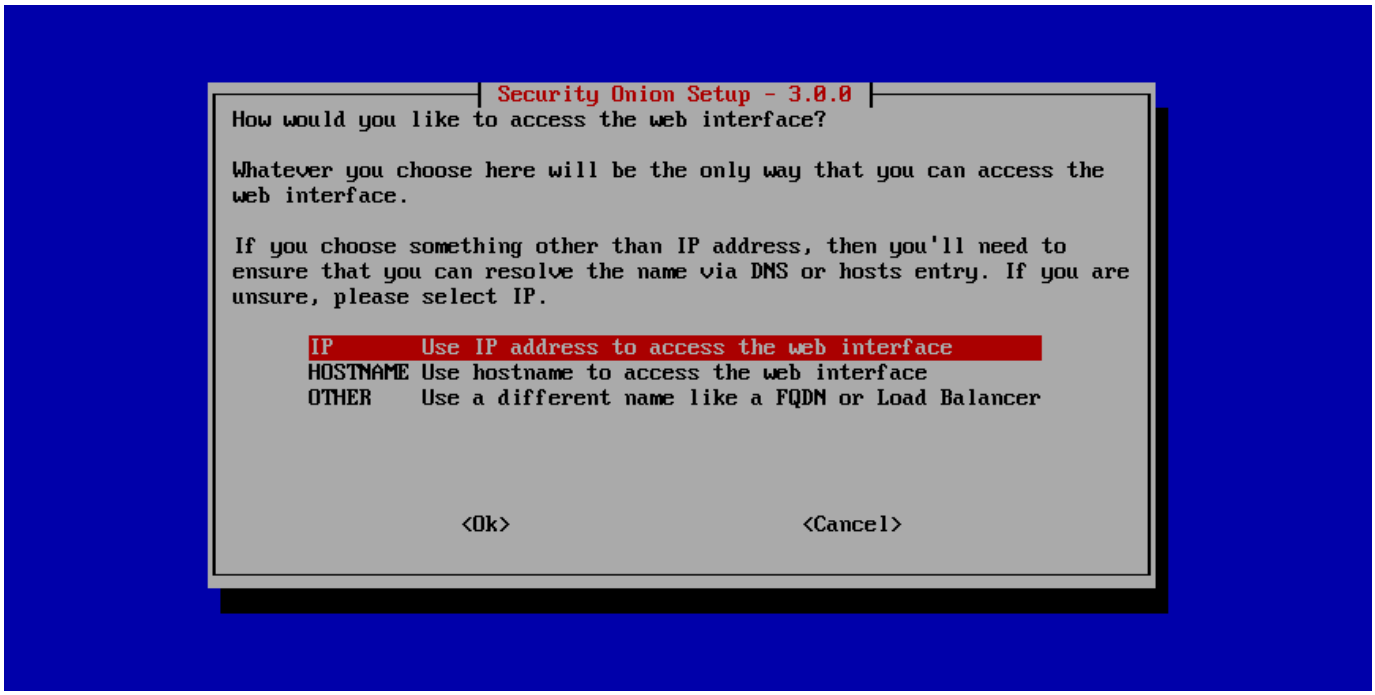

Allow connections through the host-based firewall if necessary:

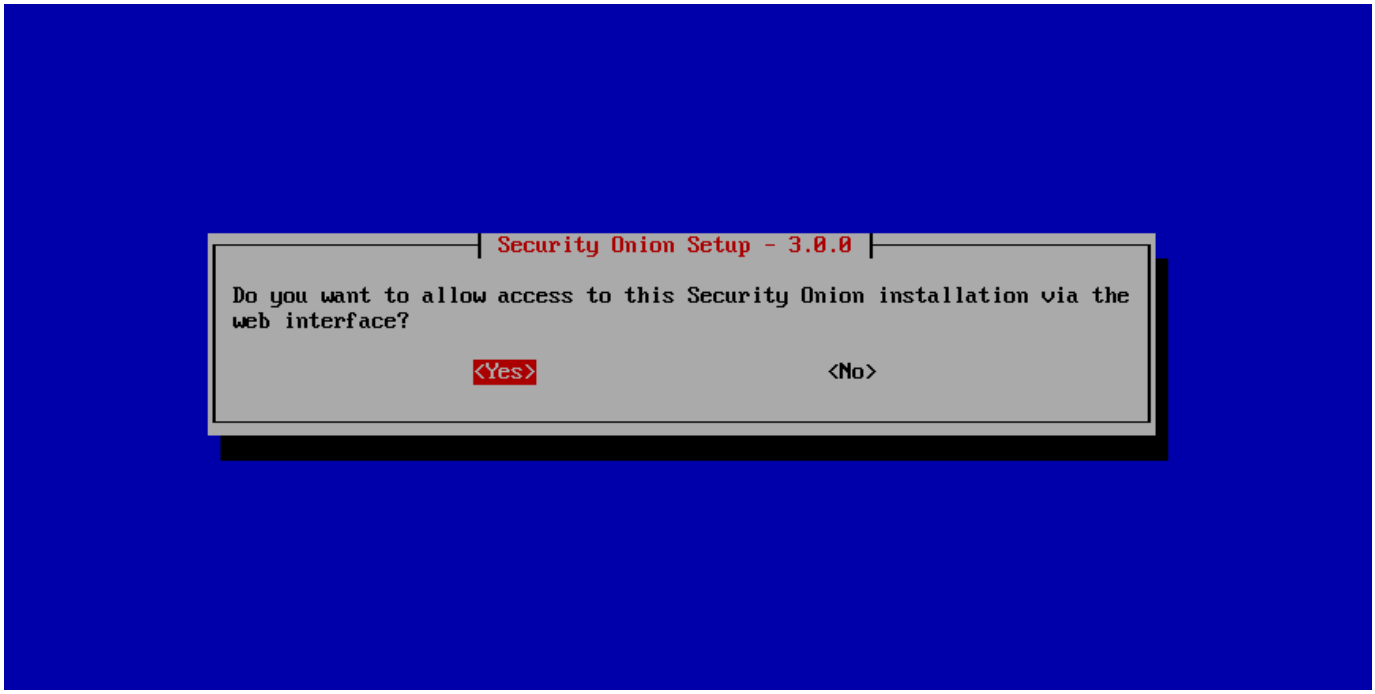

Specify an IP address or range to allow through the host-based firewall:

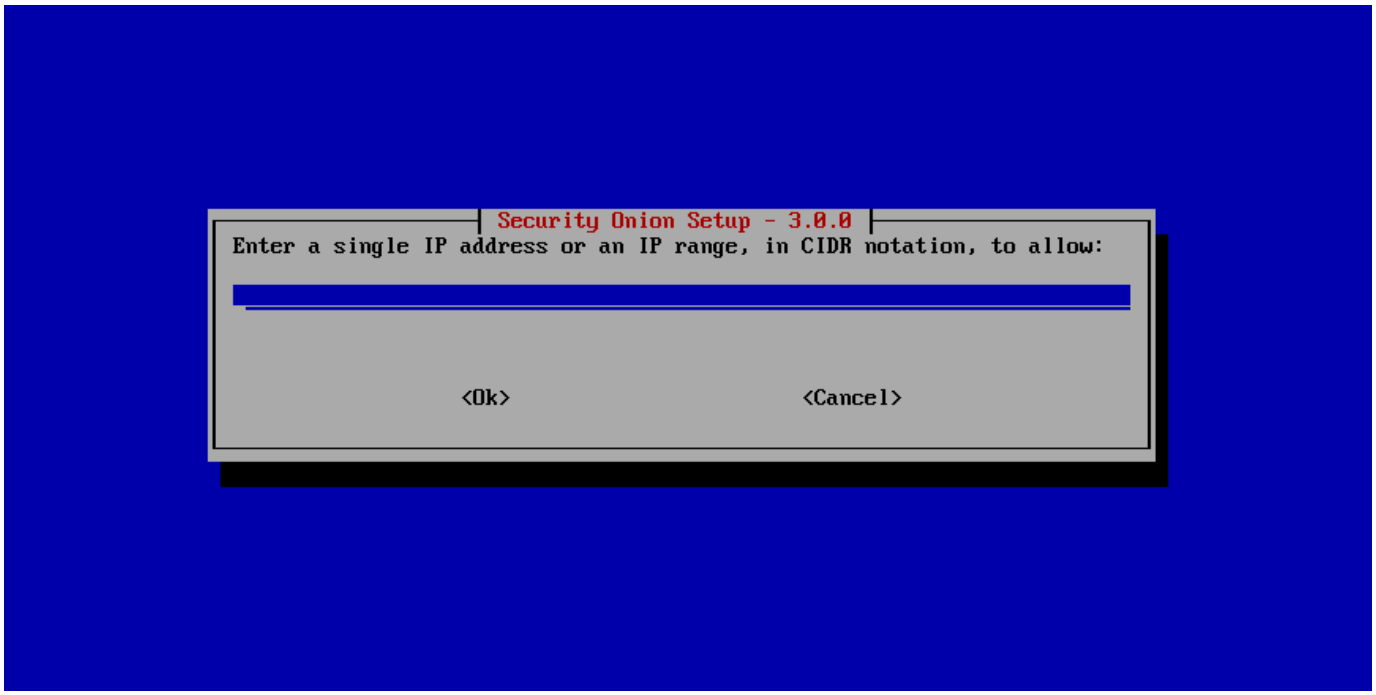

Confirm all options:

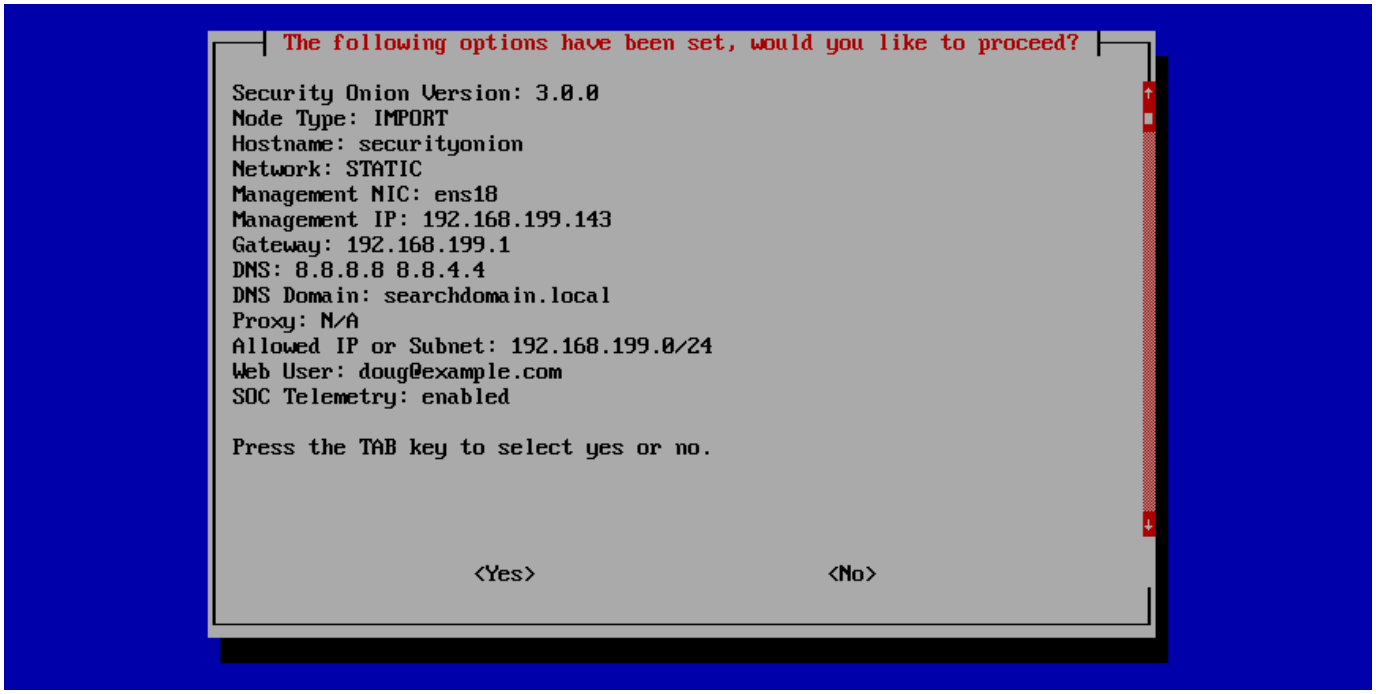

Setup complete:

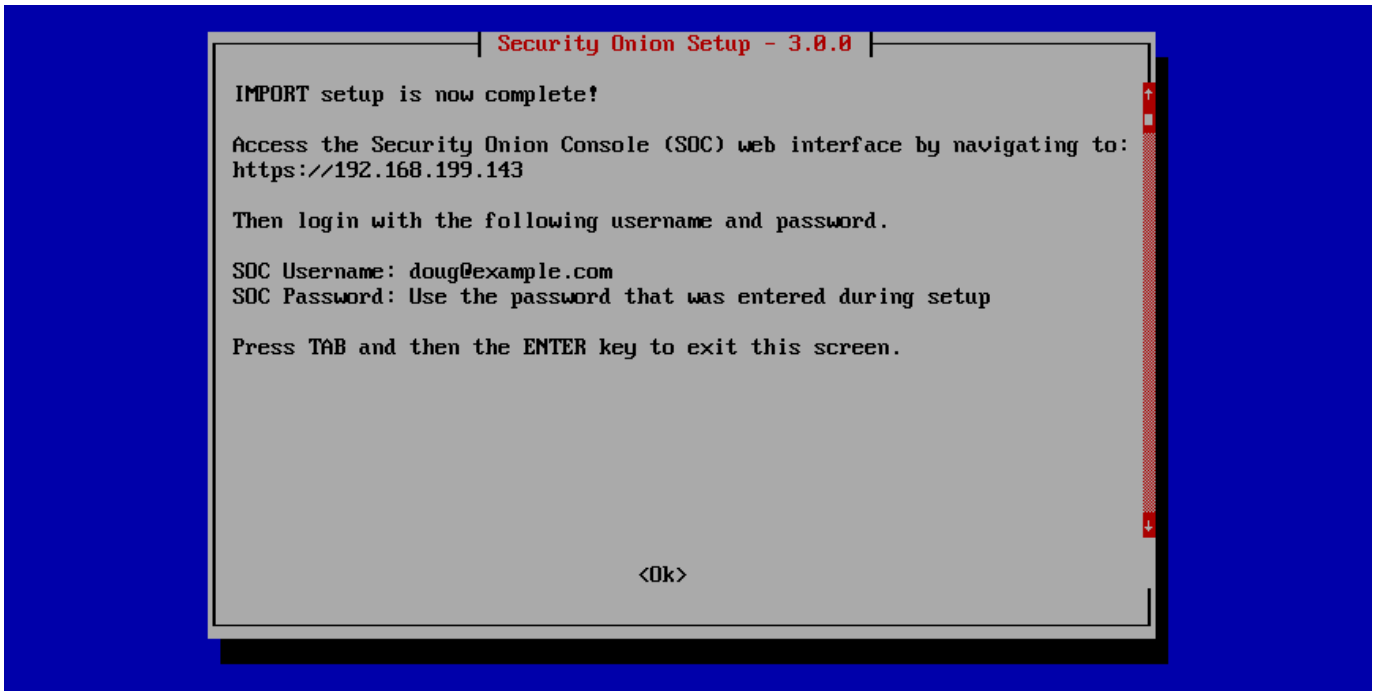

Login to Security Onion Console:

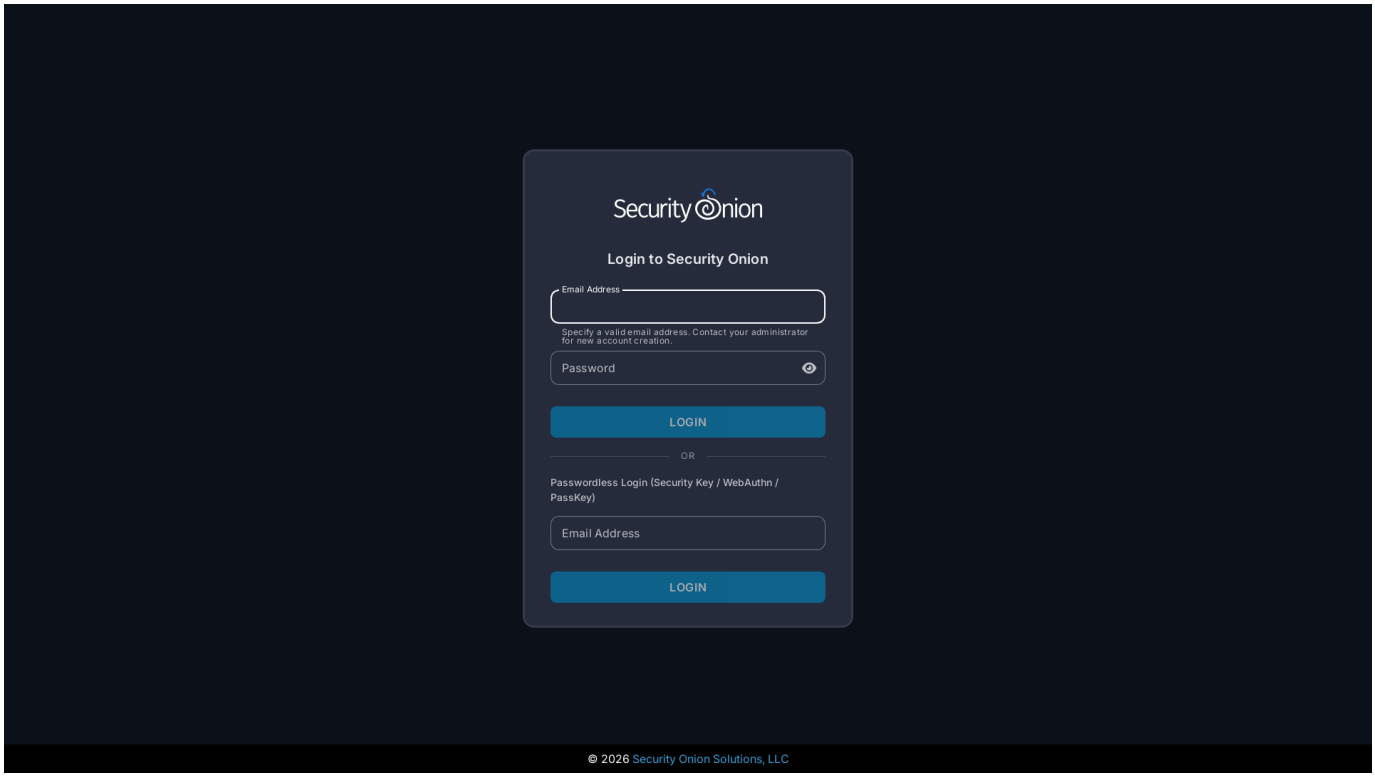

After logging in, you will see the Security Onion Console Overview page:

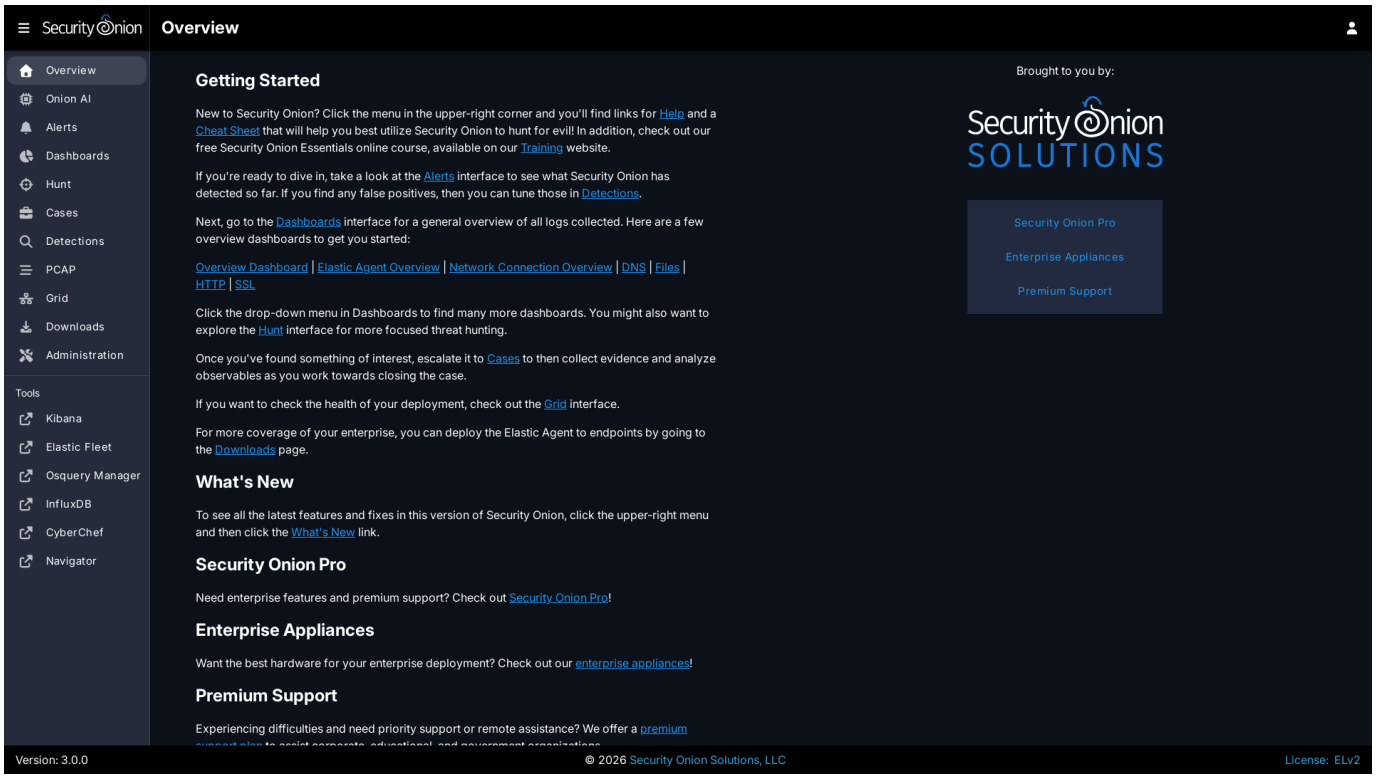

Go to the Grid page, click the button to expand the node, and then verify all services are running properly:

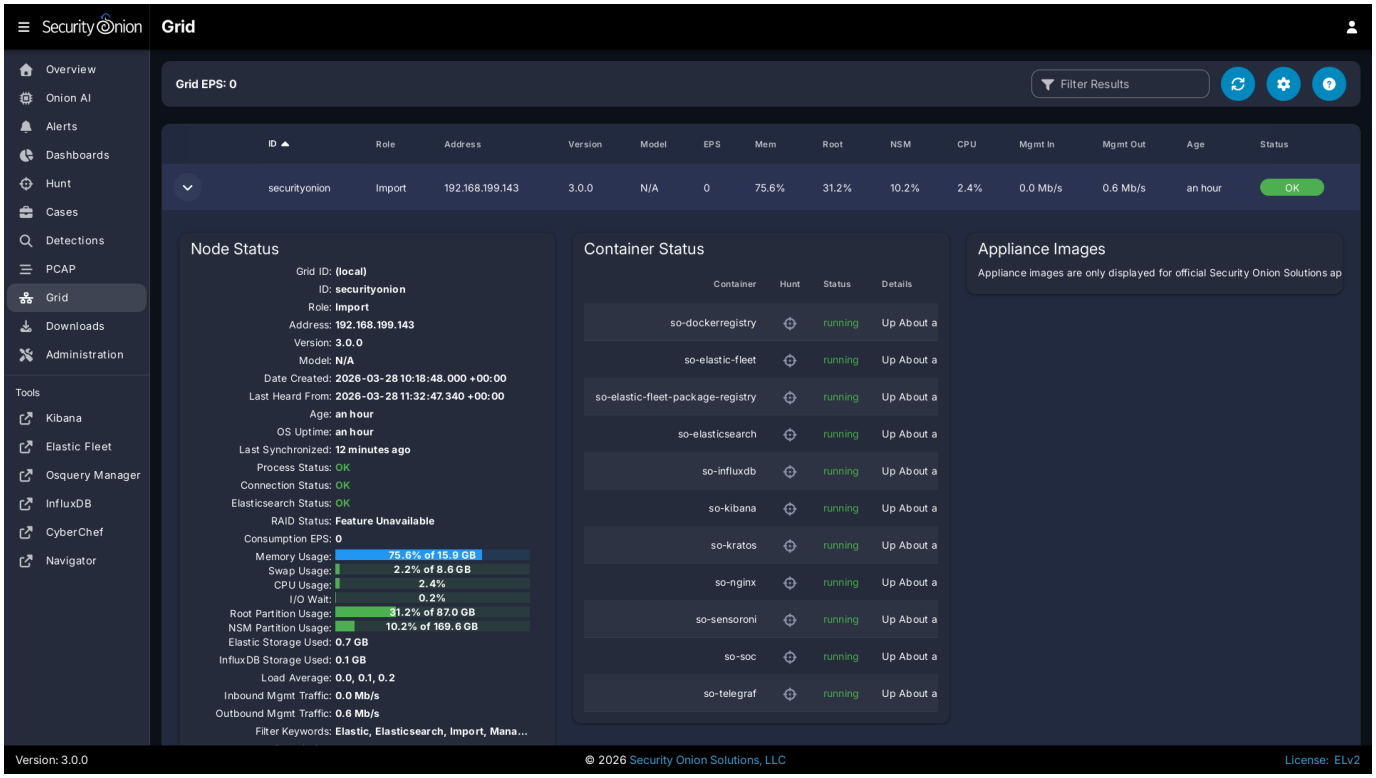

While on the Grid page, you can import a PCAP or EVTX file using the upload button at the bottom of the screen:

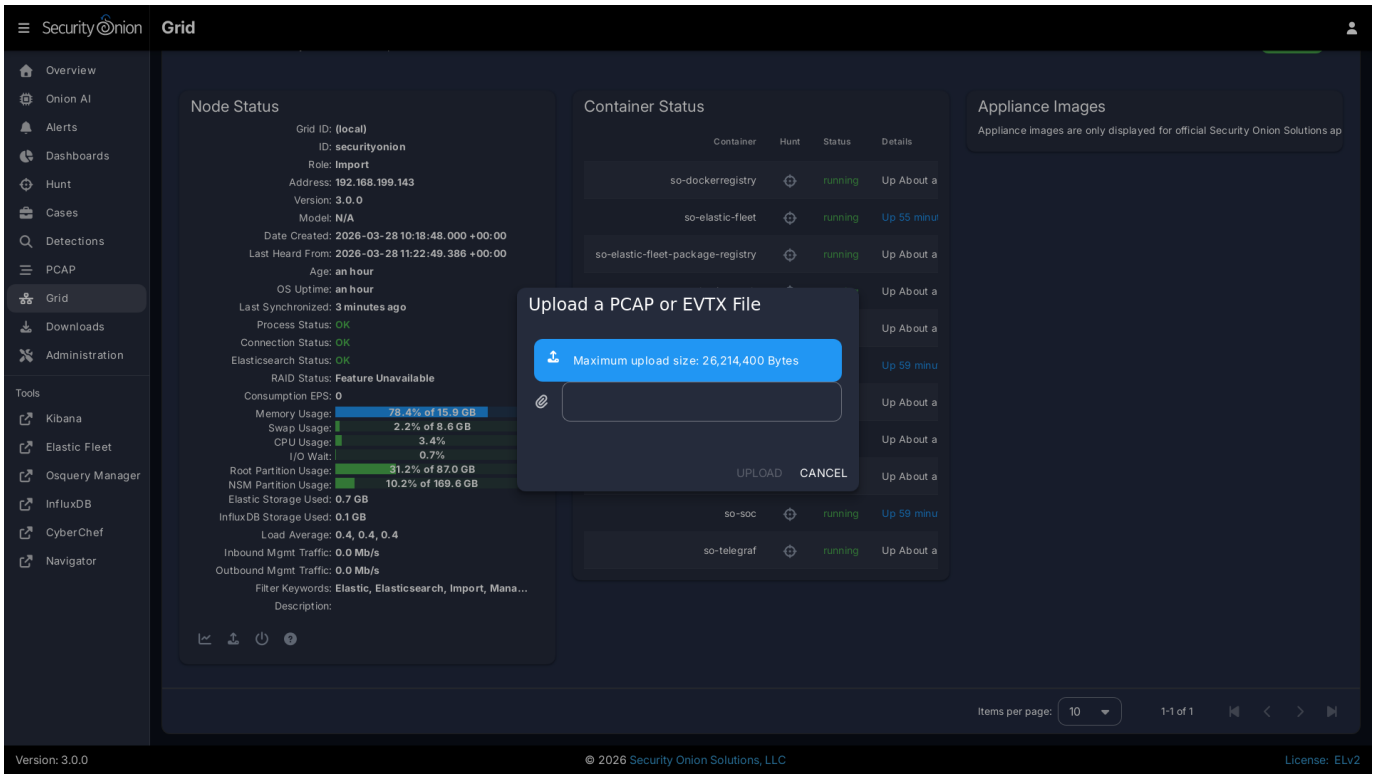

Once the import is complete, the Grid page should display a message at the top of the page and provide a link to Dashboards to view all alerts and logs from the import:

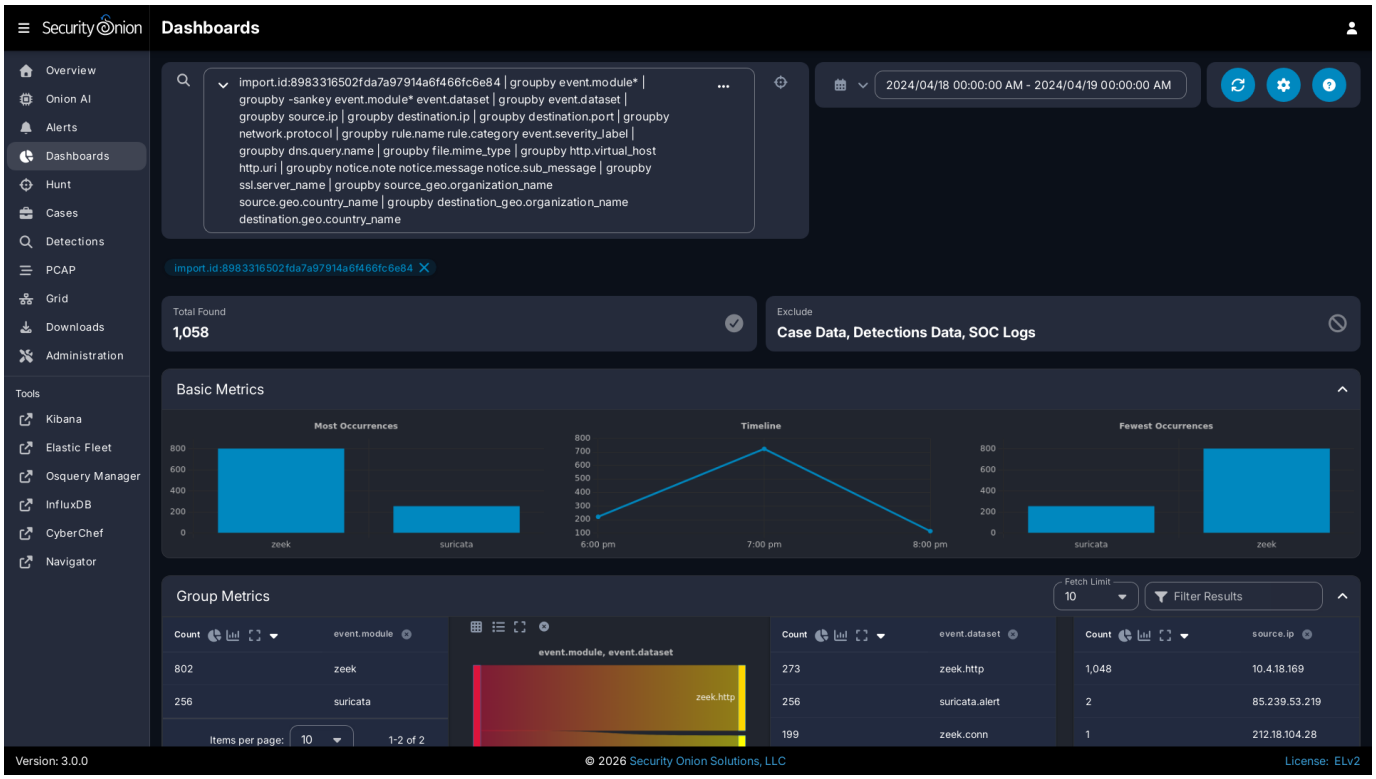

If you want to see just the alerts, you can go to the Alerts page although you may need to manually adjust the time range:

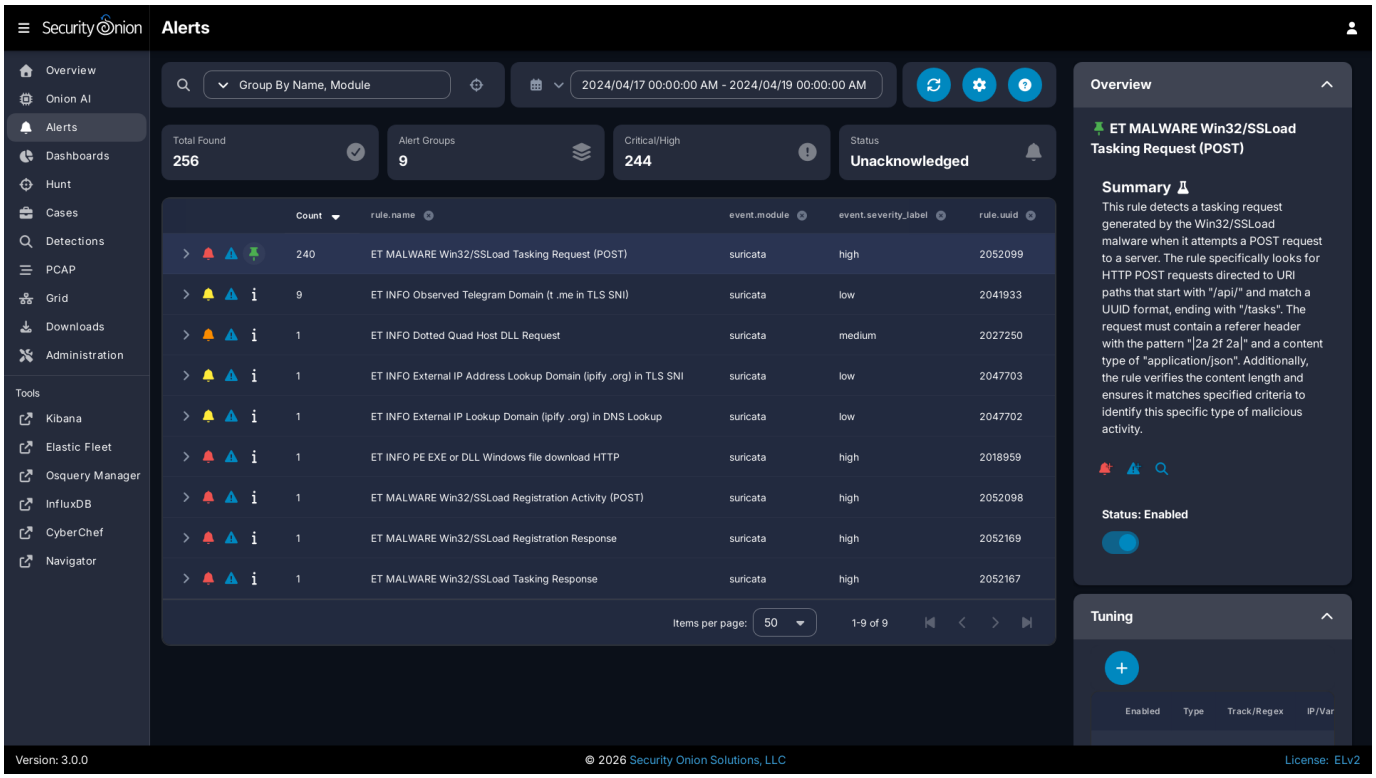

If you find something interesting on the Alerts or Dashboards pages, you may want to use the Correlate or Hunt actions to find related logs on the Hunt page:

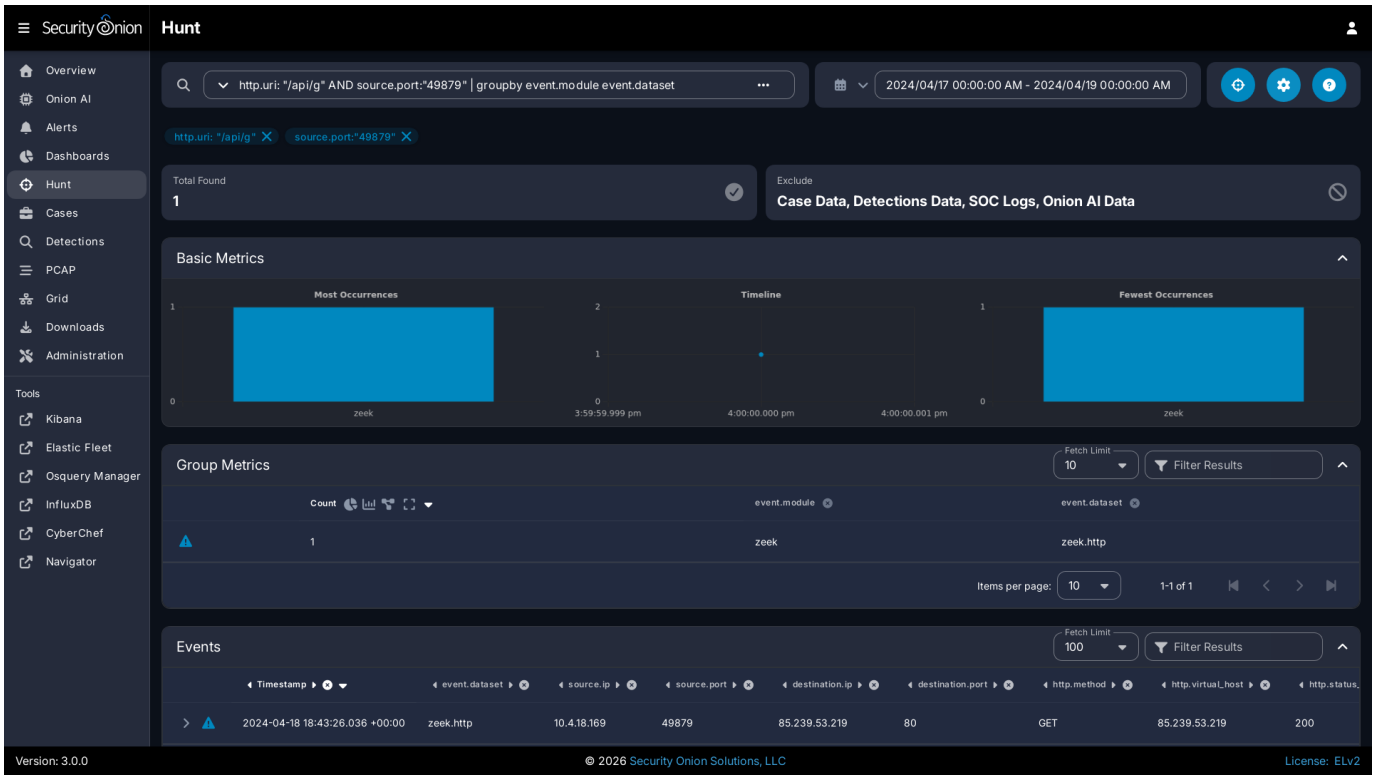

If you find interesting network traffic, you can pivot to full packet capture via the PCAP action:

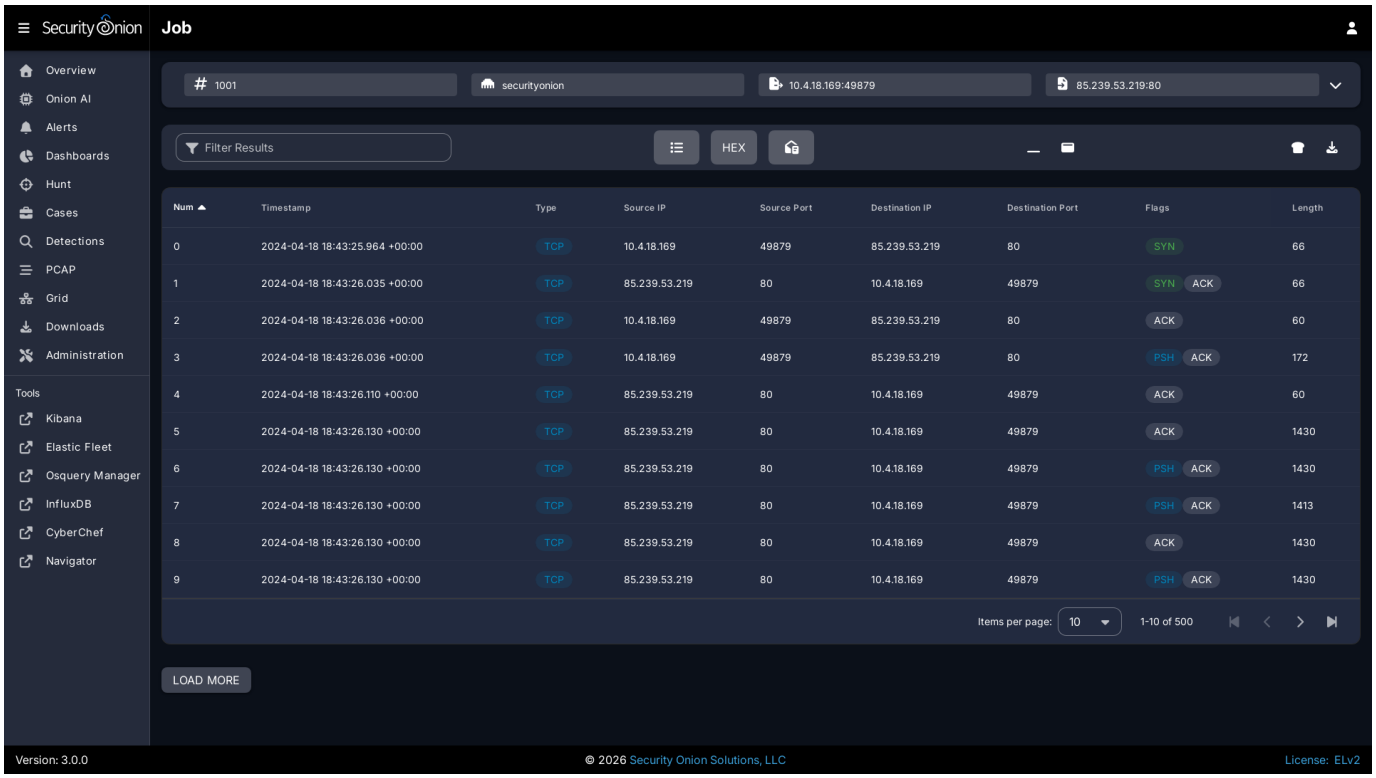

You can change the view to ASCII transcript for a more human readable view of the traffic:

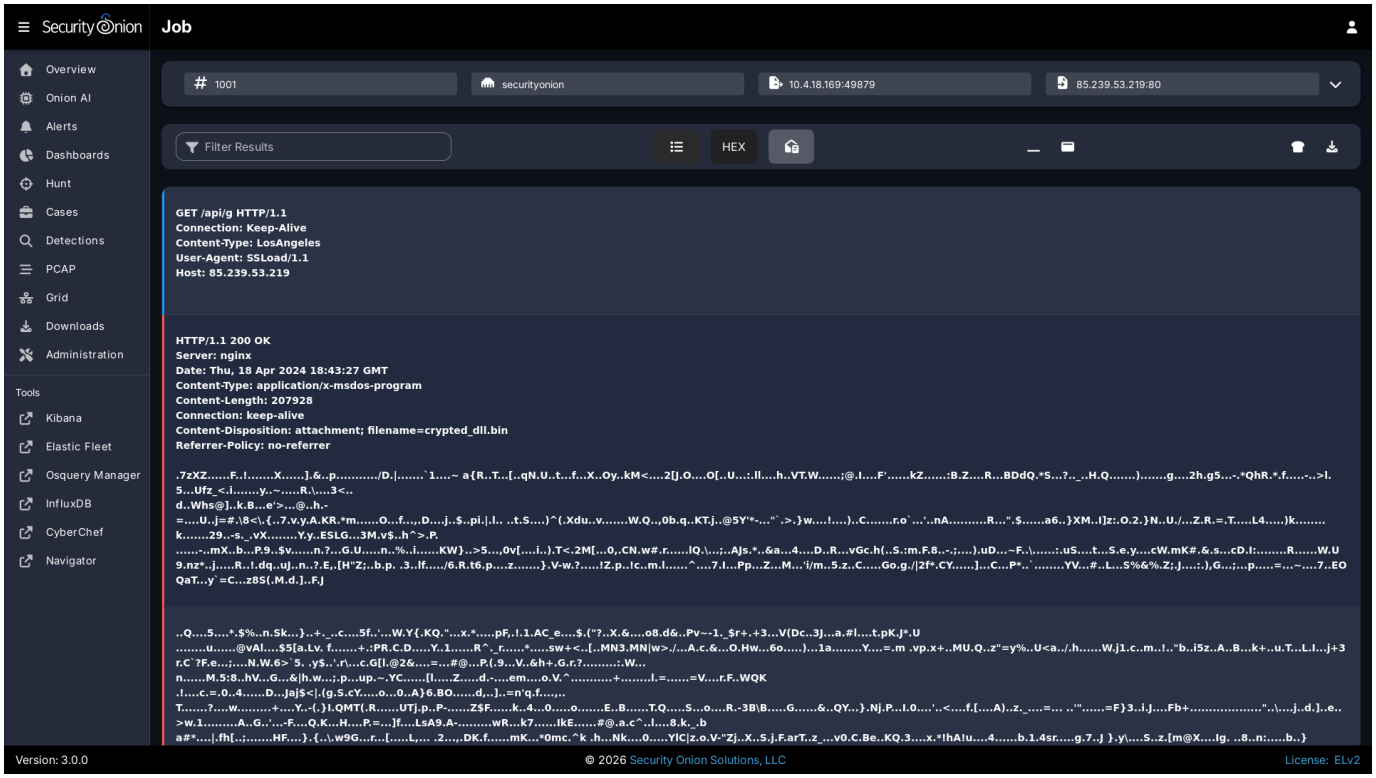

If you need to refer back to previous PCAP jobs, you can find them on the PCAP page:

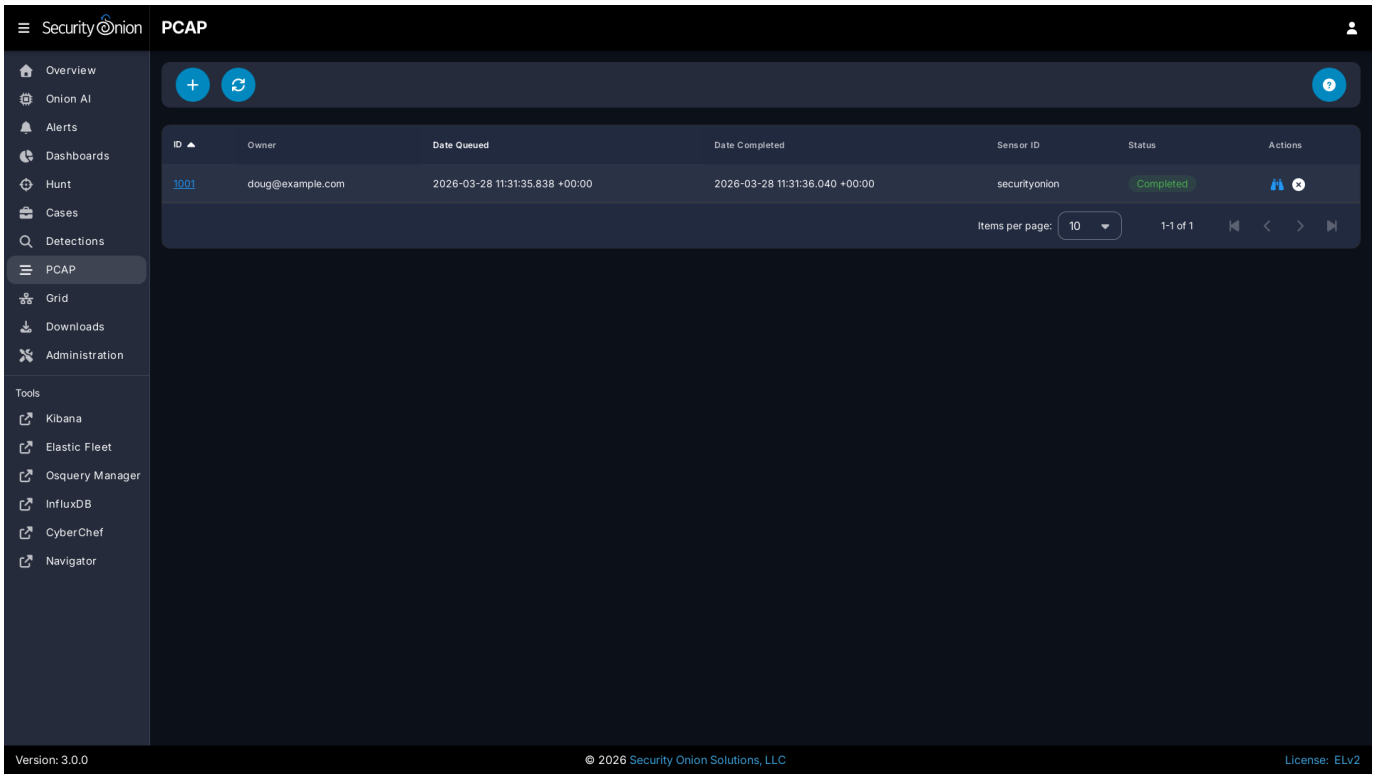

IMPORT installations do not support remote agents, but if you were running a production installation you could download the Elastic Agent installer from Downloads:

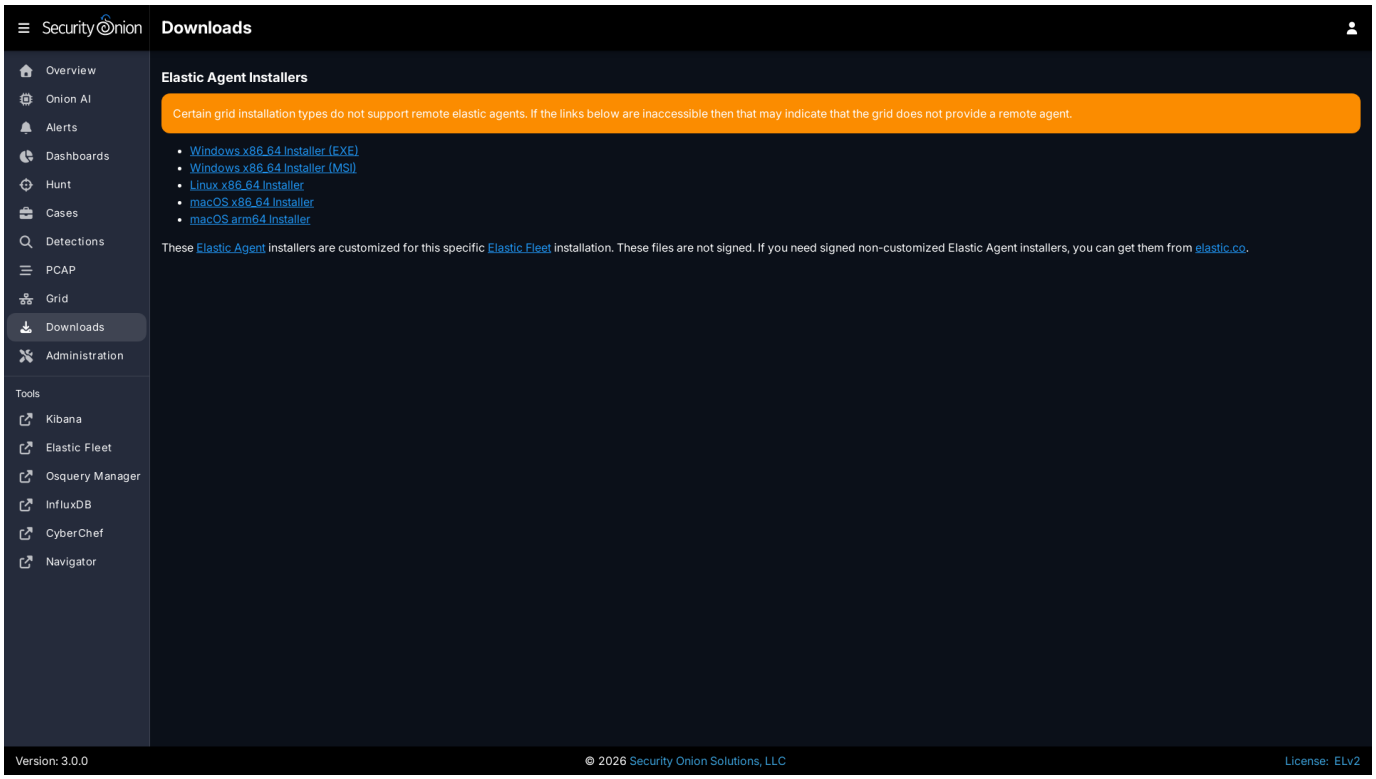

The Administration section allows you to manage user accounts:

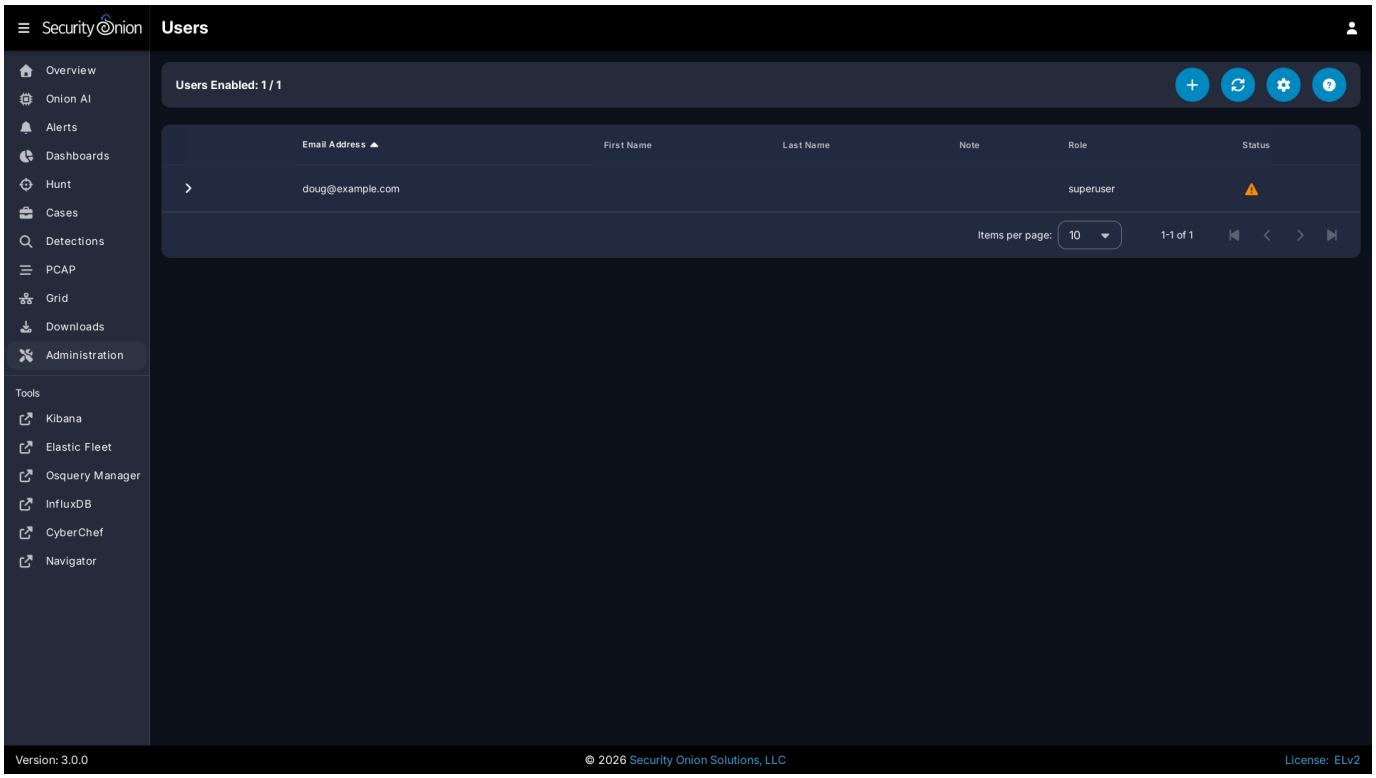

It also allows you to manage Grid members:

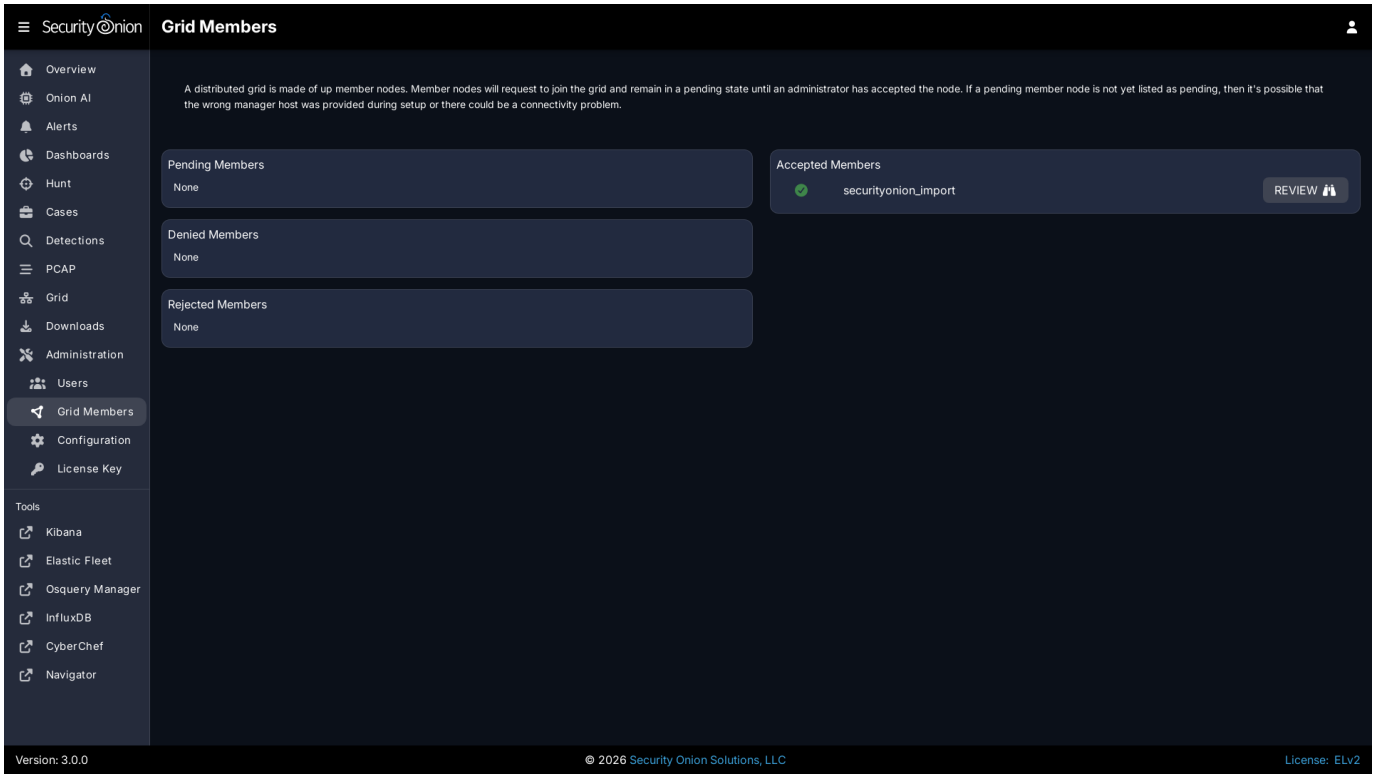

The Administration section also allows you to configure various aspects of the system:

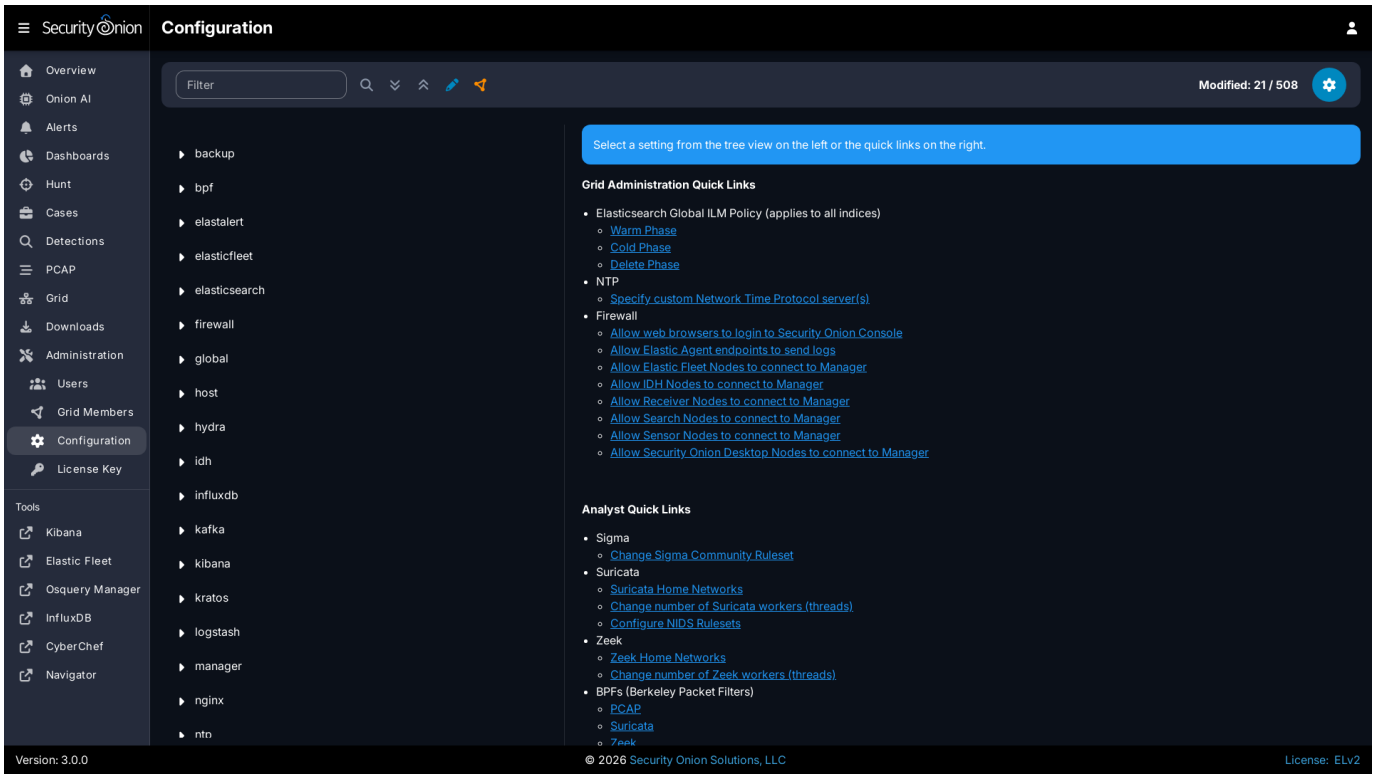

It also allows you to upload a license key for additional enterprise features:

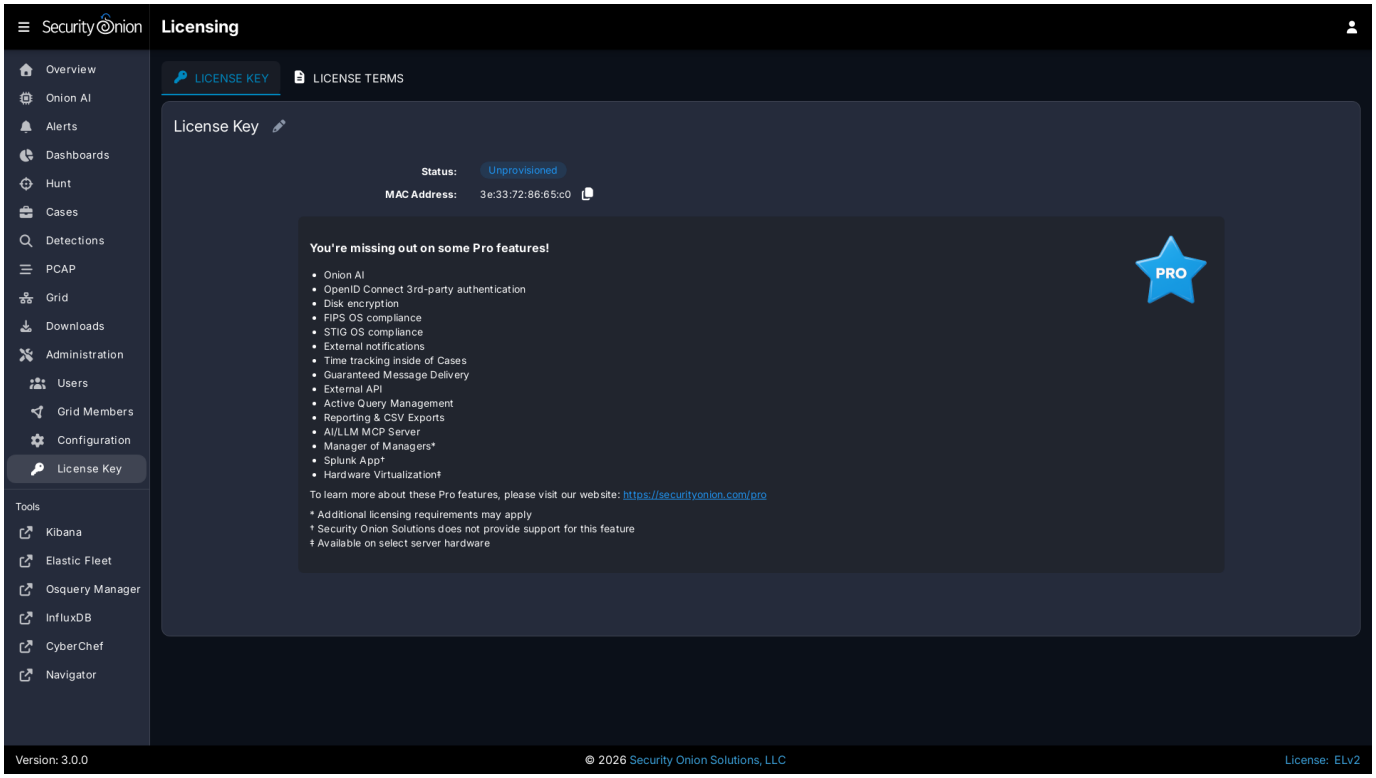

If you made it to the end of this First Time Users section, congratulations! If you have any questions or problems, please see the Help section. If you like Security Onion, please consider sharing on social media about Security Onion to help spread the word. Thanks!

5. Getting Started

5.1 Getting Started Overview

If you're ready to get started with Security Onion, you may have questions like:

What are the recommended best practices?

See the Best Practices section.

How many machines do I need?

Depending on what you're trying to do, you may need anywhere from one machine to thousands of machines. The Use Cases and Architecture sections will help you decide.

What kind of hardware does each of those machines need?

This could be anything from a small virtual machine to a large rack mount server with lots of CPU cores, lots of RAM, and lots of storage. The Hardware section provides further details.

If I just want to try Security Onion in a virtual machine, how do I create a virtual machine?

See the VMware, VirtualBox, and Proxmox sections.

How do I deploy Security Onion in the cloud?

See the Amazon Cloud, Azure Cloud, and Google Cloud sections.

What if I have trouble booting the ISO image?

Check out the Trouble Booting section.

What if I'm on an airgap network?

Review the Airgap section.

Once I've booted the ISO image, how do I install it?

See the Installation section.

After installation, how do I configure Security Onion?

The Configuration section covers many different use cases.

Is there anything I need to do after configuration?

See the Post Installation section.

5.2 Best Practices

Security Onion provides lots of options and flexibility, but for best results we recommend the following best practices.

5.2.1 Installation

- Download and verify our ISO image as shown in the Download section.
- For production deployments, prefer dedicated hardware to VMs when possible (see the Hardware section).
- If VMs must be used, ensure that resources are properly dedicated to VMs to avoid resource contention.
- Use local storage and avoid NFS, NAS, iSCSI, etc.
- Adequately spec your hardware to meet your current usage and allow for growth over time.
- When possible, we recommend using a dedicated TAP rather than SPAN ports.
- Make sure that any network firewalls have the proper firewall rules in place to allow ongoing operation and updates (see the Firewall section).

5.2.2 Configuration

- Make sure that both hostname and IP address are correct during installation.
- Avoid changing hostname and IP address after installation.
- Linux is case sensitive where other operating systems might not be, so we recommend using lowercase for things like hostnames, usernames, etc.

5.2.3 Avoid Third Party Software and Modifications

- Security Onion is a free and open platform based on standard Linux distros, but we recommend treating it as an appliance and avoid installing third party software as this may conflict with our components and cause issues when updating.
- Avoid installing automation tools such as Puppet and Chef as these may conflict with our existing Salt automation.
- Avoid installing monitoring tools such as Zabbix as this may conflict with our existing Influxdb monitoring.
- Avoid installing third-party endpoint security agents as they may break functionality or introduce unacceptable performance overhead.
- Avoid changing file permissions or umask settings.
- Hardening guidelines may break functionality, so if you must apply those hardening guidelines, we recommend testing thoroughly before deploying to production.

5.2.4 Stay Up To Date

- Join our discussion forum at https://securityonion.net/discuss or subscribe to one of our social media channels to be notified of Security Onion updates.
- Keep your deployment updated as we frequently fix bugs and add new features.
- If possible, test updates on a test deployment before deploying to production.

5.3 Use Cases

If you're going to deploy Security Onion, you should first decide what your use case is. In this section, we'll discuss some common use cases and how they map to our different kinds of architecture. This could be anything from a temporary Import installation in a small virtual machine on your personal laptop all the way to a large scalable enterprise deployment consisting of a manager node, multiple search nodes, and lots of sensor nodes.

5.3.1 Minimal Import

Suppose you just want to import PCAP or EVTX files or suppose that you have limited hardware and just want the minimal installation to get some experience with Security Onion. Install Security Onion and choose the `Import` option. This can be done in a minimal virtual machine with as little as 4GB RAM. You can read more about the `Import` option in the First Time Users section and in the Architecture section.

5.3.2 Minimal Network Visibility

Suppose you have a small network where you just want some basic network visibility. This might be monitoring traffic from a TAP or SPAN port on a homelab or other small network that doesn't require a production installation. Install Security Onion and choose the `Evaluation` option. You can read more about the `Evaluation` option in the Architecture section.

5.3.3 Minimal Network Visibility with OPNsense Firewall instead of TAP or SPAN

Suppose you have a small network segment where you just want some basic network visibility but you don't have the ability to collect traffic via TAP or SPAN port but you do have an OPNsense firewall. Install Security Onion, choose the `ManagerSearch` option, and then follow the OPNsense section to collect firewall logs, Netflow data, and Suricata NIDS alerts. You can read more about the `ManagerSearch` option in the Architecture section.

5.3.4 Minimal Netflow Collector

Suppose you have a small network where you just want to collect some Netflow data. Install Security Onion, choose the `ManagerSearch` option, and then follow the Netflow section. You can read more about the `ManagerSearch` option in the Architecture section.

5.3.5 Minimal Log Management

Suppose you have a small network where you just want some basic host visibility. This might be deploying agents to a small number of desktops and servers and/or collecting syslog from firewall or other devices. Install Security Onion, choose the `ManagerSearch` option, and then deploy the Elastic Agent to your hosts and review the Host and Third Party Integrations sections. You can read more about the `ManagerSearch` option in the Architecture section.

5.3.6 Minimal Network and Host Visibility

Suppose you have a small network where you want both network visibility and host visibility. Install Security Onion and choose the `Standalone` option. This machine will then sniff network traffic from your TAP or SPAN port and also support deploying the Elastic Agent to other hosts. You can read more about the `Standalone` option in the Architecture section.

5.3.7 Minimal Enterprise Deployment

Suppose you have a small or medium network where you want some visibility for both network and hosts. A minimal enterprise deployment would look like this:

- Install the first Security Onion instance and choose the `ManagerSearch` option.
- Deploy the Elastic Agent to hosts.
- Install Security Onion on one or more additional machines and join them to the grid as sensor nodes. They will analyze network traffic from your TAP or SPAN port.

You can read more about distributed deployments in the Architecture section.

5.3.8 More Scalable Enterprise Deployment

Suppose you have a medium or large network where you want some visibility for both network and hosts. A more scalable enterprise deployment would look like this:

- Install the first Security Onion instance and choose the `Manager` option.
- Install Security Onion on one or more additional machines and join them to the grid as search nodes. They will store logs and allow you to search them.
- Deploy the Elastic Agent to hosts. They will collect logs and send them to the grid.
- Install Security Onion on one or more additional machines and join them to the grid as sensor nodes. They will analyze network traffic from your TAP or SPAN port.

You can read more about distributed deployments in the Architecture section.

5.3.9 Comprehensive Enterprise Deployment

Suppose you have a large network where you want maximum visibility for both network and hosts. A comprehensive distributed deployment would look like this:

- Install the first Security Onion instance and choose the `Manager` option.
- Install Security Onion on one or more additional machines and join them to the grid as search nodes. They will store logs and allow you to search them.
- Install Security Onion on a machine in your DMZ and join it to the grid as a Fleet node. This node will manage your Elastic agents whether they are onsite or offsite.
- Deploy the Elastic Agent to hosts. They will collect logs and send them to the grid.
- Install Security Onion on one or more additional machines and join them to the grid as sensor nodes. They will analyze network traffic from your TAP or SPAN port.
- Install Security Onion on one or more additional machines and join them to the grid as receiver nodes. This provides load balancing and pipeline redundancy.
- Install Security Onion on one or more additional machines and join them to the grid as IDH nodes. They will provide honeypot and deception capabilities.

You can read more about distributed deployments in the Architecture section.

5.4 Architecture

In the Use Cases section, we looked at a few of the most common use cases. This section will discuss what those different use cases look like from an architecture perspective.

5.4.1 Import

The simplest architecture is an `Import` node. An import node is a single standalone box that runs just enough components to be able to import PCAP or evtx files using the Grid page. It does **not** support adding Elastic agents or additional Security Onion nodes. For a full walkthrough of the `Import` option, please see the First Time Users section.

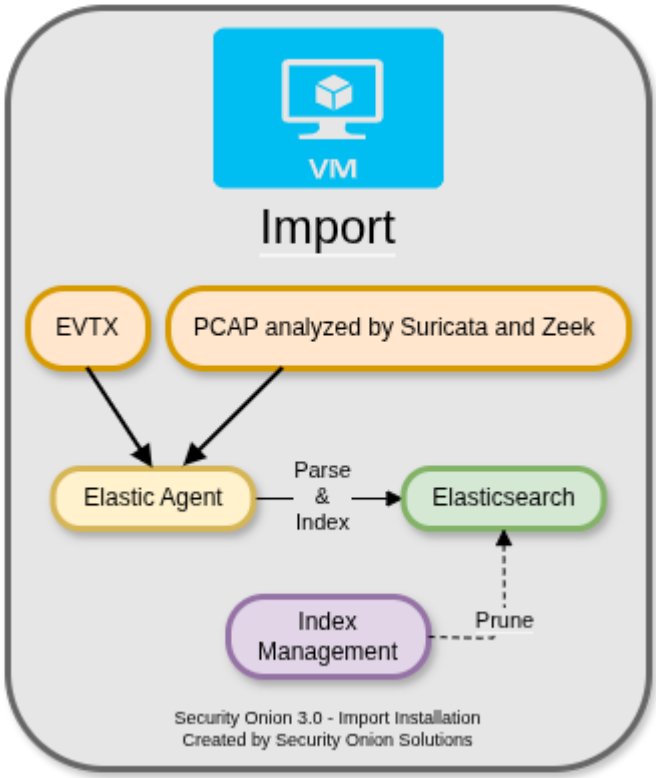

5.4.2 Evaluation

The next architecture is `Evaluation`. It's a little more complicated than `Import` because it has a network interface dedicated to sniffing live traffic from a TAP or SPAN port. Processes monitor the traffic on that sniffing interface and generate logs. Elastic Agent collects those logs and sends them directly to Elasticsearch where they are parsed and indexed. Evaluation mode is designed for a quick installation to temporarily test out Security Onion. It is **not** designed for production usage at all and it does not support adding Elastic agents or additional Security Onion nodes.

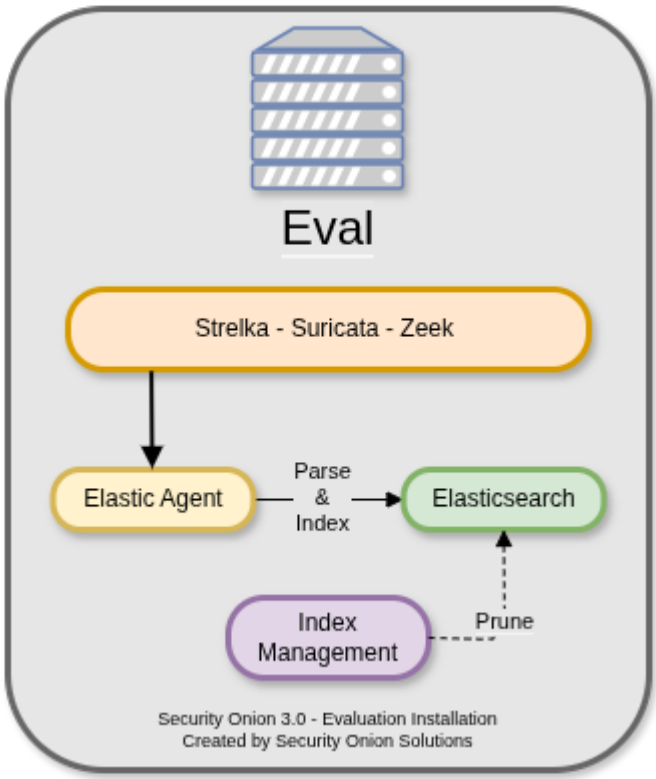

5.4.3 Standalone

`Standalone` is similar to `Evaluation` in that all components run on one box. However, instead of Elastic Agent sending logs directly to Elasticsearch, it sends them to Logstash, which sends them to Redis for queuing. A second Logstash pipeline pulls the logs out of Redis and sends them to Elasticsearch, where they are parsed and indexed.

This type of deployment is typically used for testing, labs, POCs, or **very** low-throughput environments. It's not as scalable as a distributed deployment.

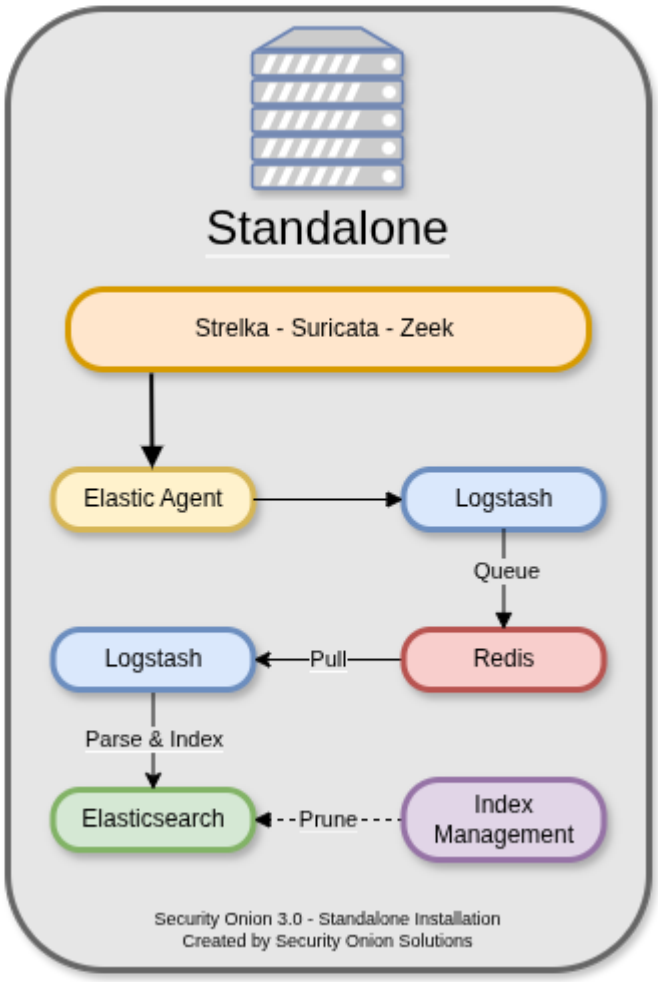

5.4.4 Security Onion Desktop

The installer includes a Security Onion Desktop option that builds a simple desktop environment. This environment includes a web browser which allows you to log into an existing Security Onion deployment. It also includes some analyst utilities like Wireshark and NetworkMiner.

For more information, please see the Security Onion Desktop section.

5.4.5 Distributed

A standard distributed deployment includes a **manager node**, one or more **sensor nodes** running network sensor components, and one or more **search nodes** running Elastic search components. This architecture may cost more upfront, but it provides for greater scalability and performance, as you can simply add more nodes to handle more traffic or log sources.

- Recommended deployment type
- Consists of a manager node, one or more sensor nodes, and one or more search nodes

> **Note**
>
> If you install a dedicated manager node, you must also deploy one or more search nodes. Otherwise, all logs will queue on the manager and have no place to be stored. If you are limited on the number of nodes you can deploy, you can install a **manager search** node so that your manager node can act as a search node and store those logs. However, please keep in mind that overall performance and scalability of a **manager search** node will be lower compared to our recommended architecture of dedicated manager node and separate search nodes.

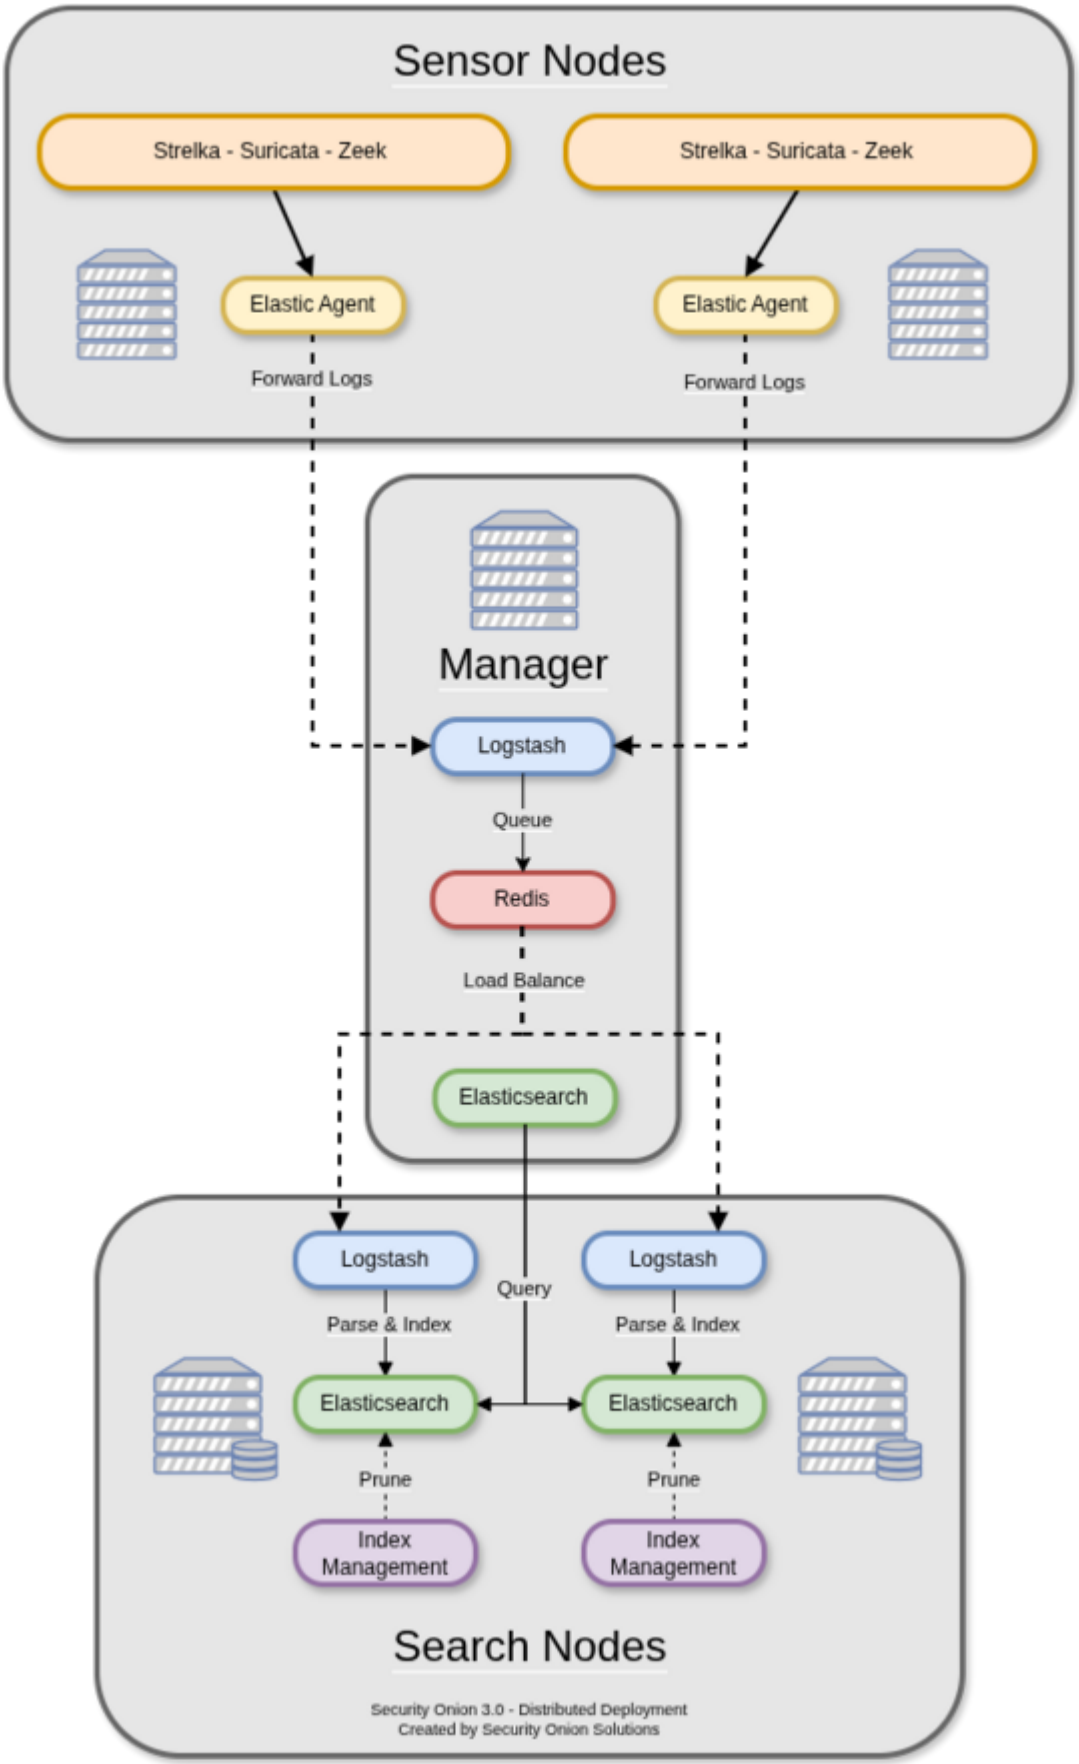
Sensor Nodes
Strelka - Suricata - Zeek
Strelka - Suricata - Zeek
Elastic Agent
Elastic Agent
Forward Logs
Forward Logs
Manager
Logstash
Queue
Redis
Load Balance
Elasticsearch
Logstash
Logstash
Query
Parse & Index
Parse & Index
Elasticsearch
Elasticsearch
Prune
Prune
Index Management
Index Management
Search Nodes
Security Onion 3.0 - Distributed Deployment
Created by Security Onion Solutions

5.4.6 Node Types

Management

The `manager node` runs Security Onion Console and Kibana. It has its own local instance of Elasticsearch, but that's mainly used for managing the Elasticsearch cluster once search nodes join the cluster. An analyst connects to the manager node from a client workstation (perhaps Security Onion Desktop) to execute queries and retrieve data.

Please note that a dedicated manager node requires separate search nodes. Also note that a dedicated manager node has to start off with the `data` node role. When you later join a separate search node, then you may want to migrate the data from the manager to the search node and then remove the `data` node role from the manager. For more information, please see the Elasticsearch section.

The manager node runs the following components:

- Security Onion Console
- Elasticsearch
- Logstash
- Kibana
- ElastAlert
- Redis

Search Node

Search nodes pull logs from the Redis queue on the manager node and then parse and index those logs. When a user queries the manager node, the manager node then queries the search nodes, and they return search results.

Search Nodes run the following components:

- Elasticsearch
- Logstash

Manager Search

A `manager search` node is both a manager node and a search node at the same time. Since it is parsing, indexing, and searching data, it has higher hardware requirements than a normal manager node.

A manager search node runs the following components:

- Security Onion Console
- Elasticsearch
- Logstash
- Kibana
- ElastAlert
- Redis

Sensor Node

A `sensor node` forwards alerts and logs from Suricata and Zeek via Elastic Agent to Logstash on the manager node, where they are stored in Elasticsearch on the manager node or a search node (if the manager node has been configured to use a search node). Full packet capture recorded by Suricata remains on the sensor node itself.

Sensor nodes run the following components:

- Zeek
- Suricata

Elastic Fleet Standalone Node

An Elastic Fleet Standalone Node is ideal when there is a large number of Elastic endpoints deployed. It reduces the amount of overhead on the Manager node by transferring the workload associated with managing endpoints to a dedicated system. It is also useful for off-network Elastic Agent endpoints that do not have remote access to the Manager node as it can be deployed to the DMZ and TCP/8220 (Elastic Agent Management network traffic) and TCP/5055 (Elastic Agent log shipping) made accessible to your off-network endpoints.

Receiver Node

Receiver nodes were designed with 2 purposes in mind:

- reduce the load on the manager
- offer pipeline redundancy

Each receiver node runs Logstash and Redis and allows for events to continue to be processed by search nodes in the event the manager node is offline. When a receiver node joins the grid, Elastic Agent on all nodes adds this new address as a load balanced Logstash output. The search nodes add this new node as another Logstash input. Receiver nodes are "active-active" and you can add as many as you want (within reason) and events will be balanced among them.

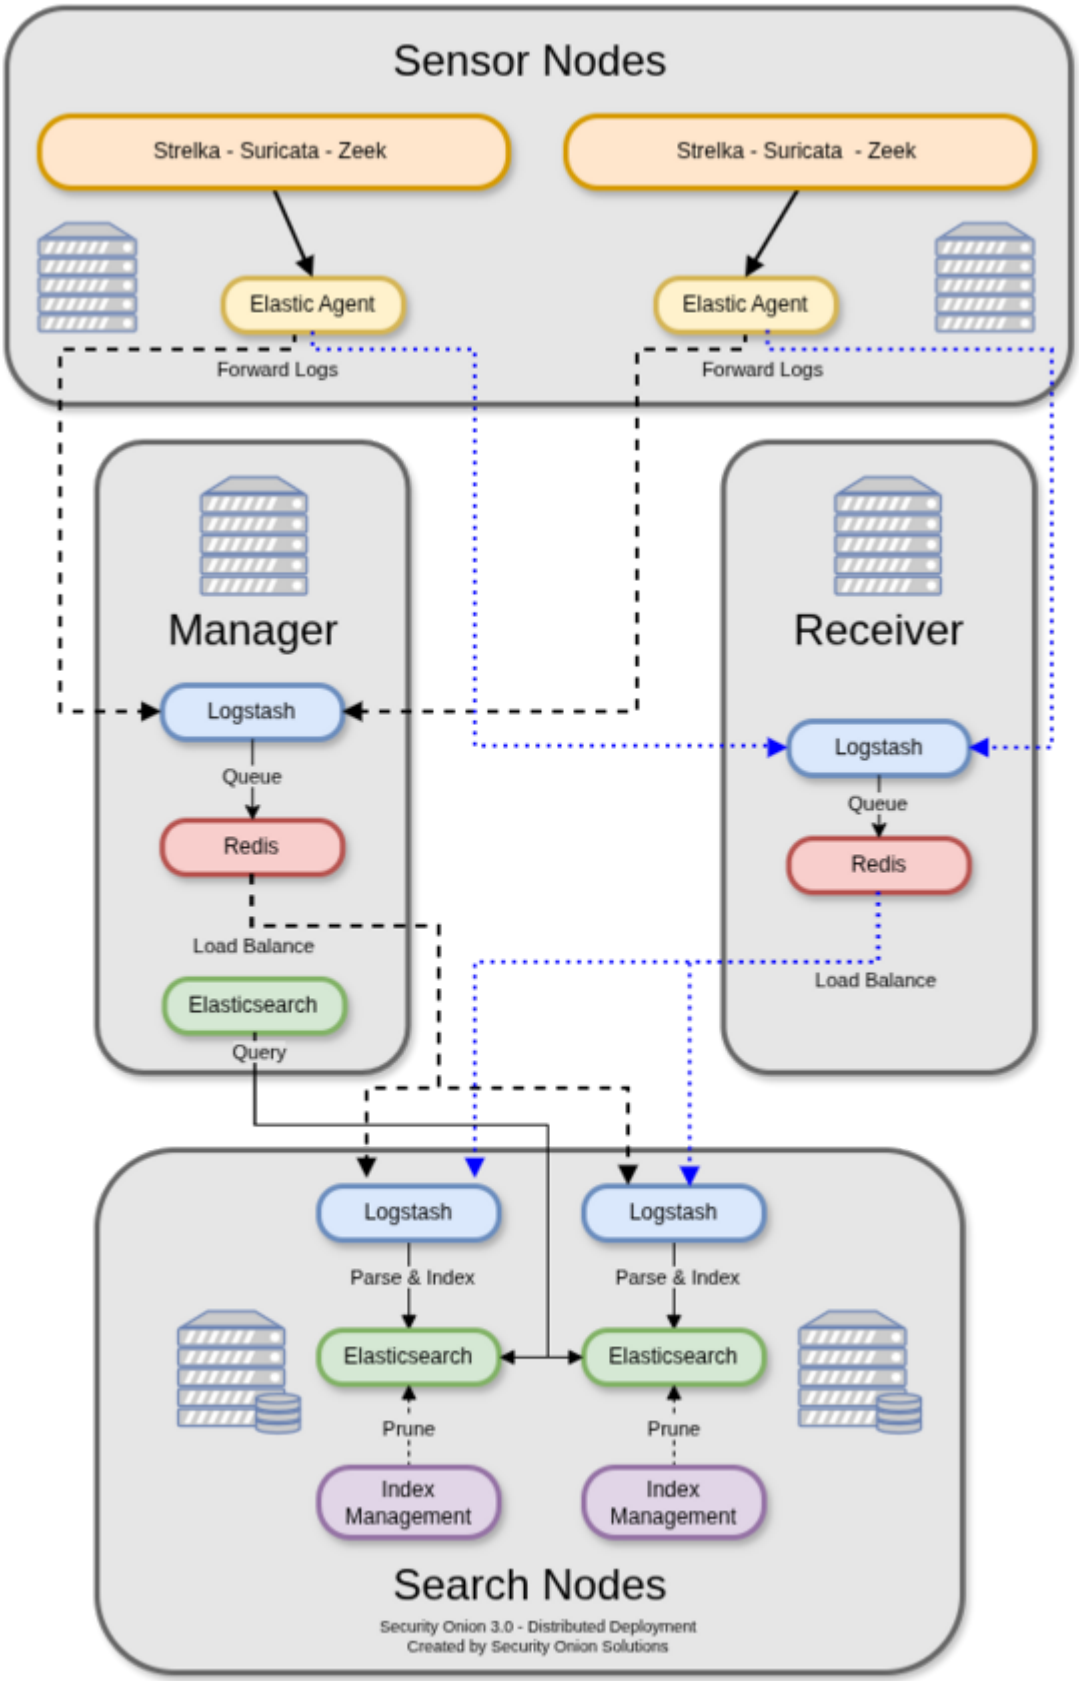
Sensor Nodes
Strelka - Suricata - Zeek
Strelka - Suricata - Zeek
Elastic Agent
Elastic Agent
Forward Logs
Forward Logs
Manager
Receiver
Logstash
Queue
Redis
Load Balance
Elasticsearch
Query
Logstash
Queue
Redis
Load Balance
Logstash
Logstash
Parse & Index
Parse & Index
Elasticsearch
Elasticsearch
Prune
Prune
Index Management
Index Management
Search Nodes
Security Onion 3.0 - Distributed Deployment
Created by Security Onion Solutions

If you don't have any receiver nodes and the manager goes down, the search nodes do not index anything because they cannot connect to Redis. The agents cannot connect to Logstash so the pipeline starts backing up on the agents.

In this same scenario with a receiver node, the agents would not be able to connect to Logstash on the manager and so they would try to connect to the receiver node. Once connected, they would send their logs to the receiver.

Search nodes connect to both the manager and receiver nodes and pull events from the Redis queue. If the manager goes down, search nodes will keep pulling the log events from the queue on the receiver node. This allows for scaling of the pipeline. More receivers + more search nodes = more event ingestion volume.

If you have a manager or managersearch that is under heavy load due to handling a high volume of events, then system resources can be freed by directing the Elastic Agent to only output events to the receiver node(s) in the environment. Once all configurable and advanced settings are enabled, this feature can be set in SOC Configuration UI under `elasticfleet > enable_manager_output`. Setting this to `False` will prevent the Elastic Agent from sending events to the manager, managersearch, or standalone nodes.

Receiver nodes need to be close to the search nodes because when you add a new receiver node to the grid, the search nodes add the Redis service as an input in their configs automatically. If you were to place a receiver node at a remote site, then ALL of your search nodes would be trying to access that Redis queue remotely. You do not save any bandwidth by placing a receiver node at a remote site.

There are a couple of things to be aware of regarding receiver nodes and Elastic Agents. The first is Fleet which handles things like updating the agents and scheduling searches. The other is the Elastic Agent log output, which in this case is Logstash running on the manager or receiver node. Due to limitations in Elastic licensing we can only have a single output policy. That means that when you add a receiver or a fleet node it gets added to a list that is distributed to the agents. The agents go down that list and stop after a successful connection. The only way to direct agents to specific receivers is to use firewall rules to block agents to certain receivers. Again keep in mind that there is no bandwidth savings here because the search nodes still need to empty the Redis queue on the receiver nodes.

Intrusion Detection Honeypot (IDH) Node

The IDH node mimics common services such as HTTP, FTP, and SSH. Any interaction with these fake services will automatically result in an alert.

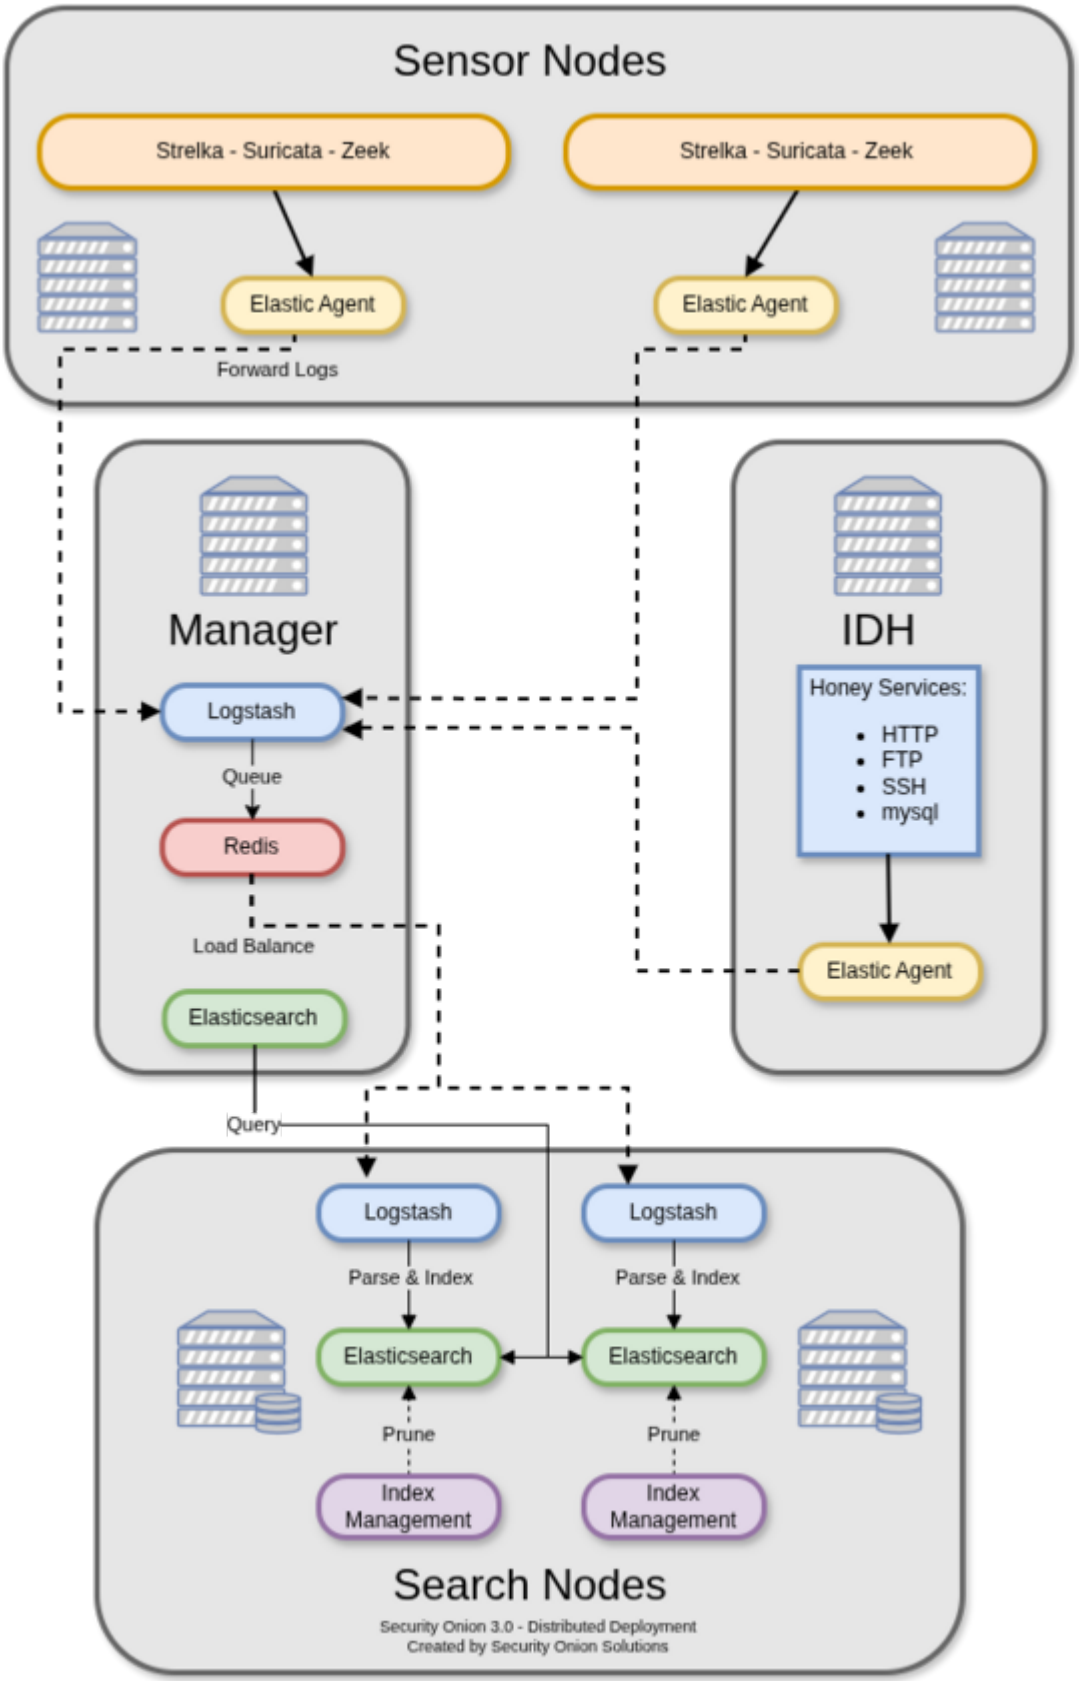
Sensor Nodes
Strelka - Suricata - Zeek
Strelka - Suricata - Zeek
Elastic Agent
Elastic Agent
Forward Logs
Manager
Logstash
Queue
Redis
Load Balance
Elasticsearch
IDH
Honey Services:
HTTP
FTP
SSH
mysql
Elastic Agent
Query
Logstash
Logstash
Parse & Index
Parse & Index
Elasticsearch
Elasticsearch
Prune
Prune
Index Management
Index Management
Search Nodes
Security Onion 3.0 - Distributed Deployment
Created by Security Onion Solutions

Heavy Node

There is also an option to have a **manager node** and one or more **heavy nodes**.

> **Warning**
>
> Heavy nodes are NOT recommended for most users due to performance reasons, and should only be used for testing purposes or in low-throughput environments.

- Recommended only if a standard distributed deployment is not possible
- Consists of a manager node and one or more heavy nodes
- Each heavy node is an independent Elastic cluster that is queried from the manager via cross-cluster search

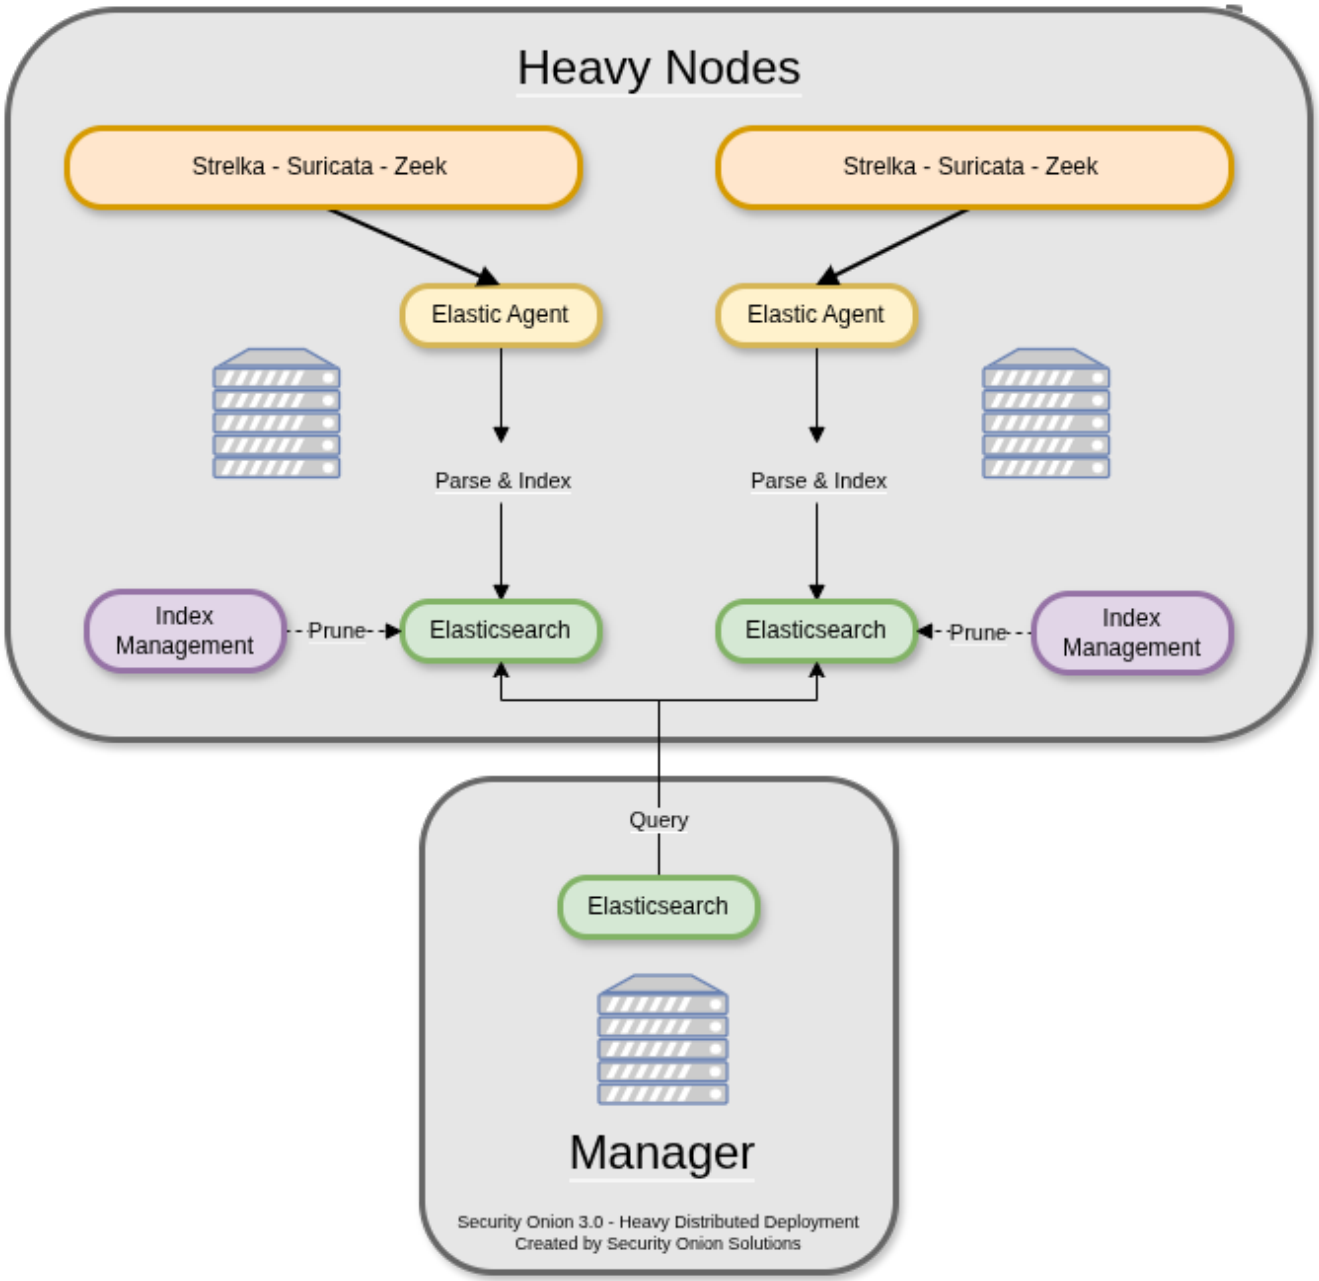

> **Note**
>
> Heavy nodes do not consume from the Redis queue on the manager. This means that if you just have a manager and heavy nodes, then the Redis queue on the manager will grow and never be drained. To avoid this, you have two options. If you are starting a new deployment, you can make your `manager` a `manager search` so that it will drain its own Redis queue. Alternatively, if you have an existing deployment with a `manager` and want to avoid rebuilding, then you can add a separate search node (NOT heavy node) to consume from the Redis queue on the manager.

Heavy nodes perform sensor duties and store their own logs in their own local Elasticsearch instance. This results in higher hardware requirements and lower performance. Heavy nodes do NOT pull logs from the Redis queue on the manager like search nodes do.

Heavy Nodes run the following components:

- Elasticsearch
- Zeek
- Suricata

There are two instances of Elastic Agent that run on a Heavy Node:

Instance 1 - Not connected to Fleet (runs standalone), runs in a container, picks up /nsm/ logs and other local logs (SOC) and sends them to the local Heavy Node ES cluster.

Instance 2 - Connected to Grid Fleet Server, runs directly on the Heavy Node. Not currently picking up any logs, but has the Osquery integration installed.

5.5 Hardware

The Architecture section should have helped you determine how many machines you will need for your deployment. This section will help you determine what kind of hardware specs each of those machines will need.

5.5.1 CPU Architecture

Security Onion only supports x86-64 architecture (standard Intel or AMD 64-bit processors).

> **Warning**
>
> We do not support ARM or any other non-x86-64 processors!

5.5.2 Minimum Specs

Node Type	CPU cores	RAM	Storage	NICs
Import	2	4GB	50GB	1
Eval	4	8GB	200GB	2
Standalone	4	24GB	200GB	2
Manager	4	16GB	200GB	1
ManagerSearch	8	16GB	200GB	1
Search node	4	16GB	200GB	1
Sensor	4	12GB	200GB	2
Heavy node	4	16GB	200GB	2
IDH node	2	1GB	12GB	1
Fleet node	4	4GB	200GB	1
Receiver node	2	8GB	200GB	1

> **Warning**
>
> Please note these are the absolute bare minimum requirements. Your requirements may increase drastically as you enable more services, monitor more traffic, and consume more logs. For more information, please see the detailed sections below.

5.5.3 Import

An Import installation runs the minimal processes required to import PCAP or EVTX files and view the results. As such, it has the lowest hardware requirements as shown in the table above. You can read more about Import in the First Time Users section.

5.5.4 Eval

An Eval installation runs the minimal processes required for a single machine to sniff live network traffic from a TAP or SPAN port and view the results. Therefore, its hardware requirements are higher than Import as shown in the table above. Eval is designed for temporary installations or homelab installations on a budget. Unlike a full Standalone installation, Evaluation is NOT designed for production usage. In order to minimize RAM usage, Eval does not run Logstash or Redis at all.

5.5.5 Production Deployments

For best results, we recommend purchasing new hardware that meets the hardware requirements detailed below.

> **Tip**
>
> If you're planning to purchase new hardware, please consider official Security Onion appliances from Security Onion Solutions at https://securityonionsolutions.com. Our custom appliances have already been designed for certain roles and traffic levels and have Security Onion pre-installed. Purchasing from Security Onion Solutions will save you time and effort and help to support development of Security Onion as a free and open platform!

5.5.6 Storage

We only support local storage. Remote storage like SAN/iSCSI/FibreChannel/NFS increases complexity and points of failure, and has serious performance implications. You may be able to make remote storage work, but we do not provide any support for it. By using local storage, you keep everything self-contained and you don't have to worry about competing for resources. Local storage is usually the most cost efficient solution as well.

5.5.7 NIC

You'll need one wired network interface dedicated to management (preferably connected to a dedicated management network). We recommend using static IP addresses where possible. To keep things simple, we recommend only one network interface with an IP address. If you have multiple network interfaces with IP addresses, then Security Onion traffic will default to the interface with the default gateway.

If you plan to sniff network traffic from a TAP or SPAN port, then you will need one or more interfaces dedicated to sniffing (no IP address). The installer will automatically disable NIC offloading functions such as `tso`, `gso`, and `gro` on sniffing interfaces to ensure that Suricata and Zeek get an accurate view of the traffic.

Make sure you get good quality network cards, especially for sniffing. Most users report good experiences with Intel cards.

Security Onion is designed to use wired interfaces. You may be able to make wireless interfaces work, but we don't recommend or support it.

5.5.8 UPS

As with most computer systems, you'll want to avoid power outages or other ungraceful shutdowns. Please consider a UPS (Uninterruptible Power Supply) for production deployments.

5.5.9 Elastic Stack

We recommend placing all Elastic storage (/nsm/elasticsearch) on SSD or fast spinning disk in a RAID 10 configuration.

Please see the Architecture section for detailed deployment scenarios.

5.5.10 Standalone Deployments

In a standalone deployment, the manager components and the sensor components all run on a single box so your hardware requirements will reflect that. You'll need at minimum 24GB RAM, 4 CPU cores, and 200GB storage. At the bare minimum of 24GB RAM, you may need swap space to avoid issues. Consider 32GB of RAM or more if you plan on monitoring even a small amount of network traffic. More network traffic means higher hardware requirements.

This deployment type is recommended for evaluation purposes, POCs (proof-of-concept) and small to medium size single sensor deployments. Although you can deploy Security Onion in this manner, it is recommended that you separate the backend components and sensor components.

- CPU: Used to parse incoming events, index incoming events, search metatadata, capture PCAP, analyze packets, and run the frontend components. As data and event consumption increases, a greater amount of CPU will be required.
- RAM: Used for Logstash, Elasticsearch, disk cache for Lucene, Suricata, Zeek, etc. The amount of available RAM will directly impact search speeds and reliability, as well as ability to process and capture traffic.
- Disk: Used for storage of indexed metadata. A larger amount of storage allows for a longer retention period. It is typically recommended to retain no more than 30 days of hot Elasticsearch indices.

Please refer to the Architecture section for detailed deployment scenarios.

5.5.11 Manager node with local log storage and search

In an enterprise distributed deployment, a manager node will store logs from itself and sensor nodes. It can also act as a syslog destination for other log sources to be indexed into Elasticsearch. An enterprise manager node should have 8 CPU cores at a minimum, 16-128GB RAM, and enough disk space (multiple terabytes recommended) to meet your retention requirements.

- CPU: Used to parse incoming events, index incoming events, and search metadata. As consumption of data and events increases, more CPU will be required.
- RAM: Used for Logstash, Elasticsearch, and disk cache for Lucene. The amount of available RAM will directly impact search speeds and reliability.
- Disk: Used for storage of indexed metadata. A larger amount of storage allows for a longer retention period. It is typically recommended to retain no more than 30 days of hot Elasticsearch indices.

Please refer to the Architecture section for detailed deployment scenarios.

5.5.12 Manager node with separate search nodes

This deployment type utilizes search nodes to parse and index events. As a result, the hardware requirements of the manager node are reduced. An enterprise manager node should have at least 4-8 CPU cores, 16GB RAM, and 200GB to 1TB of disk space. Many folks choose to host their manager node in their VM farm since it has lower hardware requirements than sensors but needs higher reliability and availability.

- CPU: Used to receive incoming events and place them into Redis. Used to run all the front end web components and aggregate search results from the search nodes.
- RAM: Used for Logstash and Redis.
- Disk: Used for general OS purposes.

Please refer to the Architecture section for detailed deployment scenarios.

5.5.13 Search Node

Search nodes increase search and retention capacity with regard to Elasticsearch. These nodes parse and index events, and provide the ability to scale horizontally as overall data intake increases. Search nodes should have at least 4-8 CPU cores, 16-64GB RAM, and 200GB of disk space or more depending on your logging requirements.

- CPU: Used to parse incoming events and index incoming events. As consumption of data and events increases, more CPU will be required.
- RAM: Used for Logstash, Elasticsearch, and disk cache for Lucene. The amount of available RAM will directly impact search speeds and reliability.
- Disk: Used for storage of indexed metadata. A larger amount of storage allows for a longer retention period. It is typically recommended to retain no more than 30 days of hot Elasticsearch indices.

Please refer to the Architecture section for detailed deployment scenarios.

5.5.14 Sensor Node

A sensor node runs sensor components only, and forwards metadata to the manager node. All PCAP stays local to the sensor, and is accessed through use of an agent.

- CPU: Used for analyzing and storing network traffic. As monitored bandwidth increases, a greater amount of CPU will be required. See below.
- RAM: Used for write cache and processing traffic.
- Disk: Used for storage of PCAP and metadata. A larger amount of storage allows for a longer retention period.

Please refer to the Architecture section for detailed deployment scenarios.

5.5.15 Heavy Node (Sensor with Elasticsearch components)

A heavy node runs all the sensor components AND Elastic components locally. This dramatically increases the hardware requirements. In this case, all indexed metadata and PCAP are retained locally. When a search is performed through Kibana, the manager node queries this node's Elasticsearch instance. You'll need at minimum 16GB RAM, 4 CPU cores, and 200GB storage. At the bare minimum of 16GB RAM, you will need swap space to avoid issues. We recommend a minimum of 24GB of RAM if you plan on monitoring traffic. The more traffic you plan on monitoring this RAM requirement will also increase.

- CPU: Used to parse incoming events, index incoming events, and search metadata. As monitored bandwidth (and the amount of overall data/events) increases, a greater amount of CPU will be required.
- RAM: Used for Logstash, Elasticsearch, and disk cache for Lucene. The amount of available RAM will directly impact search speeds and reliability.
- Disk: Used for storage of indexed metadata. A larger amount of storage allows for a longer retention period. It is typically recommended to retain no more than 30 days of hot Elasticsearch indices.

Please refer to the Architecture section for detailed deployment scenarios.

5.5.16 Receiver Node

Since receiver nodes only run Logstash and Redis, they don't require much CPU or disk space. However, more RAM means you can set a larger queue size for Redis.

5.5.17 Intrusion Detection Honeypot (IDH) Node

For an IDH node, the overall system requirements are low: 1GB RAM, 2 CPU cores, 1 NIC, and 100GB disk space.

5.5.18 Sensor Hardware Considerations

The following hardware considerations apply to sensors. If you are using a heavy node or standalone deployment type, please note that it will dramatically increase CPU/RAM/Storage requirements.

Virtualization

We recommend dedicated physical hardware (especially if you're monitoring lots of traffic) to avoid competing for resources. Sensors can be virtualized, but you'll have to ensure that they are allocated sufficient resources.

CPU

Suricata and Zeek are very CPU intensive. The more traffic you are monitoring, the more CPU cores you'll need. A very rough ballpark estimate would be 200Mbps per Suricata worker or Zeek worker. So if you have a fully saturated 1Gbps link and are running Suricata for NIDS alerts and Zeek for metadata, then you'll want at least 5 Suricata workers and 5 Zeek workers. This means you'll need at least 10 CPU cores for Suricata and Zeek with additional CPU cores for Full Packet Capture and/or other services. If you are monitoring a high amount of traffic and/or have a small number of CPU cores, you might consider using Suricata for both alerts and metadata. This eliminates the need for Zeek and allows for more efficient CPU usage.

RAM

RAM usage is highly dependent on several variables:

- the services that you enable
- the **kinds** of traffic you're monitoring
- the **actual amount of traffic** you're monitoring (example: you may be monitoring a 1Gbps link but it's only using 200Mbps most of the time)
- the amount of packet loss that is "acceptable" to your organization

For best performance, over provision RAM so that you can fully disable swap.

The following RAM estimates are a rough guideline and assume that you're going to be running Suricata (including full packet capture) and Zeek and want to minimize/eliminate packet loss. Your mileage may vary!

- If you just want to quickly evaluate Security Onion in a VM, the bare minimum amount of RAM needed is 12GB. More is obviously better!
- If you're deploying Security Onion in production on a small network (100Mbps or less), you should plan on 16GB RAM or more. Again, more is obviously better!
- If you're deploying Security Onion in production to a medium network (100Mbps - 1000Mbps), you should plan on 16GB - 128GB RAM or more.
- If you're deploying Security Onion in production to a large network (1000Mbps - 10Gbps), you should plan on 128GB - 256GB RAM or more.
- If you're buying a new server, go ahead and max out the RAM (it's cheap!). As always, more is obviously better!

Storage

Sensors that have full packet capture enabled need LOTS of storage. For example, suppose you are monitoring a link that averages 50Mbps, here are some quick calculations: 50Mb/s = 6.25 MB/s = 375 MB/minute = 22,500 MB/hour = 540,000 MB/day. So you're going to need about 540GB for one day's worth of pcaps (multiply this by the number of days of PCAP you want to keep). The more disk space you have, the more PCAP retention you'll have for doing investigations after the fact. Disk is cheap, get all you can!

Packets

You'll need some way of getting packets into your sensor interface(s). If you're just evaluating Security Onion, you can replay pcaps. For a production deployment, you'll need a SPAN port on an existing switch or a dedicated TAP. We recommend using a dedicated TAP where possible. If collecting traffic near a NAT boundary, make sure you collect from inside the NAT boundary so that you see the true internal IP addresses.

Inexpensive TAP/SPAN options (listed alphabetically):

- Dualcomm
- Midbit SharkTap
- Netgear GS105Ev2

Enterprise TAP options (listed alphabetically):

- APCON
- Arista
- cPacket
- Garland
- Gigamon
- KeySight / Ixia / Net Optics
- Profitap

Further Reading

> **Note**
>
> For large networks and/or deployments, please also see https://github.com/pevma/SEPTun.

5.6 Download

Before downloading, we highly recommend that you review the Release Notes section so that you are aware of all recent changes!

> **Warning**
>
> **ALWAYS verify the checksum of the ISO image before booting!** This ensures that the ISO image hasn't been tampered with or corrupted during download. If it fails to verify, try downloading again. If it still fails to verify, try downloading from another computer or another network.
>
> Download and verify our ISO image as shown at https://github.com/Security-Onion-Solutions/securityonion/blob/3/main/DOWNLOAD_AND_VERIFY_ISO.md.

> **Warning**
>
> Antivirus software may alert on the ISO image but any alerts are most likely false positives. If you look at the antivirus scan details, it will most likely tell you that it alerted on a file in `SecurityOnion\agrules\`. These are rules that look for malicious activity but the rules themselves are not actually malicious.

> **Note**
>
> If you're going to create a bootable USB from the ISO image, there are many ways to do that. One popular choice that seems to work well for many folks is Balena Etcher which can be downloaded at https://www.balena.io/etcher/.

5.7 VMware

5.7.1 Overview

In this section, we'll cover creating a virtual machine (VM) for our ISO image in VMware Workstation Pro and VMware Fusion. These steps should be fairly similar for most VMware installations.

> **Note**
>
> If you want to sniff live traffic, then you will need a second network interface dedicated to sniffing. You will need to set this sniffing interface to sniff from whatever network you want to monitor. With the sniffing interface in `bridged` mode, you should be able to see all traffic to and from the host machine's physical NIC. If you would like to see **ALL** the traffic on your network, you will need a method of forwarding that traffic to the interface to which the virtual adapter is bridged. This can be achieved with a TAP or SPAN port. If you want to sniff traffic from other VMs, then the virtual sniffing interface needs to be set to the same virtual network that those VMs are set to (this may be `NAT` or `bridged` depending on how they are configured).

5.7.2 Workstation Pro

VMware Workstation is available for many different host operating systems, including Windows and several popular Linux distros. Follow the steps below to create a VM in VMware Workstation Pro for our ISO image:

- From the VMware main window, select File >> New Virtual Machine.
- Select Typical installation >> Click `Next`.
- Installer disc image file >> SO ISO file path >> Click `Next`.
- Choose Linux, then choose the closest Linux distribution and click `Next`.
- Specify virtual machine name and click `Next`.
- Specify disk size (minimum 200GB), store as single file, click `Next`.
- Customize hardware and increase Memory and Processors based on the Hardware section.
- Network Adapter (`NAT` or `Bridged` -- if you want to be able to access your Security Onion machine from other devices in the network then choose Bridged, otherwise choose NAT to leave it behind the host). This will be the management interface.
- Add >> Network Adapter (`NAT` or `Bridged`). This will be the sniffing (monitor) interface.
- Click `Close`.
- Click `Finish`.
- Power on the virtual machine and then follow the installation steps for your desired installation type in the Installation section.

5.7.3 Fusion

VMware Fusion is available for Mac OS. For more information about VMware Fusion, please see https://www.vmware.com/products/fusion.html.

Follow the steps below to create a VM in VMware Fusion for our ISO image:

- From the VMware Fusion main window, click `File` and then click `New`.
- `Select the Installation Method` appears. Click `Install from disc or image` and click `Continue`.
- `Create a New Virtual Machine` appears. Click `Use another disc or disc image...`, select our ISO image, click `Open`, then click `Continue`.
- `Choose Operating System` appears. Click `Linux`, choose the closest Linux distribution, then click `Continue`.
- `Choose Firmware Type` appears. Click `Legacy BIOS` and then click `Continue`.
- `Finish` screen appears. Click the `Customize Settings` button.
- `Save As` screen appears. Give the VM a name and click the `Save` button.
- `Settings` window appears. Click `Processors & Memory`.
- `Processors & Memory` screen appears. Increase processors and memory based on the Hardware section. Click the `Add Device...` button.
- `Add Device` screen appears. Click `Network Adapter` and click the `Add...` button.
- `Network Adapter 2` screen appears. This will be the sniffing (monitor) interface. Select your desired network adapter configuration. Click the `Show All` button.
- `Settings` screen appears. Click `Hard Disk (SCSI)`.
- `Hard Disk (SCSI)` screen appears. Increase the disk size to at least `200GB` depending on your use case. Click the `Apply` button.
- Close the `Settings` window.
- At the window for your new VM, click the `Play` button to power on the virtual machine.
- Follow the installation steps for your desired installation type in the Installation section.

5.7.4 ESXi

If you're using VMware ESXi, then you're likely familiar with VM creation and installation and so we won't detail that here. There are a few things specific to ESXi that you might want to be aware of:

- You may need to set your monitoring interface in the vSwitch to VLAN ID 4095 to allow all traffic through. You can read more about this at https://github.com/Security-Onion-Solutions/securityonion/discussions/7185.
- If you're trying to monitor multiple network interfaces, then you may need to enable the `Allow MAC Changes` option at both the vSwitch and Port Group levels. You can read more about this at https://github.com/Security-Onion-Solutions/securityonion/discussions/2676.
- If you happen to notice after rebooting that the Elastic Agent takes signifantly longer than 15 minutes to initialize, then you may need to enable the following option in ESXi: Settings > VM Options > VMWare Tools > Synchronise Guest Time. You can read more about this at https://github.com/Security-Onion-Solutions/securityonion/discussions/13285.

5.7.5 VMware Tools

If using a graphical desktop, you may want to install `open-vm-tools-desktop` to enable more screen resolution options and other features. For example, using our ISO image or standard Oracle Linux 9:

```
sudo dnf install open-vm-tools-desktop
```

5.8 VirtualBox

In this section, we'll cover installing Security Onion on VirtualBox. You can download a copy of VirtualBox for Windows, Mac OS X, or Linux at https://www.virtualbox.org.

5.8.1 Creating VM

- Launch VirtualBox and click the `New` button.
- Provide a name for the virtual machine (`Security Onion` for example) and then select the ISO image. It should automatically set type to `Linux` and version to `Oracle Linux 9.x`. Click the checkbox for `Skip Unattended Installation` and then click the `Next` button.
- Specify RAM and Processors as needed per the Hardware section and then click the `Next` button.
- Specify virtual hard disk size as needed per the Hardware section and then click the `Next` button.
- Confirm options and then click the `Finish` button.
- Virtualbox should have automatically enabled a network adapter attached to the NAT network. Depending on what kind of installation you are doing, you may want to keep that as NAT or change to something else. If you want an additional network interface for sniffing from a TAP or SPAN port, then click the `Settings` button, click `Network`, and then go to `Adapter 2`. Enable the adapter, configure the network it should attach to, and then you will most likely want to go to Advanced and set Promiscuous Mode to either `Allow VMs` or `Allow All`. Click the `OK` button.
- Click the `Start` button to start the VM.
- Follow the installation steps for your desired installation type in the Installation section.

5.8.2 Guest Additions

If you want to install VirtualBox Guest Additions, please see https://www.virtualbox.org/manual/ch04.html.

5.9 Proxmox

Proxmox Virtual Environment is a virtualization platform similar to VMware or VirtualBox. You can read more about Proxmox VE at https://www.proxmox.com/en/proxmox-ve. If you would like to run Security Onion in a Proxmox VM, then here are some additional things you need to be aware of.

5.9.1 CPU

Proxmox defaults to a VM CPU which may not include all of the features of your host CPU. You may need to change this to `host` to pass through the host CPU type.

5.9.2 Display

If you plan to use NetworkMiner or other Mono-based applications in a Proxmox VM, then you may need to set the VM Display to `VMware compatible (vmware)`.

5.9.3 NIC

If you're going to install Security Onion in Proxmox and sniff live network traffic, you may need to do some additional configuration in Proxmox itself. You can either passthrough a physical NIC to the VM or you can use a virtual NIC.

Passthrough Physical NIC

The first option is to sniff traffic from a physical NIC that has been passed through to the VM. For more information about Proxmox passthrough, please see:

https://www.servethehome.com/how-to-pass-through-pcie-nics-with-proxmox-ve-on-intel-and-amd/

https://pve.proxmox.com/wiki/PCI_Passthrough

https://pve.proxmox.com/wiki/PCI(e)_Passthrough

Once the physical NIC is passed through to the Security Onion VM, then Security Onion should be able to correctly configure the NIC for sniffing.

Virtual NIC

The second option is to sniff traffic from a Proxmox virtual NIC. For more details, please see the discussion at https://github.com/Security-Onion-Solutions/securityonion/discussions/8245.

Keep in mind you may need to manually disable NIC offloading features on any Proxmox NIC used for sniffing (the physical interface and any related bridge interface). One way to do this is to add a post-up command to each sniffing interface in /etc/network/interfaces on the Proxmox host.

For example, if you have a Proxmox physical interface called `enp2s0` with a bridge interface called `vmbr1`, then you might log into Proxmox and edit /etc/network/interfaces by adding the following to the `enp2s0` section:

```
post-up for i in rx tx sg tso ufo gso gro lro; do ethtool -K enp2s0 $i off; done
```

and the following to the `vmbr1` section:

```
post-up for i in rx tx sg tso ufo gso gro lro; do ethtool -K vmbr1 $i off; done
```

For more information about NIC offloading, please see https://blog.securityonion.net/2011/10/when-is-full-packet-capture-not-full.html.

If you are running Proxmox 9 or higher, then you may also need to set `mtu 9000` for the Proxmox physical sniffing interface and its corresponding bridge interface.

Once your Security Onion VM is receiving traffic as expecte, if Grid reports Capture Loss but no Zeek Loss and you are confident that the loss is not occuring in your tap or span port, then it may be related to the Proxmox host physical interface. Certain NICs like the Intel X710 may have pre-set

channels and you can check with the `ethtool -l` command. For more information, please see the `ethtool` man page at https://man7.org/linux/man-pages/man8/ethtool.8.html.

If this is the case for your host physical interface, then you can add an additional post-up command to run `ethtool -L` with the `combined 1` option. For example, if you have a physical interface called `enp3s0f1np1`, then the corresponding section of /etc/network/interfaces would look like this:

```
post-up ethtool -L enp3s0f1np1 combined 1; for i in rx tx sg tso ufo gso gro lro; do ethtool -K enp3s0f1np1 $i off; done
```

5.10 Trouble Booting

If you have trouble booting the ISO image, here are some troubleshooting steps:

- If you're trying to create a bootable USB from an ISO image, try using Balena Etcher which can be downloaded at https://www.balena.io/etcher/.
- Certain display adapters may require the `nomodeset` option passed to the kernel (see https://unix.stackexchange.com/questions/353896/linux-install-goes-to-blank-screen).
- If you're still having problems with our 64-bit ISO image, try downloading the standard x86-64 ISO image for Oracle Linux 9. If it doesn't run, then you should double-check your 64-bit compatibility.

> **Tip**
>
> If all else fails but standard x86-64 Oracle Linux 9 installs normally, then you can install our components on top of it as described in the Network Installation section. However, please keep in mind that network installations are not supported.

5.11 Airgap

Security Onion is committed to allowing users to run a full install on networks that do not have Internet access. Our ISO image includes everything you need to run without Internet access. Make sure that you choose the airgap option during Setup.

If your network has Internet access but has overly restrictive proxies, firewalls, or other network devices that might prevent Security Onion from connecting to the sites shown in the Firewall section, then you may want to consider the airgap option as everything will install from the ISO image itself.

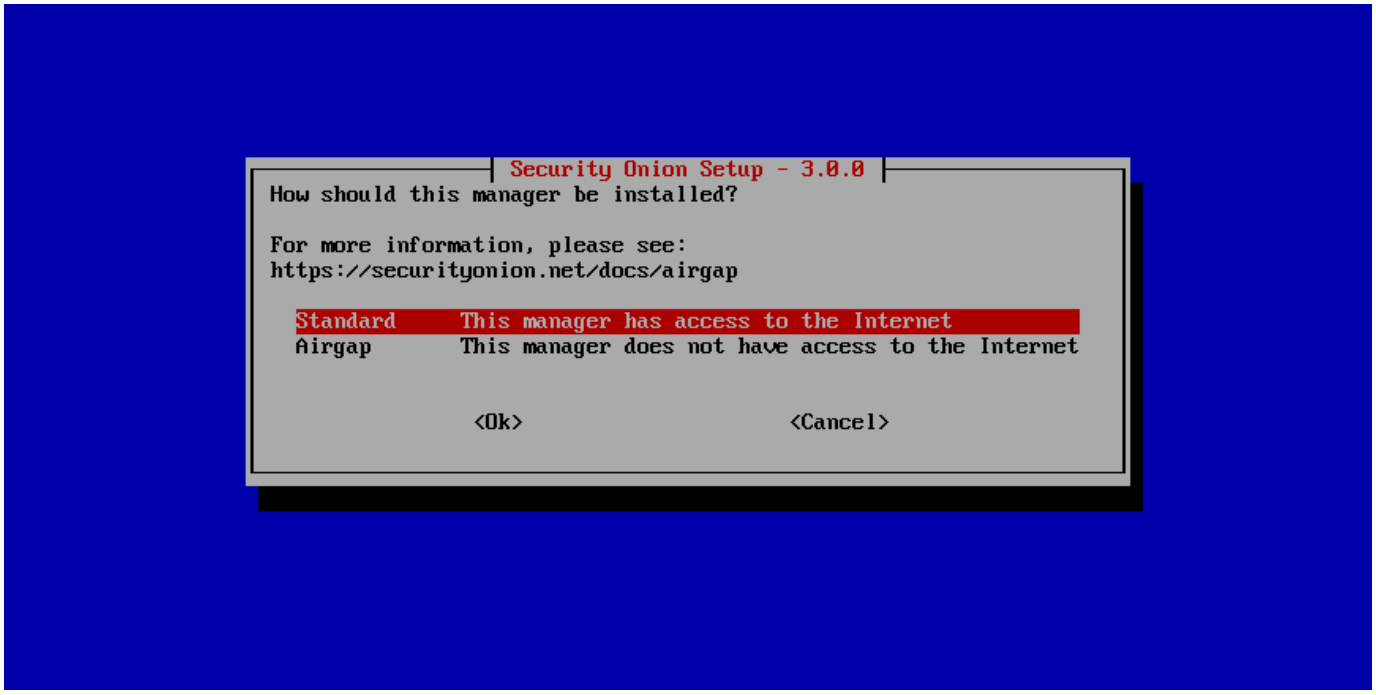

Airgap mode works as follows:

- During the install, all of the necessary RPM packages are copied from the ISO image to a new repo located in `/nsm/repo/`. All devices in the grid will now use this repo for updates to packages.
- NIDS rules for Suricata are copied to `/nsm/rules/suricata`.
- YARA rules for Strelka are copied to `/nsm/rules/yara`.
- Sigma rules for ElastAlert are copied to `/nsm/repo/rules/sigma`.
- When updating the system, soup will ask for the location of the latest ISO media and will then update using that media rather than pulling from the Internet.

5.11.1 Rule Updates

Our ISO image includes the latest version of various rulesets and will automatically install them when an airgap system is SOUP'ed via ISO:

- NIDS: Emerging Threats (ETOPEN). If you would like to switch to a different ruleset like Emerging Threats Pro (ETPRO), refer to our Ruleset config documentation NIDS
- YARA: Most recent rules from our repo
- Sigma: Most recent rule packages from the SigmaHQ repo

5.12 Installation

Warning

Please make sure that your hostname is correct during installation. Setup generates certificates based on the hostname and we do not support changing the hostname after Setup.

Note

If you want to deploy in the cloud using one of our official cloud images, you can skip to the Amazon Cloud, Azure Cloud, or Google Cloud sections.

Having downloaded our ISO image as shown in the Download section, it's now time to install!

- Review the Hardware and Release Notes sections.
- Download and verify our ISO image as shown in the Download section.
- Boot the ISO in a machine that meets the minimum hardware specs.
- Follow the prompts to complete the installation and reboot.
- You may need to eject the ISO image or change the boot order of the machine to boot from the newly installed OS.
- Login using the username and password you set in the installer.
- Security Onion Setup will automatically start. If for some reason you have to exit Setup and need to restart it, you can log out of your account and then log back in and it should automatically start. If that doesn't work, you can manually run it as follows:

```
sudo SecurityOnion/setup/so-setup iso
```

- Proceed to the Configuration section.

5.13 Amazon Cloud Image

If you would like to deploy Security Onion in Amazon Web Services (AWS), we have an Amazon Machine Image (AMI) that is already built for you: https://securityonion.net/aws/?ref=_ptnr_soc_docs_260311

Warning

Existing Security Onion cloud image installations should use the soup command to upgrade. If your grid is still running 2.4.x, use `soup` to upgrade to 2.4.210, and then use `soupto3` to proceed to Security Onion 3, after which continue using `soup` again. Attempting to switch to a newer Security Onion image from the cloud marketplace could cause loss of data and require full Grid re-installation; use the `soup` procedure to upgrade instead.

Note

This section does not cover network connectivity to the Security Onion node. This can be achieved through configuring an external IP for the node's management interface, or through the use of a VPN connection via OpenVPN. For more details about VPN connections, please see https://medium.com/@svfusion/setup-site-to-site-vpn-to-aws-with-pfsense-1cac16623bd6.

Note

This section does not cover how to set up a VPC in AWS. For more details about setting up a VPC, please see https://docs.aws.amazon.com/directoryservice/latest/admin-guide/gsg_create_vpc.html. Ensure that all Security Onion nodes can access the manager node over the necessary ports. This could require adding rules to your AWS security groups in order to satisfy the Security Onion Firewall Node Communication requirements.

5.13.1 Requirements

Before proceeding, determine the grid architecture desired. Choose from a single-node Grid versus a distributed, multi-node Grid. Additionally, determine if the lower latency of ephemeral instance storage is needed (typically when there is high-volume of traffic being monitored, which is most production scenarios), or if network-based storage, EBS, can be used for increased redundancy.

5.13.2 Single Node Grid

For simple, low-volume production monitoring, a single node Grid can be used. EBS must be used for Elasticsearch data storage if used for production purposes. Single node grids cannot use ephemeral instance storage without being at risk of data loss. However, for temporary evaluation installations, where there is little concern for data loss, ephemeral instance storage can be used.

Listed below are the minimum suggested single-node instance quantities, sizes, and storage requirements for either standalone or evaluation installations (choose one, not both). Note that when using virtual machines with the minimum RAM requirements you may need to enable memory swapping.

Standalone:

- Quantity: 1
- Type: t3a.2xlarge
- Storage: 256GB EBS (Optimized) gp3

Evaluation

- Quantity: 1
- Type: t3a.2xlarge
- Storage: 256GB EBS (Optimized) gp3
- Storage: 100GB Instance Storage (SSD/NVMe)

5.13.3 Distributed Grid

For high volume production monitoring, choose a multi-node Grid architecture. At least two search nodes must be used in this architecture. This is required due to the use of ephemeral instance storage for Elasticsearch data storage, where each of the search nodes retains a replica of another search node, for disaster recovery.

Listed below are the *minimum* suggested distributed Grid instance quantities, sizes, and storage requirements. Prefer increasing VM memory over enabling swap memory, for best performance. High volume networks will need more powerful VM types with more storage than those listed below.

VPN Node

- Quantity: 1
- Type: t3a.micro (Nitro eligible)
- Storage: 50GB EBS (Optimized) gp3

Manager

- Quantity: 1
- Type: m5a.2xlarge
- Storage: 300GB EBS (Optimized) gp3

Search Nodes

- Quantity: 2 or more
- Type: m5ad.2xlarge
- Storage: 200GB EBS (Optimized) gp3
- Storage: 150GB Instance Storage (SSD/NVMe)

Sensor monitoring the VPN ingress

- Quantity: 1
- Type: c5a.2xlarge
- Storage: 500GB EBS (Optimized) gp3

5.13.4 Create Monitoring Interface

To setup the Security Onion AMI and VPC mirror configuration, use the steps below.

5.13.5 Create a Security Group for Sniffing Interface

Security Groups act like a firewall for your Amazon EC2 instances controlling both inbound and outbound traffic. You will need to create a security group specifically for the interface that you will be using to sniff the traffic. This security group will need to be as open as possible to ensure all traffic destined to the sniffing interface will be allowed through. To create a security group, follow these steps:

- From the EC2 Dashboard Select: `Security Groups` under the Network & Security sections in the left window pane.
- Select: `Create Security Group`
- Provide a Security Group Name and Description.
- Select the appropriate VPC for the security group.
- With the inbound tab selected, select: `Add Rule`
- Add the appropriate inbound rules to ensure all desired traffic destined for the sniffing interface is allowed.
- Press the `Create security group` button.

5.13.6 Create Sniffing Interface

Prior to launching the Security Onion AMI you will need to create the interface that will be used to monitor your VPC. This interface will be attached to the Security Onion AMI as a secondary interface. To create a sniffing interface, follow these steps:

- From the EC2 Dashboard Select: `Network Interfaces` under the Network & Security section in the left window pane.
- Select: `Create Network Interface`
- Provide a description and choose the appropriate subnet you want to monitor.
- Select the security Group that you created for the sniffing interface.
- Select: `Create`

5.13.7 Create Security Onion Instances

5.13.8 Instance Creation

To configure a Security Onion instance (repeat for each node in a distributed Grid), follow these steps:

- From the EC2 dashboard select: `Launch Instance`
- Search the AWS Marketplace for `Security Onion` and make sure you get the latest version of the Security Onion official AMI.
- Choose the appropriate instance type based on the desired hardware requirements and select `Next: Configure Instance Details`. For assistance on determining resource requirements please review the AWS Requirements section above.
- From the subnet menu select the same subnet as the sniffing interface.
- Under the Network interfaces section configure the eth0 (management) interface.
- (Distributed "Sensor" node or Single-Node Grid only) Under the Network interfaces section select: `Add Device` to attach the previously created sniffing interface to the instance.
- (Distributed "Sensor" node or Single-Node Grid only) From the Network Interface menu for eth1 choose the sniffing interface you created for this instance. Please note if you have multiple interfaces listed you can verify the correct interface by navigating to the Network Interfaces section in the EC2 Dashboard.
- Select: `Next: Add Storage` and configure the volume settings.
- Select: `Next: Add Tags` and add any additional tags for the instance.
- Select: `Next: Configure Security Group` and add the appropriate inbound rules.
- Select: `Review and Launch`
- If prompted, select the appropriate SSH keypair that will be used to ssh into the Security Onion instance for administration
- The default username for the Security Onion AMI is: `onion`

5.13.9 Prepare Nodes with Ephemeral Storage

For distributed search nodes, or an evaluation node if using ephemeral storage, SSH into the node and cancel out of the setup. Prepare the ephemeral partition by executing the following command:

```
sudo so-prepare-fs
```

By default, this command expects the ephemeral device to be located at `/dev/nvme1n1` and will mount that device at `/nsm/elasticsearch`. If this fails run `lsblk` to determine which disk to use. To override either of those two defaults, specify them as arguments. For example:

```
sudo so-prepare-fs /dev/nvme3n0 /nsm/elasticsearch
```

Restart the Security Onion setup by running the following command:

```
cd /securityonion
sudo ./so-setup-network
```

5.13.10 Manager Setup

If this is an ephemeral evaluation node, ensure the node has been prepared as described in the preceding section.

After SSH'ing into the node, setup will begin automatically. Follow the prompts, selecting the appropriate install options. Most distributed installations will use the `hostname` or `other` web access method, due to the need for both cluster nodes inside the private network, and analyst users across the public Internet to reach the manager. This allows for custom DNS entries to define the correct IP (private vs public) depending on whether it's a cluster node or an analyst user. Users evaluating Security Onion for the first time should consider choosing the `other` option and specifying the node's public cloud IP.

AWS provides a built-in NTP server at IP `169.254.169.123`. This can be specified in the SOC Configuration screen after setup completes. By default the server will use the time servers at `ntp.org`.

For distributed manager nodes using ephemeral storage, go to SOC Configuration. Search for `number_of_replicas` and change to 1. This will double the storage cost but will ensure at least two VMs have the data, in case of an ephemeral disk loss.

Optionally, adjust ElastAlert indices so that they have a replica. This will cause them to turn yellow but that will be fixed when search nodes come online:

```
so-elasticsearch-query ElastAlert*/_settings -X PUT -d '{"index" : { "number_of_replicas" : 1 }}'
```

This is an optional step due to the ElastAlert indices being used primarily for short-term/recent alert history. In the event of a data loss when ElastAlert 2 restarts the indices will be regenerated.

5.13.11 Search Node Setup

Follow standard Security Onion search node installation, answering the setup prompts as applicable. If you are using ephemeral storage be sure to first prepare the instance as directed earlier in this section.

5.13.12 AWS Sensor Setup

SSH into the sensor node and run through setup to set this node up as a sensor. Choose `eth0` as the main interface and `eth1` as the monitoring interface.

5.13.13 Remote Sensor Setup

Setup the VPN (out of scope for this guide) and connect the sensor node to the VPN. During the Security Onion setup of the Sensor, when prompted to choose the management interface, select the VPN tunnel interface, typically `tun0`.

If connecting sensors through the VPN instance you will need to add the inside interface of your VPN concentrator to the `sensor` firewall hostgroup. For instance, assuming the following architecture:

```
SO Sensor        -> VPN Endpoint     -> Internet -> VPN Endpoint  -> SO Manager
Location: Remote    Location: Remote                Location: AWS    Location: AWS
192.168.33.13       192.168.33.10                   10.55.1.10       10.55.1.20
```

In order to add the Remote Network Sensor Node to the grid, you would have to add `10.55.1.10` to the `sensor` firewall hostgroup.

This change can be done in the SOC Configuration screen. Then, either wait up to 15 minutes for the scheduled configuration sync to run, or force a synchronization immediately via the SOC Configuration Options. Once the firewall hostgroup configuration has been synchronized your Manager will be ready for remote minions to start connecting.

5.13.14 AWS Traffic Mirroring

Traffic mirroring allows you to copy the traffic to/from an instance and send it to the sniffing interface of a network security monitoring sensor or a group of interfaces using a network load balancer. For more details about AWS Traffic Mirroring please see: https://docs.aws.amazon.com/vpc/latest/mirroring/what-is-traffic-mirroring.html

> **Tip**
>
> You can only mirror traffic from an EC2 instance that is powered by the AWS Nitro system. For a list of supported Nitro systems, please see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html#ec2-nitro-instances.

5.13.15 Create Mirror Target

A mirror target in AWS refers to the destination for the mirrored traffic. This can be a single interface or a group of interfaces using a network load balancer. To configure a mirror target, follow these steps:

- From the VPC dashboard select: `Mirror Targets` under the Traffic Mirroring section in the left window pane.
- Select: `Create traffic mirror target`
- Under the Choose target section select the appropriate target type and choose the sniffing interface connected to the Security Onion instance. For more details about traffic mirror targets please see: https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-targets.html
- Select: `Create`

5.13.16 Create Mirror Filter

A mirror filter allows you to define the traffic that is copied to in the mirrored session and is useful for tuning out noisy or unwanted traffic. To configure a mirror filter, follow these steps:

- From the VPC dashboard select: `Mirror Filters` under the Traffic Mirroring section in the left window pane.
- Select: `Create traffic mirror filter`
- Add the appropriate inbound and outbound rules. For mor details about traffic mirror filters please see: https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-filters.html
- Select: `Create`

5.13.17 Create Mirror Session

A traffic mirror session defines the source of the traffic to be mirrored based on the selected traffic mirror filters and sends that traffic to the desired traffic mirror target. For more details about traffic mirror sessions please see: https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-session.html

- From the VPC dashboard select: `Mirror Sessions` under the Traffic Mirroring section in the left window pane.
- Select: `Create traffic mirror session`
- Under the Mirror source section, choose the interface that you want to be mirrored.
- Under the Mirror target section, choose the interface or load balancer you want to send the mirrored traffic to.
- Assign a session number under the Additional settings section for the mirror session.
- In the filters section under Additional settings choose the mirror filter you want to apply to the mirrored traffic.
- Select: `Create`

5.13.18 Verify Traffic Mirroring

To verify the mirror session is sending the correct data to the sniffing interface run the following command on the Security Onion AWS Sensor instance:

```
sudo tcpdump -nni <interface>
```

You should see `VXLAN` tagged traffic being mirrored from the interface you selected as the Mirror Source.

To verify Zeek is properly decapsulating and parsing the VXLAN traffic you can verify logs are being generated in the `/nsm/zeek/logs/current` directory:

```
ls -la /nsm/zeek/logs/current/
```

5.14 Azure Cloud Image

Azure users can deploy an official Security Onion virtual machine image found on the Azure Marketplace: https://securityonion.net/azure

Warning

Existing Security Onion cloud image installations should use the soup command to upgrade. If your grid is still running 2.4.x, use `soup` to upgrade to 2.4.210, and then use `soupto3` to proceed to Security Onion 3, after which continue using `soup` again. Attempting to switch to a newer Security Onion image from the cloud marketplace could cause loss of data and require full Grid re-installation; use the `soup` procedure to upgrade instead.

Note

As of Septemeber, 2025, Azure has released a preview of their Virtual TAP feature. See more information here: https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-tap-overview

Note

This section does not cover network connectivity to the Security Onion node. This can be achieved through configuring an external IP for the node's management interface, or through the use of a VPN connection via OpenVPN.

Note

This section does not cover how to set up a virtual network in Azure. For more details about setting up a virtual network, please see https://docs.microsoft.com/en-us/azure/virtual-network/. Ensure that all Security Onion nodes can access the manager node over the necessary ports. This could require adding rules to your Azure Virtual Network and/or VMs in order to satisfy the Security Onion Firewall Node Communication requirements.

5.14.1 Requirements

Before proceeding, determine the grid architecture desired. Choose from a single-node Grid versus a distributed, multi-node Grid.

Security Onion recommends using either Premium SSD disks, or the more expensive Ultra SSD disks, with suitable IOPS and throughput matched to your expected network monitoring requirements.

5.14.2 Single Node Grid

For simple, low-volume production monitoring, a single node Grid can be used.

Listed below are the minimum suggested single-node instance quantities, sizes, and storage requirements for either standalone or evaluation installations (choose one, not both). Note that when using virtual machines with the minimum RAM requirements you may need to enable memory swapping.

Standalone

- Quantity: 1
- Type: Standard_D8as_v4
- Storage: 256GB Premium SSD

Evaluation

- Quantity: 1

- Type: Standard_D8as_v4
- Storage: 256GB Premium SSD

5.14.3 Distributed Grid

For high volume production monitoring, choose a multi-node Grid architecture. At least two search nodes are recommended for redundancy purposes.

Listed below are the minimum suggested distributed Grid instance quantities, sizes, and storage requirements. Prefer increasing VM memory over enabling swap memory, for best performance. High volume networks will need more powerful VM types with more storage than those listed below.

VPN Node

- Quantity: 1
- Type: Option 1: Standard_B1s - Lower cost for use with low vpn traffic volume
- Type: Option 2: Standard_D4as_v4 w/ accelerated networking - Higher cost for high vpn traffic volume
- Storage: 64GB Premium SSD

Manager

- Quantity: 1
- Type: Standard_D8as_v4
- Storage: 256GB Premium SSD

Search Nodes

- Quantity: 2 or more
- Type: Standard_D8as_v4
- Storage: 256GB Premium SSD

Sensor monitoring the VPN ingress

- Quantity: 1
- Type: Standard_D4as_v4
- Storage: 512GB Premium SSD

5.14.4 Create Monitoring Interface

To setup a Security Onion sensor node in Azure, follow the prerequisite steps below prior to creating the sensor VM.

5.14.5 Create a Security Group for Sniffing Interface

Security Groups act like a firewall for your Azure virtual machines, controlling both inbound and outbound traffic. You should consider whether a security group is needed for your virtual network, and specifically for the interface that you will be using to sniff the traffic. This security group will

need to be as open as possible to ensure all traffic destined to the sniffing interface will be allowed through. To create a security group, follow these steps:

- In the Azure Dashboard search for: `Network security groups`.
- Select: `Create`
- Provide a name, such as `so-monitoring-security-group`.
- Select the appropriate resource group and region.
- Select `Review + Create`
- Review the summary
- Select: `Create`
- Select: `Go to resource`
- Adjust the Inbound security rules to ensure that all incoming monitoring traffic is allowed.

5.14.6 Create Sniffing Interface

Prior to launching the Security Onion sensor virtual machine you will need to create the interface that will be used to monitor your virtual network. This interface will be attached to the Security Onion sensor virtual machine as a secondary interface. To create a sniffing interface, follow these steps:

- In the Azure Dashboard search for: `Network interfaces`.
- Select: `Create`
- Provide a name, such as `so-monitoring-interface`.
- Choose the resource group, region, virtual network, subnet, security group from the steps above, and IP settings.
- Select: `Review + Create`
- Review the summary
- Select: `Create`

5.14.7 Setup Virtual Network TAP

To accomplish traffic monitoring in Azure, a virtual network TAP must be created. This can be created in any resource group but must exist in the same region as the source and destination interfaces. Locate the VTAP screen in the Azure portal by searching for `vtap` and a list of results will include "Virtual network termination access points". Choose that and then click the `Create` button. Choose the region that matches your source and destination VMs. Choose the destination by locating the sniffing interface created earlier. It will show port 4789 since the monitored traffic will arrive at the sniffing interface as VxLAN traffic. Finally, choose the source network interface(s) to monitor. Multiple interfaces can be selected if needing to monitor multiple VMs.

Review the Azure VTAP documentation located at https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-tap-overview to ensure all other requirements have been satisfied.

5.14.8 Create Security Onion Instances

5.14.9 Instance Creation

To configure a Security Onion instance (repeat for each node in a distributed Grid), follow these steps:

- In the Azure Dashboard search for: `Virtual machines`
- Select: `Create` and then `Virtual machine`
- Choose or create a new Resource group.
- Enter a suitable name for this virtual machine, such as `so-vm-manager`.
- Choose the desired Region and Availability options. (Use `East US 2` for Ultra SSD support, if needed.)
- Choose the `Security Onion VM Image`. If this option is not listed on the Image dropdown, select `See all images` and search for `onion`.
- Choose the appropriate Size based on the desired hardware requirements. For assistance on determining resource requirements please review the Requirements section above.
- Change the Username to `onion`. Note that this is not mandatory -- if you accidentally leave it to the default `azureuser`, that's ok, you'll simply use the `azureuser` username any place where the documentation states to use the `onion` username.
- Select an existing SSH public key if one already exists, otherwise select the option to `Generate new key pair`.
- Choose `Other` for Licensing type.
- Select `Next: Disks`
- Ensure `Premium SSD` is selected.
- For single-node grids, distributed sensor nodes, or distributed search nodes: If you would like to separate the `/nsm` partition into its own disk, create and attach a data disk for this purpose, with a minimum size of 100GB, or more depending on predicted storage needs. Note that the size of the `/nsm` partition determines the rate that old packet and event data is pruned. Separating the /nsm partition can provide more flexibility with scaling up the grid node sizes, but requires a little more setup, which is described later.
- Select `Next: Networking`
- Choose the virtual network for this virtual machine.
- Choose a public IP if you intend to access this virtual machine directly (not recommended for production grids).
- Choose appropriate security group settings. Note that this is typically not the same security group used for the sensor monitoring interface.
- Accelerated networking will be automatically enabled if the virtual machine size supports it.
- Select: `Review + create`
- Review the summary. If a `Validation failed` message appears, correct the missing inputs under each tab section containing a red dot to the right of the tab name.
- Select. `Create` and download the new public key, if you chose to generate a new key.
- If this VM is a single-node Grid, or is sensor node:
- Stop the new VM after deployment completes.
- Edit the VM and attach the monitoring network interface created earlier.
- Start the VM.

Note that you'll need to reference the SSH public key when using SSH to access the new VMs. For example:

```
chmod 600 ~/Downloads/onion.pem
ssh -i ~/Downloads/onion.pem onion@11.22.33.44
```

5.14.10 Manager Setup

After SSH'ing into the node, setup will begin automatically. Follow the prompts, selecting the appropriate install options. Most distributed installations will use the `hostname` or `other` web access method, due to the need for both cluster nodes inside the private network, and analyst users across the public Internet to reach the manager. This allows for custom DNS entries to define the correct IP (private vs public) depending on whether it's a

cluster node or an analyst user. Users evaluating Security Onion for the first time should consider choosing the `other` option and specifying the node's public cloud IP.

5.14.11 Search Node Setup

Follow standard Security Onion search node installation, answering the setup prompts as applicable.

5.14.12 Remote Sensor Setup

Setup the VPN (out of scope for this guide) and connect the sensor node to the VPN. When prompted to choose the management interface, select the VPN tunnel interface, such as `tun0`. Use the internal IP address of the manager inside Azure when prompted for the manager IP.

5.14.13 Azure Sensor Setup

SSH into the sensor node and run through setup to set this node up as a sensor. Choose `eth0` as the main interface and `eth1` as the monitoring interface.

5.14.14 Verify Monitoring Traffic

To verify the Azure sensor is receiving the correct data on the sniffing interface run the following command on the Security Onion Azure sensor instance:

```
sudo tcpdump -nni eth1
```

To verify Zeek is properly decapsulating and parsing the traffic you can verify logs are being generated in the `/nsm/zeek/logs/current` directory:

```
ls -la /nsm/zeek/logs/current/
```

5.15 Google Cloud Image

If you would like to deploy Security Onion in Google Cloud Platform (GCP), choose the Security Onion image listed on the Google Marketplace: https://securityonion.net/google/?ref=_ptnr_soc_docs_260311

Warning

Existing Security Onion AMI installations should use the soup command to upgrade. If your grid is still running 2.4.x, use `soup` to upgrade to 2.4.210, and then use `soupto3` to proceed to Security Onion 3, after which continue using `soup` again. Attempting to switch to a newer Security Onion image from the cloud marketplace could cause loss of data and require full Grid re-installation; use the `soup` procedure to upgrade instead.

Note

This section does not cover network connectivity to the Security Onion node. This can be achieved through configuring an external IP for the node's management interface, or through the use of a VPN connection via OpenVPN.

Note

This section does not cover all aspects of how to set up a VPC in GCP, as each deployments is typically unique for the user. For more details about setting up a VPC, please see https://cloud.google.com/vpc/docs/vpc. Ensure that all Security Onion nodes can access the manager node over the necessary ports. This could require adding rules to your GCP Virtual Private Cloud and/or VMs in order to satisfy the Security Onion Firewall Node Communication requirements.

5.15.1 Requirements

Before proceeding, determine the grid architecture desired. Choose from a single-node Grid versus a distributed, multi-node Grid. Additionally, determine if the lower latency of local instance storage is needed (typically when there is high-volume of traffic being monitored, which is most production scenarios), or if persistent disks can be used for increased redundancy.

5.15.2 Single Node Grid

For simple, low-volume production monitoring, a single node Grid can be used. Persistent disks must be used for Elasticsearch data storage if used for production purposes. Single node grids cannot use local disks without being at risk of losing Elasticsearch data. However, for temporary evaluation installations, where there is little concern for data loss, local disks can be used.

Listed below are the minimum suggested single-node instance quantities, sizes, and storage requirements for either standalone or evaluation installations (choose one, not both). Note that when using virtual machines with the minimum RAM requirements you may need to enable memory swapping.

Standalone

- Quantity: 1
- Type: n2-standard-8
- Storage: 256GB Balanced Persistent Disk

Evaluation

- Quantity: 1
- Type: n2-standard-8
- Storage: 256GB SSD Persistent Disk

- Assuming evaluation of performance as well as functionality, therefore higher minimums compared to standalone.

5.15.3 Distributed Grid

For high volume production monitoring, choose a multi-node Grid architecture. At least two search nodes must be used in this architecture. This is required due to the use of local disks for Elasticsearch data storage, where each of the search nodes retains a replica of another search node, for disaster recovery.

Listed below are the minimum suggested distributed Grid instance quantities, sizes, and storage requirements. Prefer increasing VM memory over enabling swap memory, for best performance. High volume networks will need more powerful VM types with more storage than those listed below.

VPN Node

- Quantity: 1
- Type: e2.micro
- Storage: 50GB Balanced Persistent Disk

Manager

- Quantity: 1
- Type: n2-standard-8
- Storage: 300GB Balanced Persistent Disk

Search Nodes

- Quantity: 2 or more
- Type: n2-standard-8
- Storage: 256GB Balanced Persistent Disk
- Storage: 375GB Local Disk (NVMe) [optional]

Sensor monitoring the VPN ingress

- Quantity: 1
- Type: n2-standard-4
- Storage: 500GB Balanced Persistent Disk

5.15.4 Setup Traffic Mirroring

To accomplish traffic mirroring in GCP, a packet mirroring policy must be created and assigned to an internal load balancer. Google supports multiple methods for selecting what traffic to mirror. For example, a special tag keyword can be configured on the mirror policy, such as "so-mirror", and any VM that should have its traffic monitored can be given that special tag. The mirrored traffic will be forwarded to the internal load balancer, and a Security Onion sensor VM will be a member of that load balancer's instance group.

Follow the steps below to setup a traffic mirroring configuration. You will need to be logged into the Google Cloud Console, and somewhat familiar with GCP and how zones and regions are used. Note that these steps are only one of many ways to do this. For example, your scenario may require more advanced configuration, such as packet filtering, or additional VPCs.

5.15.5 Create a VPC for the Monitored Network

Create a new Virtual Private Cloud (VPC) network for collection of monitored network traffic. This will be referred to below as the Monitored VPC network. Define one subnet within this VPC that will be dedicated to receiving monitored traffic.

Add a new firewall rule to this VPC network to allow all incoming mirrored traffic. Specify a target tag of `so-collector` and a source tag of `so-mirror`. This will allow all mirrored traffic originating from a VM NIC tagged with `so-mirror`, and residing in this same VPC network, to be delivered to the sensor VM's monitoring NIC tagged with `so-collector`.

5.15.6 Create a VPC for the Security Onion Network

Create a new Virtual Private Cloud (VPC) network where the Security Onion Grid will communicate. Configure the subnets as desired, however, at least one subnet is required, and this VPC cannot overlap IP space with the above Monitored VPC network. Ensure that SSH access (port TCP/22) and HTTPS (port TCP/443) is enabled so that you have the ability to connect to VMs from your external network. For security purposes it's recommended to limit inbound access from trusted IPs.

Add a new firewall rule to allow all traffic originating from any VM instance within the Security Onion VPC network. Choose a source IP range that encapsulates the IP ranges of the subnet(s) created above. This is necessary for connectivity between the manager and minion nodes. You can also choose to be more specific about traffic within the VPC however the rules must satisfy the Security Onion Firewall Node Communication requirements.

5.15.7 Create Sensor Instance Group

Create an unmanaged Instance Group. This is found under the Compute Engine section of the Google Cloud Console. Use the Security Onion VPC as the selected network. Leave the VM instances blank; later in this document the Security Onion sensor node will be added to this group. Port mapping is not required for this group.

5.15.8 Create Internal Load Balancer

Under Network services, within the Google Cloud Console, create a Load Balancer. Choose TCP Load Balancer and select the `Only between my VMs` option. Click Continue and then select the Monitoring VPC network.

For the Backend configuration, choose the Instance Group created above. Ignore the informative box that explains the need to use additional NICs in the group instances. Specify that the backend is a failover group for backup. Create a new Health check that uses port TCP/22 (SSH) as the health test, with the following timing settings:

- Check Interval: 300
- Timeout: 1
- Healthy Threshold: 1
- Unhealthy Threshold: 1

Note that this health check is put in place only to satisfy the GCP requirement that all backends have a health check assigned. Since the backend group is marked as a failover, it will always forward traffic, regardless of the health check result.

For the Frontend configuration, select the subnet in the Monitoring VPC network that you created specifically for receiving monitored traffic. Choose non-shared IP. If there you would like to forward all traffic, choose All ports and enable global access. Under Advanced Configurations, enable the `Load Balancer for Packet mirroring` checkbox.

5.15.9 Create Packet Mirroring Policy

Traffic mirroring allows you to copy the traffic to/from an instance (or multiple instances) and send it to the sniffing interface of a network security monitoring sensor or a group of interfaces using a network load balancer. For more details about GCP Traffic Mirroring please see: https://cloud.google.com/vpc/docs/packet-mirroring

Create a Packet Mirroring policy. This can be found in the Google Cloud Console under the VPC network section. When selecting the VPC network, choose the option that denotes the mirrored source and collector destination are in the same VPC network and select the Mirrored VPC network created earlier.

Under Select mirrored source, check the box next to the "Select with network tag" label. Then enter a tag named `so-mirror`. Once completed with the grid setup, you can later tag all your VMs, whose traffic you want monitored, with the same `so-mirror` tag.

Under Select collector destination, choose the front end forwarding rule that was created during the Load Balancer setup earlier.

Finally, choose to mirror all traffic, unless you prefer to filter specific traffic for mirroring.

5.15.10 Create Security Onion Instances

5.15.11 Instance Creation

To configure a Security Onion instance (repeat for each node in a distributed Grid), follow these steps:

- Access the Google Cloud Marketplace at https://console.cloud.google.com/marketplace.
- Ensure you have a means of authenticating to VM instances over SSH. One method to authenticate is via a project-wide SSH key, which can be defined in Compute Engine -> Metadata -> SSH Keys.
- Search the Marketplace for `Security Onion` and Launch the latest version of the Security Onion official VM image. This may require clicking the "Get Started" button.
- Choose the appropriate machine type based on the desired hardware requirements. For assistance on determining resource requirements please review the Requirements section above.
- Under the Networking interfaces section, expand the pre-added Network interface and select the Security Onion VPC network and desired subnet. External ephemeral IP is sufficient, unless you are planning to use a VPN to access the Security Onion Console, in which case no external ephemeral IP is necessary. Using a VPN is recommended, but setup of a VPN in GCP is out of scope of this guide.
- (Distributed "Sensor" node or Single-Node Grid only) Add a second Network interface and select the monitoring VPC network, and the appropriate subnet. No external ephemeral IP is necessary for this interface.
- (Distributed "Manager" node or Single-Node Grid only) If not using a VPN, enable the Allow SSH and HTTPS traffic from the the desired IPs/CIDRs.
- Adjust the boot disk size and type as necessary, using the guidance in the above Requirements section and elsewhere in the Security Onion documentation.
- (Distributed "Search" node or Evaluation Grid only) If high-speed local NVMe disks are needed a manual Terraform deployment will be required, with local disks added to the compute instance. Note that the local disks support 375GB only, so if larger volumes are required the system will need to be configured to present the multiple local disks as one virtual disk. This is out of scope of this document. Also, be aware that local disks are not replicated and will result in data loss if the instance is deleted.
- If requested, review GCP Marketplace Terms, and if acceptable click the corresponding checkbox.
- Select: `Create`

5.15.12 Prepare Nodes with Ephemeral, Local Disk Storage

For distributed search nodes, or an evaluation node if using local disk storage, SSH into the node and cancel out of the setup. Prepare the local disk partition by executing the following command:

> **Note**
>
> This assumes a single NVMe disk will be used for storing all search data. Combining multiple disks into a single partition mount is out of scope of this document.

```
sudo so-prepare-fs
```

By default, this command expects the local disk device to be located at `/dev/nvme1n1` and will mount that device at `/nsm/elasticsearch`. If this fails run `lsblk` to determine which disk to use. To override either of those two defaults, specify them as arguments. For example:

```
sudo so-prepare-fs /dev/nvme0n1 /nsm/elasticsearch
```

Restart the Security Onion setup by running the following command:

```
cd /securityonion
sudo ./so-setup-network
```

5.15.13 Manager Setup

If this is an evaluation node with a local disk, ensure the node has been prepared as described in the preceding section.

After SSH'ing into the node, setup will begin automatically. Follow the prompts, selecting the appropriate install options. Most distributed installations will use the `hostname` or `other` web access method, due to the need for both cluster nodes inside the private network, and analyst users across the public Internet to reach the manager. This allows for custom DNS entries to define the correct IP (private vs public) depending on whether it's a cluster node or an analyst user. Users evaluating Security Onion for the first time should consider choosing the `other` option and specifying the node's public cloud IP.

GCP provides a built-in NTP server at hostname `metadata.google.internal`. This can be specified in the SOC Configuration screen after setup completes. By default the server will use the time servers at `ntp.org`.

For distributed manager nodes using ephemeral storage, go to SOC Configuration. Search for `number_of_replicas` and change to 1. This will double the storage cost but will ensure at least two VMs have the data, in case of an ephemeral disk loss.

Optionally, adjust ElastAlert indices so that they have a replica. This will cause them to turn yellow but that will be fixed when search nodes come online:

```
so-elasticsearch-query ElastAlert*/_settings -X PUT -d '{"index" : { "number_of_replicas" : 1 }}'
```

This is an optional step due to the ElastAlert indices being used primarily for short-term/recent alert history. In the event of a data loss when ElastAlert 2 restarts the indices will be regenerated.

5.15.14 Search Node Setup

Follow standard Security Onion search node installation, answering the setup prompts as applicable. If you are using local disk storage be sure to first prepare the instance as directed earlier in this section.

5.15.15 GCP Sensor Setup

In the GCP console, under Compute Engine go to the Instance Group page and edit the instance group that was created earlier. Use the dropdown list to add the new sensor VM instance to this group.

SSH into the sensor node and run through setup to set this node up as a sensor. Choose `ens4` as the main interface and `ens5` as the monitoring interface.

5.15.16 Remote Sensor Setup

Setup the VPN (out of scope for this guide) and connect the sensor node to the VPN. When prompted to choose the management interface, select the VPN tunnel interface, such as `tun0`. Use the internal IP (not the ephemeral IP) address of the manager inside GCP when prompted for the manager IP.

If connecting sensors through the VPN instance you will need to add the inside interface of your VPN concentrator to the `sensor` firewall hostgroup. For instance, assuming the following architecture:

```
SO Sensor        -> VPN Endpoint     -> Internet -> VPN Endpoint  -> SO Manager
Location: Remote    Location: Remote                Location: Googe  Location: Google
192.168.33.13       192.168.33.10                   10.55.1.10       10.55.1.20
```

In order to add the Remote Network Sensor Node to the grid, you would have to add `10.55.1.10` to the `sensor` firewall hostgroup.

This change can be done in the SOC Configuration screen. Then, either wait up to 15 minutes for the scheduled configuration sync to run, or force a synchronization immediately via the SOC Configuration Options. Once the firewall hostgroup configuration has been synchronized your Manager will be ready for remote minions to start connecting.

5.15.17 Verifying Traffic Mirroring

Deploy a temporary test VM instance, using a e2.micro, debian-based instance in the Monitored VPC network, and in the same region used in the rest of this guide. Add the `so-mirror` network tag to the VM.

SSH into the sensor node created earlier in this guide, and run the following command to watch mirrored traffic:

```
tcpdump -nni ens5
```

While that is running, in another terminal, SSH into this new test VM and run a curl command to a popular website. You should see that HTTP/HTTPS traffic appear in the tcpdump output.

Login to Security Onion and verify that the traffic also appears in the Hunt user interface.

Delete the temporary test VM instance when the verification is completed.

5.16 Configuration

Now that you've installed Security Onion, it's time to configure it!

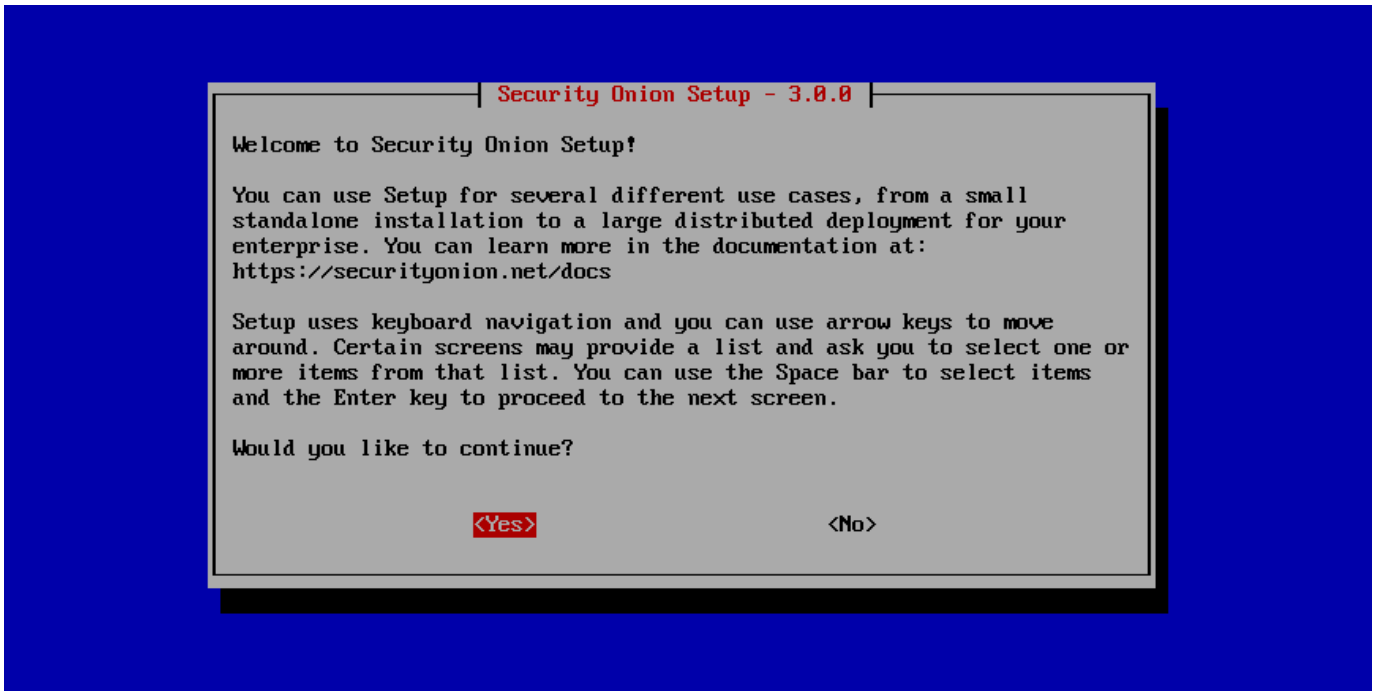

Security Onion is designed for many different use cases. Here are just a few examples!

> **Tip**
>
> If this is your first time using Security Onion and you just want to try it out, we recommend the Import option as it's the quickest and easiest way to get started.

> **Warning**
>
> If the network configuration portion displays a message like `The IP being routed by Linux is not the IP address assigned to the management interface`, then you have multiple network interfaces with IP addresses. In most cases, sniffing interfaces should not have IP addresses and there should only be an IP address on the management interface itself. Sometimes this is caused by a sniffing interface connected to a normal switch port (not a TAP/SPAN port) and acquiring an IP address via DHCP. Double-check your network interfaces, wiring, and configuration.

5.16.1 Import

One of the easiest ways to get started with Security Onion is using it to forensically analyze PCAP and log files. Simply select the `IMPORT` option, follow the prompts, and then import PCAP files or Windows event logs in EVTX format using the Grid page.

5.16.2 Evaluation

Evaluation mode is ideal for classroom or small lab environments. Evaluation is **not** designed for production usage. Choose the `EVAL` option, follow the prompts, and then proceed to the Post Installation section.

5.16.3 Production Server - Standalone

Standalone is similar to Evaluation in that it only requires a single box, but Standalone is more ready for production usage. Choose `STANDALONE`, follow the prompts, and then proceed to the Post Installation section.

5.16.4 Production Server - Distributed Deployment

If deploying a distributed environment, install and configure the manager node first and then join the other nodes to it. For best performance, the manager node should be dedicated to just being a manager for the other nodes (the manager node should not do any network sniffing, that should be handled by dedicated sensor nodes).

Build the manager by running Setup, selecting the `DISTRIBUTED` deployment option, and choosing the `New Deployment` option. You can choose either `MANAGER` or `MANAGERSEARCH`. If you choose `MANAGER`, then you must join one or more search nodes (this is optional if you choose `MANAGERSEARCH`) and you will want to do this before you start joining other node types.

Build nodes by running Setup, selecting the `DISTRIBUTED` deployment option, choosing `Existing Deployment`, and selecting the appropriate option. Please note that all nodes will need to be able to connect to the manager node on several ports and the manager will need to connect to search nodes and heavy nodes. You'll need to make sure that any network firewalls have firewall rules to allow this traffic as defined in the Firewall section. In addition to network firewalls, you'll need to make sure the manager's host-based firewall allows the connections. You can do this in two ways. The first option is going to Administration --> Configuration --> firewall --> hostgroups, selecting the appropriate node type, and adding the IP address. The second option is to wait until the node tries to join and it will prompt you to run a specific command on the manager. Regardless of which of the two options you choose, it will eventually prompt you to go to Administration --> Grid Members, find the node in the Pending Members list, click the `Review` button, and then click the `Accept` button.

Proceed to the Post Installation section.

5.17 Post Installation

5.17.1 Adjust firewall rules

Depending on what kind of installation you did, the Setup wizard may have already walked you through adding firewall rules to allow your analyst IP address(es). If you need to make other adjustments to firewall rules, you can do so by going to Administration --> Configuration --> firewall --> hostgroups.

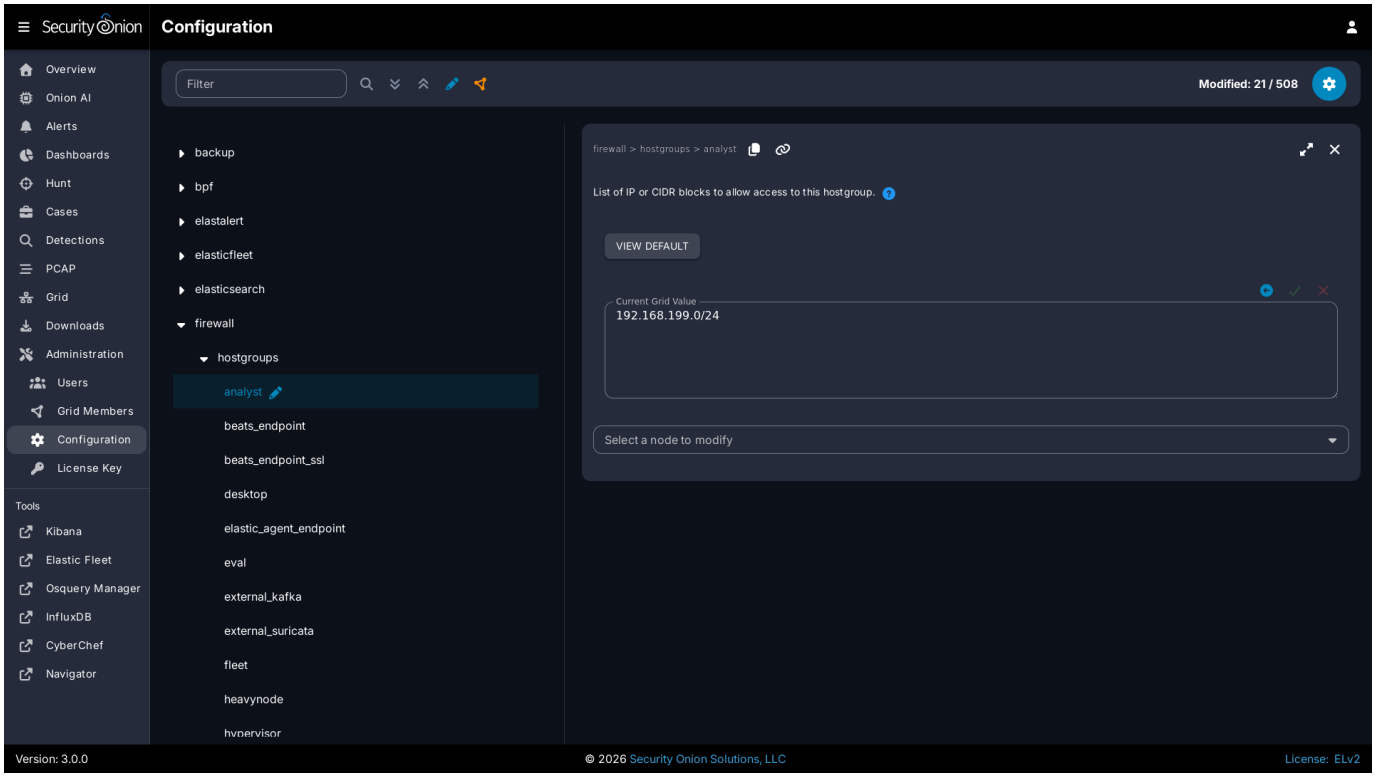

If for some reason you can't access SOC at all, you can use the so-firewall command to allow the IP address of your web browser to connect (replacing `<IP ADDRESS>` with the actual IP address of your web browser):

```
sudo so-firewall includehost analyst <IP ADDRESS>
```

For more information, please see the Firewall section.

5.17.2 Services

You can check the Grid page to see if all services are running correctly.

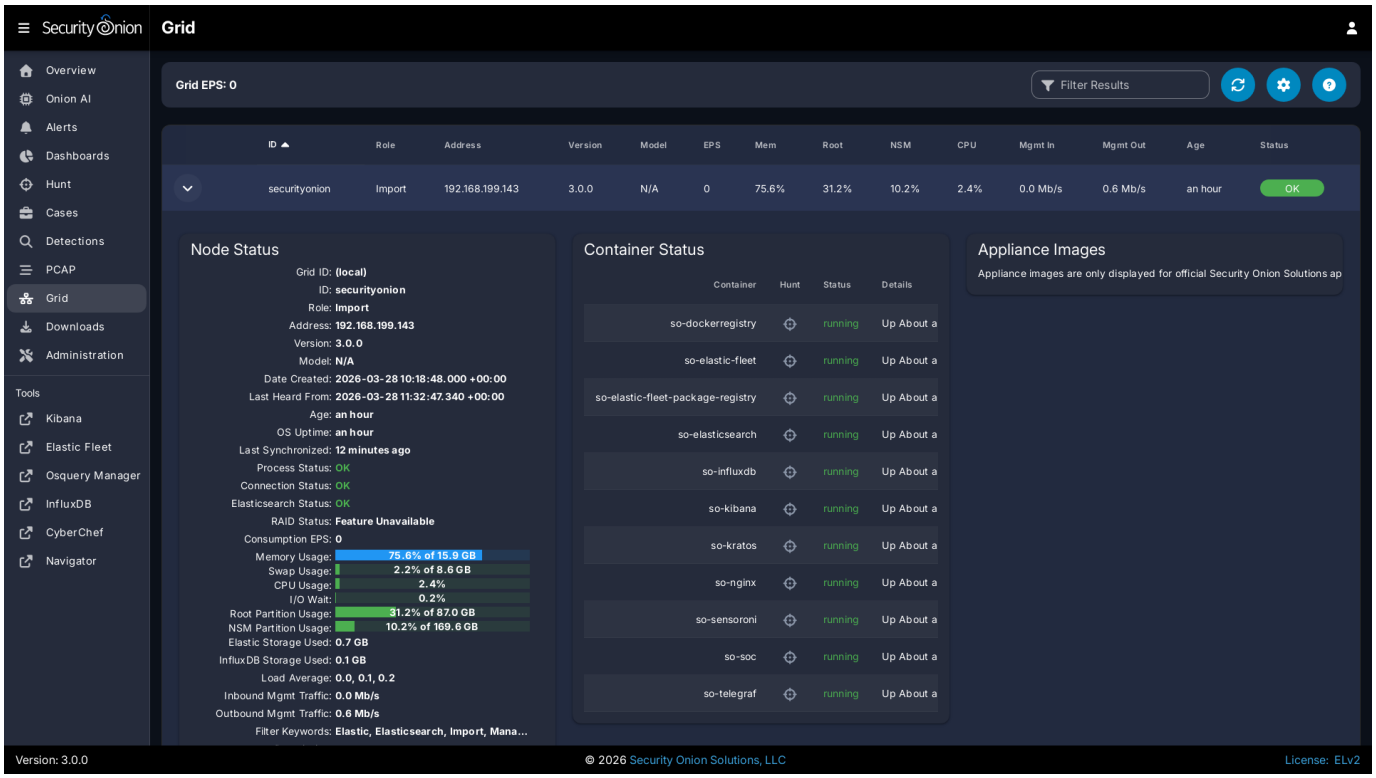

Note

Please note that new nodes start off showing a red Fault and may take a few minutes to fully initialize before they show a green OK.

You can also verify services are running from the command line with the so-status command:

```
sudo so-status
```

5.17.3 SSH

You should be able to do most administration from SOC but if you need access to the command line then we recommend using SSH rather than the Console.

5.17.4 Data

- Review the Elasticsearch section to see if you need to change any of the default settings. In particular, if you have a multi-node deployment with one or more search nodes, we HIGHLY recommend configuring ILM to delete indices before Elasticsearch reaches its watermark setting and stops ingesting new data.
- Review the Full Packet Capture and Suricata sections to see if you need to change the PCAP retention settings.

5.17.5 Other

- Go to Administration and then click Configuration to see some of the options that you may want to configure. For example, you may want to enable reverse DNS lookups when viewing IP addresses in SOC. For more information, please see the SOC Customization section.
- While on the Administration page, you may want to set your preferred NTP server.
- Full-time analysts may want to connect using a dedicated Security Onion Desktop.
- Any IDS/NSM system needs to be tuned for the network it's monitoring. Please see the Detections and Rules sections.

6. Security Onion Console

6.1 Security Onion Console Overview

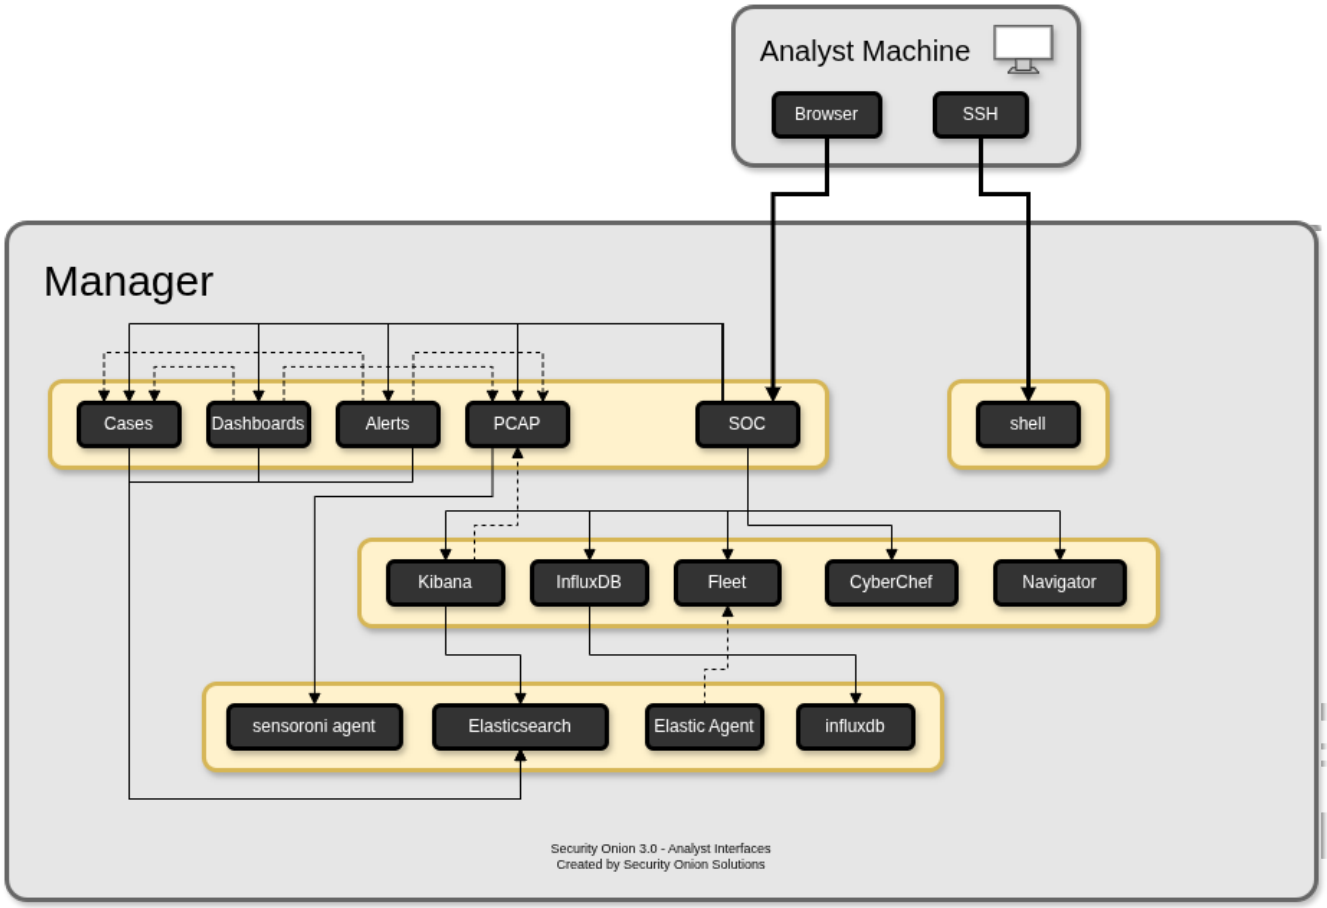

Once all configuration is complete, you can then connect to Security Onion Console (SOC) with your web browser. We recommend chromium-based browsers such as Google Chrome. Other browsers may work, but fully updated chromium-based browsers provide the best compatibility.

Depending on the options you chose in the installer, connect to the IP address or hostname of your Security Onion installation. Then login using the email address and password that you specified in the installer.

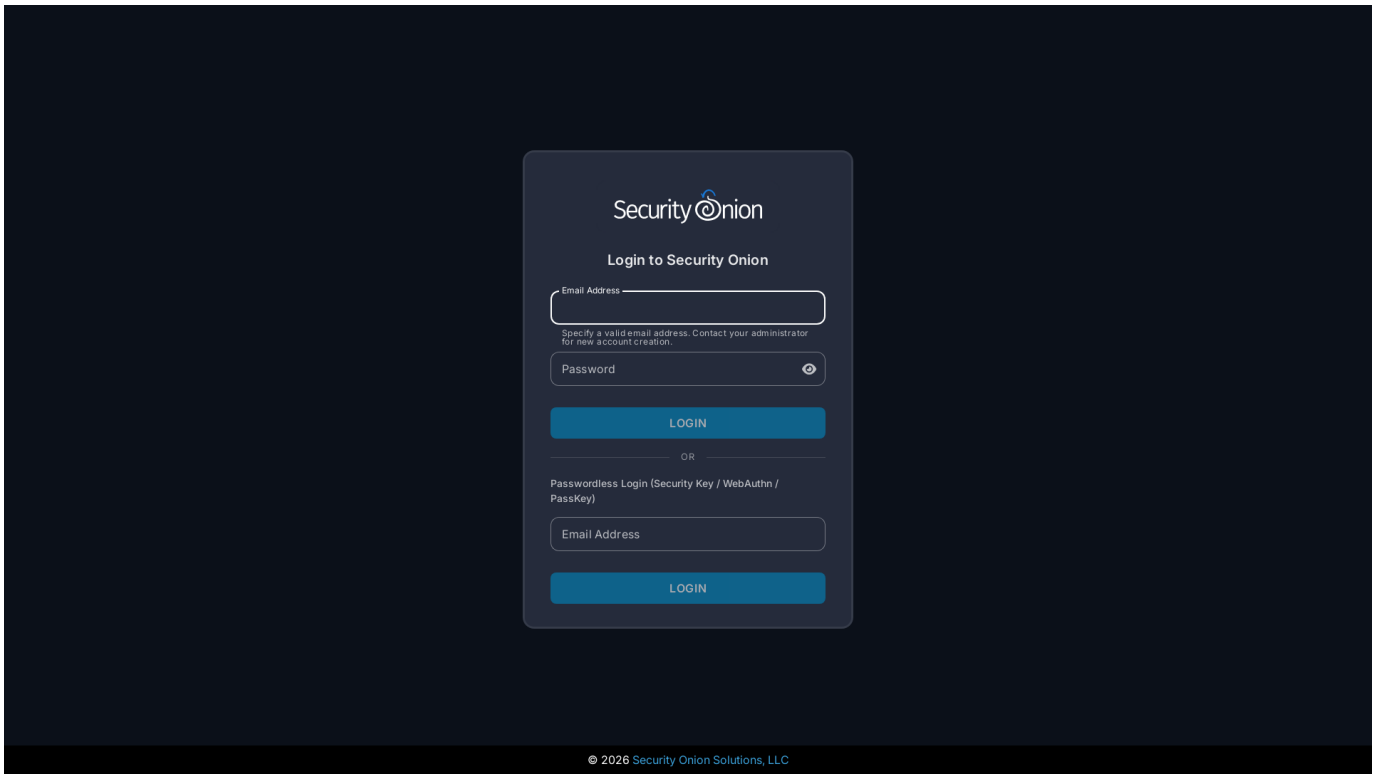

Once logged in, you'll notice the user menu in the upper-right corner. This allows you to manage your user settings and access documentation and other resources.

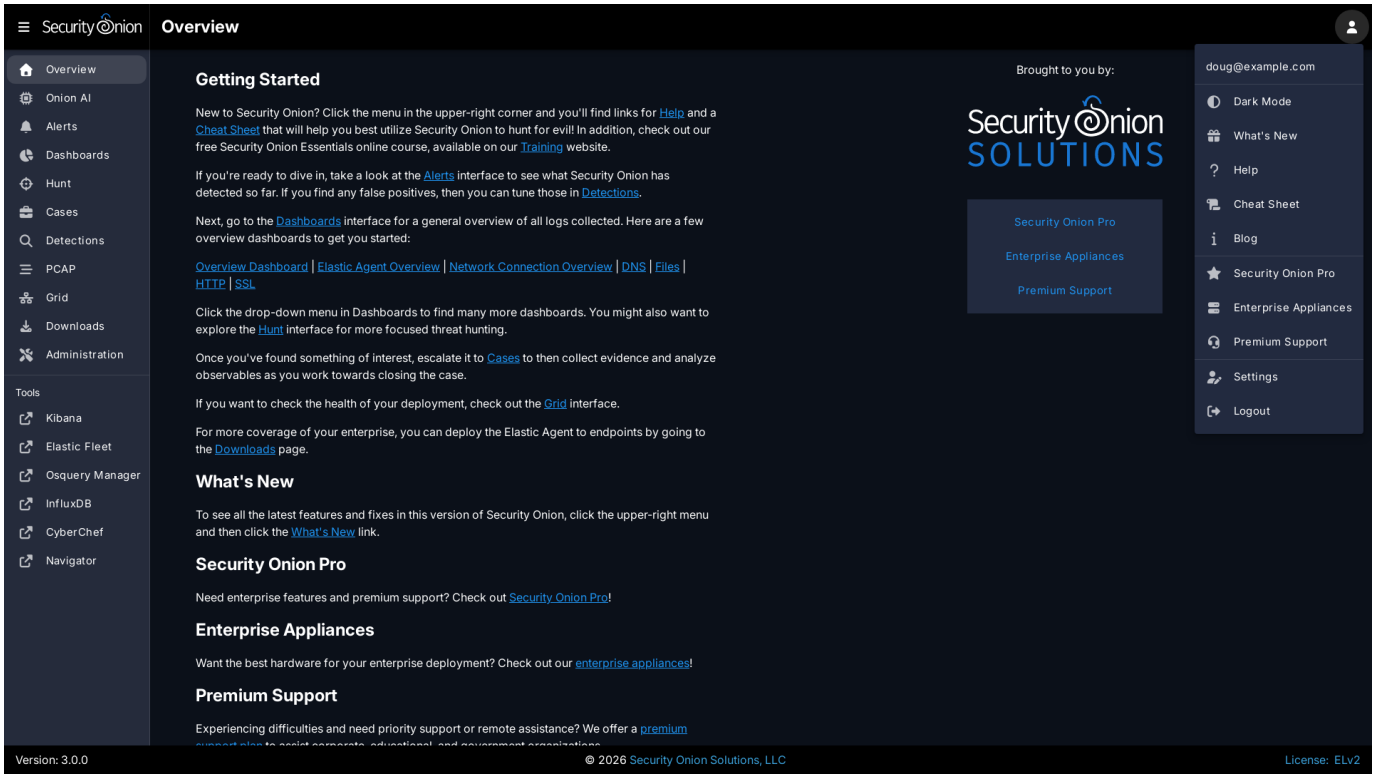

On the left side of the page, you'll see links for analyst tools like Alerts, Dashboards, Hunt, Cases, Detections, PCAP, Kibana, CyberChef, and Attack Navigator. While Alerts, Dashboards, Hunt, Cases, Detections, and PCAP are built into SOC itself, the remaining tools are external and will spawn separate browser tabs.

If you'd like to customize SOC, please see the SOC Customization section. If you'd like to learn more about SOC logs, please see the SOC Logs section.

6.2 Alerts

Security Onion Console includes an Alerts interface which gives you an overview of the alerts that Security Onion is generating. You can then quickly drill down into details, pivot to Hunt or the PCAP interface, and escalate alerts to Cases.

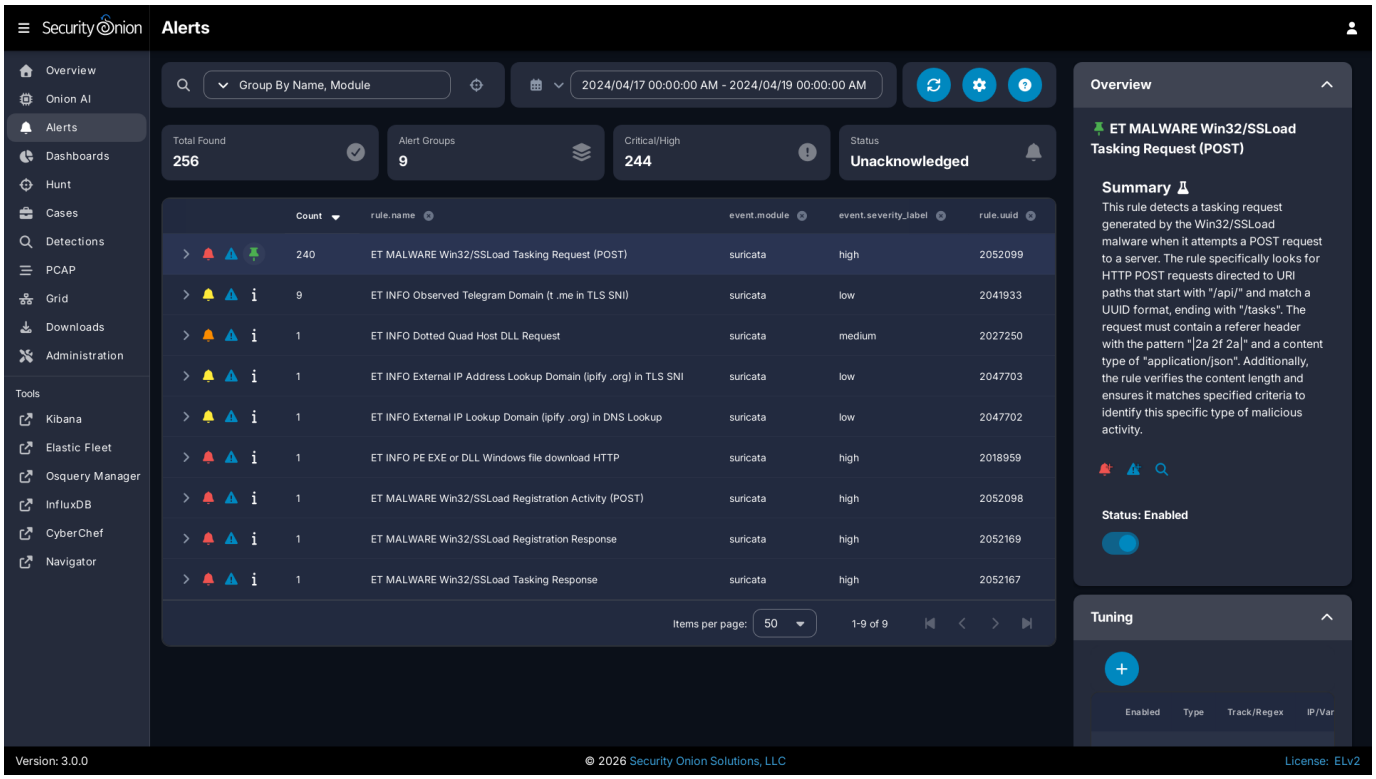

6.2.1 Options

At the top of the page, there is an Options menu that allows you to set several different options for the Alerts page.

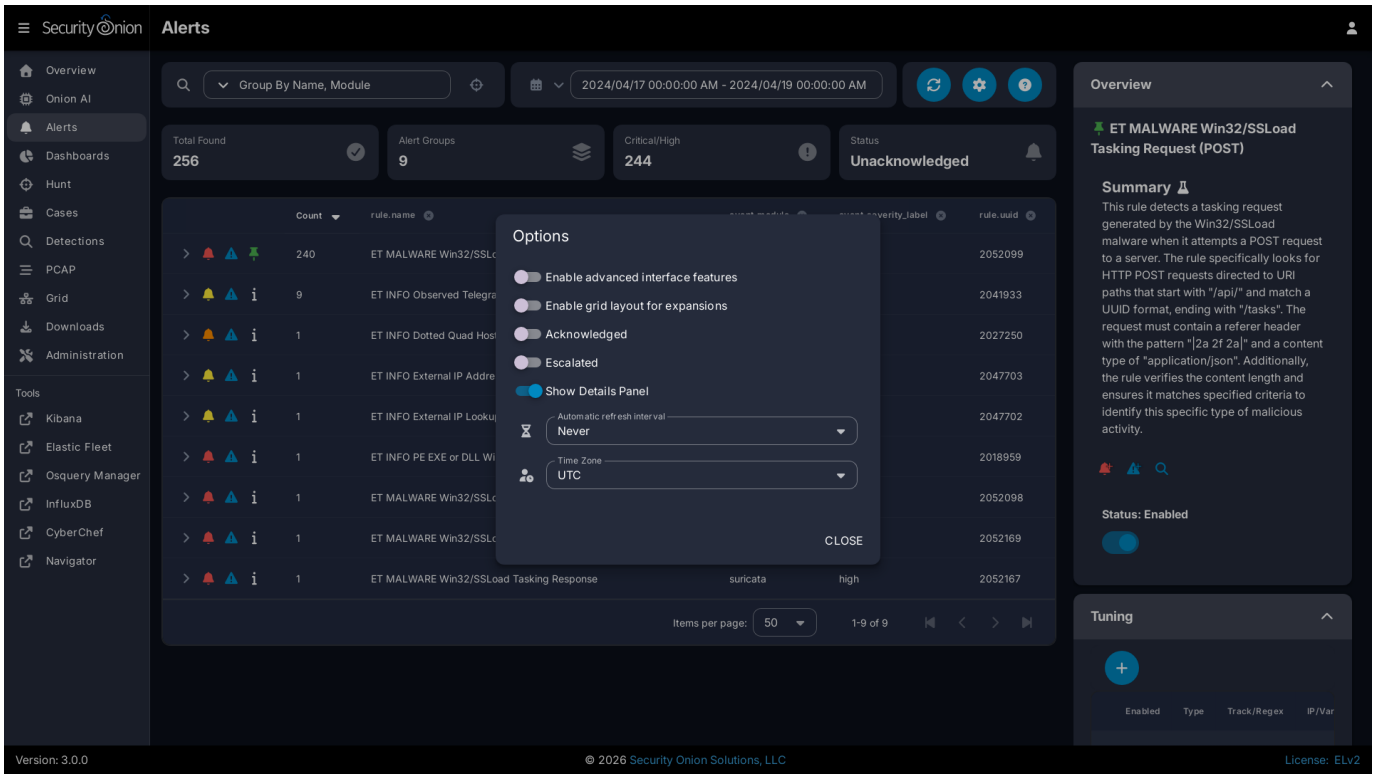

Toggles

The first toggle is labeled `Enable advanced interface features`. If you enable this option, then the interface will show more advanced features similar to Dashboards and Hunt. This includes an extra toggle labeled `Automatically apply filters, groupings, and date ranges`.

Enabling the `Acknowledged` toggle will only show alerts that have previously been acknowledged by an analyst.

Enabling the `Escalated` toggle will only show alerts that have previously been escalated by an analyst to Cases.

Finally, the `Show Details Panel` toggle controls the Details panel on the right side. For more information, see the Details Panel section below.

Automatic Refresh Interval

Another option is the Automatic Refresh Interval setting. When enabled, the Alerts page will automatically refresh at the time interval you select.

Time Zone

Alerts will try to detect your local time zone via your browser. You can manually specify your time zone if necessary.

6.2.2 Query Bar

The query bar defaults to `Group By Name, Module` which groups the alerts by `rule.name` and `event.module`. You can click the dropdown box to select other queries which will group by other fields.

On the right side of the query bar is a hunt button that will start a new hunt based on the current query.

If you would like to save your own personal queries, you can bookmark them in your browser. If you would like to customize the default queries for all users, please see the SOC Customization section.

6.2.3 Time Picker

By default, Alerts searches the last 24 hours. If you want to search a different time frame, you can change it in the upper-right corner of the screen.

6.2.4 Details Panel

The details panel on the right side shows details for the currently selected alert. This includes an AI summary (if available) and options for tuning the rule that generated the alert. This functionality is part of Detections so you can read more in that section. This panel can be disabled under the Options dropdown (see above).

6.2.5 Data Table

The remainder of the page is a data table that starts in the grouped view and can be switched to the detailed view. Both views have some functionality in common:

- Clicking the table headers allows you to sort ascending or descending.
- Starting from the left side of each row, there is an arrow which will expand the row to show more details.
- To the right of that arrow is a bell icon that acknowledges the alert. That alert can then be seen by selecting the `Acknowledged` toggle at the top of the page. In the `Acknowledged` view, clicking the bell icon removes the acknowledgement.
- To the right of that is a blue exclamation icon that escalates the alert to Cases and allows you to create a new case or add to an existing case. If you need to find that original escalated alert in the Alerts page, you can enable the `Escalated` toggle (which will automatically enable the `Acknowledged` toggle as well).
- To the right of that is an information icon that populates the Details Panel on the right with information about the alert.
- Security Onion Pro users will find one more button. The last is a computer chip icon that investigates the alert with OnionAI.
- Clicking a value in the table brings up a context menu of actions for that value. This allows you to refine your existing search, start a new search, or even pivot to external sites like Google and VirusTotal.
- You can adjust the `Rows per page` setting in the bottom right and use the left and right arrow icons to page through the table.

> **Note**
>
> If you are in the grouped view and looking at an alert that has multiple instances (the Count column is greater than 1) and you then click the arrow to expand the alert, it will show you the details for the most recent instance of that alert. If you need to see details for the other instances of the alert, then you will need to switch from the grouped view to the detailed view.

Guided Analysis

When you click the arrow to expand the alert, it starts on the `ALERT DETAILS` tab but you can switch to the `GUIDED ANALYSIS` tab which leverages Playbooks to show you plays associated with the alert. These plays include questions which help guide your investigation. You can expand all questions at once using the maximize button on the right side of the `GUIDED ANALYSIS` tab.

Each question has an associated query and the results of that query will be displayed to help you answer the question. A maximum of 5 query results will be displayed but if you want to see more, then you can click the crosshairs icon to open the query in Hunt. This also allows you to tweak the query if necessary. If you don't get any results, you could try changing the date range or other query parameters. In some cases, the query may be looking for data that you don't currently collect. For example, the query may be looking for endpoint data and so you may need to deploy the Elastic Agent to start collecting this information.

> **Note**
>
> To see Guided Analysis in action, check out our sneak peek video at https://youtu.be/SLGRB3PxB-o.

For more information about Playbooks, please see the Detections section.

> **Warning**
>
> Some playbooks were generated by AI and it's possible that they may not be 100% accurate. Please let us know if you see any issues.

Grouped View

By default, alerts are grouped by whatever criteria is selected in the query bar. Clicking a field value and then selecting the Drilldown option allows you to drill down into that value which switches to the detailed view. You can also click the value in the Count column to perform a quick drilldown. Note that this quick drilldown feature is only enabled for certain queries.

If you'd like to remove a particular field from the grouped view, you can click the trash icon at the top of the table to the right of the field name.

Detailed View

If you click a value in the grouped view and then select the Drilldown option, the display will switch to the detailed view. This shows all search results and allows you to then drill into individual search results as necessary. Clicking the table headers allows you to sort ascending or descending. Starting from the left side of each row, there is an arrow which will expand the result to show all of its fields. To the right of that arrow is the `Timestamp` field. Next, a few standard fields are shown: `rule.name`, `event.severity_label`, `source.ip`, `source.port`, `destination.ip`, and `destination.port`. Depending on what kind of data you're looking at, there may be some additional data-specific fields as well.

When you click the arrow to expand a row in the Events table, it will show all of the individual fields from that event. Field names are shown on the left and field values on the right. When looking at the field names, there are two icons to the left. The Groupby icon, the left most icon, will add a new groupby table for that field. The Toggle Column icon, to the right of the Groupby icon, will toggle that column in the Events table, and the icon will be a blue color if the column is visible. You can click on values on the right to bring up the context menu to refine your search or pivot to other pages.

6.2.6 Context Menu

Clicking a value in the page brings up a context menu that allows you to refine your existing search, start a new search, or even pivot to external sites like Google and VirusTotal.

Include

Clicking the `Include` option will add the selected field:value pair to your existing search with an `AND`. This will only show search results that include that value in that field.

Exclude

Clicking the `Exclude` option will add the selected field:value pair to your existing search with an `AND NOT`. This will only show search results that do not include that value in that field.

Only

Clicking the `Only` option will start a new search for the selected value in any field. It will remove any existing filters but retain any existing groupby terms.

Drilldown

Clicking the `Drilldown` option will drill down into a group of alerts to show each individual alert.

Tune Detection

Clicking the `Tune Detection` option will take you to Detections and allow you disable or modify the detection that fired the alert.

Group By

Clicking the `Group By` option will update the existing query and aggregate the results based on the selected field.

New Group By

Clicking the `New Group By` option will create a new data table for the selected field.

Numeric Ops

If the value you clicked is numeric, then the `Numeric Ops` sub-menu allows you to choose operations like less than, less than or equal, greater than, greater than or equal, or Between. Choosing the Between option displays a window so that you can specify a range of values.

Clipboard

The `Clipboard` sub-menu has several options that allow you to copy selected data to your clipboard in different ways.

Actions

The `Actions` sub-menu has several different options. Please note that some of these actions will only display on the Actions menu if you click on a specific log type.

- Clicking the `Hunt` option will start a new search for the selected value and will give you a good overview of what types of data are available for that indicator.
- Clicking the `Add to Case` option will add an observable to a new or existing case.
- Clicking the `Correlate` option will find related logs based on Community ID, uid, fuid, etc.
- Clicking the `PCAP` option will pivot to the PCAP interface to retrieve full packet capture for the selected stream. This option will only appear if you click on a log that contains source IP, source port, destination IP, destination port, etc.
- Clicking the `Google` option will search Google for the selected value.
- Clicking the `VirusTotal` option will search VirusTotal for the selected value.
- Clicking the `Process Info` option will show all logs that include this process's entity_id in the `process.entity_id` field. This option will only appear if you click on a log that contains the `process.entity_id` field.
- Clicking the `Process and Child Info` option will show all logs that include this process's entity_id in either the `process.entity_id` or `process.parent.entity_id` fields. Depending on the process, this may show the same logs as the `Process Info` option or it may show more. This option will only appear if you click on a log that contains the `process.entity_id` field.
- Clicking the `Process All Info` option will show all logs that include this process's entity_id in any field. Depending on the process, this may show the same logs as the `Process and Child Info` option or it may show more. This option will only appear if you click on a log that contains the `process.entity_id` field.
- Clicking the `Process Ancestors` option will show all parent processes for the selected process. This option will only appear if you click on a log that contains the `process.Ext.ancestry` field.

If you'd like to add your own custom actions, see the SOC Customization section.

6.3 Dashboards

Security Onion Console includes a Dashboards interface which includes an entire set of pre-built dashboards for our standard data types.

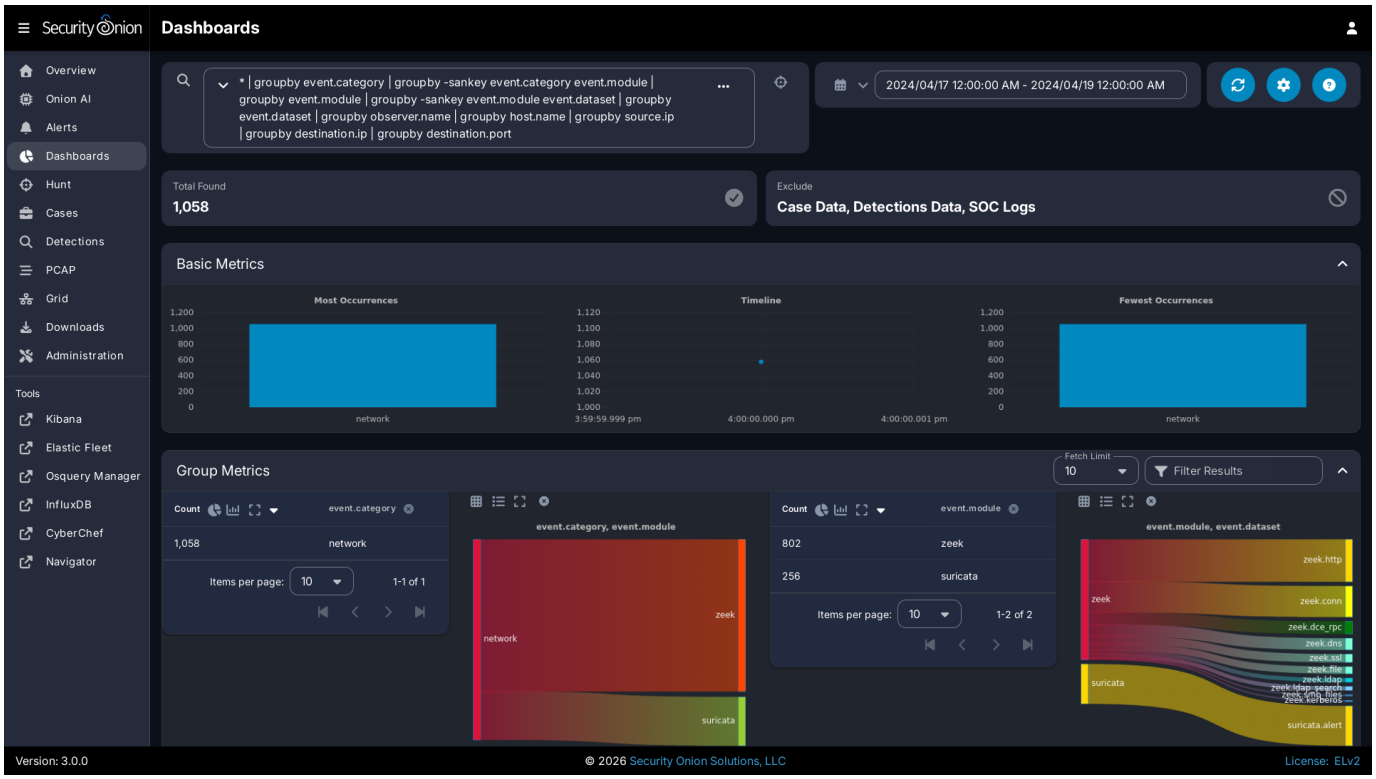

> **Note**
>
> Check out our Dashboards video at https://youtu.be/xUBhyF7se8s!

6.3.1 Options

At the top of the page, there is an Options menu that allows you to set options such as Auto Apply, Exclude case data, Exclude Detections data, Exclude SOC Logs, Automatic Refresh Interval, and Time Zone.

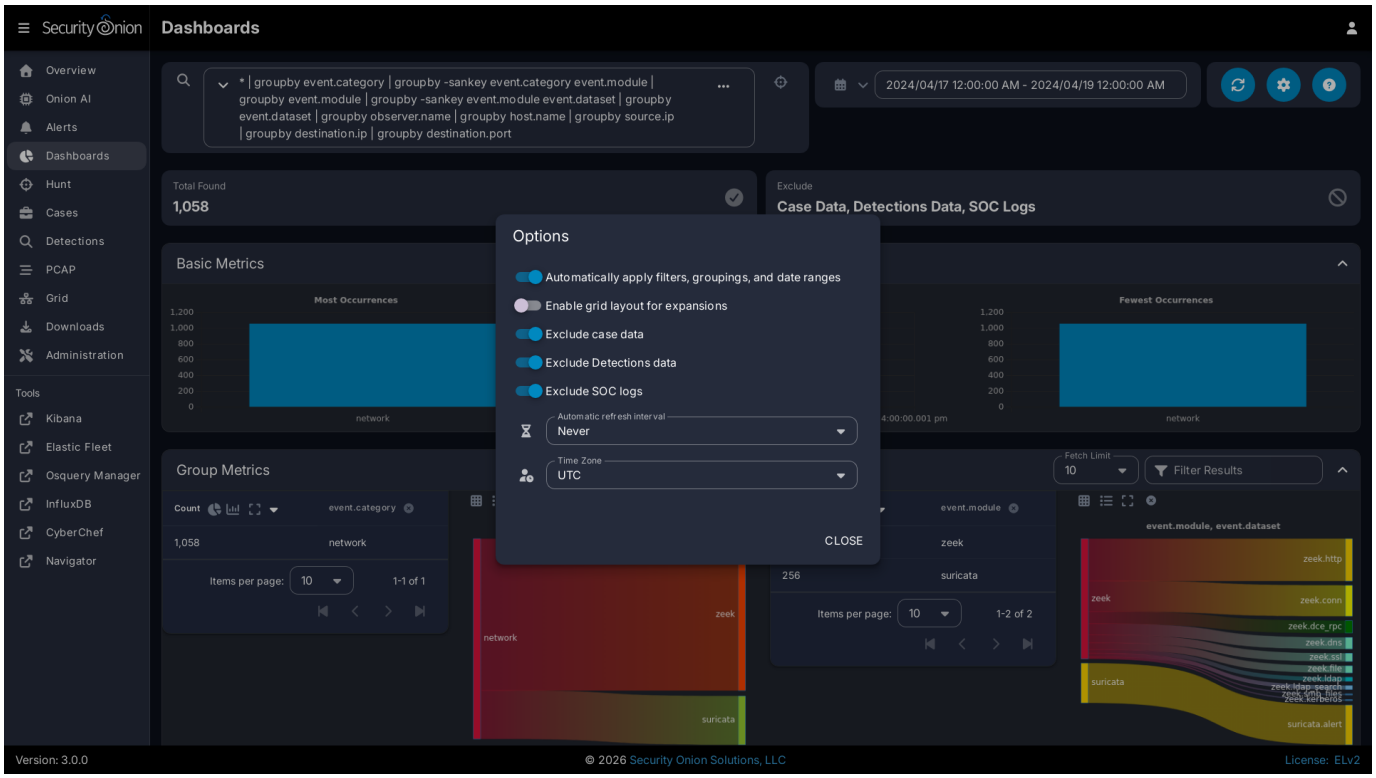

Auto Apply

The Auto Apply option defaults to enabled and will automatically submit your query any time you change filters, groupings, or date ranges.

Exclude case data

Dashboards excludes Cases data by default. If you disable this option, then you can use Dashboards to query your Cases data.

Exclude Detections data

Dashboards excludes Detections data by default. If you disable this option, then you can use Dashboards to query your Detections data.

Exclude SOC Logs

Dashboards also excludes SOC diagnostic logs by default. If you disable this option, then you can use Dashboards to query your SOC diagnostic logs.

Automatic Refresh Interval

The Automatic Refresh Interval setting will automatically refresh your query at the time interval you select.

Time Zone

Dashboards will try to detect your local time zone via your browser. You can manually specify your time zone if necessary.

6.3.2 Query Bar

The easiest way to get started is to click the query drop down box and select one of the pre-defined dashboards. These pre-defined dashboards cover most of the major data types that you would expect to see in a Security Onion deployment: NIDS alerts from Suricata, protocol metadata logs from Zeek or Suricata, Elastic Agent logs, and firewall logs.

On the right side of the query bar are two buttons. The first button is an ellipsis (three dots) which toggles between showing the full query or only the search filter. The second button is a hunt button that will start a new hunt based on the current filters.

Under the query bar, you'll notice colored bubbles that represent the individual components of the query. If you want to remove part of the query, you can click the X in the corresponding bubble to remove it and run a new search.

If you would like to save your own personal queries, you can bookmark them in your browser. If you would like to customize the default queries for all users, please see the SOC Customization section.

6.3.3 Time Picker

By default, Dashboards searches the last 24 hours. If you want to search a different time frame, you can change it in the upper-right corner of the screen. You can use the default relative time or click the clock icon to change to absolute time.

6.3.4 Basic Metrics

The Basic Metrics section of the page contains a Most Occurrences visualization, a timeline visualization, and a Fewest Occurrences visualization. Bar charts are clickable, so you can click a value to update your search criteria. Aggregation defaults to 10 values, so Most Occurrences is the Top 10 and Fewest Occurrences is the Bottom 10 (long tail). The number of aggregation values is controlled by the Fetch Limit setting in the Group Metrics section.

6.3.5 Group Metrics

The Group Metrics section of the page consists of one or more data tables or visualizations that allow you to stack (aggregate) arbitrary fields.

Group metrics are controlled by the `groupby` parameter in the search bar. You can read more about the `groupby` parameter in the OQL section below.

Clicking the table headers allows you to sort ascending or descending. Refreshing the page will retain the sort, but only for the first table.

Clicking a value in the Group Metrics table brings up a context menu of actions for that value. This allows you to refine your existing search, start a new search, or even pivot to external sites like Google and VirusTotal. The default Fetch Limit for the Group Metrics table is `10`. If you need to see more than the top 10, you can increase the Fetch Limit and then page through the output using the left and right arrow icons or increase the `Rows per page` setting.

You can use the buttons in the Count column header to convert the data table to a pie chart or bar chart. If the data table is grouped by more than one field, then you will see an additional button that will convert the data table to a sankey diagram. There is a Maximize View button that will maximize the table to fill the pane (you can press the Esc key to return to normal view). Each of the groupby field headers has a trash button that will remove the field from the table.

Once you have switched to a chart, you will see different buttons at the top of the chart. You can use the Show Table button to return to the data table, the Toggle Legend button to toggle the legend, and the Remove button to remove the chart altogether. There is a Maximize View button that will maximize the chart to fill the pane (you can press the Esc key to return to normal view).

6.3.6 Events

The third and final section of the page is a data table that contains all search results and allows you to drill into individual search results as necessary. Clicking the table header labels allows you to sort ascending or descending. You can also move a column to the right or left, or remove the column, by clicking the appropriate icons surrounding the column header labels. Starting from the left side of each row, there is an arrow which will expand the result to show all of its fields. To the right of that arrow is the `Timestamp` field. Next, a few standard fields are shown: `source.ip`, `source.port`, `destination.ip`, `destination.port`, `log.id.uid` (Zeek unique identifier), `network.community_id` (Community ID), and `event.dataset`. Depending on what kind of data you're looking at, there may be some additional data-specific fields as well.

Clicking a value in the Events table brings up a context menu of actions for that value. This allows you to refine your existing search, start a new search, or even pivot to external sites like Google and VirusTotal.

The default Fetch Limit for the Events table is `100`. If you need to see more than 100 events, you can increase the Fetch Limit and then page through the output using the left and right arrow icons or increase the `Rows per page` setting.

When you click the arrow to expand a row in the Events table, it will show all of the individual fields from that event. Field names are shown on the left and field values on the right. When looking at the field names, there are two icons to the left. The Groupby icon, the left most icon, will add a new groupby table for that field. The Toggle Column icon, to the right of the Groupby icon, will toggle that column in the Events table, and the icon will be a

blue color if the column is visible. Additionally, clicking the Toggle Column icon will add a new `| table xxx yyy zzz` segment to the active query. You can click on values on the right to bring up the context menu to refine your search or pivot to other pages.

6.3.7 Statistics

The bottom left corner of the page shows statistics about the current query including the speed of the backend data fetch and the total round trip time.

6.3.8 Context Menu

Clicking a value in the page brings up a context menu that allows you to refine your existing search, start a new search, or even pivot to external sites like Google and VirusTotal.

Include

Clicking the `Include` option will add the selected field:value pair to your existing search with an `AND`. This will only show search results that include that value in that field.

Exclude

Clicking the `Exclude` option will add the selected field:value pair to your existing search with an `AND NOT`. This will only show search results that do not include that value in that field.

Only

Clicking the `Only` option will start a new search for the selected value in any field. It will remove any existing filters but retain any existing groupby terms.

Group By

If one or more `Group By` data tables already exists, clicking the `Group By` option will add the field to the most recent data table. If there are no existing `Group By` data tables, clicking the `Group By` option will create a new data table for the selected field.

New Group By

Clicking the `New Group By` option will create a new data table for the selected field.

Numeric Ops

If the value you clicked is numeric, then the `Numeric Ops` sub-menu allows you to choose operations like less than, less than or equal, greater than, greater than or equal, or Between. Choosing the Between option displays a window so that you can specify a range of values.

Clipboard

The `Clipboard` sub-menu has several options that allow you to copy selected data to your clipboard in different ways.

Actions

The `Actions` sub-menu has several different options. Please note that some of these actions will only display on the Actions menu if you click on a specific log type.

- Clicking the `Hunt` option will start a new search for the selected value and will give you a good overview of what types of data are available for that indicator.
- Clicking the `Add to Case` option will add an observable to a new or existing case.
- Clicking the `Correlate` option will find related logs based on Community ID, uid, fuid, etc.
- Clicking the `PCAP` option will pivot to the PCAP interface to retrieve full packet capture for the selected stream. This option will only appear if you click on a log that contains source IP, source port, destination IP, destination port, etc.
- Clicking the `Google` option will search Google for the selected value.
- Clicking the `VirusTotal` option will search VirusTotal for the selected value.
- Clicking the `Process Info` option will show all logs that include this process's entity_id in the `process.entity_id` field. This option will only appear if you click on a log that contains the `process.entity_id` field.
- Clicking the `Process and Child Info` option will show all logs that include this process's entity_id in either the `process.entity_id` or `process.parent.entity_id` fields. Depending on the process, this may show the same logs as the `Process Info` option or it may show more. This option will only appear if you click on a log that contains the `process.entity_id` field.
- Clicking the `Process All Info` option will show all logs that include this process's entity_id in any field. Depending on the process, this may show the same logs as the `Process and Child Info` option or it may show more. This option will only appear if you click on a log that contains the `process.entity_id` field.
- Clicking the `Process Ancestors` option will show all parent processes for the selected process. This option will only appear if you click on a log that contains the `process.Ext.ancestry` field.

If you'd like to add your own custom actions, see the SOC Customization section.

6.3.9 OQL

Onion Query Language (OQL) starts with standard Lucene query syntax and then allows you to add optional segments that control what Dashboards does with the results from the query.

sortby

The `sortby` segment can be added to the end of a hunt query. This can help ensure that you see the most recent data, for example, when sorting by descending timestamp. Otherwise, if the search yields a dataset larger than the X Limit size selected in the UI then you will only get the first X records and then those will be sorted on the web browser.

You can specify one field to sort by or multiple fields separated by spaces. The default order is descending but if you want to force the sort order to be ascending you can add the optional caret (^) symbol to the end of the field name.

| sortby some.field another.field^

groupby

The `groupby` segment tells Dashboards to group by (aggregate) a particular field. So, for example, if you want to group by destination IP address, you can add the following to your search:

| groupby destination.ip

The `groupby` segment supports multiple aggregations so you can add more fields that you want to group by, separating those fields with spaces. For example, to group by destination IP address and then destination port in the same data table, you could use:

| groupby destination.ip destination.port

OQL supports multiple `groupby` segments so if you wanted each of those fields to have their own independent data tables, you could do:

| groupby destination.ip | groupby destination.port

In addition to rendering standard data tables, you can optionally render the data as a pie chart, bar chart, or sankey diagram.

- The pie chart is specified using the `-pie` option:

| groupby -pie destination.ip

- The bar chart is specified using the `-bar` option:

| groupby -bar destination.ip

- The sankey diagram is specified using the `-sankey` option, but keep in mind that this requires at least two fields:

| groupby -sankey destination.ip destination.port

The `-maximize` option will maximize the table or chart to fill the pane. After viewing the maximized result, you can press the Esc key to return to normal view.

By default, grouping by a particular field won't show any values if that field is missing. If you would like to include missing values, you can add an asterisk after the field name. For example, suppose you want to look for non-HTTP traffic on port 80 using a query like `event.dataset:conn AND destination.port:80 | groupby network.protocol destination.port`. If there was non-HTTP traffic on port 80, the `network.protocol` field may be null and so this query would only return port 80 traffic identified as HTTP. To fix this, add the asterisk after the `network.protocol`:

event.dataset:conn AND destination.port:80 | groupby network.protocol* destination.port

Please note that adding the asterisk to a non-string field may not work as expected. As an alternative, you may be able to use the asterisk with the equivalent `keyword` field if it is available. For example, `source.geo.ip*` may return 0 results, or a query failure error, but `source.geo.ip.keyword*` may work as expected.

table

The `table` segment tells Dashboards to include the given field names as columns in the Events table at the bottom of the dashboards screen. The columns will be ordered within the Events table following the same order used in the `| table xxx yyy zzz` segment. When no `table` segment is provided in the query, Dashboards will analyze the `event.dataset` and `event.module` values of the query results to determine which default columns would be most appropriate to represent those events. Those default columns are defined in the SOC Configuration.

Examples:

event.dataset:conn | table event.module source.ip source.protocol

Or, combined with other segments:

event.dataset:conn | groupby event.module | groupby destination.ip | sortby source.port | table event.module source.ip source.port source.protocol

> **Note**
>
> Only one `table` segment is currently supported in OQL. If multiple are provided in the query only one will be used, and the unused segments may be automatically removed.

Sankey Diagram Recursion

There's a known limitation with Sankey diagrams where the diagram is unable to render all data when multiple fields of the diagram contain the same value. This causes a recursion issue. For example, this can occur if using an OQL query of `* | groupby -sankey source.ip destination.ip` and the included events have a specific IP appearing in both the `source.ip` and `destination.ip` fields. SOC will attempt to prevent the recursion issue by omitting any data that introduces recursion. This can result in some diagrams showing partial data on the diagram, and when this occurs the Sankey diagram will have the phrase `(partial)` appended to the title. In rare scenarios, it's possible for the diagram to be completely blank, such as if all data results have the same value in each field. Following the example mentioned above, this could happen if the `source.ip` and `destination.ip` were always equal.

6.4 Hunt

Security Onion Console includes a Hunt interface which is similar to our Dashboards interface but is tuned more for threat hunting.

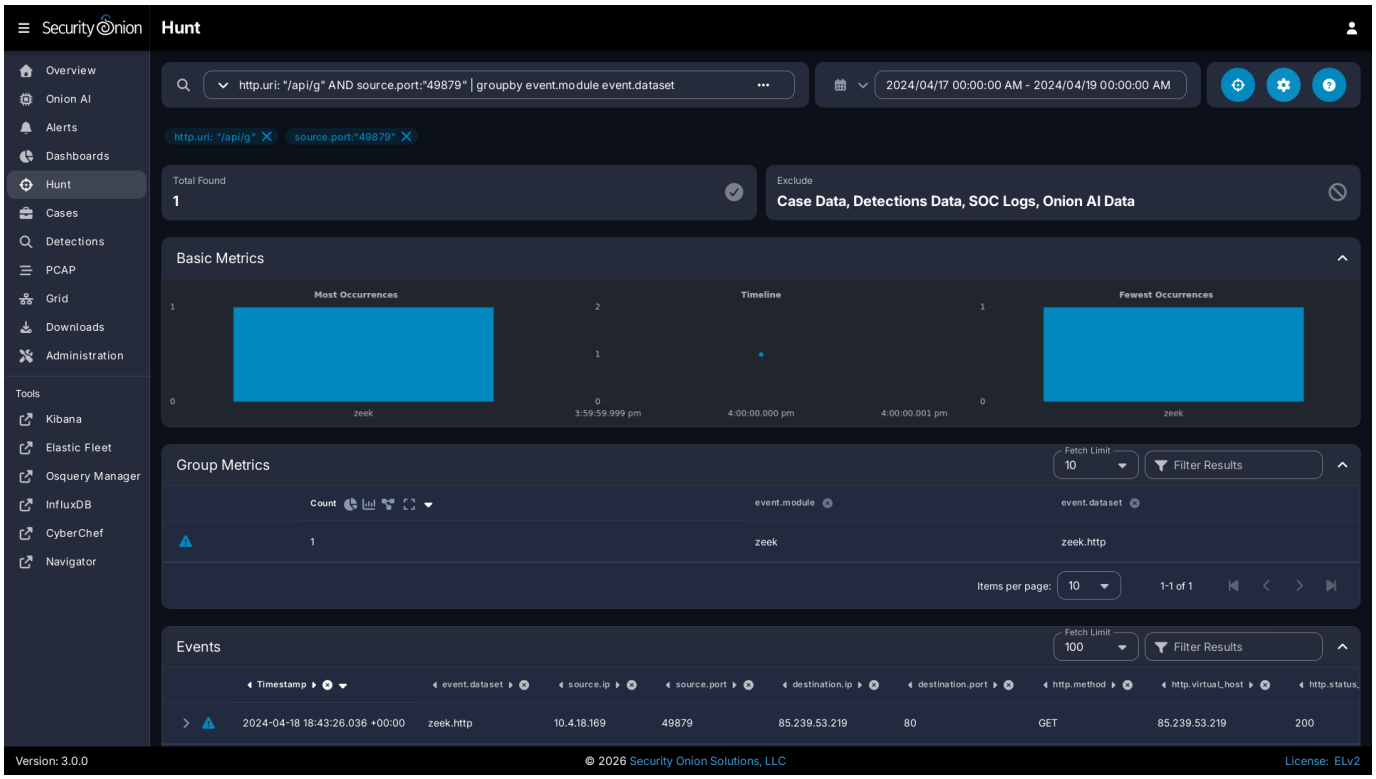

The main difference between Hunt and Dashboards is that Hunt's default queries are more focused than the overview queries in Dashboards. A second difference is that most of the default Dashboards queries display a separate table for each aggregated field, whereas many of the default queries in Hunt aggregate multiple fields in a single table which can be beneficial when hunting for more obscure activity.

Other than these two differences, Hunt and Dashboards are very similar, so for more information please see the Dashboards section.

6.5 Cases

Security Onion Console includes our Cases interface for case management. It allows you to escalate logs from Alerts, Dashboards, and Hunt, and then assign analysts, add comments and attachments, and track observables.

On a new deployment, Cases will be empty until you create a new case. Once you have one or more cases, you can use the main Cases page to get an overview of all cases.

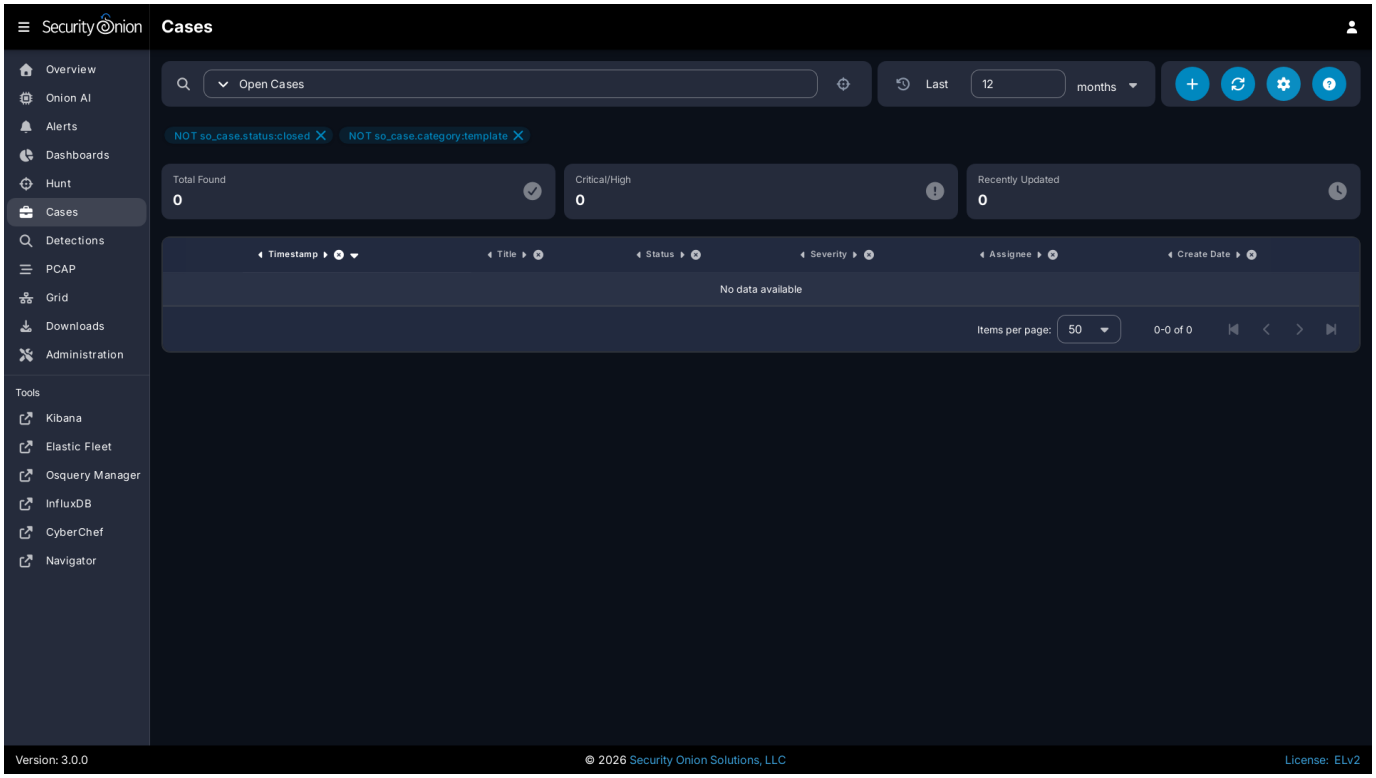

Note

Check out our Cases video at https://youtu.be/y_kr_hrtqVc!

6.5.1 Options

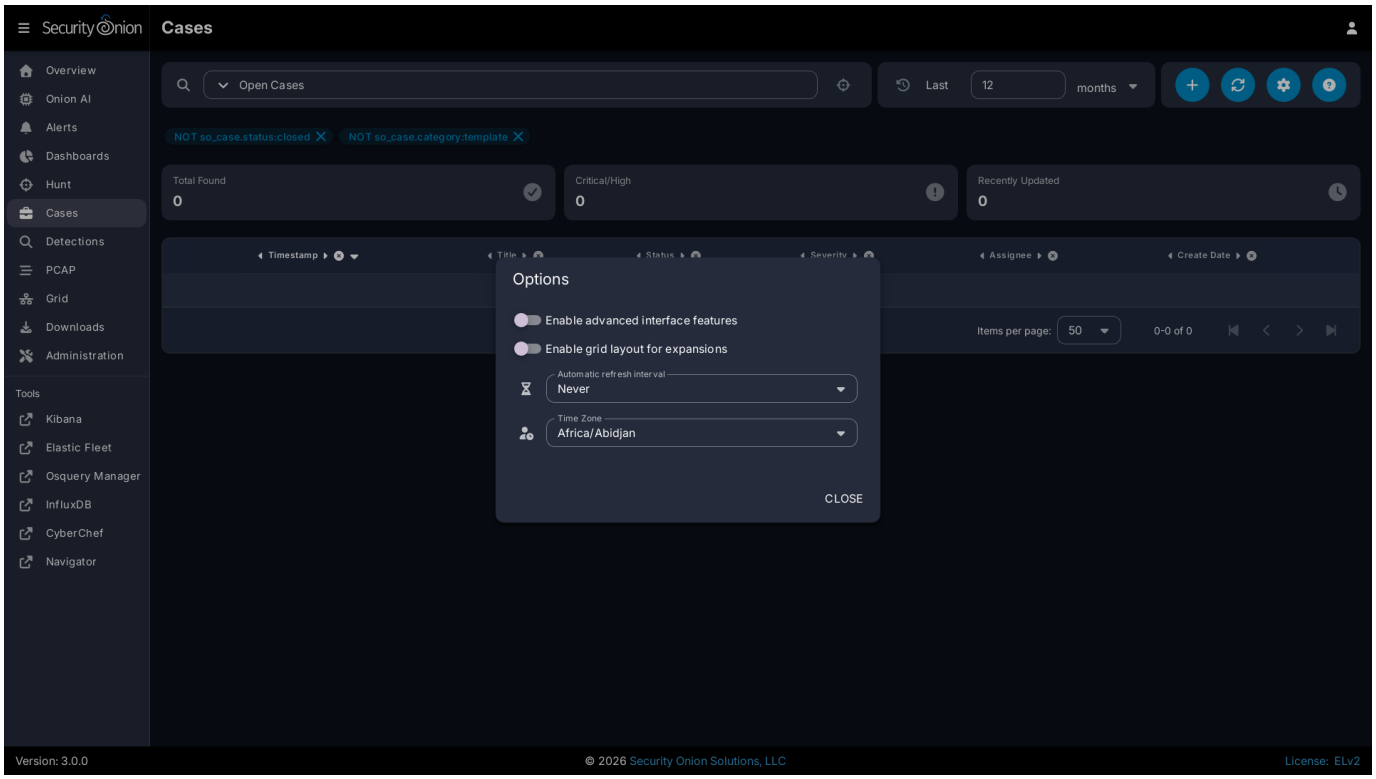

Starting at the top of the main Cases page, the Options menu allows you to set options such as Automatic Refresh Interval and Time Zone.

There is also a toggle labeled `Enable advanced interface features`. If you enable this option, then the interface will show more advanced features similar to Dashboards and Hunt.

6.5.2 Query Bar

The query bar defaults to Open Cases. Clicking the drop-down box reveals other options such as Closed Cases, My Open Cases, My Closed Cases, and Templates. If you want to send your current query to Hunt, you can click the crosshairs icon to the right of the query bar.

Under the query bar, you'll notice colored bubbles that represent the individual components of the query and the fields to group by. If you want to remove part of the query, you can click the X in the corresponding bubble to remove it and run a new search.

If you would like to save your own personal queries, you can bookmark them in your browser. If you would like to customize the default queries for all users, please see the SOC Customization section.

6.5.3 Time Picker

The time picker is to the right of the query bar. By default, Cases searches the last 12 months. If you want to search a different time frame, you can change it here.

6.5.4 Data Table

The remainder of the main Cases page is a data table that shows a high level overview of the cases matching the current search criteria.

- Clicking the table headers allows you to sort ascending or descending.
- Clicking a value in the table brings up a context menu of actions for that value. This allows you to refine your existing search, start a new search, or even pivot to external sites like Google and VirusTotal.
- You can adjust the Rows per page setting in the bottom right and use the left and right arrow icons to page through the table.
- When you click the arrow to expand a row in the data table, it will show the high level fields from that case. Field names are shown on the left and field values on the right. When looking at the field names, there is an icon to the left that will add that field to the `groupby` section of your query. You can click on values on the right to bring up the context menu to refine your search.
- To the right of the arrow is a binoculars icon. Clicking this will display the full case including the Comments, Attachments, Observables, Events, and History tabs.

6.5.5 Creating a New Case

To create a new case, click the + icon and then fill out the Title and Description and optionally the fields on the right side including Assignee, Status, Severity, Priority, TLP, PAP, Category, and Tags. Clicking the fields on the right side reveals drop-down boxes with standard options. The Assignee field will only list user accounts that are currently enabled.

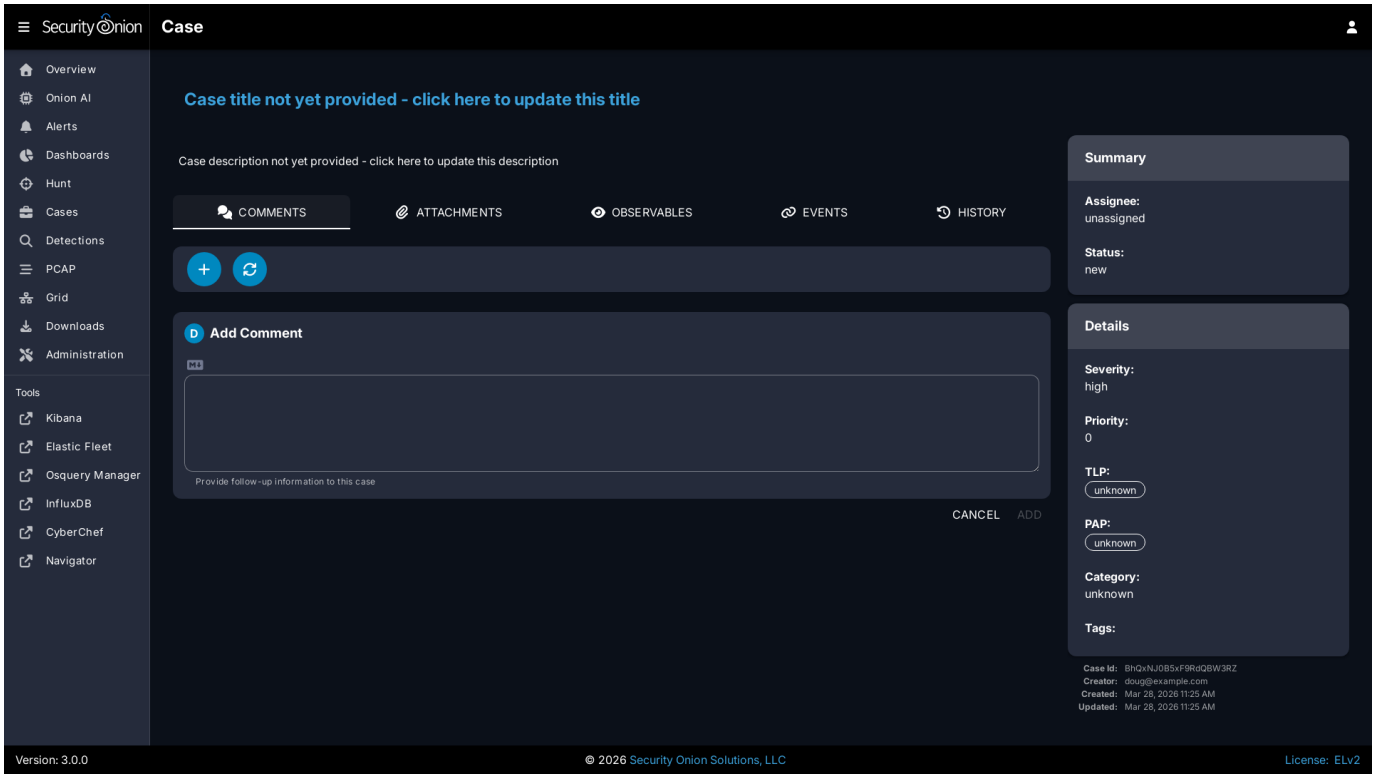

Alternatively, if you find events of interest in Alerts, Dashboards, or Hunt, you can escalate directly to Cases using the escalate button (blue triangle with exclamation point). Clicking the escalate button will escalate the data from the row as it is displayed. This means that if you're looking at an aggregated view, you will get limited details in the resulting escalated case. If you want more details to be included in the case, then first drill into the aggregation and escalate one of the individual items in that aggregation. Once you click the escalate button, you can choose to escalate to a new case or an existing case.

6.5.6 Comments

On the Comments tab, you can add comments about the case. The Comments field uses markdown syntax and you can read more about that at https://www.markdownguide.org/cheat-sheet/.

If you've enabled Security Onion Pro, then when adding a comment you can also specify how many hours you spent working on that activity. You can then see the total time spent by all analysts in the Summary in the upper-right corner.

6.5.7 Attachments

On the Attachments tab, you can upload attachments. For each attachment, you can optionally define TLP and add tags. Cases will automatically generate SHA256, SHA1, and MD5 hash values for each attachment. Buttons next to the hash values allow you to copy the value or add it as an observable.

6.5.8 Observables

On the Observables tab, you can track observables like IP addresses, domain names, hashes, etc. You can add observables directly on this tab or you can add them from the Events tab as well.

You can add multiple observables of the same type by selecting the option labeled `Enable this checkbox to have a separate observable added for each line of the provided value above`.

For each observable, you can click the icon on the far left of the row to drill into the observable and see more information about it. To the right of that is the the hunt icon which will start a new hunt for the observable. Clicking the lightning bolt icon will analyze the observable (see the Analyzers section later).

You can also add observables directly from Alerts, Dashboards, or Hunt. Click the observable and select the `Add to Case` option. You'll then have the option of adding the observable to a new case or an existing case.

6.5.9 Events

On the Events tab, you can see any events that have been escalated to the case. This could be Suricata alerts, network metadata from Suricata or Zeek, or endpoint logs.

For each event, you can click the icon on the far left of the row to drill in and see all the fields included in that event.

If you find something that you would like to track as an Observable, you can click the eye icon on the far left of the row to add it to the Observables tab. It will attempt to automatically identify well known data types such as IP addresses.

To the right of the eye icon is a Hunt icon that can be used to start a new hunt for that particular value.

6.5.10 History

On the History tab, you can see the history of the case itself, including any changes made by each user. For each row of history, you can click the icon on the far left of the row to drill in and see more information.

6.5.11 Data

Cases data is stored in Elasticsearch. You can view it in Dashboards or Hunt by clicking the Options menu and disabling the `Exclude case data` option. You can then search the `so-case` index with the following query:

```
_index:"*:so-case"
```

You can also use this query in Kibana.

You might want to backup this data as described in the backup section.

6.5.12 Analyzers

We have included analyzers which allow you to quickly gather context around an observable.

> **Note**
>
> Check out our Analyzers video at https://youtu.be/99LXr7UmtKI!

Supported Analyzers and Data Types

The following is a summary of the built-in analyzers and their supported data types:

Name	Domain	EML	Hash	IP	Mail	Other
Alienvault OTX	✓	✓				
Echotrail					✓	
Elasticsearch	✓	✓	✓	✓	✓	✓
EmailRep				✓		
Greynoise				✓		
LocalFile	✓	✓	✓			✓
Malwarebazaar			✓			
Malware Hash Registry			✓			
Pulsedive	✓	✓	✓			
Spamhaus				✓		
Sublime Platform		✓				
Threatfox	✓	✓	✓			
Urlhaus						
Urlscan						
Virustotal	✓	✓	✓			
WhoisLookup	✓					

> **Note**
>
> The `malwarehashregistry` analyzer is no longer working. This is due to a stale third-party library that is incompatible with the latest Python version. See #13571

Running Analyzers

To enqueue an analyzer job, click the lightning bolt icon on the left side of the observable menu. All configured analyzers supporting the observable's data type will then run and return their analysis.

> **Note**
>
> Observable values must be formatted to correctly match the observable type in order for analyzers to properly execute against them. For example, an IP observable type should not contain more than one IP address.

Analyzer Output

The collapsed job view for an analyzer will return a summary view of the analysis. Expanding the collapsed row will reveal a more detailed view of the analysis.

> **Warning**
>
> If you try to run the Malware Hash Registry analyzer but it results in a "Name or service not known" error, then it may be a DNS issue. Folks using 8.8.4.4 or 8.8.8.8 as their DNS resolver have reported this issue. A potential workaround is to switch to another DNS resolver like 1.1.1.1.

Configuring Analyzers

Some analyzers require authentication or other details to be configured before use. If analysis is requested for an observable and an analyzer supports that observable type but the analyzer is left unconfigured, then it will not run.

The following analyzers require users to configure authentication or other parameters in order for the analyzer to work correctly:

- AlienVault OTX
- Echotrail
- Elasticsearch
- EmailRep
- GreyNoise
- LocalFile
- Malwarebazaar
- Pulsedive
- Threatfox
- Urlscan
- VirusTotal

To configure an analyzer, navigate to Administration --> Configuration --> sensoroni.

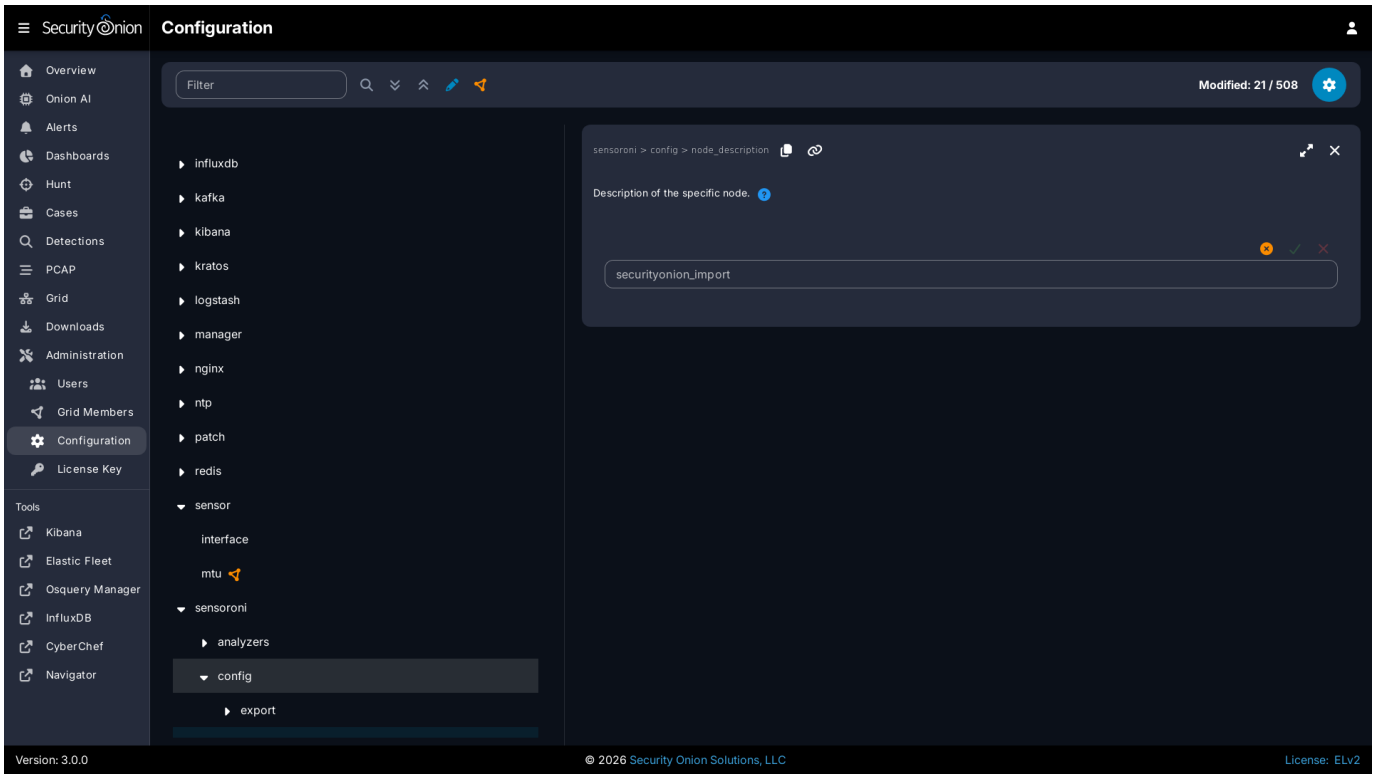

At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option. Then navigate to sensoroni --> analyzers.

Developing Analyzers

If you'd like to develop a custom analyzer, take a look at the developer's guide at https://github.com/Security-Onion-Solutions/securityonion/tree/3/main/salt/sensoroni/files/analyzers.

6.5.13 Templates

SOC can use case templates to auto-populate default values of new cases. A template is itself a case, with its category set to `template`. To utilize that template case, the new case should specify the template case ID in the `template` field of the case object.

SOC automatically populates new case template fields with the value stored in the `rule.case_template` field of the alert being escalated. This allows for specific templates to be assigned to certain detection rules. For example, if alerts triggered from a certain rule are known to require a consistent set of resolution steps then the description of a case template can be prepopulated with that checklist (in markdown format). Then, the backing rule that triggered the alert can have its `case_template` field set to that case template ID.

6.6 Detections

Security Onion Console includes our Detections interface for managing all of your rules:

- NIDS rules that get loaded into Suricata
- Sigma rules that get loaded into ElastAlert
- YARA rules that get loaded into Strelka

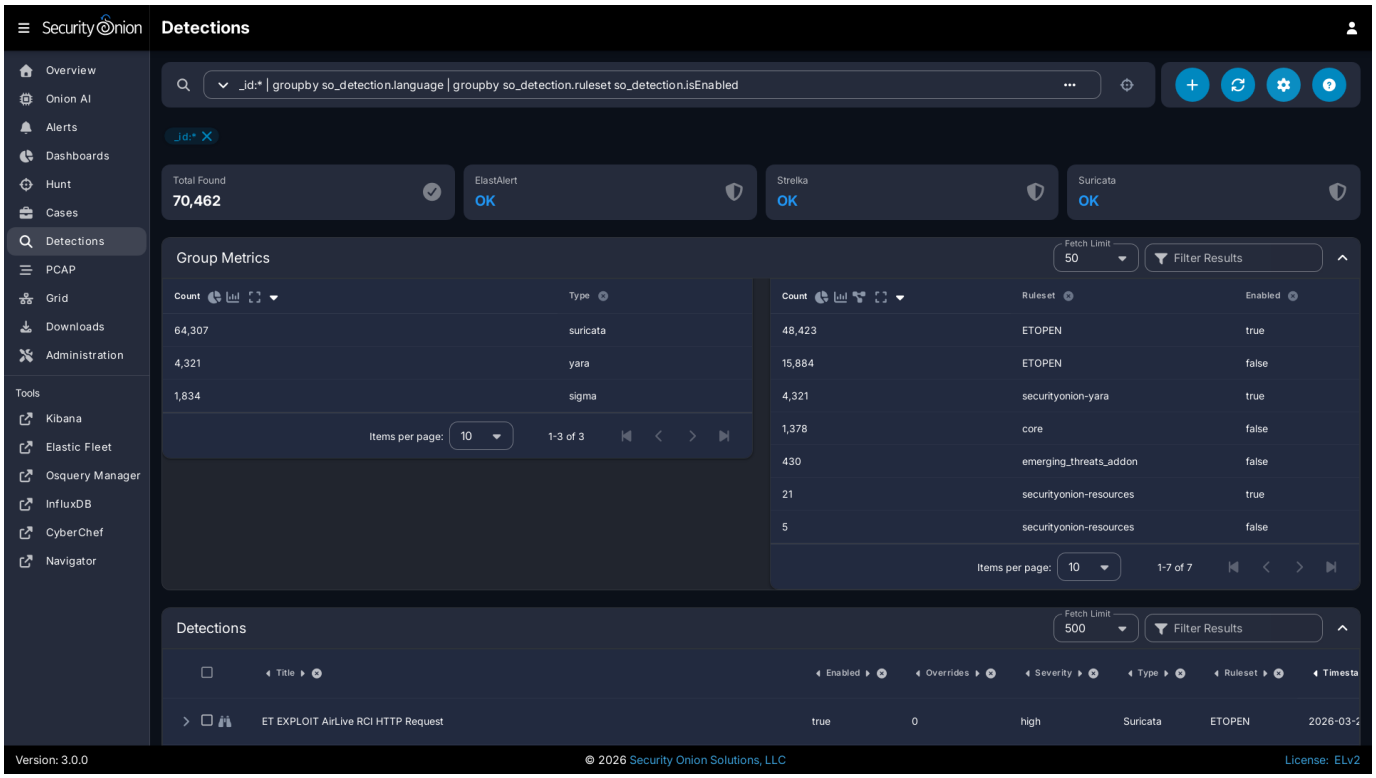

Note

Check out our Detections video at https://youtu.be/DelAmqtU2hg!

6.6.1 Rule Engine Status

The upper-right corner shows a count of detections that matched the search query. To the left of that is a status indicator for each of the detection engines. The status can show whether a sync is in process, as well as whether the engine has detected errors.

Here is the list of possible status messages and what they mean:

- **Pending**: The browser is waiting for the server to send an initial status report.
- **Import Pending**: The import will start once the system stabilizes, usually within twenty minutes.
- **Importing**: The previous version of Security Onion's rules are being imported into the new Detections system. This can take an hour or more on some systems.
- **Migrating**: Rules will be migrated between Security Onion versions following system upgrades. This can take some time if upgrading from a much older version.
- **Migration Failed**: A failure occurred during the migration. The migration will stop on the first error and will not attempt to migrate to newer versions until the issue is resolved.
- **Synchronizing**: A rule synchronization is in progress. This occurs daily, to ensure the Security Onion Grid has the latest rules.
- **Sync Failed**: A failure occurred during the synchronization procedure. The next sync will retry within a few minutes.
- **Sync Blocked**: Rule synchronization is currently blocked.
- **Rule Mismatch**: An integrity check process detected a mismatch between the deployed rules and the enabled rules. The SOC log will note the specific mismatched rules. One possible reason is that Elasticsearch has reached its disk watermark setting and is no longer allowing updates to the Detections indices.
- **OK**: No known issues with the rule engine.

> **Tip**
>
> Clicking the status text will navigate to Hunt and attempt to find related logs. If the status is reporting some kind of failure, then you might want to use Hunt to hone in on things like `integrity check failed` or other errors.

As part of the sync process, Detections checks for duplicates. If duplicates are found, Detections will log information about the duplicate.

The Detections menu option on the left side of the application will show an exclamation mark if there is a recent failure in any of the detection engines. In this situation the web browser tab will also show an exclamation indicator. If no failures are detected, and if any of the detection engines has an import pending or is performing a rule import, synchronization, or migration, then a blue hourglass will appear next to the Detections menu option.

6.6.2 Options

The Options menu allows you to synchronize a particular detection engine such as ElastAlert, Strelka, or Suricata.

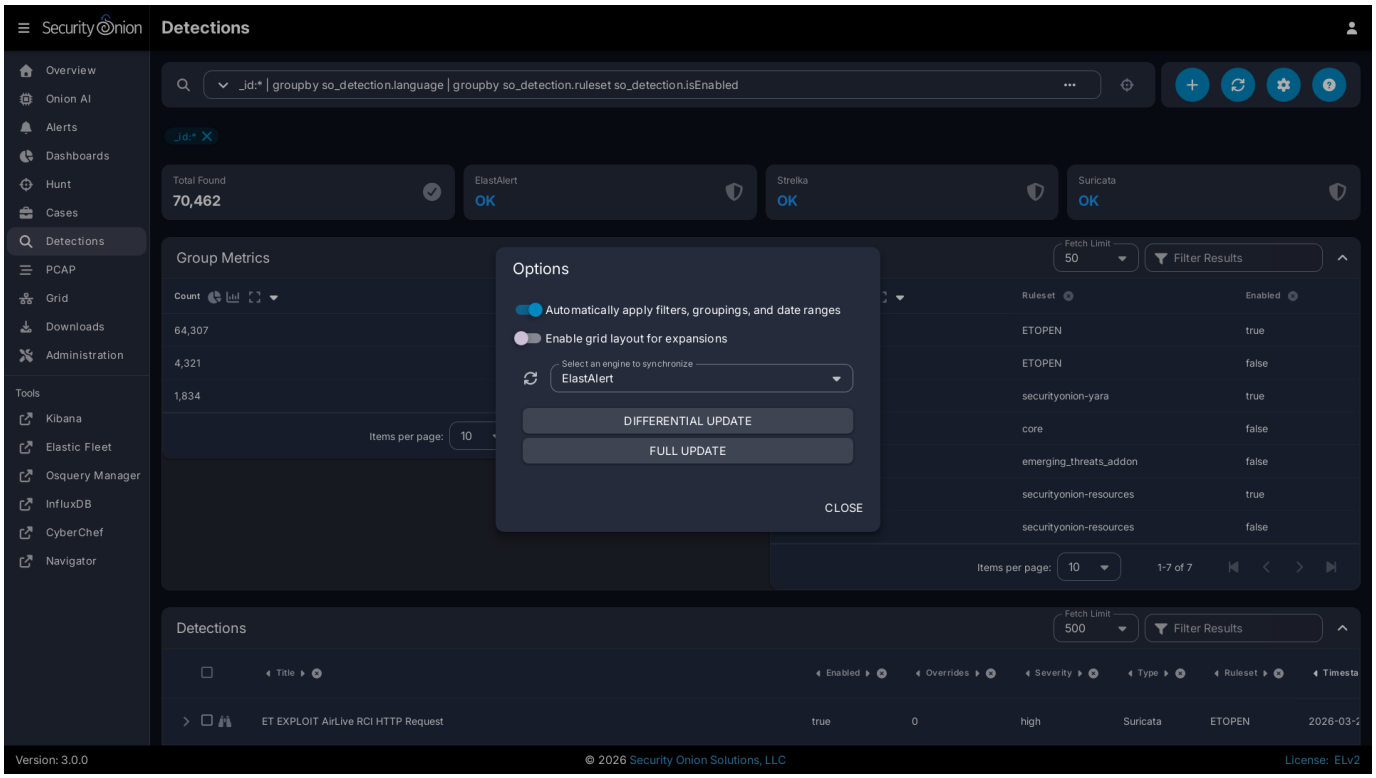

Once you've selected the detection engine that you want to synchronize, you can then click either the `DIFFERENTIAL UPDATE` or `FULL UPDATE` button.

The differential update is a lightweight sync that will skip the thorough sync and comparison of each individual rule. For example, with Suricata it will compute and compare the hash of the source rule list with the hash of the deployed rules, and only if there's a mismatch will it perform the full sync.

A full sync can involve inspecting and comparing individual rules, of which there can be thousands. This more thorough sync can take much longer than the differential sync. Note that each engine has its own unique synchronization process.

6.6.3 Query Bar

The query bar defaults to `All Detections`. Clicking the drop-down box reveals other options such as `Custom Detections`, `All Detections - Enabled`, and `All Detections - Disabled`.

Under the query bar, you'll notice colored bubbles that represent the individual components of the query. If you want to remove part of the query, you can click the X in the corresponding bubble to remove it and run a new search.

If you would like to save your own personal queries, you can bookmark them in your browser. If you would like to customize the default queries for all users, please see the SOC Customization section.

6.6.4 Group Metrics

The Group Metrics section of output consists of one or more data tables or visualizations that allow you to stack (aggregate) arbitrary fields.

6.6.5 Data Table

The remainder of the main Detections page is a data table that shows a high level overview of the detections matching the current search criteria.

- Clicking the table headers allows you to sort ascending or descending.
- Clicking a value in the table brings up a context menu of actions for that value. This allows you to refine your existing search or copy text to the clipboard.
- You can adjust the Rows per page setting in the bottom right and use the left and right arrow icons to page through the table.
- When you click the arrow to expand a row in the data table, it will show the high level fields from that detection. Field names are shown on the left and field values on the right. You can click on values on the right to bring up the context menu to refine your search.
- To the right of the arrow is a binoculars icon. Clicking this will take you to the detection details page.

6.6.6 Detection Details

There are two ways to reach the detail page for an individual detection:

- From the main Detections interface, you can search for the desired detection and click the binoculars icon.
- From the Alerts interface, you can click an alert and then click the `Tune Detection` menu item.

Once you've used one of these methods to reach the detection detail page, you can check the Status field in the upper-right corner and use the slider to enable or disable the detection.

To the left of the Status field are several tabs.

The OVERVIEW tab displays the Summary, References, and Detection Logic for the detection. The Summary field may contain an AI summary of the rule if one is available. These AI summaries are pre-generated so nothing is ever sent from your system to generate this information. That also means that AI summaries only exist for our default rules and will not exist for any of your custom rules.

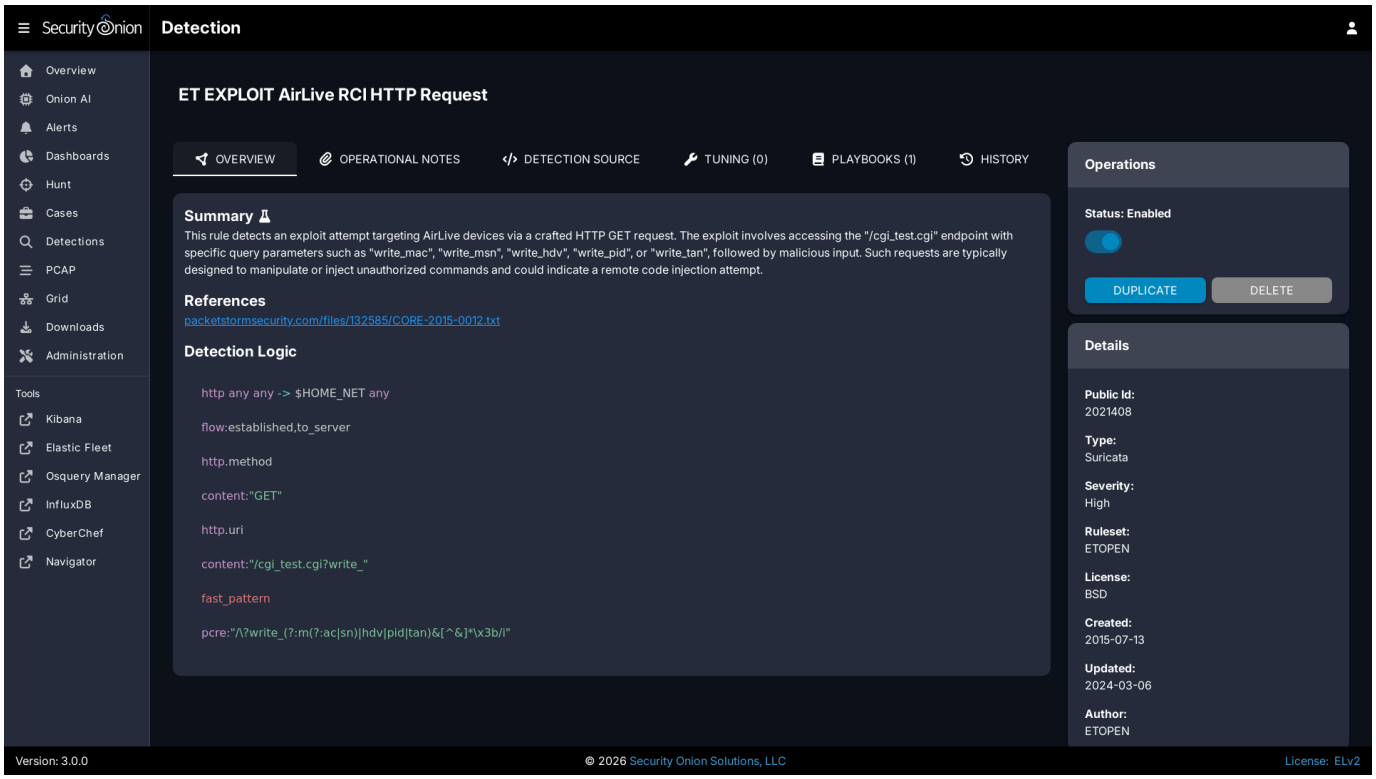

The OPERATIONAL NOTES tab allows you to add your own local notes to the detection in markdown format.

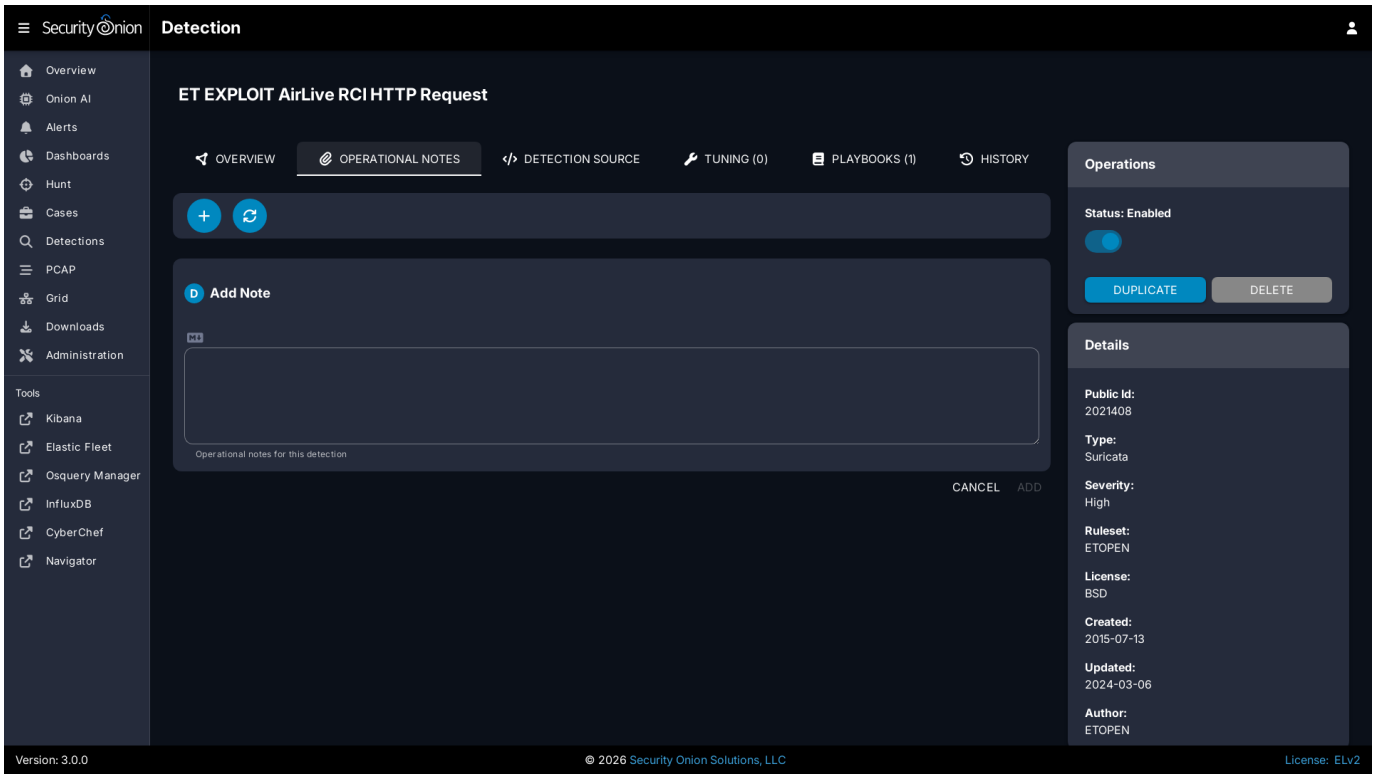

The DETECTION SOURCE tab shows the full content of the detection.

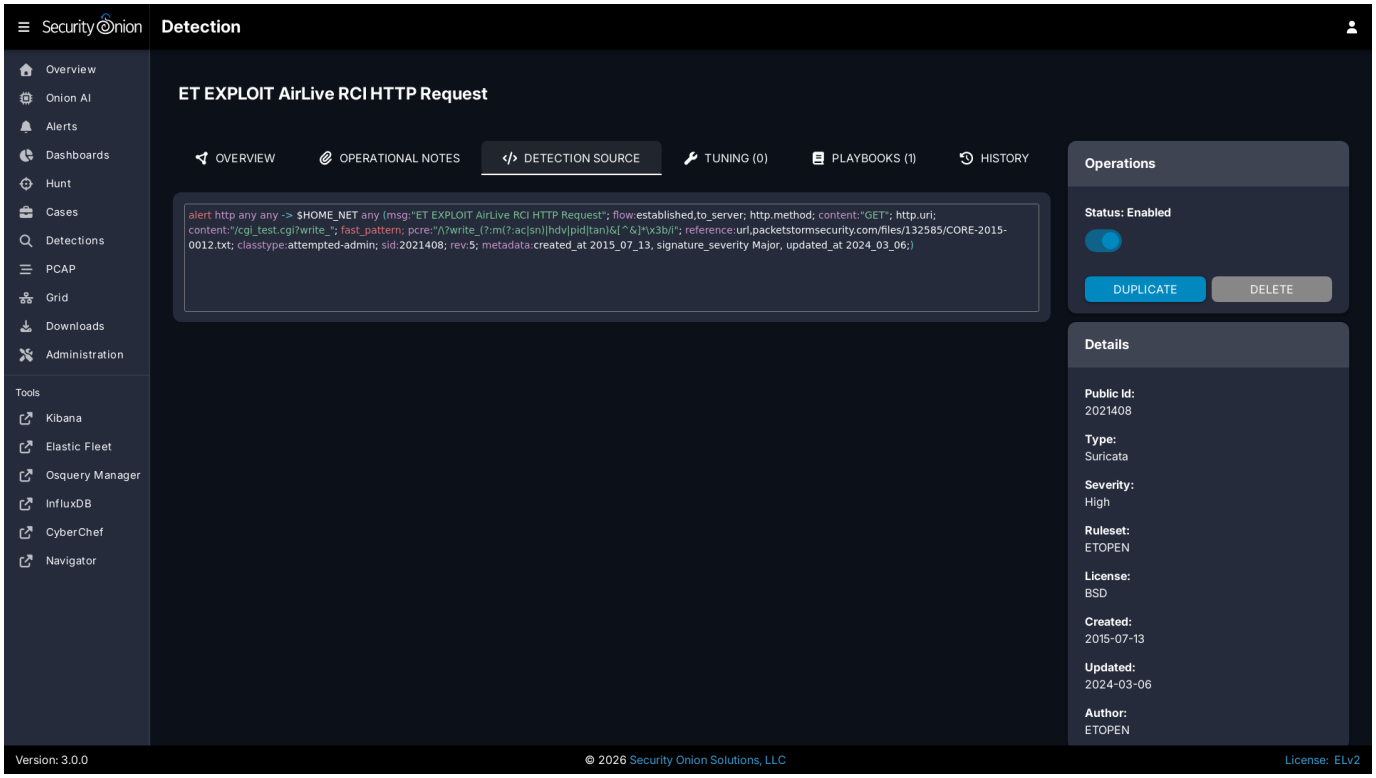

The TUNING tab allows you to tune the detection. For NIDS rules, you can modify, suppress, or threshold. For Sigma rules, you can create a custom filter.

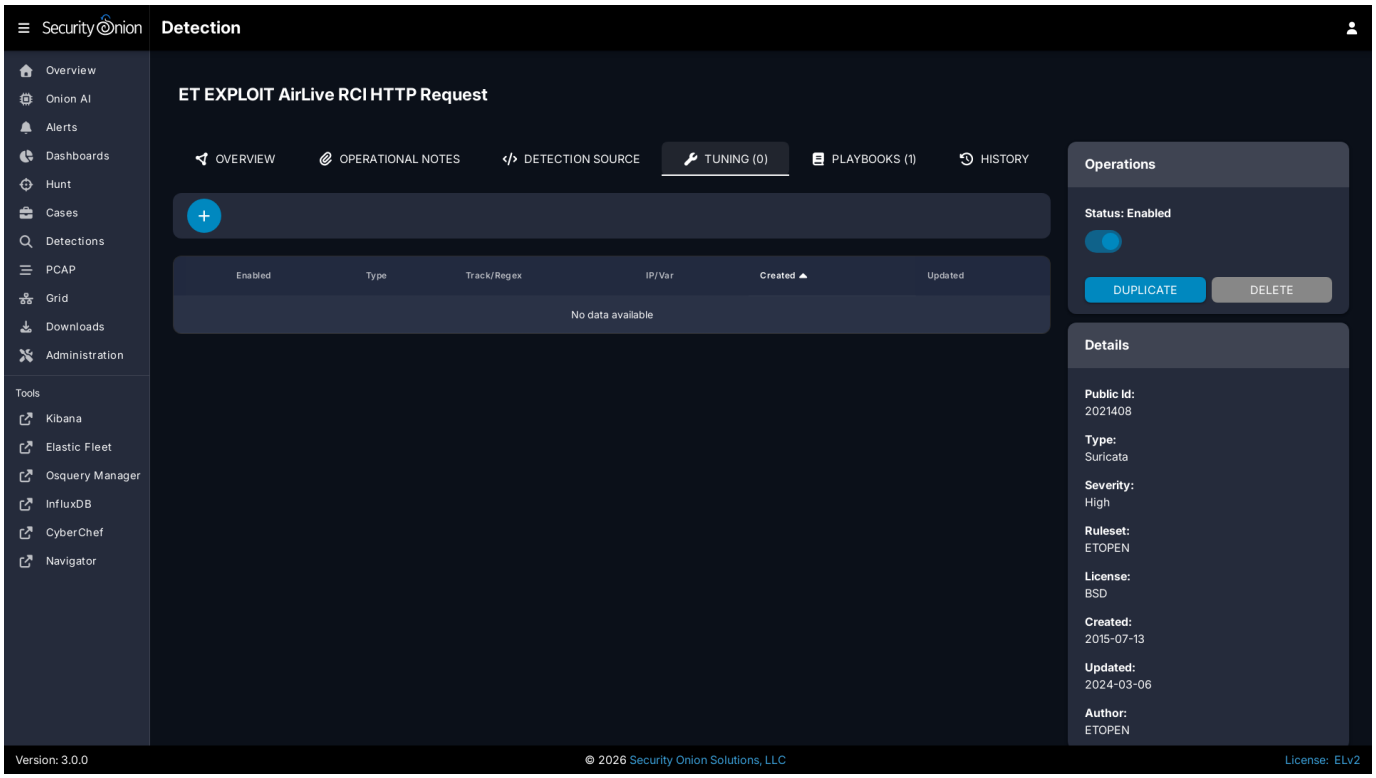

The PLAYBOOKS tab shows any applicable plays for this detection. These playbooks are used for the Guided Analysis tab in Alerts.

Warning

Some playbooks were generated by AI and it's possible that they may not be 100% accurate. Please let us know if you see any issues.

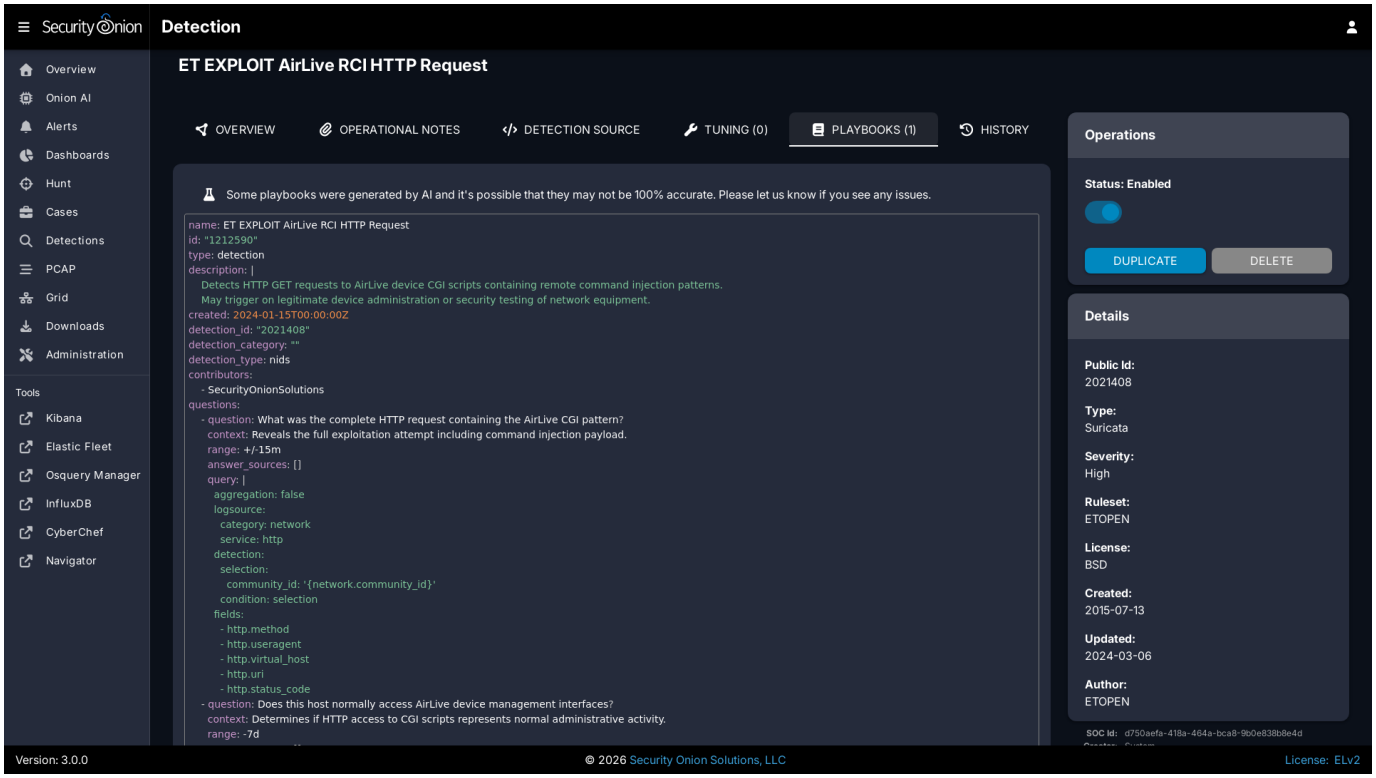

The HISTORY tab shows the history of the detection since it was added to your deployment.

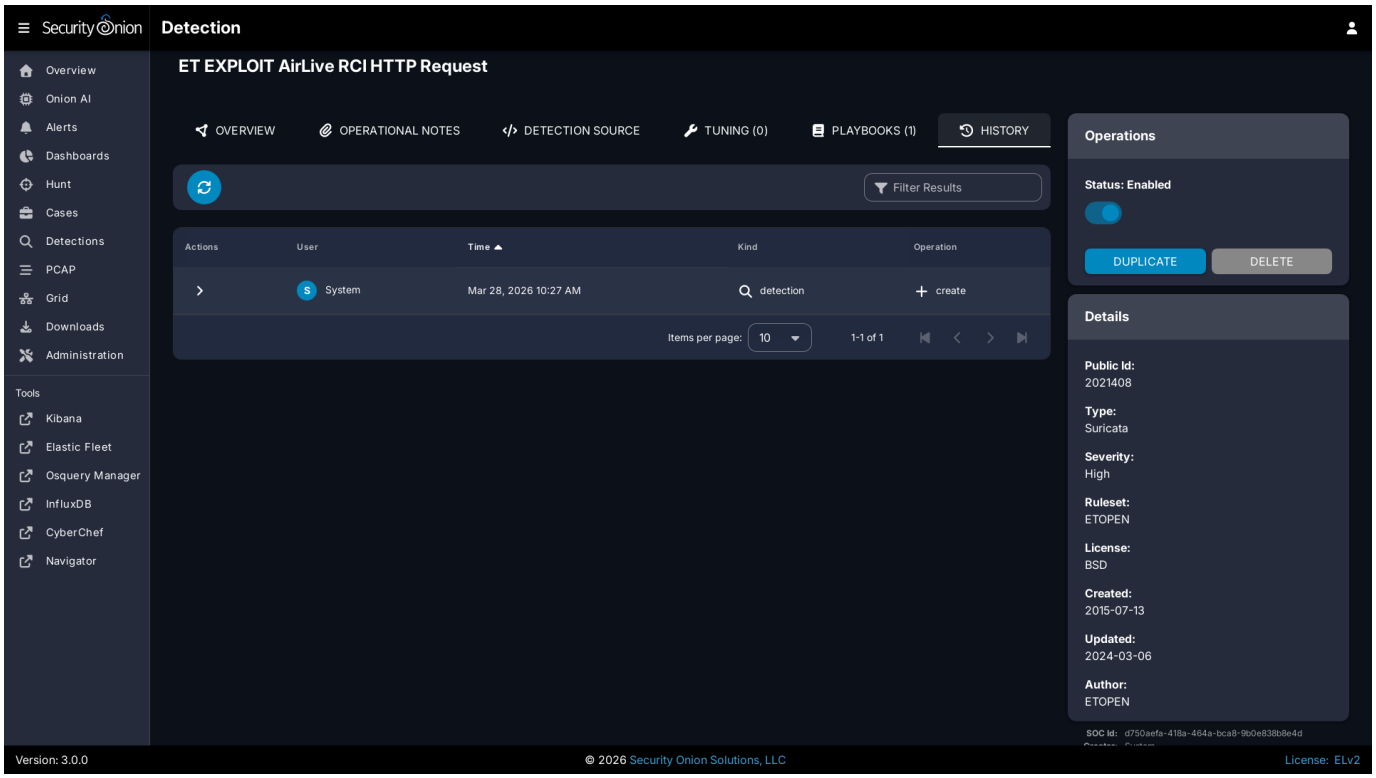

6.6.7 More Information

For more information about managing NIDS rules for Suricata, please see the NIDS section.

For more information about managing Sigma rules for ElastAlert, please see the Sigma section.

For more information about managing YARA rules for Strelka, please see the YARA section.

6.6.8 Technical Background

Detections abstracts the underlying alerting engine and simplifies writing detections for different rule types. Here's what happens behind the scenes.

Enable and Disable (Bulk and Individual) Operations

ElastAlert/Sigma - Immediate change in the UI and on disk

Suricata/NIDS - UI Bulk and Individual: Immediate change in the UI and on disk - Regex: UI and disk change once the Suricata engine syncs

Strelka/YARA - Immediate change in the UI, disk change once the `Strelka` state runs again

Tuning

ElastAlert/Sigma - Immediate change in the UI and on disk

Suricata/NIDS - Immediate change in the UI and on disk

Strelka/YARA - N/A

Ruleset Changes

ElastAlert/Sigma - Sigma Ruleset Packages: UI and disk change once the `SOC` state runs again and the ElastAlert engine syncs - Git repo (https or disk): UI and disk change once the `SOC` state runs again and the ElastAlert engine syncs

Suricata/NIDS - All ruleset sources (ETOPEN, ETPRO, custom URL, local directory): UI and disk change once the Suricata engine syncs

Strelka/YARA - Git repo (https or disk): UI and disk change once the `soc` state runs again and the Strelka engine syncs

6.7 PCAP

Security Onion Console includes a PCAP interface which allows you to access your full packet capture that was written to disk by Suricata.

You can access PCAP in two different ways. The first and most common option is to pivot to PCAP from a particular event in Alerts, Dashboards, or Hunt by choosing the PCAP action on the action menu.

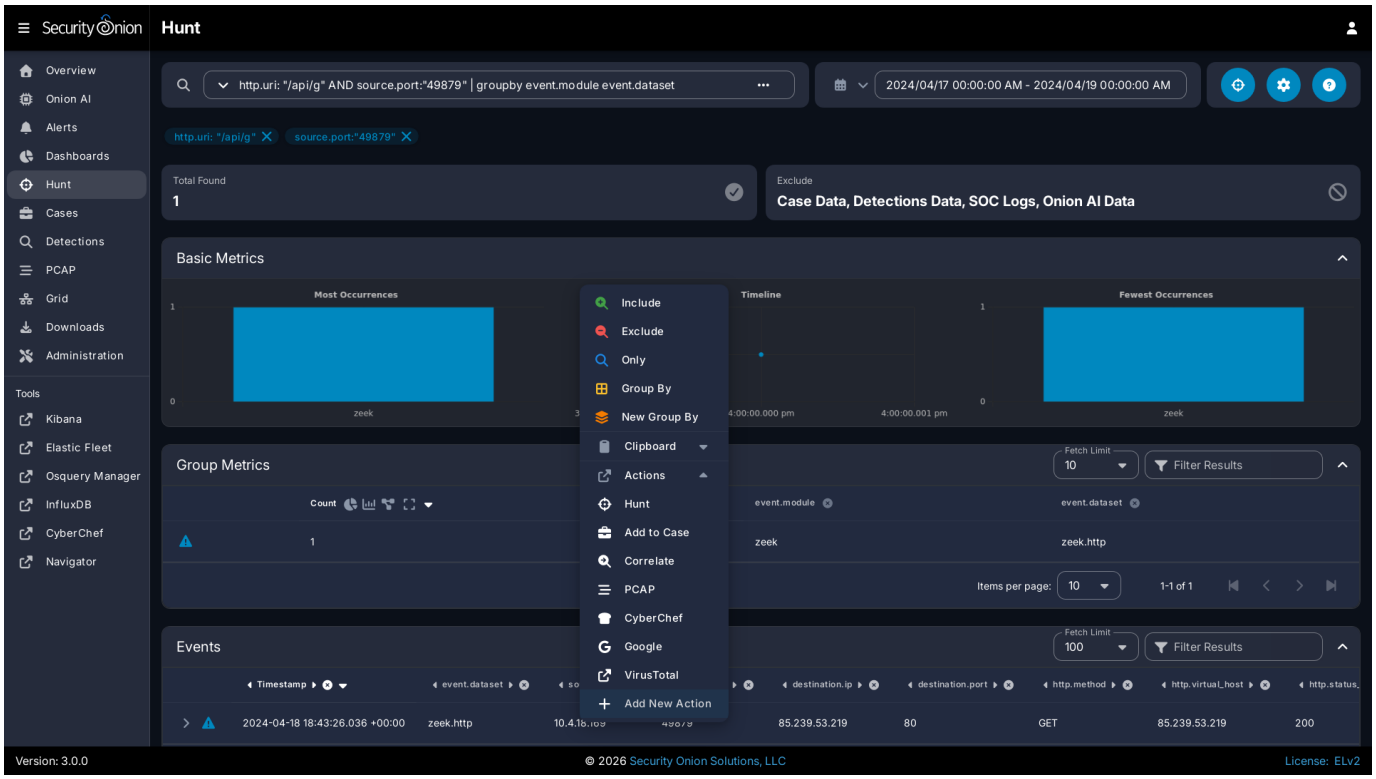

The second and less common option is to go directly to the PCAP interface, click the blue + button, and then put in your search criteria to search for a particular stream.

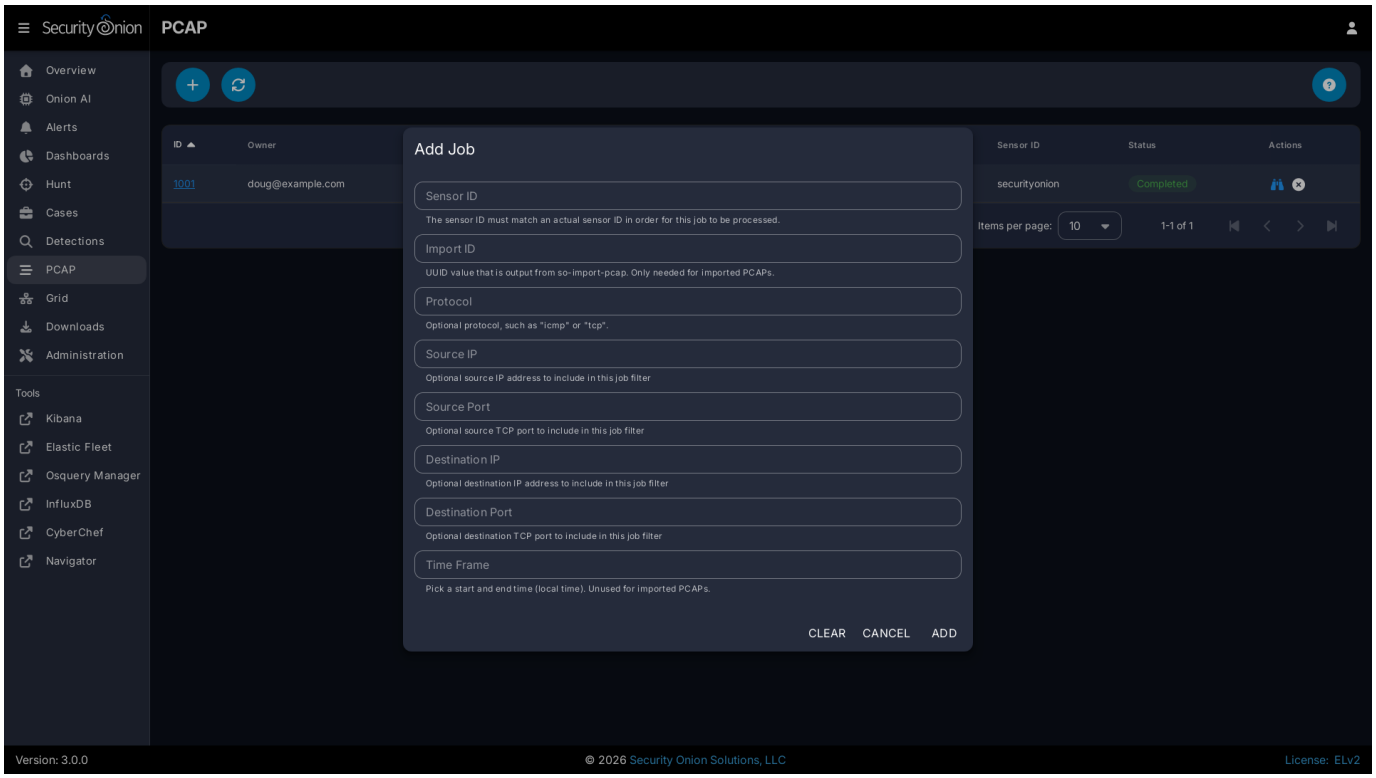

Regardless of which of these two options you choose, Security Onion should then locate the stream and render a high level overview of the packets.

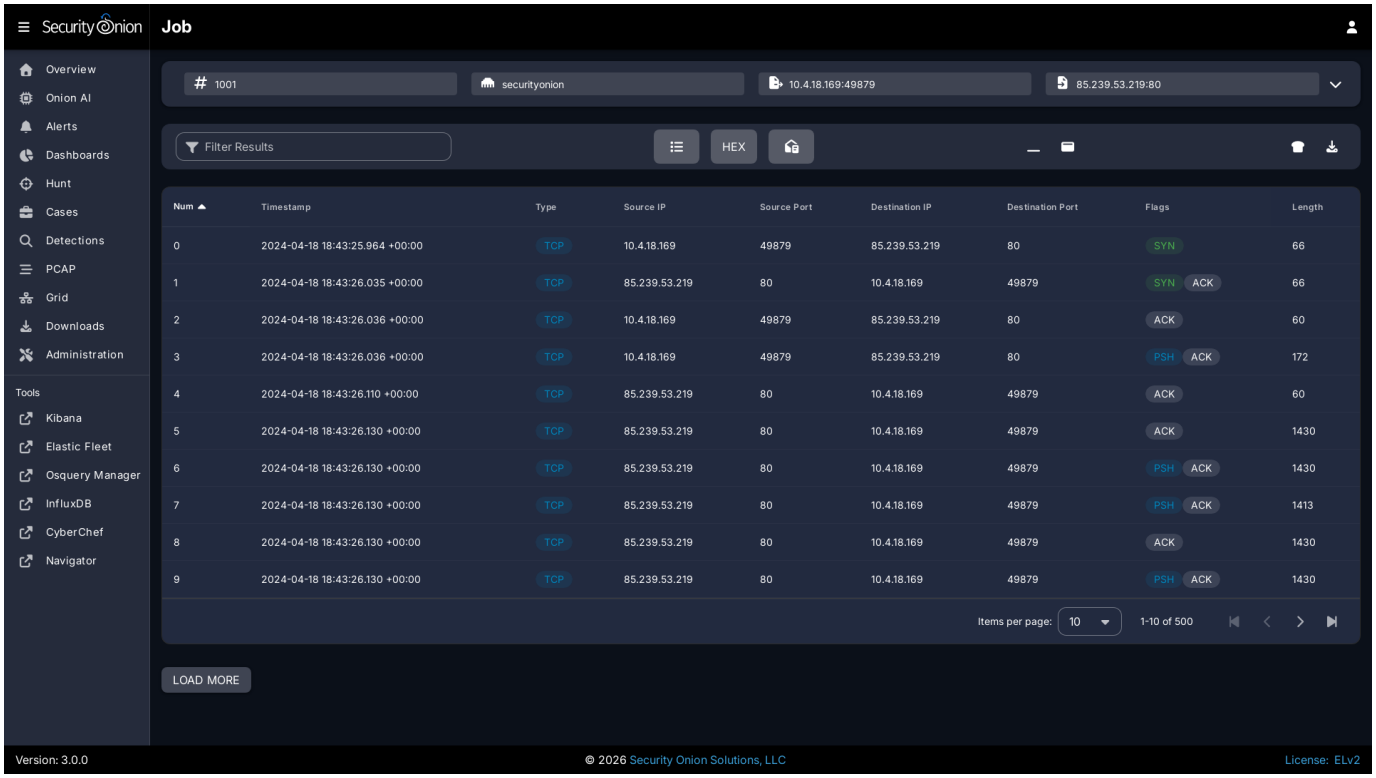

If there are many packets in the stream, you can use the `LOAD MORE` button, `Rows per page` setting, and arrows to navigate through the list of packets.

You can drill into individual rows to see the actual payload data. There are buttons at the top of the table that control what data is displayed in the individual rows. If you disable the `Show all packet data` and `HEX` buttons, then you get an ASCII transcript.

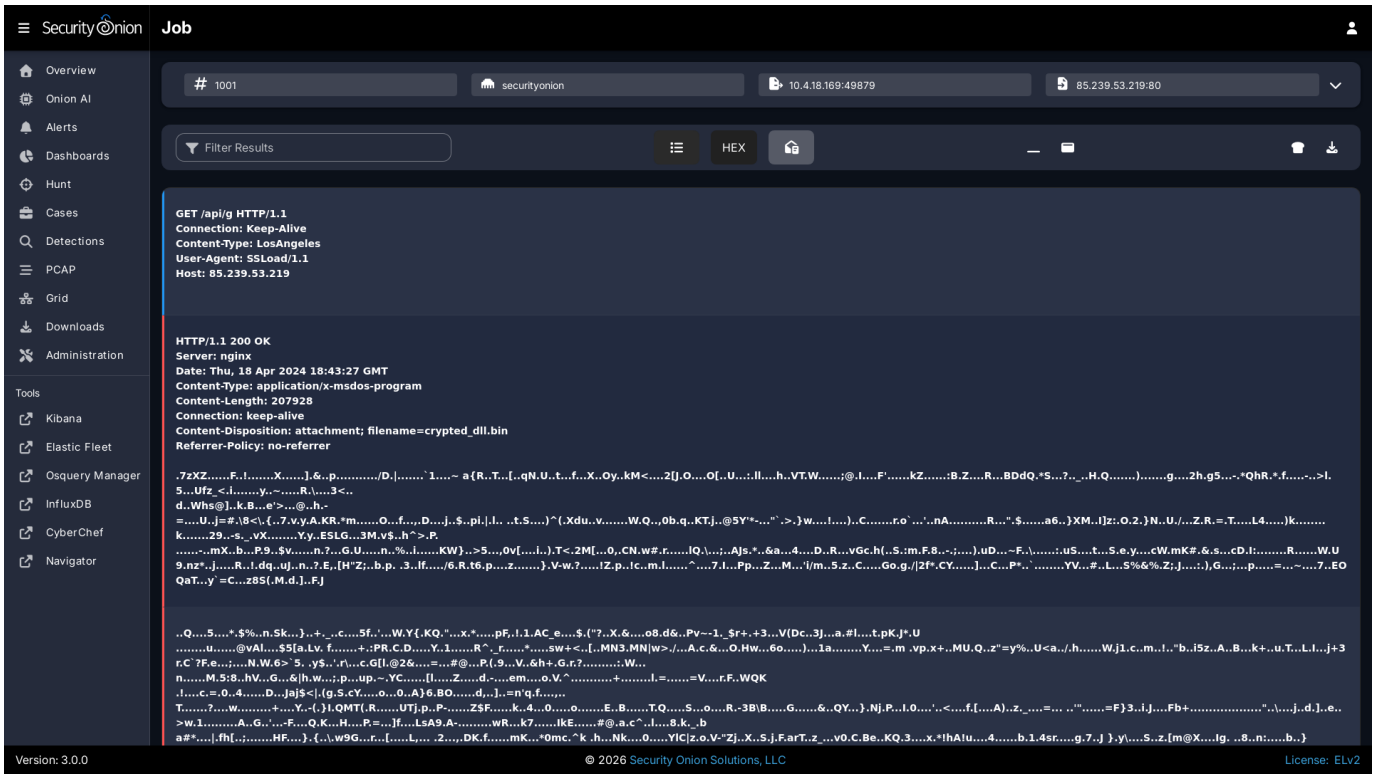

You can select text with your mouse and then use the context menu to send that selected text to CyberChef, Google, or other destinations defined in the actions list.

There are two buttons on the right side of the table header. The first will send all visible packet data to CyberChef. Please note that this only sends packet data that is currently being displayed, so if you are looking at a large stream you may need to use the `LOAD MORE` button to display all packets in the stream. The second button on the far right allows you to download the full PCAP file. You can then open the PCAP file using NetworkMiner, Wireshark, or any other standard libpcap tool. You should typically avoid doing PCAP analysis on your normal desktop environment, so you may want to consider opening the PCAP file in a Security Onion Desktop instance.

Once you've viewed one or more PCAPs, you will see them listed on the main PCAP page.

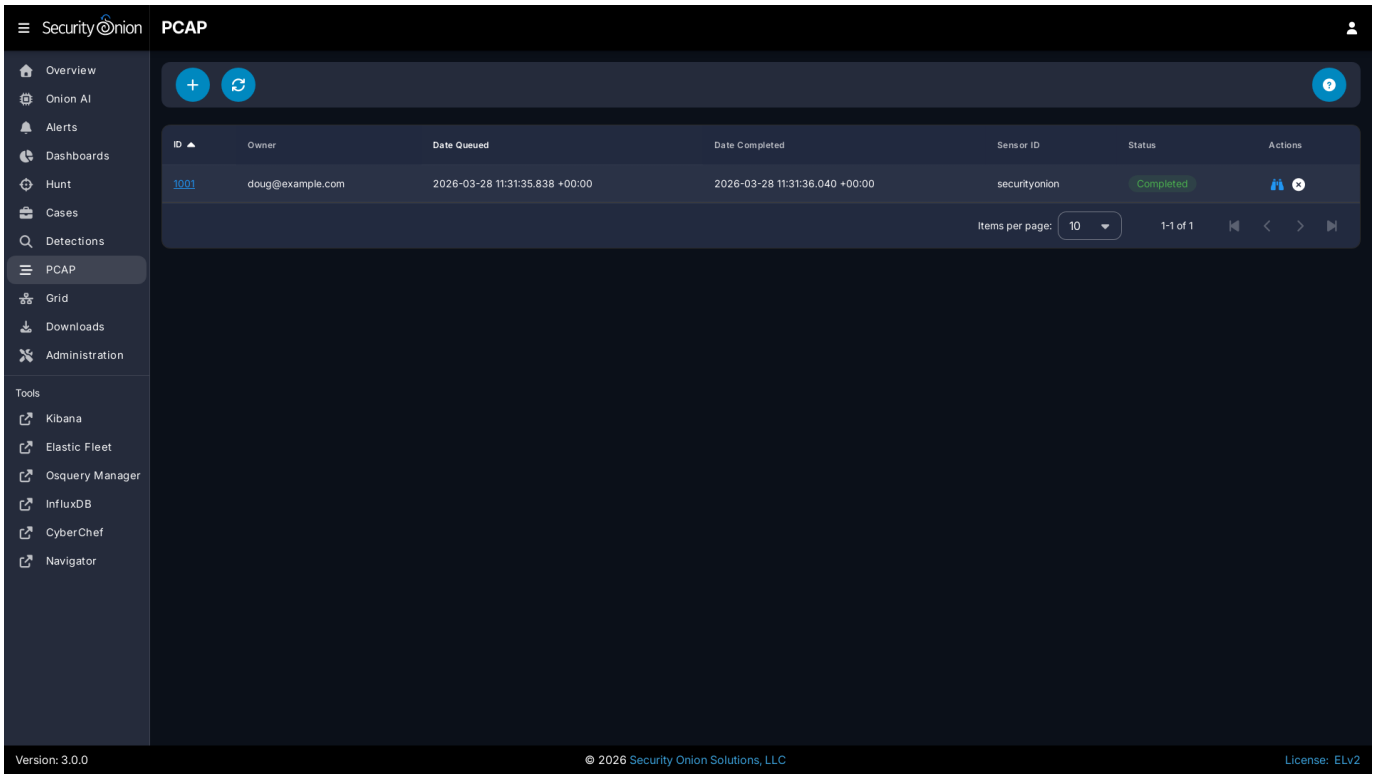

When you are done with a PCAP, you may want to delete it using the `X` button on the far right. This deletes the cached PCAP file saved at `/nsm/soc/jobs/`.

6.7.1 Troubleshooting

If you have trouble retrieving PCAP, here are some things to check:

- Verify that full packet capture is enabled via Suricata.
- Check to see if you have any BPF configuration that may cause Suricata to ignore the traffic.
- Check Grid and verify that all services are running properly.
- Check InfluxDB and verify that PCAP Retention is long enough to include the stream you're looking for.
- Make sure that there is plenty of free space on `/nsm` to carve the stream and write the output to disk.

6.8 Grid

Security Onion Console includes a grid interface which allows you to quickly check the status of all nodes in your grid.

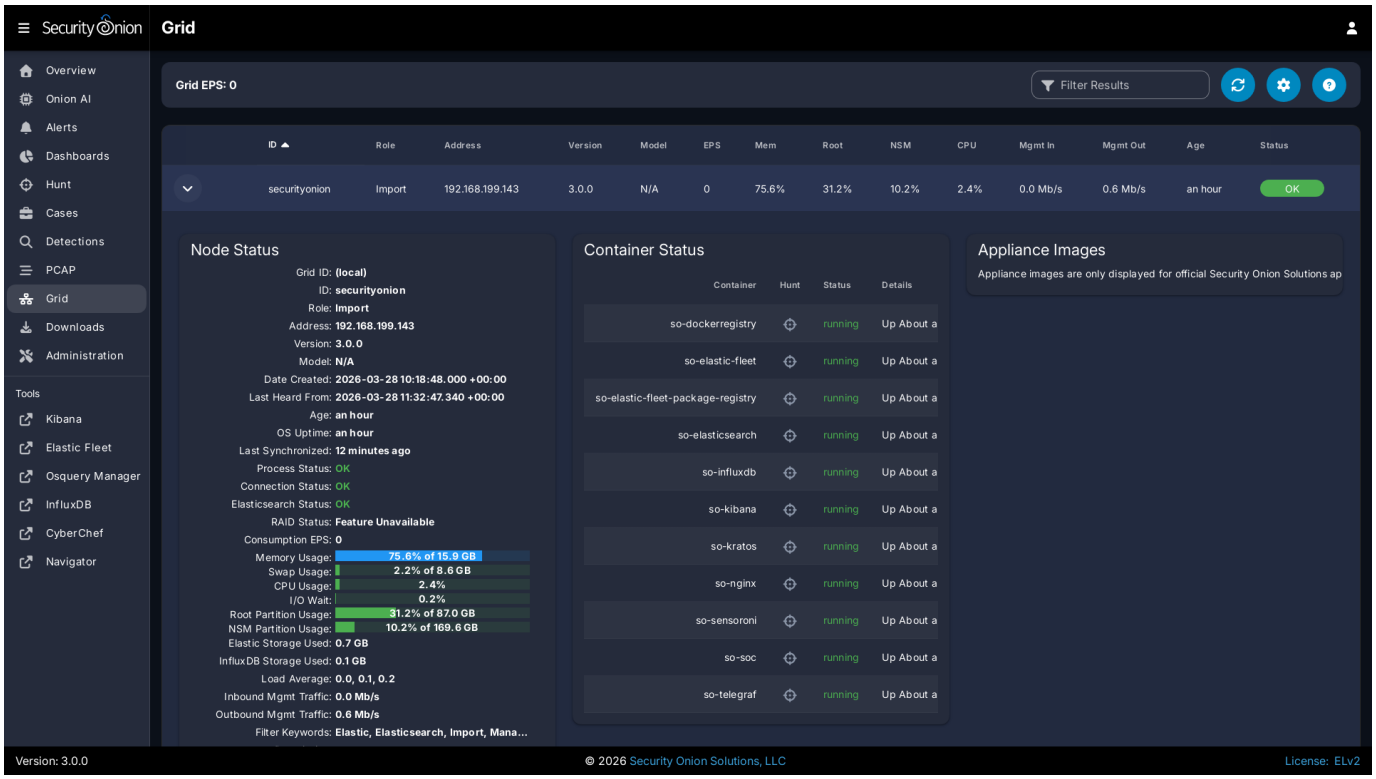

Starting at the top of the page, there is a `Grid EPS` value in the upper-right corner that shows the sum of all `Consumption EPS` measurements in the entire Grid. Below that you will find a list of all nodes in your grid.

> **Warning**
>
> Please note that new nodes start off showing a red `Fault` and may take a few minutes to fully initialize before they show a green OK.

> **Note**
>
> The `EPS` column represents Events Per Second consumed, so it will only be relevant on nodes that ingest data. Pure sensors do not ingest events, so those nodes will show 0 EPS. If you want to identify sensors that are generating large volumes of events, you can sort by the `Mgmt Out` column, which shows the outbound traffic throughput on the management network interface.

The options dropdown near the top of the page includes a checkbox which will show additional sensor-related columns in the table. You can use these sortable columns to help identify sensors that may be underperforming or due for a hardware upgrade. As these additional columns take up significant screen area, they will only be visible on wide displays where the Security Onion Console web browser window is wide enough to show a large number of tabular columns.

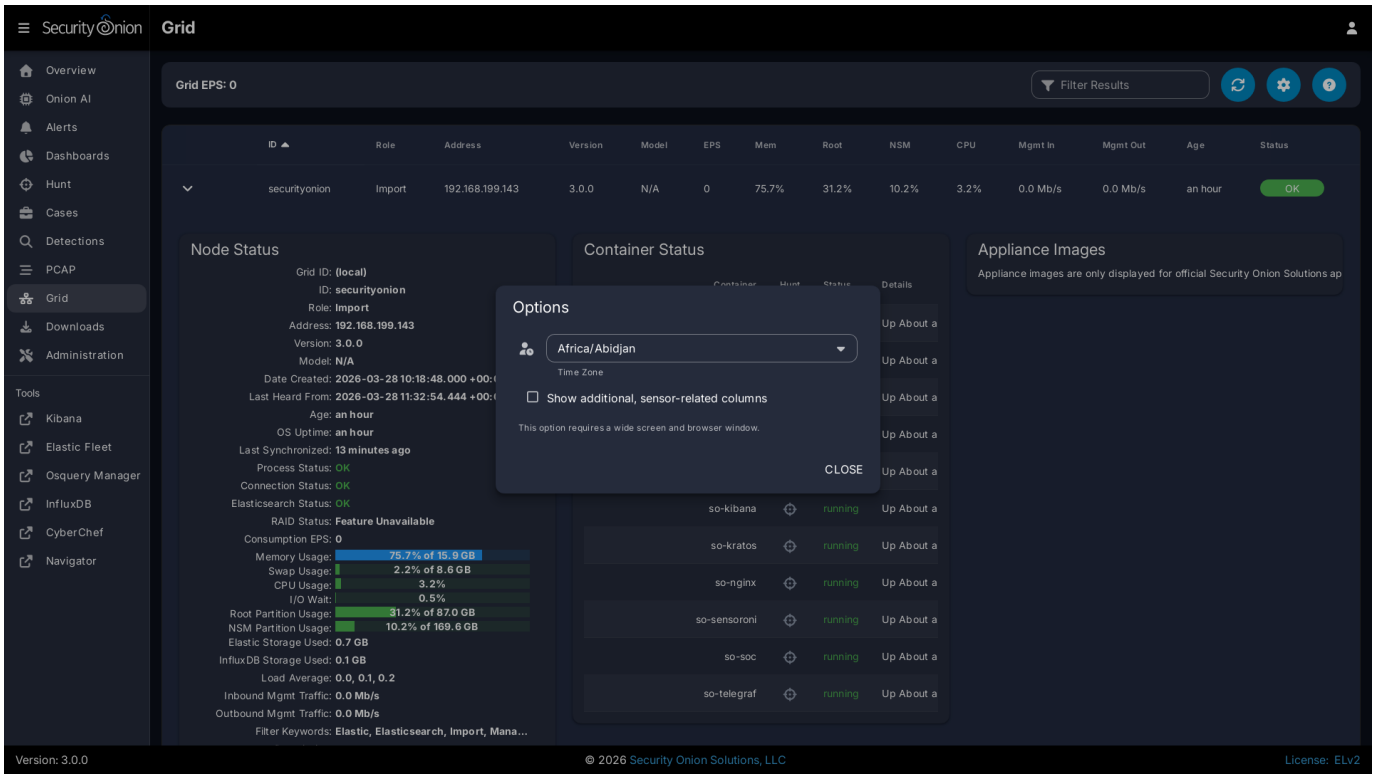

You can drill into individual nodes to see detailed information including Node Status, Container Status, and Appliance Images.

6.8.1 Node Status

The `Node Status` section displays many different fields relating to each node's status.

> **Note**
>
> If a node has not checked in recently then the metrics and statuses for that node will be slightly grayed out, to indicate that the values are stale.

ID

The `ID` field shows the hostname assigned to the node.

Role

The `Role` field shows the type of Security Onion node that was selected during Security Onion setup.

Address

The `Address` field shows the network IP address assigned to the management interface of the node.

Version

The `Version` field shows the version of Security Onion installed on this node.

Model

The `Model` field shows the official Security Onion Solutions appliance model number. For non-SOS devices, this field will show `N/A`.

Date Created

The `Date Created` field shows the date the node was created. This date is based on the node's filesystem timestamps, so replacing partition data or manually recreating core areas of the filesystem can interfere with assessing a node's true age.

Earliest PCAP

The `Earliest PCAP` field shows the earliest PCAP that is available on a sensor node and is only visible on sensor nodes which capture live packet data.

Last Heard From

The `Last Heard From` field shows the last time that the node checked-in with the manager. Note that a check-in doesn't always include updated node metrics.

Age

The `Age` field shows how long the node has been part of the grid and is based on the `Date Created` value.

OS Uptime

The `OS Uptime` field shows how long the node has been running since the last power-on or reboot event.

If the node needs to be restarted to apply kernel updates then a message will appear next to the uptime value indicating this. The reboot button at the bottom of the Grid page allows administrators to remotely reboot a node via the Security Onion Console web interface.

Last Synchronized

The `Last Synchronized` field shows how long ago the node was synchronized to the manager node. This is equivalent to the last Salt highstate run. Knowing this value can be helpful when making configuration changes to the grid and determining whether a specific node has received those changes.

Process Status

If the `Process Status` field shows `Fault`, you can check the other status indicators as well as the `Container Status` section to determine which process has failed.

Connection Status

The `Connection Status` field shows whether or not the node is currently connected to the grid.

Elasticsearch Status

If the node runs Elasticsearch, then the `Elasticsearch Status` field will show the status of it. If the status is anything other than OK, then see the Elasticsearch section to troubleshoot.

RAID Status

If you are using an official Security Onion Solutions appliance with RAID support, then you will see the corresponding status appear in this field.

Consumption EPS

The `Consumption EPS` field is the number of Events Per Second consumed.

Memory Usage

The `Memory Usage` field shows the system memory percentage used, as well as the total memory, in gigabytes. If this value is consistently in the red, then it may be time to add more system memory. Consistently red usage will likely end up causing node faults due to some services being automatically shutdown to recover memory for more critical processes.

Swap Usage

The `Swap Usage` field shows the system swap percentage used, as well as the total swap, in gigabytes. Systems that do not have swap enabled will remain at 0.0%. If this value is consistently in the red, then it may be time to increase the system memory and potentially the swap size.

CPU Usage

The `CPU Usage` field shows the system CPU percentage used, across all cores. If this value is consistently in the red, then it may be time to upgrade the node hardware or distribute the load across additional nodes.

I/O Wait

The `I/O Wait` field shows the system I/O wait percentage. Higher values indicate the system is spending more time waiting for network or disk data transfer. If this value is consistently in the red, then it may be time to replace slow disks or expand network throughput capacity.

Capture Loss

The `Capture Loss` field shows the percentage of packet capture loss reported by Zeek. Higher values indicate a reduced visibility into packets traversing the network. If Zeek is reporting capture loss but no packet loss, this usually means that the capture loss is happening upstream in the TAP or SPAN port itself.

Zeek Loss

The `Zeek Loss` field shows the percentage of dropped packets due to Zeek being unable to keep up with the flow of network data.

Suricata Loss

The `Suricata Loss` field shows the percentage of dropped packets due to Suricata being unable to keep up with the flow of network data.

Root Partition Usage

The `Root Partition Usage` field shows the percentage of the root OS disk utilization, as well as the total capacity of that disk (or partition). If this value is consistently in the red, then it can lead to problems including being unable to upgrade OS packages and Security Onion, the inability to save system logs, and other critical issues.

NSM Partition Usage

The `NSM Partition Usage` field shows the percentage of the NSM disk utilization, as well as the total capacity of that disk (or partition). If this value is consistently in the red, then it can lead to problems including being unable to ingest new events, store PCAP on disk, detect anomalous events, and other critical issues.

Elastic Storage Used

The `Elastic Storage Used` field shows the total gigabytes used by Elasticsearch to store the ingested events, across all indices.

InfluxDB Storage Used

The `InfluxDB Storage Used` field shows the total gigabytes used by InfluxDB to store the current and historic metric data collected from all nodes in the grid.

PCAP Retention

The `PCAP Retention` field shows the number of historic days of available packet capture data which can be viewed by analysts using the Security Onion Console PCAP tool.

Load Average

The `Load Average` field shows the 1 minute, 5 minute, and 15 minute load averages for the node. Note that on systems with high numbers of CPU cores, this average can be equally as high. For example, if a system has 128 cores then a load average of 128 generally indicates that all 128 cores are working at the peak capacity. Exceeding that number can indicate that some cores are bottlenecked due to waiting on I/O.

Redis Queue Size

The `Redis Queue Size` field shows the number of events queued in Redis waiting to be ingested into Elasticsearch. If this number is either steady or falling then it indicates the system is able to keep up with the current traffic flow. If this number is continually increasing then it can indicate a problem with ingest times taking too long for the amount of events that are being generated. Occasional increases are expected during traffic bursts but should eventually start to decrease once the high traffic flow period ends.

Inbound Monitor Traffic

The `Inbound Monitor Traffic` field shows the throughput of inbound bytes reaching the sensor's monitoring interface.

Dropped Monitor Traffic

The `Dropped Monitor Traffic` field shows the throughput of inbound bytes intended for the sensor's monitoring interface but are instead dropped, typically due to insufficient network capacity.

Inbound Mgmt Traffic

The `Inbound Mgmt Traffic` field shows the throughput of inbound bytes intended for the node's management interface. This is the internal interface that the node uses to communicate with other nodes in the Security Onion Grid.

Outbound Mgmt Traffic

The `Outbound Mgmt Traffic` field shows the throughput of outbound bytes being transmitted from the node's management interface. This is the internal interface that the node uses to communicate with other nodes in the Security Onion Grid.

Filter Keywords

The `Filter Keywords` fields shows the list of keywords that are associated with this node type. These keywords are useful for filtering to only show nodes of a certain type.

Description

The `Description` field shows the optional description you may have entered during Setup or set in Administration --> Configuration --> sensoroni --> config --> node_description.

Icons in Lower Left Corner

There are a few icons in the lower left of the `Node Status` section depending on what kind of node you are looking at:

- Clicking the first icon takes you to the InfluxDB dashboard for that particular node, to view historic health metrics and trends.
- If the node is a network sensor, then there will be an additional icon for sending test traffic to the sensor.
- Depending on the node type, there may be an additional icon for uploading your own PCAP or EVTX file. Clicking this icon results in an upload form. Once you've selected a file and initiated the upload, a status message appears. Uploaded PCAP files are automatically imported via so-import-pcap and EVTX files are automatically imported via so-import-evtx. Once the import is complete, a message will appear containing a hyperlink to view the logs from the import. Please note that import is not supported on heavy nodes. Also note that this import method is designed for smaller files. If you need to import files larger than the default max upload size then you will need to either change the max upload size via the Configuration screen, or manually import via so-import-pcap or so-import-evtx.

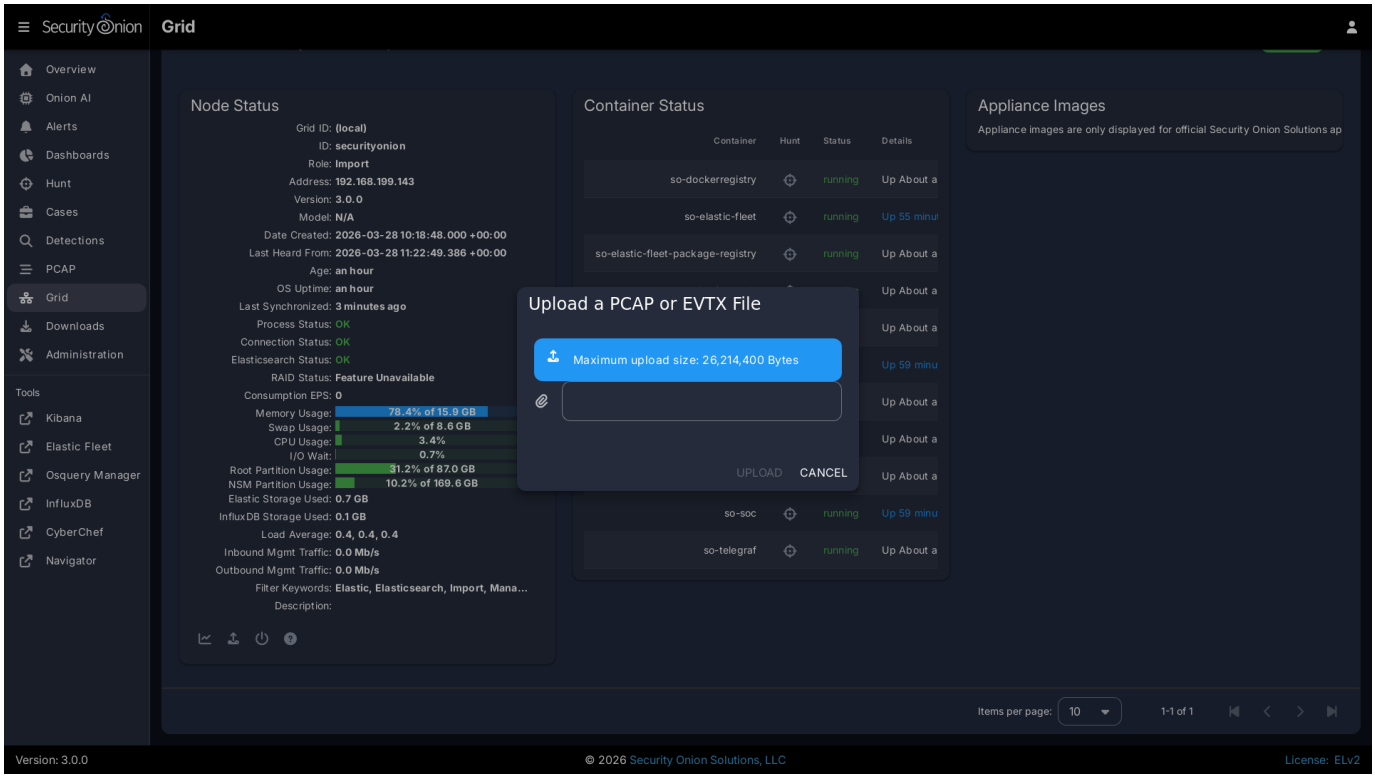

- The reboot button allows for remotely rebooting a grid node. This may be necessary when scheduled OS/kernel updates are automatically applied and require a restart to take effect. Review the notes on the confirmation dialog thoroughly before confirming a reboot. Rebooting a manager node will likely cause the Security Onion Console web interface to become temporarily unavailable.
- Clicking the question mark button takes you to this help document.

6.8.2 Container Status

> **Note**
>
> Restarting a node can take several minutes for all containers to return to a running state.

If any containers show anything other than `running` click the cross-hair icon next to the container name. This will bring up the Hunt screen showing logs specific to that container, and may help determine why the container is not running.

6.8.3 Appliance Images

If a node is running on an official Security Onion Solutions appliance then the Grid page will show pictures of the front and rear of the appliance. This is useful for walking through connectivity discussions with personnel in the data center. When not using official Security Onion Solutions appliances it will simply display a message to that effect.

6.8.4 Other Grid Pages

> **Note**
>
> You can manage Grid members and Grid configuration in the Administration section.

6.9 Downloads

Security Onion Console includes a Downloads interface that allows you to download the Elastic Agent for various operating systems.

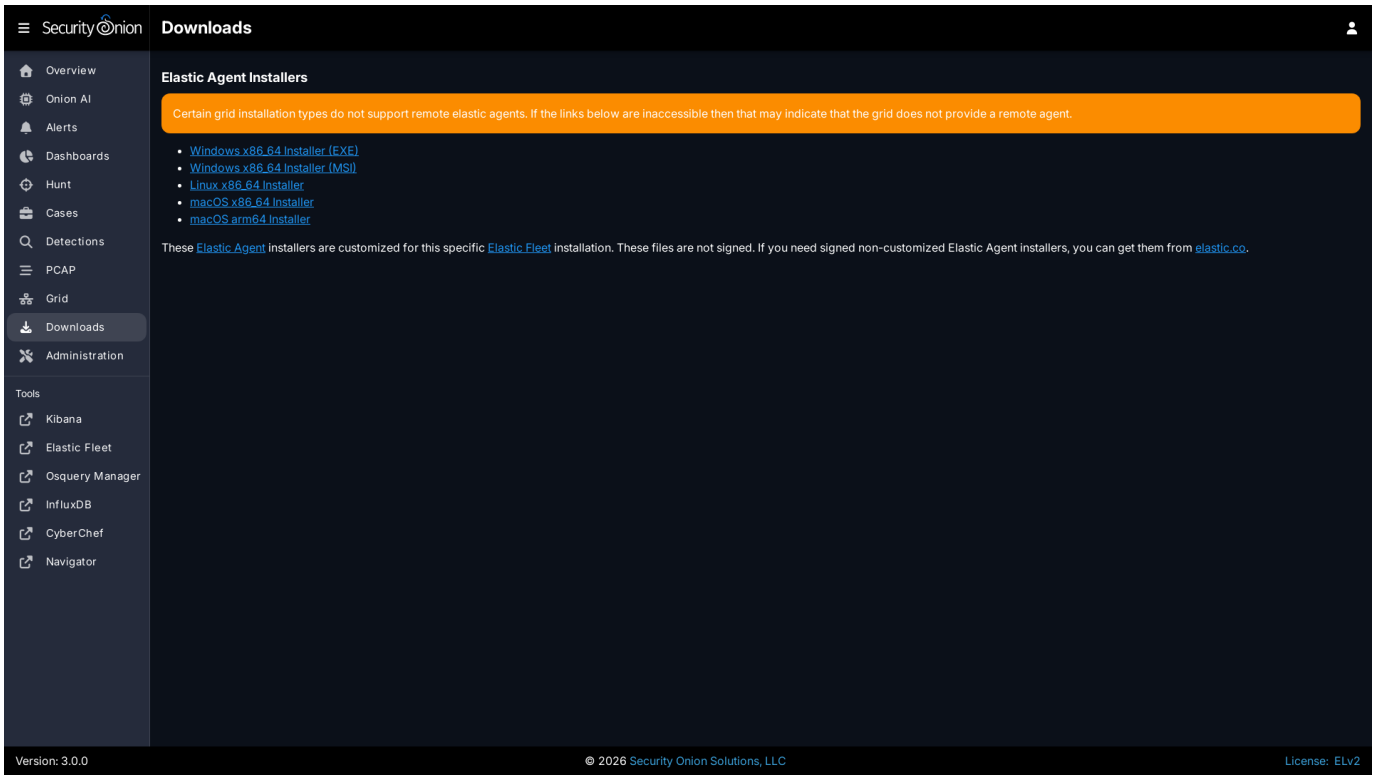

Warning

Please note that Evaluation installs and Import installs do not support remote Elastic agents, so in those cases the links are shown for demonstration purposes only.

Note

When installing the Elastic Agent onto remote systems, be sure to allow network access through the Firewall.

6.10 Administration

Security Onion Console includes an Administration section which allows you to administer Users, Grid Members, Configuration, and the License Key.

6.10.1 Users

The Users page shows all user accounts that have been created for the grid.

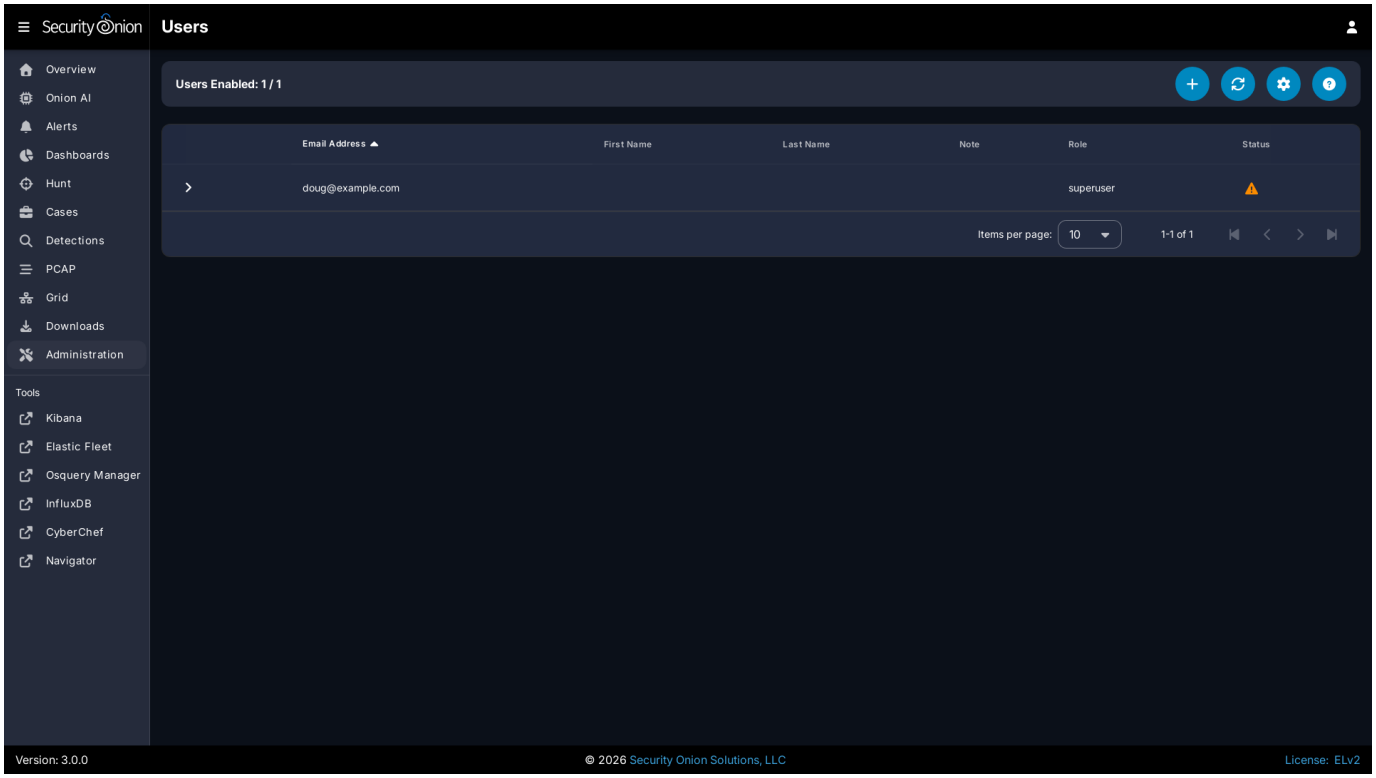

The Note column allows administrators to include a short note on a user's account.

The Role column lists roles assigned to the user as defined in the RBAC section.

The Status column will show different icons depending on the status of the account:

- orange exclamation point - account enabled but has not yet changed their password and does not have MFA enabled
- blue icon with shield - account enabled with MFA enabled
- no icon - account enabled and has changed their password but does not have MFA enabled
- grey user with slash - account locked

Hovering over the icon in the Status column will show you these details as well.

6.10.2 Grid Members

The Grid Members page shows nodes that have attempted to join the grid and whether or not they have been accepted into the grid by an administrator.

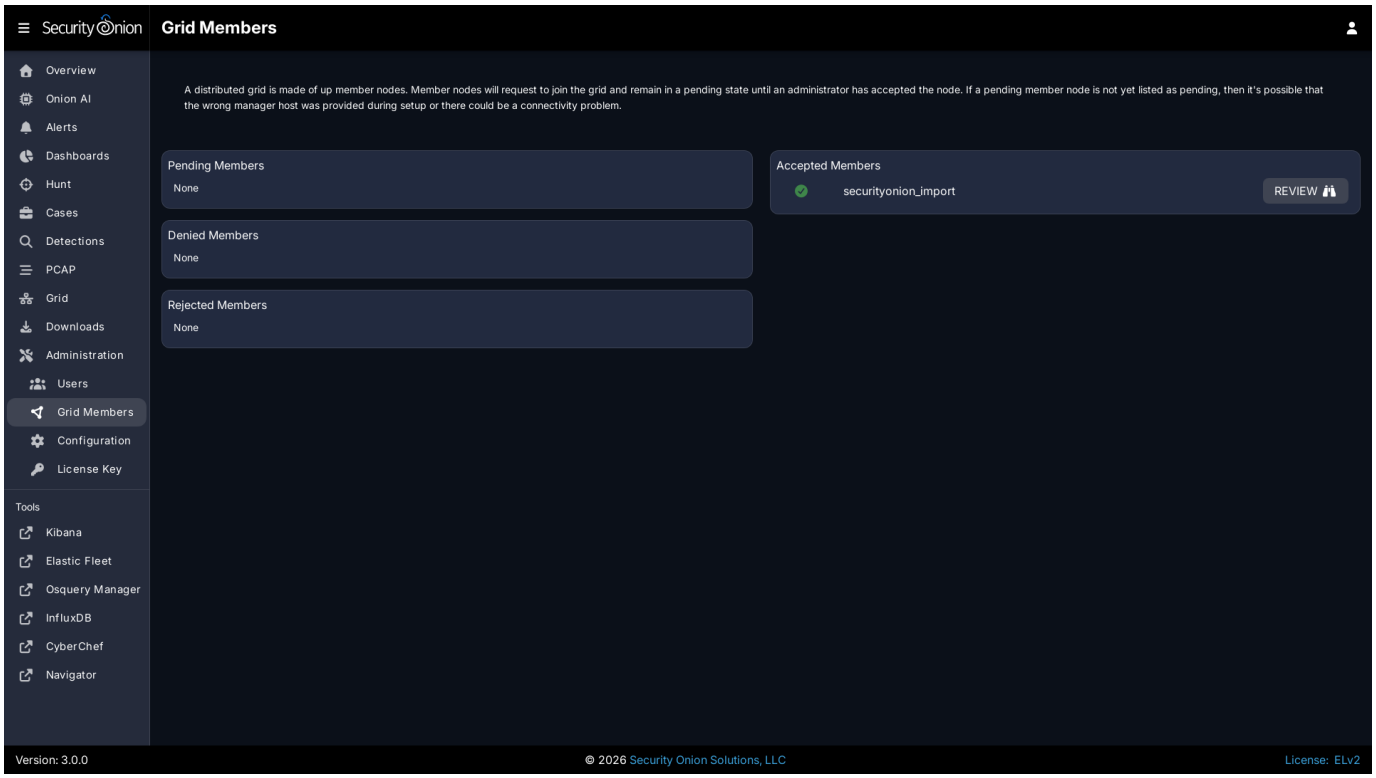

Unaccepted members are displayed on the left side and broken into three sections: Pending Members, Denied Members, and Rejected Members. When you accept a member, it will then move to the right side under Accepted Members.

For accepted members, you can click the REVIEW button to show additional information about the grid member. If you want to remove the member, you can then click the DELETE button and review the confirmation.

6.10.3 Configuration

The Configuration page allows you to configure various components of your grid.

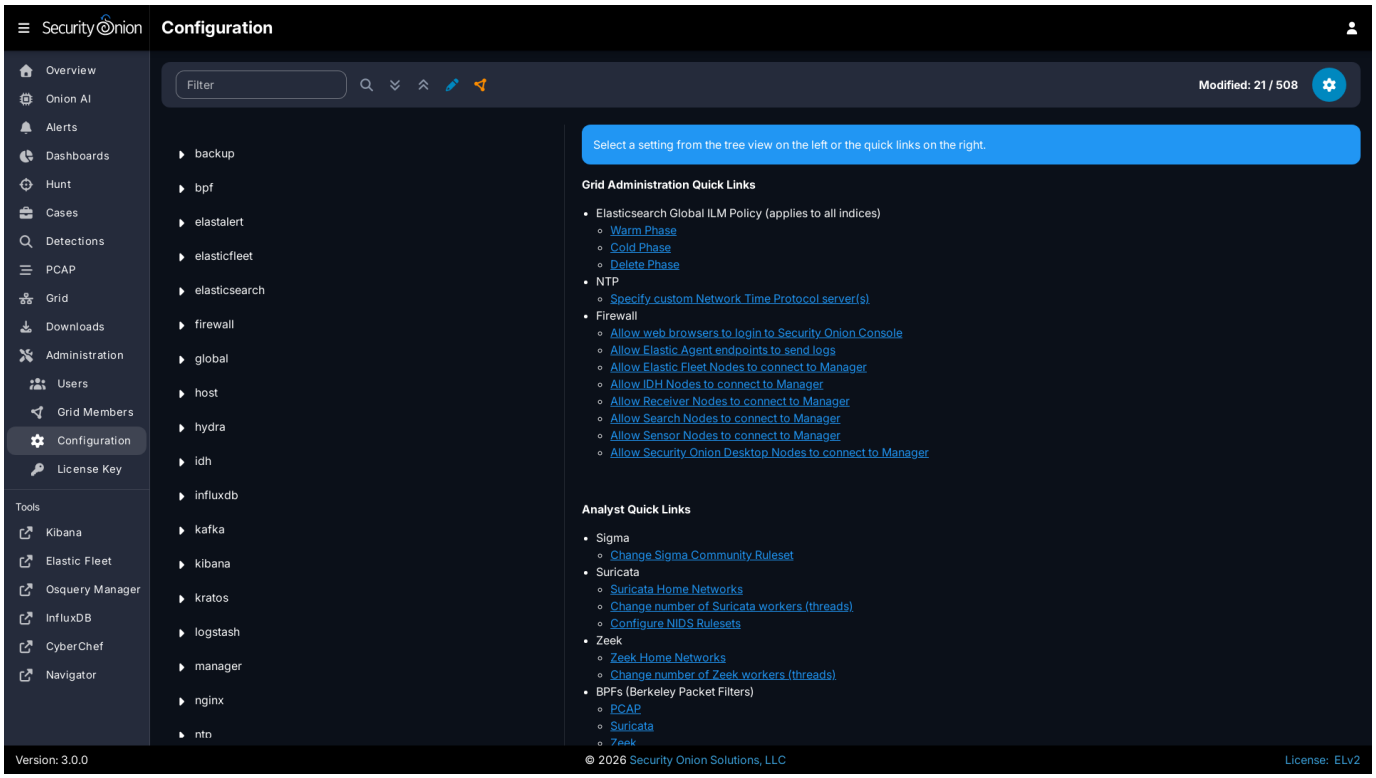

The most common configuration options are shown in the quick links on the right side. On the left side, click on a component in the tree view to drill into it and show all available settings for that component. You can then click on a setting to show the current setting or modify it if necessary. If you make a mistake, you can easily revert back to the default value. If a blue question mark appears on the setting page, click it to go to the documentation for that component.

If unsure of which component a particular setting may belong to, use the Filter at the top of the list to look for a particular setting. To the right of the Filter field are buttons that do the following:

- apply the search filter
- expand all settings
- collapse all settings
- show settings that have been modified from the default value
- show settings that have a unique value specified for one or more nodes in the grid

> **Note**
>
> Keys that include `_x_` indicate a placeholder value used to represent a period (`.`).

Some settings can be applied across the entire Grid or to specific nodes. Applying a setting to a specific node will override the grid setting.

Advanced Settings

By default, the Configuration page excludes settings that are not intended to be adjusted by most Grid administrators. These advanced settings can cause loss of data and other issues if adjusted incorrectly. To see the advanced settings, go to the Options bar at the top of the page and then click the toggle labeled `Show advanced settings` .

Enabling advanced settings will result in longer load times when viewing the Configuration screen.

> **Warning**
>
> Changing advanced settings is unsupported and could result in requiring a full cluster re-installation.

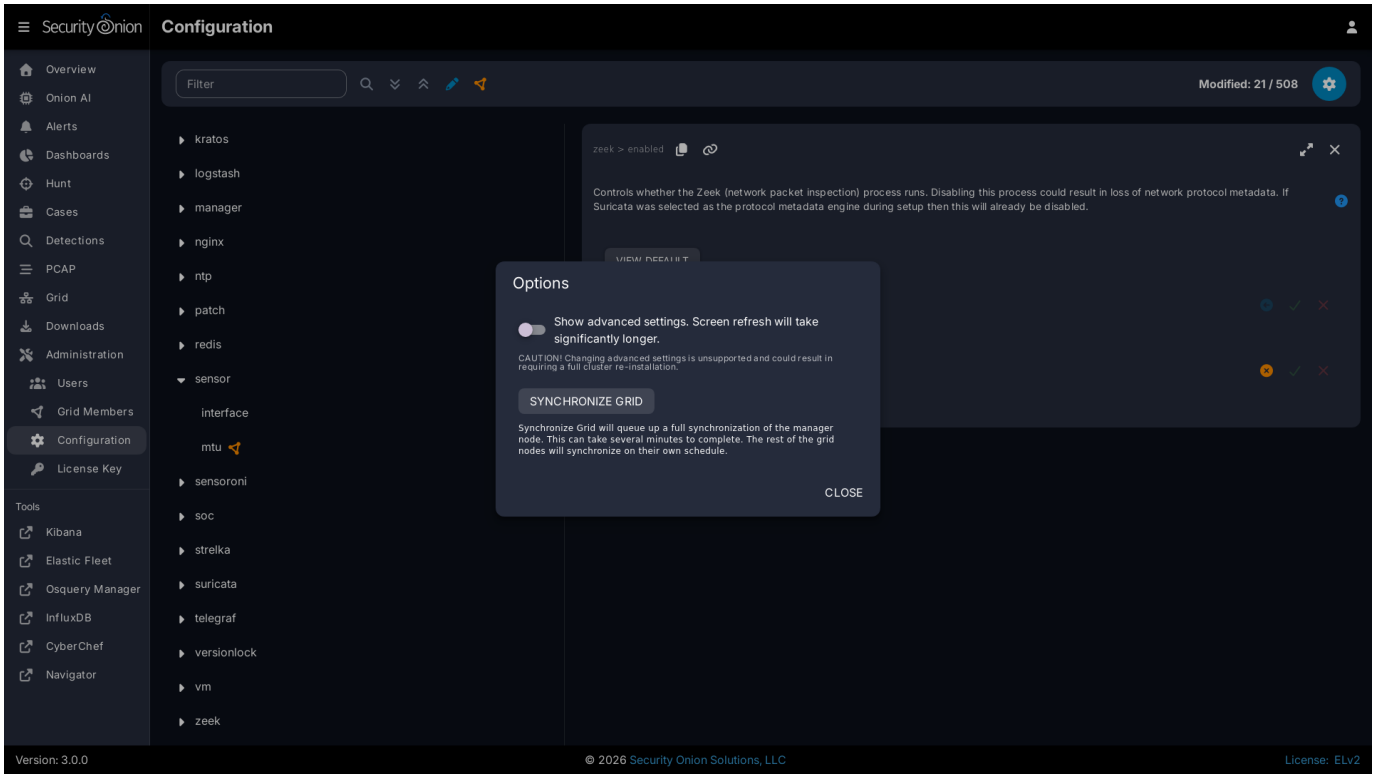

Duplicate Settings

Some settings can be duplicated to more easily create new settings. If a setting is eligible for duplication, then it will have a DUPLICATE button on the right side of the page, provided the `Show advanced settings` option is enabled at the top of the screen. Creating a duplicate setting is a TWO-STEP process.

1. Click the `DUPLICATE` button, provide a name for the new setting, and then click the `CREATE SETTING` button.
2. The new setting will automatically be shown in the Configuration screen. At this point it is not yet saved to the server. The setting's value must be modified explicitly to persist this new setting. Once the value has been modified, click the green checkmark button to save it.

> **Note**
>
> Duplicated settings do not retain their original setting's full behavior. For example, if the original setting only allowed for CIDR values, this new setting will not have the same protections on later views in the Configuration screen. Further, duplicated settings are marked as advanced settings. In order to see the new setting at a later time the `Show advanced settings` option must be enabled under the Configuration Options at the top of the Configuration screen. Finally, please note that duplicated settings cannot be removed or renamed via the SOC user interface.

6.10.4 License Key

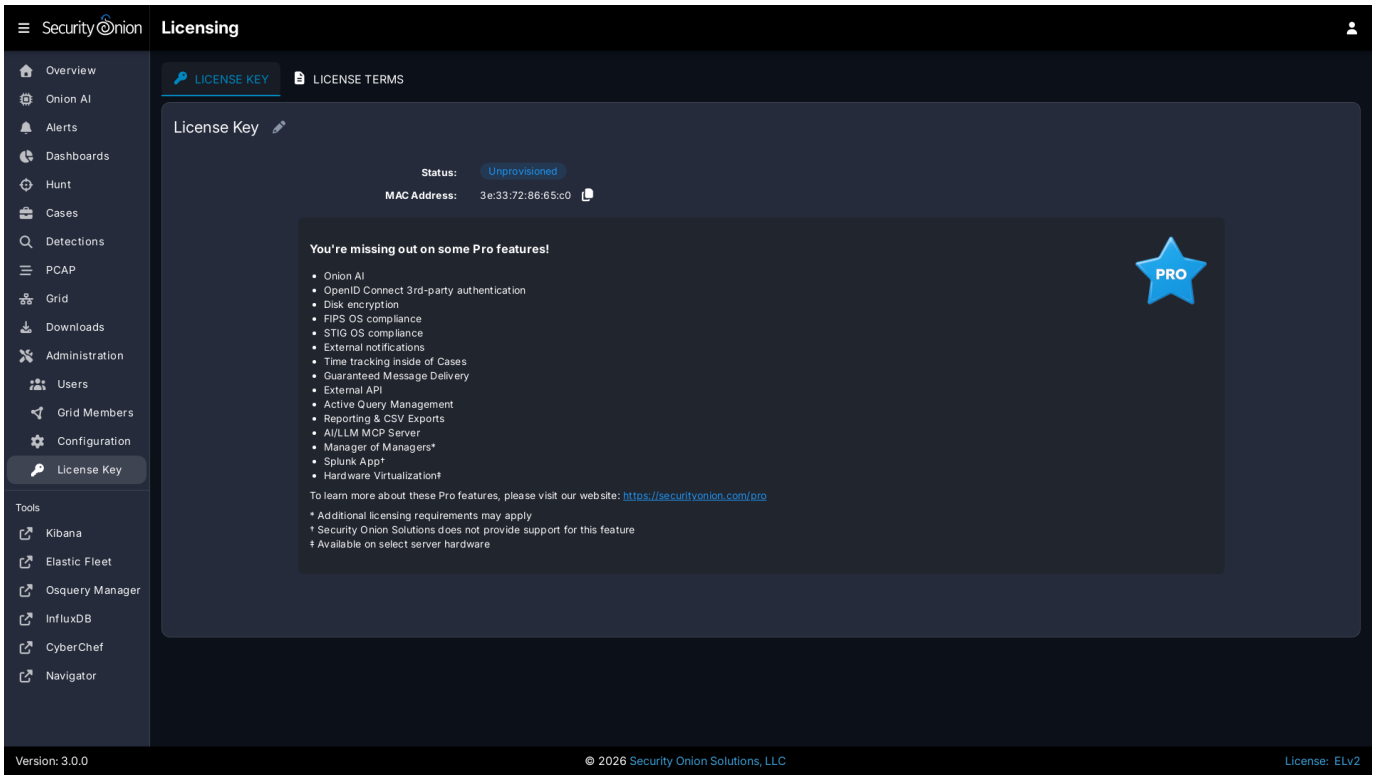

The License Key screen allows you to add a license key for Security Onion Pro. Once you've added a license key, the screen will show details about your license key.

6.11 Kibana

Security Onion Console includes a link on the sidebar that takes you to Kibana.

6.11.1 Authentication

Log into Kibana using the same username and password that you use for SOC.

You can add new user accounts to both Kibana and SOC at the same time as shown in the Adding Accounts section.

> **Warning**
>
> If you create accounts directly in Kibana (rather than in SOC), then those accounts will NOT be able to log into SOC.

6.11.2 Kibana Dashboards

We've included a simple set of dashboards in Kibana. These Kibana dashboards are not as comprehensive as those in SOC Dashboards.

Once you log into Kibana, you should start on the `Security Onion - Home` dashboard. Notice the visualization in the upper left is labeled `Security Onion - Navigation`. This navigation panel contains links to other dashboards and will change depending on what dashboard you're currently looking at. For example, when you're on the `Security Onion - Home` dashboard and click the `Alert` link, you will go to the `Security Onion - Alerts` dashboard and the Navigation panel will then contain links to more specific alert dashboards for Suricata. When you're done looking at alerts, you can click the `Home` link in the navigation panel to go back to the main `Security Onion - Home` dashboard.

If you ever need to reload Kibana dashboards, you can run the following command on your manager:

```
sudo so-kibana-config-load
```

If that doesn't resolve the issue, then you may need to run the following:

```
sudo salt-call state.apply Kibana.so_savedobjects_defaults -l info queue=True
```

If you try to modify a default Kibana dashboard, your change will get overwritten. Instead of modifying, copy the desired dashboard and edit the copy. You may also want to consider setting up Kibana Spaces as this will allow you to make whatever changes you want without them being overwritten. This includes not only dashboards but certain Kibana settings as well. You can read more about Kibana Spaces at https://www.elastic.co/guide/en/kibana/current/xpack-spaces.html.

6.11.3 Search Results

In Kibana, search results are limited to the first `100` results for a particular query. If you don't feel like this is adequate after narrowing your search, you can adjust the value for `discover:sampleSize` in Kibana by navigating to `Stack Management` --> `Advanced Settings` and changing the value. It may be best to change this value incrementally to see how it affects performance for your deployment.

6.11.4 Timestamps

By default, Kibana will display timestamps in the timezone of your local browser. If you would prefer timestamps in UTC, you can go to `Management` --> `Advanced Settings` and set `dateFormat:tz` to `UTC`.

6.11.5 Configuration

Most Kibana configuration settings are in Kibana itself. However, configuration settings that would traditionally be set in the Kibana configuration file can be configured by going to Administration --> Configuration --> Kibana.

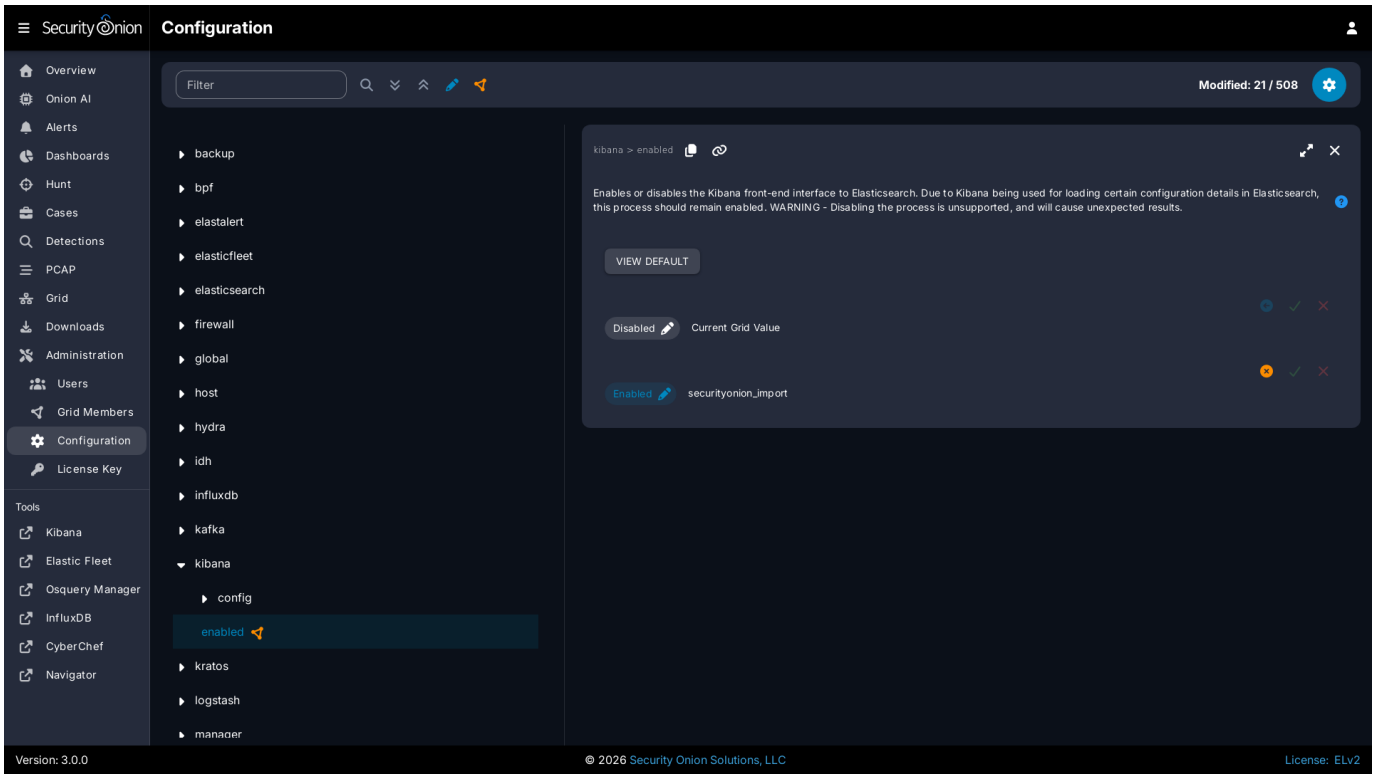

6.11.6 Diagnostic Logging

Kibana logs to `/opt/so/log/kibana/kibana.log`. Depending on what you're looking for, you may also need to look at the Docker logs for the container:

```
sudo docker logs so-kibana
```

If you try to access Kibana and it says `Kibana server is not ready yet` even after waiting a few minutes for it to fully initialize, then check `/opt/so/log/kibana/kibana.log`. You may see something like:

```
Another Kibana instance appears to be migrating the index. Waiting for that migration to complete. If no other Kibana instance is attempting migrations, you can get past this message by deleting index .kibana_6 and restarting Kibana
```

If that's the case, then you can do the following (replacing `.kibana_6` with the actual index name that was mentioned in the log):

```
curl -k -XDELETE https://localhost:9200/.kibana_6

sudo so-kibana-restart
```

If you then are able to login to Kibana but your dashboards don't look right, you can reload them as follows:

```
so-kibana-config-load
```

6.11.7 Features

You can enable or disable specific features by clicking the main menu in the upper left corner, then click `Stack Management`, then click `Spaces`, then click `Default`. For more information, please see https://www.elastic.co/guide/en/kibana/current/xpack-spaces.html#spaces-control-feature-visibility.

6.11.8 More Information

Note

For more information about Kibana, please see https://www.elastic.co/kibana.

6.12 Elastic Fleet

Security Onion Console includes a link on the sidebar that takes you to the Fleet page inside Kibana.

Elastic Fleet is pre-configured during Security Onion setup. If you need to make changes to the configuration, you can do so via the Fleet page in Kibana as detailed below.

The Fleet page has multiple tabs across the top: Agents, Agent policies, Enrollment tokens, Uninstall tokens, Data streams, and Settings. Each of these tabs is described below.

6.12.1 Agents

The Agents tab displays registered Elastic agents.

To view agent details, click the `Host` name.

To assign the agent to a new policy, unenroll, upgrade the agent, or perform other actions, click the `Actions` menu on the right side of the agent listing and select the appropriate option.

By default, Elastic Agent is installed on every Security Onion Grid node. As a result, all Grid node agents will be enrolled in the `SO-Grid-Nodes` agent policy.

> **Warning**
>
> We do not recommend removing policy settings for Security Onion Grid node agents.

Adding Agents

To add a new agent to your deployment, see the Elastic Agent section.

Upgrading Agents

Fleet automatically checks to see if agents are the latest version. If not, agents will display `Upgrade available`. If you would like to upgrade an agent, click the `Actions` menu on the right side and select the `Upgrade agent` option. You will then have the opportunity to select which version to upgrade to. Please choose the version number that matches the version of the Elastic stack that you are currently running. For example, if you are currently running Elastic 8.18.4 then you should select the 8.18.4 agent version.

> **Warning**
>
> If you try to upgrade to a version of the agent that is newer than the version of your Elastic stack, then you may run into errors. For example, if you are currently running Elastic 8.18.4 and try to upgrade to agent version 8.19.0 or 9.1.0 then you may run into errors.

Monitoring Agents

Agent health can be monitored in Elastic Fleet under the Agents tab, however you may want to generate alerts when an Agent reports a degraded or offline state. Via Administration --> Configuration --> manager -> agent_monitoring you can enable a script that will periodically poll Elastic Fleet. Agents in an offline or degraded state for longer than the configured threshold (default is 5h) will generate an Alert in the Alerts interface of SOC.

For more granular control over what Agents generate an alert, you can leverage fleet 'tags' to tag those agents. Then under Administration --> Configuration --> manager -> agent_monitoring -> config -> custom_kquery you can enter a kql search like `tags: critical`. Then only those Agents tagged as 'critical' would generate an Alert.

For more information on Elastic Fleet tags see https://www.elastic.co/docs/reference/fleet/filter-agent-list-by-tags#add-tags-in-fleet.

> **Tip**
>
> Security Onion Pro users can configure external notifications.

6.12.2 Agent Policies

Agent policies dictate what data each agent will ingest and forward to Elasticsearch. This could be through the use of an HTTP, log file, or TCP-based input.

The individual components within each agent policy are called integrations (referred to as `package policies` at the API level), and refer to a specific input and settings pertinent to a data source.

For example, the `SO-Grid-Nodes` agent policy is comprised of the following integrations:

- Elasticsearch-logs (`Elasticsearch` integration)
- import-evtx-logs (`Custom Logs` integration)
- import-Suricata-logs (`Custom Logs` integration)
- import-Zeek-logs (`Custom Logs` integration)
- kratos-logs (`Custom Logs` integration)
- Osquery-Grid-nodes (`Osquery Manager` integration)
- Redis-logs (`Redis` integration)
- Strelka-logs (`Custom Logs` integration)
- Suricata-logs (`Custom Logs` integration)
- syslog-tcp-514 (`Custom Logs` integration)
- syslog-udp-514 (`Custom Logs` integration)
- system-Grid-nodes (`System` integration)
- Zeek-logs (`Custom Logs` integration)

6.12.3 Agent Policies - endpoints-initial

Agent installers downloaded from Downloads are deployed using the `endpoints-initial` Agent Policy. This policy includes the `Elastic Defend`, `Osquery Manager`, `System`, and `Windows` integrations.

elastic-defend-endpoints (`Elastic Defend` integration)

The `Elastic Defend` integration has both free and paid features. By default, only the following free features are enabled:

- Event Collection - Windows
- Credential Access
- DLL and Driver Load
- DNS
- File
- Network
- Process
- Registry
- Security
- Event Collection - macOS
- DNS
- File
- Process
- Network
- Event Collection - Linux
- File
- Network
- Process

Osquery-endpoints (`Osquery Manager` integration)

The `Osquery Manager` integration runs Osquery as a daemon on the endpoint and makes the endpoint available for Live or Scheduled queries through the Osquery manager interface in Kibana.

system-endpoints (`System` integration)

The `System` integration collects the following logs from the endpoint, where applicable:

- System auth logs
  - `/var/log/auth.log*`
  - `/var/log/secure*`
- Syslog logs
  - `/var/log/messages*`
  - `/var/log/syslog*`
  - `/var/log/system*`
- Windows Event Log - Application channel
- Windows Event Log - Security channel
- Windows Event Log - System channel

windows-endpoints (`Windows` integration)

The `Windows` integration collects the following logs from the endpoint, where applicable:

- Windows Event Log:
- ForwardedEvents channel
- Windows Powershell channel
- Microsoft-Windows-Powershell/Operational channel
- Microsoft-Windows-Sysmon/Operational channel

6.12.4 Enrollment Tokens

An enrollment token allows an agent to enroll in Fleet, subscribe to a particular agent policy, and send data.

Each agent policy typically uses its own enrollment token. It is recommended that these tokens are NOT to be changed, especially those generated by default Security Onion agent policies.

6.12.5 Data Streams

Data collected by Elastic Agent is sent to a data stream by default. This allows data to be efficiently categorized and managed across a variety of datasets. This section within the Fleet UI allows for a quick review of data streams generated by data from Elastic Agent.

6.12.6 Settings

The section provides details such as:

- Fleet server hosts in your deployment
- Configured outputs
- specifies where data will be sent
- this should include Elasticseach for the Fleet server and Logstash for Elastic Agent
- Method in which agent binaries will be downloaded
- this will be a local artifact repository if running an airgapped deployment

> **Warning**
>
> We do NOT recommend changing these settings, as they are managed by Security Onion.

If you want more granular control over which Fleet Server an Agent will send logs to, there are two options:

- The first option is to use firewall rules to only allow certain agents. Suppose you have two Fleet Server Nodes, one at 192.168.55.25 and the other at 192.168.58.25. If you want your endpoints in the 192.168.58.0/24 subnet to only connect to the Fleet server at 192.168.58.25, you would add custom firewall rules via Administration --> Configuration --> firewall --> hostgroups --> elastic_agent_endpoint. Select the 192.168.58.25 Fleet Node and add `192.168.58.0/24`. Endpoints in that subnet will still attempt to connect to the Fleet Server Node at 192.168.55.25, but since it is not accessible (no firewall rules that enable communication), they will connect to the Fleet Node at 192.168.58.25.
- The second option is to purchase an Elastic license. A paid Elastic license offers the ability to customize different Outputs per Agent Policy.

6.12.7 Integrations

Elastic Fleet supports integrations and you can read more in the third-party-integrations section.

6.12.8 Custom FQDN URL

You can add custom FQDN for Agents to connect to (for both control traffic on port TCP/8220 and data traffic on port TCP/5055) by editing the config as follows.

First, go to Administration --> Configuration --> elasticfleet.

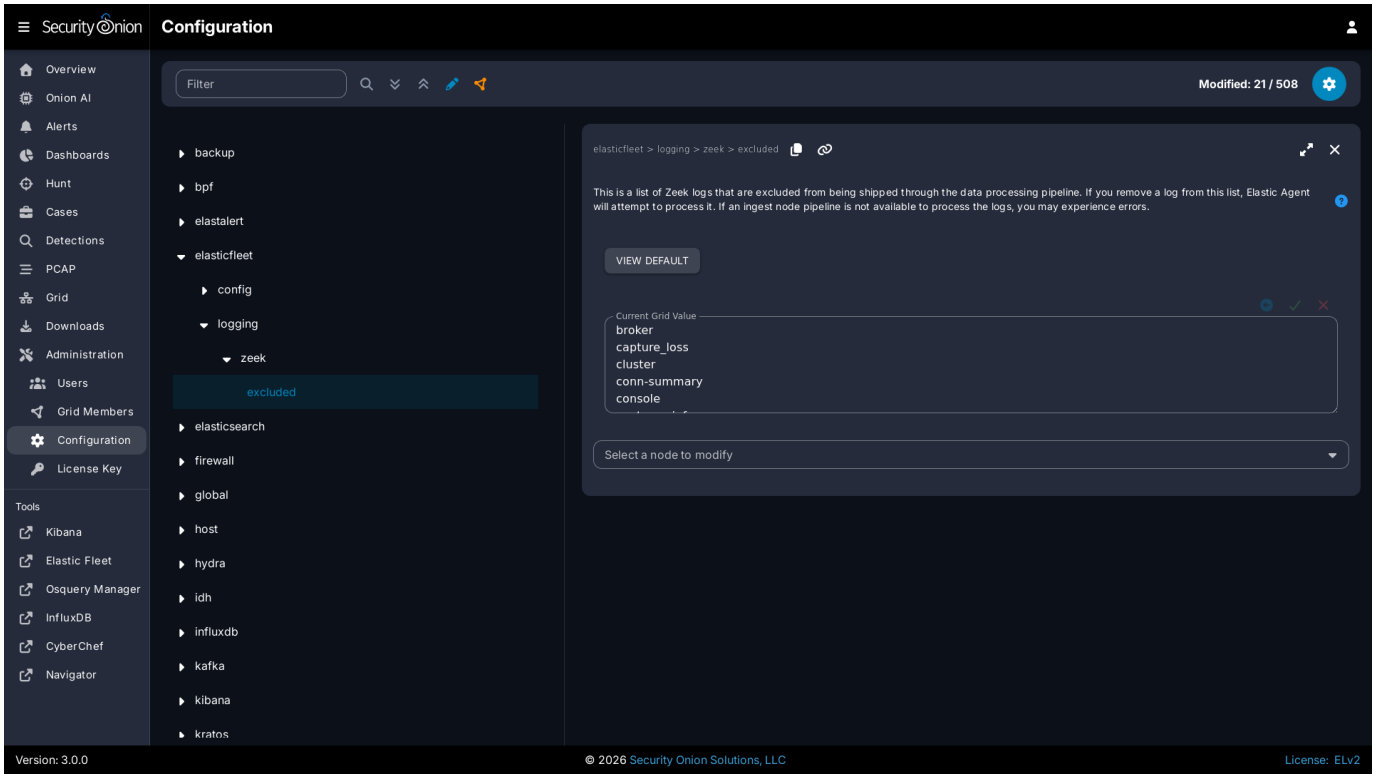

At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option. Then, navigate to elasticfleet --> config --> server --> custom_fqdn and set your custom FQDN. Within 15 minutes, the grid will apply these new settings and you should see the new FQDNs show up in Elastic Fleet settings. New agent installers will also be regenerated to use this new setting.

6.12.9 More Information

Note

For more information about Fleet, please see https://www.elastic.co/guide/en/kibana/current/fleet.html.

6.13 Osquery Manager

Security Onion Console includes a link on the sidebar which takes you to the Osquery Manager page inside Kibana.

6.13.1 More Information

Note

For more information about Osquery Manager, please see https://docs.elastic.co/en/integrations/osquery_manager.

6.14 InfluxDB

Security Onion Console includes a link on the sidebar that takes you to InfluxDB.

From https://github.com/influxdata/influxdb:

> InfluxDB is an open source time series platform. This includes APIs for storing and querying data, processing it in the background for ETL or monitoring and alerting purposes, user dashboards, and visualizing and exploring the data and more.

6.14.1 Authentication

Log into InfluxDB using the same username and password that you use for SOC.

If you need to reset your InfluxDB password, you can reset your SOC password via the Administration interface which will also update your InfluxDB password.

6.14.2 Configuration

You can configure InfluxDB by going to Administration --> Configuration --> influxdb.

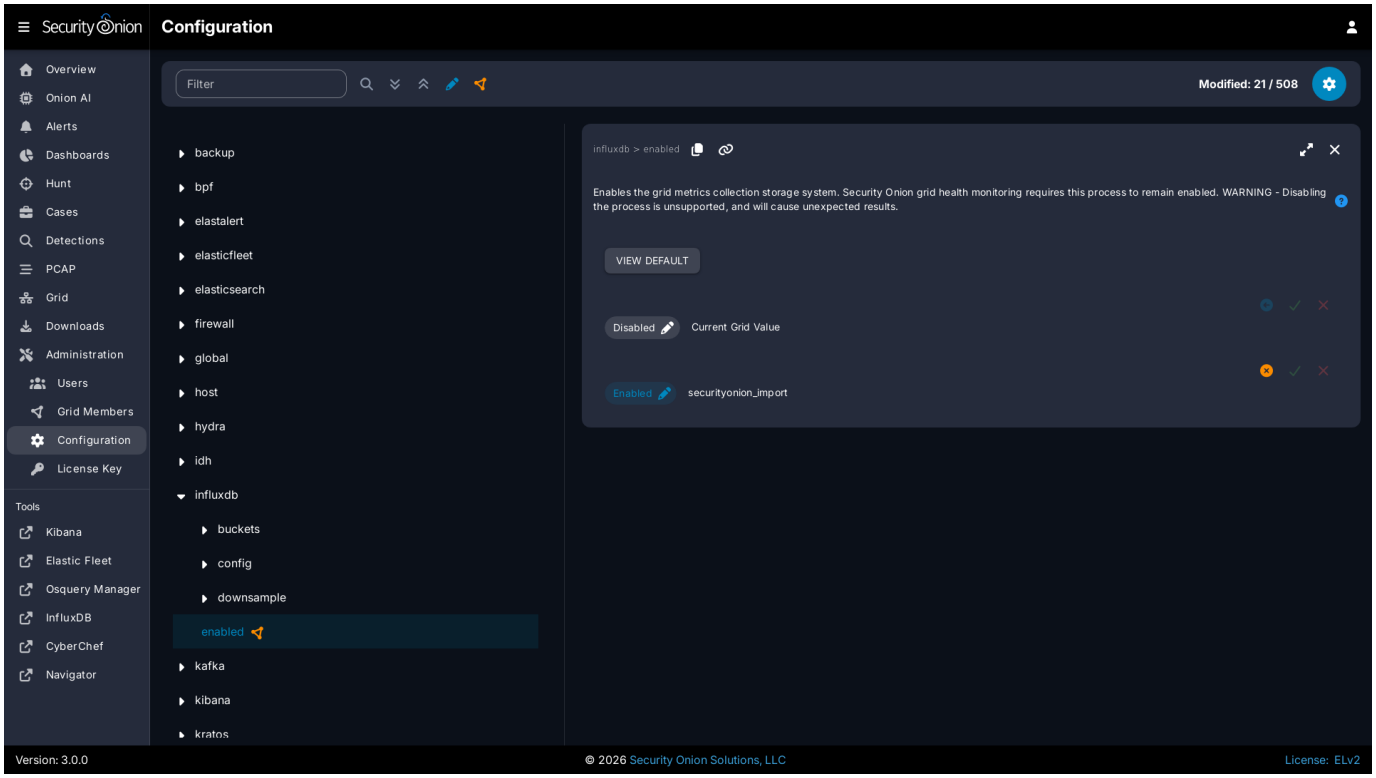

You can configure Telegraf by going to Administration --> Configuration --> telegraf.

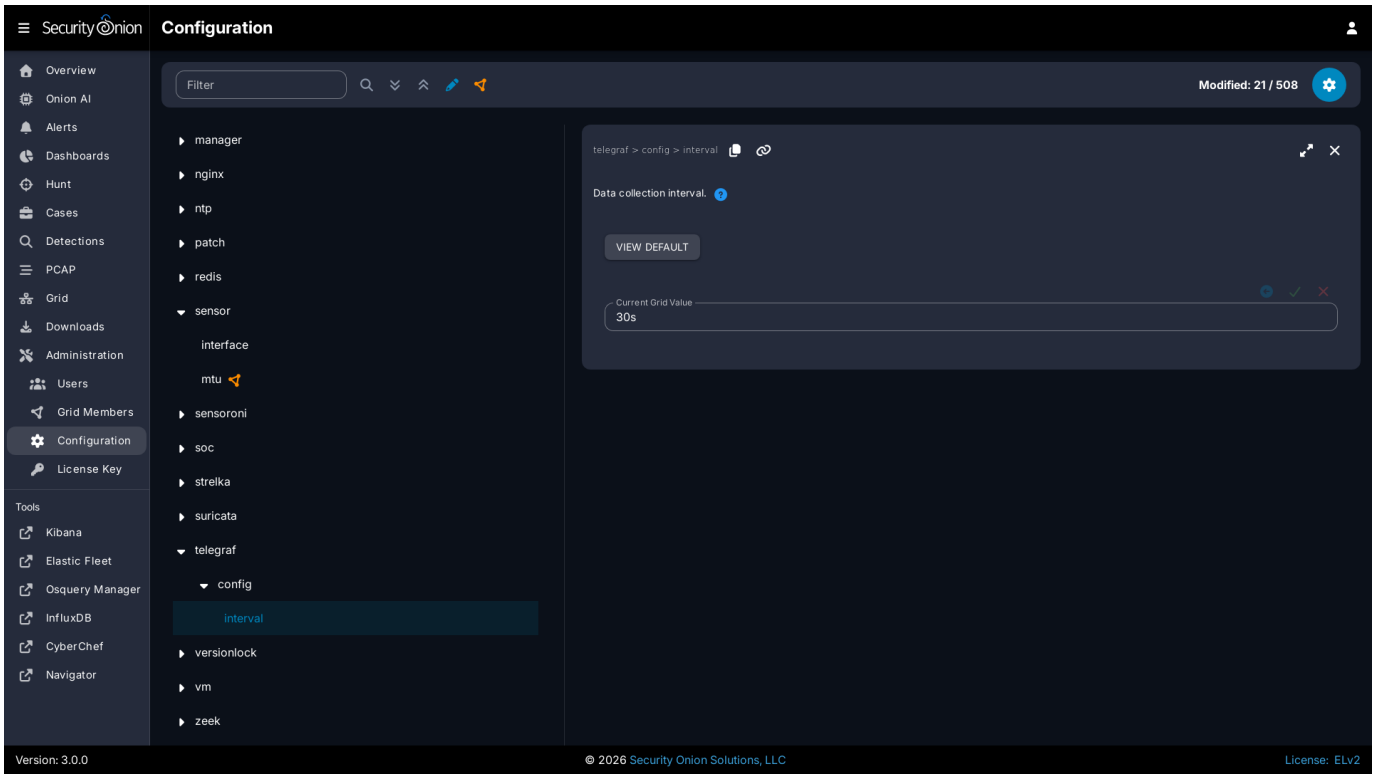

6.14.3 More Information

Note

For more information about InfluxDB, please see https://github.com/influxdata/influxdb.

6.15 CyberChef

Security Onion Console includes a link on the sidebar that takes you to CyberChef.

From https://github.com/gchq/CyberChef:

> The Cyber Swiss Army Knife
>
> CyberChef is a simple, intuitive web app for carrying out all manner of "cyber" operations within a web browser. These operations include simple encoding like XOR or Base64, more complex encryption like AES, DES and Blowfish, creating binary and hexdumps, compression and decompression of data, calculating hashes and checksums, IPv6 and X.509 parsing, changing character encodings, and much more.
>
> The tool is designed to enable both technical and non-technical analysts to manipulate data in complex ways without having to deal with complex tools or algorithms.
>
> There are four main areas in CyberChef:
>
> 1. The input box in the top right, where you can paste, type or drag the text or file you want to operate on.
> 2. The output box in the bottom right, where the outcome of your processing will be displayed.
> 3. The operations list on the far left, where you can find all the operations that CyberChef is capable of in categorised lists, or by searching.
> 4. The recipe area in the middle, where you can drag the operations that you want to use and specify arguments and options.

6.15.1 Accessing

To access CyberChef, log into SOC and click the CyberChef hyperlink.

You can send highlighted text from PCAP to CyberChef. When the CyberChef tab opens, you will see your highlighted text in both the Input box and the Output box.

You can send all visible packet data from PCAP to CyberChef. When the CyberChef tab opens, it will automatically apply the `From Hexdump` recipe to render the hexdump that was sent.

6.15.2 File Extraction

Suppose you are looking at an interesting HTTP file download in PCAP and want to extract the file using CyberChef:

- Click the PCAP CyberChef button and CyberChef will launch in a new tab. It will then show the hexdump in the Input box, automatically apply the `From Hexdump` recipe, and show the HTTP transcript in the Output box.
- You may want to apply an operation from the left column. One option is to use the `Extract Files` operation and optionally specify certain file types for extraction. Another option is to instead remove the HTTP headers using the `Strip HTTP headers` operation.
- If a magic wand appears in the Output box, then CyberChef has detected some applicable operations and you can click the magic wand to automatically apply those operations. For example, CyberChef might automatically apply `Strip HTTP headers` and then render the file.

6.15.3 More Information

Note

For more information about CyberChef, please see https://github.com/gchq/CyberChef.

6.16 ATT&CK Navigator

Security Onion Console includes a link on the sidebar that takes you to ATT&CK Navigator.

From https://github.com/mitre-attack/attack-navigator:

> The ATT&CK Navigator is designed to provide basic navigation and annotation of ATT&CK matrices, something that people are already doing today in tools like Excel. We've designed it to be simple and generic - you can use the Navigator to visualize your defensive coverage, your red/blue team planning, the frequency of detected techniques or anything else you want to do. The Navigator doesn't care - it just allows you to manipulate the cells in the matrix (color coding, adding a comment, assigning a numerical value, etc.). We thought having a simple tool that everyone could use to visualize the matrix would help make it easy to use ATT&CK.
>
> The principal feature of the Navigator is the ability for users to define layers - custom views of the ATT&CK knowledge base - e.g. showing just those techniques for a particular platform or highlighting techniques a specific adversary has been known to use. Layers can be created interactively within the Navigator or generated programmatically and then visualized via the Navigator.

6.16.1 Accessing

To access Navigator, log into SOC and then click the Navigator hyperlink on the left side.

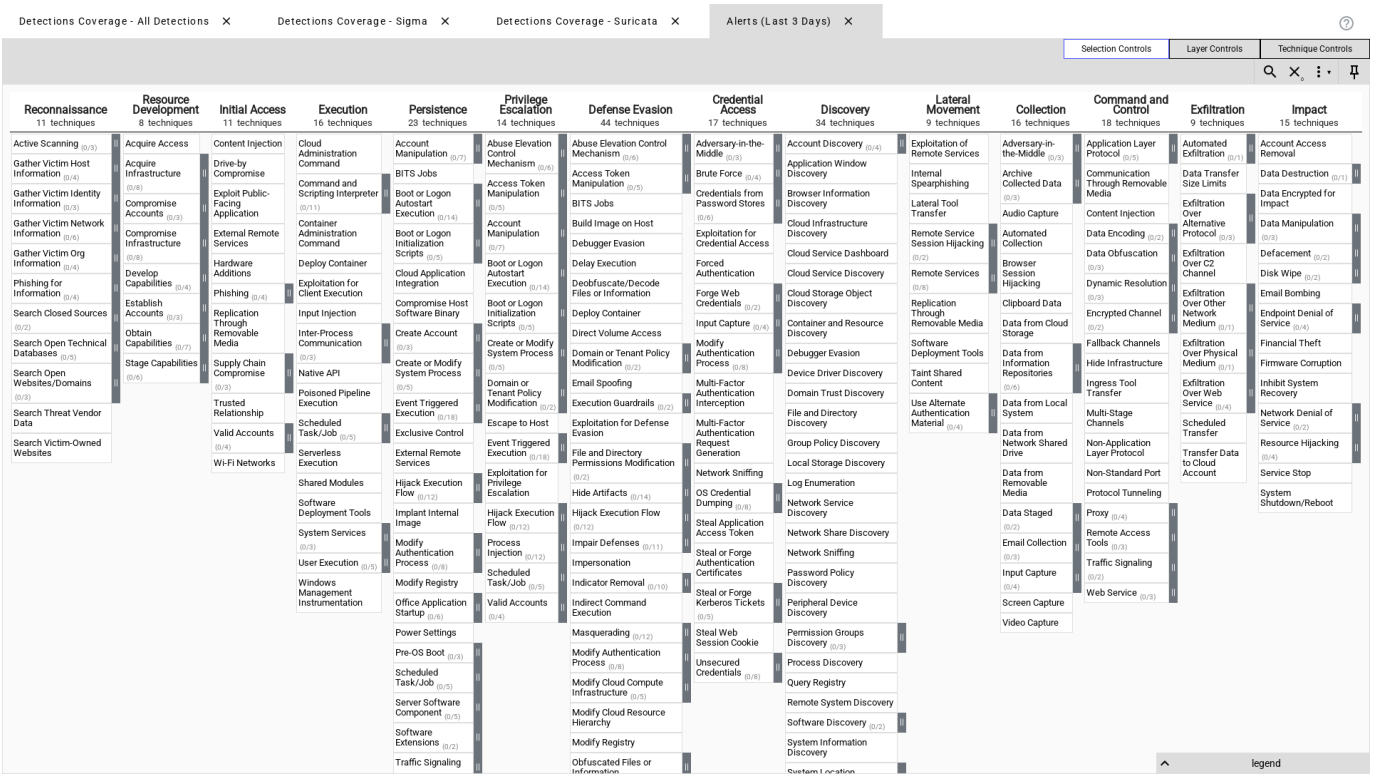

6.16.2 Configuration

Navigator reads its configuration from `/opt/so/conf/navigator/`. However, please keep in mind that if you make any changes here they may be overwritten since the config is managed with Salt.

6.16.3 More Information

Note

For more information about ATT&CK Navigator, please see: https://github.com/mitre-attack/attack-navigator

7. Security Onion Desktop

7.1 Security Onion Desktop Overview

Full-time analysts may want to use a dedicated Security Onion desktop. This allows you to investigate pcaps, malware, and other potentially malicious artifacts without impacting your Security Onion deployment or your usual desktop environment.

> **Note**
>
> Security Onion Desktop only supports Oracle Linux 9, so you'll either need to use our ISO image (recommended) or a network installation on top of Oracle Linux 9 (unsupported).

Security Onion Desktop consists of a full desktop environment including Chromium, NetworkMiner, Wireshark, and other analyst tools.

Installation

There are a few different ways to install Security Onion Desktop:

- Our ISO image includes a boot menu option for Desktop installs that will partition your disk appropriately and immediately perform a Desktop installation. The minimum disk size is 50GB.
- The `so-desktop-install` command is totally independent of the standard setup process, so you can run it before or after setup or not run setup at all if all you really want is the Analyst desktop itself.
- If you're doing a network installation on Oracle Linux 9 (NOT using our ISO image), then in our normal Setup wizard, you can choose `OTHER` and then choose `ANALYST`. Please note that network installations in general are unsupported.

> **Note**
>
> Depending on how you install, it may take a full Salt cycle before all desktop components are installed and ready for use.

Joining to Grid

You can optionally join your Desktop installation to your grid. This allows it to pull updates from the grid and automatically trust the grid's HTTPS certificate. It also updates the manager's firewall to allow the Desktop installation to connect. Desktop nodes display on the Grid page along with the other Grid nodes.

If you choose not to join your Desktop installation to your grid, then you may need to allow the traffic through the host-based firewall by going to Administration --> Configuration --> firewall --> hostgroups --> analyst.

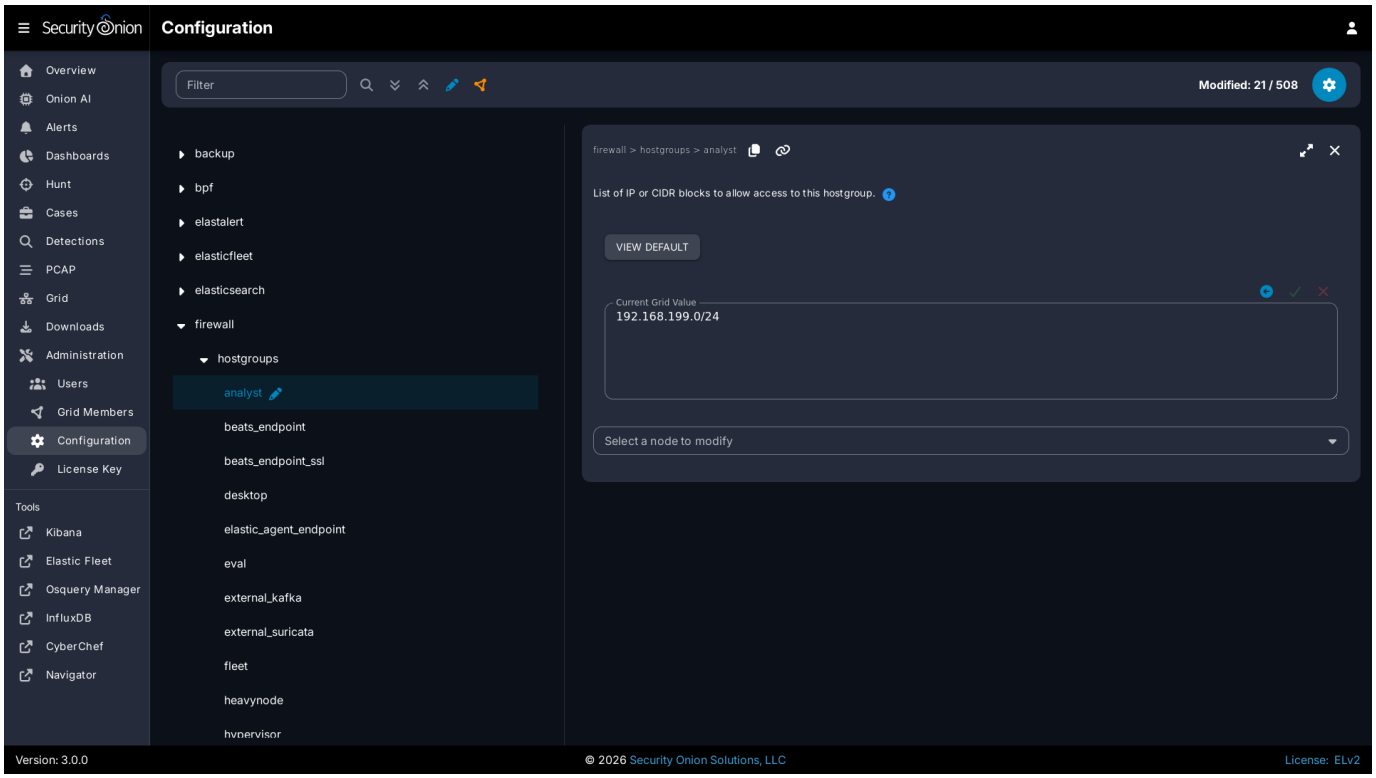

Disabling

The analyst desktop is controlled via Salt pillar. If you need to disable the Security Onion Desktop environment, find the `workstation` setting in your Salt pillar and change `enabled: true` to `enabled: false`:

```
workstation:
  gui:
    enabled: false
```

7.2 Chromium

Chromium is the web browser included in our Security Onion Desktop installation.

7.2.1 More Information

Note

For more information about Chromium, please see: https://www.chromium.org/Home/

7.3 NetworkMiner

From https://www.netresec.com/?page=NetworkMiner:

> NetworkMiner is an open source Network Forensic Analysis Tool (NFAT) for Windows (but also works in Linux / Mac OS X / FreeBSD). NetworkMiner can be used as a passive network sniffer/packet capturing tool in order to detect operating systems, sessions, hostnames, open ports etc. without putting any traffic on the network. NetworkMiner can also parse PCAP files for off-line analysis and to regenerate/reassemble transmitted files and certificates from PCAP files.
>
> NetworkMiner makes it easy to perform advanced Network Traffic Analysis (NTA) by providing extracted artifacts in an intuitive user interface. The way data is presented not only makes the analysis simpler, it also saves valuable time for the analyst or forensic investigator.

7.3.1 Usage

NetworkMiner is a part of our Security Onion Desktop installation.

7.3.2 File Extraction

Suppose you are looking at an interesting HTTP file download in PCAP and want to extract the file. Click the PCAP download button and then open the PCAP file with NetworkMiner. NetworkMiner will automatically attempt to detect and extract any files transferred. You can access these extracted files on the Files tab. If any files are images, they can be viewed on the Images tab.

7.3.3 More Information

Note

For more information about NetworkMiner, please see: https://www.netresec.com/?page=NetworkMiner

7.4 Wireshark

From https://www.wireshark.org/:

> Wireshark is the world's foremost and widely-used network protocol analyzer. It lets you see what's happening on your network at a microscopic level and is the de facto (and often de jure) standard across many commercial and non-profit enterprises, government agencies, and educational institutions. Wireshark development thrives thanks to the volunteer contributions of networking experts around the globe and is the continuation of a project started by Gerald Combs in 1998.

7.4.1 Usage

Wireshark is a part of our Security Onion Desktop installation.

7.4.2 File Extraction

Suppose you are looking at an interesting HTTP file download in PCAP and want to extract the file. Click the PCAP download button and then open the PCAP file with Wireshark. To extract files from HTTP traffic, click File - Export Objects - HTTP. Then select the file(s) to save and specify where to save them.

7.4.3 More Information

Note

For more information about Wireshark, please see https://www.wireshark.org/.

8. Network Visibility

8.1 Network Visibility Overview

When you log into Security Onion Console, you may see alerts from Suricata or IDH, protocol metadata logs from Zeek or Suricata, file analysis logs from Strelka, or full packet capture from Suricata. How is that data generated and stored? This section covers the various processes that Security Onion uses to analyze and log network traffic.

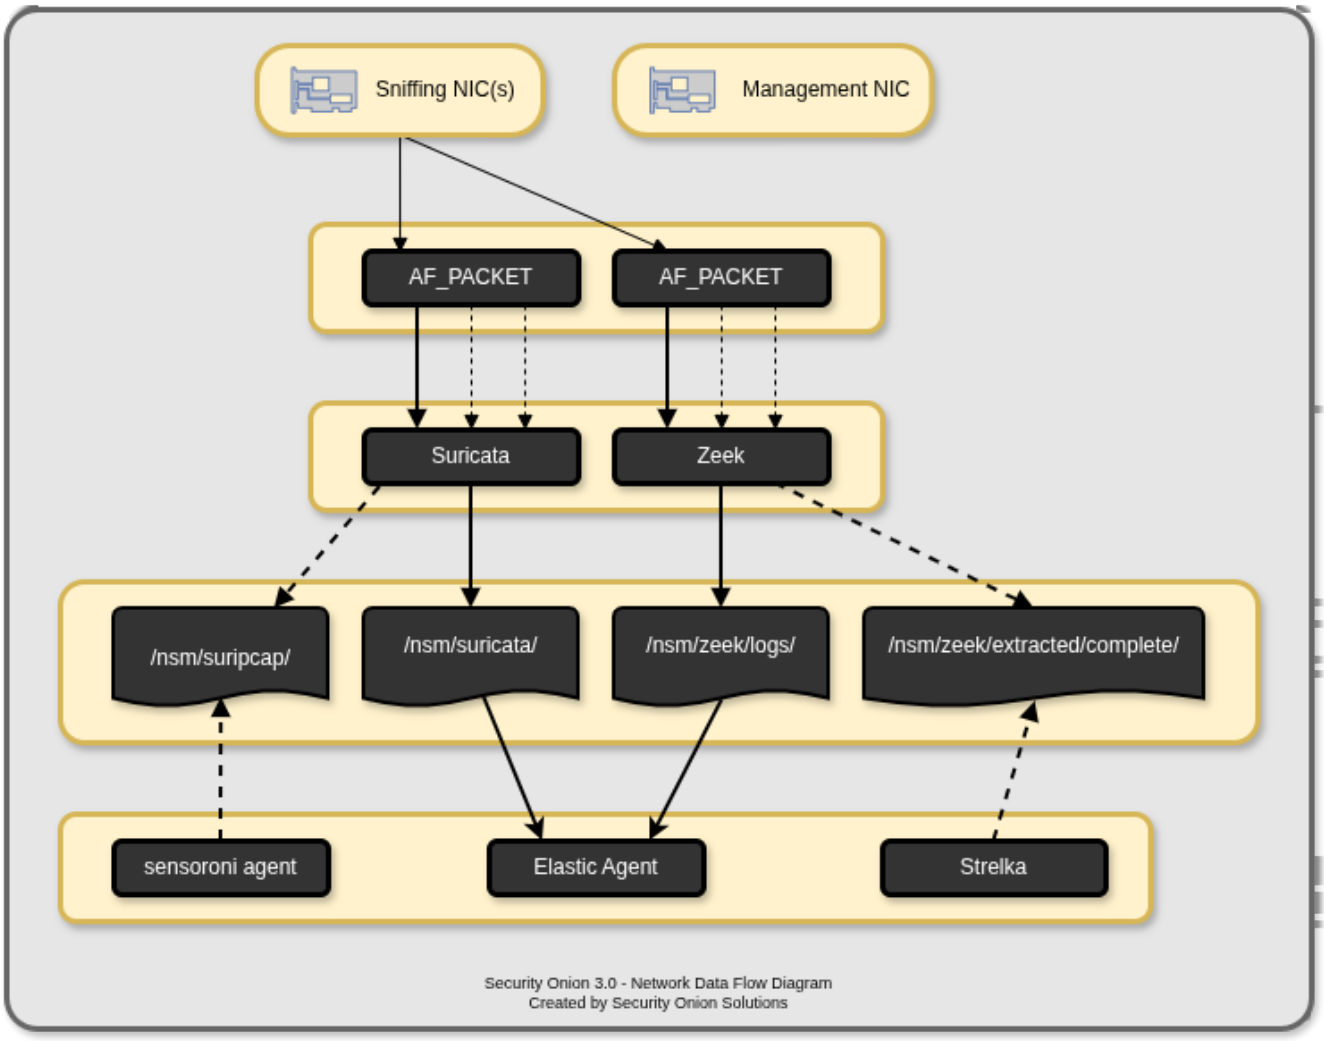

8.2 AF-PACKET

Security Onion uses AF-PACKET to collect traffic from network interfaces. AF-PACKET is built into the Linux kernel and includes fanout capabilities enabling it to act as a flow-based load balancer. This means, for example, if you configure Suricata for 4 AF-PACKET threads then each thread would receive about 25% of the total traffic that AF-PACKET is seeing.

> **Warning**
>
> If you try to test AF-PACKET fanout using tcpreplay locally, please note that load balancing will not work properly and all (or most) traffic will be handled by the first worker in the AF-PACKET cluster. If you need to test AF-PACKET load balancing properly, you can run tcpreplay on another machine connected to your AF-PACKET machine.

The following processes use AF-PACKET for traffic collection:

- Suricata
- Zeek

8.2.1 More Information

> **Note**
>
> For more information about AF-PACKET, please see: https://www.kernel.org/doc/Documentation/networking/packet_mmap.txt

8.3 BPF

BPF stands for Berkeley Packet Filter. From https://en.wikipedia.org/wiki/Berkeley_Packet_Filter:

> BPF supports filtering packets, allowing a userspace process to supply a filter program that specifies which packets it wants to receive. For example, a tcpdump process may want to receive only packets that initiate a TCP connection. BPF returns only packets that pass the filter that the process supplies. This avoids copying unwanted packets from the operating system kernel to the process, greatly improving performance.

8.3.1 Configuration

You can modify your BPF configuration by going to Administration --> Configuration --> bpf. You can apply BPF configuration to Suricata or Zeek. If you configure a BPF for Suricata, then it can apply to alerts, metadata, and/or Full Packet Capture.

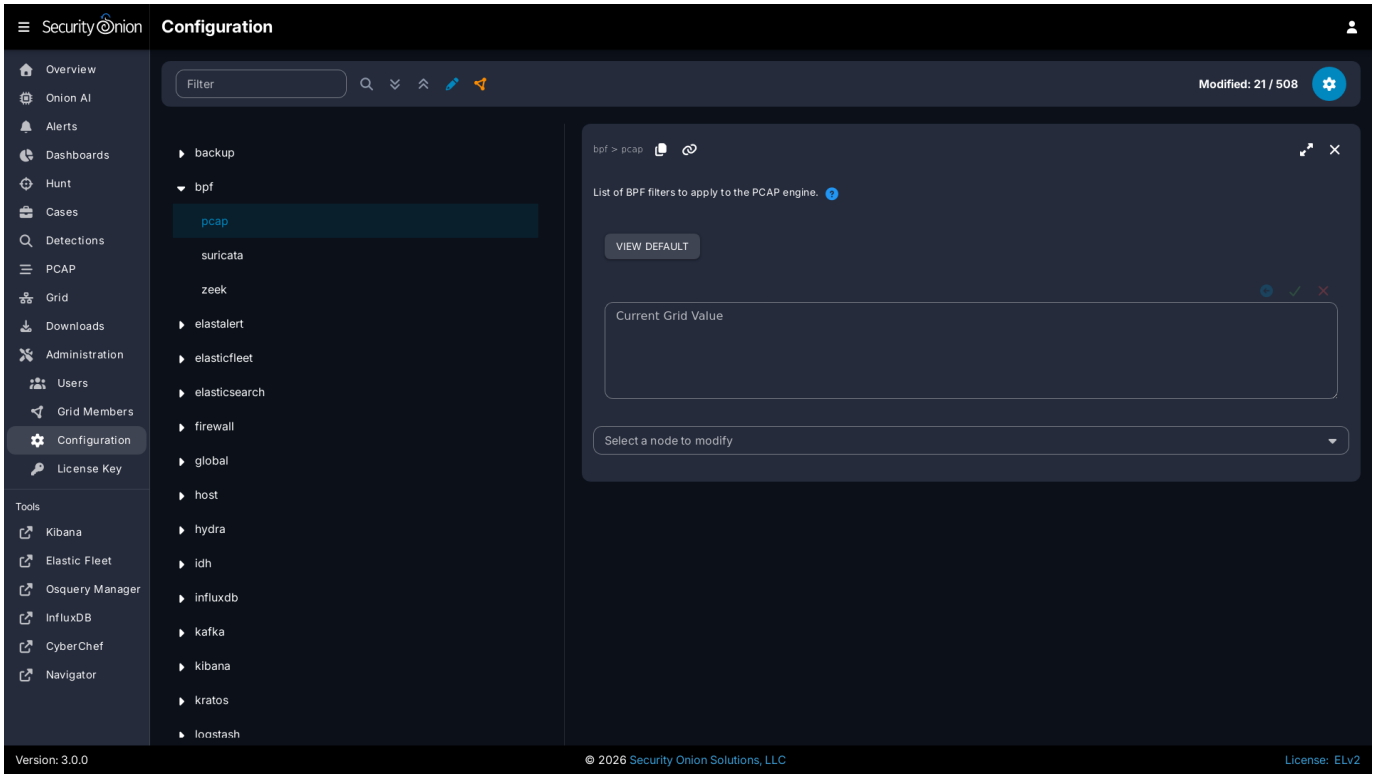

Multiple Conditions

If your BPF contains multiple conditions you can join them with a logical `and` or logical `or`.

Here's an example of joining conditions with a logical `and`:

```
not host 192.168.1.2 and not host 192.168.1.3
```

Here's an example of joining conditions with a logical `or`:

```
host 192.168.1.2 or host 192.168.1.3
```

VLAN

If you have traffic that has VLAN tags, you can craft a BPF as follows:

```
<your filter> or (vlan and <your filter>)
```

Notice that you must include your filter on both sides of the vlan tag.

For example:

```
(not (host 192.168.1.2 or host 192.168.1.3 or host 192.168.1.4)) or (vlan and (not (host 192.168.1.2 or host 192.168.1.3 or host 192.168.1.4)))
```

Adding Comments

To provide more context to your filters, you can add comments. For example:

```
# lab-east
not host 192.168.1.2 and not host 192.168.1.3 &&
# lab-west
not host 192.168.1.4 or not host 192.168.1.5 &&
# lab-central
not host 192.168.1.6 or not host 192.168.1.27
```

Troubleshooting BPF using tcpdump

If you need to troubleshoot BPF, you can use `tcpdump` as shown in the following articles: https://taosecurity.blogspot.com/2004/09/understanding-tcpdumps-d-option-have.html https://taosecurity.blogspot.com/2004/12/understanding-tcpdumps-d-option-part-2.html https://taosecurity.blogspot.com/2008/12/bpf-for-ip-or-vlan-traffic.html

8.3.2 More Information

Note

For more information about BPF, please see: https://en.wikipedia.org/wiki/Berkeley_Packet_Filter https://biot.com/capstats/bpf.html

8.4 Full Packet Capture

Security Onion writes network traffic to disk for full packet capture using Suricata.

Suricata writes standard PCAP files which can be copied off to another system and then opened with any standard libpcap tool.

8.4.1 Analysis

You can access full packet capture via the PCAP interface:

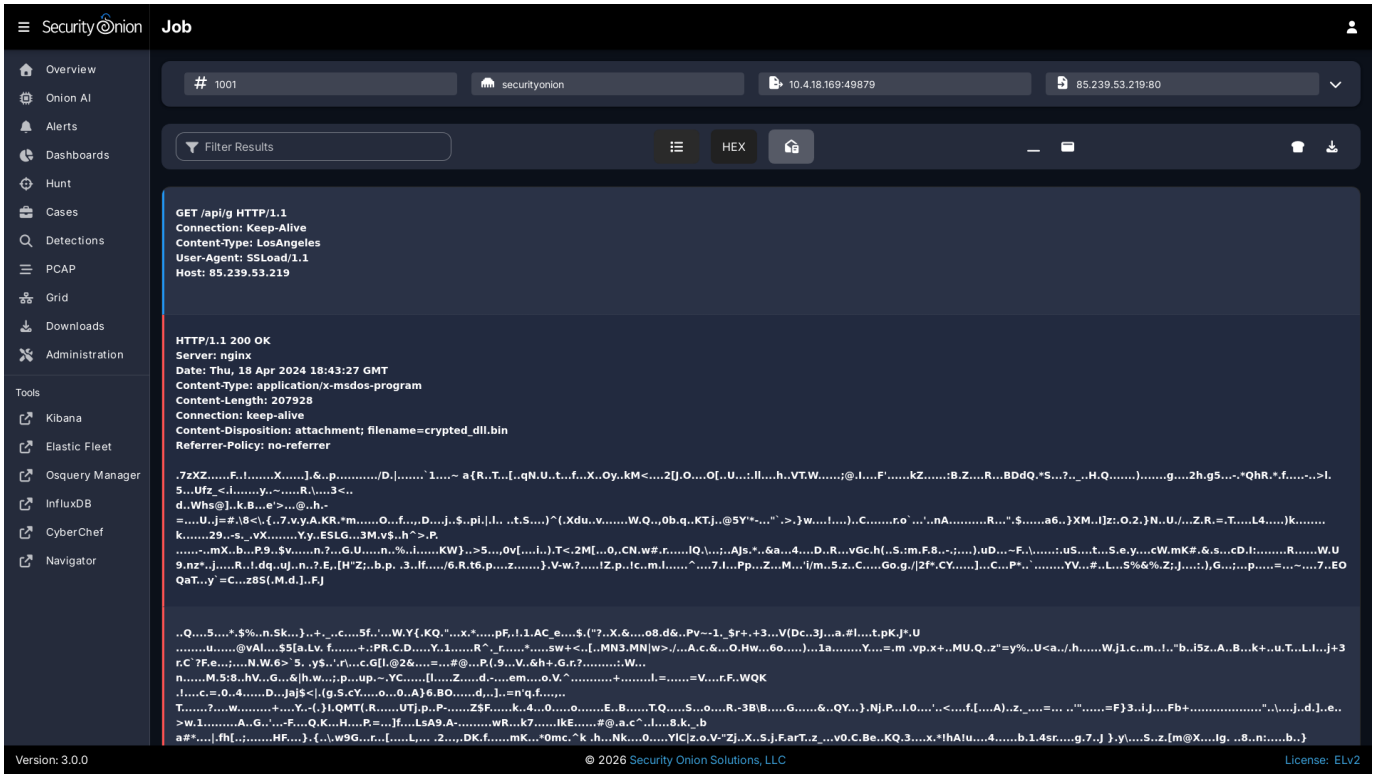

Alerts, Dashboards, and Hunt allow you to easily pivot to the PCAP interface.

8.4.2 Configuration

You can configure full packet capture by going to Administration --> Configuration --> Suricata --> config --> pcap.

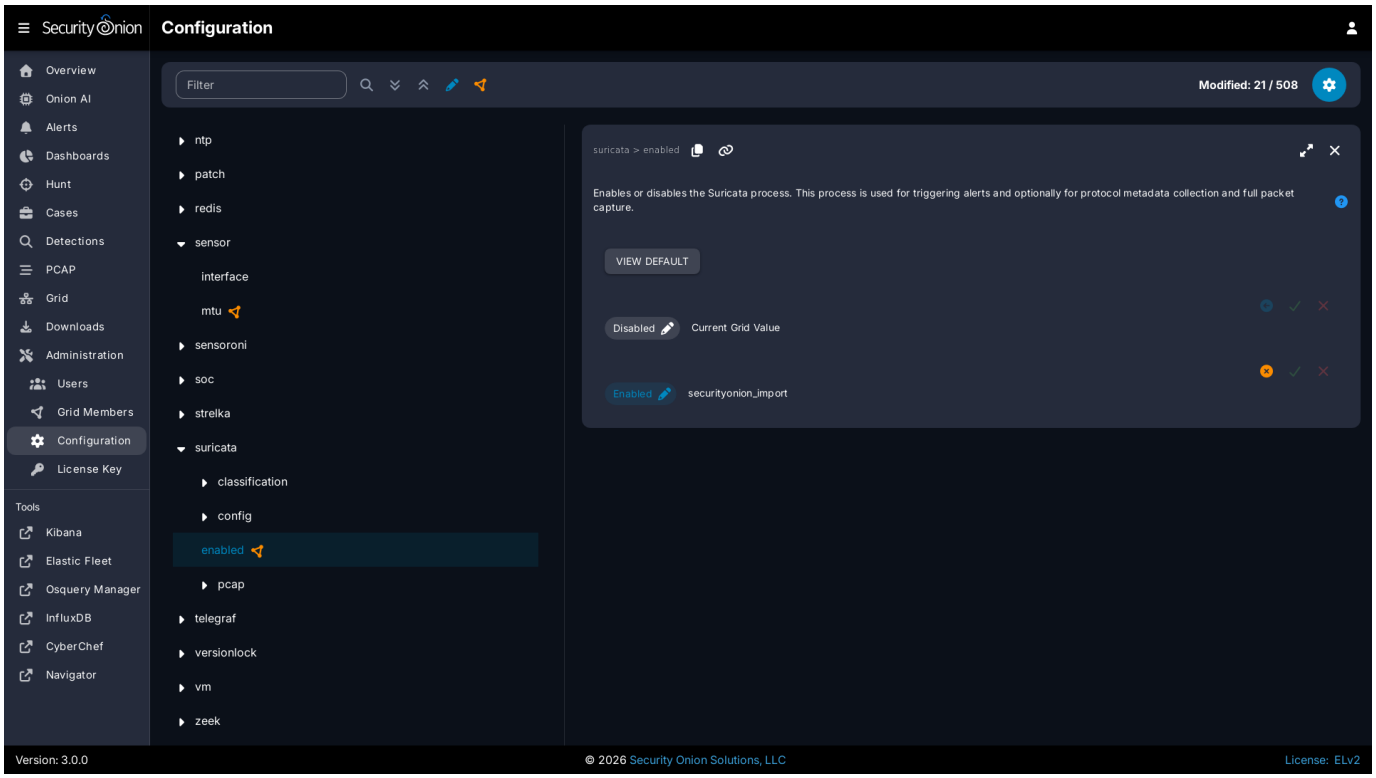

Conditional PCAP

By default, Suricata writes all network traffic to PCAP. If you would like to limit Suricata to only writing PCAP when certain conditions are met, you can go to Administration --> Configuration --> Suricata -> config -> pcap -> conditional and change it to to either `alerts` or `tag`:

- `all`: Capture all packets seen by Suricata (default).
- `alerts`: Capture only packets associated with a NIDS alert.
- `tag`: Capture packets based on a rule that is tagged.

PCAP Configuration Options

Here are some other PCAP configuration options that can be found at Administration --> Configuration --> Suricata -> config -> pcap. Some settings are considered advanced settings so you will only see them if you enable the `Show advanced settings` option.

- `enabled`: Click the slider to enable or disable Suricata packet capture.
- `compression`: Set to `none` to disable compression. Set to `lz4` to enable lz4 compression but note that this requires more CPU cycles.
- `lz4-level`: lz4 compression level of PCAP files. Set to `0` for no compression. Set to `16` for maximum compression.
- `maxsize`: Maximum size in GB for total disk usage of all PCAP files written by Suricata. You may need to adjust this value based on your disk space and desired PCAP retention.
- `filesize`: Maximum file size for individual PCAP files written by Suricata. Increasing this number could improve write performance at the expense of PCAP retrieval time.
- `use-stream-depth`: Set to `no` to ignore the stream depth and capture the entire flow. Set to `yes` to truncate the flow based on the stream depth.

8.4.3 VLAN Tags

If your network traffic has VLAN tags, then Suricata will log them. For more information, please see the Suricata section.

8.4.4 Diagnostic Logging

Diagnostic logging for Suricata can be found at `/opt/so/log/suricata/suricata.log`. Depending on what you're looking for, you may also need to look at the Docker logs for the container:

```
sudo docker logs so-suricata
```

8.5 Suricata

From https://suricata.io:

> Suricata is a free and open source, mature, fast and robust network threat detection engine. Suricata inspects the network traffic using a powerful and extensive rules and signature language, and has powerful Lua scripting support for detection of complex threats.

Suricata NIDS alerts can be found in Alerts, Dashboards, Hunt, and Kibana.

Here's an example of Suricata NIDS alerts in Alerts:

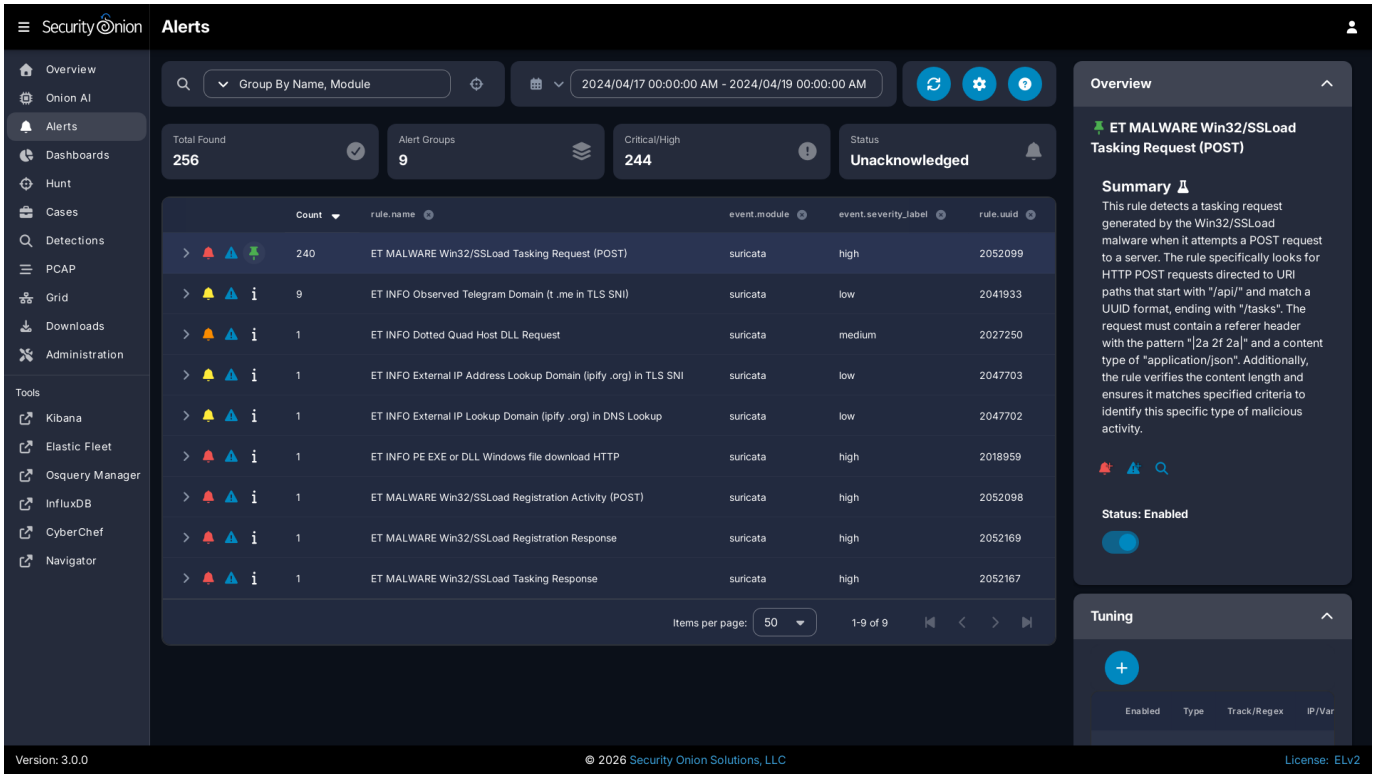

If enabled, Suricata metadata (protocol logs) can be found in Dashboards, Hunt, and Kibana.

8.5.1 Community ID

Security Onion enables Suricata's built-in support for Community ID.

8.5.2 VLAN Tags

If your network traffic has VLAN tags, then Suricata will log them. Dashboards has a VLAN dashboard which will show this data.

If your network traffic has mixed VLAN tags (VLAN tags in one direction but not the other), then you may need to do the following:

- Navigate to Administration --> Configuration.
- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- Navigate to Suricata > config > vlan > use-for-tracking and set it to `false`.

8.5.3 Configuration

You can configure Suricata by going to Administration --> Configuration --> Suricata.

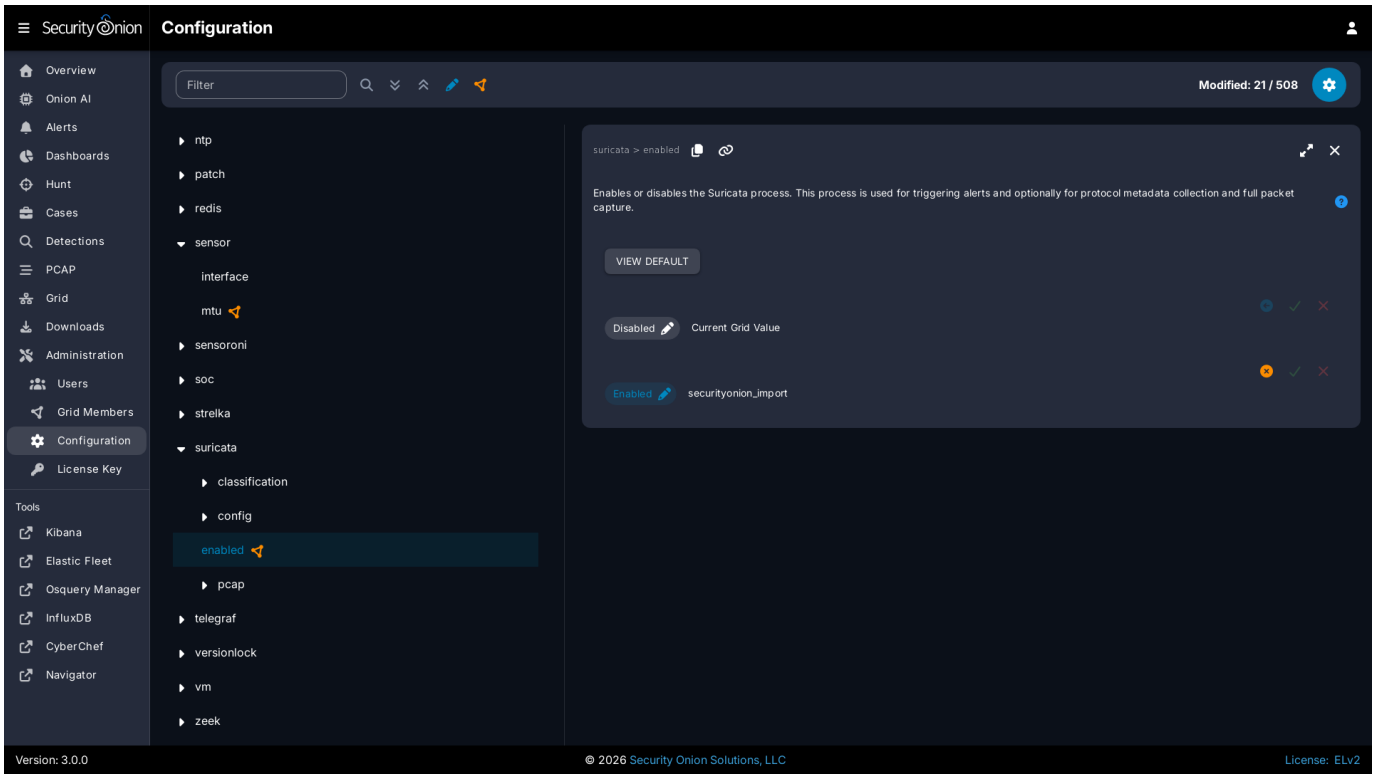

If you would like to configure NIDS rules, take a look at the Detections interface.

8.5.4 HOME_NET

The HOME_NET variable defines the networks that are considered home networks (those networks that you are monitoring and defending). The default value is RFC1918 private address space (10.0.0.0/8, 192.168.0.0/16, and 172.16.0.0/12). You can modify this default value by going to Administration --> Configuration --> Suricata --> config --> vars --> address-groups --> HOME_NET.

8.5.5 EXTERNAL_NET

By default, EXTERNAL_NET is set to `any` (which includes `HOME_NET`) to detect lateral movement inside your environment. You can modify this default value by going to Administration --> Configuration --> Suricata --> config --> vars --> address-groups --> EXTERNAL_NET.

8.5.6 Stats

For Suricata statistics, see Grid, InfluxDB, and `/opt/so/log/suricata/stats.log`.

8.5.7 Performance

If Grid shows that Suricata is experiencing packet loss, then you may need to do one or more of the following:

- Tune the NIDS ruleset.
- Apply a BPF.
- Adjust `max-pending-packets` in Administration --> Configuration --> Suricata --> config --> max-pending-packets.
- Adjust AF-PACKET workers in Administration --> Configuration --> Suricata --> config --> af-packet --> threads.

Note

For other tuning considerations, please see: https://suricata.readthedocs.io/en/latest/performance/tuning-considerations.html

If you have multiple physical CPUs, you'll most likely want to pin sniffing processes to a CPU in the same Non-Uniform Memory Access (NUMA) domain that your sniffing NIC is bound to. Accessing a CPU in the same NUMA domain is faster than across a NUMA domain.

> **Note**
>
> For more information about determining NUMA domains using `lscpu` and `lstopo`, please see: https://github.com/brokenscripts/cpu_pinning

8.5.8 Metadata

By default, Security Onion uses Zeek to record protocol metadata. If you don't need all of the protocol coverage that Zeek provides, then you can switch to Suricata metadata to save some CPU cycles. If you choose to do this, then here are some of the kinds of metadata you can expect to see in Dashboards or Hunt:

- Connections
- DHCP
- DNS
- Files
- FTP
- HTTP
- SSL

If you later find that some of that metadata is unnecessary, you can enable the SO_FILTERS ruleset to filter out unnecessary metadata. Navigate to Administration --> Configuration --> SOC --> config --> server --> modules --> suricataengine --> rulesetSources and enable the SO_FILTERS ruleset.

To change your grid's metadata engine from Zeek to Suricata, go to Administration --> Configuration --> global --> mdengine and change the value from `Zeek` to `Suricata`:

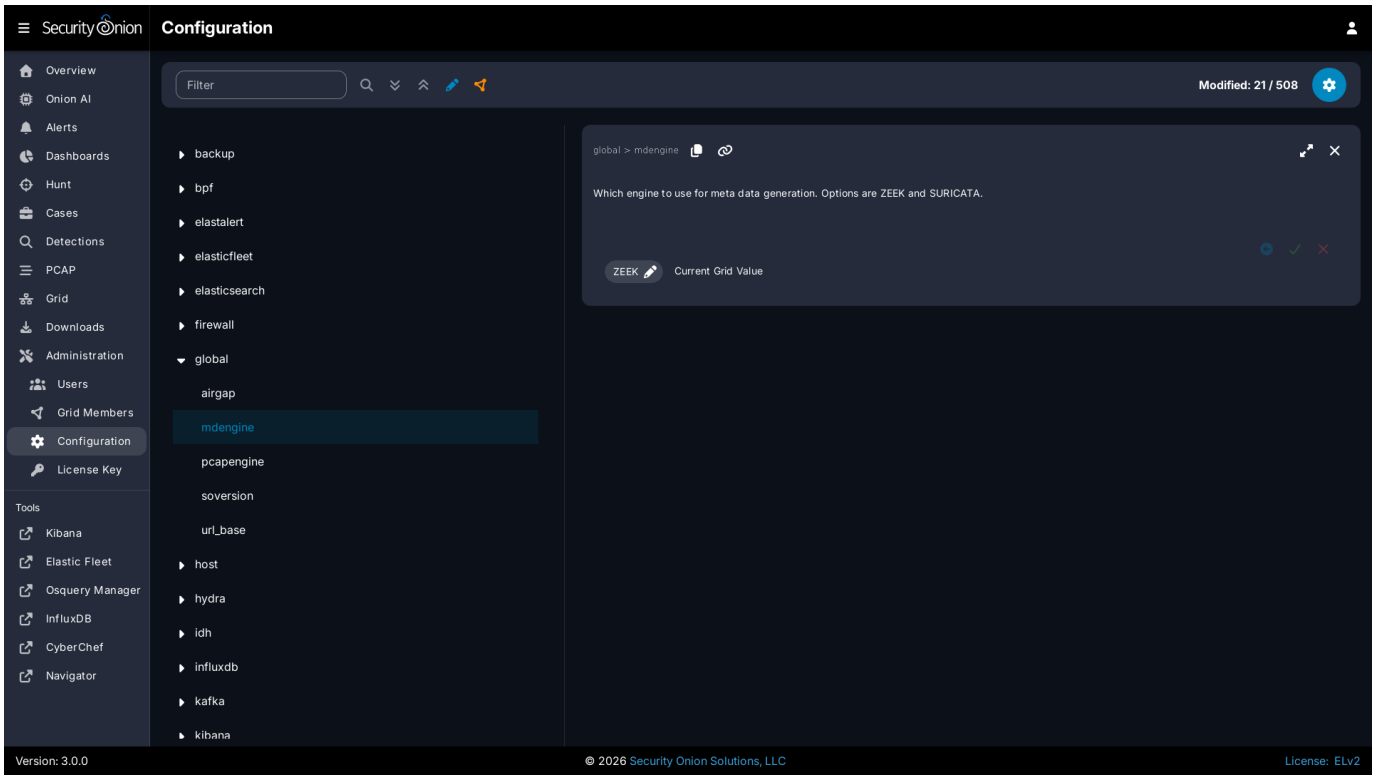

8.5.9 File Extraction

If you choose Suricata for metadata, it will extract files from network traffic and Strelka will then analyze those extracted files. The SO_EXTRACTIONS ruleset controls which file types are extracted and is enabled by default when Suricata is the metadata engine. You can customize file extraction by modifying rules in this ruleset.

8.5.10 PCAP

Security Onion uses Suricata for full packet capture. For more information, please see the Full Packet Capture section.

8.5.11 Diagnostic Logging

If you need to troubleshoot Suricata, check `/opt/so/log/suricata/suricata.log`. Depending on what you're looking for, you may also need to look at the Docker logs for the container:

```
sudo docker logs so-suricata
```

8.5.12 Testing

The first and easiest way to test Suricata is to access http://testmynids.org/uid/index.html from a machine that is being monitored by your Security Onion deployment. You can do so via the command line using `curl`:

```
curl testmynids.org/uid/index.html
```

If everything is working correctly, you should see a corresponding alert (`GPL ATTACK_RESPONSE id check returned root`) in Alerts. You should also be able to find the alert in Dashboards or Hunt.

If you do not see this alert, try checking to see if the rule is enabled by going to Detections and searching for the SID of the rule which is `2100498`. One way to search for this rule is to specify it in the URL as follows:

`#/detections?q=2100498`

Another way to test Suricata is with a utility called `tmNIDS`. You can run the tool in interactive mode like this:

```
curl -sSL https://raw.githubusercontent.com/0xtf/testmynids.org/master/tmNIDS -o /tmp/tmNIDS && chmod +x /tmp/tmNIDS && /tmp/tmNIDS
```

Finally, you can also test Suricata alerting by replaying some test PCAP files via so-test.

8.5.13 Troubleshooting Alerts

If you're not seeing the Suricata alerts that you expect to see, here are some things that you can check:

- If you have metadata enabled, check to see if you have metadata for the connections. Depending on your configuration, this could be Suricata metadata or Zeek metadata. Go to Dashboards, click the dropdown menu, select the `Connections seen by Zeek or Suricata` dashboard, and see if the connections you expect to see in your network traffic are listed there.
- If you have metadata enabled but aren't seeing any metadata, then something may be preventing the process from seeing the traffic. Check to see if you have any BPF configuration that may cause the process to ignore the traffic. If you're sniffing traffic from the network, verify that the traffic is reaching the NIC using tcpdump. If importing a PCAP file, verify that file contains the traffic you expect and that the Suricata process can read the file and any parent directories.
- Check to see if you have mixed VLAN tags (VLAN tags in one direction but not the other). If so, see the VLAN Tags section above to configure Suricata appropriately.
- Check your HOME_NET configuration to make sure it includes the networks that you're watching traffic for.
- Check to see if you have a full NIDS ruleset with rules that should specifically alert on the traffic and that those rules are enabled.
- Check to see if you have any threshold or suppression configuration that might be preventing alerts.
- Check the Suricata log for additional clues.
- Check the Elastic Agent, Logstash, and Elasticsearch logs for any pipeline issues that may be preventing the alerts from being written to Elasticsearch.
- Try installing a simple import node (perhaps in a VM) following the steps in the First Time Users section and see if you get alerts there. If so, compare the working system to the non-working system and determine where the differences are.

8.5.14 Testing Rules

To test a new rule, use the following utility on a node that runs Suricata (ie Sensor or Import).

```
sudo so-suricata-testrule <Filename> /path/to/pcap/test.PCAP
```

The file should contain the new rule that you would like to test. The PCAP should contain network data that will trigger the rule.

8.5.15 Variables

To add or modify Suricata Variables, navigate to **Suricata > config > vars > address-groups** or **port-groups**.

You can assign a list of hosts, networks, or other customizations to a Suricata variable. The variable can then be re-used within Suricata rules and/or Overrides. This allows for a single adjustment to the variable that will affect all rules referencing it.

Address Groups

Address groups define IP addresses or network ranges. Suricata comes with a number of common address groups already defined.

ADDRESS GROUP SYNTAX

Values can be specified using the following formats (single or multi-line):

- Single IP: `192.168.1.100`
- CIDR notation: `192.168.1.0/24`
- IP range: `192.168.1.1-192.168.1.50`
- Multiple values: `192.168.1.0/24,10.0.0.0/8`
- Negation: `!192.168.1.100` or `!$OTHER_VAR`
- Variable reference: `$HOME_NET`

Port Groups

Port groups define TCP/UDP ports for specific services. Suricata comes with a number of common port groups already defined.

PORT GROUP SYNTAX

Values can be specified using the following formats (single or multi-line):

- Single port: `80`
- Port range: `1024:65535`
- Multiple ports: `80,443,8080`
- Negation: `!80`
- Any port: `any`

Custom Variables

Create custom variables by selecting an existing Address Group or Port Group and clicking "Duplicate". Enter a name using uppercase naming convention, then click "Create Setting".

Note: The new variable is not saved until you modify its value and click the green "Save Changes" checkmark.

Note: Custom Port Group variables have some syntactical limitations, for example, not being able to use negation.

8.5.16 Disabling

If you need to disable Suricata, you can do so via Administration --> Configuration --> Suricata --> enabled.

8.5.17 More Information

> **Note**
>
> For more information about Suricata, please see https://suricata.io.

8.6 Zeek

Security Onion includes Zeek for network protocol analysis and metadata. From https://zeek.org:

> Zeek is a powerful network analysis framework that is much different from the typical IDS you may know. (Zeek is the new name for the long-established Bro system. Note that parts of the system retain the "Bro" name, and it also often appears in the documentation and distributions.)

Zeek logs are sent to Elasticsearch for parsing and storage and can then be found in Dashboards, Hunt, and Kibana. Here's an example of Zeek logs in Hunt:

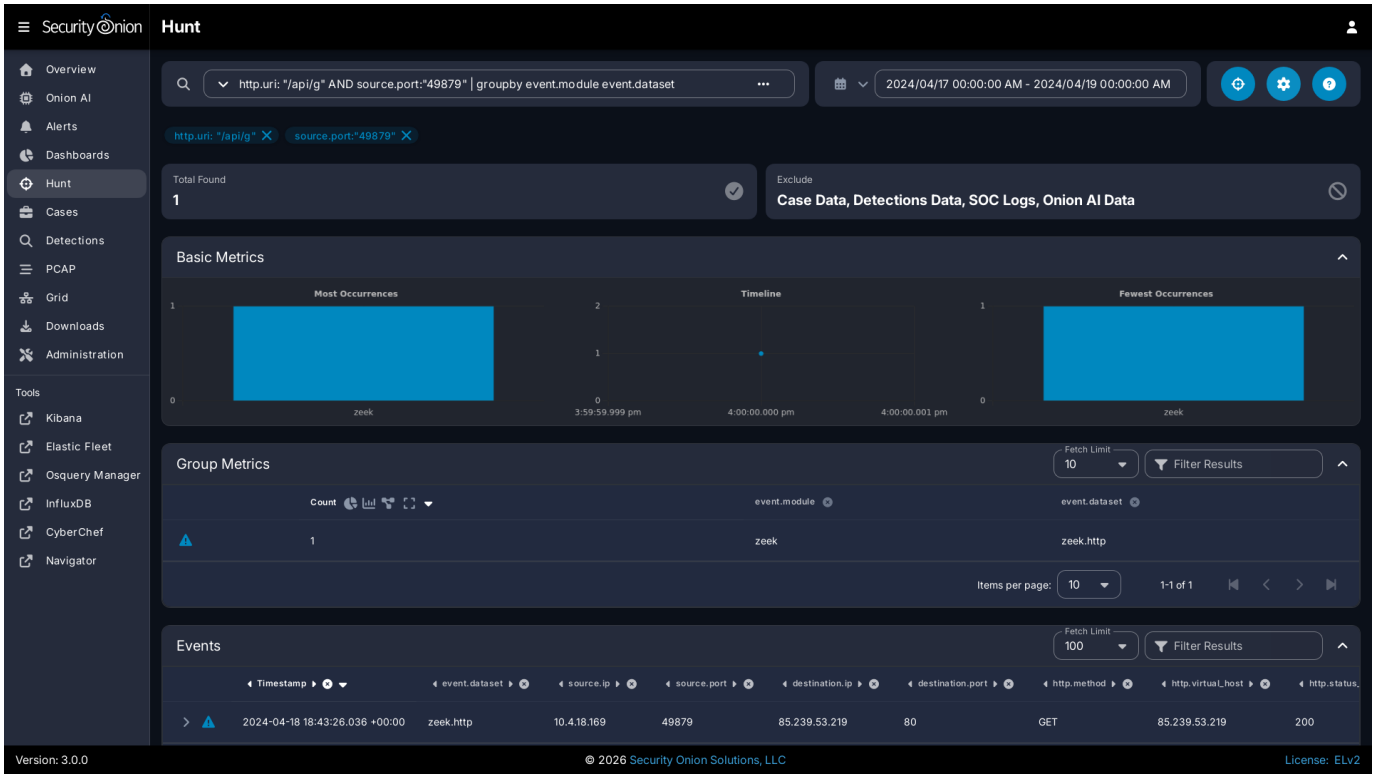

8.6.1 Community ID

Security Onion enables Zeek's built-in support for Community ID.

8.6.2 Packet Loss and Capture Loss

Zeek reports both packet loss and capture loss and you can find graphs of these in InfluxDB. If Zeek reports packet loss, then you most likely need to adjust the number of Zeek workers as shown below or filter out traffic using BPF. If Zeek is reporting capture loss but no packet loss, this usually means that the capture loss is happening upstream in the TAP or SPAN port itself.

8.6.3 Configuration

You can configure Zeek by going to Administration --> Configuration --> Zeek.

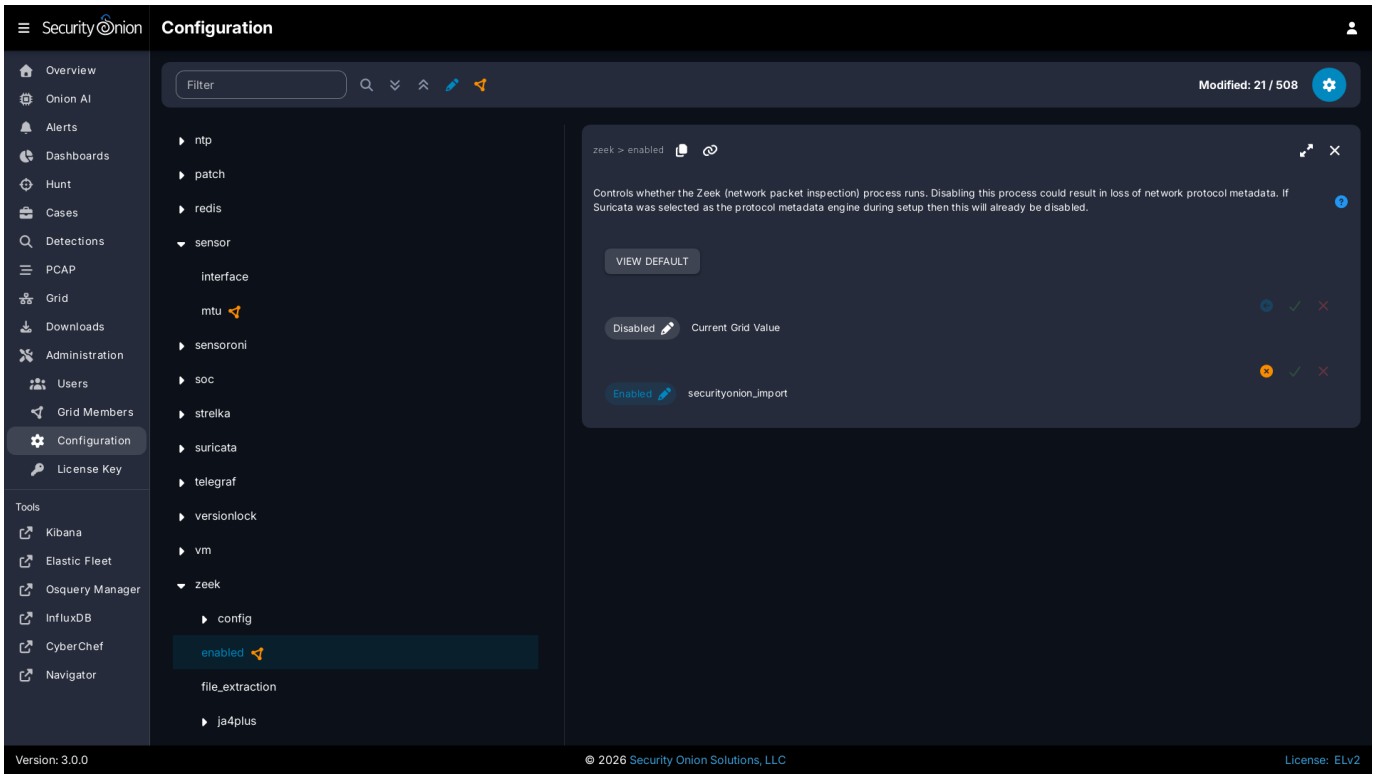

Zeek logs are consumed by the Elastic Agent (managed by Elastic Fleet) so if you want to configure which Zeek logs are excluded, you can go to Administration --> Configuration --> elasticfleet --> logging --> Zeek --> excluded.

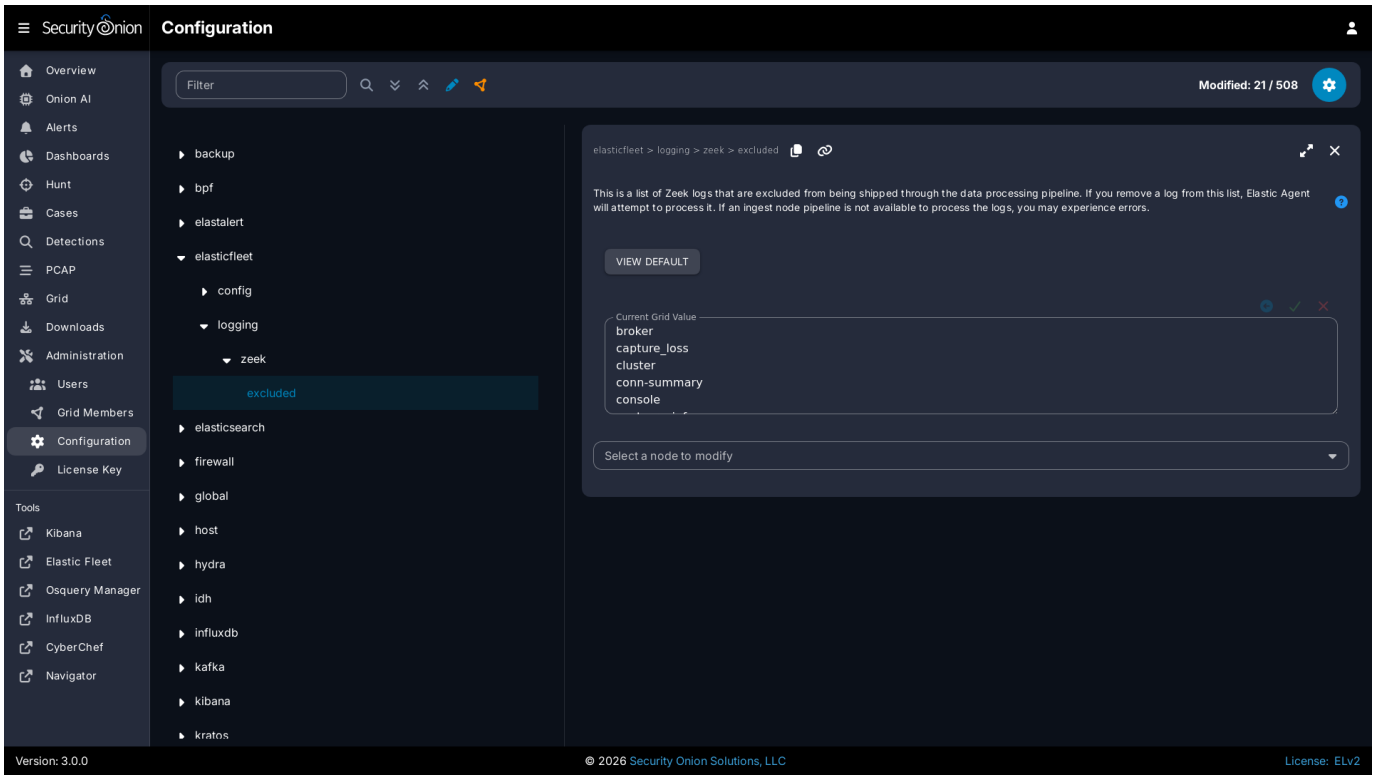

8.6.4 HOME_NET

The HOME_NET variable defines the networks that are considered home networks (those networks that you are monitoring and defending). The default value is RFC1918 private address space (10.0.0.0/8, 192.168.0.0/16, and 172.16.0.0/12). You can modify this default value by going to Administration --> Configuration --> Zeek --> config --> networks --> HOME_NET.

8.6.5 Performance

Zeek uses AF-PACKET so that you can spin up multiple Zeek workers to handle more traffic.

If you have multiple physical CPUs, you'll most likely want to pin sniffing processes to a CPU in the same Non-Uniform Memory Access (NUMA) domain that your sniffing NIC is bound to. Accessing a CPU in the same NUMA domain is faster than across a NUMA domain.

> **Note**
>
> For more information about determining NUMA domains using `lscpu` and `lstopo`, please see https://github.com/brokenscripts/cpu_pinning.

You can modify Zeek worker count by going to Administration --> Configuration --> Zeek --> config --> node --> workers.

8.6.6 Disabling

Zeek can be disabled by going to Administration --> Configuration --> Zeek --> enabled.

8.6.7 Syslog

To forward Zeek logs to an external syslog collector, please see the syslog-output section.

8.6.8 Logs

Zeek logs are stored in `/nsm/zeek/logs`. They are collected by Elastic Agent, parsed by and stored in Elasticsearch, and viewable in Dashboards, Hunt, and Kibana.

We configure Zeek to output logs in JSON format. If you need to parse those JSON logs from the command line, you can use jq.

Zeek monitors your network traffic and logs protocol metadata. Here are just a few examples.

conn.log

- TCP/UDP/ICMP connections
- For ICMP connections, ICMP type is logged as source port and ICMP code as destination port
- For more information, see https://docs.zeek.org/en/latest/scripts/base/protocols/conn/main.zeek.html#type-Conn::Info.

dns.log

- DNS activity
- For more information, see https://docs.zeek.org/en/latest/scripts/base/protocols/dns/main.zeek.html#type-DNS::Info.

http.log

- HTTP requests and replies
- For more information, see https://docs.zeek.org/en/latest/scripts/base/protocols/http/main.zeek.html#type-HTTP::Info.

ssl.log

- SSL/TLS handshake info
- For more information, see https://docs.zeek.org/en/latest/scripts/base/protocols/ssl/main.zeek.html#type-SSL::Info.

notice.log

- Zeek notices

- For more information, see https://docs.zeek.org/en/latest/scripts/base/frameworks/notice/main.zeek.html#type-Notice::Info.

Other Zeek logs

Zeek also provides other logs by default and you can read more about them at https://docs.zeek.org/en/latest/script-reference/log-files.html.

In addition to Zeek's default logs, we also include protocol analyzers for STUN, TDS, and Wireguard traffic. We also include support for ICS/SCADA protocols such as BACnet, BSAP, CIP, COTP, DNP3, ECAT, ENIP, Modbus, OPC UA, Profinet, and S7. All of these analyzers are enabled by default and you can find corresponding dashboards for each of them in Dashboards.

We also include MITRE BZAR scripts and you can read more about them at https://github.com/mitre-attack/bzar. Please note that the MITRE BZAR scripts are disabled by default. If you would like to enable them, you can do so via Administration --> Configuration --> Zeek. Once enabled, you can then check for BZAR detections by going to Dashboards and selecting the Zeek Notice dashboard.

As you can see, Zeek log data can provide a wealth of information to the analyst, all easily accessible through Dashboards, Hunt, or Kibana.

8.6.9 File Extraction

By default, Zeek will extract files from network traffic and Strelka will then analyze those extracted files.

8.6.10 VLAN Tags

If your network traffic has VLAN tags, then Zeek will log them in conn.log. Dashboards includes a VLAN dashboard which shows this data.

8.6.11 Intel

You can add your own intel to `/opt/so/saltstack/local/salt/zeek/policy/intel/intel.dat` on the manager and it will automatically replicate to all sensor nodes. If the `/opt/so/saltstack/local/salt/zeek/policy/intel/` directory is empty, you can copy the default files (both `intel.dat` and `__load__.Zeek`) from `/opt/so/saltstack/default/salt/zeek/policy/intel/` as follows:

```
sudo cp /opt/so/saltstack/default/salt/zeek/policy/intel/* /opt/so/saltstack/local/salt/zeek/policy/intel/
```

Please note that Zeek is very strict about the format of `intel.dat`. When editing this file, please follow these guidelines:

- no leading spaces or lines
- separate fields with a single literal tab character
- no trailing spaces or lines

The default `intel.dat` file follows these guidelines so you can reference it as an example of the proper format.

When finished editing `intel.dat`, run `sudo salt $SENSORNAME_$ROLE state.highstate` to sync `/opt/so/saltstack/local/salt/zeek/policy/intel/` to `/opt/so/conf/zeek/policy/intel/`. If you have a distributed deployment with separate sensor nodes, it may take up to 15 minutes for intel to sync to the sensor nodes.

If you experience an error, or do not notice `/nsm/zeek/logs/current/intel.log` being generated, try having a look in `/nsm/zeek/logs/current/reporter.log` for clues. You may also want to restart Zeek after making changes by running `sudo so-Zeek-restart`.

For more information, please see https://docs.zeek.org/en/latest/frameworks/intel.html.

8.6.12 Custom Packages

You can install custom Zeek packages using `zkg`. Place each package as a subdirectory in `/opt/so/saltstack/local/salt/zeek/zkg/` on the manager. For example, if you have a custom package called `my-package`, the directory structure would be:

```
/opt/so/saltstack/local/salt/zeek/zkg/my-package/
```

The package directory should contain a valid Zeek package structure (including a `zkg.meta` file).

> **Note**
>
> zkg requires packages from git repos to be a clone of the repo. The working tree must be clean otherwise it will not install the package. `git status` will tell you if it is clean or not.

After placing the package, run `sudo salt $SENSORNAME_$ROLE state.highstate` to sync the packages to the sensor nodes. The packages will be automatically installed with `zkg` each time the Zeek container starts. If you have a distributed deployment with separate sensor nodes, it may take up to 15 minutes for packages to sync to the sensor nodes.

You can verify that a custom package was installed by checking the Zeek container logs:

```
sudo docker logs so-zeek
```

8.6.13 Diagnostic Logging

Zeek diagnostic logs can be found in `/nsm/zeek/logs/`. Look for files like `reporter.log`, `stats.log`, `stderr.log`, and `stdout.log`. Depending on what you're looking for, you may also need to look at the Docker logs for the container:

```
sudo docker logs so-Zeek
```

8.6.14 More Information

> **Note**
>
> For more information about Zeek, please see https://zeek.org.

8.7 Strelka

From https://github.com/target/strelka:

> Strelka is a real-time file scanning system used for threat hunting, threat detection, and incident response. Based on the design established by Lockheed Martin's Laika BOSS and similar projects (see: related projects), Strelka's purpose is to perform file extraction and metadata collection at huge scale.

If you are monitoring network traffic, then either Zeek or Suricata should be extracting certain files detected in unencrypted network traffic. Strelka then analyzes those files and they end up in `/nsm/strelka/processed/`.

Security Onion checks file hashes before sending to Strelka to avoid analyzing the same file multiple times in a 48 hour period.

8.7.1 Alerts

Strelka scans files using YARA rules. If it detects a match, then it will generate an alert that can be found in Alerts, Dashboards, or Hunt. You can configure YARA rules via Detections.

8.7.2 Logs

Even if Strelka doesn't detect a YARA match, it will still log metadata about the file. You can find Strelka logs in Dashboards and Hunt.

8.7.3 Configuration

You can configure Strelka by going to Administration --> Configuration --> Strelka.

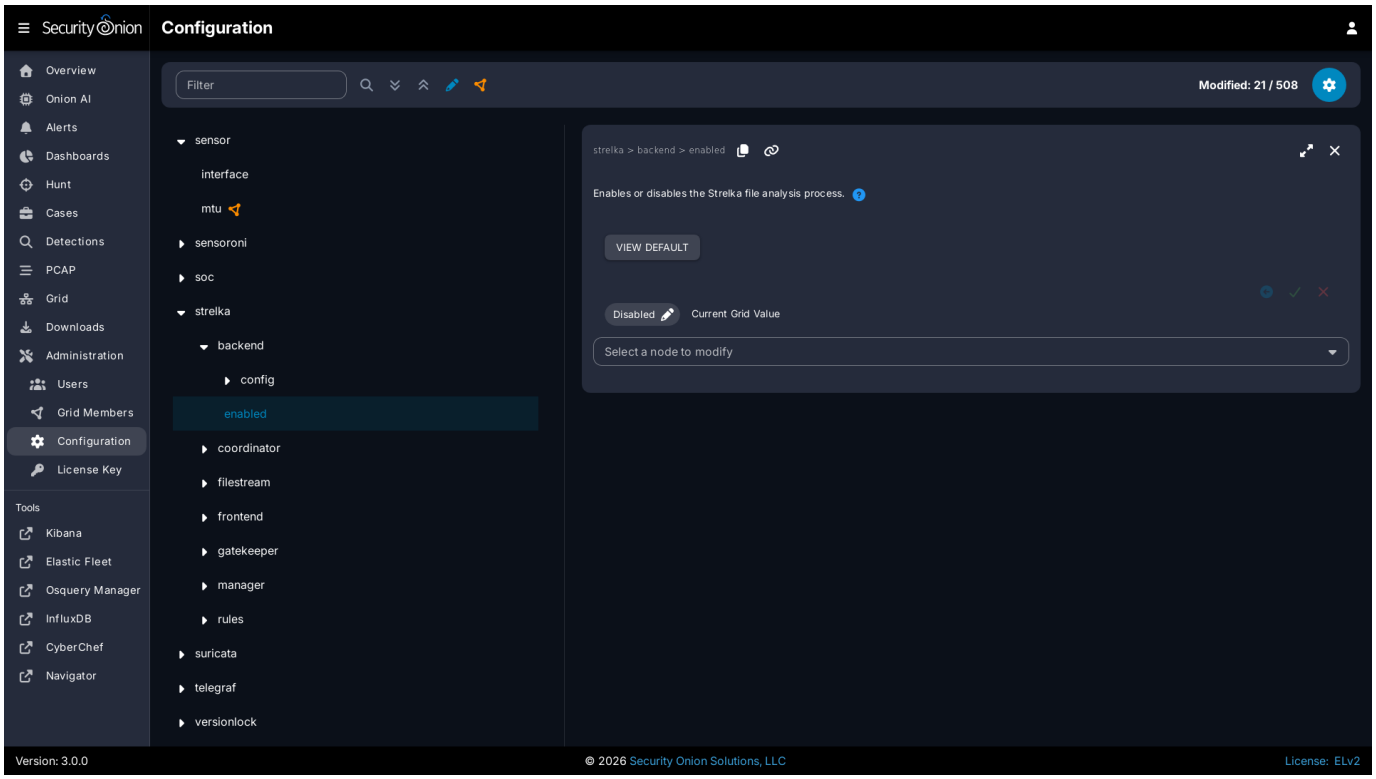

8.7.4 Diagnostic Logging

Strelka diagnostic logs are in `/nsm/strelka/log/`. Depending on what you're looking for, you may also need to look at the Docker logs for the containers:

```
sudo docker logs so-strelka-backend
sudo docker logs so-strelka-coordinator
sudo docker logs so-strelka-filestream
sudo docker logs so-strelka-frontend
sudo docker logs so-strelka-manager
```

8.7.5 More Information

Note

For more information about Strelka, please see https://github.com/target/strelka.

8.8 Intrusion Detection Honeypot

Security Onion includes an Intrusion Detection Honeypot (IDH) node option. This allows you to build a node that mimics common services such as HTTP, FTP, and SSH. Any interaction with these fake services will automatically result in an alert.

From the book, *Intrusion Detection Honeypots* (Sanders, C):

> An Intrusion Detection Honeypot (IDH) is a security resource placed inside your network perimeter that generates alerts when probed or attacked. These systems, services, and tokens rely on deception to lure attackers in and convince them to interact. Unbeknownst to the attacker, you're alerted when that interaction occurs and can begin investigating the compromise.

Chris Sanders and Josh Brower presented the IDH concept at Security Onion Conference 2021 and you can view the recording at https://www.youtube.com/watch?v=NzUhfARVfJk&list=PLljFlTO9rB17mESq7Z9OeFKvVh39vJW34&index=5.

8.8.1 Installation

IDH nodes are dedicated to just being IDH nodes and cannot run any other services. Therefore, you must have a separate manager to connect to. You can join a new IDH node to an existing Standalone deployment or full distributed deployment. Our ISO image includes a boot menu option for IDH installs that will partition your disk appropriately with lower requirements than a full installation.

> **Warning**
>
> The IDH node is designed to be placed *inside* your network perimeter! It should not be accessible from the Internet!

8.8.2 Configuration

- Run Setup and select the `DISTRIBUTED` deployment option.
- Select the `Existing Deployment` option.
- Select the `IDH` option.
- You can optionally prevent the IDH services from listening on the management interface.
- Once Setup is complete and the IDH node is fully joined to the grid, you can do additional configuration by going to Administration --> Configuration --> IDH.
- After configuration is complete, connections to honeypot services will result in `SO IDH` alerts that can be seen in Alerts.

8.8.3 Technical Background

The IDH node utilizes OpenCanary which is a modular opensource honeypot by Thinkst. You can read more about it at https://github.com/thinkst/opencanary.

OpenCanary logs can be found through Dashboards, Hunt, or Kibana using the following queries:

```
event.module: opencanary
event.dataset: IDH
```

Sigma rules within Detections look for certain logs emitted by OpenCanary to generate alerts, which can be viewed in the Alerts interface.

8.8.4 Services Configuration

The following services are available for use with the IDH node. Pay special attention to how an alert is triggered for a service as some of them require more than a simple connection request to trigger.

- FTP - a File Transfer Protocol server which alerts on login attempts
- Git - a Git server which alerts on repo cloning
- HTTP - an HTTP web server that alerts on login attempts
- HTTP Proxy - an HTTP web proxy that alerts when there is an attempt to proxy to another page
- MSSQL - an MS SQL server that alerts on login attempts
- MySQL - a MySQL server that alerts on login attempts
- Telnet - a Telnet server that alerts on login attempts
- SNMP - an SNMP server which alerts on OID requests
- SSH - a Secure Shell server which alerts on login attempts
- SIP - a SIP server which alerts on SIP requests
- VNC - a VNC server which alerts on login attempts
- Redis - a Redis server which alerts on actions
- TFTP - a TFTP server which alerts on requests
- NTP - an NTP server which alerts on NTP requests

This is based on the list at https://opencanary.readthedocs.io/en/latest/starting/configuration.html#services-configuration. RDP and SMB are not currently available for use within an IDH node.

In addition to changing the default ports, some of these services have further configuration options. For instance, the HTTP server has the ability to use custom HTML pages ("skins"). For more information, please see the OpenCanary documentation at https://opencanary.readthedocs.io/en/latest/starting/configuration.html#default-configuration.

These types of configuration changes can be made by modifying the minion pillar (see the Custom Configuration section below).

8.8.5 sshd

For IDH nodes, the local sshd is configured to listen on TCP/2222 and connections are only accepted from the Manager node. This allows TCP/22 to be used for honeypot services.

8.8.6 Custom Configuration

You can configure IDH by going to Administration --> Configuration --> IDH.

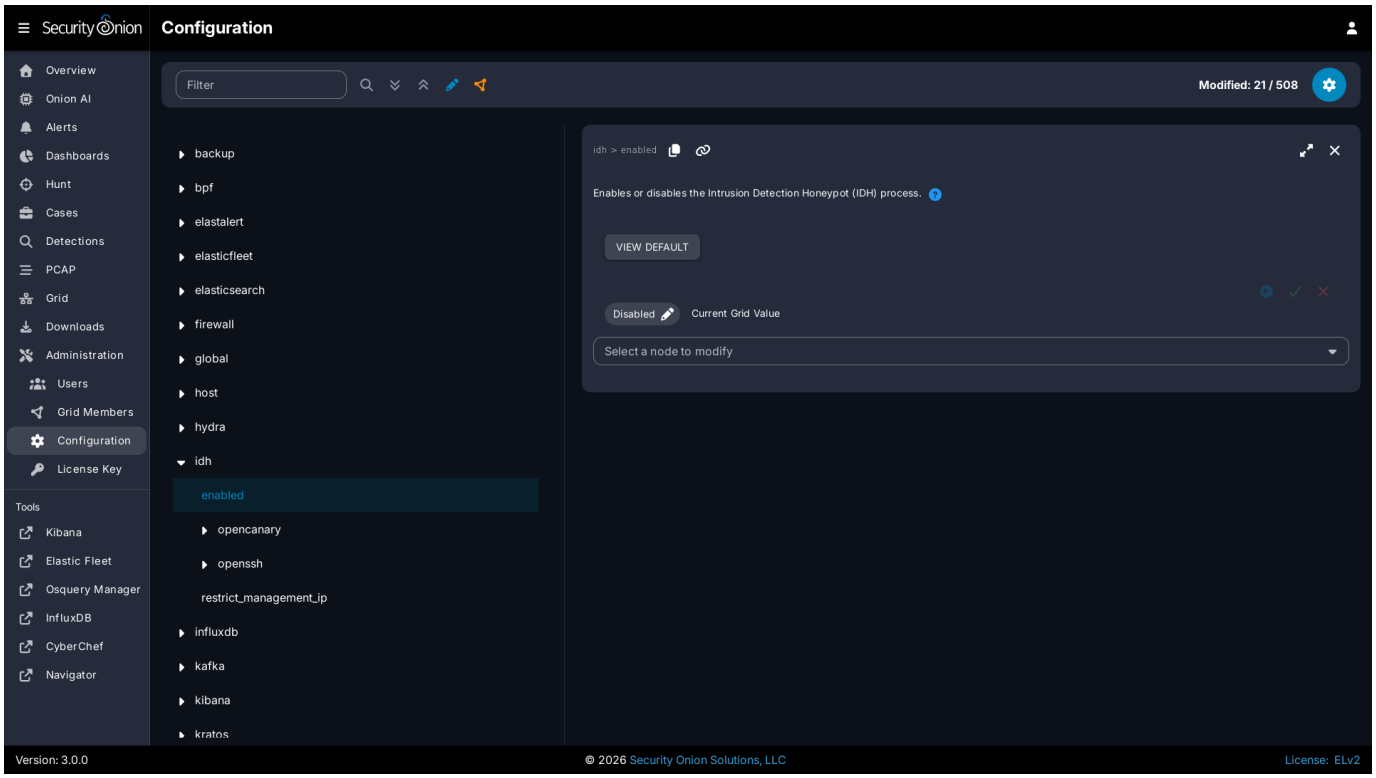

8.8.7 Custom Configuration Example

For example, suppose that we already have the HTTP service running but we want to change the default port from `80` to `8080`.

> **Warning**
>
> Please be very careful when making changes!

- Go to Administration --> Configuration.
- At the top of the page, click the `Options` menu and enable the `Show advanced settings` option.
- On the left side, navigate to IDH --> opencanary --> config --> http_x_port.
- On the right side, change the port value and then click the checkmark to save the change.
- At the top of the page, click the `SYNCHRONIZE GRID` button under the `Options` menu.

8.8.8 Activating Additional Network Interfaces

If you want to activate additional network interfaces after joining your IDH node to your grid, you can do so using standard Linux networking tools like nmtui. You can read more about nmtui at https://docs.oracle.com/en/operating-systems/oracle-linux/9/network/network-NetworkConfigurationTools.html.

9. Additional Network Visibility

9.1 Additional Network Visibility Overview

In the Network Visibility section, we looked at network visibility provided by Security Onion itself. The ideal situation would be to have Security Onion network sensors covering each and every one of your network segments. If you're able to achieve that ideal situation, then you may not need any additional network visibility. However, there may be times when you simply can't cover certain network segments with Security Onion network sensors and that's when these additional options can be beneficial. Keep in mind, though, that the data that they provide is nowhere near as comprehensive as a full Security Onion network sensor.

One option for additional network visibility would be NetFlow logs from firewalls, switches, or routers showing what traffic was observed by the network device. Another option would be firewall logs showing what traffic was allowed through the firewall and what traffic was denied. Firewall logs may be in different formats such as CEF, iptables, or a combination of the two (as in UniFi firewalls). We also have support for pfSense and OPNsense firewalls and you can find other firewall integrations in the Third Party Integrations section.

9.2 NetFlow

You may have devices on your network such as firewalls, routers, and switches that are capable of exporting NetFlow records. If you would like to collect these NetFlow records, add the Elastic integration for `NetFlow Records` and then allow the Netflow traffic through the firewall.

9.2.1 Add the NetFlow Records integration

First, add the Elastic integration for `NetFlow Records`.

> **Note**
>
> For more information about the `NetFlow Records` integration, please see https://docs.elastic.co/en/integrations/netflow.

1. Go to Elastic Fleet, click the `Agent policies` tab, and then click the desired policy (for example `so-Grid-nodes_general`).
2. Click the `Add integration` button.
3. Search for `netflow` and then click on the `NetFlow Records` integration.
4. The Elastic Integration page will show an overview of the NetFlow Integration. Review all information on the page and then click the `Add NetFlow Records` button.
5. On the `Add NetFlow Records integration` screen, go to the `UDP host to listen on` field and change `localhost` to `0.0.0.0`. Verify the `UDP port to listen on` field matches what your NetFlow exporter will be sending to. Click the `Save and continue` button and then click `Save and deploy changes`.

9.2.2 Allow NetFlow traffic through firewall

Next, allow the traffic from the NetFlow exporter through the firewall to the NetFlow listener port.

> **Note**
>
> The following instructions assume that this is the first firewall change you have made and therefore refer to `customhostgroup0` and `customportgroup0`. If those have already been used, you can select the next available hostgroup and portgroup.

1. Navigate to Administration --> Configuration.
2. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
3. On the left side, go to `firewall`, select `hostgroups`, and click the `customhostgroup0` group. On the right side, enter the IP address of the NetFlow exporter and click the checkmark to save.
4. On the left side, go to `firewall`, select `portgroups`, select the `customportgroup0` group, and then click `udp`. On the right side, enter your desired NetFlow listener port (2055 by default) and click the checkmark to save.
5. On the left side, go to `firewall`, select `role`, and then select the node type that will receive the NetFlow records. Then drill into `chain` --> `INPUT` --> `hostgroups` --> `customhostgroup0` --> `portgroups`. On the right side, enter `customportgroup0` and click the checkmark to save.
6. If you would like to apply the rules immediately, click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

9.2.3 NetFlow dashboard

Once all configuration is complete, you should be able to go to Dashboards and select the `NetFlow` dashboard to see your NetFlow records.

9.3 CEF

If you have devices on your network that can send logs in CEF format, you can send those logs to Security Onion. To get those CEF logs into Elasticsearch, you'll need to add the Elastic integration for CEF and then configure the Security Onion firewall to allow the remote device to send its logs.

9.3.1 Add the CEF integration

First, add the Elastic integration for `CEF`.

> **Note**
>
> For more information about the `CEF` integration, please see https://www.elastic.co/docs/reference/integrations/cef.

1. Go to Elastic Fleet, click the `Agent policies` tab, and then click the desired policy (for example `so-Grid-nodes_general`).
2. Click the `Add integration` button.
3. Search for `cef` and then click on the `Common Event Format (CEF)` integration.
4. The Elastic Integration page will show an overview of the CEF Integration. Review all information on the page and then click the `Add Common Event Format (CEF)` button.
5. On the `Add Common Event Format (CEF) integration` screen, disable the options labeled `Collect CEF application logs (input: logfile)` and `Collect CEF application logs (input: tcp)`. Make sure that `Collect CEF application logs (input: udp)` is enabled and then change the `Syslog Host` setting from `localhost` to `0.0.0.0`. The `Syslog Port` should be set to `9003` by default. Click the `Save and continue` button and then click `Save and deploy changes`.

9.3.2 Allow CEF logs through firewall

Next, allow the traffic from the CEF host through the firewall to the CEF integration port.

> **Note**
>
> The following instructions assume that this is the first firewall change you have made and therefore refer to `customhostgroup0` and `customportgroup0`. If those have already been used, you can select the next available hostgroup and portgroup.

1. Navigate to Administration --> Configuration.
2. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
3. On the left side, go to `firewall`, select `hostgroups`, and click the `customhostgroup0` group. On the right side, enter the IP address of the CEF host and click the checkmark to save.
4. On the left side, go to `firewall`, select `portgroups`, select the `customportgroup0` group, and then click `udp`. On the right side, enter your desired listener port (9003 by default) and click the checkmark to save.
5. On the left side, go to `firewall`, select `role`, and then select the node type that will receive the CEF logs. Then drill into `chain` --> `INPUT` --> `hostgroups` --> `customhostgroup0` --> `portgroups`. On the right side, enter `customportgroup0` and click the checkmark to save.
6. If you would like to apply the rules immediately, click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

9.3.3 CEF dashboard

Once all configuration is complete, you should be able to go to Dashboards and select the `CEF` dashboard to see your CEF logs.

9.4 iptables

If you have Linux iptables firewalls on your network, you can send those logs to Security Onion. To get those iptables logs into Elasticsearch, you'll need to add the Elastic integration for iptables and then configure the Security Onion firewall to allow the remote firewall to send its logs.

9.4.1 Add the iptables integration

First, add the Elastic integration for `iptables`.

> **Note**
>
> For more information about the `iptables` integration, please see https://www.elastic.co/docs/reference/integrations/iptables.

1. Go to Elastic Fleet, click the `Agent policies` tab, and then click the desired policy (for example `so-Grid-nodes_general`).
2. Click the `Add integration` button.
3. Search for `iptables` and then click on the `iptables` integration.
4. The Elastic Integration page will show an overview of the iptables Integration. Review all information on the page and then click the `Add iptables` button.
5. On the `Add iptables integration` screen, disable the options labeled `Collect iptables application logs (input: logfile)` and `Collect iptables application logs (input: journald)`. Make sure that `Collect iptables application logs (input: udp)` is enabled and then change the `Syslog host` setting from `localhost` to `0.0.0.0`. The `Syslog Port` should be set to `9001` by default. Click the `Save and continue` button and then click `Save and deploy changes`.

9.4.2 Allow iptables logs through firewall

Next, allow the traffic from the iptables host through the firewall to the iptables integration port.

> **Note**
>
> The following instructions assume that this is the first firewall change you have made and therefore refer to `customhostgroup0` and `customportgroup0`. If those have already been used, you can select the next available hostgroup and portgroup.

1. Navigate to Administration --> Configuration.
2. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
3. On the left side, go to `firewall`, select `hostgroups`, and click the `customhostgroup0` group. On the right side, enter the IP address of the iptables host and click the checkmark to save.
4. On the left side, go to `firewall`, select `portgroups`, select the `customportgroup0` group, and then click `udp`. On the right side, enter your desired listener port (9001 by default) and click the checkmark to save.
5. On the left side, go to `firewall`, select `role`, and then select the node type that will receive the iptables logs. Then drill into `chain` --> `INPUT` --> `hostgroups` --> `customhostgroup0` --> `portgroups`. On the right side, enter `customportgroup0` and click the checkmark to save.
6. If you would like to apply the rules immediately, click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

9.4.3 iptables dashboard

Once all configuration is complete, you should be able to go to Dashboards and select the `Firewall - iptables` dashboard to see your iptables logs.

9.5 UniFi

If you have UniFi firewalls on your network, you can send their logs to Security Onion. Typically, UniFi firewalls can send two different kinds of logs. The first is iptables firewall logs and the second is system logs in CEF format. To get all of these logs into Elasticsearch, you'll need to add the Elastic integrations for iptables and CEF, configure the UniFi device to send those logs, and then configure the Security Onion firewall to allow those logs.

9.5.1 Add the iptables and CEF integrations

First, add the Elastic integrations for iptables and CEF.

> **Note**
>
> For more information about the iptables integration, see https://www.elastic.co/docs/reference/integrations/iptables.
>
> For more information about the CEF integration, see https://www.elastic.co/docs/reference/integrations/cef.

1. Go to Elastic Fleet, click the `Agent policies` tab, and then click the desired policy (for example `so-Grid-nodes_general`).
2. Click the `Add integration` button.
3. Search for `iptables` and then click on the `iptables` integration.
4. The Elastic Integration page will show an overview of the iptables Integration. Review all information on the page and then click the `Add iptables` button.
5. On the `Add iptables integration` screen, disable the options labeled `Collect iptables application logs (input: logfile)` and `Collect iptables application logs (input: journald)`. Make sure that `Collect iptables application logs (input: udp)` is enabled and then change the `Syslog host` setting from `localhost` to `0.0.0.0`. The `Syslog Port` should be set to `9001` by default. Click the `Save and continue` button and then click `Save and deploy changes`.
6. Back at the desired policy screen, click the `Add integration` button.
7. Search for `cef` and then click on the `Common Event Format (CEF)` integration.
8. The Elastic Integration page will show an overview of the CEF Integration. Review all information on the page and then click the `Add Common Event Format (CEF)` button.
9. On the `Add Common Event Format (CEF) integration` screen, disable the options labeled `Collect CEF application logs (input: logfile)` and `Collect CEF application logs (input: tcp)`. Make sure that `Collect CEF application logs (input: udp)` is enabled and then change the `Syslog Host` setting from `localhost` to `0.0.0.0`. The `Syslog Port` should be set to `9003` by default. Click the `Save and continue` button and then click `Save and deploy changes`.

9.5.2 Configure UniFi

Next, configure UniFi to send both types of logs to Security Onion.

> **Note**
>
> UniFi configuration may be different depending on what specific UniFi device you have and what software it is running. These instructions are based on a Cloud Gateway Fiber device running control plane version 4.2.12 and Network application version 9.3.43.

To configure UniFi to send iptables firewall logs to the Elastic integration for iptables:

1. In the UniFi web interface, navigate to Settings --> CyberSecure --> Traffic Logging.
2. Next to `Activity Logging (Syslog)`, choose the `SIEM Server` option.
3. Set the `Server Address` to the IP address of the Security Onion node to send the logs to.
4. Set the `Port` to 9001.
5. Click the `Apply Changes` button.

To configure UniFi to send system logs to the Elastic integration for CEF:

1. In the UniFi web interface, navigate to Settings --> Control Plane --> Integrations.
2. Next to `Activity Logging (Syslog)`, choose the `SIEM Server` option.
3. Set the `Server Address` to the IP address of the Security Onion node to send the logs to.
4. Set the `Port` to 9003.
5. Click the `Apply Changes` button.

While in UniFI, check your UniFi firewall rules and update if necessary:

1. In the UniFi web interface, navigate to Settings --> Policy Engine.
2. For any firewall rule that you want to see in Security Onion, make sure that `Syslog Logging` is enabled and the description starts with either `Block` or `Allow`.

9.5.3 Allow UniFi logs through Security Onion firewall

Finally, allow the traffic from the UniFi device through the Security Onion firewall to the Elastic integration ports.

> **Note**
>
> The following instructions assume that this is the first firewall change you have made and therefore refer to `customhostgroup0` and `customportgroup0`. If those have already been used, you can select the next available hostgroup and portgroup.

1. Navigate to Administration --> Configuration.
2. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
3. On the left side, go to `firewall`, select `hostgroups`, and click the `customhostgroup0` group. On the right side, enter the IP address of the UniFi host and click the checkmark to save.
4. On the left side, go to `firewall`, select `portgroups`, select the `customportgroup0` group, and then click `udp`. On the right side, enter `9001` and `9003` and then click the checkmark to save.
5. On the left side, go to `firewall`, select `role`, and then select the node type that will receive the UniFi logs. Then drill into `chain` --> `INPUT` --> `hostgroups` --> `customhostgroup0` --> `portgroups`. On the right side, enter `customportgroup0` and click the checkmark to save.
6. If you would like to apply the rules immediately, click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

9.5.4 UniFi dashboards

Once all configuration is complete, you should be able to go to Dashboards and select one of the `Firewall - UniFi` dashboards to see your UniFi logs.

9.6 pfSense

pfSense is a free and open firewall that can be found at https://www.pfsense.org/. There are a few different ways in which you can integrate pfSense into your Security Onion deployment.

The simplest method of integrating pfSense into your Security Onion deployment is to configure pfSense to send its firewall logs to Security Onion. Security Onion has a couple of options for ingesting logs from pfSense firewalls: a simple parser and the more comprehensive Elastic Integration for pfSense. We recommend using the more comprehensive option by following the steps in the Elastic Integration section below. You can also follow along with our video at https://www.youtube.com/watch?v=aoH8qZwAxek.

You can also configure pfSense to send NetFlow data to your Security Onion deployment. This may provide some additional visibility if you can't deploy a full Security Onion network sensor to monitor that network segment.

9.6.1 Simple Parser

The first option for collecting pfSense firewall logs is to use our simple parser. Please note that this only supports `filterlog` (actual firewall logs) and no other logs.

> **Warning**
>
> This simple parser will be phased out over time in favor of the more comprehensive Elastic Integration for pfSense below.

To use the simple parser, first go to Administration --> Configuration --> firewall --> hostgroups.

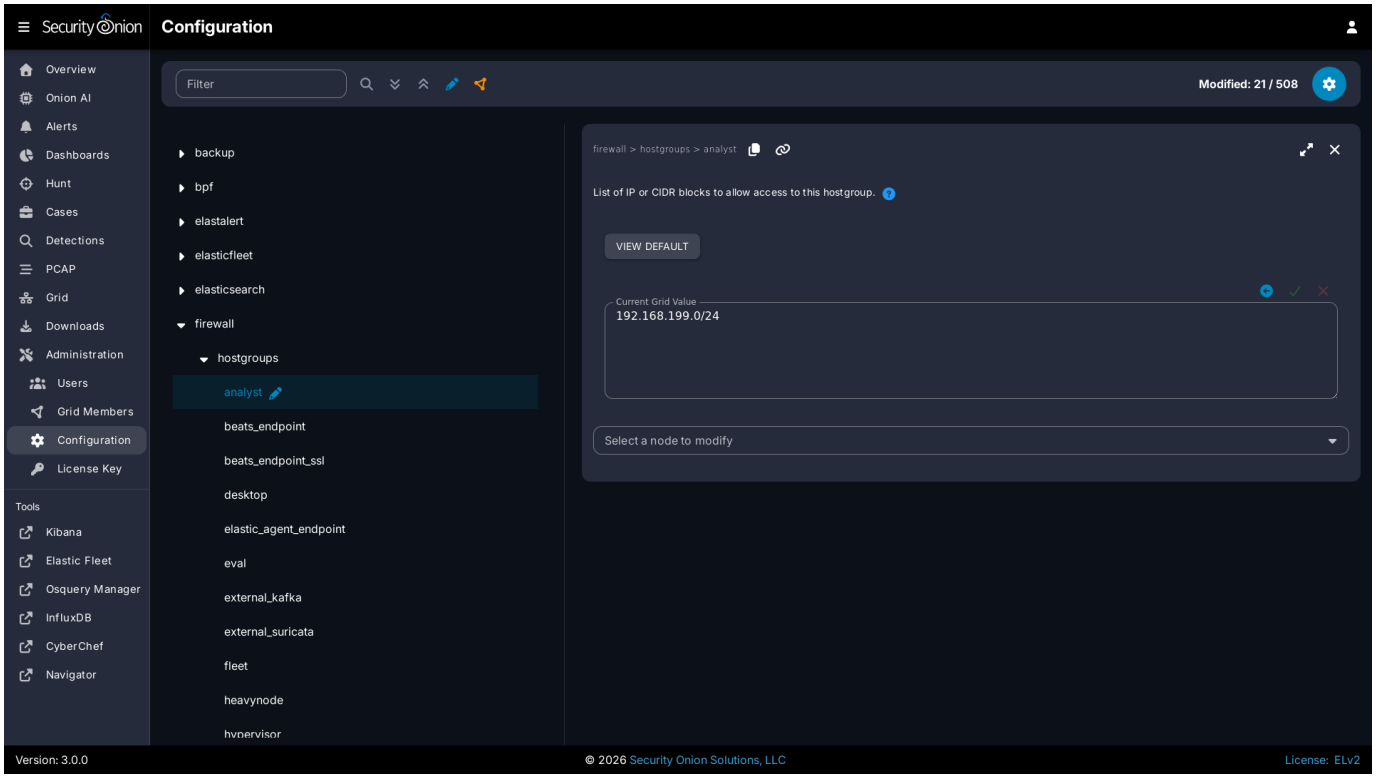

Once there, select the `syslog` option, specify the IP address of the pfSense firewall, and click the checkmark to save. Then click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

Next, configure your pfSense firewall to send `syslog` to the IP address of your Security Onion box. If you are using pfSense 2.6.0 or higher, make sure that `Log Message Format` is set to `BSD (RFC 3164, default)`.

Once all configuration is complete, you should be able to go to Dashboards and select the Firewall dashboard to see your firewall logs.

9.6.2 Elastic Integration for pfSense

The second option for collecting pfSense firewall logs is using the Elastic Integration for pfSense (https://docs.elastic.co/integrations/pfsense). This integration is more comprehensive than the simple parser above and supports more log types.

First, add the pfSense integration and configure the pfSense firewall:

1. Go to Elastic Fleet, click the `Agent policies` tab, and then click the desired policy (for example `so-Grid-nodes_general`).
2. Click the `Add integration` button.
3. Search for `pfSense` and then click on the `pfSense` integration.
4. The Elastic Integration page will show instructions for configuring pfSense. Follow these instructions but please note that the Elastic Integration expects to receive pfSense logs on port 9001 by default.
5. Once you've configured pfSense, then go back to the Elastic screen and click the `Add pfSense` button.
6. On the `Edit pfSense integration` screen, go to the `Syslog Host` field and change `localhost` to `0.0.0.0`.
7. Click the `Save and continue` button and then click `Save and deploy changes`.

Next, allow the traffic from the pfSense firewall to port 9001. These instructions assume that this is the first firewall change you have made and therefore refer to `customhostgroup0` and `customportgroup0`. If those have already been used, select the next available hostgroup and portgroup.

1. Navigate to Administration --> Configuration.
2. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
3. On the left side, go to `firewall`, select `hostgroups`, and click the `customhostgroup0` group. On the right side, enter the IP address of the pfSense firewall and click the checkmark to save.
4. On the left side, go to `firewall`, select `portgroups`, select the `customportgroup0` group, and then click `udp`. On the right side, enter `9001` and click the checkmark to save.
5. On the left side, go to `firewall`, select `role`, and then select the node type that will receive the pfSense logs. Then drill into `chain` --> `INPUT` --> `hostgroups` --> `customhostgroup0` --> `portgroups`. On the right side, enter `customportgroup0` and click the checkmark to save.
6. If you would like to apply the rules immediately, click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

Once all configuration is complete, you should be able to go to Dashboards and select the `Firewall - pfSense` dashboard to see your firewall logs.

9.7 OPNsense

OPNsense is a free and open firewall that can be found at https://opnsense.org/. Similar to pfSense, there are a few different ways in which you can integrate OPNsense into your Security Onion deployment. First, you can configure OPNsense to send firewall logs to your Security Onion deployment. You can also configure OPNsense to send NetFlow data to your Security Onion deployment. Finally, a third option is to integrate OPNsense's Suricata into Security Onion so that it pulls NIDS rules from your Security Onion deployment and then sends those Suricata alerts to Security Onion. All three options are detailed below.

9.7.1 Sending OPNsense Firewall Logs to Security Onion

Centralize your logging by sending OPNsense firewall logs to your grid.

Steps:

1. Navigate to **System → Settings → Logging** in the OPNsense GUI.
2. Click on the **Remote Logging** tab.
3. Click the **+** button to add a new remote logging destination.
4. **Configure Remote Logging:**
5. **Transport**: Select **TCP**.
6. **Application Levels**: Leave at default to send all logs or specify as needed.
7. **Facilities**: Leave at default to include all facilities or specify as needed.
8. **Hostname/IP Address**: Enter the IP address of the grid system where you set up the logging input.
9. **Port**: Enter the port number configured on the grid system.
10. **Format**: Choose the appropriate format (e.g., Syslog).
11. Click **Save** to apply the settings.

9.7.2 Sending OPNsense NetFlow data to Security Onion

To collect network flow data (similar to Zeek connection logs), configure NetFlow on OPNsense to send data to your grid.

Steps:

1. **Prepare Your Grid to Receive NetFlow Data:**
2. Refer to the NetFlow section to set up your grid for receiving NetFlow data.
3. **Configure NetFlow on OPNsense:**
4. Navigate to **Reporting → NetFlow** in the OPNsense GUI.
5. Under **Capture**, select the internal interfaces you wish to monitor.
6. Also, select your **WAN** interface to monitor external traffic.
7. Under **Destinations**, add a new destination:
    - **Hostname/IP Address**: Enter the IP address of the grid node configured to accept NetFlow data.
    - **Port**: Enter the port number you set up on the grid node.
    - **Format**: Choose the appropriate NetFlow version (e.g., NetFlow v5 or v9).
8. Click **Apply** to save the settings.

9.7.3 Integrating OPNsense Suricata into Security Onion

You can integrate your OPNsense firewall with your Security Onion Grid system to utilize it as a makeshift sensor, allowing management of Suricata NIDS rules through Detections. This is particularly useful for organizations that find it challenging to install a network TAP on their egress connection.

> **Note**
>
> When OPNsense downloads rules from Security Onion, it can only enable and disable rules. Threshold settings are ignored.

Prerequisites

- **Security Onion Manager Access**: Access to the manager of your grid.
- **Administrative Rights**: Ability to modify settings in the OPNsense GUI.
- **SSH Access**: You must have SSH access to your OPNsense firewall.

Share NIDS Rules

The Detections module can now manage NIDS rules for external Suricata instances. Please refer to the NIDS section for detailed instructions on how to set this up.

Import Security Onion CA

To establish a secure connection between your OPNsense firewall and the grid manager, your firewall will need to trust the grid's Certificate Authority (CA) certificate.

Steps:

1. **Copy the grid CA Certificate:**
2. SSH into your grid manager.
3. Run the command `cat /etc/pki/ca.crt` to display the CA certificate.
4. Copy the entire output of that command.
5. **Import the CA Certificate into OPNsense:**
6. Log into the OPNsense GUI.
7. Navigate to **System** → **Trust** → **Authorities**.
8. Click the **+** button to add a new certificate authority.
9. Set **Method** to **Import an existing Certificate Authority**.
10. **Descriptive Name**: Enter a name like "Security Onion CA".
11. **Certificate Data**: Paste the copied CA certificate content.
12. Click **Save**.

Setting Up the Suricata Rules Repository

Since OPNsense doesn't allow enabling third-party repositories through the GUI, you'll need to modify the configuration manually.

Steps:

1. **Remove Existing Rule Repositories:**
2. SSH into your OPNsense firewall and run the following:

```
rm -rf /usr/local/opnsense/scripts/suricata/metadata/rules/*
```

1. **Create a New Repository File:**
2. Create and edit the `onion.xml` file:

```
vi /usr/local/opnsense/scripts/suricata/metadata/rules/onion.xml
```

- Paste the following content into the file, replacing `YOURMANAGER` with the hostname or IP address of your grid manager:

```
<?xml version="1.0"?>
<ruleset documentation_url="http://docs.opnsense.org/">
    <location url="https://YOURMANAGER:7789/" prefix="SecurityOnion"/>
    <files>
        <file description="SecurityOnion rules">all.rules</file>
        <file description="SecurityOnion" url="inline::all.rules">all.rules</file>
    </files>
</ruleset>
```

- Save and exit the editor.
- **Refresh Rule Sets in OPNsense:**
- Navigate to **Services** → **Intrusion Detection** → **Administration** → **Download** in the OPNsense GUI.
- You should see **Security Onion** listed as a ruleset.
- Select **Security Onion** and click **Download & Update Rules**.
- Once updated, the rules will appear under the **Rules** tab.

Scheduling Rule Updates

To keep your Suricata rules up to date, schedule regular updates.

Steps:

1. Navigate to **Services** → **Intrusion Detection** → **Administration** → **Schedule**.
2. Click the **+** button to add a new schedule.
3. **Configure the Schedule:**
4. **Description**: Enter a name like "Suricata Rule Update".
5. **Cron Expression**: Set the frequency to every 15 minutes.
6. **Type**: Choose **Update and reload intrusion detection rules**.
7. Click **Save**.

OPNsense will now automatically download and reload the rules every 15 minutes.

10. Host Visibility

10.1 Host Visibility Overview

More and more of our network traffic is encrypted these days and that's a good thing for privacy but it's somewhat of a blind spot for us as defenders. Host visibility can help fill in those blind spots. You can send host logs to Security Onion via your choice of either Elastic Agent or Syslog:

- Choose Elastic Agent for comprehensive telemetry if you can install an agent on the host.
- Choose Syslog if you can't install an agent but the device supports sending standard syslog. Examples include firewalls, switches, routers, and other network devices.

For Windows endpoints, you can optionally augment the standard Windows logging with Sysmon.

10.2 Elastic Agent

From https://www.elastic.co/elastic-agent:

> With Elastic Agent you can collect all forms of data from anywhere with a single unified agent per host. One thing to install, configure, and scale. Each Security Onion node uses the Elastic Agent to transport logs to Elasticsearch. You can also deploy the Elastic Agent to your endpoints to transport logs to your Security Onion deployment.

Note

In order to receive logs from the Elastic Agent, Security Onion must be running Logstash. Evaluation Mode and Import Mode do not run Logstash, so you'll need Standalone or a full Distributed Deployment. In a Distributed Deployment, sensor nodes do not run Logstash, so you'll need to configure agents to send to your manager or receiver nodes. For more information, please see the Architecture section.

To deploy an Elastic agent to an endpoint, go to the SOC Downloads page and download the proper Elastic agent for the operating system of that endpoint.

Warning

Within the Elastic Fleet interface, there is an `Add Agent` button - it is not recommended to use this particular method to install the Elastic Agent, as it requires much more manual configuration.

Don't forget to allow the agent to connect through the firewall by going to Administration --> Configuration --> firewall --> hostgroups.

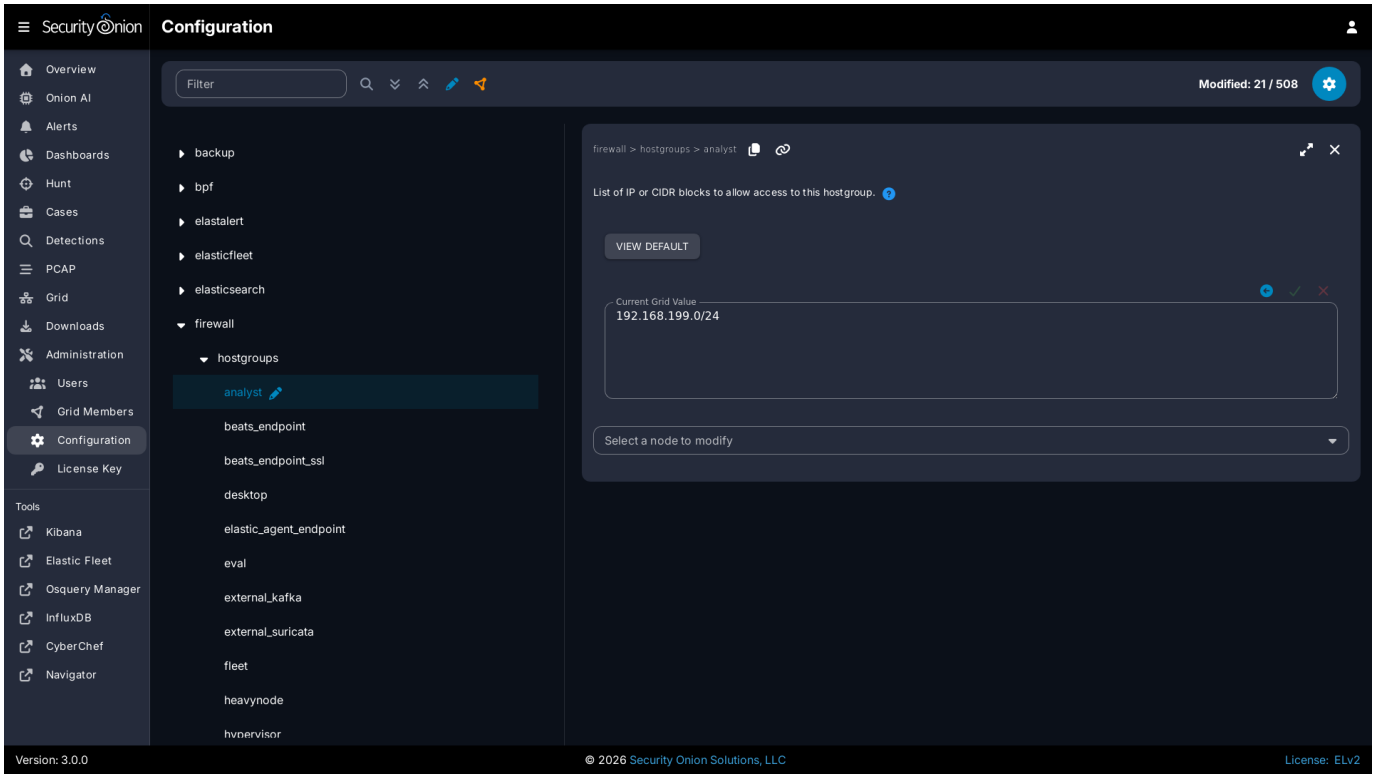

Once there, select the `elastic_agent_endpoint` option.

> **Note**
>
> Check out our Elastic Agent video at https://youtu.be/cGmQMsFuAvw!

10.2.1 Elastic Agent Options (Non-MSI)

There are additional installer runtime options:

`-token=$TOKEN`

This option allows you to override the agent policy that the installer uses by default. The token comes from the Elastic Fleet interface under `Enrollment Tokens`.

`-fleet=$FLEETHOST`

This option allows you to override the default Fleet host used for enrollment.

`-delay-enroll=true|false`

Defaults to false. If set to true, it adds the builtin `delay-enroll` flag when enrolling the agent.

`-timeout=$MINUTES`

Defaults to 5 minutes.

10.2.2 Elastic Agent Options (MSI)

When using the MSI installer runtime options are:

`TOKEN=$TOKEN`

This option allows you to override the agent policy that the installer uses by default. The token comes from the Elastic Fleet interface under `Enrollment Tokens`.

`FLEET=$FLEETHOST`

This option allows you to override the default Fleet host used for enrollment.

`DELAYENROLL=true|false`

This option defaults to false. If set to true, it adds the builtin `delay-enroll` flag when enrolling the agent.

`TIMEOUT=$MINUTES`

Defaults to 5 minutes.

10.2.3 Installing Elastic Agent on Linux

If deploying the Elastic Agent to a Linux host, make the file executable and then execute using sudo:

```
chmod +x ./so-elastic-agent_linux_amd64
sudo ./so-elastic-agent_linux_amd64
```

10.2.4 Installing Elastic Agent on macOS

If deploying the Elastic Agent to macOS, you will need to take a few steps. First, remove the quarantine attribute. Then, make the file executable. Finally, execute the file using sudo:

```
xattr -d com.apple.quarantine ./so-elastic-agent_darwin_amd64
chmod +x ./so-elastic-agent_darwin_amd64
sudo ./so-elastic-agent_darwin_amd64
```

After the installer has completed, review the Elastic docs for your version of macOS and approve the required settings (system extension and full drive access) as shown at https://www.elastic.co/guide/en/security/current/elastic-endpoint-deploy-reqs.html.

10.2.5 Installing Elastic Agent on Windows via MSI

When deploying via MSI, installation can be as simple as double-clicking the MSI installer. If you require additional runtime flags, use `msiexec /i` :

```
msiexec /i so-elastic-agent_windows_amd64_msi TOKEN=$TOKEN TIMEOUT=10m
```

10.2.6 Logs

Once the agent starts sending logs, you should be able to find them in Dashboards, Hunt, or Kibana.

10.2.7 Management

You can manage your agents using Elastic Fleet.

10.2.8 Live Queries

You can query your agents in realtime using Osquery Manager.

10.2.9 Integrations

You can read more about integrations in the Elastic Fleet section and at https://docs.elastic.co/integrations.

10.2.10 Reinstalling

If for some reason you need to uninstall and reinstall the Elastic Agent on one of your Security Onion Grid members, you can do so as follows:

```
sudo elastic-agent uninstall
sudo salt-call state.apply elasticfleet.install_agent_grid
```

10.2.11 More Information

Note

For more information about the Elastic Agent, please see https://www.elastic.co/guide/en/fleet/current/fleet-overview.html.

10.3 Syslog

If you want to send syslog from other devices, you should check to see if the device has an existing Elastic Agent integration. If so, using the Elastic Agent integration should provide some parsing by default.

If your device does not have an existing Elastic Agent integration, you can still collect standard syslog. Start by going to Administration --> Configuration --> firewall --> hostgroups.

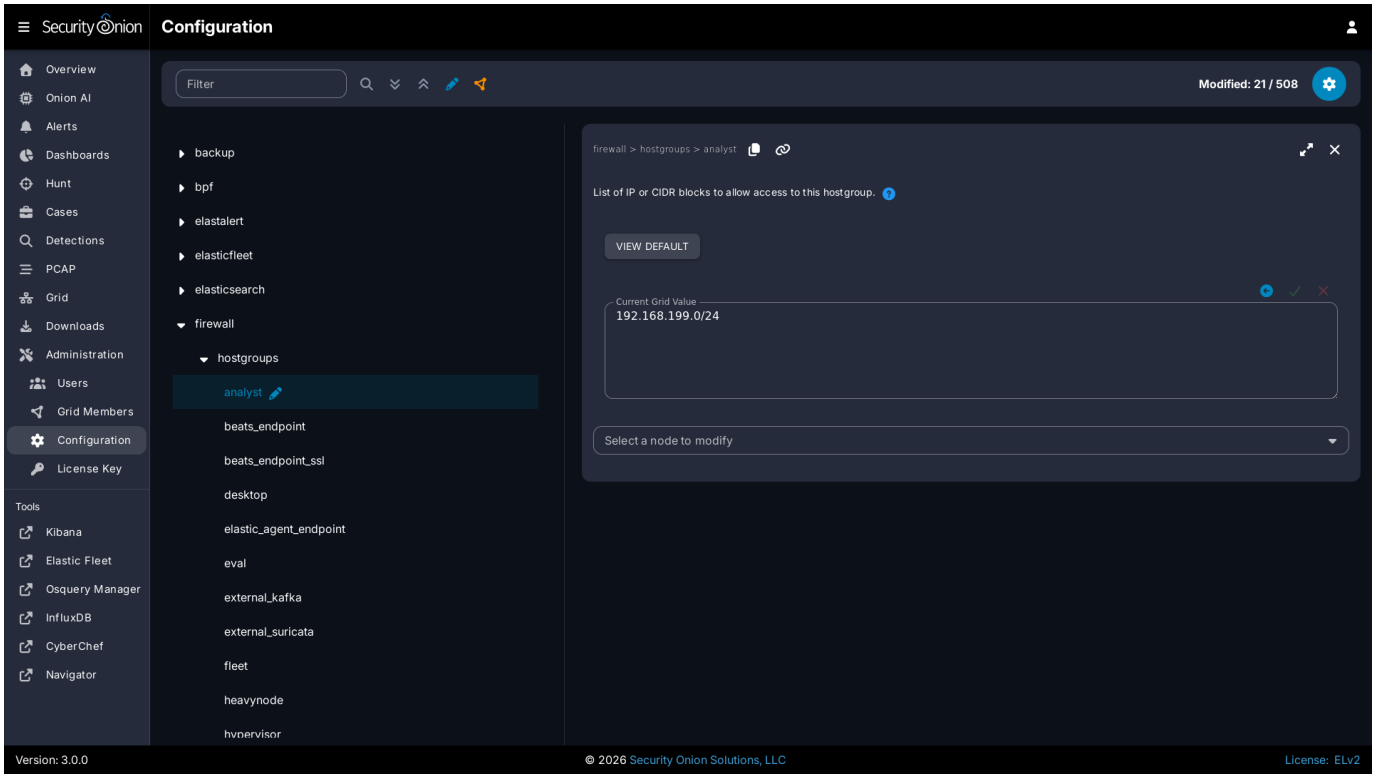

Then choose the `syslog` option to allow the port through the firewall. If sending syslog to a sensor, please see the Examples in the Firewall section. If you need to add custom parsing for those syslog logs, we recommend using Elasticsearch ingest parsing.

Also note that if you're monitoring network traffic with Zeek, then by default it will detect any syslog in that network traffic and log it even if that syslog was not destined for that particular Security Onion node.

10.4 Sysmon

From https://learn.microsoft.com/en-us/sysinternals/downloads/sysmon:

> System Monitor (Sysmon) is a Windows system service and device driver that, once installed on a system, remains resident across system reboots to monitor and log system activity to the Windows event log. It provides detailed information about process creations, network connections, and changes to file creation time. By collecting the events it generates using Windows Event Collection or SIEM agents and subsequently analyzing them, you can identify malicious or anomalous activity and understand how intruders and malware operate on your network.

10.4.1 Downloads

You can download Sysmon from Microsoft at https://download.sysinternals.com/files/Sysmon.zip.

Once you've downloaded Sysmon, you probably also want to download a Sysmon configuration to use as a starting point. Here are a few options to choose from:

- https://github.com/Neo23x0/sysmon-config
- https://github.com/SwiftOnSecurity/sysmon-config
- https://github.com/olafhartong/sysmon-modular

10.4.2 Transport

Sysmon logs can be collected and transported using Elastic Agent. Confirm that your configuration does NOT use the Elastic Sysmon module. Security Onion will do all the necessary parsing.

10.4.3 Visualizations

Once Security Onion is receiving and parsing Sysmon data, you can search for that data and visualize it via Dashboards, Hunt, or Kibana. Each of these interfaces have at least one dashboard or query specifically designed for Sysmon data.

10.4.4 More Information

Note

For more information about Sysmon, please see: https://learn.microsoft.com/en-us/sysinternals/downloads/sysmon

TrustedSec has a great Community Guide on Sysmon: https://github.com/trustedsec/SysmonCommunityGuide

11. Third Party Integrations

In addition to network visibility and host visibility, you may want to pull in data from other third party systems. You can do that via Elastic integrations which support many of the most common products and services. You can read more about Elastic integrations at https://docs.elastic.co/integrations.

> **Warning**
>
> Third party integrations are provided by Elastic and are not specifically tested by the Security Onion team. Support provided by the Security Onion team for third party integrations is considered best-effort.

11.1 Adding an Integration

New integrations can be added to existing policies to provide increased visibility and more comprehensive monitoring.

> **Tip**
>
> When adding a new integration, it is important that you add it to an appropriate policy.
>
> If an integration pulls the data, you should add it to the Fleet Server policy. Depending on complexity and log volume, it might make sense to stand up a Fleet Node and add your integrations to it.
>
> If an integration receives data pushed to it (for example: receiving syslog), consider adding it to the Fleet Server policy. If that is not feasible, then you can add it to the grid Nodes policy but make sure to set the firewall rules correctly so that you are not opening ports on all of your nodes.

To add an integration to an existing policy:

- From the main Fleet page, click the `Agent policies` tab.
- Select the desired agent policy.
- Click the `Add Integration` button.
- Follow the steps for adding the integration.

> **Note**
>
> If the integration is designed to listen on a port to receive data, it will most likely default to listening on `localhost` only. Depending on how you are sending data to the integration, you may need to change that to `0.0.0.0` so that it can receive data from other hosts.

For examples of this process, please see the NetFlow and pfSense sections. The pfSense section includes a link to a video which illustrates the process.

11.2 Adding a Custom Integration

A custom integration can be added by adding an integration such as the `Custom Logs` integration. You can specify various settings relative to the data source and define additional actions to be performed.

11.3 Managing Integration Upgrades

> **Tip**
>
> By default, integrations are not automatically kept up to date. This avoids potential log ingest downtime if there is an issue with the latest package or if the latest package requires a manual update to your integration configuration. If you would like to automatically upgrade integrations, you can change this behavior via Administration -> Configuration -> elasticfleet -> config -> auto_upgrade_integrations.

To find integrations that have upgrades available:

- Navigate to Elastic Fleet.
- At the top left corner, click the menu.
- Under `Management`, select `Integrations`.
- Click the `Installed Integrations` tab.
- Review any integrations listed under `Updates available`.

11.4 Managing Third Party Integration Index Templates

Index templates for third party integrations can be managed as described in the Elasticsearch section, but first `managed_integrations` must be updated by navigating to Administration --> Configuration --> manager --> managed_integrations.

11.5 Supported Integrations

The current release of Security Onion supports all standard Elastic integrations as shown at https://docs.elastic.co/integrations.

11.6 More Information

> **Note**
>
> You can read more about Elastic integrations at https://docs.elastic.co/integrations.

12. Rules

12.1 Rules Overview

Security Onion supports three main types of rules: NIDS, Sigma, and YARA. You can manage all three types via Detections.

12.2 NIDS

NIDS (Network Intrusion Detection System) rules are loaded into Suricata to monitor network traffic for suspicious or noteworthy activity. Active NIDS rules generate alerts that can be found in Alerts.

12.2.1 Managing Existing NIDS Rules

You can manage existing NIDS rules using Detections. There are two ways to do so:

- From the main Detections interface, you can search for the desired detection and click the binoculars icon.
- From the Alerts interface, you can click an alert and then click the `Tune Detection` menu item.

Once you've used one of these methods to reach the detection detail page, you can check the Status field in the upper-right corner and use the slider to enable or disable the detection.

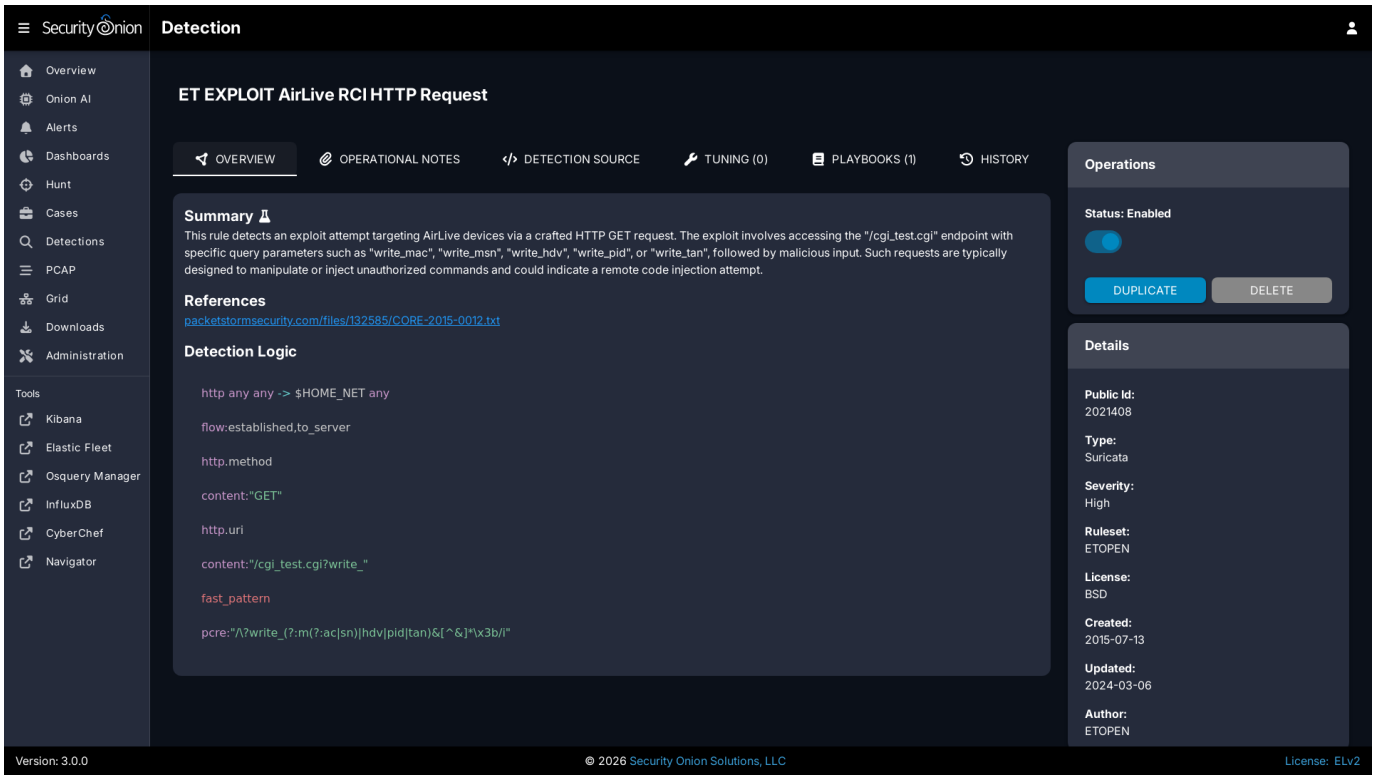

To tune the detection:

- Click the Tuning tab
- Click the blue + button
- Select the type of tuning (Modify, Suppress, or Threshold)
- Fill out the requested values
- Click the `CREATE` button

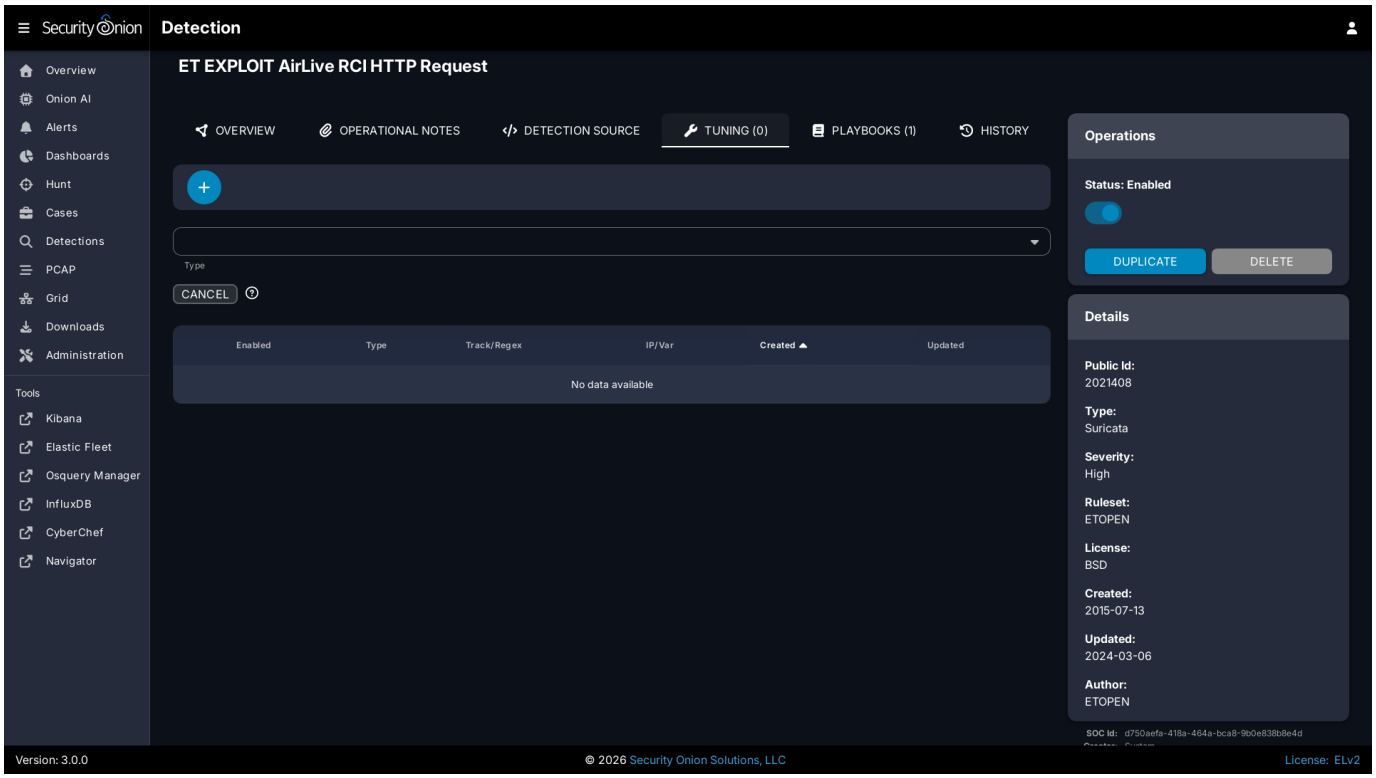

12.2.2 Enabling and Disabling with Regex

NIDS rules can be enabled or disabled in Detections using regex patterns. Navigate to SOC Administration - Configuration and filter for `regex`, then drill down into `SOC` --> `config` --> `server` --> `modules` --> `suricataengine` --> `disableRegex` or `enableRegex`.

The regex flavor is Google RE2: https://github.com/google/re2/wiki/Syntax

In ETOPEN, categories are prepended to the rule name. For example, the `ET EXPLOIT PHP-Live-Chat Get Shell Attempt Inbound` rule is in the `ET EXPLOIT` category. So suppose you want to disable the `ET EXPLOIT` and `ET MALWARE` categories but NOT the `ET EXPLOIT_KIT` category. You would use the following regex patterns:

```
ET EXPLOIT\s
ET MALWARE\s
```

The `\s` is a shortcut for whitespace and is useful in this situation to make sure we are only matching the specific categories that we want to disable.

If a detection would be matched by both an enable and disable regex, it is enabled. If a detection's status is changed via the Detections interface but it is currently matched by a regex pattern, the change initiated from the Detections interface is reverted and a message is shown.

Enable and disable operations that are based on regex patterns are actioned during the daily rule update. If you have made a change to the regex patterns and would like to have it implemented more immediately:

- Under Grid Configuration, click the `SYNCHRONIZE GRID` button and wait about 5 minutes for it to complete.
- Navigate to Detections, click the Options menu, select Suricata in the dropdown menu, click the `FULL UPDATE` button, and then wait for it to complete.
- Refresh the Detections page and you should see the relevant rule statuses have changed.

Note

If a disable regex is applied to a setter flowbit rule and that rule is still required, it will be written out to the rules file as enabled, but `noalert`

12.2.3 Tuning Overrides

Overrides allow you to tune rule behavior without modifying the rule itself. Security Onion supports three types of overrides for NIDS rules.

Threshold Override

The threshold override limits how often a rule generates alerts. Use this when a rule is generating too many alerts and you want to reduce volume without disabling the detection.

Thresholds are written to Suricata's `threshold.conf` file and control alert frequency based on occurrence count and time window.

- **Threshold Type**: `limit`, `threshold`, or `both`
- `limit`: Alert at most N times per time window
- `threshold`: Alert only after N occurrences per time window
- `both`: Alert once per time window after N occurrences
- **Track**: `by_src` or `by_dst` - whether to track per source IP or destination IP
- **Count**: Number of occurrences for the threshold logic (must be > 0)
- **Seconds**: Time window in seconds (must be > 0)

Example: Limit alerts to once per hour per source IP:

Threshold Type: `limit` **Track:** `by_src` **Count:** `1` **Seconds:** `3600`

Result in threshold.conf:

```
threshold gen_id 1, sig_id 2001219, type limit, track by_src, count 1, seconds 3600
```

Suppress Override

The suppress override silences alerts from specific IP addresses or networks. Use this when a rule generates false positives from known-good sources.

Suppressions are written to Suricata's `threshold.conf` file. When traffic matches the rule and originates from (or is destined to) the specified IP, no alert is generated.

- **Track**: `by_src` or `by_dst` - suppress based on source IP or destination IP
- **IP**: IP address, CIDR network, or Suricata variable (e.g., `192.168.1.100`, `10.0.0.0/8`, or `$SCANNERS`)

Example: Suppress alerts from a vulnerability scanner:

Track: `by_src` **IP:** `192.168.50.10`

Result in threshold.conf:

```
suppress gen_id 1, sig_id 2001219, track by_src, ip 192.168.50.10
```

Example: Suppress using a Suricata variable:

Track: `by_src` **IP:** `$SCANNERS`

Result in threshold.conf:

```
suppress gen_id 1, sig_id 2001219, track by_src, ip $SCANNERS
```

Modify Override

The modify override allows you to change the content of a Suricata rule using regular expression pattern matching. This is useful for tuning rules without creating custom copies.

How It Works

The modify override applies a regex find-and-replace to the rule content at sync time. The original rule in the database is unchanged; the modification is applied when writing the rules file that Suricata reads.

- **Regex**: A Go-compatible regular expression pattern to match
- **Value**: The literal replacement string

> **Note**
>
> The replacement string is treated as literal text. Backreferences (`\1` , `\2` , etc.) are not supported. If you need to capture part of the match, use multiple specific overrides instead.

Examples

Exclude IP Range from Variable

To exclude `$DC_SERVERS` from `$EXTERNAL_NET` :

Regex: `\$EXTERNAL_NET` **Value:** `[$EXTERNAL_NET,!$DC_SERVERS]`

Before:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET any (msg:"Example"; sid:1001;)
```

After:

```
alert tcp [$EXTERNAL_NET,!$DC_SERVERS] any -> $HOME_NET any (msg:"Example"; sid:1001;)
```

Change Threshold Seconds

To change a rule's threshold from 60 seconds to 3600:

Regex: `seconds \d+` **Value:** `seconds 3600`

Before:

```
alert http any any -> any any (msg:"Test"; threshold:type limit,track by_src,count 1,seconds 60; sid:1001;)
```

After:

```
alert http any any -> any any (msg:"Test"; threshold:type limit,track by_src,count 1,seconds 3600; sid:1001;)
```

Modify Content Match

To change a specific content match value:

Regex: `content:"67:98:30` **Value:** `content:"88:98:30`

Before:

```
alert tls any any -> any any (msg:"SSL Cert"; content:"67:98:30:81:90"; sid:1001;)
```

After:

```
alert tls any any -> any any (msg:"SSL Cert"; content:"88:98:30:81:90"; sid:1001;)
```

Limitations

- **No backreferences**: Python-style backreferences (`\1` , `\2`) in the replacement string are not supported. The sync will fail with an error if these are detected.
- **PCRE exception**: Backslash sequences inside `pcre:"..."` sections are allowed, as these are valid PCRE syntax.
- **Literal replacement**: The replacement value is always treated as literal text. Special regex characters in the replacement do not have special meaning.

12.2.4 Adding New NIDS Rules

To add a new NIDS rule, go to the main Detections page and click the blue + button between Options and the query bar. A form will appear where you will:

1. Click the Language drop-down and select `Suricata` .
2. Optionally specify a license.
3. Add the signature.
4. Click the `CREATE` button and the detection should deploy to your grid at the next 15-minute cycle.

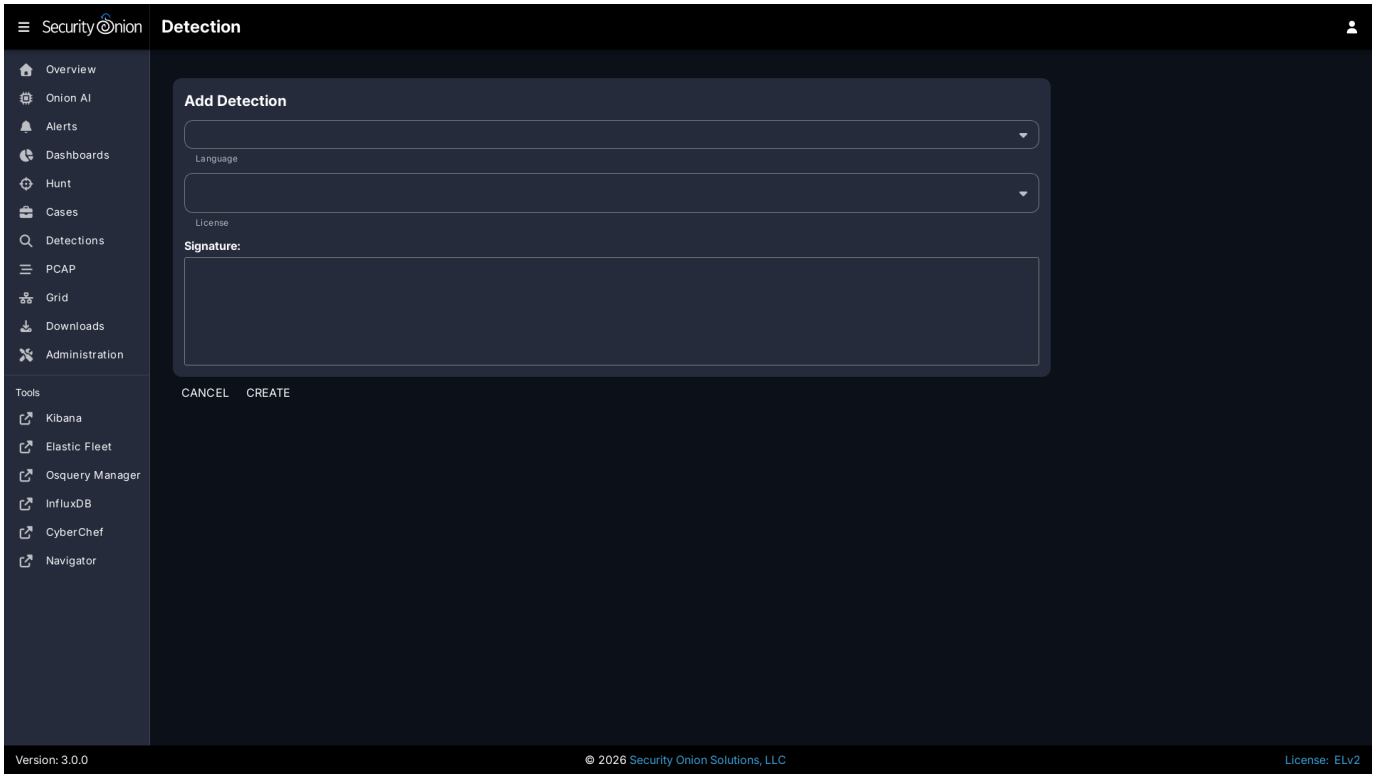

12.2.5 Update Frequency

By default, Security Onion checks for new NIDS rules every 24 hours. You can change this value as follows:

- Navigate to Administration --> Configuration.
- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- Navigate to `SOC` --> `config` --> `server` --> `modules` --> `suricataengine` --> `communityRulesImportFrequencySeconds` .

12.2.6 Allow External Access to NIDS Rules

You can enable external access to NIDS rules managed by Detections. This is useful when configuring OPNsense or other network devices to pull NIDS rules from your Security Onion deployment. You can do this as follows:

- Navigate to Administration --> Configuration.
- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- Navigate to Nginx --> config --> external_suricata.
- On the right side of the page, change the value to `true` and then click the checkmark to save the new setting.
- You can wait for the next Grid update or click the `SYNCHRONIZE GRID` button under Options.
- Once the grid is fully synchronized, the Manager should listen on port 7789 for https connections from hosts defined in the `external_suricata` host group.

12.2.7 Configuring Rulesets

Security Onion allows you to configure multiple NIDS rulesets. You can manage these rulesets by navigating to Administration --> Configuration --> `SOC` --> `config` --> `server` --> `modules` --> `suricataengine` --> `rulesetSources`. This setting is also available via the Configuration quicklinks.

There are two configuration profiles:

- **`default`** : Used for standard (non-Airgap) deployments
- **`airgap`** : Used for Airgap deployments

If your system is in Airgap mode, the Airgap configuration profile will automatically be used - otherwise the default is in use.

Within this configuration, you can enable additional rulesets, add custom rulesets, or disable existing ones. When you save a ruleset configuration change and apply the SOC state, Security Onion will detect the change and automatically sync all configured rulesets within 15 minutes.

OISF-maintained list of Suricata-compatible rulesets: https://github.com/OISF/suricata-intel-index

> **Note**
>
> Each ruleset must have a unique name. Duplicate names will cause sync failures.

Ruleset Configuration Options

Each ruleset source has the following configuration options:

- **Ruleset Name**: Required. The unique name for this ruleset (e.g., "Emerging-Threats", "ABUSECH-SSLBL", "local-rules"). This is the name displayed in the UI.
- **Description**: Optional description of the ruleset.
- **Enabled**: Required. If set to false, existing rules and overrides from this ruleset will be removed.
- **License Key**: Optional. Required for commercial rulesets like ET Pro.
- **Source Type**: Required. Either `url` (downloads rules from HTTPS) or `directory` (reads rules from local filesystem).
- **Source Path**: Required. The full URL or directory/file path depending on Source Type. See Supported Source Path Formats below.
- **Exclude Files**: Optional. List of rule file names to exclude, separated by commas (e.g., `*deleted*, *retired*`).
- **Ruleset License**: Required. The license type for this ruleset (e.g., "BSD", "Commercial", "CC0-1.0").
- **Read Only**: Optional, defaults to false. When enabled, prevents modification of rule content via the UI - users can still enable/disable rules and add tuning overrides (suppress, threshold, modify). Use this for vendor-managed rulesets where you want to preserve the original rule text.
- **Delete Unreferenced**: Optional, defaults to false. Controls what happens to rules in Elasticsearch when they are removed from the source.
- **false** (default): Rules removed from the source remain in Elasticsearch. This preserves user modifications and prevents accidental data loss.
- **true**: Rules removed from the source are automatically deleted from Elasticsearch. Use this when the source is authoritative (e.g., git-managed rulesets).

> **Warning**
>
> Changing this setting from `false` to `true` will delete any rules that no longer exist in the source. If you have orphaned rules (rules in ES but not in source), they will be permanently removed on the next sync.

Supported Source Path Formats

For **url** Source Type:

- URL to a `.rules` file (e.g., `https://rules.emergingthreats.net/open/suricata-7.0.3/emerging-all.rules`)
- URL to a `.tar.gz` archive (e.g., `https://rules.emergingthreats.net/open/suricata-7.0.3/emerging-all.rules.tar.gz`)

For **directory** Source Type:

- Directory containing multiple `.rules` files (e.g., `/nsm/rules/custom-local-repos/local-Suricata-import/`)
- Directory containing a single `.rules` file
- Direct path to a `.rules` file (e.g., `/nsm/rules/custom-local-repos/local-Suricata-import/import.rule`)
- Direct path to a `.tar.gz` archive (e.g., `/nsm/rules/custom-local-repos/local-Suricata-import/import.tar.gz`)

URL Source Options

When using `url` as the Source Type, additional options are available:

- **URL Hash**: URL to a hash file (.md5 or .sha256) for verifying the downloaded ruleset.
- **Proxy URL**: HTTP/HTTPS proxy URL for downloading the ruleset. (e.g., `http://192.168.1.50:3128`)
- **Proxy Username**: Proxy authentication username.
- **Proxy Password**: Proxy authentication password.
- **Proxy CA Path**: Path to CA certificate file for MITM proxy verification (e.g., `/opt/so/saltstack/local/salt/suricata/files/ruleset_ca.crt`)

Default Rulesets

Emerging Threats (ET Open / ET Pro) Security Onion includes the Emerging Threats Open ruleset by default. To switch to ET Pro (commercial), edit the Emerging-Threats ruleset and enter your license key in the License Key field. Click the green checkmark to save, then apply the SOC state. Leave the License Key empty for ET Open (free) rules.

- Optimized for Suricata
- ET Open is **free**, ET Pro requires a license fee per sensor

For more information, see: - https://rules.emergingthreats.net/open/ - https://www.proofpoint.com/us/threat-insight/et-pro-ruleset

Abuse.ch SSL Blacklist (ABUSECH-SSLBL) SSL certificate blacklist from Abuse.ch. Only available in non-Airgap, disabled by default.

For more information, see: - https://sslbl.abuse.ch/

Local Rules A directory-based ruleset source for custom local rules. Rules are read from `/nsm/rules/custom-local-repos/local-Suricata`. This ruleset is enabled by default with Read Only set to false, allowing you to edit rules directly once they are imported. Keep in mind that if the local .rules file remains on disk in this location, each sync will attempt to re-import the rules and potentially overwrite any changes made within the web interface.

Suricata Metadata Rulesets

When Suricata is configured as the metadata engine (instead of Zeek), two additional rulesets become available:

SO_EXTRACTIONS Extraction rules that control which file types Suricata extracts from network traffic for analysis by Strelka. This ruleset is imported and **enabled by default** when Suricata is the metadata engine.

SO_FILTERS Filter rules that control which metadata Suricata logs. Use these to reduce unnecessary metadata logging. This ruleset is imported but **disabled by default** when Suricata is the metadata engine.

12.2.8 Common Ruleset Configurations

This section provides configuration examples for common deployment scenarios.

12.2.9 ET Pro in Airgap Environments

For Airgap deployments using ET Pro (commercial) rules, you must manually transfer the ruleset to your Security Onion Manager since it cannot download from the internet.

Prerequisites:

- Valid ET PRO license key
- A system with internet access to download the ruleset

Procedure:

- **Download the ET PRO ruleset (on internet-connected system)**

Use your license key to download the latest ruleset:

```
# Replace YOUR_LICENSE_KEY with your actual ET PRO license key
curl -o etpro.rules.tar.gz \
  "https://rules.emergingthreatspro.com/YOUR_LICENSE_KEY/suricata-7.0.3/etpro.rules.tar.gz"
```

- **Transfer to Airgapped Manager**

Use your approved file transfer method to copy the archive to the Manager node.

- **Place the ruleset on the Manager**

Copy the archive to a directory accessible by SOC:

```
sudo cp etpro.rules.tar.gz /nsm/rules/custom-local-repos/local-etpro-Suricata/
sudo chown -R socore:socore /nsm/rules/custom-local-repos/local-etpro-Suricata
```

- **Configure the ruleset source**

Navigate to Administration --> Configuration --> `SOC` --> `config` --> `server` --> `modules` --> `suricataengine` --> `rulesetSources`.

Modify the existing `Emerging-Threats` ruleset (recommended):

- **License Key**: `YOUR_LICENSE_KEY`

You can also create a new ruleset source (make sure to disable the existing Emerging-Threats ruleset):

- **Ruleset Name**: `ETPRO-Airgap`
- **Source Type**: `directory`
- **Source Path**: `/nsm/rules/custom-local-repos/local-etpro-Suricata/etpro.rules.tar.gz`
- **Read Only**: `true` (recommended - preserves vendor rule content)
- **Delete Unreferenced**: `true` (recommended - removes outdated rules when you update the archive)
- **Ruleset License**: `Commercial`
- **Enabled**: `true`
- **Apply configuration and sync**

Save the configuration and apply the SOC state. Then either wait for the next automatic sync (up to 15 minutes) or trigger a manual sync:

- Navigate to Detections
- Click Options menu
- Select Suricata engine
- Click `FULL UPDATE`

Updating Rules:

To update your ET Pro rules in an Airgap environment:

- Download the latest `etpro.rules.tar.gz` on an internet-connected system
- Transfer to the airgapped Manager
- Replace the existing archive:

```
sudo cp etpro.rules.tar.gz /nsm/rules/custom-local-repos/local-etpro-Suricata/
```

- Wait for the next automatic sync or trigger a manual `FULL UPDATE`

With `Delete Unreferenced: true`, rules that were removed in the new version will be automatically cleaned up from Elasticsearch.

12.2.10 Flowbit Dependency Handling

Overview

Suricata rules can use **flowbits** to share state between rules. A common pattern is for one rule to detect an initial condition and "set" a flowbit, while other rules check if that flowbit is set before alerting. This creates a dependency between rules.

Security Onion automatically manages these dependencies to ensure your enabled rules function correctly, even when you disable related rules.

How Flowbits Work

Flowbits allow rules to communicate within a single network flow:

- **Setter rules** use `flowbits:set,name` to mark a flow
- **Getter rules** use `flowbits:isset,<name>` to check if a flow was marked

For example, a malware detection might work like this:

- **Rule A** (setter): Detects initial malware handshake, sets `flowbits:set,malware.detected`
- **Rule B** (getter): Detects follow-up command, requires `flowbits:isset,malware.detected`

Rule B will only alert if Rule A has already matched on the same flow. If Rule A is disabled, Rule B can never trigger.

Automatic Dependency Resolution

When you disable a rule that sets a flowbit needed by other enabled rules, Security Onion automatically handles this:

- **Your preference is preserved** - The rule remains marked as "disabled" in Elasticsearch & SOC
- **The rule still runs** - It is included in the active ruleset so dependent rules can function
- **No alerts are generated** - The `noalert` option is automatically added so the disabled rule runs silently

Example

Consider these rules:

SID	Rule Name	Flowbit	Your Setting
2012236	x0Proto Init	`flowbits:set,et.x0proto`	**Disabled**
2012237	x0Proto Client Info	`flowbits:isset,et.x0proto`	Enabled
2012238	x0Proto Pong	`flowbits:isset,et.x0proto`	Enabled

Even though you disabled rule 2012236, it will still run because rules 2012237 and 2012238 depend on it. However:

- Rule 2012236 will **not** generate alerts
- Rules 2012237 and 2012238 will alert normally when their conditions match

In the rules file, you will see a comment explaining the automatic inclusion:

```
# AUTO-ENABLED (flowbit: et.x0proto, required by 2 rule(s)): This disabled rule runs with noalert
alert tcp $EXTERNAL_NET any -> $HOME_NET any (msg:"x0Proto Init"; ... noalert; sid:2012236; ...)
```

When Disabled Rules Are Excluded

A disabled setter rule is only auto-enabled if at least one getter rule depends on it. If you disable **all** rules that check a particular flowbit, the setter rule will be excluded from the active ruleset entirely.

Using the example above, if you disable all three rules (2012236, 2012237, and 2012238), then rule 2012236 will not be included in the rules file since no enabled rules need its flowbit.

12.2.11 Sync Block

For the upgrade to 2.4.200, the dependency on idstools has been removed and all functionality has been moved directly into SOC. Because of the complexity of this change, if SOUP detects a non-default Suricata ruleset configuration, it creates a block file that stops any further Suricata ruleset changes until the block file has been removed.

> **Warning**
>
> This block is critical because if the Suricata rulesets are synced without the configuration properly migrated, all current rules and overrides in SOC Detections will be removed and will need to be recreated.

To resolve this block, use the following procedure:

- **Review the `syncBlock` file**

Login to the Manager and view the `syncBlock` file to see what non-default configuration was detected:

```
sudo cat /opt/so/conf/soc/fingerprints/suricataengine.syncBlock
```

Example output showing ET Pro was detected:

```
Suricata ruleset sync is blocked until this file is removed.
**CRITICAL** Make sure that you have manually added any custom Suricata
rulesets via SOC config before removing this file - review the documentation
for more details: <https://securityonion.net/docs/nids#sync-block>
Custom so-rule-update detected (hash: 207d8918a2d963bb7dcc0f1ebf28d6f7b5778019fedf0cc36d5d0850cbd8a529)
ET Pro code found: YOUR_LICENSE_KEY
```

Note any license codes or custom configurations mentioned - you will need to enter these in the next step.

- **Configure Suricata rulesets in SOC**

Navigate to SOC Configuration:

- Administration --> Configuration --> Quicklinks --> **Configure NIDS Rulesets**

There are two configuration profiles:

- **`default`** : Used for standard (non-Airgap) deployments
- **`airgap`** : Used for Airgap deployments

Select the appropriate profile for your environment.

For ET Pro configurations - Non-Airgap:

- Find the `Emerging-Threats` ruleset entry
- Copy and paste your ET Pro license code (shown in the syncBlock file) into the `License Key` field

For ET Pro configurations - Airgap:

Following the procedure outlined here: ET Pro in Airgap Environments

During this migration, it is important to use the built-in `Emerging Threats` Airgap config profile.

For proxy configurations:

If your environment requires a proxy to download rulesets, configure the proxy settings on the ruleset entry:

- **Proxy URL**: Your proxy server URL (e.g., `http://192.168.1.50:3128`)
- **Proxy Username** / **Proxy Password**: If proxy authentication is required
- **Proxy CA Path**: Path to CA certificate if using a MITM proxy

For custom rulesets:

If you had custom rulesets configured, add new ruleset entries with the appropriate Source Type, Source Path, and other settings. See Configuring Rulesets for details.

- **Save and synchronize configuration**
- Click the green checkmark to save your changes
- Click `SYNCHRONIZE SOC` and wait for it to complete
- **Remove the syncBlock file**

Once the configuration is saved and synchronized, remove the block file:

```
sudo rm /opt/so/conf/soc/fingerprints/suricataengine.syncBlock
```

- **Trigger a full ruleset sync**
- Navigate to Detections
- Click the `Options` menu
- In the engine dropdown, select `Suricata`
- Click `FULL UPDATE`
- **Verify successful sync**

Within a minute or so, you should see a success message: `Synchronized Suricata rules successfully.`

The engine status indicator should clear to `OK`.

If the sync fails, click the `Sync Failure` crosshair icon in the top-right corner of the Suricata engine to view the error details.

12.3 Sigma

Sigma rules are loaded into ElastAlert to monitor incoming logs for suspicious or noteworthy activity. Active Sigma rules generate alerts that can then be found in Alerts.

From https://github.com/SigmaHQ/sigma:

> Sigma is a generic and open signature format that allows you to describe relevant log events in a straightforward manner. The rule format is very flexible, easy to write and applicable to any type of log file. The main purpose of this project is to provide a structured form in which researchers or analysts can describe their once developed detection methods and make them shareable with others. Sigma is for log files what Snort is for network traffic and YARA is for files.

12.3.1 Managing Existing Sigma Rules

You can manage existing Sigma rules via Detections. There are two ways to do so:

- From the main Detections interface, you can search for the desired detection and click the binoculars icon.
- From the Alerts interface, you can click an alert and then click the `Tune Detection` menu item.

Once you've used one of these methods to reach the detection detail page, you can check the Status field in the upper-right corner and use the slider to enable or disable the detection.

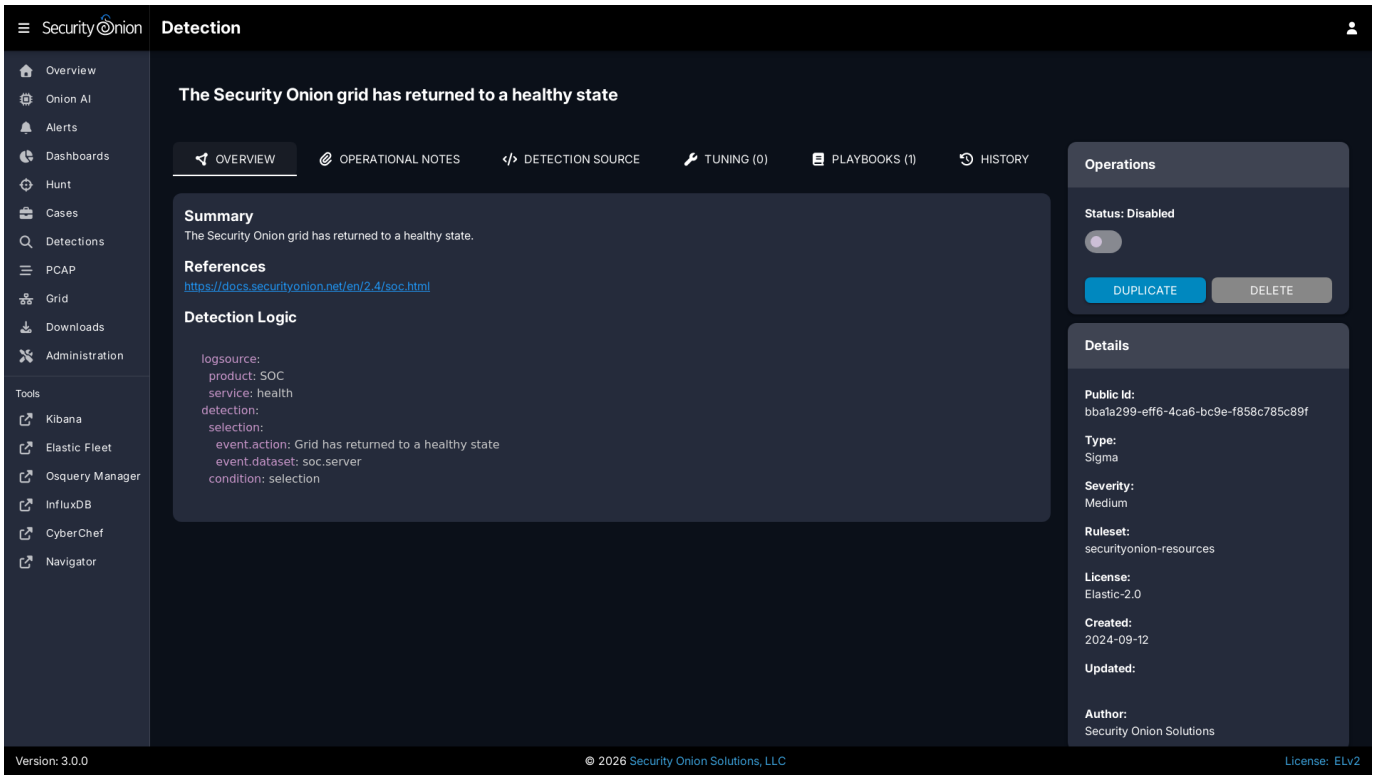

To tune the detection:

- Click the Tuning tab
- Click the blue + button
- Select the type of tuning (Custom Filter)
- Enter your custom filter in the Custom Filter field
- Click the `CREATE` button to create and enable the Override

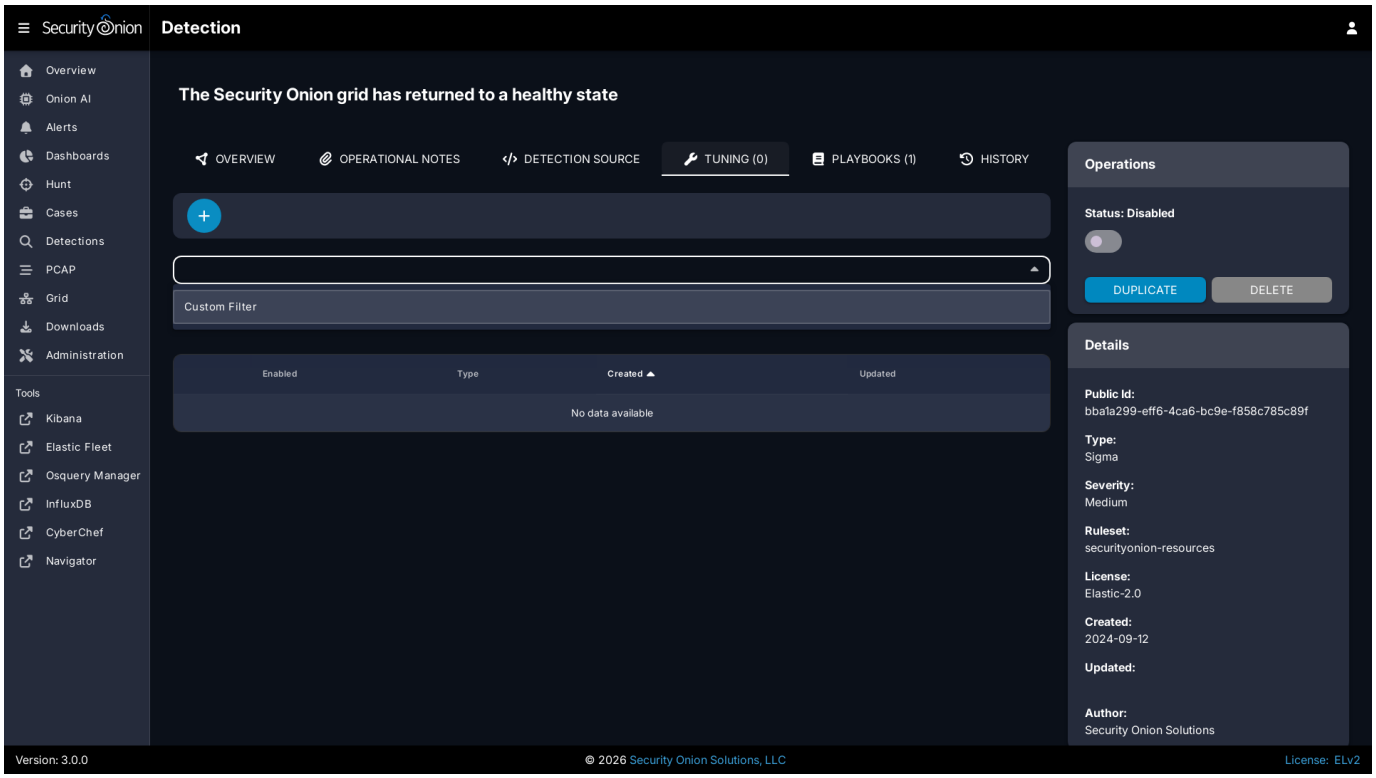

Custom Filters are Sigma Search Identifiers and will be applied like so: `"($ORIGINAL_CONDITION) and not 1 of sofilter*"`

For example, suppose that you have an IDH node installed with the HTTP webserver enabled. Your nightly vulnerability scan is connecting to it and generating an alert from the `Security Onion IDH - HTTP Access` detection. To filter out connection attempts from this scanner, you would add the following Custom Filter to this detection:

```
sofilter:
  src_ip|cidr: 192.168.55.45/32
```

Once you save this filter, it is enabled by default for this detection. Clicking on the `Detection Source` tab and then on `Convert` will show you what the new EQL query looks like, which should include a filter for the IP address.

For more information on Sigma rule syntax, please see the Sigma documentation at https://sigmahq.io/docs/basics/rules.html#detection.

12.3.2 Adding New Sigma Rules

To add a new Sigma rule, go to the main Detections page and click the blue + button between Options and the query bar. A form will appear where you will:

1. Click the Language drop-down and select `Sigma`.
2. Optionally specify a license.
3. Add the signature.
4. Click the `CREATE` button and the detection should deploy to your grid at the next 15-minute cycle.

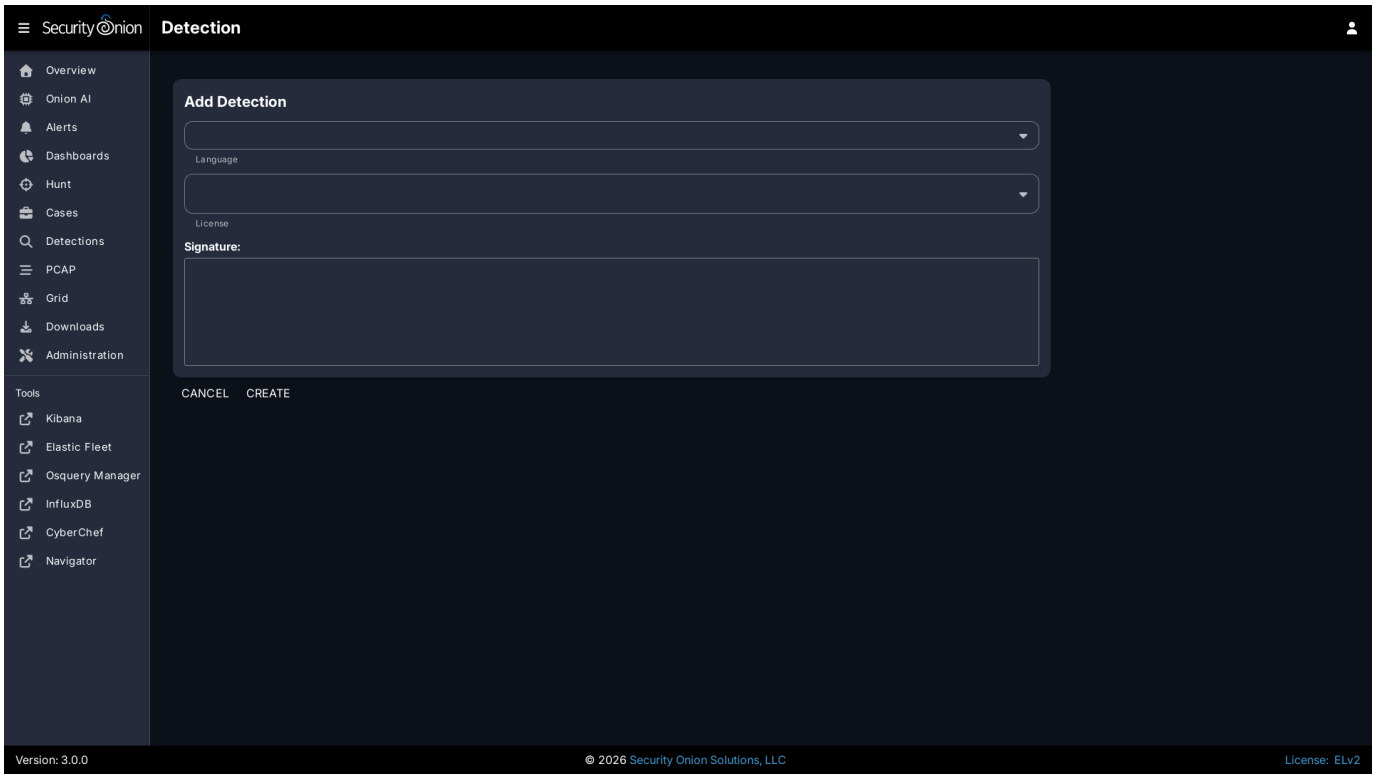

12.3.3 Sigma Configuration

- Navigate to Administration --> Configuration.
- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- Navigate to `SOC` --> `config` --> `server` --> `modules` --> `elastalertengine`.

Once you've reached this location, here are some common settings.

Sigma Update Frequency

By default, Security Onion checks for new Sigma rules every 24 hours. You can change this value at `SOC` --> `config` --> `server` --> `modules` --> `elastalertengine` --> `communityRulesImportFrequencySeconds`.

Sigma Packages

You can choose from different Sigma packages:

https://github.com/SigmaHQ/sigma/blob/master/Releases.md

You can modify this setting via `SOC` --> `config` --> `server` --> `modules` --> `elastalertengine` --> `sigmaRulePackages`.

Custom Sigma Repositories

You can configure Security Onion to pull Sigma rules from custom git repos via `SOC` --> `config` --> `server` --> `modules` --> `elastalertengine` --> `rulesRepos` --> `default`.

Repos can be accessed via https or from the local filesystem. For example:

```
file:///nsm/rules/detect-sigma/repos/my-custom-rep
```

Enable Sigma Rules on Import

```
`SOC` > `config` > `server` > `modules` > `elastalertengine` > `enabledSigmaRules` > `default`
```

This configuration options allows you to specify which rules are automatically enabled upon initial import. The format for this filter is a YAML list that supports flexible filtering criteria based on a number of fields in a Sigma rule. A rule is enabled only if it matches all specified filters - if there is more than one filter for a field, then it has to match at least one.

Configuration Format

Each item in the YAML list represents a set of filters, using the following fields:

- `ruleset` : List of strings. Specifies the ruleset(s) to filter by (e.g., "core", "securityonion-resources", "*" for any ruleset).
- `level` : List of strings. Specifies the severity level(s) (e.g., "critical", "high", "*" for any level. This is not a greater than or equal check - just a string match).
- `product` : List of strings. Specifies the product(s) to filter by (e.g., "windows", "*" for any products).
- `category` : List of strings. Specifies the event category or categories (e.g., "process_creation", "registry_event", "*" for any category).
- `service` : List of strings. Specifies the service(s) to filter by (e.g., "security", "dns-client", "*" for any service).

For example:

```
# Enable all critical and high rules from the "securityonion-resources" ruleset
- ruleset: ["securityonion-resources"]
  level: ["critical", "high"]
  product: ["*"]
  category: ["*"]
  service: ["*"]
```

12.4 YARA

YARA rules are loaded into Strelka to monitor files for suspicious or noteworthy characteristics. Active YARA rules generate alerts that can be found in Alerts.

From https://virustotal.github.io/yara/:

> YARA is a tool aimed at (but not limited to) helping malware researchers to identify and classify malware samples. With YARA you can create descriptions of malware families (or whatever you want to describe) based on textual or binary patterns. Each description, a.k.a rule, consists of a set of strings and a boolean expression which determine its logic.

12.4.1 Managing Existing YARA Rules

You can manage existing YARA rules via Detections. There are two ways to do so:

- From the main Detections interface, you can search for the desired detection and click the binoculars icon.
- From the Alerts interface, you can click an alert and then click the `Tune Detection` menu item.

Once you've used one of these methods to reach the detection detail page, you can check the Status field in the upper-right corner and use the slider to enable or disable the detection.

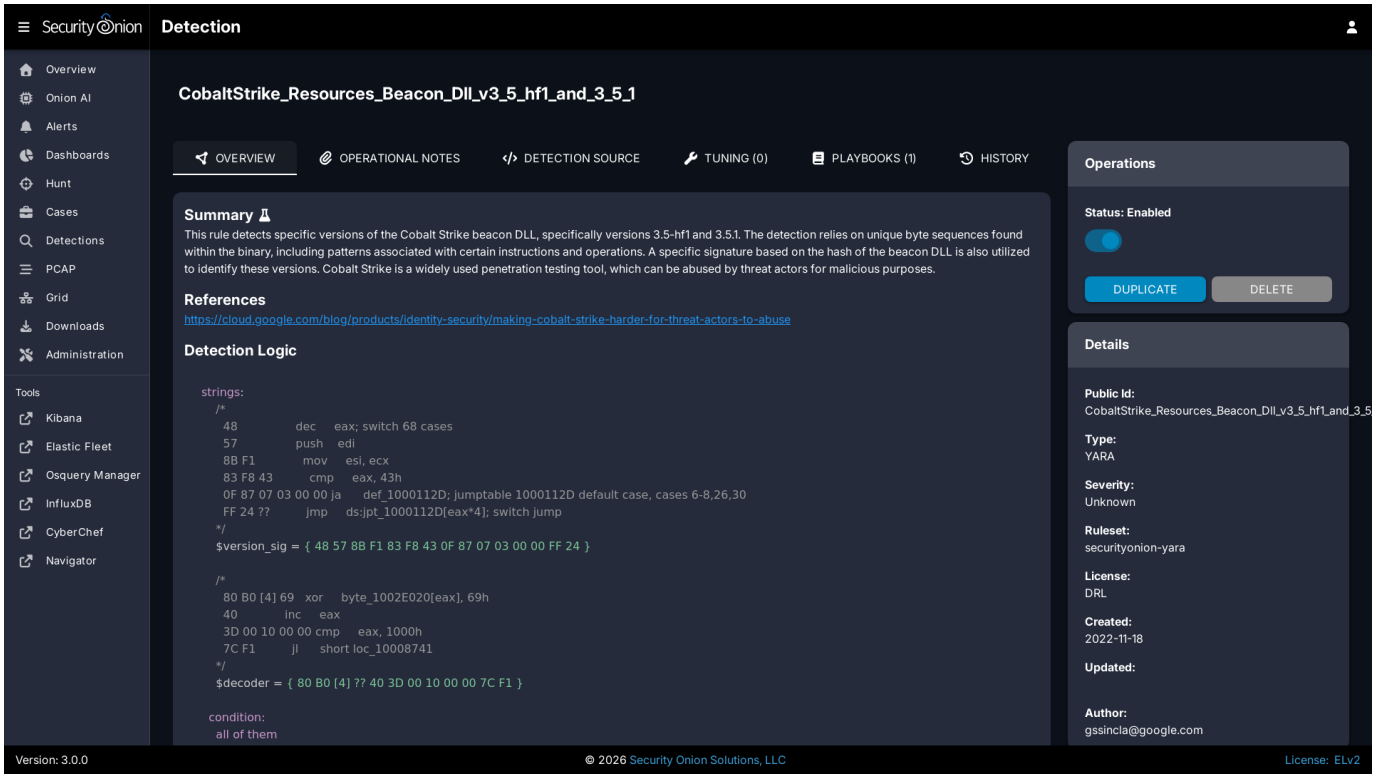

12.4.2 Adding New YARA Rules

To add a new YARA rule, go to the main Detections page and click the blue + button between Options and the query bar. A form will appear where you will:

1. Click the Language drop-down and select `YARA`.
2. Optionally specify a license.
3. Add the signature.
4. Click the `CREATE` button and the detection should deploy to your grid at the next 15-minute cycle.

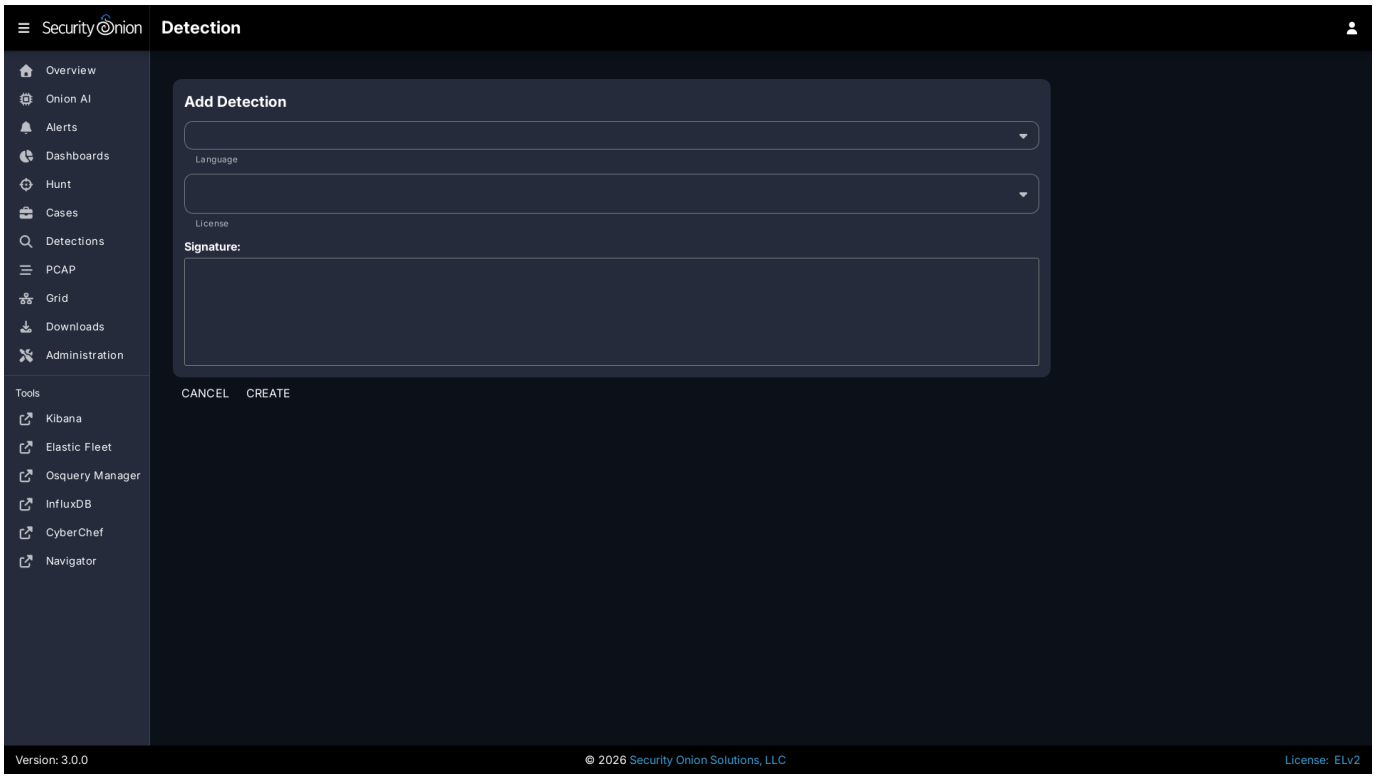

12.4.3 YARA Rules Options

You can configure YARA rules options as follows:

- Navigate to Administration --> Configuration.
- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- Navigate to `SOC` --> `config` --> `server` --> `modules` --> `strelkaengine`.

Once you've reached this location, here are some common settings.

YARA Update Frequency

By default, Security Onion checks for new YARA rules every 24 hours. You can change this value at `SOC` --> `config` --> `server` --> `modules` --> `strelkaengine` --> `communityRulesImportFrequencySeconds`.

Custom YARA Repositories

You can configure Security Onion to pull YARA rules from custom git repos via `SOC` --> `config` --> `server` --> `modules` --> `strelkaengine` --> `rulesRepos` --> `default`.

Repos can be accessed via https or from the local filesystem. For example:

```
file:///nsm/rules/detect-yara/repos/my-custom-rep
```

13. Logs

13.1 Logs Overview

Once logs are generated by network sniffing processes or endpoints, where do they go? How are they parsed? How are they stored? That's what we'll discuss in this section.

13.2 Ingest

Here's an overview of how logs are ingested in various deployment types.

13.2.1 Import

Core Pipeline: Elastic Agent [IMPORT Node] --> Elasticsearch Ingest [IMPORT Node] **Logs:** Zeek, Suricata

13.2.2 Eval

Core Pipeline: Elastic Agent [EVAL Node] --> Elasticsearch Ingest [EVAL Node] **Logs:** Zeek, Suricata

13.2.3 Standalone

Core Pipeline: Elastic Agent [SA Node] --> Logstash [SA Node] --> Redis [SA Node] <--> Logstash [SA Node] --> Elasticsearch Ingest [SA Node] **Logs:** Zeek, Suricata, syslog

Elastic Agent: Elastic Agent [Windows Endpoint] --> Logstash [SA Node] --> Redis [SA Node] <--> Logstash [SA Node] --> Elasticsearch Ingest [SA Node] **Logs:** WEL, Sysmon

13.2.4 Fleet Standalone

Pipeline: Elastic Agent [Fleet Node] --> Logstash [M | MS] --> Elasticsearch Ingest [S | MS] **Logs:** Elastic Agent

13.2.5 Manager (separate search nodes)

Core Pipeline: Elastic Agent [Fleet | Sensor] --> Logstash [Manager] --> Redis [Manager] **Logs:** Zeek, Suricata, syslog

Elastic Agent: Elastic Agent [Windows Endpoint] --> Logstash [Manager] --> Redis [Manager] **Logs:** WEL, Sysmon

13.2.6 Manager Search

Core Pipeline: Elastic Agent [Fleet | Sensor] --> Logstash [MS] --> Redis [MS] <--> Logstash [MS] --> Elasticsearch Ingest [MS] **Logs:** Zeek, Suricata, syslog

Pipeline: Elastic Agent [MS] --> Logstash [MS] --> Elasticsearch Ingest [MS] **Logs:** Local Elastic Agent

Elastic Agent: Elastic Agent [Windows Endpoint] --> Logstash [MS] --> Elasticsearch Ingest [MS] **Logs:** WEL, Sysmon

13.2.7 Heavy

Pipeline: Elastic Agent [Heavy Node] --> Elasticsearch Ingest [Heavy] **Logs:** Zeek, Suricata, syslog

13.2.8 Search

Pipeline: Redis [Manager] --> Logstash [Search] --> Elasticsearch Ingest [Search] **Logs:** Zeek, Suricata, syslog

13.2.9 Sensor

Pipeline: Elastic Agent [Sensor] --> Logstash [M | MS] --> Elasticsearch Ingest [S | MS] **Logs:** Zeek, Suricata, syslog

13.3 Logstash

From https://www.elastic.co/products/logstash:

> Logstash is a free and open server-side data processing pipeline that ingests data from a multitude of sources, transforms it, and then sends it to your favorite "stash."

When Security Onion is running in Standalone mode or in a full distributed deployment, Logstash transports unparsed logs to Elasticsearch which then parses and stores those logs. It's important to note that Logstash does NOT run when Security Onion is configured for Import or Eval mode. You can read more about that in the Architecture section.

13.3.1 Configuration

You can configure Logstash by going to Administration --> Configuration --> Logstash.

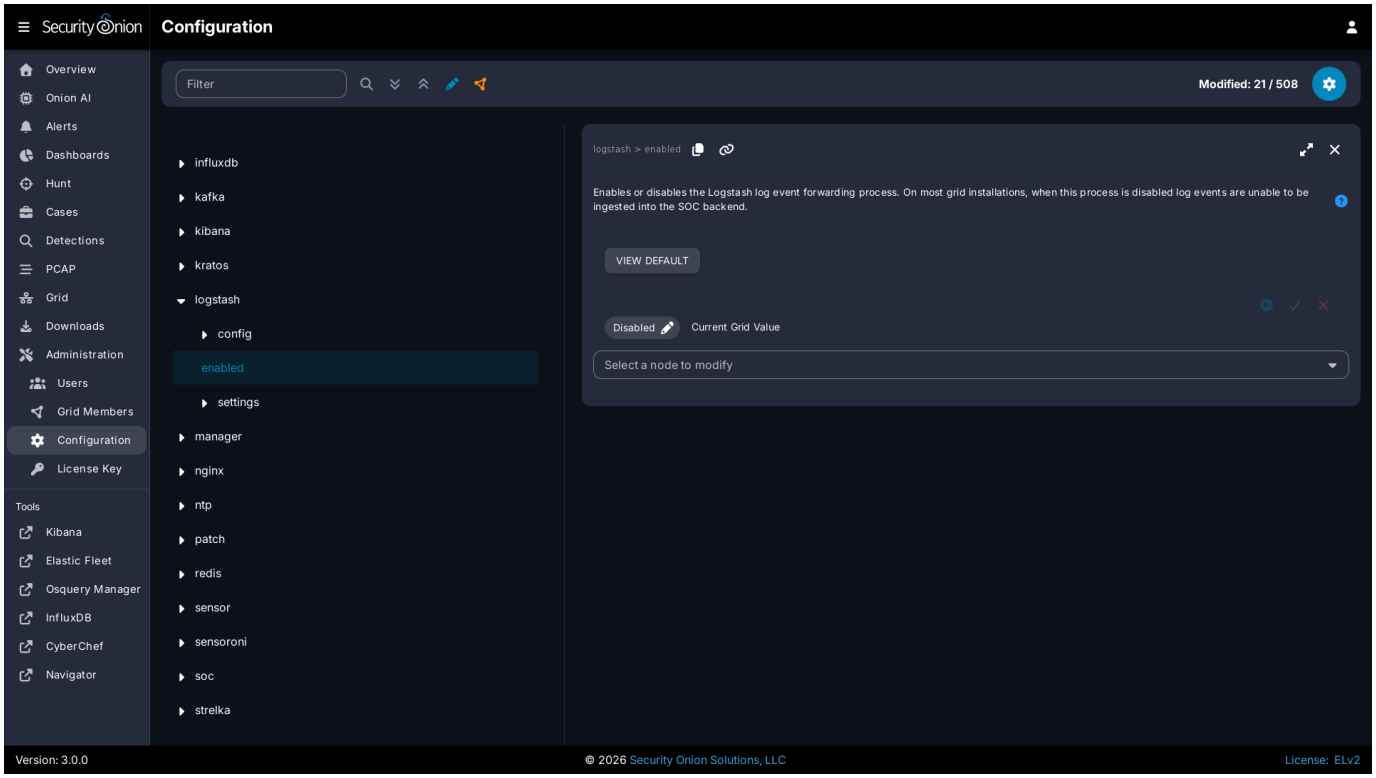

ls_pipeline_batch_size

> The maximum number of events an individual worker thread will collect from inputs before attempting to execute its filters and outputs. Larger batch sizes are generally more efficient, but come at the cost of increased memory overhead. This is set to 125 by default.

ls_pipeline_workers

> The number of workers that will, in parallel, execute the filter and output stages of the pipeline. If you find that events are backing up, or that the CPU is not saturated, consider increasing this number to better utilize machine processing power. By default this value is set to the number of cores in the system.

For more information, please see https://www.elastic.co/guide/en/logstash/current/logstash-settings-file.html.

lsheap

If total available memory is 8GB or greater, Setup sets the Logstash heap size to 25% of available memory, but no greater than 4GB.

For more information, please see https://www.elastic.co/guide/en/elasticsearch/guide/current/heap-sizing.html#compressed_oops.

You may need to adjust the value depending on your system's performance. The changes will be applied the next time the minion checks in. You can force it to happen immediately by running `sudo salt-call state.apply Logstash` on the actual node or by running `sudo salt $SENSORNAME_$ROLE state.apply Logstash` on the manager node.

13.3.2 Parsing

Logstash does not parse logs in Security Onion, so modifying existing parsers or adding new parsers should be done via Elasticsearch.

13.3.3 Forwarding Events to an External Destination

Please keep in mind that we don't provide free support for third party systems, so this section will be just a brief introduction to how you would send syslog to external syslog collectors. If you need commercial support, please see https://www.securityonionsolutions.com.

13.3.4 Original Event Forwarding

To forward events to an external destination with minimal modifications to the original event, create a new custom configuration file on the manager in `/opt/so/saltstack/local/salt/logstash/pipelines/config/custom/` for the applicable output. We recommend using either the `http`, `tcp`, `udp`, or `syslog` output plugin. At this time we only support the default bundled Logstash output plugins.

For example, to forward all Zeek events from the `dns` dataset, we could use a configuration like the following:

```
output {
  if [event][module] == "Zeek" and [pipeline] == "dns" {
    udp {
      id => "cloned_events_out"
      host => "192.168.x.x"
      port => 1001
      codec => "json_lines"
    }
  }
}
```

> **Warning**
>
> When using the `tcp` output plugin, if the destination host or port is down, it will cause the Logstash pipeline to be blocked. To avoid this, try using the other output options or consider having forwarded logs use a separate Logstash pipeline.
>
> Also keep in mind that when forwarding logs from the manager, some fields may not be set as expected since the events have not yet been processed by the Ingest Node configuration.

In SOC, navigate to Administration --> Configuration. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option. Then navigate to Logstash --> defined_pipelines --> manager and append the name of your newly created file to the list of config files used for the `manager` pipeline:

```
custom/myfile.conf
```

The configuration will be applied at the next 15-minute interval or you can apply it immediately by clicking the `SYNCHRONIZE GRID` button under the `Options` menu.

You can monitor events flowing through the output by running the following command on the manager:

```
curl -s localhost:9600/_node/stats | jq .pipelines.manager
```

13.3.5 Modified Event Forwarding

To forward events to an external destination AFTER they have traversed the Logstash pipelines (NOT ingest node pipelines), perform the same steps as above but instead of adding the reference for your Logstash output to the `manager` pipeline add it to `search` pipeline instead. The configuration will be applied at the next 15-minute interval or you can apply it immediately by clicking the `SYNCHRONIZE GRID` button under the `Options` menu.

You can monitor events flowing through the output by running the following command on the search nodes:

```
curl -s localhost:9600/_node/stats | jq .pipelines.search
```

Please keep in mind that events will be forwarded from all applicable search nodes, as opposed to just the manager.

13.3.6 Queue

Memory-backed

From https://www.elastic.co/guide/en/logstash/current/persistent-queues.html:

> By default, Logstash uses in-memory bounded queues between pipeline stages (inputs → pipeline workers) to buffer events. The size of these in-memory queues is fixed and not configurable.

Persistent

If you experience adverse effects using the default memory-backed queue, you might consider a disk-based persistent queue. From https://www.elastic.co/guide/en/logstash/current/persistent-queues.html:

> In order to protect against data loss during abnormal termination, Logstash has a persistent queue feature which will store the message queue on disk. Persistent queues provide durability of data within Logstash.

Queue Max Bytes

> The total capacity of the queue in number of bytes. Make sure the capacity of your disk drive is greater than the value you specify here. If both `queue.max_events` and `queue.max_bytes` are specified, Logstash uses whichever criteria is reached first.

Dead Letter Queue

If you want to check for dropped events, you can enable the dead letter queue. This will write all records that are not able to make it into Elasticsearch into a sequentially-numbered file (for each start/restart of Logstash).

This can be achieved by adding the following to the Logstash configuration:

```
dead_letter_queue.enable: true
```

and restarting Logstash:

```
sudo so-logstash-restart
```

The dead letter queue files are located in `/nsm/logstash/dead_letter_queue/main/`.

More information: - https://www.elastic.co/guide/en/logstash/current/dead-letter-queues.html

Redis

When using search nodes, Logstash on the manager node outputs to Redis (which also runs on the manager node). Redis queues events from the Logstash output (on the manager node) and the Logstash input on the search node(s) pull(s) from Redis. If you notice new events aren't making it into Elasticsearch, you may want to first check Logstash on the manager node and then the Redis queue.

13.3.7 Diagnostic Logging

The Logstash log file is located at `/opt/so/log/logstash/logstash.log`. Log file settings can be adjusted in `/opt/so/conf/logstash/etc/log4j2.properties`. By default, logs are set to rollover daily and purged after 7 days. Depending on what you're looking for, you may also need to look at the Docker logs for the container:

```
sudo docker logs so-logstash
```

13.3.8 Errors

Read-Only

```
[INFO ][Logstash.outputs.Elasticsearch] retrying failed action with response code: 403 ({"type"=>"cluster_block_exception", "reason"=>"blocked by: [FORBIDDEN/
12/index read-only / allow delete (api)];"})
```

This error is usually caused by the `cluster.routing.allocation.disk.watermark` (`low`, `high`) being exceeded.

You may want to check `/opt/so/log/elasticsearch/<hostname>.log` to see specifically which indices have been marked as read-only.

Additionally, you can run the following command to allow writing to the affected indices:

```
curl -k -XPUT -H 'Content-Type: application/json' https://localhost:9200/<your_index>/_settings -d'{ "index.blocks.read_only": false }'
```

13.3.9 More Information

> **Note**
>
> For more information about Logstash, please see https://www.elastic.co/products/logstash.

13.4 Redis

From https://redis.io/:

> Redis is an open source (BSD licensed), in-memory data structure store, used as a database, cache and message broker. It supports data structures such as strings, hashes, lists, sets, sorted sets with range queries, bitmaps, hyperloglogs and geospatial indexes with radius queries.

On Standalone (non-Eval) installations and distributed deployments, Logstash on the manager node outputs to Redis. Search nodes can then consume from Redis.

13.4.1 Queue

To see how many logs are in the Redis queue:

```
sudo so-redis-count
```

If the queue is backed up and doesn't seem to be draining, try stopping Logstash on the manager node:

```
sudo so-logstash-stop
```

Then monitor the queue to see if it drains:

```
watch 'sudo so-redis-count'
```

If the Redis queue looks okay, but you are still having issues with logs getting indexed into Elasticsearch, you will want to check the Logstash statistics on the search node(s).

13.4.2 Tuning

Security Onion configures Redis to use 812MB of your total system memory. If you have sufficient RAM available, you may want to increase the `redis_maxmemory` setting by going to Administration --> Configuration --> Redis. This value is in Megabytes so to set it to use 8 gigs of ram you would set the value to 8192.

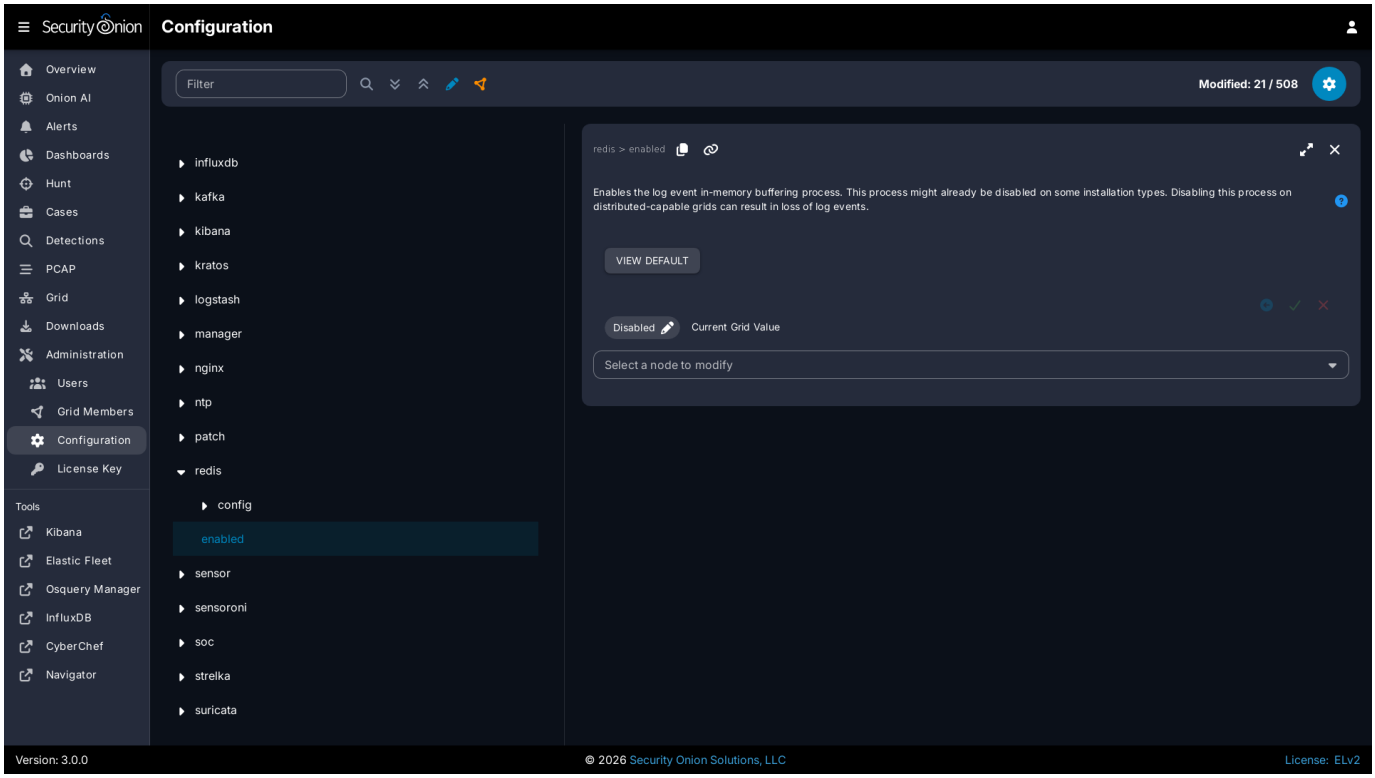

Logstash on the manager node is configured to send to Redis. For best performance, you may want to tune the `ls_pipeline_batch_size` value at Administration --> Configuration --> logstash_settings to find the sweet spot for your deployment.

> **Note**
>
> For more information about the Logstash output plugin for Redis, please see: https://www.elastic.co/guide/en/logstash/current/plugins-outputs-Redis.html

Logstash on search nodes pulls from Redis. For best performance, you may want to tune `ls_pipeline_batch_size` and `ls_input_threads` at Administration --> Configuration --> logstash_settings to find the sweet spot for your deployment.

> **Note**
>
> For more information about the Logstash input plugin for Redis, please see: https://www.elastic.co/guide/en/logstash/current/plugins-inputs-Redis.html

13.4.3 Diagnostic Logging

Redis logs can be found at `/opt/so/log/redis/`. Depending on what you're looking for, you may also need to look at the Docker logs for the container:

```
sudo docker logs so-redis
```

13.4.4 More Information

> **Note**
>
> For more information about Redis, please see https://redis.io/.

13.5 Elasticsearch

From https://www.elastic.co/products/elasticsearch:

> Elasticsearch is a distributed, RESTful search and analytics engine capable of addressing a growing number of use cases. As the heart of the Elastic Stack, it centrally stores your data for lightning fast search, fine-tuned relevancy, and powerful analytics that scale with ease.

13.5.1 Storage

All of the data Elasticsearch collects is stored under `/nsm/elasticsearch/`.

> **Warning**
>
> Do not manually delete any files in `/nsm/elasticsearch`! If you need to delete Elasticsearch indices, this should be done through Elasticsearch itself rather than deleting files from the filesystem.

13.5.2 Schema

Security Onion tries to adhere to Elastic Common Schema (ECS) wherever possible. In some cases, additional fields or slight modifications to native Elastic field mappings may be found within the data. You can learn more about ECS at https://www.elastic.co/elasticsearch/common-schema.

13.5.3 Querying

You can query Elasticsearch using web interfaces like Alerts, Dashboards, Hunt, and Kibana. You can also query Elasticsearch from the command line using so-elasticsearch-query.

13.5.4 Authentication

You can authenticate to Elasticsearch using the same username and password that you use for SOC.

You can add new user accounts to both Elasticsearch and SOC at the same time as shown in the Adding Accounts section. Please note that if you instead create accounts directly in Elastic, then those accounts will only have access to Elastic and not SOC.

13.5.5 Indexing

Most data is associated with a data stream, which is an abstraction from traditional indices that leverages one or more backing indices to manage and represent the data within the data stream. The usage of data streams allows for greater flexibility in data management.

Data streams can be targeting during search or other operations directly, similar to how indices are targeted.

For example, a CLI-based query against Zeek connection records would look like the following:

```
so-elasticsearch-query logs-zeek-so/_search?q=event.dataset:conn
```

When this query is run against the backend data, it is actually targeting one or more backing indices, such as:

```
.ds-logs-zeek-so-2022-03-07.0001
.ds-logs-zeek-so-2022-03-08.0001
.ds-logs-zeek-so-2022-03-08.0002
```

Similarly, you can target a single backing index with the following query:

```
so-elasticsearch-query .ds-logs-zeek-so-2022-03-08.001/_search?q=event.dataset:conn
```

You can learn more about data streams at https://www.elastic.co/guide/en/elasticsearch/reference/current/data-streams.html.

13.5.6 Configuration

You can configure Elasticsearch by going to Administration --> Configuration --> Elasticsearch.

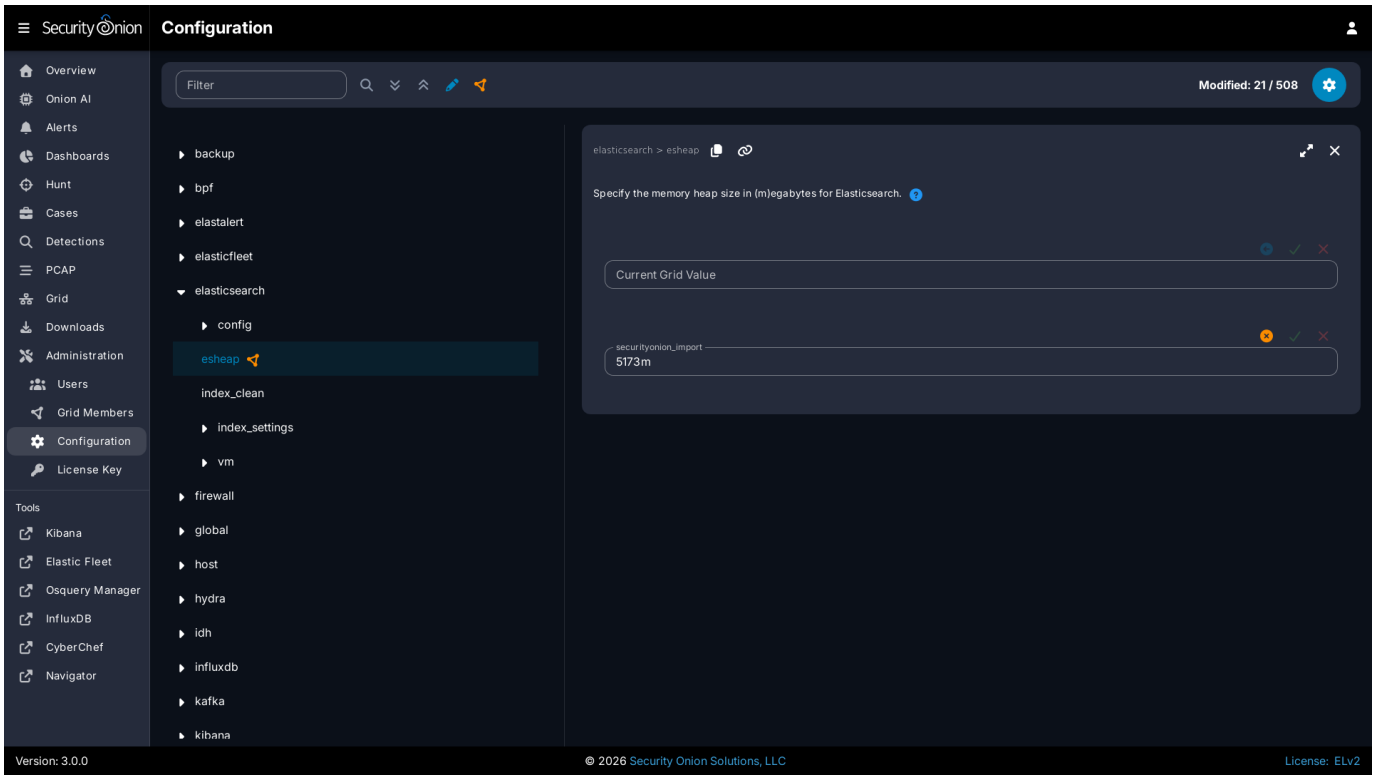

13.5.7 Parsing

Elasticsearch receives unparsed logs from Logstash or Elastic Agent. Elasticsearch then parses and stores those logs. Parsers are stored in `/opt/so/conf/elasticsearch/ingest/`. Custom ingest parsers can be placed in `/opt/so/saltstack/local/salt/elasticsearch/files/ingest/`. To make these changes take effect, restart Elasticsearch using `so-elasticsearch-restart`.

Elastic Agent may pre-parse or act on data before the data reaches Elasticsearch, altering the data stream or index to which it is written, or other characteristics such as the event dataset or other pertinent information. This configuration is maintained in the agent policy or integration configuration in Elastic Fleet.

> **Note**
>
> You can learn more about Elasticsearch ingest parsing at: https://www.elastic.co/guide/en/elasticsearch/reference/current/ingest.html

13.5.8 Cluster

In a distributed deployment with a manager and one or more search nodes, the manager and search nodes are joined together as a multi-node Elastic cluster.

For more information, please see the cluster information at https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-node.html#modules-node.

13.5.9 Elasticsearch Node Roles

Please see the Elasticsearch node roles documentation at https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-node.html.

When building a distributed deployment, the Security Onion manager has to start off with the `data` node role. If you later join a separate search node, then you may want to migrate the data from the manager to the search node and then remove the `data` node role from the manager. For more information, please see:

https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-cluster.html#cluster-shard-allocation-filtering

If you want to set certain search nodes to the `data_hot`, `data_warm`, or `data_cold` roles, make sure you remove the `data` role from them. You will also want to ensure that `data_content` is assigned to your hot nodes.

> **Warning**
>
> Elasticsearch node roles is an advanced setting and you should be careful to avoid disruption to your cluster!

To see and modify Elasticsearch node roles, first navigate to Administration --> Configuration, click the `Options` menu at the top of the page, and enable the `Show advanced settings` option. Then navigate to Elasticsearch --> so_roles and select the desired role. Finally, navigate to config --> node --> roles and the list of roles should appear on the right side of the page.

13.5.10 Templates

Fields are mapped to their appropriate data type using templates. When making changes for parsing, it is necessary to ensure fields are mapped to a data type to allow for indexing, which in turn allows for effective aggregation and searching in Dashboards, Hunt, and Kibana. Elasticsearch leverages both component and index templates.

> **Note**
>
> For managing templates for third party integrations, please see the third-party-integrations section.

Component Templates

From https://www.elastic.co/guide/en/elasticsearch/reference/current/index-templates.html:

> Component templates are reusable building blocks that configure mappings, settings, and aliases. While you can use component templates to construct index templates, they aren't directly applied to a set of indices. Also see https://www.elastic.co/guide/en/elasticsearch/reference/current/indices-component-template.html.

Index Templates

From https://www.elastic.co/guide/en/elasticsearch/reference/current/index-templates.html:

> An index template is a way to tell Elasticsearch how to configure an index when it is created. Templates are configured prior to index creation. When an index is created - either manually or through indexing a document - the template settings are used as a basis for creating the index. Index templates can contain a collection of component templates, as well as directly specify settings, mappings, and aliases.

In Security Onion, component templates are stored in `/opt/so/saltstack/default/salt/elasticsearch/templates/component/`.

These templates are specified to be used in the index template definitions in `/opt/so/saltstack/default/salt/elasticsearch/defaults.yml`.

13.5.11 Community ID

For logs that don't naturally include Community ID, we use the Elasticsearch Community ID processor: https://www.elastic.co/guide/en/elasticsearch/reference/current/community-id-processor.html

13.5.12 field expansion matches too many fields

If you get errors like `failed to create query: field expansion for [*] matches too many fields, limit: 3500, got: XXXX`, then this usually means that you're sending in additional logs and so you have more fields than our default `max_clause_count` value. To resolve this, you can go to

Administration --> Configuration --> Elasticsearch --> config --> indices --> query --> bool --> max_clause_count and adjust the value for any boxes running Elasticsearch in your deployment.

13.5.13 Shards

Here are a few tips from https://www.elastic.co/blog/how-many-shards-should-i-have-in-my-elasticsearch-cluster:

> **TIP**: Avoid having very large shards as this can negatively affect the cluster's ability to recover from failure. There is no fixed limit on how large shards can be, but a shard size of 50GB is often quoted as a limit that has been seen to work for a variety of use-cases.
>
> **TIP**: Small shards result in small segments, which increases overhead. Aim to keep the average shard size between a few GB and a few tens of GB. For use-cases with time-based data, it is common to see shards between 20GB and 40GB in size.
>
> **TIP**: The number of shards you can hold on a node will be proportional to the amount of heap you have available, but there is no fixed limit enforced by Elasticsearch. A good rule-of-thumb is to ensure you keep the number of shards per node below 20 to 25 per GB heap it has configured. A node with a 30GB heap should therefore have a maximum of 600-750 shards, but the further below this limit you can keep it the better. This will generally help the cluster stay in good health.

To see your existing shards, run the following command and the number of shards will be shown in the fifth column:

```
sudo so-elasticsearch-indices-list
```

If you want to view the detail for each of those shards:

```
sudo so-elasticsearch-shards-list
```

Given the sizing tips above, if any of your indices are averaging more than 50GB per shard, then you should probably increase the shard count until you get below that recommended maximum of 50GB per shard.

The number of shards for an index can be adjusted by going to Administration --> Configuration --> Elasticsearch --> index_settings --> so-INDEX-NAME --> index_template --> template --> settings --> index --> number_of_shards.

Please keep in mind that old indices will retain previous shard settings and the above settings will only be applied to newly created indices.

13.5.14 Heap Size

If total available memory is 8GB or greater, Setup configures the heap size to be 33% of available memory, but no greater than 25GB. You may need to adjust the value for heap size depending on your system's performance. You can modify this by going to Administration --> Configuration --> Elasticsearch --> esheap.

For more information, please see:

- https://www.elastic.co/guide/en/elasticsearch/guide/current/heap-sizing.html#compressed_oops
- https://www.elastic.co/guide/en/elasticsearch/reference/current/important-settings.html#heap-size-settings

13.5.15 Field limit

Security Onion currently defaults to a field limit of 5000. If you receive error messages from Logstash, or you would simply like to increase this, you can do so by going to Administration --> Configuration --> Elasticsearch --> index_settings --> so-INDEX-NAME --> index_template --> template --> settings --> index --> mapping --> total_fields --> limit.

Please note that the change to the field limit will not occur immediately, only on index creation.

13.5.16 Re-indexing

Re-indexing may need to occur if field data types have changed and conflicts arise. This process can be VERY time-consuming, and we only recommend this if keeping data is absolutely critical.

For more information about re-indexing, please see: https://www.elastic.co/guide/en/elasticsearch/reference/current/docs-reindex.html

13.5.17 Clearing

If you want to clear all Elasticsearch data including documents and indices, you can run the `so-elastic-clear` command.

13.5.18 GeoIP

Elasticsearch 8 no longer includes GeoIP databases by default. We include GeoIP databases for Elasticsearch so that all users will have GeoIP functionality. If your search nodes have Internet access and can reach geoip.elastic.co and storage.googleapis.com, then you can opt-in to database updates if you want more recent information. To do so, run the following command on your manager:

```
sudo so-elasticsearch-query _cluster/settings -d '{"persistent":{"ingest.geoip.downloader.enabled":true}}' -XPUT
```

13.5.19 Health

To check Elasticsearch health, go to the Grid interface and check the Elasticsearch Status field. If it shows anything other than OK, then run the following command from the CLI on the manager node to check for additional clues:

```
sudo so-elasticsearch-query _cluster/health?pretty
```

13.5.20 Status Pending

If the Grid interface shows Elasticsearch Status as `Pending`, check for unassigned shards by running the following command from the CLI on the manager node:

```
sudo so-elasticsearch-query _cat/shards | grep UN
```

The result of the query should display affected indices. Older metrics indices for Elastic Endpoint logs may have been assigned a replica, so if you are running a single-node Elastic cluster there will be nowhere for the replica to exist. To resolve the issue, run the following command for each affected index (replacing `$index` with the actual index name):

```
sudo so-elasticsearch-query $index/_settings -d '{"number_of_replicas":0}' -XPUT
```

After running the command, the index should no longer use replicas and the status should change from "Pending" to "OK" once all indices have been successfully modified.

13.5.21 Index Management

Elasticsearch indices are managed by both the `so-elasticsearch-indices-delete` utility and Index Lifecycle Management (ILM).

> **Note**
>
> Check out our Index Lifecycle Management video at https://youtu.be/Y6HVein7nP8!

> **Warning**
>
> `so-elasticsearch-indices-delete` is primarily designed for single-node deployments (IMPORT, EVAL, and STANDALONE). Running it on a multi-node deployment with one or more search nodes has the possibility of getting into a corner case state where more data is deleted than intended. Therefore, this script is disabled on multi-node deployments. If you have a multi-node deployment, then you will need to ensure that ILM is configured properly to delete indices before disk usage reaches the Elasticsearch watermark setting. Otherwise, Elasticsearch may stop ingesting new data.

so-elasticsearch-indices-delete

`so-elasticsearch-indices-delete` manages size-based deletion of Elasticsearch indices based on the value of the `Elasticsearch.retention.retention_pct` setting. This setting is checked against the total disk space available for `/nsm/elasticsearch` across all nodes in the Elasticsearch cluster. If your indices are using more than `retention_pct`, then `so-elasticsearch-indices-delete` will delete old

indices until disk space consumed by indices is back under `retention_pct`. The default value for this setting is `50` percent so that standalone deployments have sufficient space for not only Elasticsearch but also full packet capture and other logs. For distributed deployments with dedicated search nodes where Elasticsearch is main consumer of disk space, you may want to increase this default value.

To modify the `retention_pct` value, first navigate to Administration --> Configuration. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option. Then navigate to Elasticsearch --> retention --> retention_pct. Once you make the change and save it, the new setting will take effect at the next 15 minute interval. If you would like to make the change immediately, you can click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

ILM

Index Lifecycle Management (ILM) manages the following:

- size-based index rollover
- time-based index rollover
- time-based content tiers
- time-based index deletion

Time-based index deletion is based on the `min_age` setting within the global policy or the individual policy for the index itself. Please note that size-based deletion via `so-elasticsearch-indices-delete` takes priority over time-based deletion, as disk usage may reach `retention_pct` and indices will be deleted before the `min_age` value is reached.

ILM settings can be found by navigating to Administration --> Configuration --> Elasticsearch --> index_settings:

- To edit the global policy that applies to ALL indices, navigate to global_overrides --> policy --> phases and there you will see the cold, delete, hot, and warm ILM phases.
- To edit the policy for an individual index, first click the `Options` menu at the top of the page and then enable the `Show advanced settings` option. Then navigate to $index --> policy --> phases. There you will see the cold, delete, hot, and warm ILM phases for that particular index.
- It's important to note that settings like `min_age` are calculated relative to the rollover date (NOT the original creation date of the index). For example, if you have an index that is set to rollover after 30 days and delete `min_age` is set to 30 then there will be 30 days from index creation to rollover and then an additional 30 days before deletion.
- When modifying ILM settings, note that some settings will only take effect after a new index is created.

Now that you have an overview of all that ILM can do, here's a very high level overview of how you would configure ILM deletion for your deployment:

- Determine your data retention requirements. This might be 1 week, 1 month, or more. It may also be different for different kinds of data.
- Determine your current daily ingestion. One way to do this is to go to Kibana, select the menu on the left, select Stack Management, and then go to Index Management to see what your current indices look like. Another option is to run `sudo so-elasticsearch-indices-growth` from the command line.
- Now that you have your data retention requirements and current daily ingestion, use those values to determine your storage requirements. Keep in mind that Elasticsearch's default watermark setting of 80% means that you will want to keep 20% of your disk free and this will need to be accounted for in your storage requirements. If your storage requirements are greater than the amount of storage that you have available, then you may need to add additional search nodes.
- Configure ILM Deletion to delete logs before hitting the Elasticsearch 80% watermark. This can be done globally for all indices by going to Administration -> Configuration -> Elasticsearch > index_settings > global_overrides > policy > phases > delete > min_age. Again, keep in mind that the `min_age` setting is calculated relative to the index rollover date and NOT the original creation date of the index. If you want to specify different deletion values for different kinds of data, then you can enable advanced settings and then drill into specific policies to do so.

> **Tip**
>
> You might want to run `sudo so-elasticsearch-indices-growth` on a regular basis to keep an eye on the size of your indices.

In addition to `sudo so-elasticsearch-indices-growth`, you can also run `sudo so-elasticsearch-retention-estimate` which will give you an approximation of how many days' worth of logs you can store. For example:

```
DISCLAIMER: Script output is based on current data patterns, but are approximations solely
  intended to assist with getting a general ILM policy configured.

================ Storage Overview ================

Indexed data size:       141.12 GB (Elasticsearch)
Cluster capacity:        498.69 GB total
Cluster used:            245.32 GB
Low watermark:           80% (398.96 GB threshold)
Remaining space:         153.63 GB before low watermark
Cluster shards:          498 / 4000 (12.4%)
Cluster data nodes:      4
  example-search1: data_hot,data_warm,data_content
  example-search2: data_hot,data_warm,data_content
  example-search3: data_cold
  example-man: data_content

================ ES Growth ================

Daily growth rate:       4.08 GB/day
ILM deletion rate:       0.47 GB/day (scheduled)
Net growth rate:         3.62 GB/day
Daily shard creation:    ~3 shards/day
Storage to be freed (30d): 2 indices (~13.97 GB, 4 shards)

================ Retention Projection ================

Oldest index:            ~55 days (.ds-logs-auditd_manager.auditd-default-2025.09.30-000001)
Estimated retention:     ~98 days (until configured low watermark setting)

 Low watermark breach estimated in ~42.47 days (2026-01-06)
```

For maximum retention, our goal is to get the cluster balanced as close to the low watermark setting as possible. In this example, it appears the cluster is gaining about 4GB worth of logs per day. However, ILM is currently deleting roughly 0.5GB per day, so overall storage usage is increasing. Without tweaking ILM configuration, this cluster will hit the watermark in roughly 42 days (total retention being roughly 98 days). To combat this, one option is to set a global_overrides for the delete phase as described above setting the delete phase to something like 90 days. This gives us a bit of space between the estimated retention and our actual delete phase.

In addition to global overrides, which apply to all indices, it is possible to tweak per index ILM policies. For example, perhaps Suricata alert data is something you need to keep in storage for 120 days. You can configure the so-suricata.alerts policy to have a delete phase of 120d. This comes at the cost of needing to reduce retention on other indices in order to free up the needed storage for Suricata alerts.

Tip

Once you have ILM configured, you can consider increasing the cluster low / high watermark settings to allow Elasticsearch to use more of the available disk space. This can be done by going to Administration --> Configuration --> Elasticsearch --> config --> cluster --> routing --> allocation --> disk --> watermark.

Note

You can learn more about ILM at: https://www.elastic.co/guide/en/elasticsearch/reference/current/index-lifecycle-management.html

13.5.22 Diagnostic Logging

- Elasticsearch logs can be found in `/opt/so/log/elasticsearch/`.
- Logging configuration can be found in `/opt/so/conf/elasticsearch/log4j2.properties`.

Depending on what you're looking for, you may also need to look at the Docker logs for the container:

```
sudo docker logs so-elasticsearch
```

13.5.23 More Information

Note

For more information about Elasticsearch, please see: https://www.elastic.co/products/elasticsearch

13.6 ElastAlert 2

From https://elastalert2.readthedocs.io/en/latest/elastalert.html#overview:

> ElastAlert 2 is a simple framework for alerting on anomalies, spikes, or other patterns of interest from data in Elasticsearch.

ElastAlert queries Elasticsearch and provides an alerting mechanism with multiple output types, such as Slack, Email, JIRA, OpsGenie, and many more.

13.6.1 Sigma Rules

The Detections module will generate ElastAlert 2 compatible rules automatically for all Sigma detections. There is no need to manually modify the generated rules on disk. Further, any modifications will be overwritten during the next Sigma rule synchronization.

Adjusting a Sigma rule should always be done via the Detections screen.

See the notifications section for information on how to enable outbound notifications via the Detections module.

13.6.2 Custom Rules

Custom ElastAlert 2 rules, which are not associated to the Detections module, can be added to the Security Onion Manager node inside a custom subdirectory under the `/opt/so/rules/elastalert/rules/` directory. For example, create a subdirectory called `/opt/so/rules/elastalert/rules/custom/` and place custom rules within that directory.

> **Warning**
>
> Do not modify or add rules to the `/opt/so/rules/elastalert/rules/` directory itself as those rules are overwritten by the Detections module.

Refer to the ElastAlert 2 documentation, linked above, for detailed information on how to write custom rules. Be aware that writing rules requires an in-depth understanding of Elasticsearch document records, their data structure, and other related concepts.

13.6.3 Diagnostic Logging

ElastAlert diagnostic logs are in `/opt/so/log/elastalert/` and may also appear in the Docker logs for the container. To view container logs run the following command on the Manager:

```
sudo docker logs so-elastalert
```

ElastAlert 2 stores rule status information, such as number of hits, times each rule last ran, etc. to Elasticsearch indices. This data can be helpful in assisting with troubleshooting custom rules. You can search for this data in Dashboards, Hunt, or Kibana. SOC does not automatically include the `ElastAlert` indices by default. If you would like to include them, you can adjust the appropriate configuration setting. In SOC, navigate to Administration --> Configuration. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option. Then filter for `elastic.index` to locate the setting. On the right side of the screen, add `*:Elastalert*` to the existing `index` setting. The updated setting should resemble the following:

```
*:so-*,*:endgame-*,*:logs-*,*:ElastAlert*
```

so-elastalert-create

`so-elastalert-create` can be used to help ease the pain of ensuring correct syntax and creating ElastAlert rules from scratch. It will walk you through various questions, and eventually output an ElastAlert rule file that you can deploy in your environment to start alerting quickly and easily.

so-elastalert-test

`so-elastalert-test` is a wrapper script for ElastAlert's `elastalert-test-rule` tool. The script allows you to test an ElastAlert rule and get results immediately. Simply run `so-elastalert-test`, and follow the prompt(s).

> **Note**
>
> `so-elastalert-test` does not yet include all options available to `elastalert-test-rule`.

Performance

For better performance, avoid writing rules that return large numbers of records. Instead, use the `use_count_query: true` in each rule file. This will only return counts of matching records and not the records themselves.

Timeframe

For queries that span greater than a minute back in time, you may want to add the following fields to your rule to ensure searching occurs as planned (for example, for 10 minutes):

```
buffer_time:
    minutes: 10
```

```
allow_buffer_time_overlap: true
```

For more information, please see: - https://elastalert2.readthedocs.io/en/latest/ruletypes.html#buffer-time - https://github.com/Yelp/elastalert/issues/805

13.6.4 Configuration

You can modify ElastAlert 2 configuration by going to Administration --> Configuration --> ElastAlert.

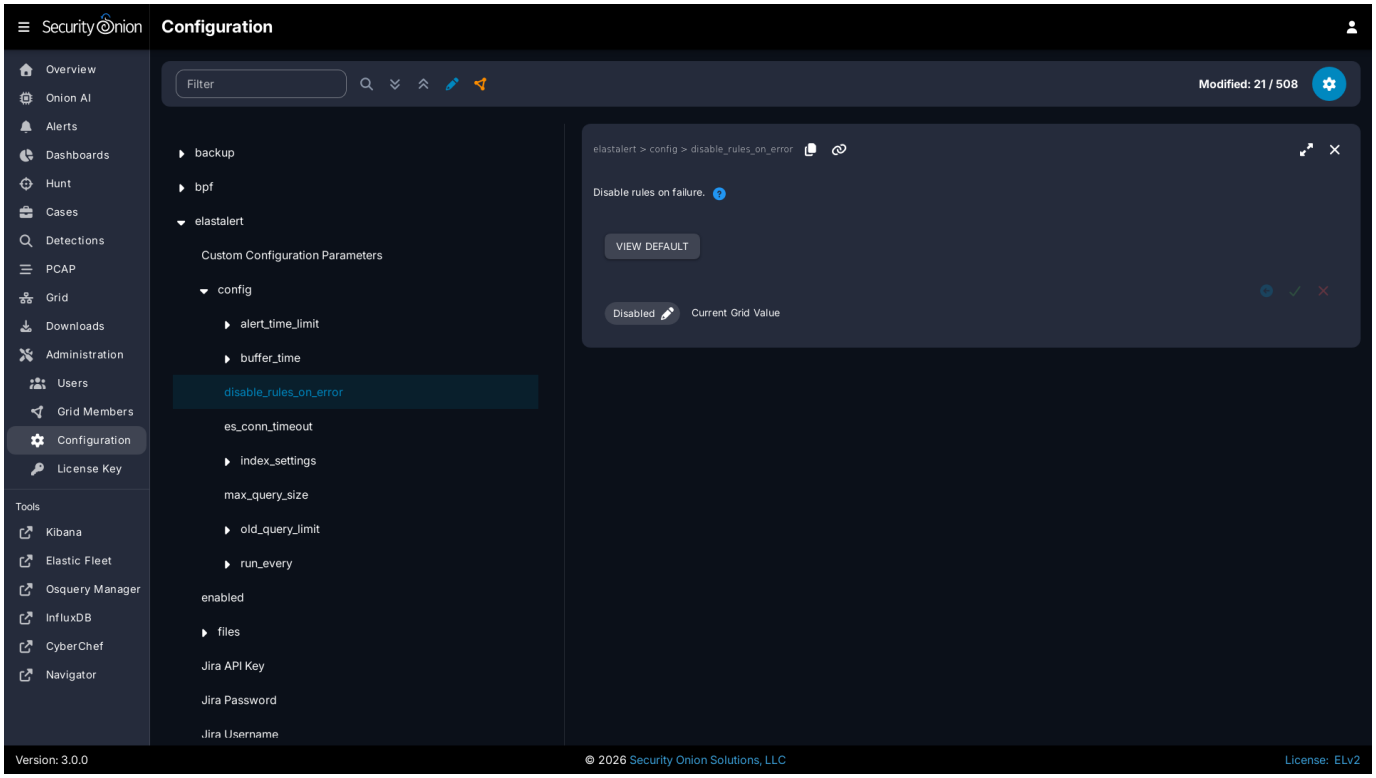

13.6.5 More Information

Note

For more information about ElastAlert, please see https://elastalert2.readthedocs.io/en/latest/.

13.7 Data Fields

This page references the various types of data fields utilized by the Elastic Stack in Security Onion.

13.7.1 ECS

We try to align with Elastic Common Schema (ECS) where possible.

> **Note**
>
> For more information about ECS, please see https://www.elastic.co/guide/en/ecs/current/ecs-reference.html.

13.7.2 Fields

- Alert Data Fields
- ElastAlert Fields
- Zeek Fields

13.7.3 Template files

Fields are mapped to their proper type using template files found in `/opt/so/conf/elasticsearch/templates/`.

13.8 Alert Data Fields

Elasticsearch receives NIDS alerts from Suricata via Elastic Agent or Logstash and parses them using:

- `/opt/so/conf/elasticsearch/ingest/suricata.alert`
- `/opt/so/conf/elasticsearch/ingest/common.NIDS`
- `/opt/so/conf/elasticsearch/ingest/common`

You can find these online at:

- https://github.com/Security-Onion-Solutions/securityonion/blob/3/main/salt/elasticsearch/files/ingest/suricata.alert
- https://github.com/Security-Onion-Solutions/securityonion/blob/3/main/salt/elasticsearch/files/ingest/common.nids
- https://github.com/Security-Onion-Solutions/securityonion/blob/3/main/salt/elasticsearch/files/ingest-dynamic/common

You can find parsed NIDS alerts in Alerts, Dashboards, Hunt, and Kibana via their predefined queries and dashboards or by manually searching for:

- `event.module:"Suricata"`
- `event.dataset:"alert"`

Those alerts should have the following fields:

- `source.ip`
- `source.port`
- `destination.ip`
- `destination.port`
- `network.transport`
- `rule.gid`
- `rule.name`
- `rule.rule`
- `rule.rev`
- `rule.severity`
- `rule.uuid`
- `rule.version`

13.9 ElastAlert Fields

The following list includes field names as they are formatted in Elasticsearch. ElastAlert provides its own template to use for mapping into ElastAlert, so we do not currently utilize a config file to parse data from ElastAlert.

`index:*:elastalert_status`

- alert_info.type
- alert_sent
- alert_time
- endtime
- hist
- matches
- match_body.@timestamp
- match_body.num_hits
- match_body.num_matches
- rule_name
- starttime
- time_taken

13.10 Zeek Fields

Zeek logs are sent to Elasticsearch where they are parsed using ingest parsing. Most Zeek logs have a few standard fields and they are parsed as follows:

- `ts` => `@timestamp`
- `uid` => `log.id.uid`
- `id.orig_h` => `source.ip`
- `id.orig_p` => `source.port`
- `id.resp_h` => `destination.ip`
- `id.resp_p` => `destination.port`

The remaining fields in each log are specific to the log type. To see how the fields are mapped for a specific Zeek log, take a look at its ingest parser.

You can find ingest parsers in your local filesystem at `/opt/so/conf/elasticsearch/ingest/` or you can find them online at:

https://github.com/Security-Onion-Solutions/securityonion/tree/3/main/salt/elasticsearch/files/ingest

For example, suppose you want to know how the Zeek conn.log is parsed. You could take a look at `/opt/so/conf/elasticsearch/ingest/zeek.conn` or view it online at:

https://github.com/Security-Onion-Solutions/securityonion/blob/3/main/salt/elasticsearch/files/ingest/zeek.conn

You'll see that `Zeek.conn` then calls the `Zeek.common` pipeline (`/opt/so/conf/elasticsearch/ingest/zeek.common`):

https://github.com/Security-Onion-Solutions/securityonion/blob/3/main/salt/elasticsearch/files/ingest/zeek.common

which in turn calls the `common` pipeline (`/opt/so/conf/elasticsearch/ingest-dynamic/common`):

https://github.com/Security-Onion-Solutions/securityonion/blob/3/main/salt/elasticsearch/files/ingest-dynamic/common

13.11 Community ID

From https://github.com/corelight/community-id-spec:

> When processing flow data from a variety of monitoring applications (such as Zeek and Suricata), it's often desirable to pivot quickly from one dataset to another. While the required flow tuple information is usually present in the datasets, the details of such "joins" can be tedious, particular in corner cases. This spec describes "Community ID" flow hashing, standardizing the production of a string identifier representing a given network flow, to reduce the pivot to a simple string comparison.

Security Onion enables the built-in Community ID support in both Zeek and Suricata.

For logs that don't naturally include Community ID, we use the Elasticsearch Community ID processor: https://www.elastic.co/guide/en/elasticsearch/reference/current/community-id-processor.html

Dashboards includes a Community ID dashboard that will show all logs with that value.

13.11.1 More Information

Note

For more information about Community ID, please see: https://github.com/corelight/community-id-spec

13.12 Security Onion Console Logs

Standard Security Onion Console logs can be found at `/opt/so/log/soc/` .

13.12.1 SOC Auth Logs

SOC auth is handled by Kratos and you can read more about that at https://github.com/ory/kratos. SOC auth logs can be found at `/opt/so/log/kratos/` . Those logs are ingested into Elasticsearch and available for searching in Dashboards, Hunt, and Kibana. Both Dashboards and Hunt have pre-defined queries for SOC auth logs.

14. Updating

14.1 Updating Overview

In this section, we'll cover keeping Security Onion up-to-date via soup and list important EOL dates for older versions of Security Onion.

14.2 soup

`soup` stands for `Security Onion UPdater` and you can use it to update your Security Onion deployment.

14.2.1 SSH

> **Warning**
>
> If you run `soup` via an SSH session and that SSH session terminates, then any processes running in that session would terminate. You should avoid leaving `soup` unattended especially if the machine you are SSHing from is configured to sleep after a period of time. You might also consider using something like screen or tmux so that if your SSH session terminates, the processes will continue running on the server.

14.2.2 Production Deployments

> **Warning**
>
> If you have a production deployment, we recommend that you test the upgrade process on a test deployment if possible before deploying to production.

14.2.3 Updating

To update your Security Onion deployment, run the `soup` command with sudo:

```
sudo soup
```

If necessary, `soup` will update itself and then ask you to run `soup` again. Once `soup` is fully updated, it will then check for other updates. This includes Security Onion version updates, Security Onion hotfixes, and operating system (OS) updates.

After running `soup` or rebooting a Security Onion node, it may take a few minutes for services to display an `OK` status on the Grid screen. This may be due to the intial on-boot Salt highstate running. If services do not appear to be fully up and running within 15 minutes, try running the following command:

```
sudo so-checkin
```

14.2.4 Security Onion Version Updates

When we release a new version of Security Onion, we update the Release Notes section and publish a blog post to https://blog.securityonion.net. You'll want to review these for any relevant information about the individual updates.

If `soup` finds a full version update, then it will update the Security Onion version in `/etc/soversion`, all Salt code, and all Docker images.

`soup` automatically keeps the previous version of Docker images. These older unused Docker images will be automatically removed at the next version update. If you need to remove these older Docker images immediately, first verify that the upgrade completed successfully and that everything is working properly. You could then remove the older images individually or all at once using a command like:

```
sudo docker system prune -a
```

However, please note that this an aggressive option and you should exercise caution if you have any non-standard Docker images or configuration. You may want to test it on a test system first.

14.2.5 Security Onion Hotfixes

`soup` checks for Security Onion hotfixes. Hotfixes typically include updates to the Salt code and small configuration changes that do not warrant a full version update. This does not include Docker images since that would require a full version update.

After applying a hotfix, you may notice that the Security Onion version in `/etc/soversion` stays the same. The application of the hotfix is tracked on the manager in the `/etc/sohotfix` file.

14.2.6 OS Updates

In addition to Security Onion docker image updates, `soup` also checks for missing OS updates and asks if you want to install them.

Automatic OS Updates

You can configure automatic OS updates by going to Administration --> Configuration --> patch.

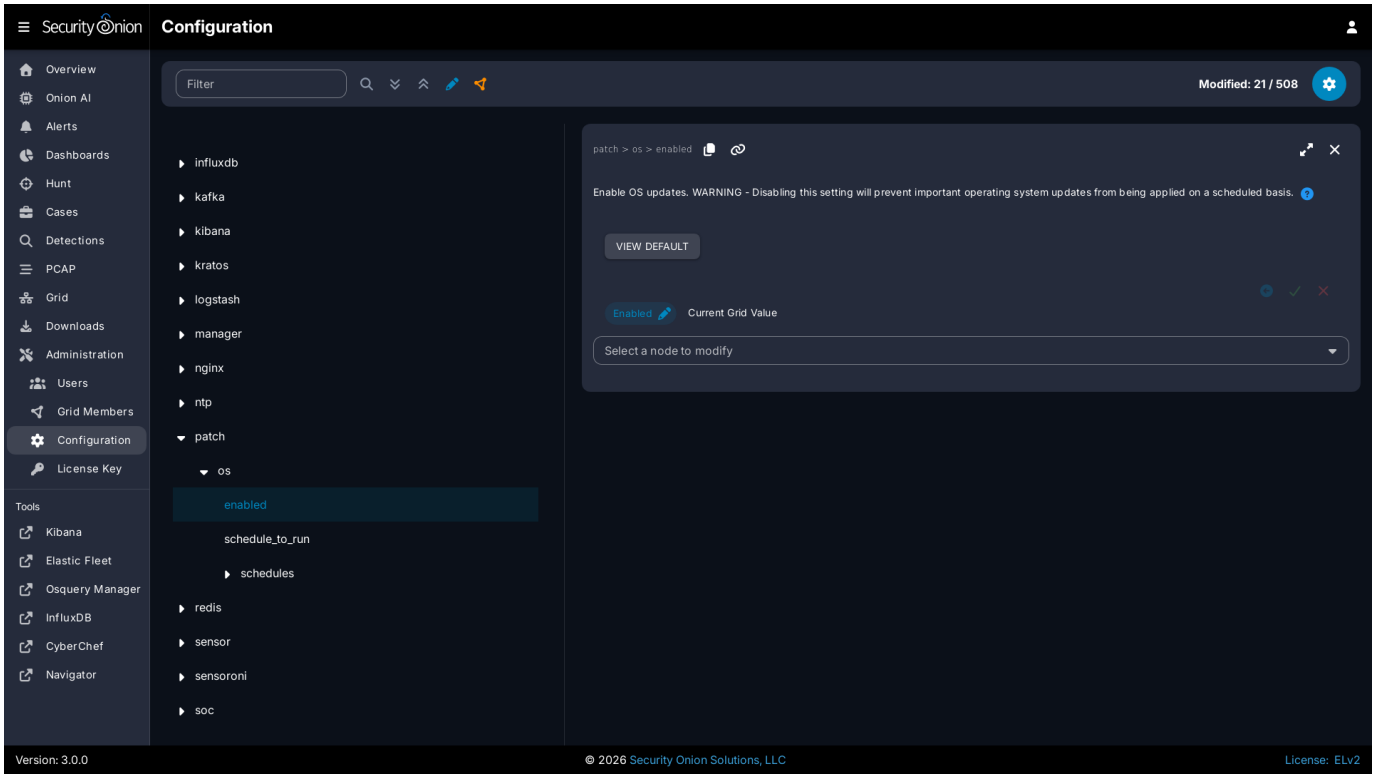

Holding OS Updates

If you would like to prevent certain packages from being upgraded automatically (commonly referred to as "locking" or "holding" packages), then you can do that via the `hold` setting. The most frequent use case is holding the kernel to prevent frequent system reboots. To hold a package, add the package name to the `hold` list in Administration --> Configuration --> versionlock --> hold.

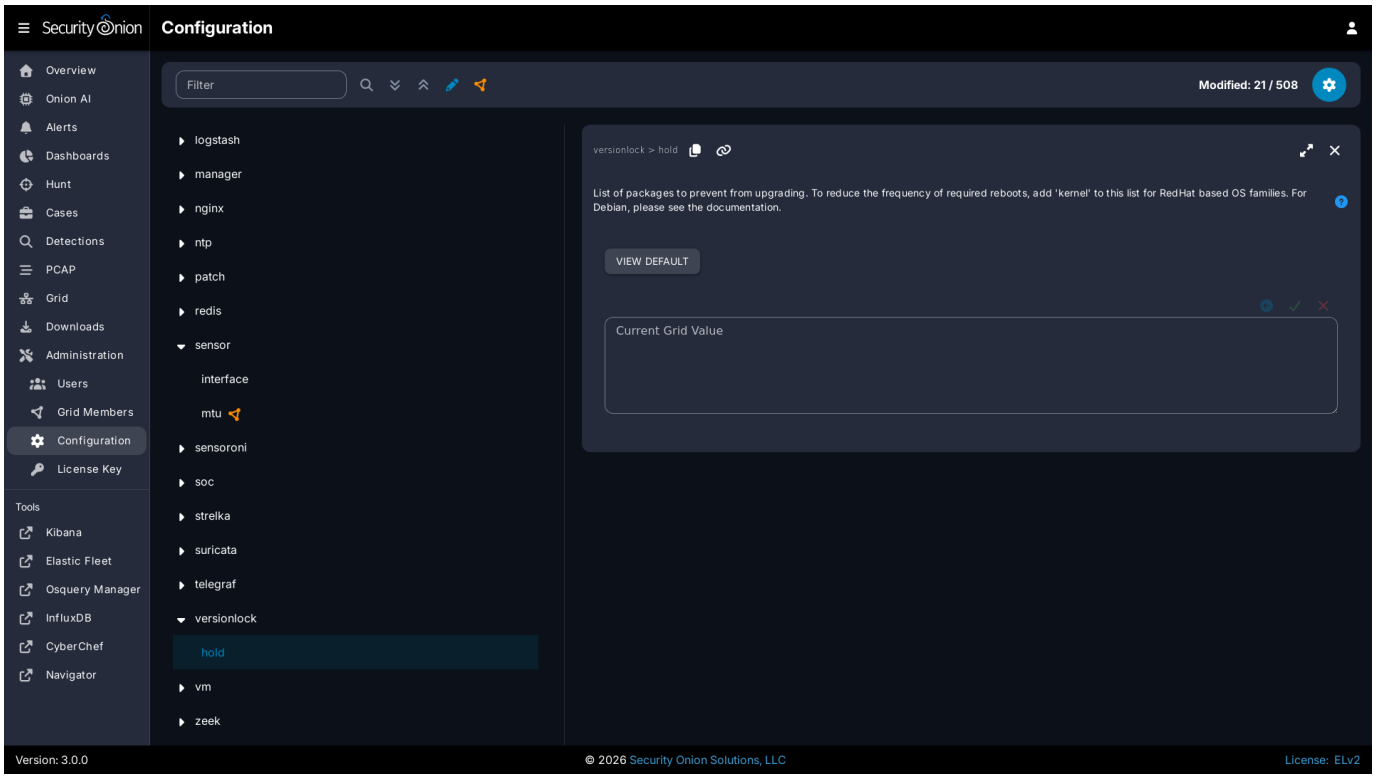

14.2.7 Local Configurations

`soup` will check for local configurations in `/opt/so/saltstack/local/` that may cause issues and flag them with the message `Potentially breaking changes found in the following files`. Please examine the output of `soup` and review any local configurations for possible issues.

14.2.8 Log

If `soup` displays any errors, you can check `/root/soup.log` for additional clues.

14.2.9 Airgap

To update an Airgap deployment, you'll need to get the latest ISO image to the airgapped manager and then run `soup` which will ask where to find it:

- burn the latest ISO image to a DVD and insert it in the DVD drive of the manager (example: `/dev/cdrom`)
- flash the ISO image to a USB drive and connect that USB drive to the manager (example: `/dev/sdb`)
- simply copy the ISO file itself to the manager (example: `/home/YourUser/securityonion-3.X.Y-YYYYMMDD.iso`)

Instead of waiting for soup to prompt for the location, you can also specify the path on the command line using the `-f` option. For example (change this to reflect the actual path to the ISO file or disk device containing the ISO media):

```
sudo soup -y -f /home/YourUser/securityonion-3.X.Y-YYYYMMDD.iso
```

14.2.10 Elastic

If `soup` updated to a new version of the Elastic stack, then you'll want to go to Elastic Fleet and:

- drill into each of your active agent policies, check the Agent Binary Download setting, and adjust if necessary for your deployment
- check for any integrations that need to be upgraded
- check for any agents that need to be upgraded (Grid node agents should automatically upgrade so you should just need to look for any additional endpoint agents that you've deployed)

14.2.11 Kibana

After `soup` completes, if Kibana says `Kibana server is not ready yet` even after waiting a few minutes for it to fully initialize, then take a look at the Diagnostic Logging section of the Kibana page.

If Kibana loads but the dashboards display errors that they didn't before the upgrade, first shift-reload your browser to make sure there are no cache issues. If that doesn't resolve the issue, then you may need to reload the dashboards on your manager:

```
sudo rm /opt/so/state/kibana_*.txt
sudo salt-call state.apply Kibana.so_savedobjects_defaults -l info queue=True
```

14.2.12 Automation

`soup` can be automated as follows:

```
sudo soup -y
```

This will make `soup` proceed unattended, automatically answering `yes` to any prompt. If you have an airgap installation, you can specify the path to the ISO image using the `-f` option as follows:

```
sudo soup -y -f /home/user/securityonion.iso
```

14.2.13 Errors

Data failed to compile

Occasionally, `soup` may output a `Data failed to compile` error that says something like `Rendering SLS failed: Jinja variable 'None' has no attribute`. In most cases, this error corrects itself on the next Salt run.

Pillars and sls files

`soup` will check Salt pillars to make sure they can be rendered. If not, it will output a message like this:

```
There is an issue rendering the manager's pillars. Please correct the issues in the sls files mentioned below before running SOUP again.
```

This usually means that somebody has modified the Salt sls files and introduced a typo.

Downloading images

As `soup` is downloading container images, it may encounter errors if there are Internet connection issues or if the disk runs out of free space. Once you've resolved the underlying condition, you can manually refresh your container images using `so-docker-refresh`.

Docker Registry

If you see errors relating to `so-dockerregistry` (Docker Registry), then please take a look at the following discussions to see if your symptoms match and if their solutions may help you:

https://github.com/Security-Onion-Solutions/securityonion/discussions/12078

https://github.com/Security-Onion-Solutions/securityonion/discussions/12635

Highstate already running

Here are some other errors that you may see when running `soup`:

```
local:
    Data failed to compile:
----------
    Rendering SLS 'base:common' failed: Jinja variable 'list object' has no attribute 'values'
```

and/or

```
There is a problem downloading the so-xyz:3.X.Y image. Details:
gpg: Signature made Thu 18 Feb 2021 02:26:10 PM UTC using RSA key ID FE507013 gpg: BAD signature from "Security Onion Solutions, LLC
<info@securityonionsolutions.com>"
```

If you see these errors, it most likely means that a salt highstate process was already running when `soup` began. You can wait a few minutes and then try `soup` again. Alternatively, you can run `sudo so-checkin` and wait for it to complete before running `soup` again.

14.2.14 Distributed deployments

If you have a distributed deployment with a manager node and separate sensor nodes and/or search nodes, you **only** need to run `soup` on the manager. Once `soup` has completed, other nodes should update themselves at the next Salt highstate (typically within 15 minutes).

> **Warning**
>
> Just because the update completed on the manager does NOT mean the upgrade is complete on other nodes in the grid. Do not manually restart anything until you know that all the search nodes and heavy nodes are updated.
>
> Each minion is on a random 15 minute check-in period and things like network bandwidth can be a factor in how long the actual upgrade takes. If you have a heavy node on a slow link, it is going to take a while to get the containers to it. Depending on what changes happened between the versions, Elasticsearch might not be able to talk to said heavy node until the update is complete.
>
> If it looks like you're missing data after the upgrade, please avoid restarting services and instead make sure at least one search node has completed its upgrade. The best way to do this is to run `sudo so-checkin` from a search node and make sure there are no errors. Typically if it works on one node it will work on the rest. Sensor nodes are less complex and will update as they check in so you can monitor those from the Grid section of SOC.

When you run `soup` on the manager, it does the following:

- Checks to see if it is running on a manager.
- Checks to see if the grid is in Airgap mode. If so, it will then ask for the location of the ISO or mount point.
- Checks to see if we're running the latest version of `soup`. If not, it will put the latest in the correct place and ask you to re-run `soup`.
- Compares the installed version with what is available on github or the ISO image.
- Checks to see if Salt needs to be updated (more on this later).
- Downloads the new Docker images or, if airgap, copies them from the ISO image.
- Stops the Salt master and minion and restarts it in a restricted mode. This mode only allows the manager to connect to it so that we make sure the manager is done before any of the minions are updated.
- Updates Salt if necessary. This will cause the master and minion services to restart but still in restricted mode.
- Makes any changes to pillars that are needed such as adding new settings or renaming values. This varies from release to release.
- If the grid is in Airgap mode, then it copies the latest ET Open rules and yara rules to the manager.
- The new Salt code is put into place on the manager.
- Runs a highstate on the manager which is the actual upgrade where it will use the new Salt code and Docker containers.
- Unlocks the Salt master service and allows minions to connect again.
- Issues a command to all minions to update Salt if necessary. This is important to note as it takes time to to update the Salt minion on all minions. If the minion doesn't respond for whatever reason, it will not be upgraded at this time. This is not an issue because the first thing that gets checked when a minion talks to the master is if Salt needs to be updated and will apply the update if it does.
- Nodes connect back to the manager and actually perform the upgrade to the new version.

14.3 End Of Life

This page lists End Of Life (EOL) dates for older versions of Security Onion and older components.

- Security Onion 2.4 reaches EOL on October 1, 2026 (please migrate to Security Onion 3)
- Security Onion 2.3 reached EOL on April 6, 2024:
  https://blog.securityonion.net/2023/10/6-month-eol-notice-for-security-onion-23.html
- Ubuntu 18.04 reached End of Ubuntu Standard Support in April 2023:
  https://blog.securityonion.net/2023/02/ubuntu-1804-reaches-end-of-ubuntu.html
- TheHive 3 reached EOL on December 31, 2021. TheHive and Cortex were fully removed from Security Onion in Security Onion 2.3.120:
  https://blog.securityonion.net/2022/04/security-onion-23120-now-available.html
- Security Onion 16.04 reached EOL on April 16, 2021:
  https://blog.securityonion.net/2021/04/security-onion-1604-has-reached-end-of.html
- Security Onion 14.04 reached EOL on November 30, 2018:
  https://blog.securityonion.net/2018/06/6-month-eol-notice-for-security-onion.html

15. Accounts

15.1 Accounts Overview

In Security Onion, there are two main types of accounts:

- operating system (OS) accounts
- application accounts used when authenticating to SOC

OS accounts are controlled by standard Linux account utilities. SOC accounts are maintained via the Administration interface. If for some reason you can't log into SOC, you can use so-user from the command line.

15.2 Adding Accounts

15.2.1 OS

If you need to add a new OS user account, you can use the `adduser` command. For example, to add a new account called `tom`:

```
sudo adduser tom
```

> **Tip**
>
> We recommend creating OS usernames in lower case for consistency.

For more information about adding OS user accounts, please see the adduser manual by typing `man adduser`.

15.2.2 SOC

If you need to add a new account to SOC, navigate to the Administration interface, and then click `Users`.

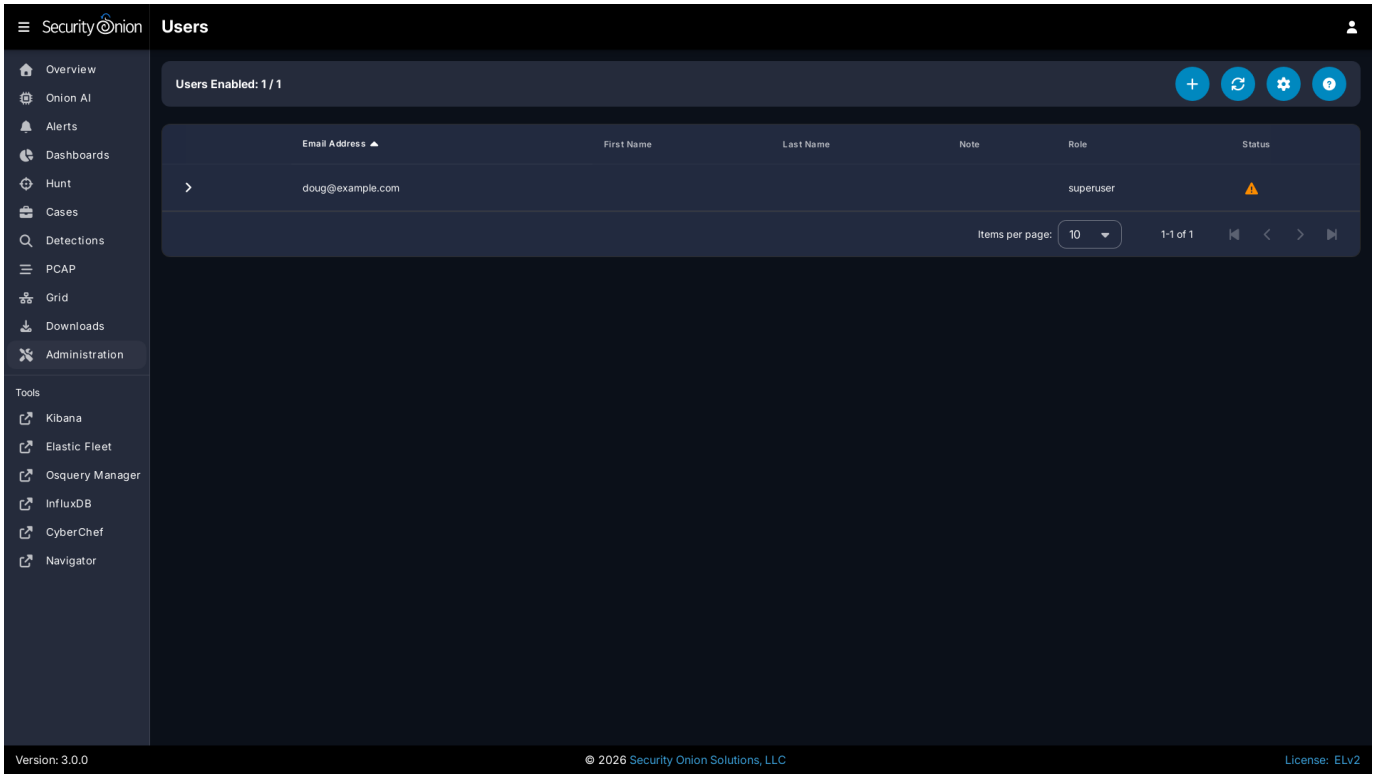

Click the `+` button, fill out the necessary information, and then click the `ADD` button.

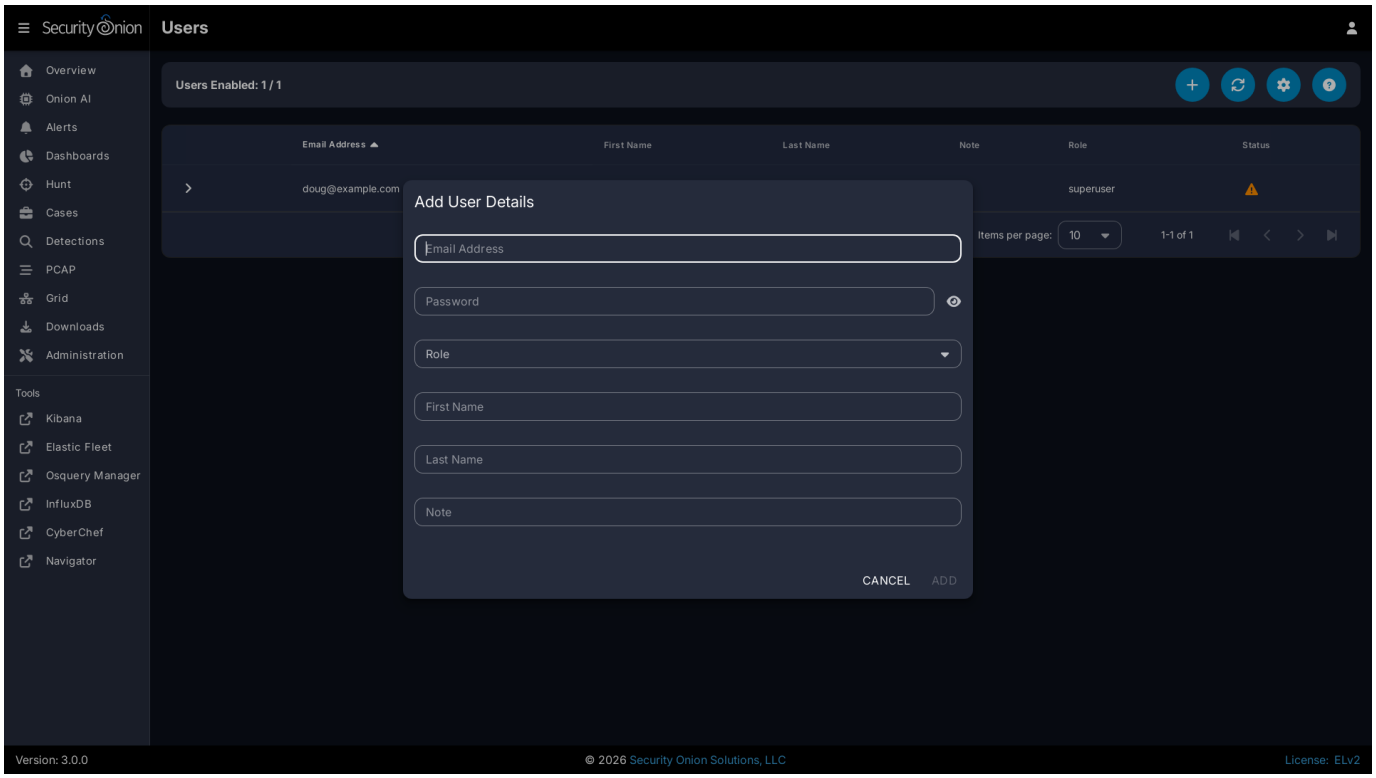

Tip

We recommend specifying email addresses in lower case for consistency.

For more information about the Users page, please see the Administration section.

15.3 Disabling Accounts

15.3.1 OS

If you need to disable an OS user account, you can expire the account using `usermod --expiredate 1`. For example, to disable the account for user `tom`:

```
sudo usermod --expiredate 1 tom
```

For more information, please see passwd manual by typing `man passwd` and the usermod manual by typing `man usermod`.

15.3.2 SOC

If you need to disable an account in SOC, you can go to Administration --> Users, expand the user account, and click the `LOCK USER` button.

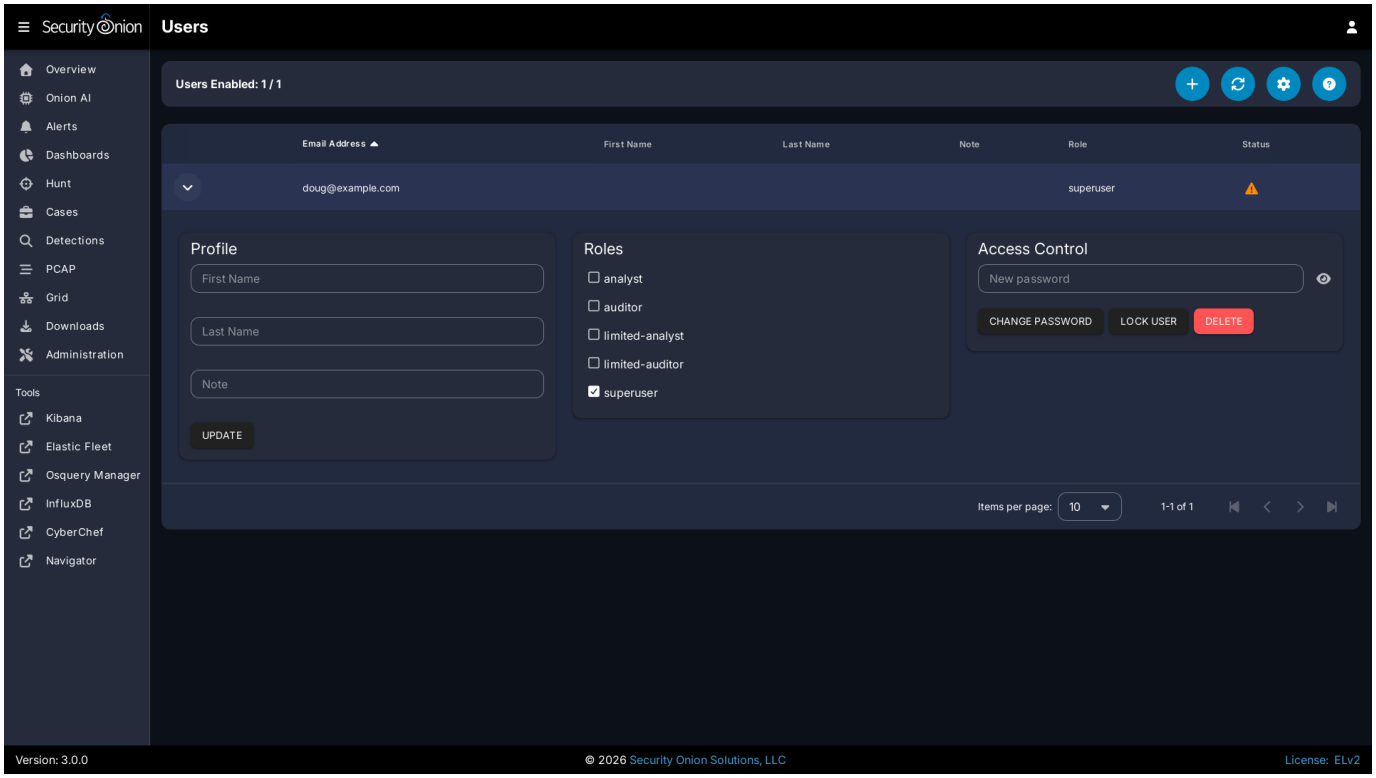

After disabling a user account, it will be shown with a disabled icon in the Status column.

For more information about the Users page, please see the Administration section.

15.4 Listing Accounts

15.4.1 OS

Operating System (OS) user accounts are stored in `/etc/passwd`. You can get a list of all OS accounts using the following command:

```
cut -d: -f1 /etc/passwd
```

If you want a list of user accounts (not service accounts), then you can filter `/etc/passwd` for accounts with a UID greater than 999 like this:

```
cat /etc/passwd | awk -F: '$3 > 999 {print ;}' | cut -d: -f1
```

15.4.2 SOC

You can get a list of users in SOC by navigating to the Administration interface and then clicking `Users`:

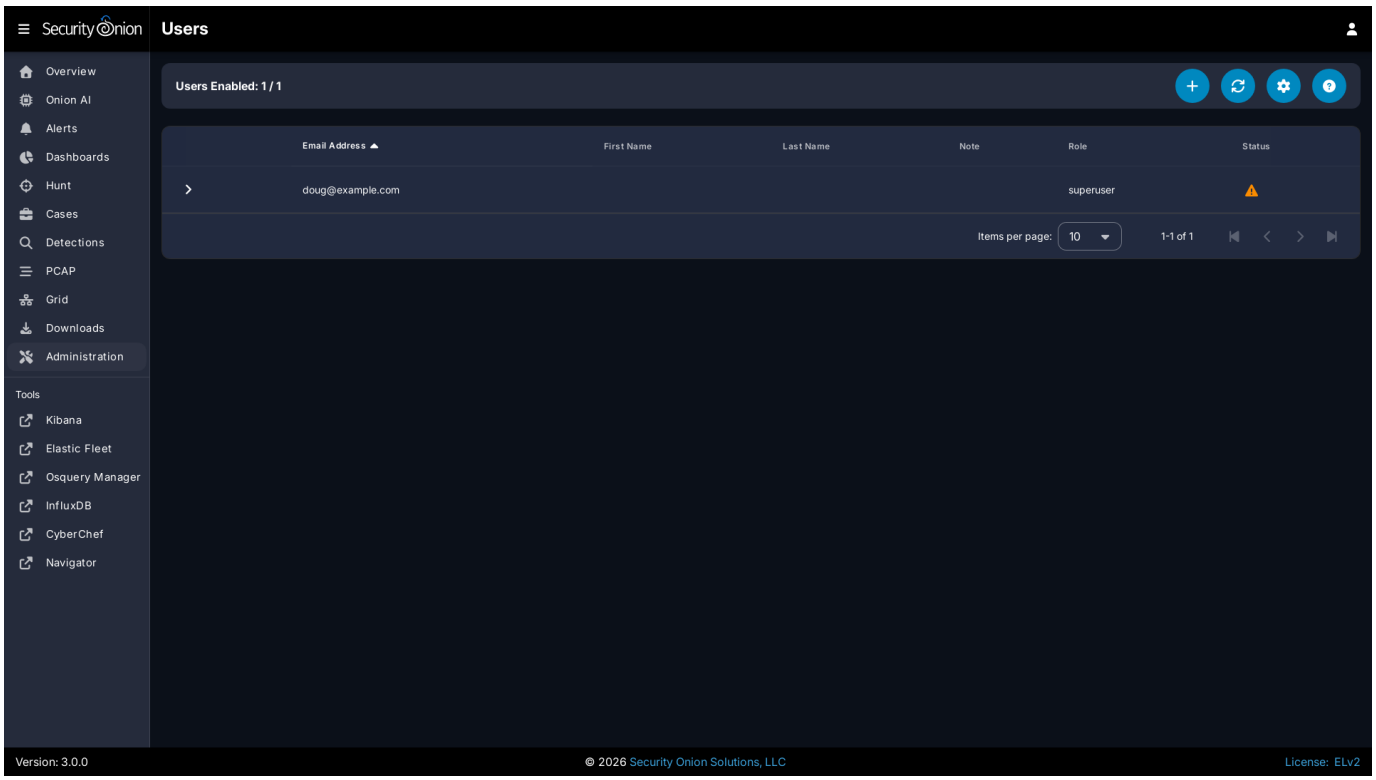

To see detail on an individual user account, click the button on the left side of the row to expand the user account:

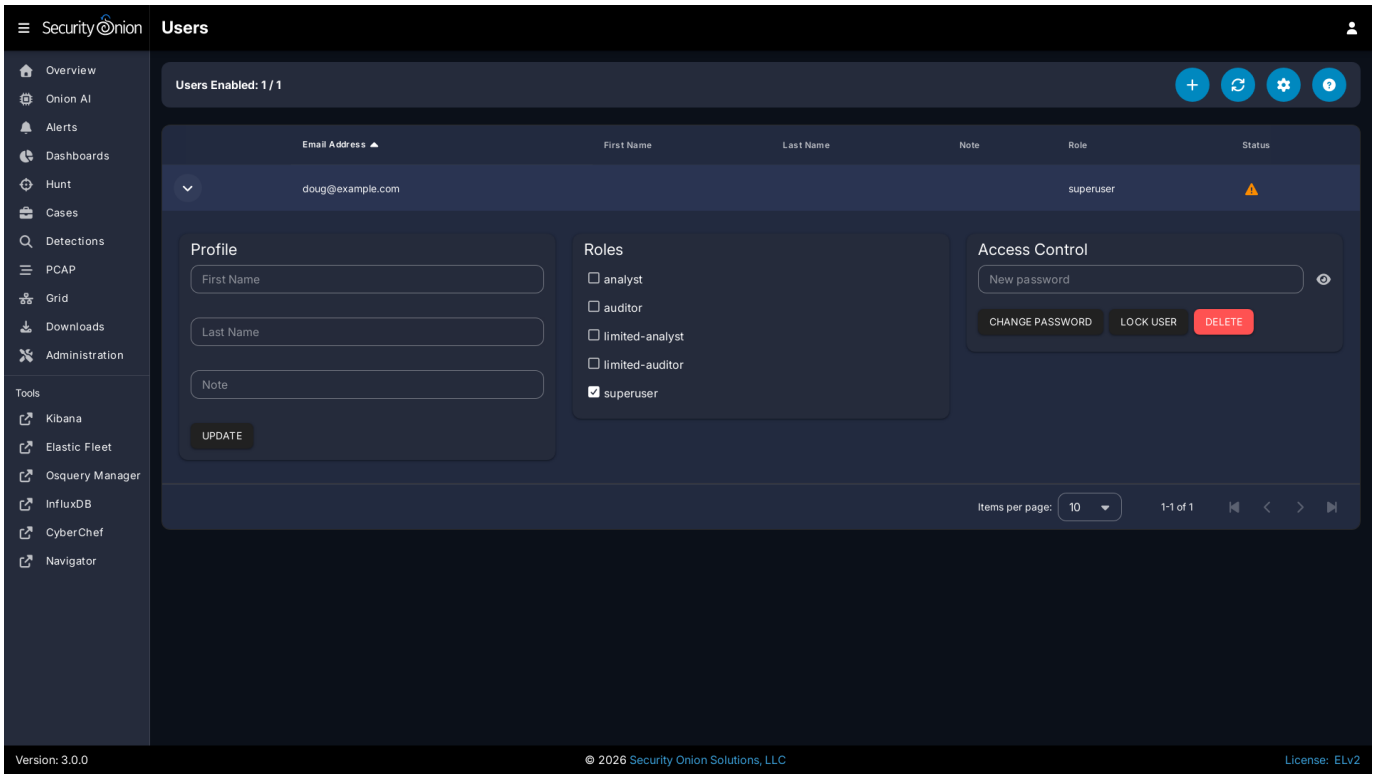

For more information about the Users page, please see the Administration section.

15.5 Passwords

15.5.1 OS user account

When you first install Security Onion, you create a standard OS user account for yourself. If you need to change your OS user password, you can use the `passwd` command:

```
passwd
```

15.5.2 OS root account

Your default user account should have sudo permissions. Command-line utilities that require administrative access can be prefixed with `sudo`. For example, the so-status command requires administrative access so you can run it with `sudo` as follows:

```
sudo so-status
```

15.5.3 Password Logins to SOC

Log into SOC using the username and password you created in the Setup wizard or the username and password provided by your administrator.

You can change your password in SOC by clicking the user icon in the upper-right corner, clicking `Settings`, and then going to the `Security` tab. Please note that, due to technical limitations, if you change your SOC password here it will not update your password in InfluxDB. However, resetting your password via Administration will reset your InfluxDB password.

If you've forgotten your SOC password, an administrator can change it using the Administration interface.

15.5.4 Passwordless Logins to SOC

Once logged in to SOC using the username and password method, users can optionally enable passwordless logins, provided the setting is enabled. The login screen will show a separate section for passwordless logins, if it is enabled. Note that it is enabled by default on new installations.

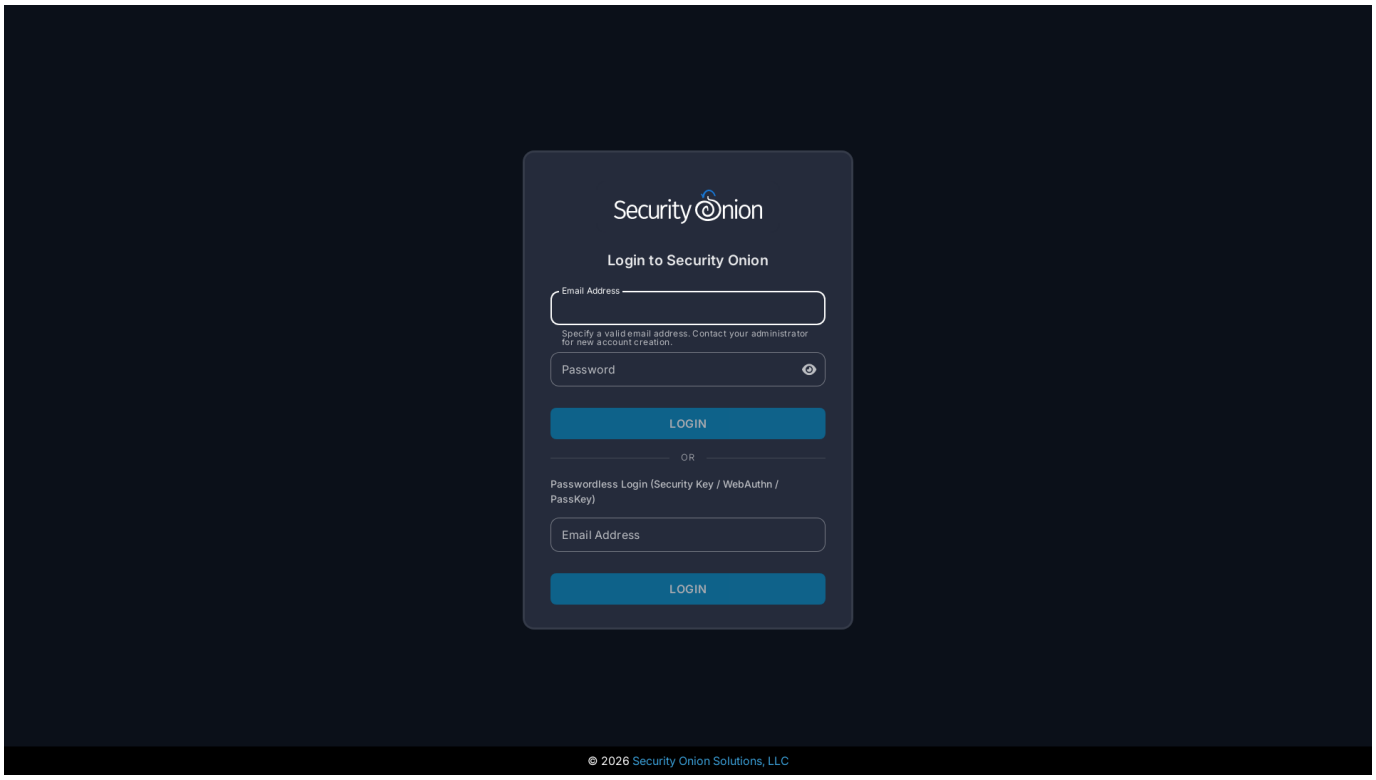

Activate passwordless login for your SOC user by clicking the user icon in the upper-right corner, clicking `Settings`, and then going to the `Security` tab. Scroll down to the `Security Keys` section and follow the provided instructions.

Similarly, disable passwordless logins by returning to the `Security` tab and clicking the delete icon next to any previously-created Security Key.

Note

While it is possible to use TOTP MFA as a second authentication factor in combination with passwordless logins, it is not possible to use a second security key as a second authentication factor with passwordless logins.

Important

The webauthn specification requires that the web server be accessed via a hostname. Therefore, IP addresses cannot be used to access SOC when utilizing webauthn. Also, the server's TLS certificate must not have any errors. Consequently, self-signed certificates will only be permitted provided the certificate authority (CA) has also been imported into analyst's browsers and/or operating systems, and marked as trusted.

15.6 MFA

You can enable Multi-Factor Authentication (MFA) to further protect your account. This can be enabled in SOC by clicking the user icon in the upper-right corner, clicking `Settings`, and then going to the `Security` tab.

15.6.1 TOTP

Time-based One-Time Passwords (TOTP) can be activated on a user account. TOTP requires the use of an authenticator app. Currently only Google Authenticator has been tested, however other authenticator apps that implement the time-based one-time password (TOTP) specification could also work.

To require all users setup TOTP upon login, enable the `Require TOTP` configuration setting, located on the Configuration screen: `SOC > config > server > Require TOTP`.

Warning

Please note that TOTP requires that both the Security Onion manager and the device supplying the TOTP code to have their system time set correctly. Otherwise, the TOTP code may be seen as invalid and rejected.

Note

If you lose access to your authenticator app, an administrator can reset your password using the Administration interface which will also remove the TOTP from your account.

Customizing the MFA Name

If you utilize multiple Security Onion environments, such as one for testing and one for production, and both are setup with TOTP MFA, SOC users may have trouble distinguishing them in their authenticator app. There are two options for handling this situation:

1. Most authenticator apps allow the user to edit or rename the entry. For example, in Google Authenticator on Android, swiping right on the entry provides an Edit screen. This is useful if it's just affecting one or two users.
2. Edit the TOTP issuer via the SOC Configuration screen, specifically the `kratos > config > selfservice > methods > totp > config > issuer` setting. This should be done prior to enabling TOTP since it will not help users that already setup TOTP.

15.6.2 WebAuthn Security Keys

WebAuthn allows the use of built-in mobile device biometric sensors, USB security devices, and other PKI-based security devices to authenticate users during the login process.

If the Security Onion installation has been configured to use security keys for MFA instead of passwordless logins then you can add one or more security keys to your account as a second authentication factor.

Note

If you lose access to your security key device, an administrator can reset your password using the Administration interface which will also remove the security keys from your account.

Important

The webauthn specification requires that the web server be accessed via a hostname. Therefore, IP addresses cannot be used to access SOC when utilizing webauthn. Also, the server's TLS certificate must not have any errors. Consequently, self-signed certificates will only be permitted provided the certificate authority (CA) has also been imported into analyst's browsers and/or operating systems, and marked as trusted.

15.7 Role-Based Access Control (RBAC)

The ability to restrict or grant specific privileges to a subset of users is covered by role-based access control, or "RBAC" for short. RBAC is an authorization technique in which users are assigned one of a small set of roles, and then the roles are associated to many low-level privileges. This provides the ability to build software with fine-grained access control, but without the need to maintain complex associations of users to large numbers of privileges. Users are traditionally assigned a single role, one which correlates closely with their role in the organization. However, it's possible to assign a user to multiple roles, if necessary.

RBAC in Security Onion covers both Security Onion privileges and Elastic stack privileges. Security Onion privileges are only involved with functionality specifically provided by the components developed by Security Onion, while Elastic stack privileges are only involved with the Elasticsearch, Kibana, and related Elastic stack. For example, Security Onion will check if a user has permission to create a PCAP request, while Elastic will check if the same user has permission to view a particular index or document stored in Elasticsearch.

15.7.1 Default Roles

Security Onion ships with the following user roles: `superuser`, `analyst`, `limited-analyst`, `auditor`, and `limited-auditor`.

See the table below which explains the specific Security Onion privileges granted to each role.

	superuser	analyst	limited-analyst	auditor	limited-auditor
View alerts	X	X	X	X	X
Acknowledge alerts	X	X	X		
Escalate alerts and events	X	X	X		
View detections	X	X	X	X	X
Modify (and Delete) detections	X	X			
View events in Hunt	X	X	X	X	X
View own PCAP jobs	X	X	X	O	O
View all PCAP jobs	X	X		X	
Pivot to PCAP job from event	X	X	X		
Request arbitrary PCAP jobs	X	X			
Delete own PCAP job	X	X	X	O	O
Delete any PCAP job	X	X			
View all nodes in Grid	X	X	X	X	X
View all users	X	X		X	
View all users' roles	X	X		X	
View own user	X	X	X	X	X
View own user roles	X	X	X	X	X
Change own password	X	X	X	X	X
Add, update, and reset SOC users	X				
Modify and synchronize Grid config	X				
Manage Grid membership of nodes	X				
	X				

	superuser	analyst	limited-analyst	auditor	limited-auditor
Initiate node actions (Ex: reboot)					
Manage API clients	X				
View and list existing API clients	X		X		
Manage Active Queries	X				
View Playbooks	X	X	X	X	X
Chat with Onion AI	X	X			
Delete Own Onion AI Sessions	X	X			
View Own Onion AI History	X	X			
View All Users' Onion AI History	X				

Note

Both `auditor` and `limited-auditor` roles can interact with previously created PCAPs if they were created before a user was converted to that role (e.g. user was downgraded from `analyst` to `auditor`). This is denoted by **O** in the above table.

Note

A system role called `agent` is used by the Security Onion agent that runs on each node of the Security Onion Grid. This special role is given the *jobs/process*, *nodes/read*, and *nodes/write* permissions (defined at the bottom of this page). Avoid creating custom roles that share the same name as Security Onion-provided roles.

15.7.2 Superusers

After a new installation of Security Onion completes, a single administrator user will be created and assigned the `superuser` role. Additional users can also be assigned to the `superuser` role, if desired.

15.7.3 Adding a User With a Specific Role

In the Administration interface, navigate to the Users screen and click the + icon to add a new user. In the popup dialog you can check the roles you would like to assign to the new user.

15.7.4 Modifying User Roles

In the Administration interface, navigate to the Users screen and click the > icon to the left of the email address needing adjusting. Check or uncheck the desired roles.

15.7.5 Default Role Assignment

When a user is created, they can optionally be automatically assigned to a specific role. To enable this automatic assignment, locate the `defaultRole` configuration setting and specify the desired default role name. This automatic assignment will occur when the user logs into SOC for the first time.

> **Warning**
>
> If an administrator removes all role assignments from a user and the user logs back in that user will again be automatically assigned the default role. Always lock inactive users instead of removing role assignments from a user.

15.7.6 Creating Custom Roles

> **Warning**
>
> The creation of custom RBAC roles is an advanced feature that is recommended only for experienced administrators.

These steps will guide you through an example where we wish to introduce a new role called `eastcoast-analyst`, which will inherit the same Security Onion permissions as a limited-analyst, but will be restricted to only view a subset of documents in the Elastic stack. We base this role on the `limited-analyst` instead of the `analyst` role so that the user does not have the ability to create arbitrary PCAPs on any sensor.

1. For the Security Onion role: Follow the instructions in the next section entitled "Defining Security Onion Roles" to create a new role named `eastcoast-analyst`.
2. For the Elastic stack role: Create a new json role file named `eastcoast-analyst.json` under `/opt/so/saltstack/local/salt/elasticsearch/roles`. In this example we will define the new role that only allows access to documents from sensors on the east coast of the United States. Specifically, the role will include a query filter that limits search results to only include documents originating from sensors having a name prefixed with `nyc` (New York City) or `atl` (Atlanta).

`eastcoast-analyst.json`:

```
{
  "cluster": [
    "cancel_task",
    "create_snapshot",
    "monitor",
    "monitor_data_frame_transforms",
    "monitor_ml",
    "monitor_rollup",
    "monitor_snapshot",
    "monitor_text_structure",
    "monitor_transform",
    "monitor_watcher",
    "read_ccr",
    "read_ilm",
    "read_pipeline",
    "read_slm"
  ],
  "indices": [
    {
      "names": [
        "so-*"
      ],
      "privileges": [
        "index",
        "maintenance",
        "monitor",
        "read",
        "read_cross_cluster",
        "view_index_metadata"
      ],
      "query": "{ \"bool\": { \"should\": [ { \"prefix\": { \"observer.name\": \"nyc\" }}, { \"prefix\": { \"observer.name\": \"atl\" }} ]}}"
    }
  ],
  "applications": [
    {
      "application": "Kibana-.Kibana",
      "privileges": [
        "feature_discover.all",
        "feature_dashboard.all",
        "feature_canvas.all",
```

```
        "feature_maps.all",
        "feature_ml.all",
        "feature_logs.read",
        "feature_visualize.all",
        "feature_infrastructure.read",
        "feature_apm.read",
        "feature_uptime.read",
        "feature_siem.read",
        "feature_dev_tools.read",
        "feature_advancedSettings.read",
        "feature_indexPatterns.read",
        "feature_savedObjectsManagement.read",
        "feature_savedObjectsTagging.read",
        "feature_fleet.all",
        "feature_actions.read",
        "feature_stackAlerts.read"
      ],
      "resources": [
        "*"
      ]
    }
  ],
  "run_as": []
}
```

Note

Elasticsearch requires that a subscription be purchased in order to use field or document-level security, as referenced in the above example.

Note

The format of the json in this file must match the request body outlined in the Elastic docs here: https://www.elastic.co/guide/en/elasticsearch/reference/current/security-api-put-role.html#security-api-put-role-request-body.

The available cluster and indices permissions are explained in the Elastic docs here: https://www.elastic.co/guide/en/elasticsearch/reference/current/security-privileges.html.

The available Kibana permissions can be obtained by running the following command on the manager node:

```
sudo so-elasticsearch-query _security/privilege/kibana-.Kibana | jq '. | map_values(keys)'
```

1. Run so-checkin from the manager:

```
sudo so-checkin
```

15.7.7 Defining Security Onion Roles

There are two ways to define a custom Security Onion role:

- Building it from scratch using the built-in permissions and default roles available as outlined later in this document, or
- Inheriting the permissions of another role, and optionally adding more permissions to the new custom role.

Note

The `custom_roles` file contains further instructions on modifying roles that are not within the scope of this documentation.

The common syntax for either method of defining a role is as such:

```
<existing role or permission>:<new role>
```

1. Creating the role for the above east coast analyst using the first method, building the custom role from scratch, would be written like so:
2. `case-admin:eastcoast-analyst`
3. `event-admin:eastcoast-analyst`
4. `node-monitor:eastcoast-analyst`
5. `user-monitor:eastcoast-analyst`
6. `job-user:eastcoast-analyst`
7. Alternatively, the `eastcoast-analyst` role could be created by inheriting the permissions of the analyst role:
8. `limited-analyst:eastcoast-analyst`

SECURITY ONION PRIVILEGES AND DEFAULT ROLES

The available low-level Security Onion privileges are listed in the table below:

cases/read	Read all case-related information for all cases
cases/write	Create and update cases, and escalate events to cases
clients/read	List and view existing API clients. Client secrets are inaccessible.
clients/write	Create and update API clients, and regenerate secrets
clients/delete	Delete API clients
config/read	Read system configuration parameters
config/write	Update and in some cases duplicate system configuration parameters
detections/read	Read all detection related details
detections/write	Create and update detections and overrides
events/read	Read from Elasticsearch
events/read	Read from Elasticsearch
events/write	Write to Elasticsearch
events/ack	Acknowledge alerts
Grid/read	Read information about the grid and its node memberships
Grid/write	Accept and reject Grid memberships from new and existing nodes
jobs/read	View all PCAP jobs
jobs/pivot	Pivot to PCAP job from event
jobs/write	Request arbitrary PCAP jobs
jobs/delete	Delete any PCAP job
jobs/process	Update, read, and attach packets to all pending PCAP jobs †
nodes/read	View all nodes in Grid
nodes/write	Update node information †
roles/read	View all users' roles
roles/write	Change any user's role
queries/delete	Cancel active queries
queries/read	View active queries
users/read	View all users
users/write	Change any user's password
users/delete	Delete any user
playbooks/read	View all playbooks
playbooks/write	Currently unused
playbooks/delete	Currently unused
assistant/read_authored	View own Onion AI conversation history
assistant/write_authored	Chat with Onion AI
assistant/delete_authored	Delete own Onion AI conversation history

assistant/read_shared	View shared Onion AI conversation history
assistant/read_all	View all Onion AI conversation history
assistant/write_all	Currently unused
assistant/delete_all	Currently unused

These discrete privileges are then collected into privilege groups as defined below:

case-admin	*cases/read*, *cases/write*
case-monitor	*cases/read*
client-admin	*clients/read*, *clients/write*, *clients/delete*
client-monitor	*clients/read*
config-admin	*config/read*, *config/write*
config-monitor	*config/read*
detections-admin	*detections/read*, *detections/write*
detections-monitor	*detections/read*
event-admin	*events/read*, *events/write*, *events/ack*
event-monitor	*events/read*
Grid-admin	*Grid/read*, *Grid/write*
Grid-monitor	*Grid/read*
job-admin	*jobs/read*, *jobs/pivot*, *jobs/write*, *jobs/delete*
job-monitor	*jobs/read*
job-user	*jobs/pivot*
job-processor	*jobs/process* †
node-admin	*nodes/read*, *nodes/write*
node-monitor	*nodes/read*
query-admin	*queries/read*, *queries/delete*
query-monitor	*queries/read*
user-admin	*roles/read*, *roles/write*, *users/read*, *users/write*, *users/delete*
user-monitor	*roles/read*, *users/read*
playbook-monitor	*playbooks/read*
playbook-admin	*playbooks/read*, *playbooks/write*, *playbooks/delete*
assistant-user	*assistant/read_authored*, *assistant/write_authored*, *assistant/delete_authored*, *assistant/read_shared*
assistant-admin	*assistant/read_authored*, *assistant/write_authored*, *assistant/delete_authored*, *assistant/read_shared*, *assistant/read_all*, *assistant/write_all*, *assistant/delete_all*
assistant-monitor	*assistant/read_authored*, *assistant/read_shared*, *assistant/read_all*

† intended for use by Sensoroni agents only

15.8 Kratos

SOC authentication is handled by Kratos. You can read more about Kratos at https://github.com/ory/kratos.

15.8.1 Configuration

You can configure Kratos by going to Administration --> Configuration --> kratos.

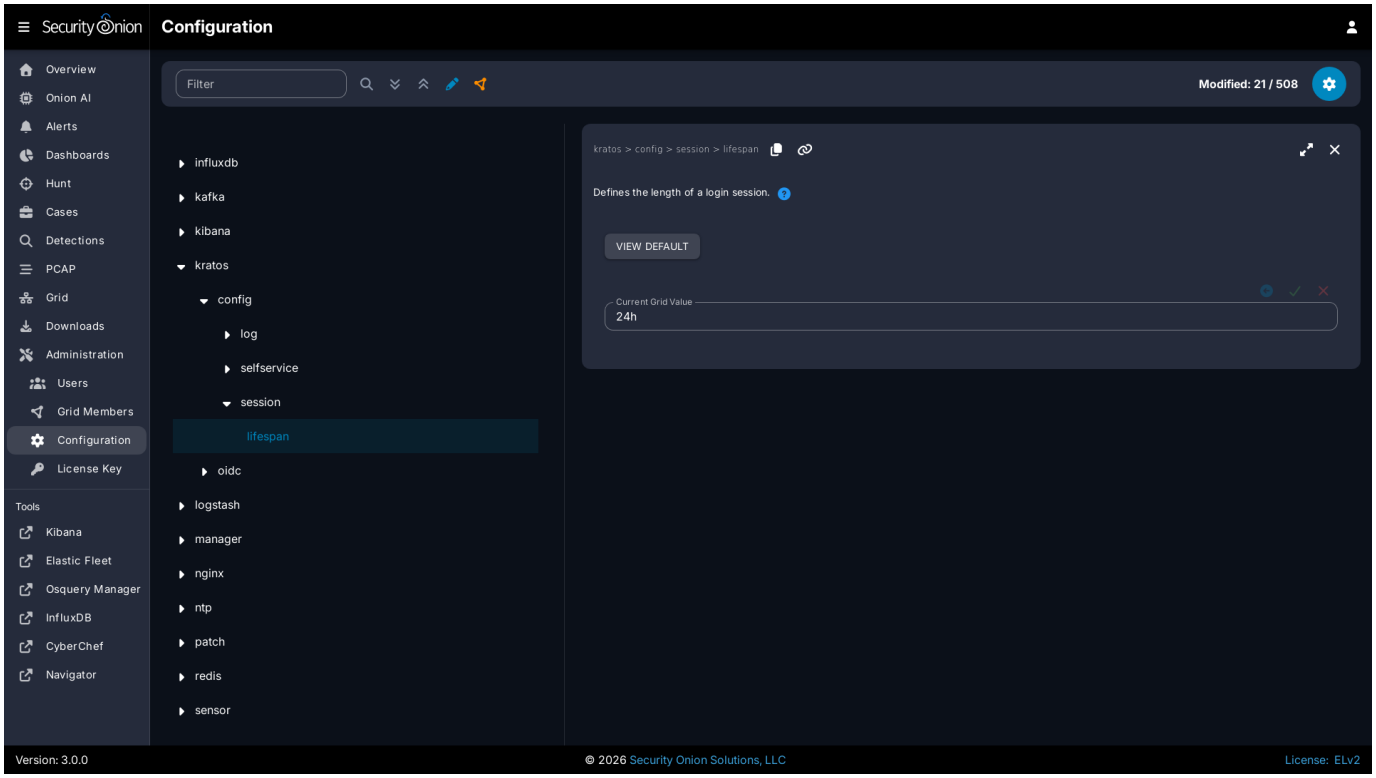

15.8.2 More Information

Note

For more information about Kratos, please see https://github.com/ory/kratos.

16. Services

You can control individual services with the `so-<component>-<verb>` scripts. You can see a list of all of these scripts with the following command:

```
ls /usr/sbin/so-*
```

The following examples are for Zeek, but you could substitute whatever service you're trying to control (Logstash, Elasticsearch, etc.).

Start Zeek:

```
sudo so-zeek-start
```

Stop Zeek:

```
sudo so-zeek-stop
```

Restart Zeek:

```
sudo so-zeek-restart
```

17. Customizing

17.1 Customizing Overview

This section covers how to customize Security Onion for your environment.

17.2 Security Onion Console Customization

You can customize Security Onion Console by going to Administration --> Configuration --> SOC.

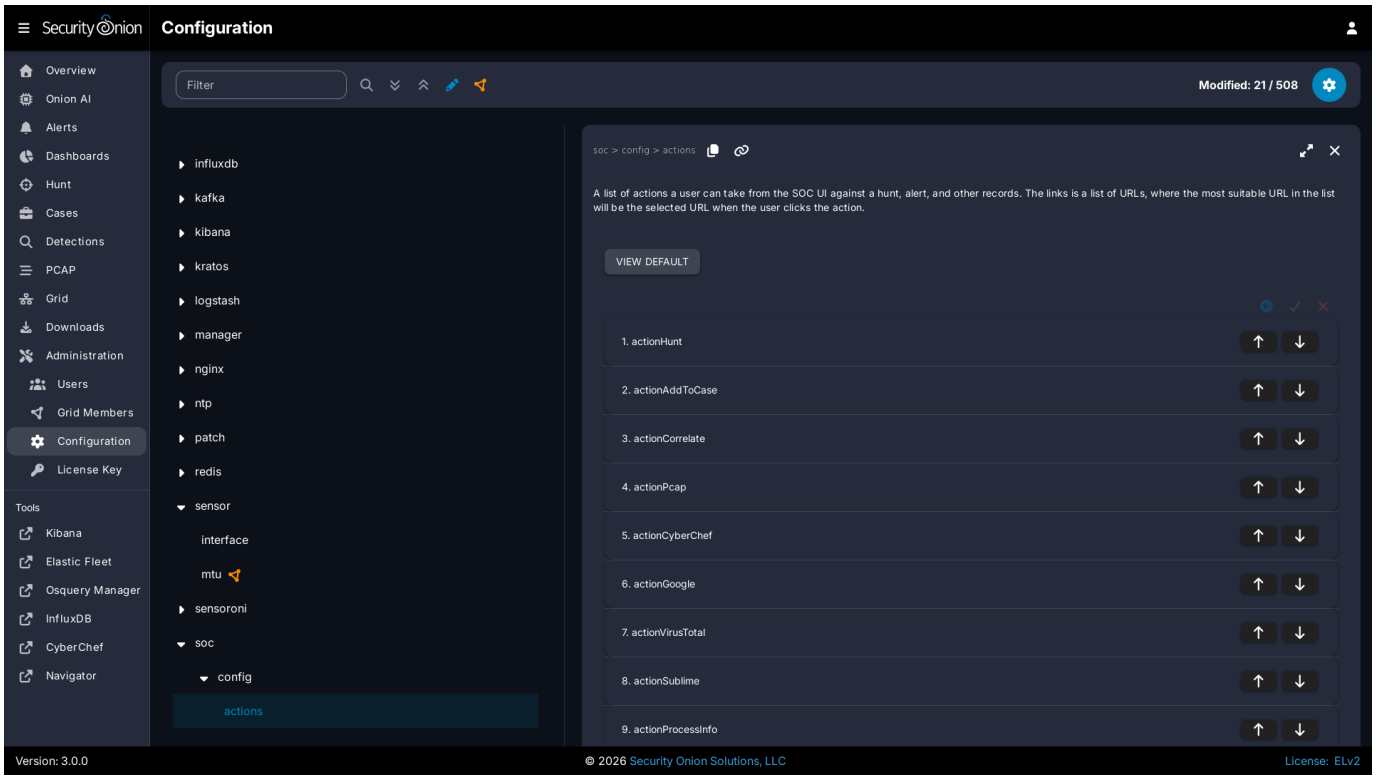

Below are some ways in which you can customize SOC. Once all customizations are complete, you can make the changes take effect by clicking the `Options` bar at the top and then clicking the `SYNCHRONIZE GRID` button.

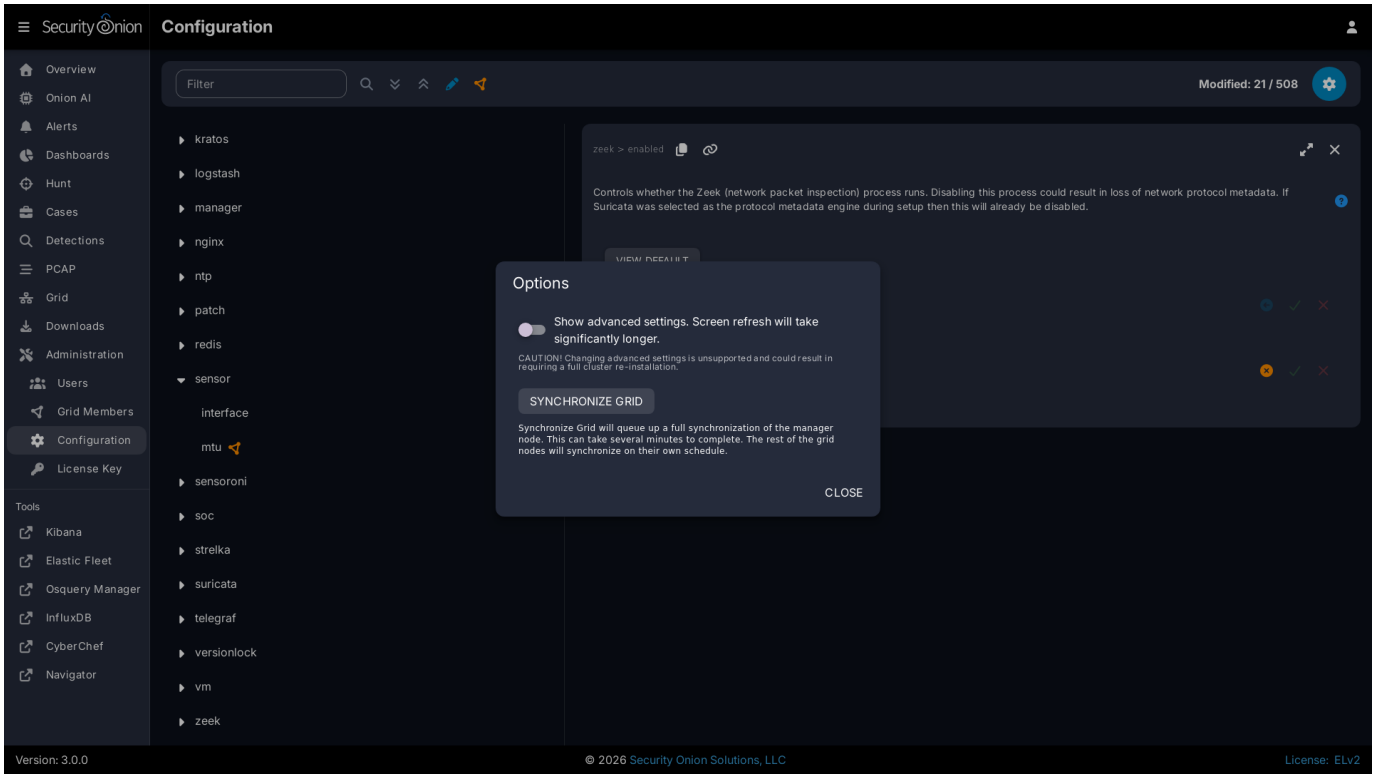

17.2.1 Login Page

You can customize the SOC login page with a login banner by going to Administration --> Configuration --> SOC --> files --> SOC --> Login Banner. The login banner uses the common Markdown (.md) format and you can learn more about that at https://markdownguide.org.

17.2.2 Overview Page

After logging into SOC, you'll start on the main SOC Overview page which can be customized as well. You can customize this by going to Administration --> Configuration --> SOC --> files --> SOC --> Overview Page. This uses Markdown format as mentioned above. You can add images but they must be hosted from another host that is accessible by the user's browser.

17.2.3 Links

You can also customize the links on the left side. To do so, go to Administration --> Configuration --> SOC --> server --> client --> tools.

17.2.4 Reverse DNS

When you are viewing IP addresses in Alerts, Dashboards, or Hunt, you might want to enable automatic reverse DNS lookups to provide more information. You can do so by going to Administration --> Configuration --> SOC --> config --> server --> enableReverseLookup.

17.2.5 Local Lookups

If you don't want to enable reverse DNS lookups for all IP addresses but do have a subset of IP addresses that you would like to resolve to hostnames in SOC, then you can create a CSV file at `/nsm/custom-mappings/ip-descriptions.csv` on your Manager and populate the file with IP addresses and descriptions as follows:

```
IP,Description
```

Elasticsearch will then ingest the CSV and use the contents to populate a new index called `so-ip-mappings`.

When you are viewing IP addresses in Alerts, Dashboards, or Hunt, Security Onion Console will check the local mappings first. If it doesn't find a match, then it will attempt a reverse DNS lookup (if enabled).

If you later need to make changes to your local IP/Descriptions mappings, make the changes in `/nsm/custom-mappings/ip-descriptions.csv` and Elasticsearch will automatically update the `so-ip-mappings` index.

17.2.6 Cases

Cases comes with presets for things like category, severity, TLP, PAP, tags, and status. You can modify these presets by going to Administration --> Configuration --> SOC --> server --> client --> case --> presets.

17.2.7 Session Timeout

The default timeout for user login sessions is 24 hours. This is a fixed timespan and will expire regardless of whether the user is active or idle in SOC. You can configure this by going to Administration --> Configuration --> kratos --> config --> session --> lifespan.

17.2.8 Custom Queries

If you'd like to add your own custom queries to Alerts, Cases, Dashboards, Detections or Hunt, you can go to Administration --> Configuration --> SOC --> config --> server --> client and then select the specific app you'd like to modify.

> **Warning**
>
> When you save your custom queries, SOC saves the entire list of queries (including our default queries included in the product). So if you update to a new version which includes new or updated default queries, you won't see the new or updated default queries since your custom query list is overriding the default.

To see all available fields for your queries, go down to the Events table and then click the arrow to expand a row. It will show all of the individual fields from that particular event.

For example, suppose you want to add GeoIP information like `source.geo.region_iso_code` or `destination.geo.region_iso_code` to Alerts. You would go to Administration --> Configuration --> SOC --> config --> server --> client --> alerts --> queries and insert the following line wherever you want it show up in the query list:

```
{ "name": "Group By Source IP/Port/Geo, Destination IP/Port/Geo, Name", "query": "* | groupby source.ip source.port source.geo.region_iso_code destination.ip destination.port destination.geo.region_iso_code rule.name" },
```

Please note that some events may not have GeoIP information and this query would only show those alerts that do have GeoIP information.

17.2.9 Action Menu

Alerts, Dashboards, and Hunt have an action menu with several default actions. If you'd like to add your own custom HTTP GET or POST actions, you can go to Administration --> Configuration --> SOC --> actions and click the plus sign at the bottom of the list. Then fill out the fields and save the new action.

17.2.10 Escalation

Alerts, Dashboards, and Hunt display logs with a blue triangle that allows you to escalate the event. This defaults to our Cases interface. If for some reason you want to escalate to a different case management system, you can change this setting. You can go to Administration --> Configuration --> SOC --> server --> modules --> cases and specify one of the following values:

- `SOC` - Enables the built-in Case Management, with our Escalation menu (default).
- `elasticcases` - Enables escalation to the Elastic Cases tool. Escalations will always open a new case; there will not be an advanced escalation menu popup. This module will use the same user/pass that SOC uses to talk to Elastic. Note, however, that Elastic cases is actually a Kibana feature, therefore, when this setting is used, SOC will be communicating with the local Kibana service (via its API) for case escalations.

17.3 nginx

nginx is the main web server for Security Onion.

17.3.1 Configuration

You can modify nginx configuration by going to Administration --> Configuration --> nginx.

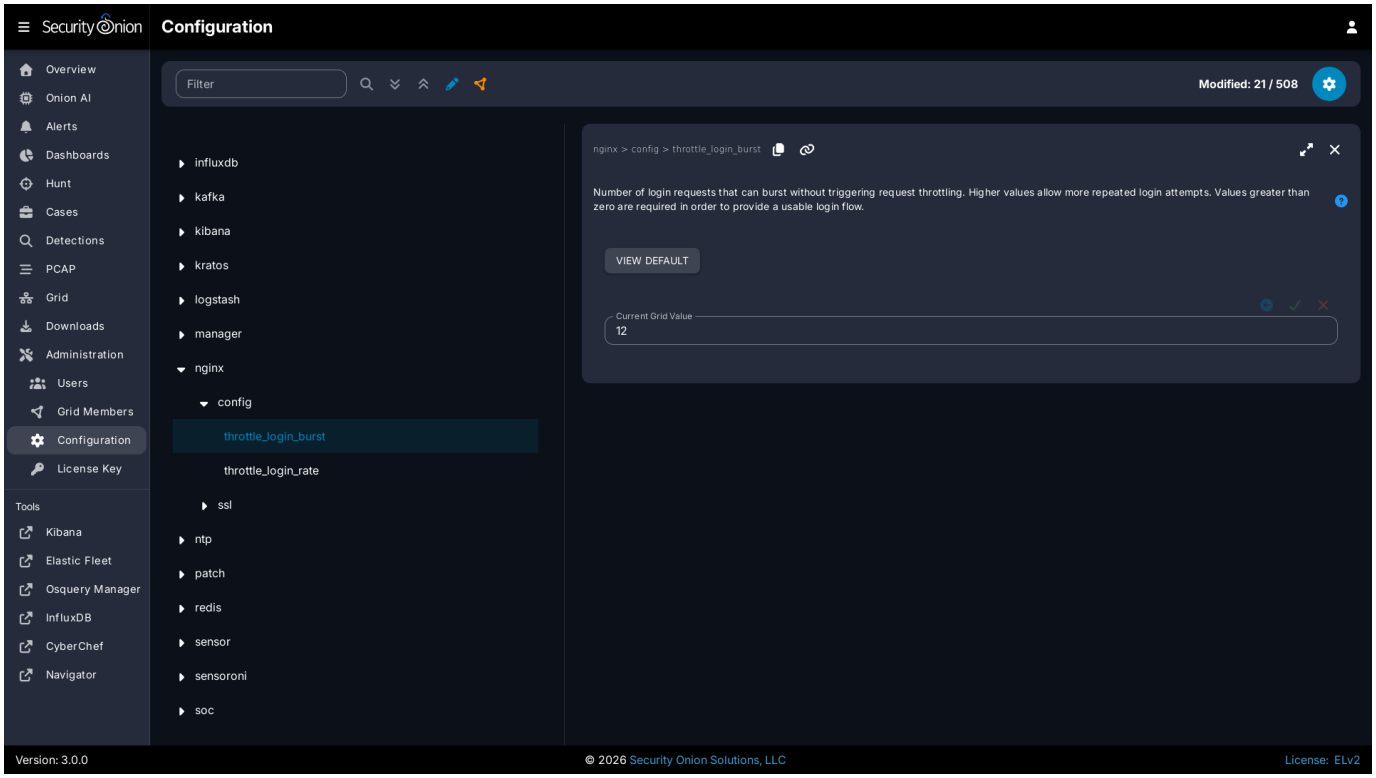

17.3.2 Replacing Default Cert

If you'd like to replace the default certificate with your own cert, then you can do so as shown below.

> **Warning**
>
> Please be very careful when modifying advanced settings like this!

- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- On the left side, go to `nginx`, expand `ssl`, and then select the `Replace Default Cert` setting.
- On the right side, change the setting to `true` and then click the checkmark to save the value.
- On the left side, select the `SSL/TLS Cert File` setting.
- On the right side, paste your new cert file and then click the checkmark to save it.
- On the left side, select the `SSL/TLS Key File` setting.
- On the right side, paste your new key file and then click the checkmark to save it.

17.3.3 More Information

> **Note**
>
> For more information about nginx, please see https://nginx.org/.

17.4 Proxy

Setup will ask if you want to connect through a proxy server and, if so, it will automatically configure the system for you.

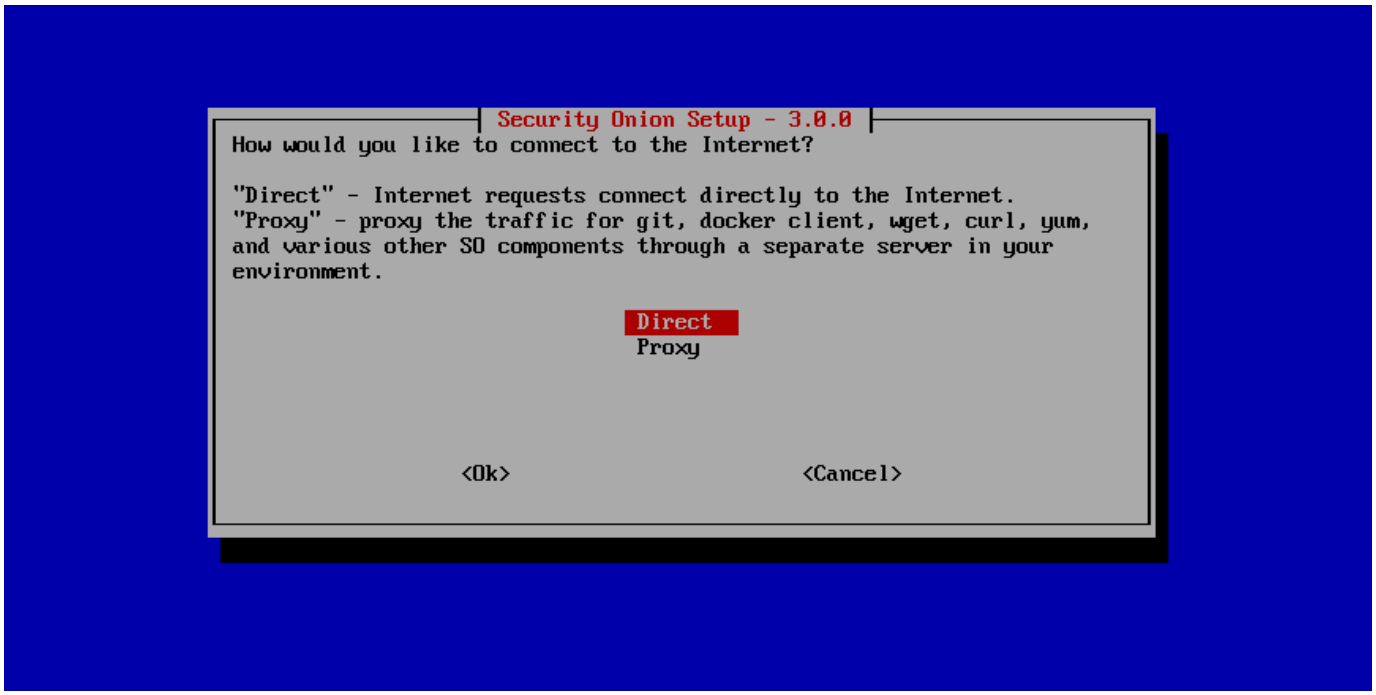

If you have problems installing via your proxy server, you may want to consider the Airgap option as everything will install via the ISO image.

17.4.1 Configuration

If you need to make changes after Setup, please see the proxy settings in Administration --> Configuration --> manager.

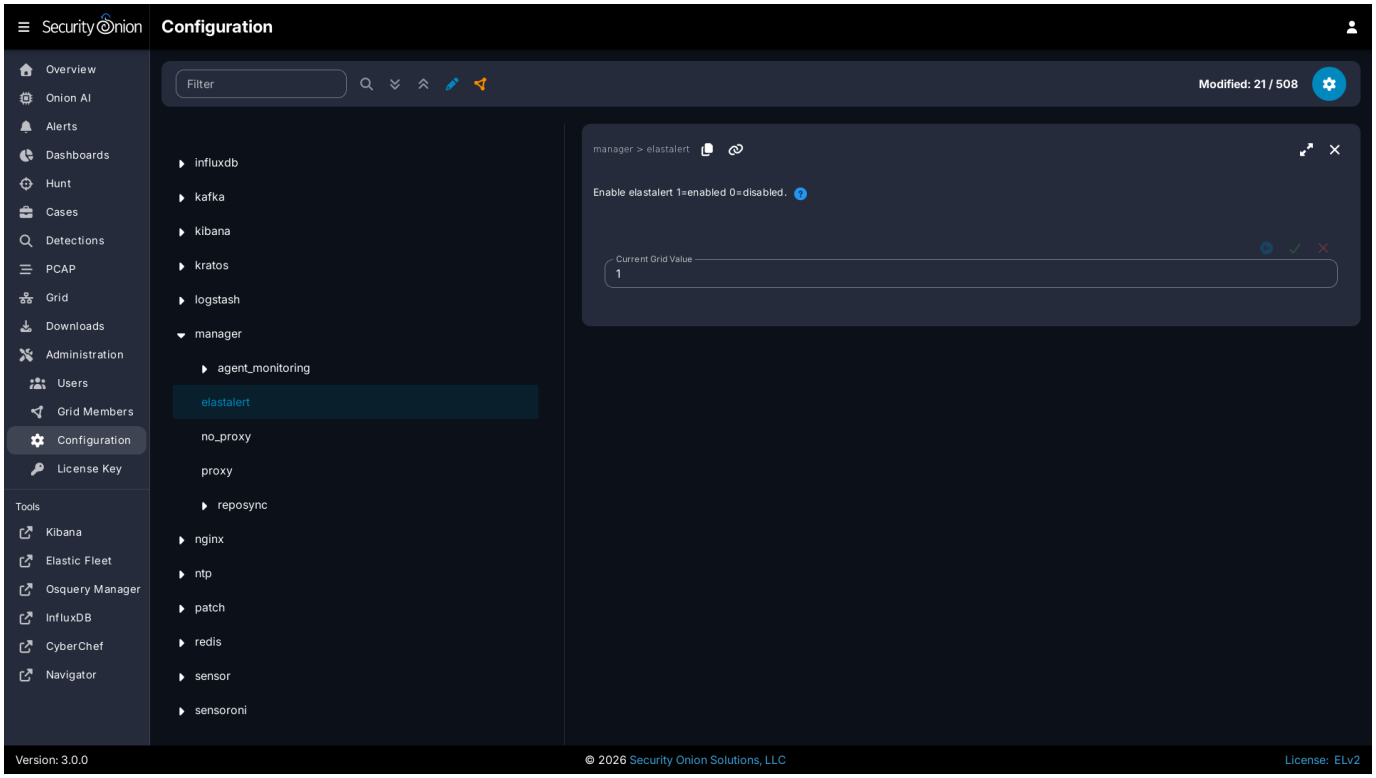

Once there, select the `proxy` or `no_proxy` options.

17.4.2 General Information

There is no way to set a global proxy on Linux, but several tools will route their traffic through a proxy if the following lines are added to `/etc/environment`:

```
http_proxy=<proxy_url>
https_proxy=<proxy_url>
ftp_proxy=<proxy_url>
no_proxy="localhost, 127.0.0.1, <management_ip>, <hostname>"
```

Where:

- `<proxy_url>` is the url of the proxy server. (For example, `http://10.0.0.2:3128` or `https://user:password@your.proxy.url`)
- `<management_ip>` is the IP address of the Security Onion box.
- `<hostname>` is the hostname of the Security Onion box.

> **Note**
>
> You may also need to include the IP address and hostname of the manager in the `no_proxy` variable above if configuring the proxy on a sensor node.

To configure Docker proxy settings, please see https://docs.docker.com/network/proxy/.

To configure git to use a proxy for all users, add the following to `/etc/gitconfig`:

```
[http]
  proxy = <proxy_url>
```

17.4.3 sudo

If you're going to run something using sudo, remember to use the `-i` option to force it to process the environment variables. For example:

```
sudo -i so-suricata-restart
```

> **Warning**
>
> Using `sudo su -` will ignore `/etc/environment`, instead use `sudo su` if you need to operate as root.

17.4.4 NIDS Rules

If you are using a proxy and need to download NIDS rulesets, you will also need to configure proxy settings for the NIDS ruleset downloads. These settings are separate from the system-wide proxy configuration above. See the NIDS documentation for details on configuring the Proxy URL, Proxy Username, Proxy Password, and Proxy CA Path for ruleset downloads.

17.5 Firewall

This section will cover both network firewalls outside of Security Onion and the host-based firewall built into Security Onion.

17.5.1 Network Firewalls

This first sub-section will discuss network firewalls outside of Security Onion.

Internet Communication

When configuring network firewalls for Internet-connected deployments (non-Airgap), you'll want to ensure that the deployment can connect outbound (TCP/443) to the following:

- raw.githubusercontent.com (Security Onion public key)
- sigs.securityonion.net (Signature files for Security Onion containers)
- ghcr.io (Container downloads)
- pkg-containers.githubusercontent.com (Container downloads)
- rules.emergingthreatspro.com (Emerging Threats IDS rules)
- rules.emergingthreats.net (Emerging Threats IDS open rules)
- github.com (Strelka and Sigma rules updates)
- objects.githubusercontent.com (Strelka and Sigma rules updates)

If you are using our ISO image, you will also need access to the following:

- repo.securityonion.net (primary repo for Oracle Linux package updates)
- repo-alt.securityonion.net (secondary repo for Oracle Linux package updates)
- so-repo-east.s3.us-east-005.backblazeb2.com (secondary repo for Oracle Linux package updates)

If you choose to enable GeoIP updates for Elasticsearch, you will also need access to the following:

- geoip.elastic.co
- storage.googleapis.com

If you choose to enable the Snort Talos ruleset, you will also need access to the following:

- www.snort.org

Node Communication

When configuring network firewalls for distributed deployments, you'll want to ensure that nodes can connect as shown in the table below. Please note that some of the sources and destinations listed in the table have specific definitions:

- `Security Onion Grid nodes` includes any node joined to your manager. This includes search nodes, sensors, fleet nodes, receiver nodes, and IDH nodes.
- `Search nodes` are Grid nodes that run Elasticsearch and join to the manager to enlarge its Elastic cluster.
- `Elastic cluster nodes` include the search nodes and the manager itself.
- `Endpoint Elastic Agents` includes any endpoint where you have deployed the Elastic Agent and want to send the data to your Security Onion Grid.

Source (SRC)	Destination (DST)	Destination Port(s) (TCP)	Description
Security Onion Grid nodes	Manager	443, 4505, 4506, 5000, 5055, 8086, 8220, 8443	Management, Registry, Salt, Updates
Security Onion Grid nodes	Fleet node	5055, 8220	Elastic Agent data and management
Security Onion Grid nodes	Receiver node	5055	Elastic Agent data
Search nodes	Manager	443, 4505, 4506, 5000, 5055, 8086, 8220, 8443, 9696	Management, Registry, Salt, Updates, Redis
Elastic cluster nodes	Elastic cluster nodes	9200, 9300	Logstash to Elasticsearch and Elasticsearch node-to-node
Endpoint Elastic Agents	Manager	8220, 8443, 5055	Elastic Agent management, binary updates, data
Endpoint Elastic Agents	Fleet node	5055, 8220	Elastic Agent management and data
Endpoint Elastic Agents	Receiver node	5055	Elastic Agent data
Fleet node	Receiver node	5056	Logstash-to-Logstash
Fleet node	Manager	5056, 9200	Logstash-to-Logstash and Elasticsearch node-to-node
Manager	IDH node	2222	SSH for management

17.5.2 Host Firewall

The remainder of this section will cover the host firewall built into Security Onion.

> **Note**
>
> Security Onion locks down the firewall by default.

17.5.3 Configuring Host Firewall

You can configure the firewall by going to Administration --> Configuration --> firewall --> hostgroups.

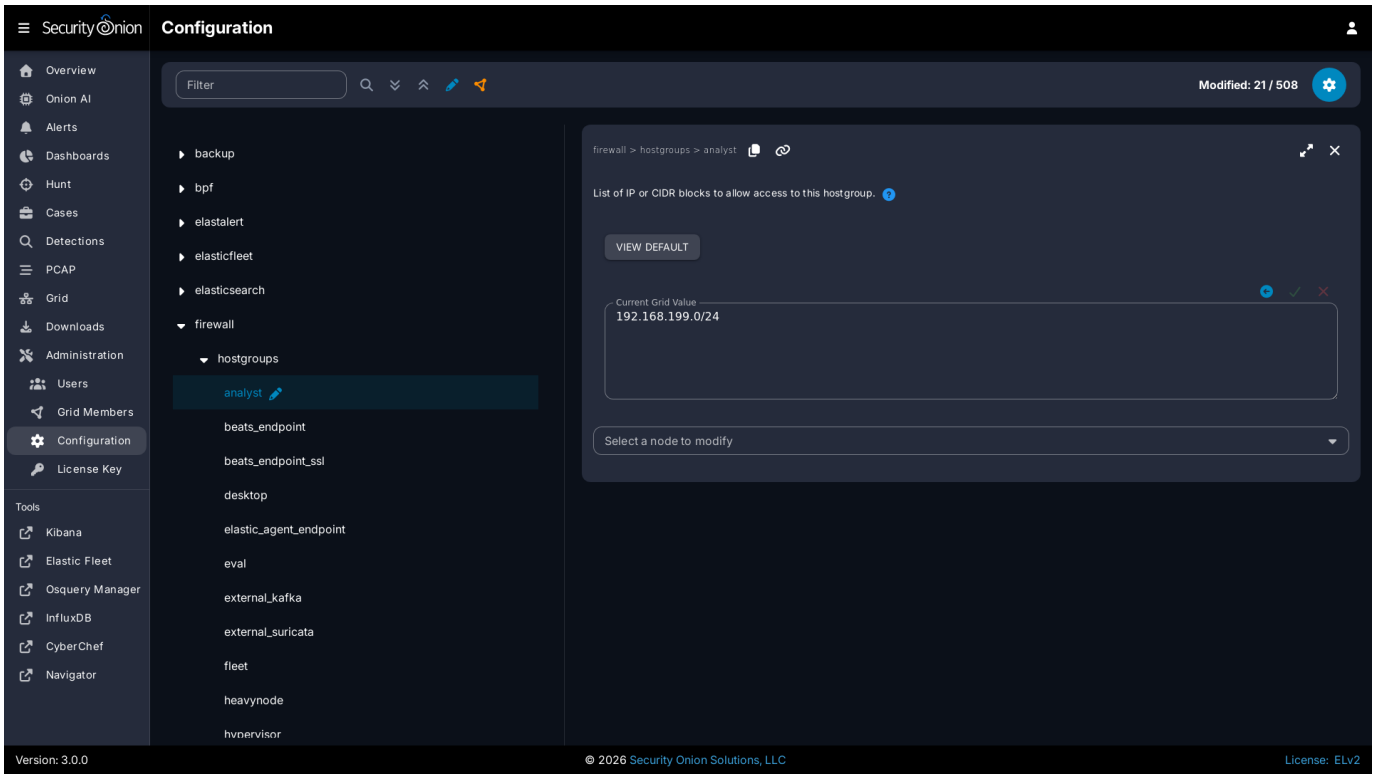

If for some reason you can't access SOC at all, you can use the so-firewall command to allow the IP address of your web browser to connect (replacing `<IP ADDRESS>` with the actual IP address of your web browser):

```
sudo so-firewall includehost analyst <IP ADDRESS>
```

17.5.4 Reviewing Host Firewall

You can view the entire firewall configuration from the command line using the `iptables` command like this:

```
sudo iptables -nvL
```

> **Warning**
>
> You can use this command to view the iptables configuration, but please do not modify the firewall manually using iptables as it is managed by Salt. You should only make changes via the Configuration screen as shown above.

17.5.5 Port Groups

Port groups are a way of grouping together ports similar to a firewall port/service alias. For example, if you have a web server you might add ports 80 and 443 into a port group.

17.5.6 Host Groups

Host groups are similar to port groups but for storing lists of hosts that will be allowed to connect to the associated port groups.

17.5.7 Function

The firewall state is designed with the idea of creating port groups and host groups, each with their own alias or name, and associating the two in order to create an allow rule. A node that has a port group and host group association assigned to it will allow those hosts to connect to those ports on that node.

The default allow rules for each node are defined by its role (manager, searchnode, sensor, heavynode, etc) in the grid. Host groups and port groups can be created or modified from the manager node by going to Administration --> Configuration --> firewall --> hostgroups. When setup is run on a new node, it will ask the manager to add itself to the appropriate host groups. All node types are added to the minion host group to allow Salt communication. If you were to add a search node, you would see its IP appear in both the `minion` and the `search_node` host groups.

17.5.8 Advanced Firewall Config

When you go to Administration --> Configuration --> firewall, you will only see `hostgroups` by default. If you need to modify port groups, then you will need to click the `Options` menu and then enable the `Show advanced settings` option.

Modifying a default port group

The analyst hostgroup is allowed access to the nginx ports which are 80 and 443 by default. In this example, we will extend the default nginx port group to include a custom port.

- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- On the left side, go to `firewall`, select `portgroups`, locate the `nginx` portgroup, and then select `tcp`.
- On the right side, select the manager node, specify your custom port to be added, and then click the checkmark to save the value.
- If you would like to apply the rules immediately, click the `SYNCHRONIZE GRID` button under the `Options` menu at the top of the page.

Creating a custom host group with a custom port group

In this example, we will add a new custom hostgroup to allow a custom set of hosts to connect to a custom port on an IDH node.

- At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option.
- On the left side, go to `firewall`, select `hostgroups`, and then select `customhostgroup0`.
- On the right side, select the IDH node that you want to allow access to, add the list of hosts that require access, and then click the checkmark to save the value.
- On the left side, go to `firewall`, select `portgroups`, select `customportgroup0`, and then select the appropriate protocol.
- On the right side, select the IDH node that you want to allow access to, add your custom port, and then click the checkmark to save the value.
- On the left side, go to `firewall`, `role`, and then select `IDH`, `chain`, `DOCKER-USER`, `hostgroups`, `customhostgroup0`, `portgroups`.
- On the right side, select the IDH node that you want to allow access to, add the portgroup `customportgroup0`, and then click the checkmark to save the value.
- The next time the IDH node checks in, it should get the appropriate firewall rules.

17.6 Email

Some applications rely on having a mail server in the OS itself and other applications have their own mail configuration and so they don't rely on a mail server in the OS itself.

17.6.1 Operating System

You can install and configure your favorite mail server. Depending on your needs, this could be something simple like `nullmailer` or something more complex like `exim4`.

17.6.2 ElastAlert

Follow the steps in the ElastAlert section.

17.7 NTP

Depending on how you installed, the underlying operating system may be configured to pull time updates from the NTP Pool Project and perhaps others as a fallback. You may want to change this default NTP config to your preferred NTP provider by going to Administration --> Configuration --> ntp.

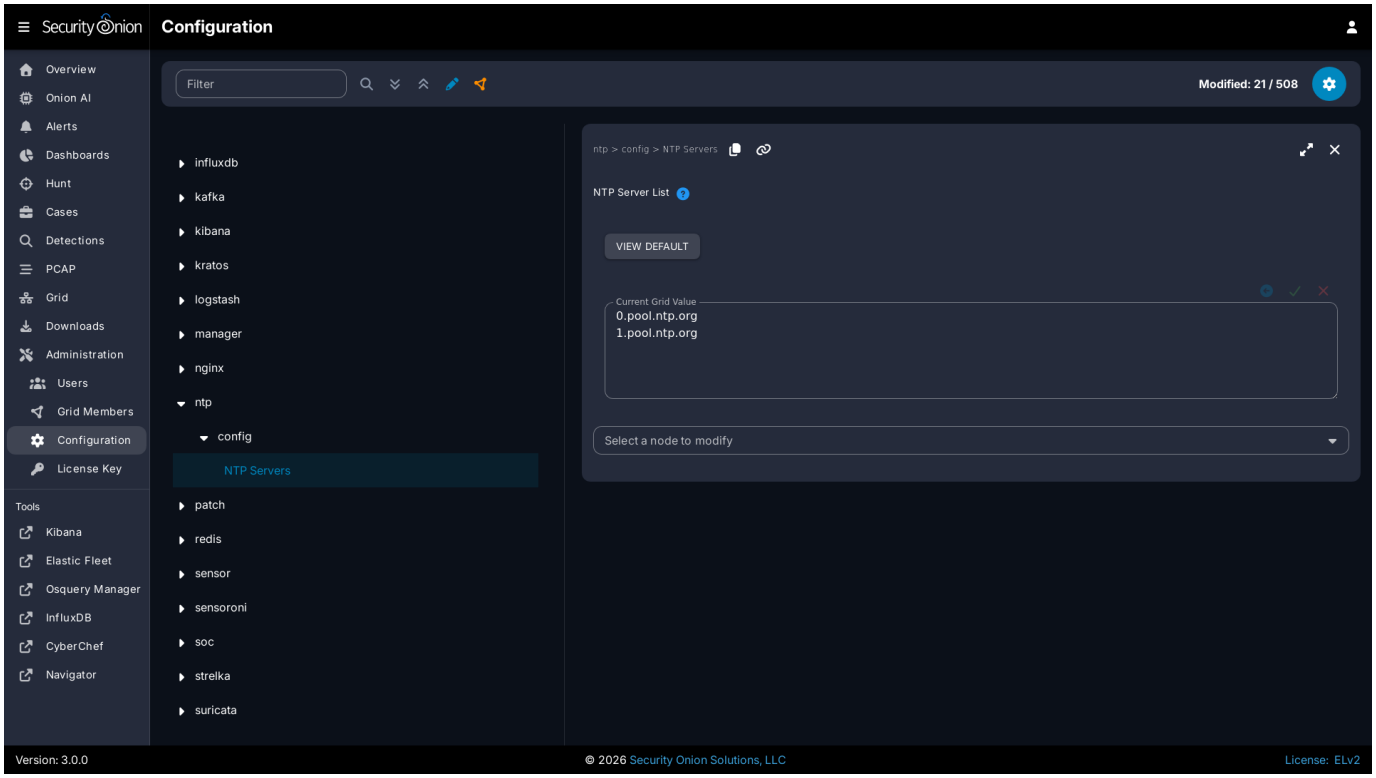

For a distributed deployment, it's vitally important that all nodes have their clock synchronized. Otherwise, you may run into issues where logs or other types of data appear to be missing.

For more information about the operating system time, please see the Time Zones section.

17.8 Console

The current version of Security Onion automatically disables kernel messages in the local console (tty). If you are running an older version of Security Onion and log into the local console, you may see lots of messages from the Linux kernel. To avoid these kernel messages, you have a few options:

- You can use SSH instead of the local console.
- If you really need to use the local console, you can temporarily disable console messages with `sudo dmesg -D`. For more information about dmesg, please see https://man7.org/linux/man-pages/man1/dmesg.1.html. Also see https://man7.org/linux/man-pages/man8/sysctl.8.html and https://www.kernel.org/doc/html/next/core-api/printk-basics.html.
- Upgrade to the latest version of Security Onion using soup.

17.9 SSH

Security Onion uses the latest SSH packages. It does not manage the SSH configuration in `/etc/ssh/sshd_config` with Salt. This allows you to add any PAM modules or enable two factor authentication (2FA) of your choosing.

17.10 Hostname

Setup generates certificates based on the hostname and we do not support changing the hostname after Setup. Please make sure that your hostname is correct during installation as mentioned in the Best Practices section.

17.11 IP Address

The Best Practices section recommends that you avoid changing IP addresses after installation. If for some reason you must do so, you can try the experimental utility `so-ip-update`.

> **Warning**
>
> `so-ip-update` is an experimental utility and only supports standalone machines, not distributed deployments.

> **Warning**
>
> You may still need to manually update the IP address in some settings like the Elastic Fleet Agent Binary Download setting.

17.12 DNS

DNS is normally configured during initial setup. If you need to later change your DNS settings, you can use NetworkManager's console utilities, nmtui (text based user interface) and nmcli (command line inteface). You can learn more at https://docs.oracle.com/en/operating-systems/oracle-linux/9/network/network-NetworkConfigurationTools.html.

17.13 Web Access URL

If you need to change the URL for web access to Security Onion (for example, from IP to FQDN), go to Administration --> Configuration --> global --> url_base. Enter the new URL in the field on the right and then click the checkmark to save the new setting.

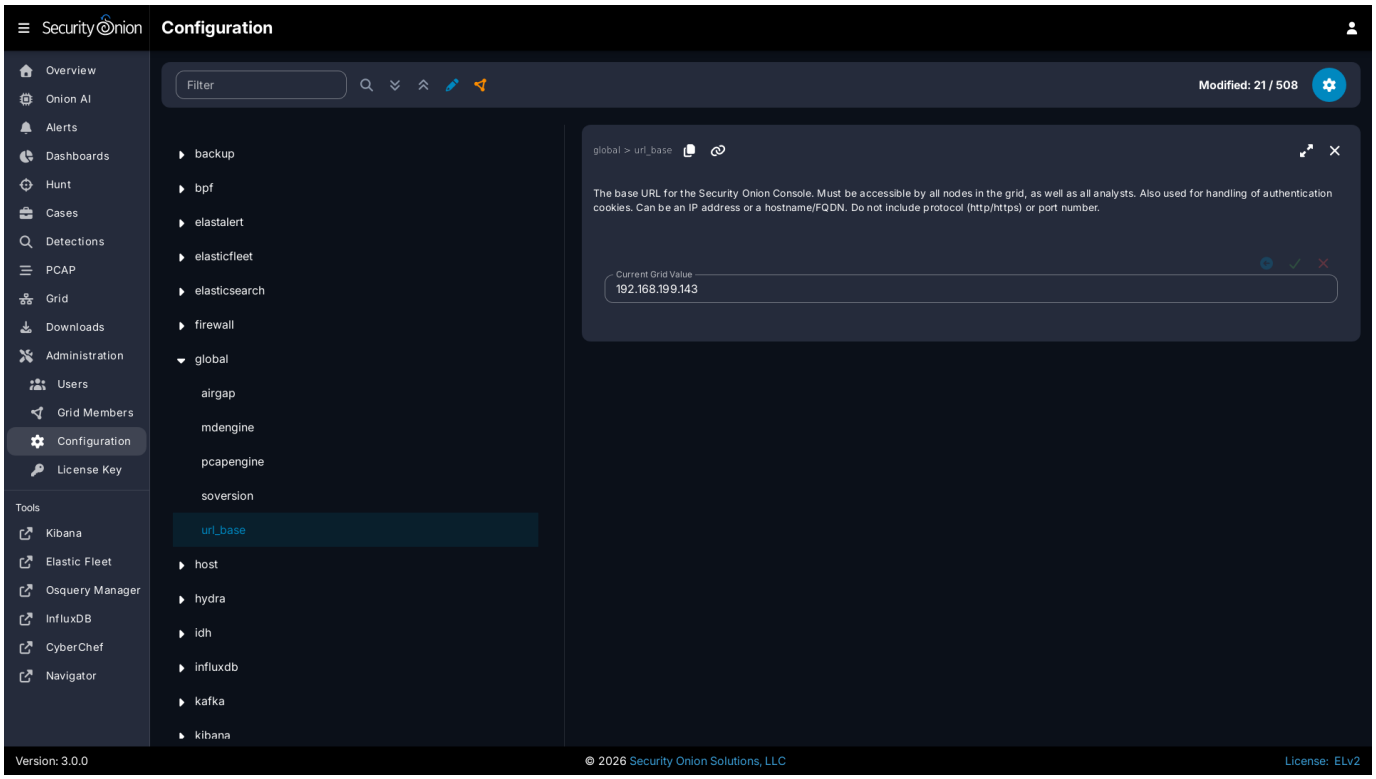

18. Tricks and Tips

18.1 Tricks and Tips Overview

This section is a collection of miscellaneous tricks and tips for Security Onion.

18.2 Backup

Security Onion performs a daily backup of some critical files so that you can recover your grid from a catastophic failure of the manager. Daily backups create a tar file located in the `/nsm/backup/` directory located on the manager. You may want to replicate this backup directory to a location outside of your manager in case the manager ever needs to be rebuilt.

Here is what gets backed up automatically by default:

- `/etc/pki/` - All of the certs including the CA.
- `/etc/salt/` - Configuration for the Salt manager and minions.
- `/nsm/kratos/` - Configuration for Kratos.
- `/nsm/hydra/` - Configuration for Hydra (used for Connect).
- `/opt/so/saltstack/local/` - Customizations done via Administration --> Configuration.

If you need to restore one or more files from backup, locate the tar backup file from the desired date and use the standard `tar` command to expand the file. For example, to expand the backup file from March 17, 2025:

```
tar xvf so-config-backup-2025_03_17.tar
```

This will extract the config files from the tar file into subdirectories for `etc`, `nsm`, and `opt`. You can then copy the needed files from those expanded subdirectories to the actual directories.

18.2.1 Configuration

You can configure backups by going to Administration --> Configuration --> backup.

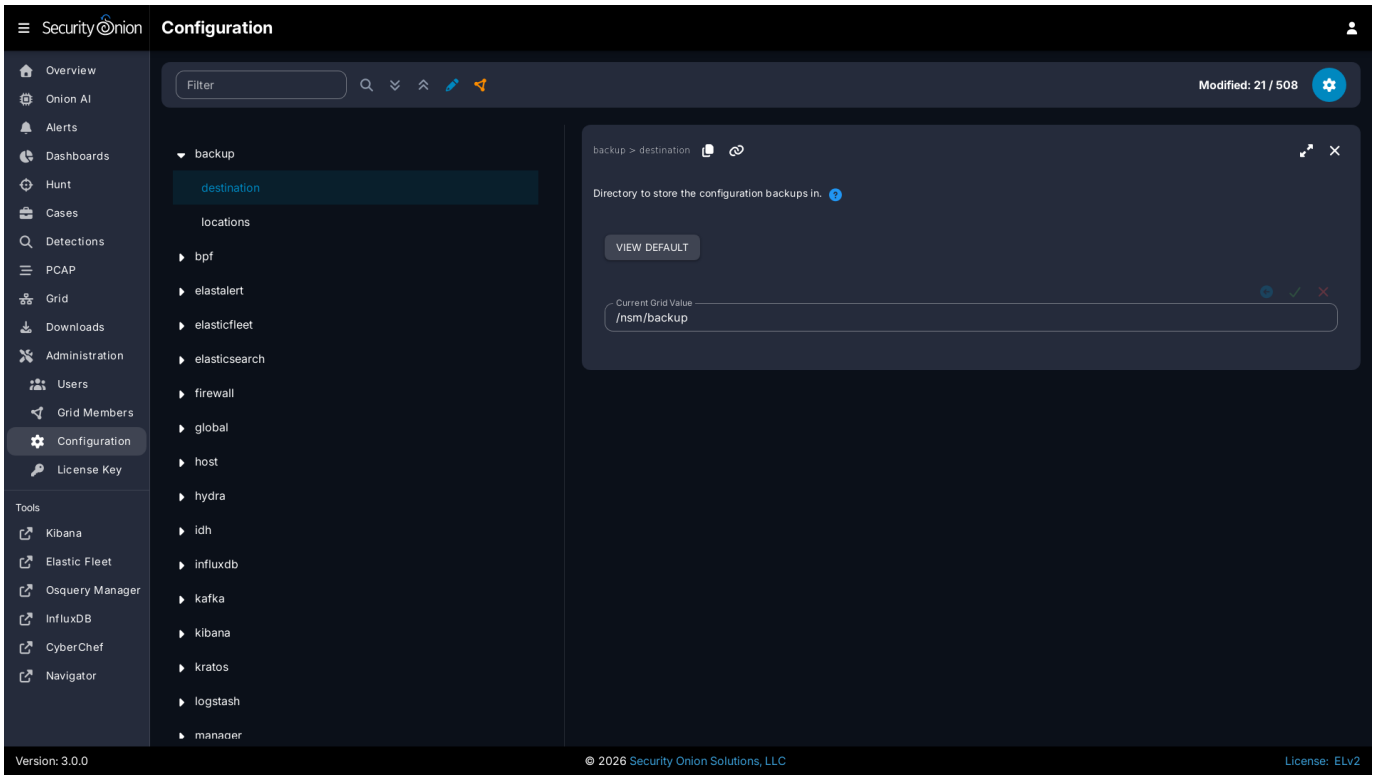

18.2.2 Elasticsearch

Elasticsearch data is not automatically backed up. This includes things that may be important to you like Kibana customizations and Cases data. Kibana customizations are located in the `.Kibana` indices and Cases data is stored in the `so-case` and `so-casehistory` indices. If you have a

distributed deployment with Elasticsearch clustering, then you can enable replicas to have redundancy in case of a single node failure. Of course, please keep in mind that enabling replicas doubles your storage needs.

Another option is to use Elasticsearch's built-in support for snapshots: https://www.elastic.co/guide/en/elasticsearch/reference/current/snapshot-restore.html

This option requires that you configure Elasticsearch with a `path.repo` setting where it can store the snapshots. Once Elasticsearch has the `path.repo` setting, you should be able to log into Kibana and configure snapshots as shown in the link above. Those snapshots will then be accessible in `/nsm/elasticsearch/repo/`.

18.3 Docker

From https://www.docker.com/what-docker:

> Docker is the world's leading software container platform. Developers use Docker to eliminate "works on my machine" problems when collaborating on code with co-workers. Operators use Docker to run and manage apps side-by-side in isolated containers to get better compute density. Enterprises use Docker to build agile software delivery pipelines to ship new features faster, more securely and with confidence for both Linux, Windows Server, and Linux-on-mainframe apps.

18.3.1 Download

Our ISO image includes the Docker engine and all of our Docker images.

18.3.2 Security

To prevent tampering, our Docker images are signed using GPG keys. Soup verifies GPG signatures any time Docker images are updated.

18.3.3 Elastic

To maintain a high level of stability, reliability, and support, our Elastic Docker images are based on the Docker images provided by Elastic.co.

18.3.4 Images

After installation, you can see all Docker images with the following command:

```
sudo docker images
```

18.3.5 Logs

If a service is not writing its logs to `/opt/so/log`, then you may need to check the Docker logs for more detail. For example, to check the Docker logs for Kibana:

```
sudo docker logs so-kibana
```

18.3.6 Registry

The manager node runs a Docker registry. From https://docs.docker.com/registry/recipes/mirror/:

```
If you have multiple instances of Docker running in your environment (e.g., multiple physical or virtual machines, all running the Docker daemon), each time
one of them requires an image that it doesn't have it will go out to the internet and fetch it from the public Docker registry. By running a local registry
mirror, you can keep most of the redundant image fetch traffic on your local network.
```

If you see errors relating to `so-dockerregistry` (Docker Registry), then please take a look at the following discussions to see if your symptoms match and if their solutions may help you:

https://github.com/Security-Onion-Solutions/securityonion/discussions/12078

https://github.com/Security-Onion-Solutions/securityonion/discussions/12635

18.3.7 Networking and Bridging

By default, Docker configures its network bridge with an IP address of `172.17.0.1`. This works fine for networks that aren't already using the `172.17.0.0/16` range. If you are using this range in your network, then you can change the Docker range during installation.

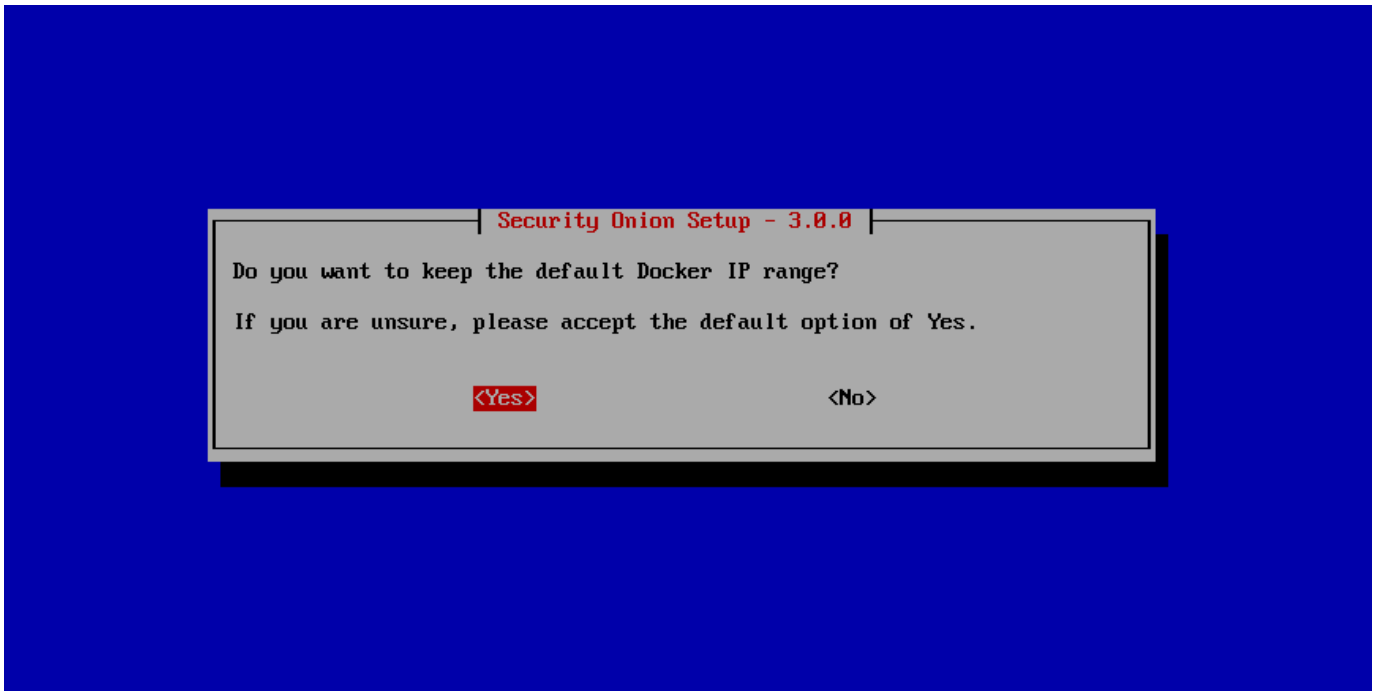

18.3.8 Containers

Our Docker containers all belong to a common Docker bridge network, called `sobridge`. Each container is also aliased, so that communication can occur between the different docker containers using said alias. For example, communication to the `so-elasticsearch` container would occur through an alias of `Elasticsearch`.

You may come across interfaces in `ifconfig` with the format `veth*`. These are the external interfaces for each of the Docker containers. These interfaces correspond to internal Docker container interfaces (within the Docker container itself).

To identify which external interface belongs to which container, we can do something like the following:

From the host, type:

```
sudo docker exec so-elasticsearch cat /sys/class/net/eth0/iflink
```

This should provide you with a value with which you can grep the host `net` class `ifindex(es)`:

Example:

```
grep 25 /sys/class/net/veth*/ifindex | cut -d'/' -f5
```

You should then receive some output similar to the following:

```
vethc5ff027
```

where `vethc5ff027` is the external interface of `eth0` within the `so-elasticsearch` container.

18.3.9 VMware Tools

If you have VMware Tools installed and you suspend and then resume, the Docker interfaces will no longer have IP addresses and the Elastic stack will no longer be able to communicate. One workaround is to remove `/etc/vmware-tools/scripts/vmware/network` to prevent VMware suspend/resume from modifying your network configuration.

18.3.10 More Information

Note

For more information about Docker, please see https://www.docker.com/what-docker.

18.4 Jupyter Notebook

18.4.1 Overview

This section is a brief overview of connecting a Jupyter notebook/server instance to Elasticsearch to slice and dice the data as you wish. It will not cover the setup of a Jupyter instance, which has been thoroughly documented (using Docker) at https://jupyter-docker-stacks.readthedocs.io/en/latest/index.html.

18.4.2 Security Onion Setup

Create Jupyter User

As a best practice, you'll want to create a dedicated Jupyter notebook user with just read-only access to the data inside of Elasticsearch. In Kibana, navigate to Stack Management -> Users and create the user with appropriate permissions.

Security Onion Firewall

In order to allow network-based access to Elasticsearch, you'll need to allow the traffic through the host-based firewall by going to Administration --> Configuration --> firewall --> hostgroups.

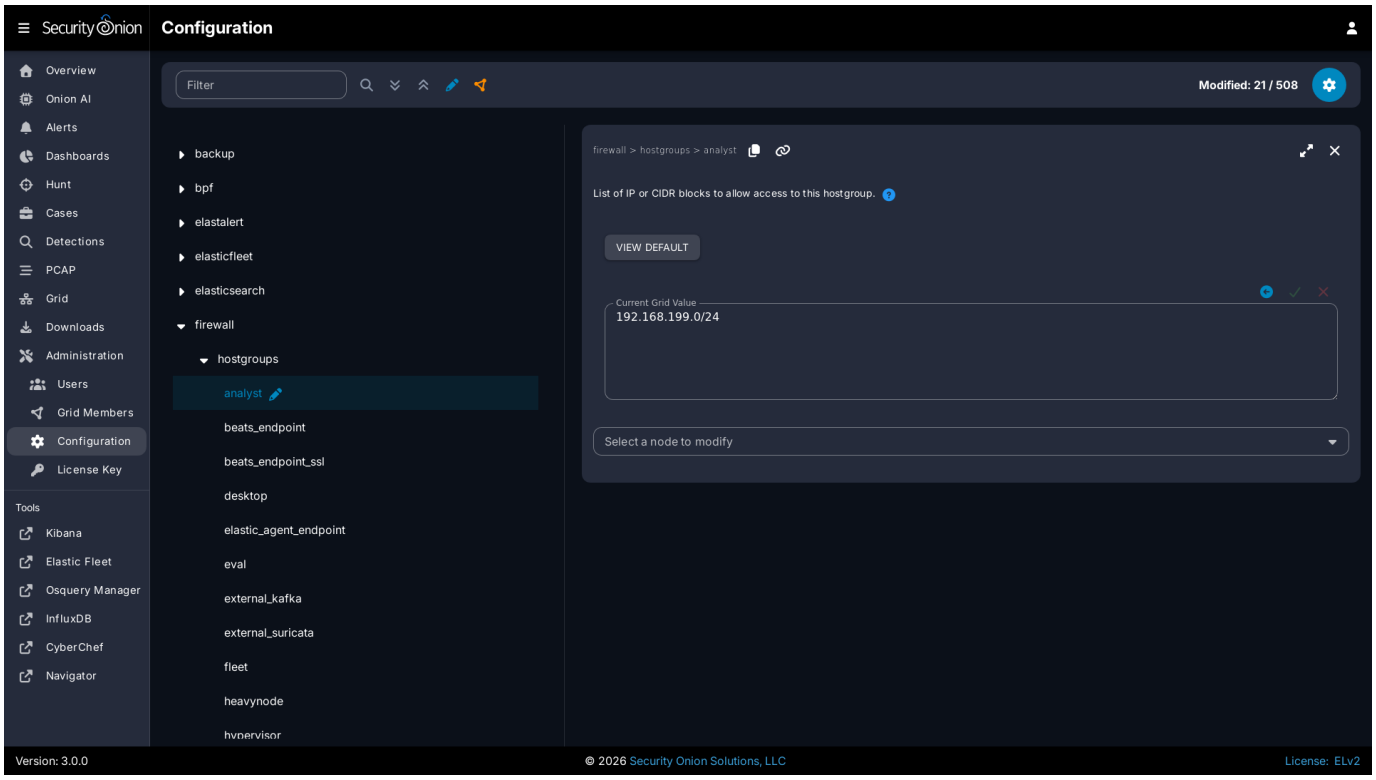

At the top of the page, click the `Options` menu and enable the `Show advanced settings` option. On the left side, select the `elasticsearch_rest` option. On the right side, add your IP address or CIDR blocks and click the checkmark to save.

Once complete, you should be able to connect to the Elasticsearch instance. You can confirm connectivity using tools like curl or Powershell's Test-NetConnection.

18.4.3 Jupyter Notebook

> **Note**
>
> The following steps are heavily inspired by Roberto Rodriguez's Medium post:
>
> https://medium.com/threat-hunters-forge/jupyter-notebooks-from-sigma-rules-%EF%B8%8F-to-query-Elasticsearch-31a74cc59b99

The Jupyter environment will need to have at least the following Python libraries installed:

- Elasticsearch
- elasticsearch_dsl
- pandas

You can install these using the following commands on the Jupyter host, or within the Jupyter Docker container:

```
pip3 install Elasticsearch
pip3 install elasticsearch_dsl
pip3 install pandas
```

Once the Python prerequisites are installed, we can start executing commands within our notebook.

We'll start with importing the libraries we just mentioned. In the first cell, we'll paste the following:

```
from Elasticsearch import Elasticsearch
from elasticsearch_dsl import Search
import pandas as pd
```

Then, we'll press **Shift+ENTER** to execute the command(s) within the cell (can also click to run the cell from the Run menu).

In the next cell, we'll specify the Elasticsearch instance address and port (`192.168.6.100:9200`) and the username (`jupyter`) and password (`password`) we created within Security Onion, as well as the index filter we would like to use for searching (`*:so-*`):

```
es = Elasticsearch(['https://192.168.6.100:9200'],
ca_certs=False,verify_certs=False, http_auth=('jupyter','password'))
searchContext = Search(using=es, index='*:so-*', doc_type='doc')
```

> **Note**
>
> We are choosing to use `verify_certs=False` here to avoid complications with self-signed certificates during testing. Ideally, we would want to make sure we are performing verification wherever possible.

Again, we'll execute the code within the cell, by pressing **Shift+ENTER**.

We may see a certificate warning due to the fact that we are not performing verification for certificates.

For convenience during our testing, we can disable the warning in future runs, by pasting the following the next cell and executing it with **Shift+ENTER**:

```
import urllib3
urllib3.disable_warnings(urllib3.exceptions.InsecureRequestWarning)
```

In the following cell, we'll paste the following:

```
s = searchContext.query('query_string', query='event.module:sysmon')
```

In this example, we are looking for logs that contain a field called `event.module` and a value of `sysmon` (Sysmon logs). Once more, we'll press **Shift+ENTER** and continue on.

Finally, we'll submit our query in the next cell using the following:

```
response = s.execute()
if response.success():
    df = pd.DataFrame((d.to_dict() for d in s.scan()))
df
```

The above code simply takes the results and converts them to a Python dict.

We can select a few fields, and modify the column values if we like:

```
response = s.execute()
if response.success():
    df = pd.DataFrame(([d['event']['dataset'], d['process']['executable'], d['file']['target']] for d in s))
df.columns=['Dataset','Executable', 'Target']
df
```

Then we end up with something a little bit more targeted (you may need to adjust `pd.options.display.max_colwidth` for it to display appropriately).

Obviously, there is much more we can do with this data other than just running the above example code. Happy hunting!

18.5 Adding Disk Space

If you ever need to add disk space to expand your partitions, there are a few different ways to do this.

> **Warning**
>
> Before doing this in production, make sure you practice this on a non-production system!

18.5.1 Recommend Method: LVM (Logical Volume Management)

If your disk is partitioned using LVM (the default for our Security Onion ISO image), then you should be able to use LVM to add new disk space to your LVM partitions.

> **Note**
>
> For more information about LVM, please see https://docs.oracle.com/en/operating-systems/oracle-linux/9/stordev/stordev-WorkingWithLogicalVolumeManager.html#using-lvm.

LVM Example

For a simple LVM example, suppose that you have installed our Security Onion ISO image in a virtual machine (VM) using a virtualization solution like Proxmox that allows you to increase storage on the fly. Our Security Onion ISO image automatically uses LVM and should use the XFS filesystem for `/nsm`, so if you want to add space to `/nsm` here is a brief overview of the steps you would use.

First, expand the storage. If using Proxmox, select the VM, select `Hardware`, select the `Hard Disk`, click `Disk Action`, click `Resize`, enter the amount that you would like to add to the existing disk size, and then click the `Resize disk` button.

Next, log into the VM, become root, and use a partition tool like `cfdisk` to resize the LVM partition to take advantage of the additional free space:

```
cfdisk
```

Now that the LVM partition is larger, we need to enlarge the LVM physical volume (PV). Start by getting a list of all PVs:

```
pvdisplay
```

Determine the PV to be enlarged and then use the `pvresize` command (replacing `/dev/sda3` with the desired PV):

```
pvresize /dev/sda3
```

Verify that the desired PV has been enlarged:

```
pvdisplay
```

Verify that the volume group (VG) now shows free space:

```
vgdisplay
```

Now that the VG has free space, we need to enlarge the logical volume (LV). Start by getting a list of all LVs:

```
lvdisplay
```

Determine the LV to be enlarged and then use the `lvresize` command (replacing `/dev/system/nsm` with the desired LV):

```
lvresize /dev/system/nsm -l +100%FREE
```

Verify that the desired LV has been enlarged:

```
lvdisplay
```

Now that the LV is larger, we need to enlarge the filesystem. Start by getting a list of all filesystems:

```
df -Th
```

Determine the filesystem to be enlarged and then use the appropriate command (replacing `xfs_growfs` with the appropriate tool based on the filesystem type and replacing `/dev/mapper/system-nsm` with the appropriate filesystem):

```
xfs_growfs /dev/mapper/system-nsm
```

Finally, verify that the filesystem has been enlarged:

```
df -Th
```

18.5.2 Unsupported Methods

> **Warning**
>
> There are other methods for expanding your partitions, but they are UNSUPPORTED. You should only attempt these methods if you really know what you are doing.

If you aren't using LVM but you need to expand your `/nsm` partition, then you can mount a separate physical drive directly to `/nsm`. If doing this after installation, you will need to stop services, move the data to the new drive, and then restart services.

A variation on this method is to make `/nsm` a symbolic link to the new logging location. Certain services like AppArmor may need special configuration to handle the symlink.

18.6 Network Installation

> **Warning**
>
> Network installations are NOT supported and should only be used as a last resort in case there is some reason you can't use our official Security Onion images as shown in the Installation section.

Our official Security Onion images (ISO image and cloud images) are the ONLY supported installation method and you should use them if any of the following apply to you:

- You are deploying in an enterprise environment.
- You are deploying in an airgap environment.
- You are performing a distributed deployment.
- You want the quickest and easiest installation with the fewest issues.
- You need any kind of support.

If NONE of the above apply to you, then you MAY be able to install Oracle Linux 9 and then perform a network installation.

18.6.1 Partitioning

Our official Security Onion images take care of partitioning for you. However, if you choose to perform a network installation then it's your responsibility to make sure that partitions are configured correctly to avoid filling up a partition.

Minimum Storage

As the Hardware section mentions, the MINIMUM requirement is 200GB storage. This is to allow 100GB for `/nsm` and 100GB for the rest of `/`.

LVM

You may want to consider Logical Volume Management (LVM) as it will allow you to more easily change your partitioning in the future if you need to.

/boot

You probably want a dedicated `/boot` partition of at least 1GB at the beginning of the drive.

/nsm

The vast majority of data will be written to `/nsm`, so you'll want to dedicate the vast majority of your disk space to that partition. You'll want at least 100GB.

/

`/` (the root partition) currently contains `/var/lib/docker/` (more on that below) and thus you'll want at least 100GB.

Docker

Docker images are stored in `/var/lib/docker/`. The current set of Docker images uses 30GB on disk. If you're planning a production deployment, you should plan on having enough space for another set of those Docker images for in-place updates.

Other

The OS installer may try to dedicate a large amount of space to `/home`. You may need to adjust this to ensure that it is not overly large and wasting valuable disk space.

Example

Here's an example of how our current ISO image partitions a 1TB disk:

- 1GB `/boot` partition at the beginning of the drive
- the remainder of the drive is an `LVM` volume that is then partitioned as follows:
- 630GB `/nsm`
- 300GB `/`
- 2GB `/tmp`
- 8GB `swap`

18.6.2 Installing via the network

> **Warning**
>
> Please keep in mind that network installations are NOT supported and should only be used as a last resort.

If you understand all of the warnings above and still want to perform a network installation, then you can follow the steps below.

- Review the Hardware and Release Notes sections.
- Download the ISO image for your desired x86-64 operating system. Verify the ISO image and then boot from it.
- Follow the prompts in the installer. If you're building a production deployment, you'll probably want to use LVM and dedicate most of your disk space to `/nsm` as discussed in the Partitioning section above.
- Reboot into your new installation.
- Login using the username and password you specified during installation.
- Install prerequisites:

```
sudo dnf -y install git
```

- Download our repo and start the Setup process:

```
git clone -b 3/main https://github.com/Security-Onion-Solutions/securityonion
cd securityonion
sudo bash so-setup-network
```

- Proceed to the Configuration section.

18.7 PCAPs for Testing

The easiest way to download PCAP files for testing is our so-test tool. Alternatively, you could manually download pcaps from one or more of the following locations:

- https://www.malware-traffic-analysis.net/
- https://digitalcorpora.org/corpora/network-packet-dumps
- https://www.netresec.com/?page=PcapFiles
- https://www.netresec.com/?page=MACCDC
- https://github.com/zeek/zeek/tree/master/testing/btest/Traces
- https://www.ll.mit.edu/r-d/datasets/2000-darpa-intrusion-detection-scenario-specific-datasets
- https://wiki.wireshark.org/SampleCaptures
- https://www.stratosphereips.org/datasets-overview
- https://ee.lbl.gov/anonymized-traces.html
- https://redmine.openinfosecfoundation.org/projects/suricata/wiki/Public_Data_Sets
- https://forensicscontest.com/puzzles
- https://github.com/markofu/hackeire/tree/master/2011/pcap
- https://www.defcon.org/html/links/dc-ctf.html
- https://github.com/chrissanders/packets

You can download PCAP files from the links above using a standard web browser or from the command line using a tool like `wget` or `curl`.

18.7.1 Replay

You can use `tcpreplay` to replay any standard PCAP to the sniffing interface of your Security Onion sensor.

18.7.2 Import

A drawback to using tcpreplay is that it's replaying the PCAP as new traffic and thus the timestamps that you see in SOC and other interfaces do not reflect the original timestamps from the PCAP. To avoid this, you can import the PCAP using the Grid page.

18.8 High Performance Tuning

18.8.1 CPU Affinity/Pinning

For best performance, CPU intensive processes like Zeek and Suricata should be pinned to specific CPUs. In most cases, you'll want to pin sniffing processes to the same CPU that your sniffing NIC is bound to. For more information, please see the Performance subsection in the appropriate Suricata and Zeek sections.

18.8.2 Misc

Consider adopting some of the suggestions from here:

- https://suricata.readthedocs.io/en/latest/performance/packet-capture.html
- https://github.com/pevma/SEPTun
- https://github.com/pevma/SEPTun-Mark-II

18.8.3 RSS

Check your sniffing interfaces to see if they have Receive Side Scaling (RSS) queues. If so, you may need to reduce to 1: https://suricata.readthedocs.io/en/latest/performance/packet-capture.html#rss

18.8.4 Disk/Memory

If you have plenty of RAM, disable swap altogether: https://www.elastic.co/guide/en/elasticsearch/reference/current/setup-configuration-memory.html#disable-swap-files

Use `hdparm` to gather drive statistics and alter settings, as described here: https://www.linux-magazine.com/Online/Features/Tune-Your-Hard-Disk-with-hdparm

`vm.dirty_ratio` is the maximum amount of system memory that can be filled with dirty pages before everything must get committed to disk.

`vm.dirty_background_ratio` is the percentage of system memory that can be filled with "dirty" pages, or memory pages that still need to be written to disk -- before the pdflush/flush/kdmflush background processes kick in to write it to disk.

More information: https://lonesysadmin.net/2013/12/22/better-linux-disk-caching-performance-vm-dirty_ratio/

18.8.5 Elastic

You will want to make sure that each part of the pipeline is operating at maximum efficiency. Depending on your configuration, this may include Elastic Agent, Logstash, Redis, and Elasticsearch.

18.9 Removing a Node

There may come a time when you need to remove a node from your distributed deployment. To do this, you'll need to remove the node's configuration from a few different components.

18.9.1 Removing a Search Node

To remove a search node, the data stored on the node needs to be migrated off before other node removal actions.

Use the following command in Kibana Dev Tools to stop shard allocation to the node (replacing 10.0.0.1 with the actual IP address of the search node to be removed):

```
PUT _cluster/settings
{
  "persistent" : {
    "cluster.routing.allocation.exclude._ip" : "10.0.0.1"
  }
}
```

For more information about this command, please see https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-cluster.html#cluster-shard-allocation-filtering.

Once all data has migrated off the search node, then you can continue with other node removal actions.

18.9.2 Removing from Salt

You can remove a node from Salt by going to Administration --> Grid Members.

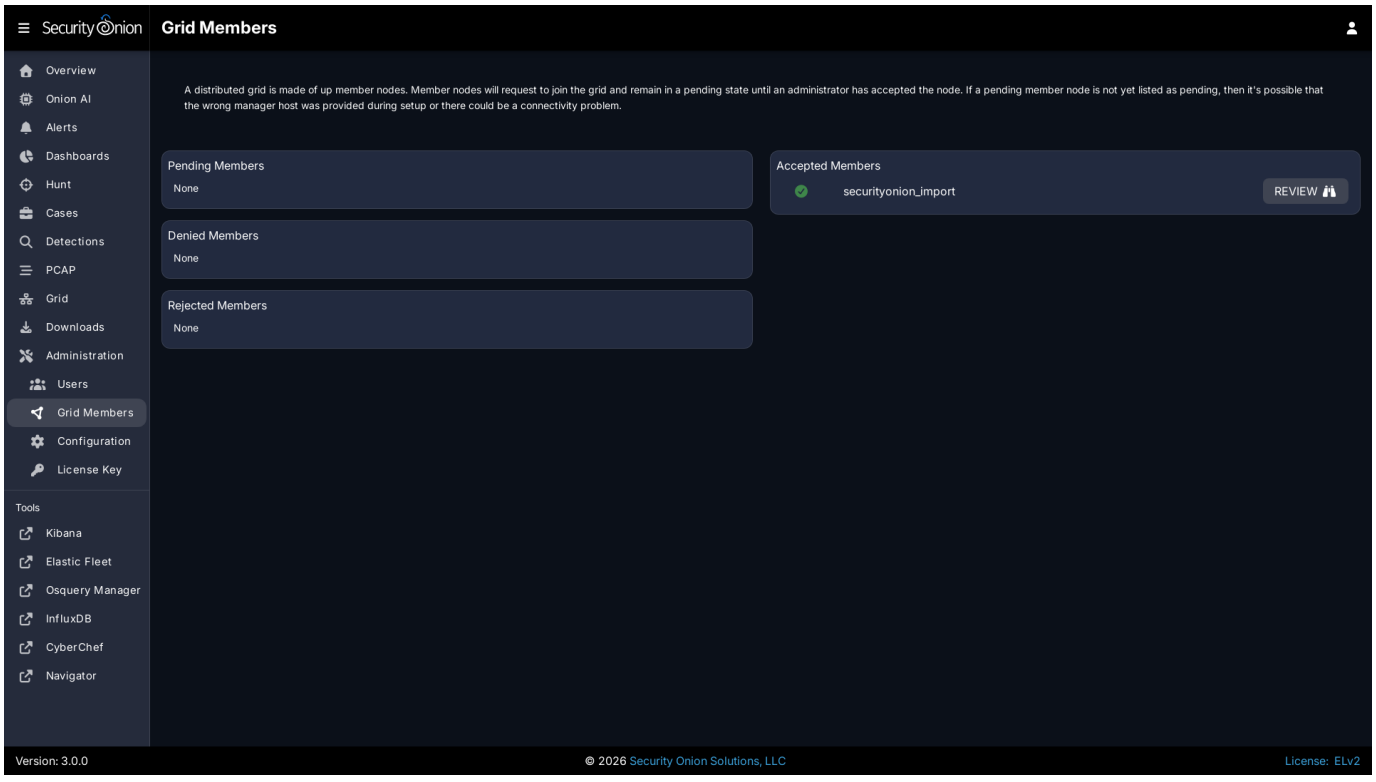

Find the grid Member you would like to remove, click the `REVIEW` button, and then click the `DELETE` button.

18.9.3 Removing from SOC

To remove the node from the SOC Grid page, make sure the node is powered off and then restart SOC:

```
sudo so-soc-restart
```

18.9.4 Removing from Fleet

To remove the node from Elastic Fleet, go to the Agents tab and find the node. Then click the checkbox to the left of the node. Click the `Actions` button and then click `Unenroll 1 agent`. Select the `Remove agent immediately` option and then click the `Unenroll agent` button.

18.10 Salt

From https://docs.saltproject.io/en/latest/topics/about_salt_project.html#about-salt:

> Built on Python, Salt is an event-driven automation tool and framework to deploy, configure, and manage complex IT systems. Use Salt to automate common infrastructure administration tasks and ensure that all the components of your infrastructure are operating in a consistent desired state.

Note

Salt is a core component of Security Onion as it manages all processes on all nodes. In a distributed deployment, the manager node controls all other nodes via salt. These non-manager nodes are referred to as salt minions.

18.10.1 Firewall Requirements

Salt minions must be able to connect to the manager node on ports `4505/tcp` and `4506/tcp`.

18.10.2 Checking Status

You can use salt's `test.ping` to verify that all your nodes are up:

```
sudo salt \* test.ping
```

18.10.3 Remote Execution

Similarly, you can use salt's `cmd.run` to execute a command on all your nodes at once. For example, to check disk space on all nodes:

```
sudo salt \* cmd.run 'df'
```

18.10.4 Node checkin

If you want to force a node to do a full update of all salt states, you can run `so-checkin`. This will execute `salt-call state.highstate -l info` which outputs to the terminal with the log level set to `info` so that you can see exactly what's happening:

```
sudo so-checkin
```

18.10.5 Configuration

Many of the options that are configurable in Security Onion are done by going to Administration and then Configuration.

18.10.6 Diagnostic Logs

Diagnostic logs can be found in `/opt/so/log/salt/`.

18.10.7 Known Issues

You may see the following error in the salt-master log located at `/opt/so/log/salt/master`:

```
[ERROR   ][24983] Event iteration failed with exception: 'list' object has no attribute 'items'
```

The root cause of this error is a state trying to run on a minion when another state is already running. This error now occurs in the log due to a change in the exception handling within Salt's event module. Previously, in the case of an exception, the code would just pass. However, the exception is now logged. The error can be ignored as it is not an indication of any issue with the minions.

18.10.8 More Information

Note

For more information about Salt, please see https://docs.saltproject.io/en/latest/contents.html.

18.11 Syslog Output

If you want to send logs to an external system, you can configure Logstash to output to syslog.

> **Note**
>
> For more information about Logstash's syslog output plugin, please see: https://www.elastic.co/guide/en/logstash/current/plugins-outputs-syslog.html

Please keep in mind that we don't provide free support for third party systems.

18.12 Time Zones

When you run Security Onion Setup, it sets the operating system time zone to UTC/GMT. Logging in UTC is considered a best practice across the cybersecurity industry because it makes it that much easier to correlate events across different systems, organizations, or time zones. Additionally, it avoids issues with time zones that have daylight savings time which would result in a one-hour time warp twice a year.

Web interfaces like Alerts, Dashboards, Hunt, and Kibana should try to detect the time zone of your web browser and then render those UTC timestamps in local time. Alerts, Dashboards, and Hunt allow you to manually set your time zone under Options.

For more information about the operating system time, please see the NTP section.

18.13 Endgame

> **Warning**
>
> Endgame support has not been tested yet!

You can ingest Endgame data by following the steps below.

> **Note**
>
> Please keep in mind that we currently use the `*:endgame-*` index pattern for Endgame data. This means the data will not be visible using the normal Security Onion dashboards/index pattern in Kibana. However, Endgame data will be viewable and aggregatable using Hunt and Elastic Security.

18.13.1 Configuration

To configure Endgame ingestion during setup, ensure the `ENDGAMEHOST` variable is set to the IP address of the Endgame SMP that you want to send data from:

```
sudo ENDGAMEHOST=192.168.1.100 ./so-setup-network
```

This will open the Security Onion host-based firewall for access from the SMP to Security Onion on TCP port 3765.

18.13.2 Pivot to Endgame Console

If Endgame support is enabled, then Dashboards and Hunt will have an `Endgame` action on the Actions menu. Clicking that action will pivot to Endgame Console based on the `agent.id` field.

19. Utilities

19.1 Utilities Overview

This section covers some of the utilities in Security Onion.

19.2 jq

From https://stedolan.github.io/jq/:

> jq is like sed for JSON data - you can use it to slice and filter and map and transform structured data with the same ease that sed, awk, grep and friends let you play with text.

19.2.1 Usage

We configure Zeek and Suricata to write logs to `/nsm/` in JSON format. If you want to parse those logs from the command line, then you can use `jq`. Here's a basic example:

```
jq '.' /nsm/zeek/logs/current/conn.log
```

This command will parse all of the records in `/nsm/zeek/logs/current/conn.log`. For each of the records, it will then output every field and its value.

19.2.2 More Information

Note

For more information about `jq`, please see https://stedolan.github.io/jq/.

19.3 so-allow

In previous versions of Security Onion, so-allow was used to allow traffic through the host-based firewall. This is now done by going to Administration --> Configuration --> firewall --> hostgroups.

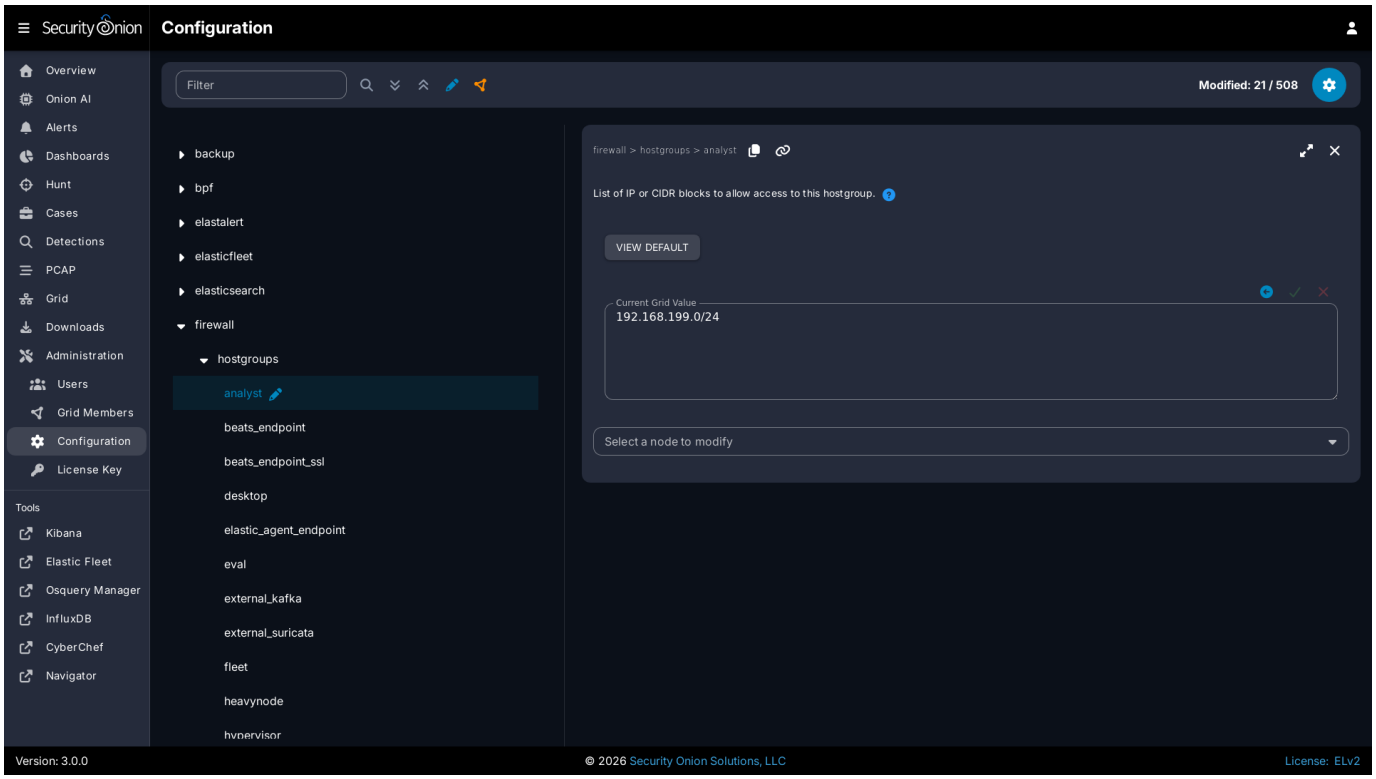

19.4 so-elastic-auth-password-reset

Elastic service accounts use randomly generated passwords that are 72 characters in length. If you need to reset these passwords, you can use the `so-elastic-auth-password-reset` utility.

19.5 so-elasticsearch-query

You can use `so-elasticsearch-query` to submit a cURL request to the local Security Onion Elasticsearch host from the command line.

19.5.1 Usage

```
so-elasticsearch-query <PATH> [ARGS,...]
```

Where:

- PATH represents the elastic function being requested.
- ARGS is used to specify additional, optional curl parameters.

19.5.2 Examples

Here's a basic example:

```
sudo so-elasticsearch-query /
```

Here's a more complicated example that includes piping the output to jq:

```
sudo so-elasticsearch-query '*:so-*/_search' -d '{"query": {"match_all": {}},"size": 1}' | jq
```

19.6 so-import-pcap

`so-import-pcap` will import one or more pcaps into Security Onion and preserve original timestamps. It will do the following:

- generate IDS alerts using Suricata
- generate network metadata using Zeek
- store IDS alerts and network metadata in Elasticsearch with original timestamps
- store pcaps where SOC can find them
- provide a hyperlink for you to view all alerts and logs in SOC

> **Tip**
>
> You can run this command manually, but for most use cases it's easier to upload a PCAP via Grid and it will automatically run `so-import-pcap` for you.

19.6.1 Configuration

so-import-pcap requires you to run through Setup and choose a configuration that supports so-import-pcap. This includes Import Node and other nodes that include sensor services like Eval and Standalone. The quickest and easiest option is to choose Import Node which gives you the minimal services necessary to import a PCAP.

> **Warning**
>
> Please note that so-import-pcap is not supported on heavy nodes.

19.6.2 Usage

Once Setup completes, you can then run `sudo so-import-pcap` and supply the full path to at least one PCAP file. For example, to import a single PCAP named `import.PCAP`:

```
sudo so-import-pcap /full/path/to/import.PCAP
```

To import multiple pcaps:

```
sudo so-import-pcap /full/path/to/import1.PCAP /full/path/to/import2.PCAP
```

Please note that if you import multiple pcaps at one time, so-import-pcap currently only provides a hyperlink for the last PCAP in the list. If you need a hyperlink for each PCAP, then you can run one PCAP file per so-import-pcap and use a for-loop to iterate over your collection of PCAP files.

so-import-pcap calculates the MD5 hash of the imported PCAP and creates a directory in `/nsm/import/` for that hash. This is where so-import-pcap stores the alerts and logs generated by the traffic in the PCAP. If you try to import that same PCAP again, it will tell you that it has already imported that PCAP. If for some reason you really do need to import that PCAP again, you can remove that PCAP's directory in `/nsm/import/` and then try again.

19.6.3 Examples

If you don't already have some PCAP files to import, see PCAPs for a list of sites where you can download sample pcaps.

Our Quick Malware Analysis series at https://blog.securityonion.net/search/label/quick%20malware%20analysis uses so-import-pcap to import pcaps from https://www.malware-traffic-analysis.net/ and other sites. Following along with these blog posts in your own so-import-pcap VM is a great way to practice your skills!

19.7 so-import-evtx

`so-import-evtx` will import one or more evtx files into Security Onion.

> **Tip**
>
> You can run this command manually, but for most use cases it's easier to upload an evtx file via Grid and it will automatically run `so-import-evtx` for you.

19.7.1 Usage

Run `sudo so-import-evtx` and supply the full path to at least one evtx file. For example, to import a single evtx file named `import.evtx`:

```
sudo so-import-evtx /full/path/to/import.evtx
```

To import multiple evtx files:

```
sudo so-import-evtx /full/path/to/import2.evtx /full/path/to/import2.evtx
```

so-import-evtx then provides a hyperlink for you to view all logs in SOC.

19.8 so-monitor-add

If you've already run through Setup but later find that you need to add a new monitor (sniffing) interface, you can run `so-monitor-add`. This will allow you to add network interfaces to `bond0` so that their traffic is monitored.

> **Warning**
>
> Cloud images sniff directly from network interfaces rather than using `bond0` so this utility won't work in those environments.

19.9 so-status

To check the status of Security Onion services, you can either run `sudo so-status` or simply view the Status panel on the Grid page.

so-status reads the list of enabled services from `/opt/so/conf/so-status/so-status.conf` and checks the status of each. If you ever disable a service, you may need to remove it from that file.

19.9.1 Quiet Mode

so-status supports a quiet mode:

```
so-status -h
Usage: /usr/sbin/so-status [OPTIONS]
  Options:
    -h                  - Prints this usage information
    -q                  - Suppress output; useful for automation of exit code value
    -j                  - Output in JSON format
    -i                  - Consider the installation outcome regardless of whether the system appears healthy

  Exit codes:
    0                   - Success, system appears to be running correctly
    1                   - Error, one or more subsystems are not running
    2                   - System is starting
    99                  - Installation in progress
    100                 - System installation encountered errors
```

```
sudo so-status -q
echo $?
0
```

19.10 so-test

`so-test` will run `so-tcpreplay` to replay some PCAP samples to your sniffing interface.

Warning

Please note that this action could trigger alerts on other enterprise monitoring systems, depending on your network and interface configuration.

Warning

You will need to have Internet access in order to download the PCAP samples. Also, if you have a distributed deployment, make sure you run `so-tcpreplay` on the manager first to download the necessary Docker image.

```
so-test
Replay functionality not enabled; attempting to enable now (may require Internet access)...

Pulling so-tcpreplay image
================================================================================
Starting tcpreplay...

This could take a while if another Salt job is running.
Run this command with --force to stop all Salt jobs before proceeding.
================================================================================
local:
----------
    ID: so-tcpreplay
    Function: docker_container.running
    Result: True
    Comment: Created container 'so-tcpreplay'
    Started: 15:55:48.390107
    Duration: 1460.452 ms
    Changes:
    ----------
        container_id:
            ----------
            added:
                f035103cd8bf43134b56d4b19d77a0ae9e7c09fcb54ef6da67cf89bef5fc4019
        state:
            ----------
            new:
                running
            old:
                None

Summary for local
------------
Succeeded: 1 (changed=1)
Failed:    0
------------
Total states run:     1
Total run time:   1.460 s
Replaying PCAP(s) at 10 Mbps on interface bond0...
Actual: 111557 packets (12981286 bytes) sent in 10.38 seconds
Rated: 1249997.6 Bps, 9.99 Mbps, 10742.07 pps
Flows: 4102 flows, 394.99 fps, 2074477 flow packets, 45106 non-flow
Statistics for network device: bond0
    Successful packets:        55304
    Failed packets:            444
    Truncated packets:         0
    Retried packets (ENOBUFS): 0
    Retried packets (EAGAIN):  0
Replay completed. Warnings shown above are typically expected.
```

Once this completes, you can then go to Alerts, Dashboards, and Hunt to review data.

19.11 so-user

SOC user management should normally be done via Administration as shown in the Accounts section. However, if for some reason you can't log into SOC, you can use `so-user` from the command line to manage SOC user accounts.

`so-user` has many different operations. You can see them all by running `so-user` with no options:

```
sudo so-user
```

19.11.1 Listing SOC Users

To see a list of all SOC users, use the `list` operation:

```
sudo so-user list
```

19.11.2 Changing SOC User Password

If you've forgotten your password, you can reset it using the `password` operation:

```
sudo so-user password --email onionuser@example.com
```

Once you've reset your password, you should be able to log into SOC and go back to managing user accounts via Administration as shown in the Accounts section.

20. Help

20.1 Help Overview

Having problems? Try the suggestions below.

- Have you run soup to ensure that you're on the latest version?
- Check the FAQ.
- Search the Community Support forum.
- Search the documentation and support forums of the tools contained within Security Onion.
- Check log files in `/opt/so/log/` or other locations for any errors or possible clues:
- Setup `/root/sosetup.log`
- Suricata `/opt/so/log/suricata/suricata.log`
- Zeek `/nsm/zeek/logs/current/`
- Elasticsearch `/opt/so/log/elasticsearch/<hostname>.log`
- Kibana `/opt/so/log/kibana/kibana.log`
- Logstash `/opt/so/log/logstash/logstash.log`
- ElastAlert `/opt/so/log/elastalert/elastalert_stderr.log`
- Are you able to duplicate the problem on a fresh Security Onion installation?
- Check the Known Issues to see if this is a known issue that we are working on.
- If all else fails, please feel free to reach out for support.

20.2 FAQ

20.2.1 Table of Contents

20.2.2 Install / Update / Upgrade

Why won't the ISO image boot on my machine?

Please see the Help section.

What's the recommended procedure for installing Security Onion?

Please see the Installation section.

What's the recommended procedure for configuring Security Onion?

Please see the Configuration section.

What if I receive a `The IP being routed by Linux is not the IP address assigned to the management interface` error message?

Please see the warning about this in the Configuration section.

What languages are supported?

We only support the English language at this time.

How do I install Security Onion updates?

Please see the soup section.

What connectivity does Security Onion need to stay up to date?

Please see the Firewall section.

What do I need to do if I'm behind a proxy?

Please see the Proxy section.

Can I run Security Onion on Raspberry Pi or some other non-x86 box?

No, we only support x86-64 (standard Intel/AMD 64-bit architectures). Please see the Hardware section.

Back to Top

20.2.3 Users / Passwords

What is the password?

Please see the Passwords section.

How do I add a new user account?

Please see the Adding Accounts section.

Back to Top

20.2.4 Support / Help

Where do I send questions/problems/suggestions?

Please see the Community Support section.

Is commercial support available for Security Onion?

Yes, we offer commercial support at https://securityonionsolutions.com.

Back to Top

20.2.5 IDS engines

Can Security Onion run in `IPS` mode?

No, Security Onion does not support blocking traffic. Most organizations have some sort of Next Generation Firewall (NGFW) with IPS features and that is the proper place for blocking to occur. Security Onion is designed to monitor the traffic that makes it through your firewall.

Back to Top

20.2.6 Security Onion internals

Where can I read more about the tools contained within Security Onion?

Please see the Tools section.

What's the directory structure of `/nsm`?

Please see the Directory section.

Why does Security Onion use `UTC`?

Please see the Time Zones section.

Why are the `timestamps` in Kibana not in UTC?

Please see the Time Zones section.

Why is my disk filling up?

In general, Security Onion attempts to make use of as much disk space as you give it. Depending on installation type, it should continue writing data to disk until disk usage reaches 80-90% at which point it should begin purging old data. Most disk space is used by Elasticsearch or Full Packet Capture written to disk via Suricata.

How is my data kept secure?

Standard network connections to or from Security Onion are encrypted. This includes SSH, HTTPS, Elasticsearch network queries, and Salt minion traffic. All endpoint agent (Elastic Agent) traffic is encrypted except for binary updates, which are served from the Manager over http - these update files are cryptographically signed by Elastic and are verified before they are used. There is also the option to pull these updates via https directly from Elastic. SOC user account passwords are hashed via bcrypt in Kratos and you can read more about that at https://github.com/ory/kratos.

Back to Top

20.2.7 Tuning

How do I configure email for alerting and reporting?

Please see the Email section.

How do I configure a `BPF` ?

Please see the BPF section.

How do I filter traffic?

Please see the BPF section.

How do I exclude traffic?

Please see the BPF section.

What are the default firewall settings and how do I change them?

Please see the Firewall section.

What do I need to modify in order to have the log files stored on a different mount point?

Please see the New Disk section.

Back to Top

20.2.8 Common Problems

Why do containers go missing?

Docker containers that stop running, due to exiting from errors or other reasons, will be automatically removed by the scheduled cleanup process.

Most container logs are redirected to their application log directory, located in `/opt/so/log`. In some cases the logs may not get written to disk, and instead must be viewed via `docker logs <container-name>` before the container is cleaned up.

Why does ElastAlert often go missing on my Grid?

ElastAlert 2 will exit upon encountering syntax errors with rules or when Elasticsearch is not in a healthy state.

Why does Elasticsearch go to the unhealthy state?

Elasticsearch will become unhealthy for a variety of reasons, but the most common reasons are running out of disk space and having indices with unallocated shards.

Back to Top

20.2.9 Miscellaneous

Where can I find interesting pcaps to replay?

Please see the PCAPs section.

Why is Security Onion connecting to an IP address on the Internet over port 123?

Please see the NTP section.

Should I backup my Security Onion box?

Security Onion automatically backs up some important configuration as described in the Backup section. However, there is no automated data backup. Network Security Monitoring as a whole is considered "best effort". It is not a "mission critical" resource like a file server or web server. Since we're dealing with "big data" (potentially terabytes of full packet capture) of a transient nature, backing up the data would be prohibitively expensive. Most organizations don't do any data backups and instead just rebuild boxes when necessary.

What happened to Filebeat?

Filebeat has been replaced by Elastic Agent.

What happened to Grafana?

Grafana has been replaced by Grid.

What happened to Playbook?

Playbook has been replaced by Detections.

What happened to Wazuh?

Wazuh has been replaced by Elastic Agent.

What happened to Stenographer?

Stenographer has been replaced by Suricata full packet capture.

How can I add local rules?

Please see the Detections section.

Can I connect Security Onion to Active Directory or another OIDC provider?

Please see the OIDC section.

Back to Top

20.3 Directory Structure

20.3.1 /opt/so/conf

Applications read their configuration from `/opt/so/conf/`. However, please keep in mind that most config files are managed with Salt, so if you manually modify those config files, your changes may be overwritten at the next Salt update.

20.3.2 /opt/so/log

Debug logs are stored in `/opt/so/log/`.

20.3.3 /opt/so/rules

ElastAlert and Suricata rules are stored in `/opt/so/rules/`.

20.3.4 /opt/so/saltstack/local

Custom Salt settings can be added to `/opt/so/saltstack/local/`.

20.3.5 /nsm

The vast majority of data is stored in `/nsm/`.

20.3.6 /nsm/zeek

Zeek writes its protocol logs to `/nsm/zeek/`.

20.3.7 /nsm/elasticsearch

Elasticsearch stores its data in `/nsm/elasticsearch/`.

20.3.8 /nsm/suripcap

Suricata stores full packet capture in `/nsm/suripcap/`.

20.4 Community Support

20.4.1 Check Documentation First

First, check to see if your question has already been answered in the Help or FAQ sections.

20.4.2 Forum Guidelines

Before posting, please review the forum guidelines at https://github.com/Security-Onion-Solutions/securityonion/discussions/1720.

20.4.3 Forum

Once you've read and understand all of the above, you can post your question to the community support forum at https://securityonion.net/discuss.

20.5 Support

20.5.1 Paid Support

If you need private or priority support, please consider purchasing hardware appliances or support from Security Onion Solutions:

https://securityonionsolutions.com/support

> **Tip**
>
> Purchasing from Security Onion Solutions helps to support development of Security Onion as a free and open platform!

20.5.2 Community Support

If you need free support, you can reach out to our Community Support.

20.6 Help Wanted

Folks frequently ask how they can give back to the Security Onion community. Here are a few of our community teams that you can help with.

20.6.1 Marketing Team

We need more folks to help spread the word about Security Onion by blogging, tweeting, and other social media.

20.6.2 Support Team

If you'd like help out other Security Onion users, please join the forum and start answering questions!

https://securityonion.net/discuss

20.6.3 Documentation Team

If you find that some information in our Documentation is incorrect or lacking, please feel free to submit Pull Requests to the `3/dev` branch of the `docs` repo at https://github.com/Security-Onion-Solutions/docs.

20.6.4 Core Development

Most of our code is on GitHub. Please feel free to submit pull requests!

https://github.com/Security-Onion-Solutions

20.6.5 Thanks

The following folks have made significant contributions to Security Onion over the years. Thanks!

- Lawrence Abrams
- Jack Blanchard
- Kevin Branch
- Josh Brower
- Pete Di Giorgio
- Dennis Distler
- Jason Ertel
- Seth Hall
- Paul Halliday
- Joe Hargis
- Mark Hillick
- Wes Lambert
- Dustin Lee
- Josh More
- Corey Ogburn
- Eric Ooi
- Josh Patterson
- Phil Plantamura
- Liam Randall
- Mike Reeves
- Jorge Reyes
- Scott Runnels
- Jon Schipp
- Brad Shoop
- Bryant Treacle
- William Wernert
- Matthew Wright

21. Security Onion Pro

21.1 Security Onion Pro Overview

Note

Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable these features.

21.2 OpenID Connect (OIDC)

SOC supports single sign-on (SSO) authentication via OpenID Connect (OIDC) to one of several OIDC-compatible identity providers. For example, users can login to Security Onion using an Active Directory user, a GitHub user, a Google account, an Auth0 account, etc. Only one OIDC provider can be active at a time.

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

> **Warning**
>
> LDAP and SAML integrations are not supported.

> **Warning**
>
> Integrating Security Onion into an organization's global identity management platform is generally not recommended. If an attacker compromises the identity management platform, which is typically a high priority target, then that attacker could use compromised SSO credentials to access Security Onion and potentially undermine the benefits provided by Security Onion. This integration is made available for those who understand these risks and have appropriate mitigations in place.

21.2.1 Configuration

OIDC configuration can be complex and we recommend taking advantage of the official Security Onion support team. Note that purchases of a Security Onion license include some level of support. This will help avoid time-consuming problems that can occur when configuring OIDC.

The first step in configuration OIDC is to determine which provider the grid will use, and collecting the required configuration inputs necessary for that specific provider.

Next, in Security Onion Console, while logged in as an administrator, navigate to the Administration -> Configuration screen and enter `oidc` into the filter field. Then click the *Expand All* icon.

Review the following instructions for the applicable provider.

MICROSOFT ENTRA ID (AZURE ACTIVE DIRECTORY)

Locate the `provider` setting in the SOC configuration screen. Specify the value `microsoft` for this setting.

In a separate browser tab, login to the Microsoft Azure account you plan to use for the integration. Navigate to the Microsoft Entra ID service and find the `Tenant ID`, which will resemble a UUID similar to `abcdef12-1234-abcd-5678-a1b2c3d4e5f6`.

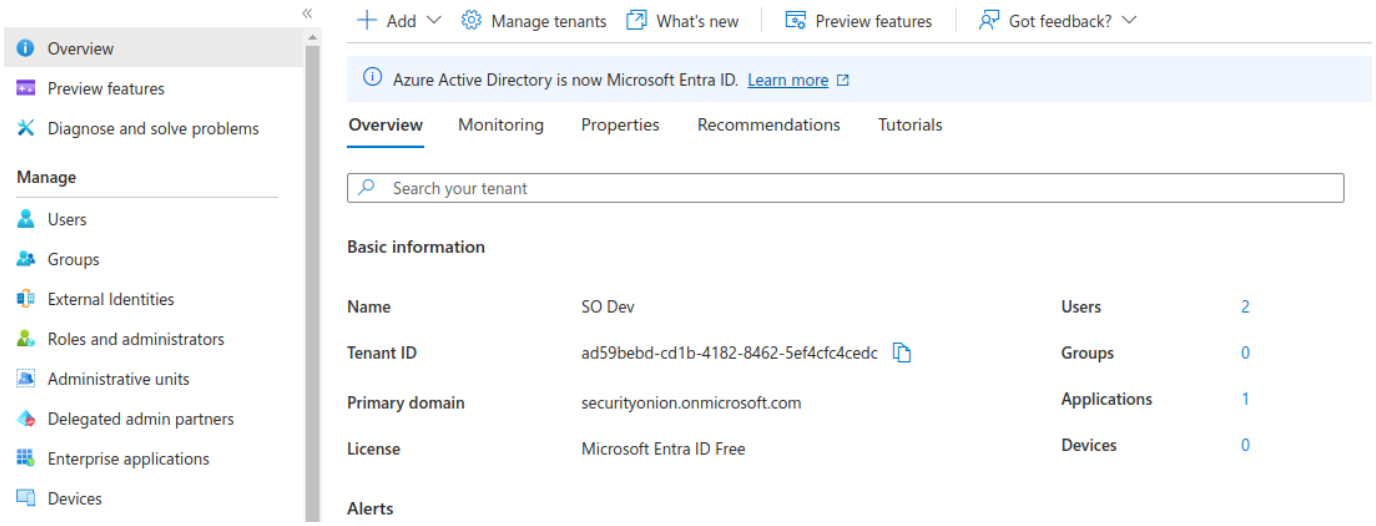

Locate the `microsoft_tenant` setting in the SOC configuration screen back on the SOC browser tab. Specify the UUID value for this setting.

Back in the Azure tab, under the desired Azure Tenant, register a new App named `Security Onion`. Most organizations will only desire organization accounts to have access to Security Onion so be sure to choose the correct account type option. Failure to choose this correctly could expose your Security Onion installation to users outside of your organization. Specify the application as Web, with Redirect URI using the URL that the analysts will use to access SOC after finalizing their login to Azure. This is typically going to resemble the following pattern: `https://<my-SOC-base-url>/auth/self-service/methods/oidc/callback/SSO`. Omit the `/SSO` suffix if forcing PKCE (Proof Key Code Exchange). Click *Register*, and on the resulting screen find the application ID for this new app registration. It will also resemble a UUID.

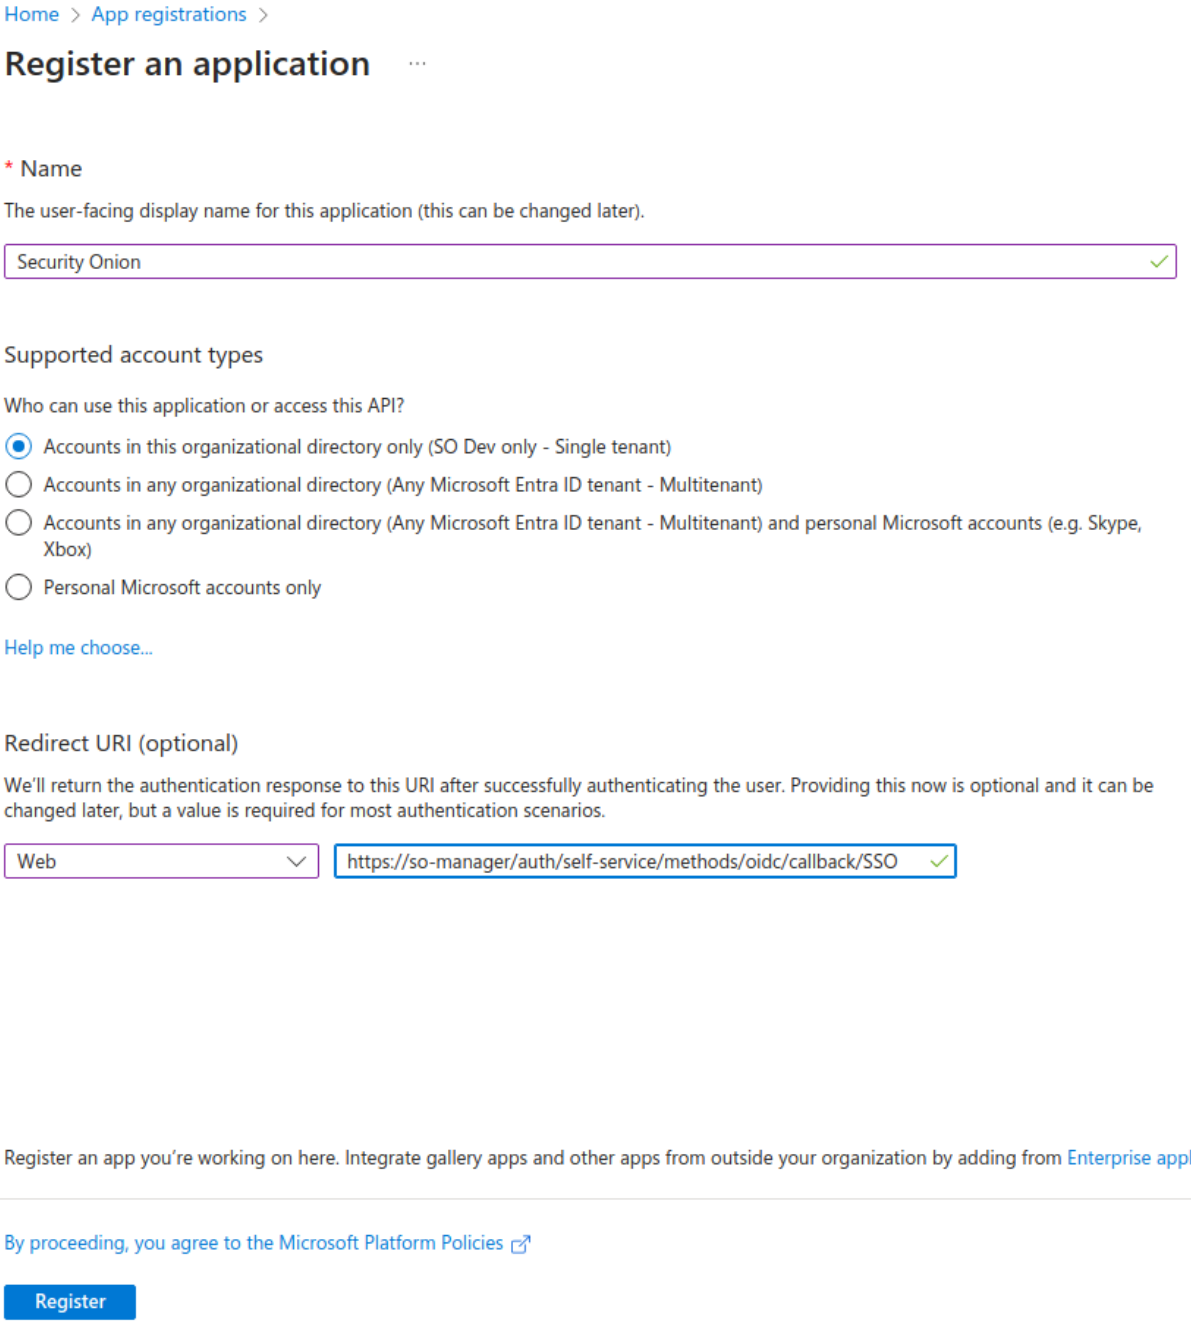

Locate the `client_id` setting in the SOC configuration screen back on the SOC browser tab. Specify the above application ID for this setting.

Add a new client secret to the app registration created above. Specify the secret name as `so-oidc` and choose the expiration that makes the most sense for your organization. If you choose to use a short or medium term expiration, a good practice is to make it very short, so that it forces your rotation processes to be well-known and documented. Choosing a medium expiration of two years will likely cause more trouble when the secret expires and the knowledge of how to resolve it is lost among the administrative team. Copy the generated secret to your clipboard. You will only have this one chance to copy the secret. Returning to this secret later will not provide access to the original secret.

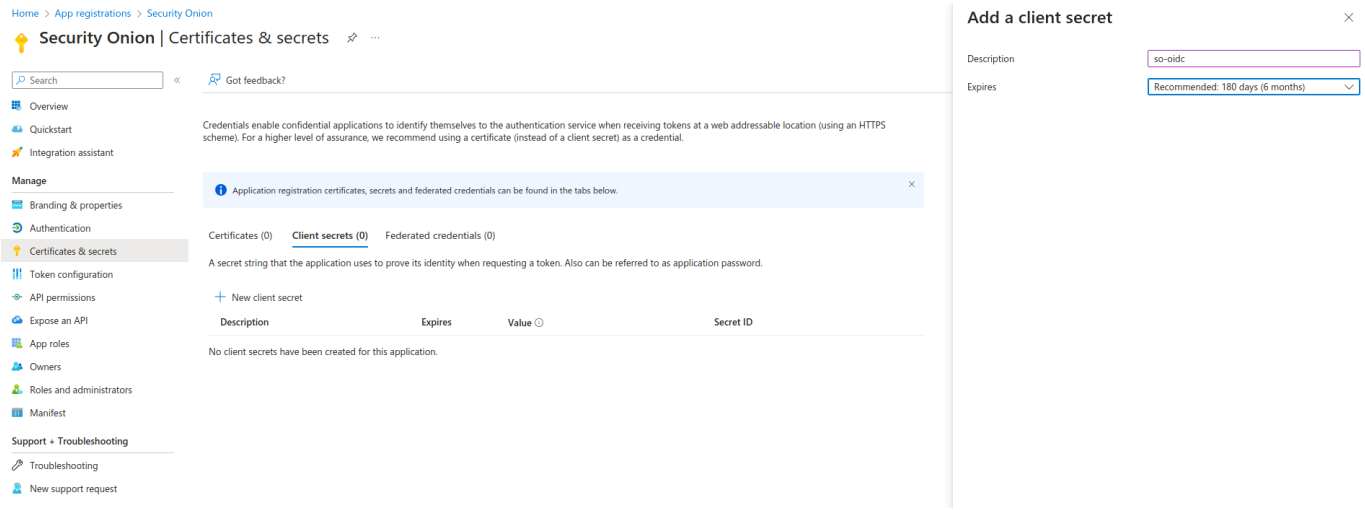

Locate the `client_secret` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client secret for this setting.

If forcing PKCE (Proof Key Code Exchange) to be enabled, set the `pkce` setting in the SOC configuration screen to `force`. When forcing PKCE, the redirect URI should omit the trailing `/SSO` suffix.

Next, skip to the *Enabling OIDC* section to enable the newly configured OIDC authentication.

ACTIVE DIRECTORY (SELF-HOSTED)

Contact the Security Onion Solutions support team to determine the specific configuration changes required to integrate your Security Onion Grid with your organization's Active Directory installation. They will review your current Windows Server version, assist with TLS certificate configurations applicable to your organization, and walk you through the steps needed to complete the integration.

Integration with on-premise Active Directory has several prerequisites:

- Fully-functioning installation of Active Directory on Windows Server 2022
- Administrator access to Windows Active Directory server
- Administrator access to Security Onion manager via SSH and SOC
- Ability to obtain TLS certificates that are trusted on client and SO systems
- Ability to activate Active Directory Federated Services (ADFS) on Active Directory server
- HTTPS access from Security Onion analyst browsers to the ADFS server.
- HTTPS access from the Security Onion manager to the ADFS server

GOOGLE

Locate the `provider` setting in the SOC configuration screen. Specify the value `google` for this setting.

In a separate browser tab, login to the Google Cloud Console and select or create a project under your Google organization containing the workspace users you plan to use for the integration. Navigate to the Credentials screen and add a new OAuth 2.0 Client ID named `Security Onion`, of type `Web`. Specify the web Redirect URI using the URL that the analysts will use to access SOC after finalizing their login to Azure. This is typically going to resemble the following pattern: `https://<my-SOC-base-url>/auth/self-service/methods/oidc/callback/SSO`.

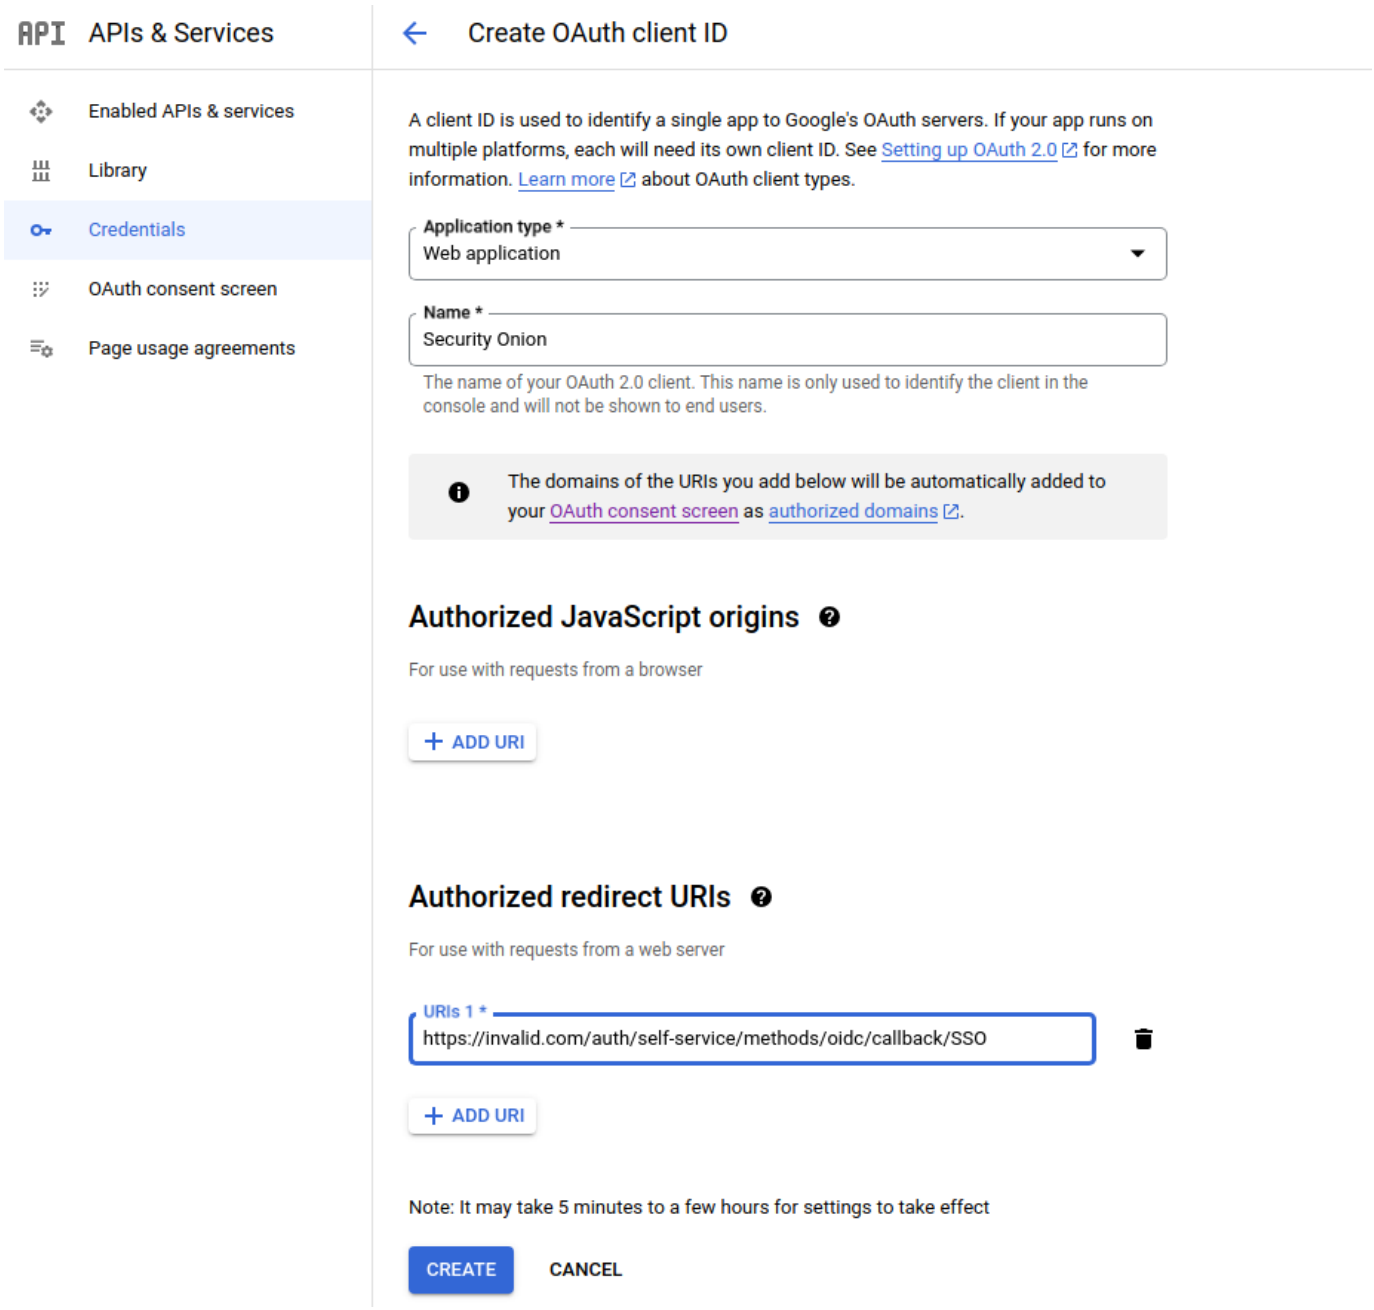

When the client ID is added a popup will appear with the new Client ID and Secret. These values must be entered into the SOC configuration screen.

Locate the `client_id` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client ID for this setting.

Locate the `client_secret` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client secret for this setting.

Next, skip to the *Enabling OIDC* section to enable the newly configured OIDC authentication.

GITHUB

Locate the `provider` setting in the SOC configuration screen. Specify the value `github` for this setting.

In a separate browser tab, login to the GitHub account you plan to use for the integration. Navigate to the Organization Settings and then Developer Settings -> OAuth Apps. Click the *New Org OAuth App* button. Enter `Security Onion` for the Application name, the login URL to your SOC Grid for the Homepage URL, and optional description, and then the authorization callback URL, which will resemble the following pattern: `https://<my-SOC-base-url>/auth/self-service/methods/oidc/callback/SSO`. Click *Register Application*.

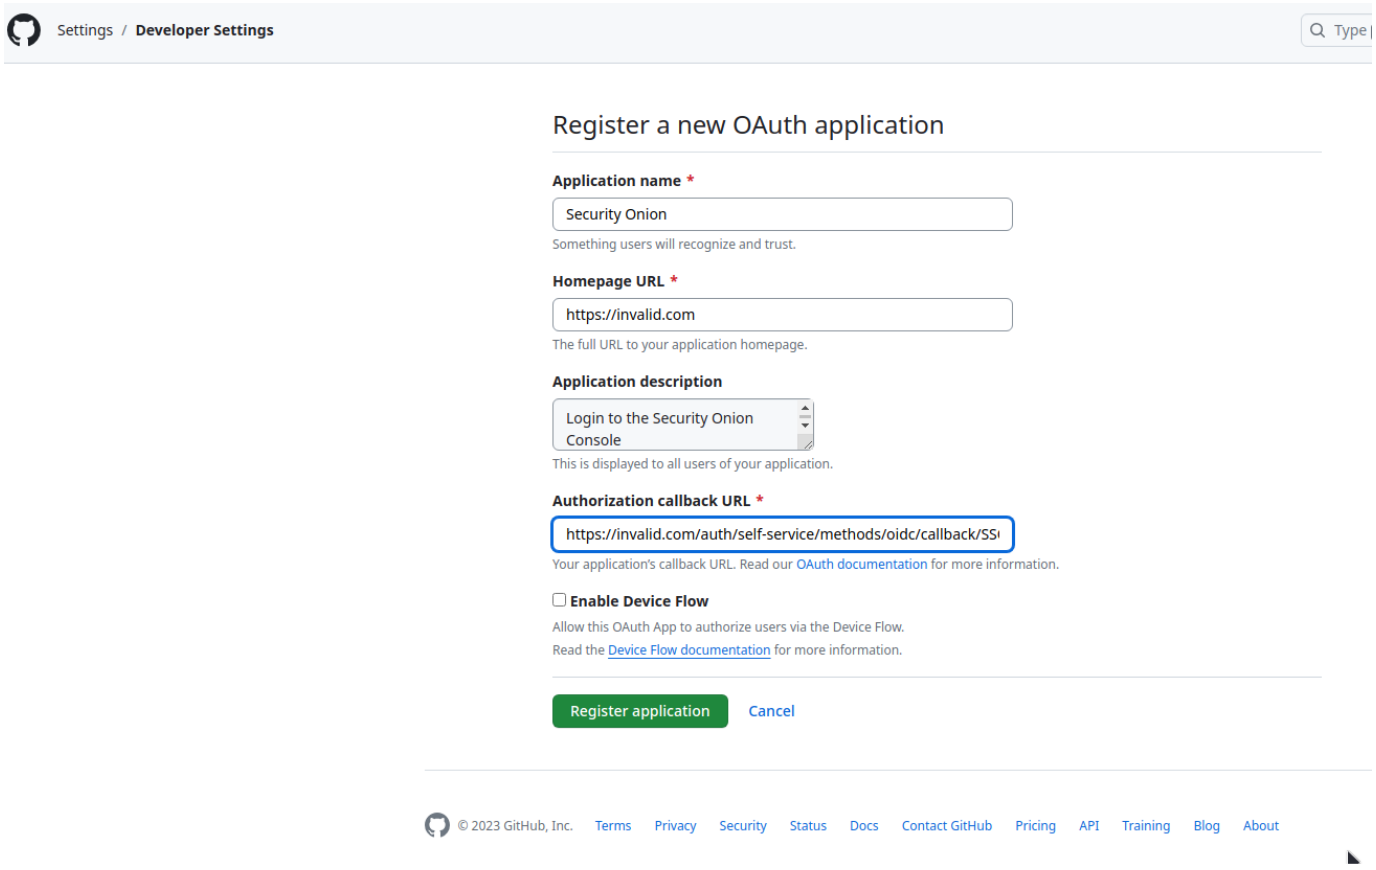

Once the app is created a new screen will show the newly create OAuth application settings, including the generated client ID and secret.

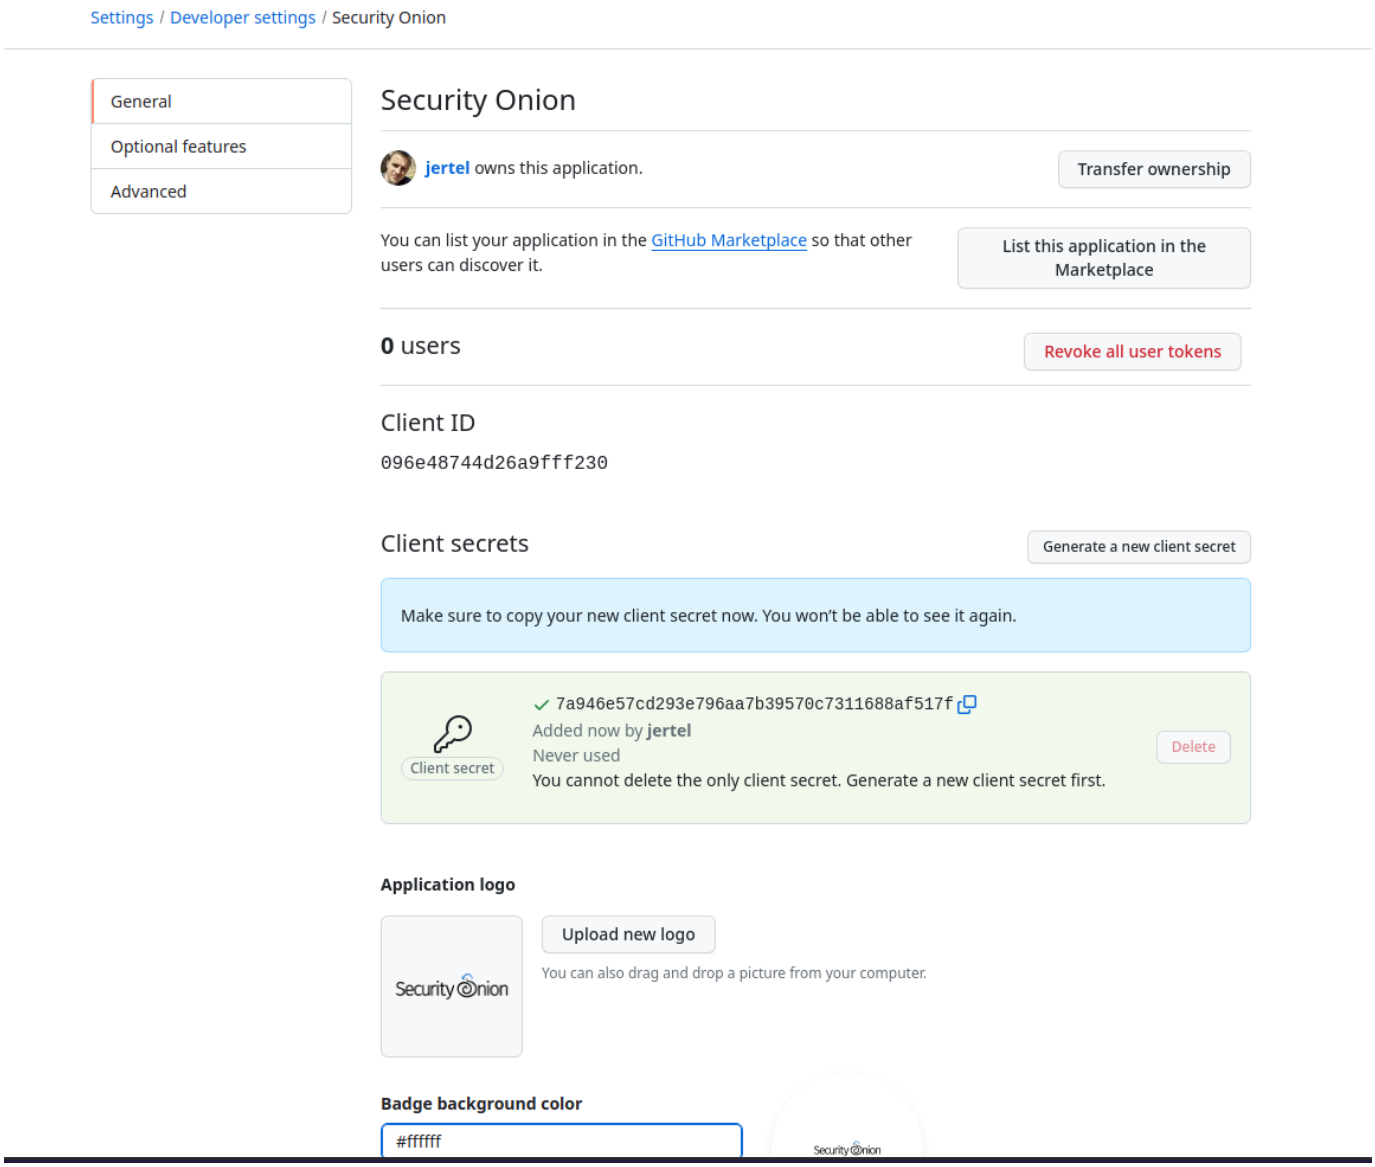

Be sure to copy the secret before refreshing or navigating away from this screen. These two values must be entered into the SOC configuration screen.

Locate the `client_id` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client ID for this setting.

Locate the `client_secret` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client secret for this setting.

Next, skip to the *Enabling OIDC* section to enable the newly configured OIDC authentication.

AUTH0

Locate the `provider` setting in the SOC configuration screen. Specify the value `auth0` for this setting.

In a separate browser tab, login to the Auth0 account you plan to use for the integration. Create a new application named `Security Onion`. After it's created, navigate to the Settings tab. Scroll down to the Application URIs section and enter `https://<my-SOC-base-url>` for the Application Login URI and Logout URL, and then enter the callback URL, which will resemble the following pattern: `https://<my-SOC-base-url>/auth/self-service/methods/oidc/callback/SSO`. Click *Save Changes*.

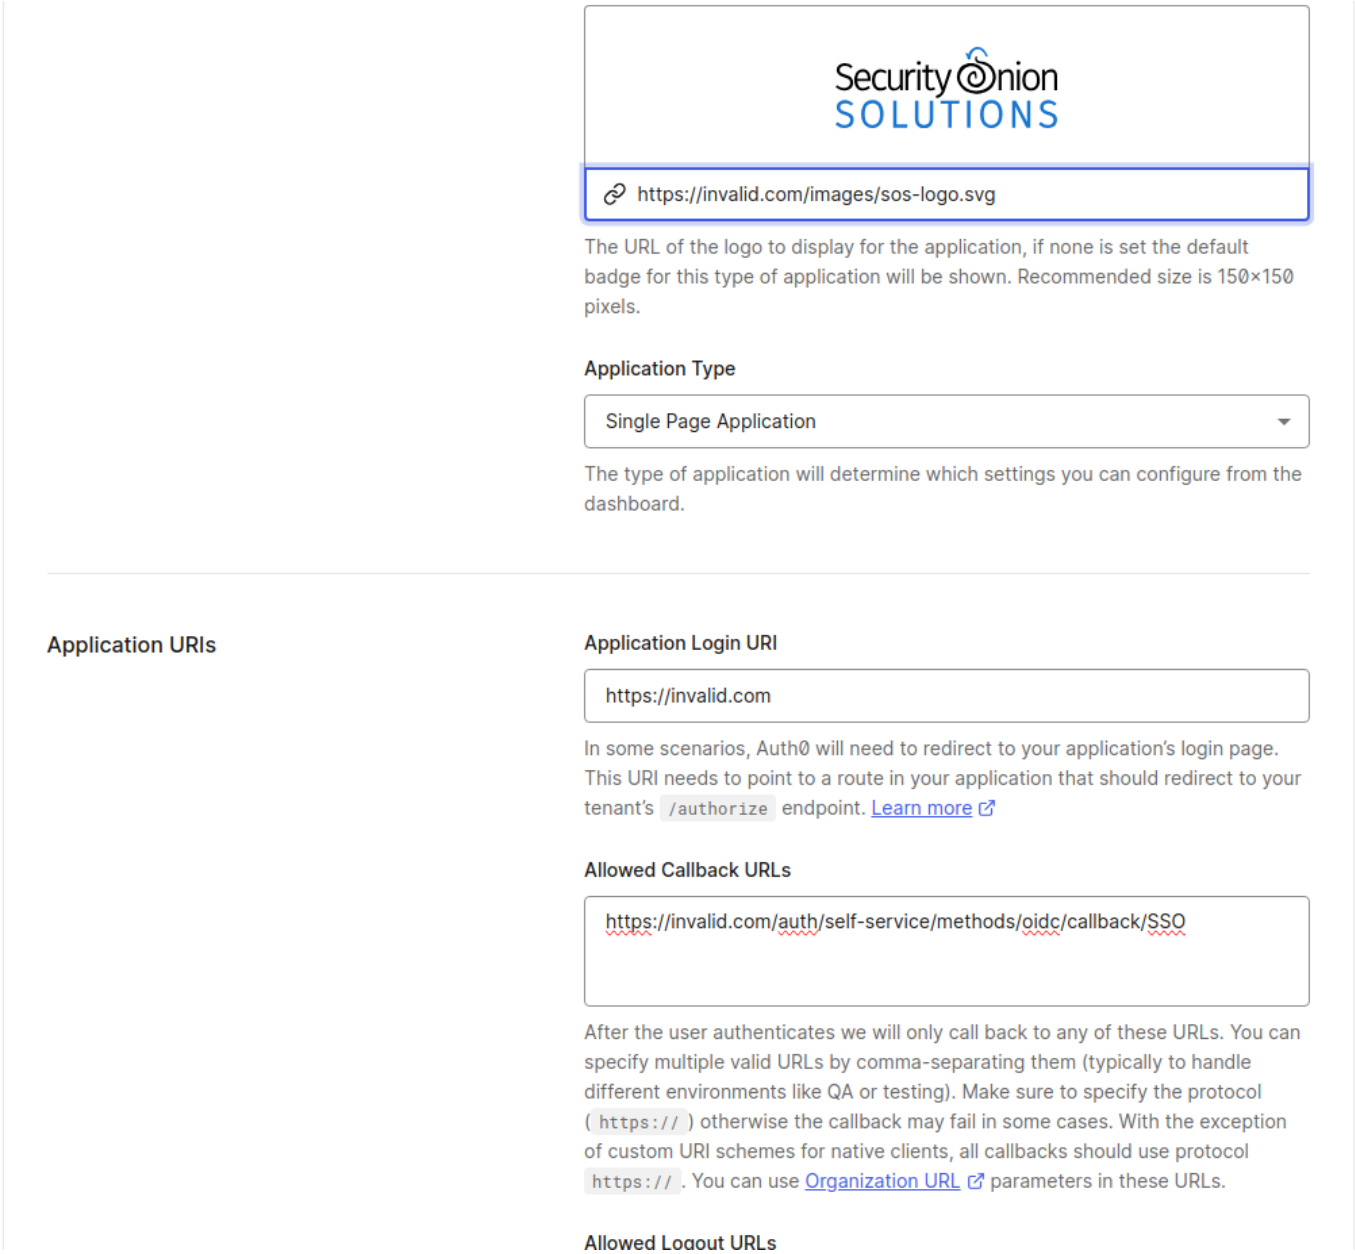

Scroll back to the top of the Auth0 Settings page where the Client ID and Secret are shown.

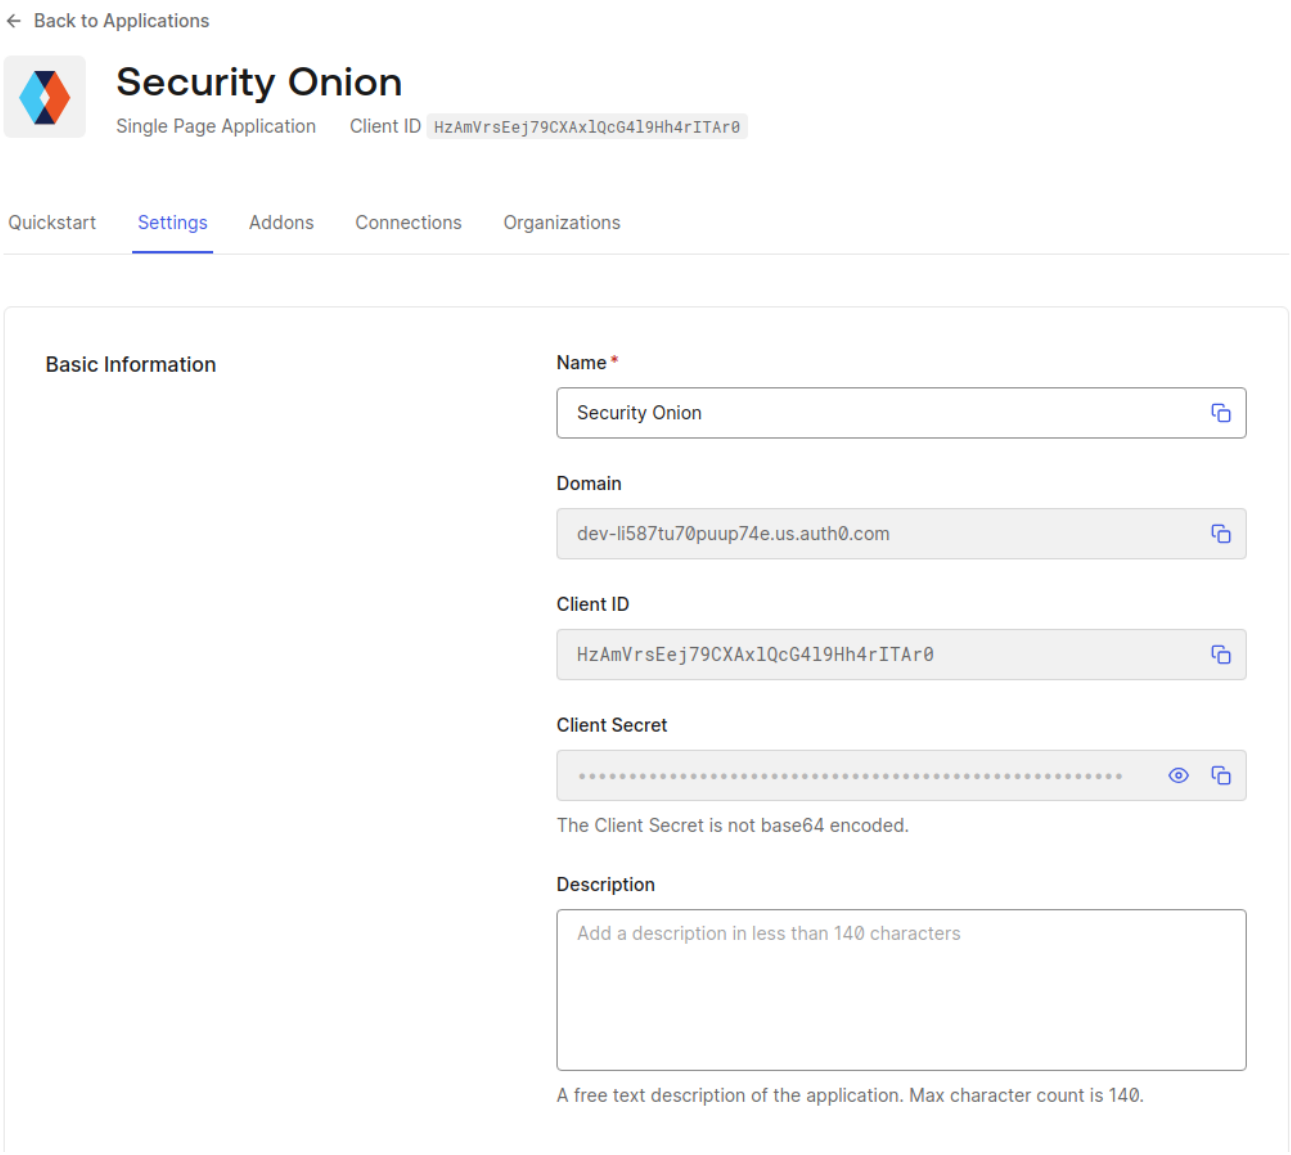

Be sure to copy the secret before refreshing or navigating away from this screen. These two values must be entered into the SOC configuration screen.

Locate the `client_id` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client ID for this setting.

Locate the `client_secret` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client secret for this setting.

Back in the Auth0 tab, scroll down to the Advance Settings section, and click on *Endpoints*. Copy the OAuth Authorization URL, but without the `/authorize` path. Locate the `issuer_url` setting in the SOC configuration screen back on the SOC browser tab. Paste the copied URL into this setting. It should resemble the following: `https://dev-xyz123abc456.us.auth0.com`

Next, skip to the *Enabling OIDC* section to enable the newly configured OIDC authentication.

GENERIC (EX: PING)

Security Onion can work with most OIDC providers, even if not mentioned as an explicitly-supported provider above. To show an example of how to configure a generic provider the below instructions will show how `[Ping SSO](https://pingone.com)` can be used as an OIDC provider for Security Onion.

Locate the `provider` setting in the SOC configuration screen. Specify the value `generic` for this setting.

In a separate browser tab, login to a Ping Identity console and, under the desired workspace environment, create a new application called Security Onion. Choose the `OIDC Web App` and click Save. On the Configuration tab, specify a Redirect URI using the following pattern: `https://<my-SOC-base-url>/auth/self-service/methods/oidc/callback/SSO`.

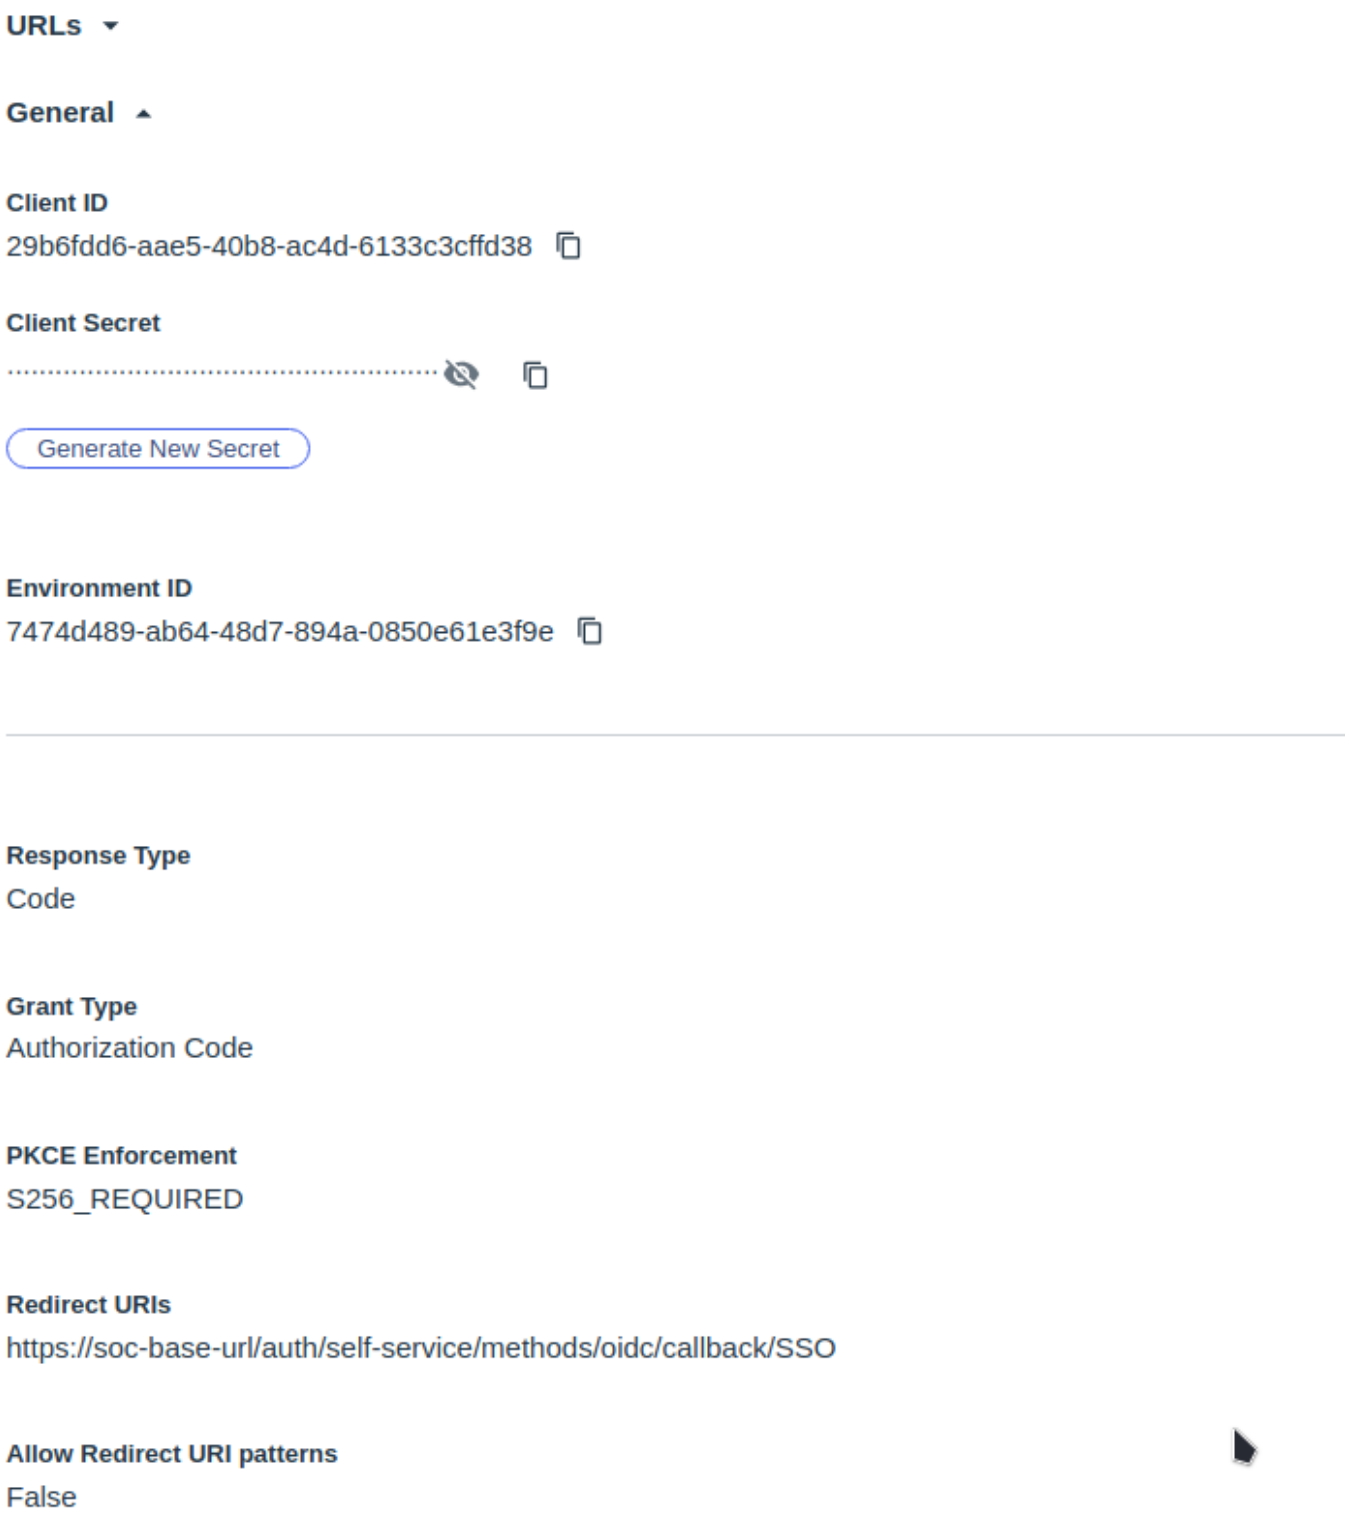

Locate the `client_id` setting in the SOC configuration screen back on the SOC browser tab. Specify the Ping `Client ID` for this setting. That ID can be located on the Ping configuration tab, as shown above.

Generate a new client secret under the Ping `Client ID` field. Copy the generated secret to your clipboard.

Locate the `client_secret` setting in the SOC configuration screen back on the SOC browser tab. Specify the above client secret for this setting.

On the Ping console browser tab, under the configuration tab, expand the URLs section, near the top. Copy and paste the three following URLs into the appopriate SOC configuration screen settings:

- Authorization URL -> auth_url
- Issuer -> issuer_url
- Token Endpoint -> token_url

If forcing PKCE (Proof Key Code Exchange) to be enabled, set the `PKCE Enforcement` setting in the Ping console's configuration tab to `S256_REQUIRED`.

In the Ping console browser tab, navigate to the Resources tab and add `email` as an additional scope.

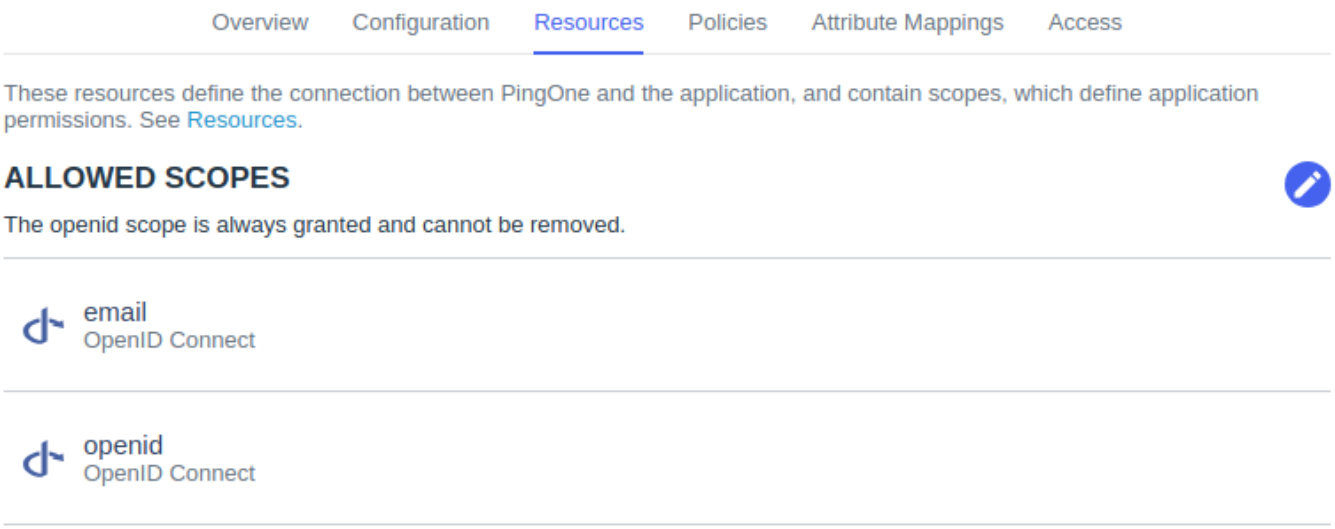

Locate the `scope` setting in the SOC configuration screen back on the SOC browser tab. Change the default `profile` scope to `openid`. There should now be both `email` and `openid` scopes listed.

Next, skip to the *Enabling OIDC* section to enable the newly configured OIDC authentication.

ENABLING OIDC

Finally, enable OIDC by locating the `enabled` setting in the SOC configuration and specify the value of `true` for this setting.

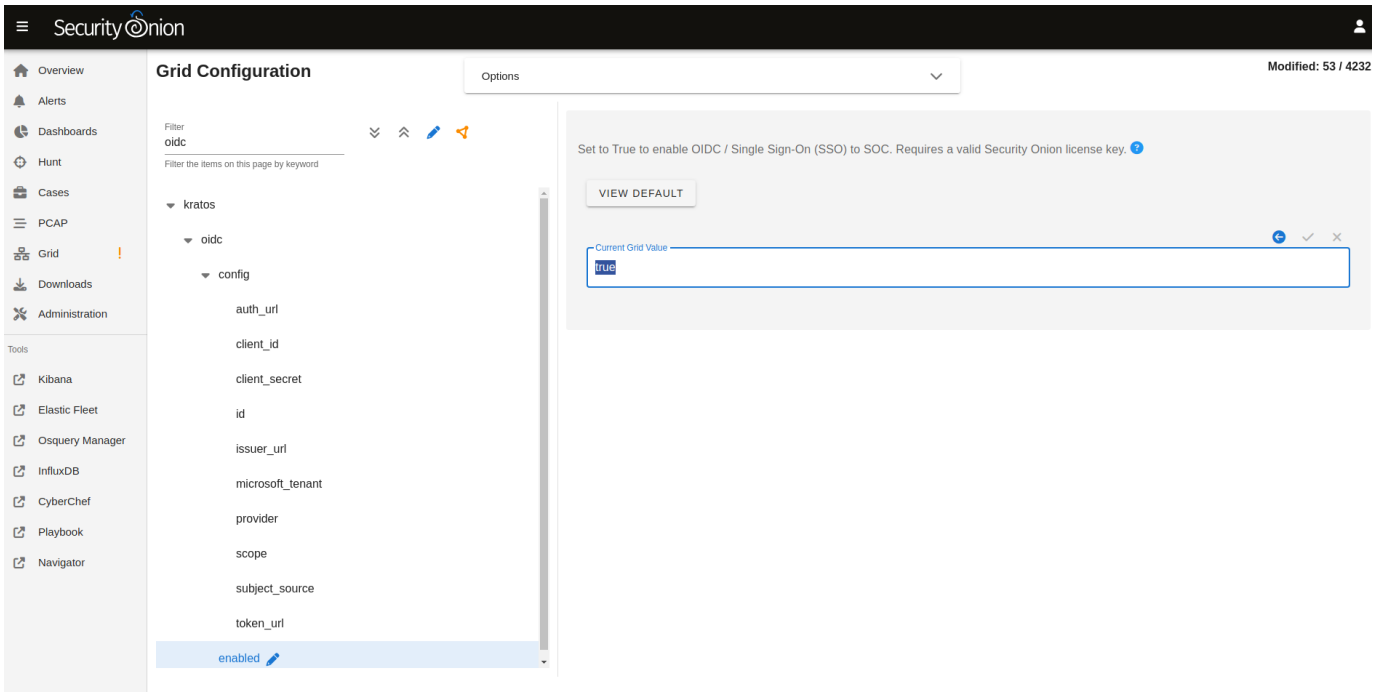

Note

Do not enable OIDC until all required configuration settings have been entered and double-checked for accuracy. Once enabled the backend system will automatically synchronize the settings across the grid, typically within 15 minutes. If some settings are incorrect or missing the backend authentication services could be left in an error state and make it impossible to fix via the Configuration screen, as the SOC UI may no longer be accessible. If this occurs an SSH session will be required to access the underlying configuration files on the manager node. Contact support for assistance if needed.

Warning

Once OIDC is enabled, any user of the selected external identity provider will be able to login to SOC, provided they have network access to do so. However, once logged in the new user will have no assigned roles and cannot view or modify sensitive SOC data. See the *Roles* section below for more information.

21.2.2 Initial Login

Upon the first login via OIDC the user will likely be returned back to the login screen. However, clicking on the *Continue with* the second time will take the newly linked user to the SOC interface. This additional login click is only required once.

21.2.3 Roles

When a new OIDC user logs into SOC, that user will not be assigned any roles, unless a default system role has been configured. This greatly limits what functions the user will be capable of performing within SOC. For example, new users will be unable to see any alerts, hunt for events, view dashboard data, view or create cases, manage the grid, or view other users. Attempting to view those role-protected screens will result in an error message.

If a default system role is not configured, an administrator will need to login to SOC and assign roles to OIDC users via the Adminstration -> Users screen. This is a one time operation, per user.

21.2.4 Managing OIDC Users

Users created via an OIDC login should not have their credentials managed within SOC. When an administrator views an OIDC user in the Administration -> Users screen, they will notice a message appears near Access Control panel, and cautions them against changing authentication settings for that user.

> **Note**
>
> Authentication relates to obtaining access to a system, whereas authorization relates to permissions a user has within the system. While *authentication* settings of OIDC users should not be managed within SOC, *authorization* settings can be managed within SOC for OIDC users. See the *Roles* section above for more information about granting roles to OIDC users.

21.2.5 OIDC Self Service

Users will continue to have access to their own Security Settings via the User Settings -> Security screen. A user could set a local SOC password via this screen, which would allow logins to SOC for that user without using SSO. After setting a local password, a user could then unlink the SSO account, which would disallow the user from logging in via SSO but still allow the user to login via the local password.

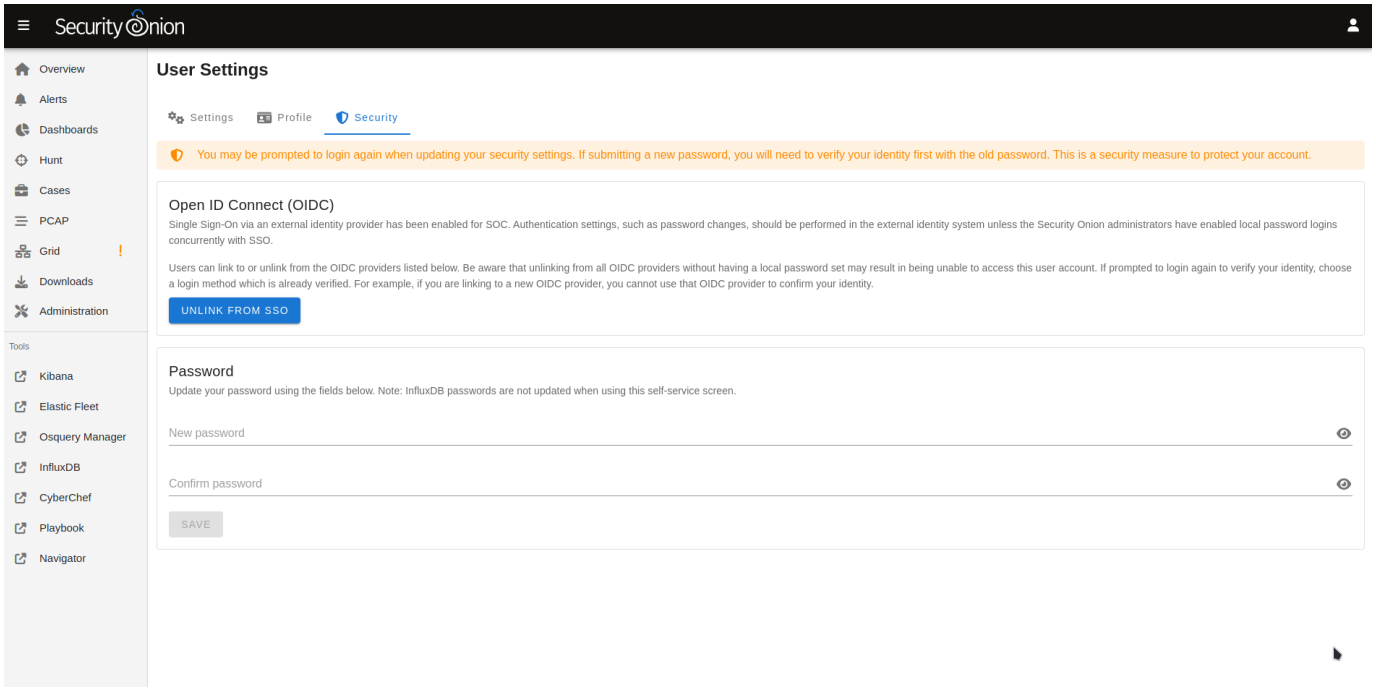

Conversely, locally logged in users that have not logged in via SSO yet can link to their SSO user.

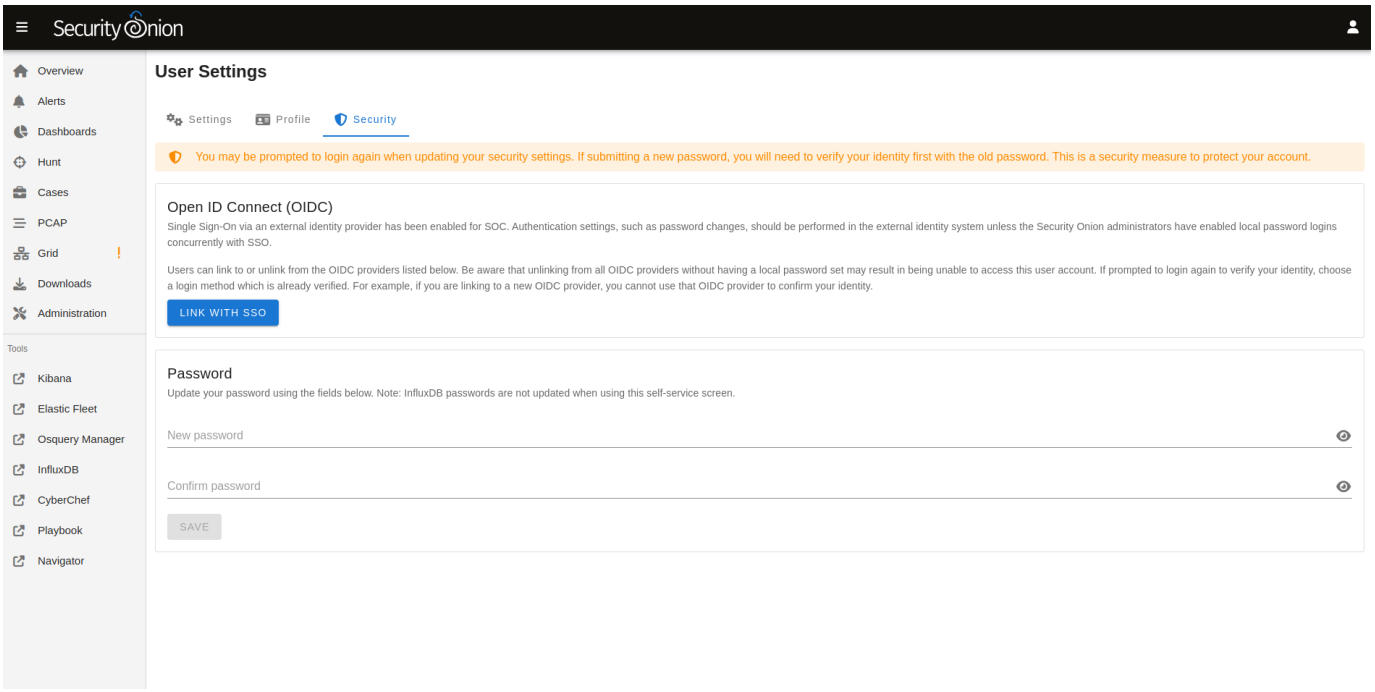

If you would like to ensure that all logins go through the external OIDC provider, then you can disable password logins. In SOC, navigate to Administration --> Configuration. At the top of the page, click the `Options` menu and then enable the `Show advanced settings` option. Then filter for `password.enabled` to locate the setting.

Similarly, the TOTP MFA and Passwordless options can also be disabled, if there is a desire to prevent users from altering all local authentication methods. Search for `totp.enabled` and `webauthn.enabled`, respectively, to disable those authentication methods.

When all local authentication methods have been disabled, users will have no security settings to modify in their self-service screen:

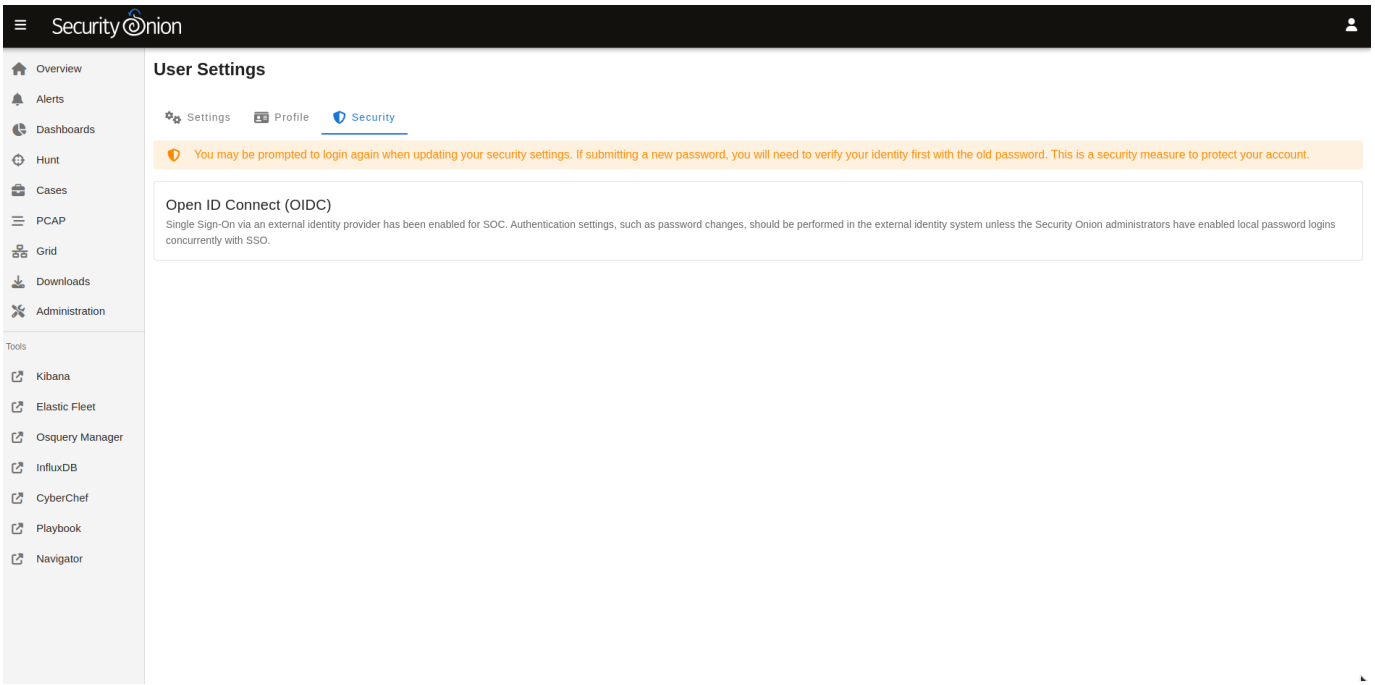

21.2.6 External Tools

Tools included with Security Onion, but provided by other vendors, will not utilize SOC single sign-on. This includes tools such as InfluxDB, Kibana and other Elastic-provided tools. If users need to access these tools the password authentication method must be enabled and a local password setup. The users can then login to those tools using their SSO email address and the local SOC password.

21.3 LUKS

LUKS stands for Linux Unified Key Setup and you can read more about it at https://en.wikipedia.org/wiki/Linux_Unified_Key_Setup.

LUKS disk encryption is a feature that requires the use of the Security Onion Pro license.

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.3.1 Enabling LUKS During the ISO Install

The recommended way to use LUKS with Security Onion is to install via our Security Onion ISO image. At the ISO boot menu, you'll need to modify the boot command. This can be done using the `e` key in UEFI boot mode or the `Tab` key in BIOS boot mode before it automatically boots. Then append `luks=1` (and possibly other options like FIPS and STIG) to the boot parameters and continue the boot process.

21.3.2 LUKS Install without a TPM

During the initial install of the ISO, the user will be prompted to enter a password to use to encrypt the LUKS partitions. If multiple drives are detected then the user has the option of just encrypting /nsm. Please note that this password will be required at each boot.

21.3.3 LUKS Install Options with a TPM

If a TPM module is installed in the system you will have the option of storing the key in the TPM to unlock the drives automatically at boot. This process uses clevis in order to manage this process.

21.3.4 LUKS TPM enrollment / re-enrollment

There may be a case where you have already installed Security Onion with LUKS enabled, but did not opt-in to use your TPM for automatic decryption at boot. In this case, you can use the `so-luks-tpm-regen` script to enroll the TPM and configure it for automatic decryption.

SSH to the Security Onion node and run the following command:

```
sudo so-luks-tpm-regen --enroll-tpm
```

Similarly, if for any reason automatic decryption was previously enabled using the ISO and has now stopped working you can re-enroll the TPM.

```
sudo so-luks-tpm-regen
```

21.4 FIPS

FIPS stands for Federal Information Processing Standards and you can read more about it at https://en.wikipedia.org/wiki/Federal_Information_Processing_Standards.

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.4.1 Enabling FIPS During the ISO Install

The recommended way to use FIPS with Security Onion is to install via our Security Onion ISO image. At the ISO boot menu, you'll need to modify the boot command. This can be done using the `e` key in UEFI boot mode or the `Tab` key in BIOS boot mode before it automatically boots. Then append `fips=1` (and possibly other options like LUKS and STIG) to the boot parameters and continue the boot process.

21.5 STIG

STIG stands for Security Technical Implementation Guide. For more information about STIGs, please see https://public.cyber.mil/stigs/.

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.5.1 STIG During the ISO Install

The recommended way to use STIG with Security Onion is to install via our Security Onion ISO image. From the installation menu you'll select the `Install Security Onion Pro` option.

Installing using the Security Onion Pro menu options will create additional partitions on your system to meet the STIG requirements. The partitions created include:

Partition	Storage
/home	25GB
/tmp	2GB
/var	50GB
/var/log	5GB
/var/log/audit	2GB
/var/tmp	2GB

In addition to the required partitions, using the STIG menu option will also configure the system to use FIPS mode, and enable LUKS disk encryption. Both of these options can be used independently of the STIG menu option depending on your requirements.

21.5.2 Enabling STIG

> **Warning**
>
> Before enabling STIGs on your production Security Onion deployment, we recommend testing in a development environment. With different environments and configurations, there may be unexpected errors.

To enable STIGs you'll first need setup your Security Onion Grid and apply your Security Onion Pro license. You can then navigate to Administration --> Configuration --> stig --> enabled and set the value to `true`.

> **Note**
>
> You will need to enable the Administration --> Show advanced settings option to modify this setting.

21.5.3 OpenSCAP

In order to apply STIGs on Security Onion we use a combination of our existing Saltstack configuration managment and OpenSCAP. Currently, OpenSCAP is using a draft version of STIGs for Oracle Linux 9.

OpenScap can be configured to run at different time intervals. By default, OpenSCAP will run a remediation every 12 hours meaning any changes made to the system that bring it out of compliance will be reverted back to the STIG compliant state. This setting can be lowered or increased by modifying the `run_interval` setting found under Administration --> Configuration --> stig

With the STIG feature enabled, you can find OpenSCAP reports under `/opt/so/log/stig`. Currently, the expected compliance score is 86%.

21.5.4 More information

For more information about OpenSCAP see: https://www.open-scap.org/.

For more information about STIGs see: https://public.cyber.mil/stigs/.

21.6 Notifications

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

The Detections module, specifically Sigma rules, can be enabled to send outbound notifications upon an alert being created. By default, no outbound notifications are enabled in a Security Onion installation. However, with the Pro license applied to a grid, notifications can be quickly configured via the Configuration screen.

21.6.1 Configuration

Configuring notifications involves adjusting configuration in two areas:

1. ElastAlert 2 Alerters
2. SOC Detections

ElastAlert 2 Alerters

ElastAlert includes a large number of alerters that can reach out to remote systems to deliver notifications. As each alerter supports a unique protocol the alerter requires its own set of supporting parameters in order for the alerter to know how to reach out to the remote endpoint. For example, to send a notification to a Slack channel, a webhook URL must be provided.

Navigate to the Administration -> Configuration screen. Next, locate the `ElastAlert` settings.

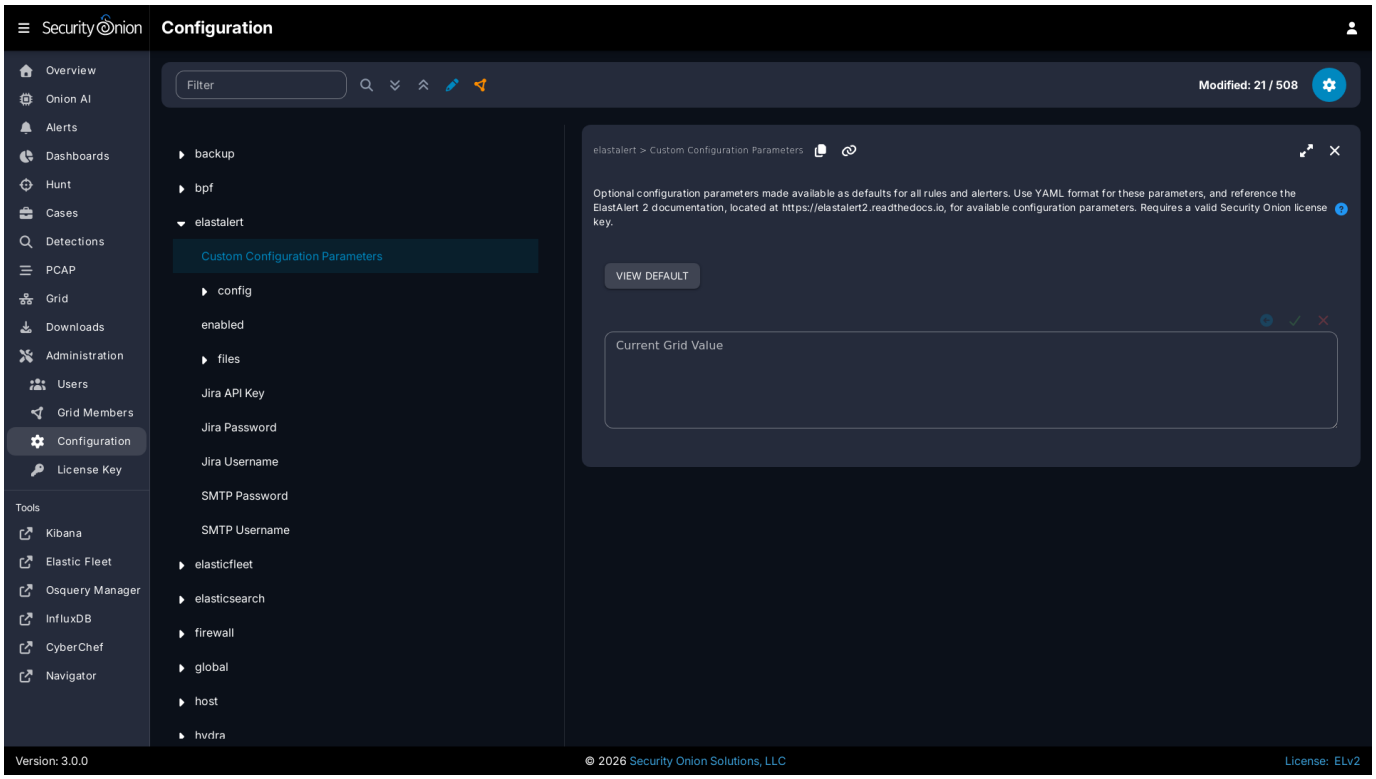

Notice there are special settings for Jira and SMTP notifications. These are unique in that ElastAlert requires those two alerters to read their credentials from a file. Security Onion has simplified this process by presenting these Configuration fields to enter the optional credential data, and the backend process will take care of generating the required files for ElastAlert.

The files subtree includes a list of several file settings, which allows for populating the contents of certain files that the alerters can optionally utilize. Most alerters use the files for specifying a custom Certificate Authority, so that ElastAlert can securely and confidently connect to remote servers that may be using custom SSL/TLS certificates. Again, Security Onion's backend process will handle generating these files from the supplied configuration data provided in the user interface.

Next, the **SOC > config > server > modules > elastalertengine > Notifications: Sev 0/Default Parameters** setting is used to customize each alerter's own parameters. As ElastAlert already provides detailed documentation on the required parameters for each alerter, this documentation will not cover the same information, but instead will focus on two popular alerters: Slack and SMTP.

> **Note**
>
> Reference the alerter parameters at https://elastalert2.readthedocs.io/en/latest/alerts.html#alert-types.

Slack

To have Sigma rules send notifications to Slack, add the following line to the **SOC > config > server > modules > elastalertengine > Notifications: Sev 0/Default Parameters** configuration setting:

```
slack_webhook_url: "[https://hooks.slack.com/services/YOUR_WEBHOOK_URI](https://hooks.slack.com/services/YOUR_WEBHOOK_URI)"
```

Email (SMTP)

To have Sigma rules send notifications via email, add the following lines to the **SOC > config > server > modules > elastalertengine > Notifications: Sev 0/Default Parameters** configuration setting:

```
email: youremail@yourcompany.com
smtp_host: "your_company_smtp_server"
from_addr: "ElastAlert@yourcompany.com"
```

If the SMTP server requires authentication make sure the special **SMTP Username** and **SMTP Password** configuration settings are also specified.

SOC Detections

Once the alerter parameters are configured, as described above, the next step is to configure Detections in order to activate one or more notification alerters.

Navigate to the Administration -> Configuration screen. Next, locate the `SOC > config > server > modules > elastalertengine` settings.

In the **Notifications: Sev 0/Default Alerters** configuration setting, add the name of each alerter that should be activated, one alerter name per line. For example, to add both slack and email:

slack
email

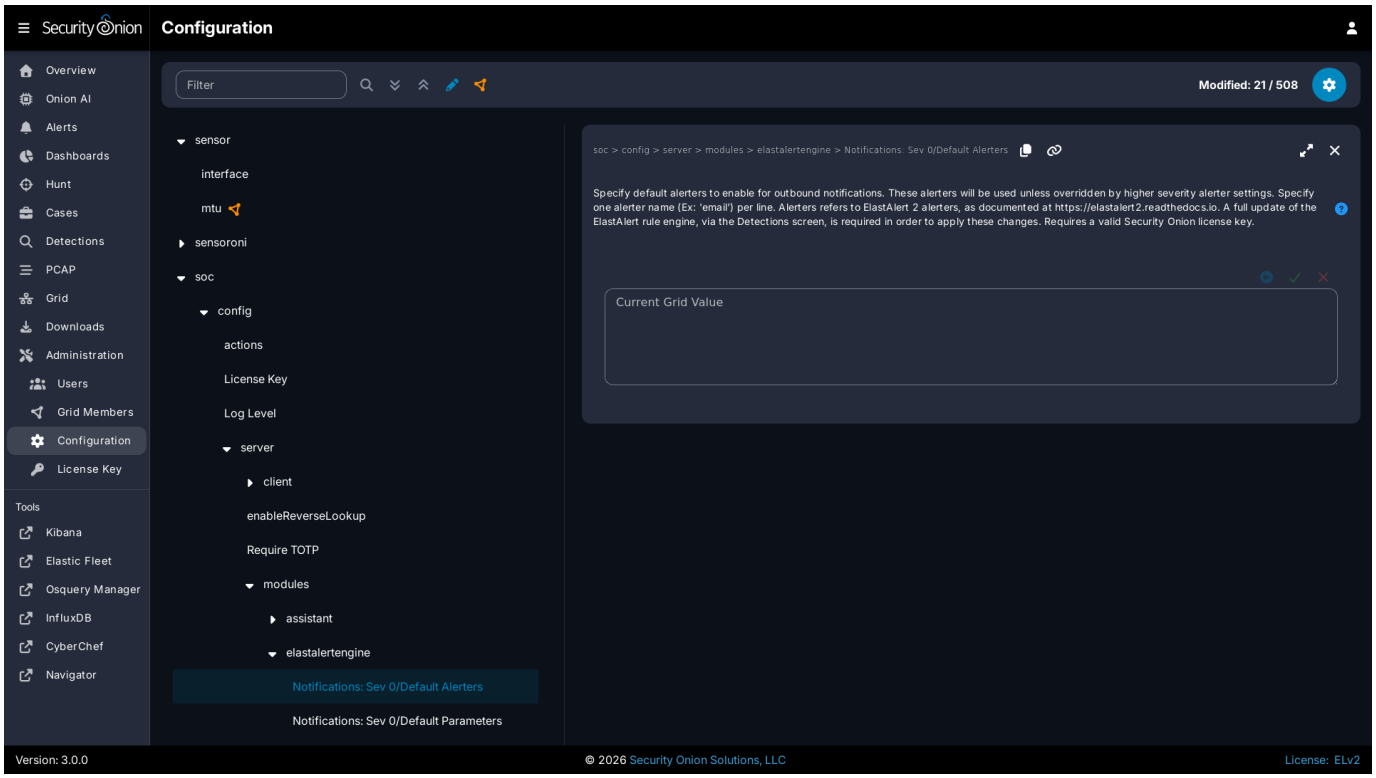

Important! After activating (or removing) an alerter from this setting, the ElastAlert engine must be fully updated. This can be done via the Detections screen, under the Options dropdown.

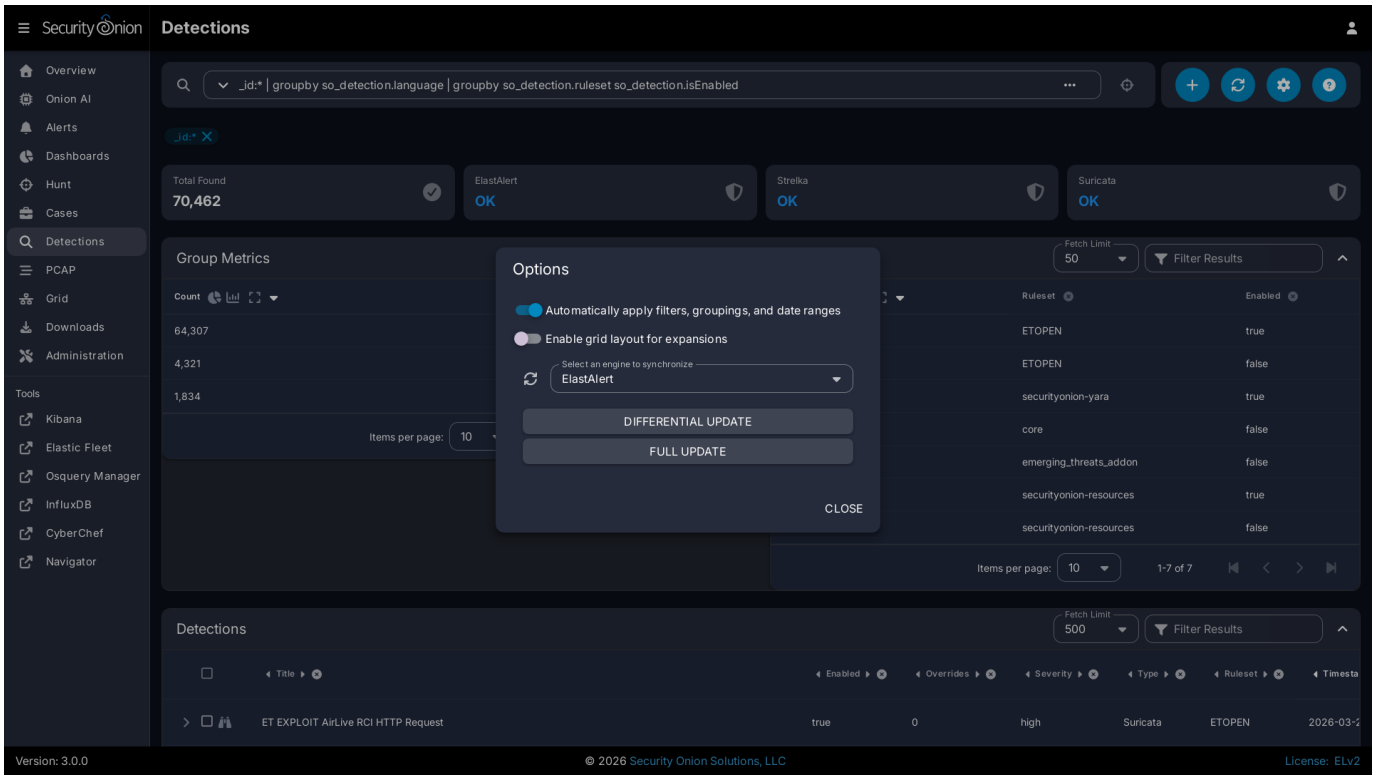

Severity-Based Notifications

The instructions above setup the default notification settings, for all outbound notifications. However, notification settings can be customized for higher level severities. Severities are specified in Sigma Detections.

Severity levels progress as follows, starting with the lowest, least significant severity:

1. Unknown Severity
2. Informational Severity
3. Low Severity
4. Medium Severity
5. High Severity
6. Critical Severity

If notification settings are not specified for a particular severity level then it will use whatever settings are specified at the next lower severity. If that severity is also not specified, then it continues looking for lower severity settings.

> **Note**
>
> Higher severity levels do not inherit parameters or alerters from lower severities. Consequently, if `email` is specified as the default (Severity 0) alerter, and it's desired to have both `email` and `slack` notifications sent with **High/Sev 4** severity or above, then both `email` and `slack` will need to be specified for the `Notifications: Sev 4/Default Alerters` setting, one per line. This same principle applies to the parameters, which are also not inherited. In order to inherit default parameters across all severities, the parameters can be specified in the `ElastAlert > Custom Configuration Parameters` setting.

User-Defined Notifications

Individual Sigma detections can be tagged to change the detection's alerting behavior. The tags are set inside the detection source. Tag details are defined below:

- `so.notification`: When this tag is present inside of a Sigma tag list, the detection will only perform outbound notifications. It will not add an alert to the SOC Alerts screen.
- `so.alerters.customAlerters`: When this tag is present inside of a Sigma tag list, the detection will perform notifications for an alternate set of ElastAlert 2 alerters. More information on how to choose these alerters is provided below.
- `so.params.customAlertersParams`: When this tag is present inside of a Sigma tag list, and when the above tag is also included, then an alternate set of custom parameters will be applied to the ElastAlert 2 alerters.

To customize the alerters and parameters to use when these tags are specified in a Sigma detection, navigate to the Configuration screen. Find the `SOC > config > server > modules > elastalertengine > additionalUserDefinedNotifications > customAlerters` setting and add the custom alerters, one per line, similar to what is done for the Severity-Based notifications above. Similarly, find the sibling setting to define custom alerter parameters: `SOC > config > server > modules > elastalertengine > additionalUserDefinedNotifications > customAlertersParams`.

> **Note**
>
> User-defined alerters will override severity-based alerters, provided the user-defined alerters are properly configured. If the Sigma tags specify custom alerters but the corresponding setting does not exist in the Configuration then the severity-based notifications will continue to be used.

To create additional user-defined alerter configurations, enable Advanced mode and navigate to the same `customAlerters` and `customAlertersParams` settings mentioned above. With Advanced mode enabled there will be a "Create Duplicate" button that allows for duplicating these settings. Follow the on-screen instructions to create the duplicate settings. Then, to make use of these new settings, in the Sigma tag list replace the `so.alerters.customAlerters` tag suffix with the name (case-sensitive) of the duplicated setting. For example, if the duplicated settings are named `SysAdminAlerters` and `SysAdminParams` then the two tags to specify in the Sigma detection source are `so.alerters.SysAdminAlerters` and `so.params.SysAdminParams`. Only one user-defined alerters and parameters setting will be used if multiple tags match the `so.alerters.` and `so.params.` prefixes. In other words, attempting to specify multiple user-defined alerters within a single Sigma detection will result in an ambiguous outcome.

Example:

```
title: Security Onion - Grid Node Login Failure (SSH) (copy)
id: 0c880a39-f2cc-4e80-af26-eb08e2fe4b0a
status: experimental
description: Detects when a user fails to login to a grid node via SSH. Review associated logs for username and source IP.
```

```
author: Security Onion Solutions
date: 2024/08/27
logsource:
    product: linux
    service: auth
detection:
    selection:
        event.outcome: failure
        process.name: sshd
        tags|contains: so-Grid-node
    filter:
        system.auth.ssh.method: '*'
    condition: selection and not filter
tags:
    - so.alerters.SysAdminAlerters
    - so.params.SysAdminParams
    - so.notification
falsepositives:
    - none
level: high
license: Elastic-2.0
```

21.6.2 Notification Formatting

There are a wide range of capabilities to format notification messages to the various endpoints supported by ElastAlert 2. Refer to the ElastAlert 2 documentation for all available formatting parameters: https://elastalert2.readthedocs.io/en/latest/alerts.html#alert-subject.

Below is an example of customizing the notification message. This format is compatible with most of the ElastAlert 2 alerters but may only work with specific connection-related alerts, due to it referencing specific connection fields. To use, paste these settings into the desired configuration alerter params field, as discussed earlier in this section. Change the hostname as necessary in the included URLs.

```
alert_subject: "Alert: {0} {1}"
alert_subject_args:
- rule.name
- "@timestamp"
alert_text: |
  Alert details are available in Security Onion Console: https://manager/#/hunt?q=log.id.uid%3A{0}&rt=1&rtu=days

  Source: {1}:{2}
  Destination: {3}:{4}

  Investigate the network community ID: https://manager/#/hunt?q=network.community_id%3A"{5}"&rt=1&rtu=days
alert_text_type: alert_text_only
alert_text_args: ["log.id.uid", "source.ip", "source.port", "destination.ip", "destination.port", "network.community_id"]
```

21.6.3 Applying Changes

In order for alerters and parameters to take effect, multiple synchronizations must occur. These are done automatically on a set schedule, but it is possible to force them earlier, if needed. Specifically, the following must take place for the changes to be applied to the ElastAlert 2 rules:

1. Changes are saved in Configuration screen by the SOC Admin.
2. Configuration is synchronized across the grid. To manually force a grid sync, go to the Configuration screen, open the `Options` dropdown at the top, and click `Synchronize`.
3. Sigma Detection edits are saved, such as adding the user-defined notification tags, or changing the severity.
4. Sigma Detections are synchronized. Click Full Synchronize for ElastAlert rules, or to force a single detection sync go to the Detection Source tab, make an edit to the source, and click `Update`.

> **Note**
>
> It may take a minute or two for the ElastAlert 2 process to detect the changed rules, and then another few minutes for ElastAlert 2 to run that rule.

21.7 Kafka

From https://kafka.apache.org :

> Apache Kafka is an open-source distributed event streaming platform used by thousands of companies for high-performance data pipelines, streaming analytics, data integration, and mission-critical applications.

If you need guaranteed message delivery, then you can enable Kafka which replaces Redis and Logstash on the Security Onion Manager node and Receiver nodes.

Note

This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.7.1 Guaranteed Message Delivery

By leveraging Kafka, you can ensure that messages sent to the Kafka cluster are written to disk on the partition leader and replicated to other physical brokers before acknowledging receipt to the producer.

For more information, please see https://kafka.apache.org/documentation/#producerconfigs_acks.

21.7.2 High Availability

With a properly configured Kafka cluster, data integrity and data availability are maintained even in the event of a Kafka broker or controller failure. In order to start taking advantage of Kafka, we'd recommend three Kafka controllers and a minimum of three Kafka brokers. With this configuration, you can increase your replication factor to ensure that messages are replicated to multiple brokers.

For more information about replication, please see https://kafka.apache.org/documentation/#replication.

21.7.3 Configuration

Important

Before configuring Kafka, it is recommened you build your grid as you would normally. This includes adding any receiver nodes that will later be repurposed as Kafka brokers or controllers.

Also note that Guaranteed Message Delivery which leverages Kafka, requires a valid Security Onion Pro license. See the Security Onion Pro section.

You can modify your Kafka configuration by going to Administration --> Configuration --> Kafka.

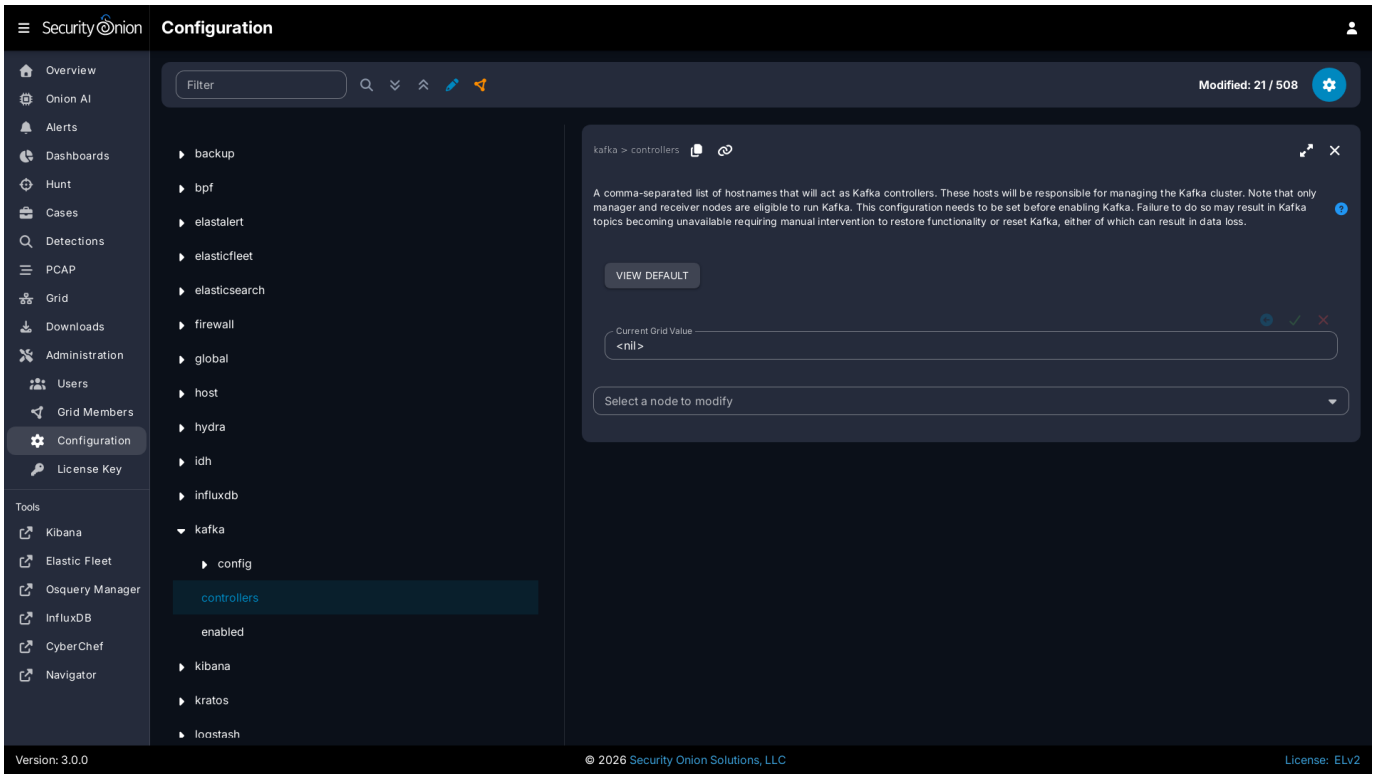

Controllers

Controllers are responsible for managing the Kafka cluster. This includes electing a leader and managing the cluster metadata.

Controllers can be relatively lightweight virtual machines or physical machines. A system with 4 CPU cores, 8GB of RAM, and 200GB storage should be sufficient for a single Kafka controller.

Controllers are assigned by adding the hostnames to the `controllers` configuration option separated by a comma.

```
hostname1,hostname2,hostname3
```

We recommend that you have at least three controllers. This allows for a single controller to fail and the cluster to continue to operate.

For more information about controllers, please see https://kafka.apache.org/documentation/#kraft_voter.

Brokers

Brokers are responsible for the storage and replication of messages. With Kafka enabled, the Elastic Agent will begin to act as a producer and write its messages to topics stored on the Kafka broker(s).

Brokers require much more resources than controllers as they are responsible for managing the data and providing the data to consumers. Sizing brokers depends heavily on expected message volume. A system with 8 CPU cores, 32GB of RAM, and 500GB storage should be sufficient for a single Kafka broker.

> **Warning**
>
> The above hardware recommendations should be used as a minimum. Increasing the number of brokers and the resources available to each broker will increase the overall performance of the Kafka cluster. Additionally, without sufficient storage space on each broker, the cluster may run out of space and stop accepting messages. The brokers `log.retention.hours` setting can be configured to delete messages after a certain amount of time to prevent this from happening.

Broker configuration can be modified by going to Administration --> Configuration --> Kafka --> config --> broker.

For more information about broker configuration, please see https://kafka.apache.org/documentation/#brokerconfigs.

Enabling Kafka

Once you have the appropriate configuration in place, you can enable Kafka by navigating to Administration --> Configuration --> global --> pipeline and setting the value to `KAFKA`.

There is no need to click on the `SYNCHRONIZE GRID` button. Once you have set the global pipeline value to `KAFKA`, the changes will begin to take effect in the background before finally switching the grid to the new pipeline.

> **Note**
>
> In order to change the global pipeline you will need to enable the Administration --> Show advanced settings option.

21.7.4 More information

> **Note**
>
> | For more information about Kafka, please see: https://kafka.apache.org/documentation/#gettingStarted

21.8 Connect API

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

The Security Onion Connect API allows other servers to integrate with Security Onion, and access the same functionality that the Security Onion Console user-interface provides. Access to the Connect API is permitted through API Clients, which can be created by SOC administrators via the SOC UI -> Administration -> API Clients screen.

The Connect API currently provides functionality exposed by the Security Onion Console server. It does not provide full access to third-party applications included with the Security Onion platform. Specifically, while you can read events from Elasticsearch, you cannot manipulate Kibana settings via the Security Onion Connect API, unless those settings are already exposed via the SOC Configuration system.

21.8.1 Enabling Connect API

By default, newly setup grids will not be configured for API client access. To enable API client access, the following steps must be taken:

1. A license key must be applied to the grid. The license key must include the API feature.
2. The Hydra feature must be enabled via the `hydra > enabled` setting in the Configuration screen.
3. Synchronize the grid to apply the license key and configuration changes. This can be done via the Configuration screen options dropdown.

21.8.2 API Client Credentials

In order to communicate with the Connect API, an API Client must be created. Navigate to the Administration menu using a superuser account. Under the Administration menu click the API Clients menu option. Create a new API client using a short name that reflects the intended usage of this client. Use the Notes field to provide more information, if desired. Upon saving the new client a generated secret will be issued. This client ID and secret pair is needed to authenticate to the Connect API. Protect these credentials using industry best practices.

21.8.3 Authorization / RBAC

API clients are permitted access to various components within Security Onion using the same RBAC system for users. However, rather than assign *roles* to API clients, the more granular *permissions* are assigned. For example, while a *user* might be assigned the `analyst` role, an API *client* would be assigned the `events/read`, `events/write`, `cases/read`, etc. This ensures that remote systems will only have access to the minimum necessary permissions required for the integration.

Currently OAuth 2.0 scopes are not utilized, since these permissions are assigned outside of the OAuth 2.0 flow.

21.8.4 OAuth 2.0 Authentication Flow

API clients must use the OAuth 2.0 client credentials flow to authenticate to the Security Onion manager node.

Exchange Client Credentials for an Access Token

Obtain an access token by submitting a POST request to https://BASE_URL/oauth2/token, and providing the client ID and client secret via the *Basic* authentication scheme. The body of the request must contain `grant_type=client_credentials`.

Example:

```
curl --cacert ca.crt -X POST -u socl_my_new_client:hwKHspsX2bMuoIs7kGwN https://BASE_URL/oauth2/token -d grant_type=client_credentials
```

Where you will replace:

- `ca.crt` with your manager's certificate authority. If a custom certificate has been applied to your grid after setup completed, you can access it via the Configuration screen (requires superuser role) from the `nginx > ssl > SSL/TLS Cert File [adv]` config setting, or if using the default generated certificate authority, retrieve the `/etc/pki/ca.crt` certificate file via SSH from the manager node.
- `socl_my_new_client` with your client ID (generated by SOC during API client creation)
- `hwKHspsX2bMuoIs7kGwN` with your API client's generated secret
- `BASE_URL` with your manager's IP or hostname, depending on which option you selected during Security Onion setup

The response will resemble the following:

```
{"access_token":"ory_at_xI1_2FVvoWR60GHAXZXAcDW7V3qEi2mIB8RKnpqN0fk.Hy5LaHPqh9sfWVEtDXDhs8Gj-9YZ85FJHp6pyD0eeNw","expires_in":
3599,"scope":"","token_type":"bearer"}
```

The access token will expire in 2 hours, by default, after which a new access token must be requested using this same credential exchange method again.

Authorize API Requests with an Access Token

Now that the access token has been retrieved, the API requests can be submitted. These requests will utilize the access token via the HTTP Authorization header, using the *Bearer* scheme.

Example:

```
curl --cacert ca.crt -X GET --oauth2-bearer ory_at_U74544Scqho5KGOci-qemWsOjxOU8TALqddAnrfxAGg.7GlO4SYPUAllO23LVqs9e_FXl0tAdRlUk3AH9IplWRU https://BASE_URL/
connect/info
```

Where the provided bearer token above must be replaced with the access token extracted from the client credential exchange response.

21.8.5 Manager of Managers

To interact with subgrid data, while still communicating with the primary **Manager of Managers (MoM)** node, include an additional query string parameter on the API URL. The parameter key is `gridId` and the value should be set to the desired subgrid ID.

21.8.6 API Reference

> **Warning**
>
> New releases of Security Onion may contain additional fields in API responses. Consequently, it is important that the API output be properly parsed by official libraries that can handle these scenarios. Using custom parsing of API outputs may lead to upgrade-related malfunctions.

An interactive API view is available for browser-based viewing: Interactive API

21.9 Active Query Management

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

Security Onion Pro customers can now view and cancel long-running Elasticsearch queries.

This screen is located under the Administration menu on the left side of the Security Onion Console. This menu option will only be visible for users having the `superuser` role.

21.9.1 Viewing Active Queries

By default, the active query list will filter out internal, non-cancelable, and child queries. Therefore, in most situations it will be normal for the list to be empty. If there is a desire to view all active queries without any filtering the superuser can click the Options bar at the top of the screen and it will expand to show a filter slider that can be toggled off. Note that the query to show this list will be shown in the active query list. However, by the time the list is rendered on the screen that query will have already finished.

While it would be great to show the user that initiated the query on this active query list, unfortunately Elasticsearch does not provide this information.

> **Note**
>
> Be aware that queries not initiated by a user can appear in the filtered list. The filter only filters out a specific type of query that are always known to be internally initiated.

21.9.2 Canceling Active Queries

Sometimes a poorly written query or an indexing problem can result in a query running for an abnormally long time. This long-running query can adversely affect other analysts' Security Onion experience, since their queries will begin performing slower than usual. If a long-running query needs to be stopped the superuser can click the X icon at the right side of the active query listing. A popup will appear asking for confirmation.

Canceling queries can take several seconds. Consequently, refreshing the active queries may still show that query in the list for a short time.

> **Note**
>
> Not all queries are cancelable. If a query cannot be canceled then no cancel icon will appear next to that query listing.

> **Warning**
>
> Canceling internal, system queries can disrupt the Elasticsearch internal processes. Avoid canceling queries that are not confirmed to be user-created.

21.10 Manager of Managers

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license with the necessary subgrid allocations to use this feature.

If you are an enterprise customer responsible for managing multiple grids, then you may benefit from the Manager of Managers feature!

This feature allows you to elect a special grid manager (MoM) to be configured with knowledge of remote grids (subgrids) allowing the MoM to reach out to the subgrid to query and/or modify Grid state, events, configuration, etc.

Appropriately privileged users logging into the MoM will be able to interact with those subgrids from a single user interface. A Subgrid selection list will be available in the top-right of the SOC UI. Users can choose which Grid they would like to interact with. Additionally, fault states will automatically propagate up to the MoM user interface giving those users the quick updates of new fault states within subgrids, regardless of which subgrid is currently selected. If the subgrid list becomes disabled that will indicate the current screen does not support subgrid interaction.

Users logging into a subgrid SOC user interface will have no knowledge or visibility into other sibling grids or the MoM Grid.

21.10.1 MoM Requirements

The Manager of Managers feature has been designed such that only the MoM Grid node requires network access to the subgrid manager nodes. This simplifies network and firewall configuration by ensuring end-user connections are isolated to a single endpoint. Therefore, users logged into the MoM SOC only need web access to the MoM manager.

Please note that Manager of Managers is intended for production deployments so it is only supported on MANAGER, MANAGERSEARCH, and STANDALONE installations. It is not supported for IMPORT or EVAL installations.

21.10.2 Data Isolation

In a Manager of Managers configuration the data from each subgrid resides at rest within the subgrid only. If a MoM user requests information of a subgrid then that returned data will transit through the MoM node and from there move to the user's web browser.

> **Note**
>
> Manager of Managers currently does not attempt to merge data from multiple grids into a single SOC view.
>
> The suggested method for accessing data across multiple grids, such as when searching for alerts across all grids, is to utilize cross-cluster search (CCS). Keep in mind that this should be configured such that internal Security Onion indices are excluded from CCS, otherwise there is a risk of data duplication and/or corruption of internal SOC entities such as Cases, Detections, etc.

21.10.3 Subgrid Outages

The MoM node will expect all subgrids to be reachable. In the event a subgrid is not reachable, or is in a state where it cannot respond successfully to incoming MoM requests, the MoM SOC interface will display an error notifying the logged in users that it cannot complete the request. Subgrid outages could delay retrieving state information from other subgrids.

If a subgrid is expected to remain in a disconnected state for a longer period it is recommended to disable that subgrid in the Configuration screen and then synchronize the MoM Grid state. This scenario can occur more frequently with subgrid "kits" that are transported between field locations for forensic collection purposes.

21.10.4 Configuration

> **Warning**
>
> Do not proceed with the configuration step until an appropriate Security Onion Pro license has been applied to each Grid. Configuring subgrids on a Security Onion Pro license that does not have sufficient subgrid allocations can cause that license to be exceeded which can disable all other Pro features.

Configuration of Manager of Managers requires two steps:

1. API Client: Creation of an API Client on each subgrid manager via the subgrid's Connect Client screen
2. Subgrid Config: Configuration of the subgrids on the MoM Configuration screen

API Client

Create a new, dedicated Connect Client on each subgrid and assign all permissions that the MoM will need to perform the desired functions. For example, if the MoM users will need complete management and visibility (effectively `superuser` access) then grant all permissions. However, if the MoM users will only be monitoring the subgrid health then the API Client permissions can be limited to `read` permission on specific resources, such as `Grid`, `nodes`, etc.

The API Client ID and Secret will be needed on the next step, so ensure those are recorded for each of the subgrids.

While on the API Client screen, click the ⤓ icon to download the Certificate Authority (CA) certificate file. This file is in a text format and the contents will be pasted into the subgrid setting in the next step.

> **Note**
>
> Please be aware that any users in the MoM Grid will be able to connect to the subgrid using the permissions defined for the API client. For example, suppose that you create an API client ID in the subgrid called `supermom` and you grant it all permissions. Once the MoM is configured to connect to the subgrid as shown in the next section, then any users in the MoM Grid will connect to the subgrid as `supermom` and have all permissions to the subgrid regardless of whether the user has equivalent permissions in the MoM.

Subgrid Config

Once the subgrid API Client credentials are known that subgrid can then be added to the MoM's `Subordinate Grids` Configuration screen.

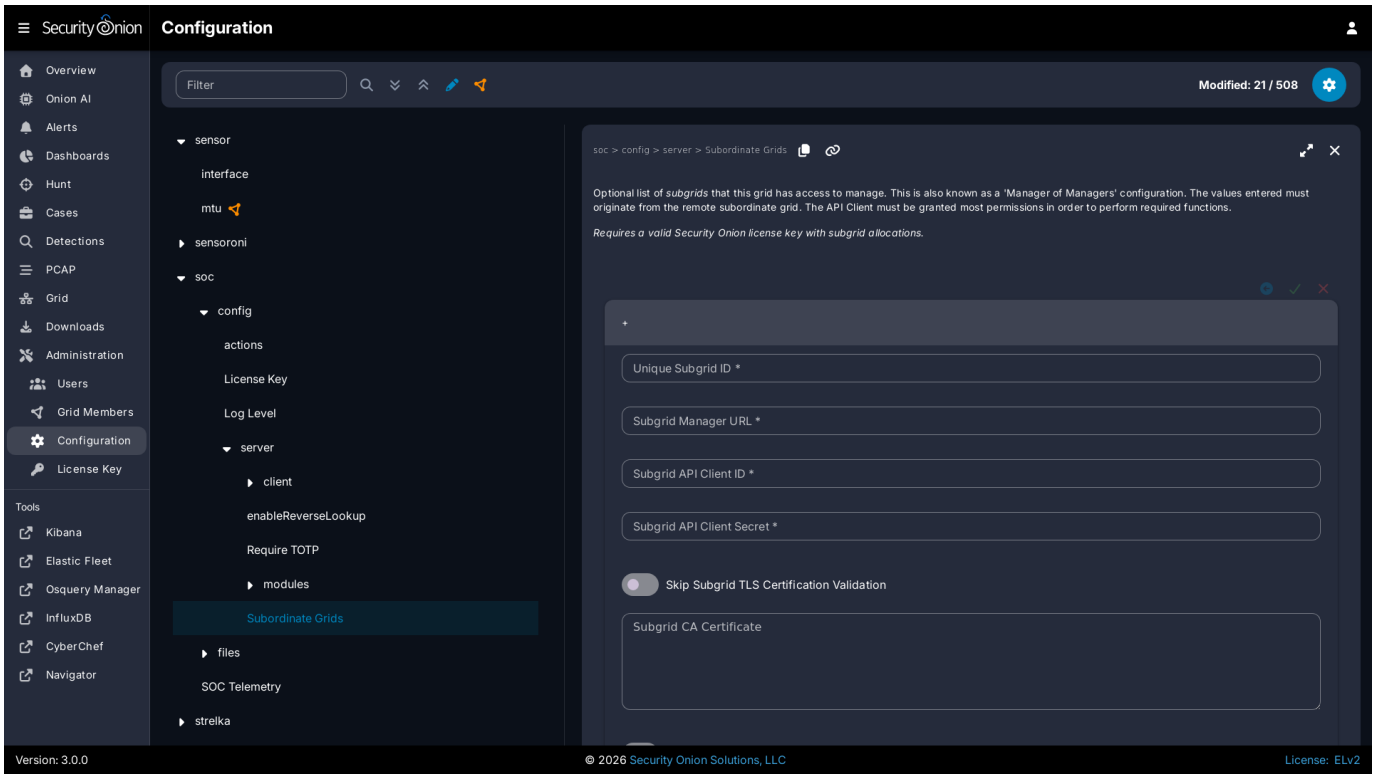

- As a superuser, log into the MoM SOC interface and navigate to Administration -> Configuration.
- Find the `SOC > config > server > Subordinate Grids` setting.
- Click the `+` icon to add a new subgrid.
- Give the new subgrid a unique ID that accurately describes this subgrid from the MoM's perspective.
- Enter in the subgrid's Manager URL (reference the subgrid's `base_url` value in this fully formed URL). Ex: `https://mysubgrid`. Note that this URL must be accessible from the MoM Grid node.
- Paste or enter the subgrid's API Client credentials. This refers to the Client API ID and generated secret.
- Paste the subgrid's API Client Certificate Authority (CA) contents into the `Subgrid CA Certificate` field.
- If this subgrid is ready, enable it.

Add additional subgrids as your Security Onion Pro license allows, and then click the green checkmark to save the configuration.

The configuration will be applied at the next 15-minute interval or you can apply it immediately on the MoM Grid by clicking the `SYNCHRONIZE GRID` button under the `Options` menu.

21.10.5 Licensing

The Manager of Managers feature requires that all involved grids have a valid Security Onion Pro license applied.

The MoM Grid will require a special Security Onion Pro license with an allocation of subgrids encoded into the license that meets or exceeds the number of configured subgrids. Exceeding the licensed subgrid allocation will cause the MoM SOC to show an "Exceeded" license state which will disable Security Onion Pro features.

Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license with the appropriate subgrid allocations.

21.11 MCP Server

Enterprise customers utilizing their own AI/LLM platform can now interface that platform directly to their Security Onion Grid by utilizing the Security Onion Model Context Protocol (MCP) server.

The MCP server gives AI the ability to query alerts, playbooks, events, etc. which provides the LLM with additional information needed to make informed decisions.

> **Note**
>
> The MCP server utilizes the Security Onion Connect API, which is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.11.1 Configuration

The MCP server is a low-overhead server application that is maintained separately from the Security Onion software. It's available in a GitHub repository, which includes the installation and configuration instructions.

See https://github.com/Security-Onion-Solutions/securityonion-mcp to get started.

> **Note**
>
> A Connect API Client must be created in the Security Onion API Clients screen. The API Client should be granted sufficient permissions needed to perform the tasks that the LLM will need to execute.
>
> For example, if the LLM will be querying events and playbooks then the API Client will need *events/read*, *playbooks/read*, and *detections/read* permissions.

21.11.2 Capabilities

The MCP server currently supports the following operations:

- Query Events via Onion Query Language (OQL)
- Query Playbooks

Future releases will likely bring more functionality, such as the ability to acknowledge alerts.

21.12 Security Onion App for Splunk

Enterprise customers utilizing Splunk can now install the Security Onion App for Splunk.

Warning

The Security Onion App for Splunk is not officially supported at this time.

21.12.1 Requirements

Note

The Security Onion App for Splunk utilizes the Security Onion Connect API, which is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.12.2 Configuration

See https://splunkbase.splunk.com/app/7887 to get started.

Note

A Connect API Client must be created in the Security Onion API Clients screen. The API Client should be granted sufficient permissions needed to perform the tasks that the Splunk app will need to execute.

21.13 Hypervisor

Security Onion Pro customers can create a hypervisor node that can run virtualized instances of Security Onion. If you have eligible machines with extra horsepower, you can use this feature to spin up additional Security Onion virtual machines (VMs) to take advantage of that extra power. This supports most major Security Onion node types and is especially helpful if you want to optimize your hardware's potential and expand your Elastic performance and retention. Contact your account manager to see if your hardware is supported.

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.13.1 Minimum Requirements

Please note the following MINIMUM requirements for a Hypervisor or Managerhype node:

- 32 CPU cores
- 64GB RAM

21.13.2 Host Hardware Reservation

The following resources will be reserved for the host machine and will be subtracted from the `Available` row under `Resource Summary`:

- **Hypervisor**
- 8 CPU cores
- 16GB RAM
- **Managerhype**
- 16 CPU cores
- 32GB RAM

21.13.3 Airgap

If you are in an Airgap environment, you will need to perform these steps prior to accepting the new hypervisor node in SOC Grid Members:

1. Download the Oracle 9 Qcow2 image from https://download.securityonion.net/file/securityonion/OL9U5_x86_64-kvm-b253.qcow2
2. Place OL9U5_x86_64-kvm-b253.qcow2 into /nsm/libvirt/boot/ on your manager node

21.13.4 Adding a Manager + Hypervisor

Manager nodes can have hypervisor capabilities. This node type is called a `managerhype`.

Install a new node, select the `DISTRIBUTED` deployment option, choose `New Deployment`, and then select the `Managerhype` option:

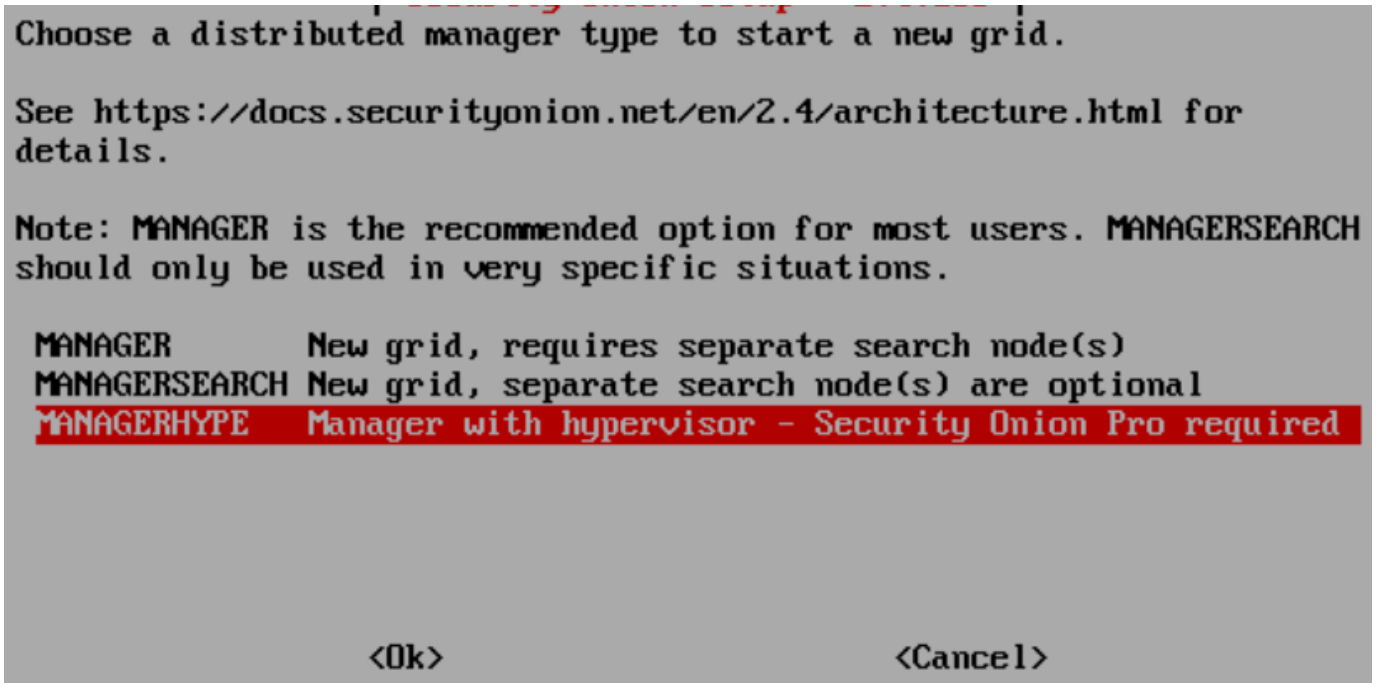

Once installation has completed, the new node will act as a manager node until a license key is added in SOC. After adding the license, it will take approximately 45 minutes for the hypervisor portion of the managerhype to be ready. This is due to the three highstates that are required to complete setup of the node.

If this is an airgap installation, then the instructions above will need to be followed prior to adding the license to SOC. The user will need to run `mkdir -p /nsm/libvirt/boot`.

The details under `Adding a Hypervisor`, regarding when the base domain is ready and VMs can be created, applies to the managerhype node as well as `Adding a Security Onion VM`.

21.13.5 Adding a Hypervisor

Install a new node, select the `DISTRIBUTED` deployment option, choose `Existing Deployment`, and then select the `Hypervisor` option:

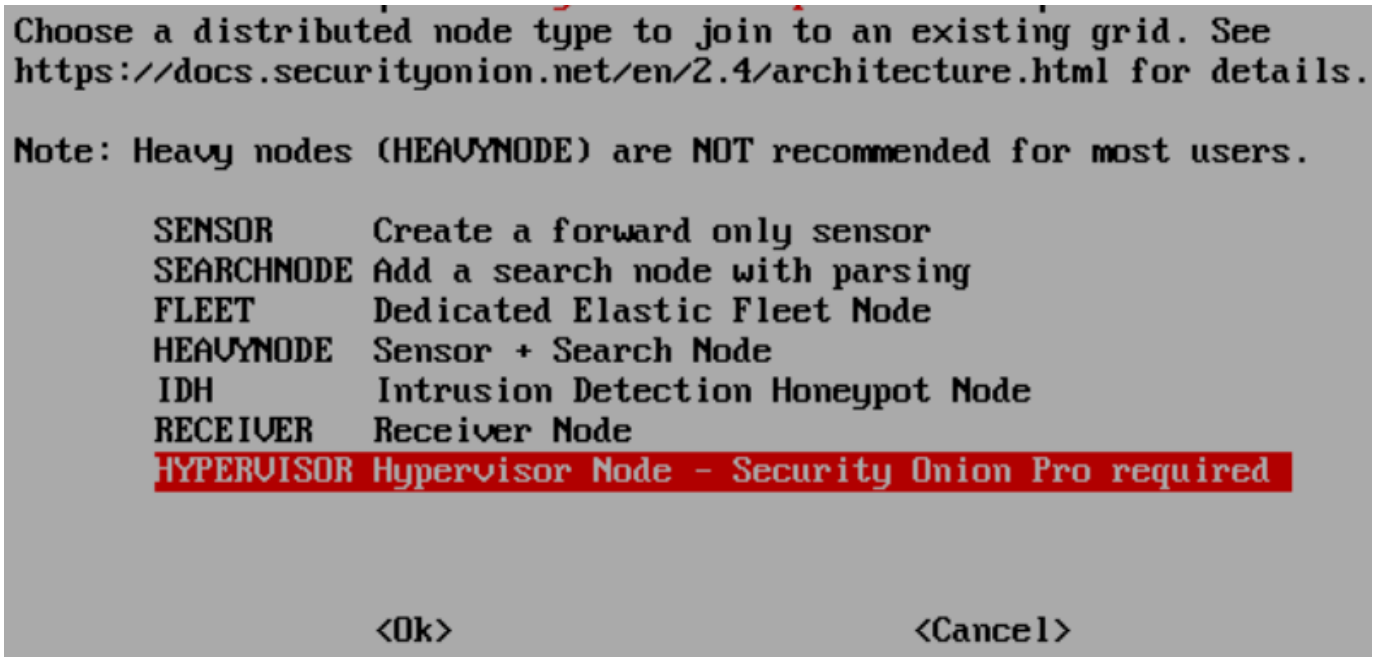

The manager will need to be able to connect to the hypervisor node by name so it will either need a DNS entry or you can manually add an entry in /etc/hosts on the manager.

Once the new hypervisor node has been accepted into the grid, go to SOC Configuration, click the Options menu, enable advanced settings, and then navigate to `hypervisor` settings. It should look like this:

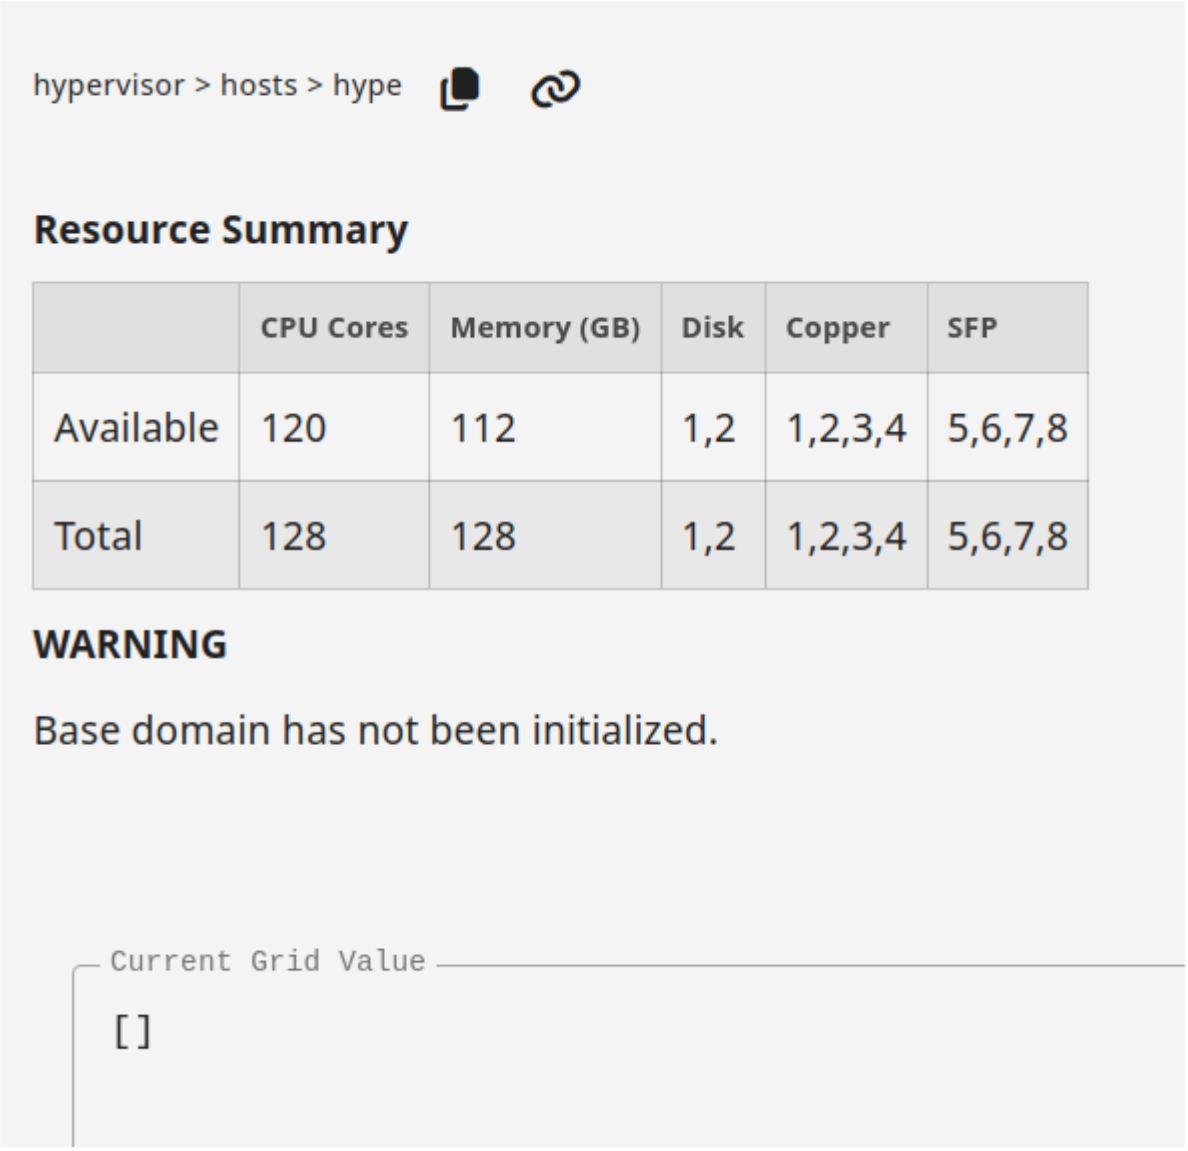

Once the base domain has been configured on the hypervisor (allowing VMs to be created), it should look like this:

21.13.6 VM Storage Options for /nsm

The vast majority of data, for all node types, is stored in /nsm/. For a VM, there are three options available for /nsm storage.

21.14 . Directory under /

This option uses a simple directory structure within the root filesystem. This is the default option if neither disk pass through or virtual disk are selected. The default image used for the VMs is 220GB, so this leaves about 200GB for storage on a fresh VM.

21.15 . Disk pass through

This method passes a physical disk directly to the VM, providing a dedicated disk to the node. It provides near-native disk performance and is ideal for production environments with high throughput requirements.

21.16 . Virtual disk

A virtual disk is created based on the size specified by the user in the SOC Grid Configuration and the space is pre-allocated on the hypervisor. The disk image file is not removed when the VM is deleted. A user may decide to leave this data around for a while, or delete it manually from the hypervisor where it is stored under `/nsm/libvirt/volumes`.

Note

If a user selects both a disk pass through and assigns a size for the virtual disk, then the disk pass through will be used for /nsm and the virtual disk will be ignored.

21.16.1 Adding a Security Onion VM

Please note that hardware for a vm cannot be modified once a vm is created.

To create a new VM, click the plus sign. You should see a form like this. If using DHCP, please pay special attention to the DHCP notes:

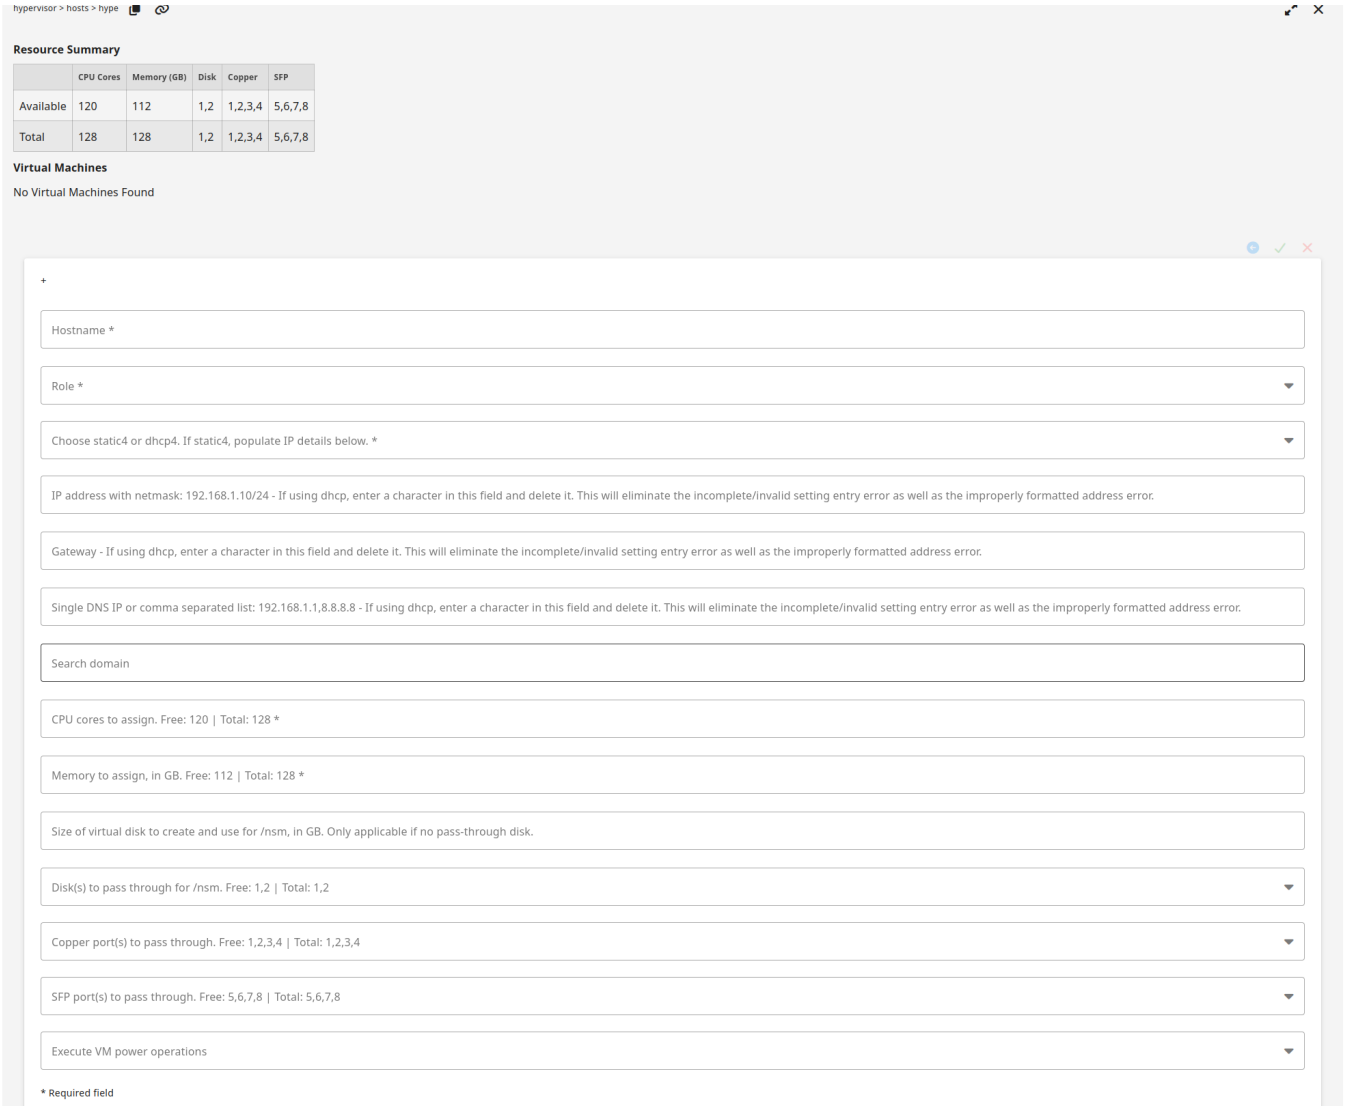

Fill out the form and then click the green check to save and create the VM:

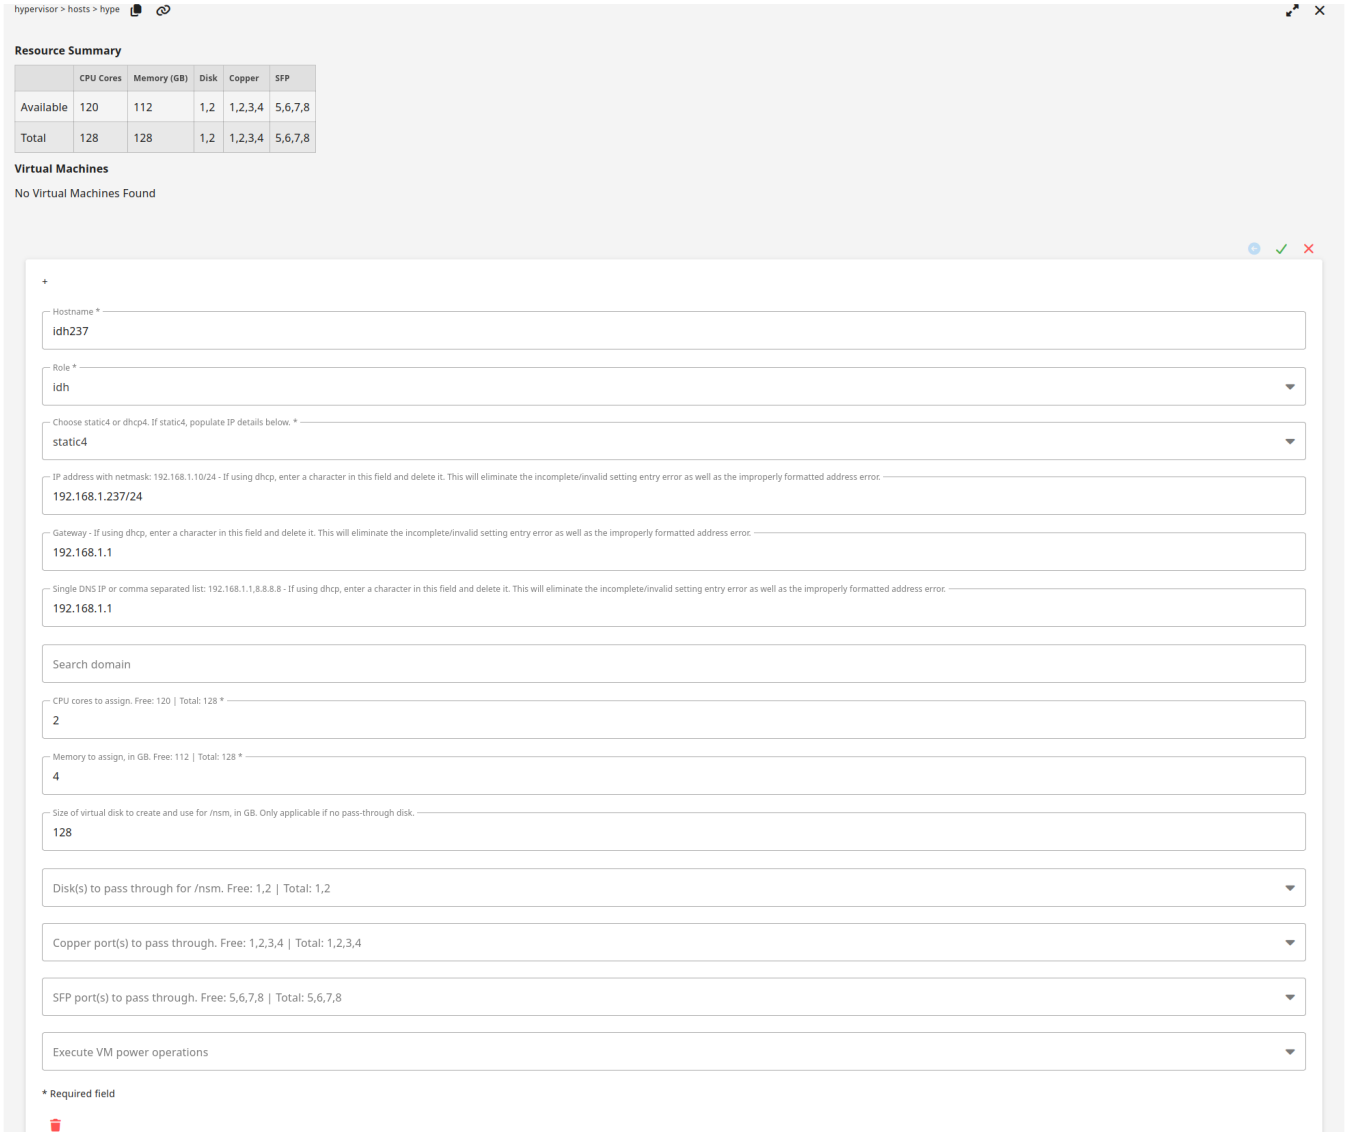

Once the VM is created and the first highstate is initiated, it should look like this:

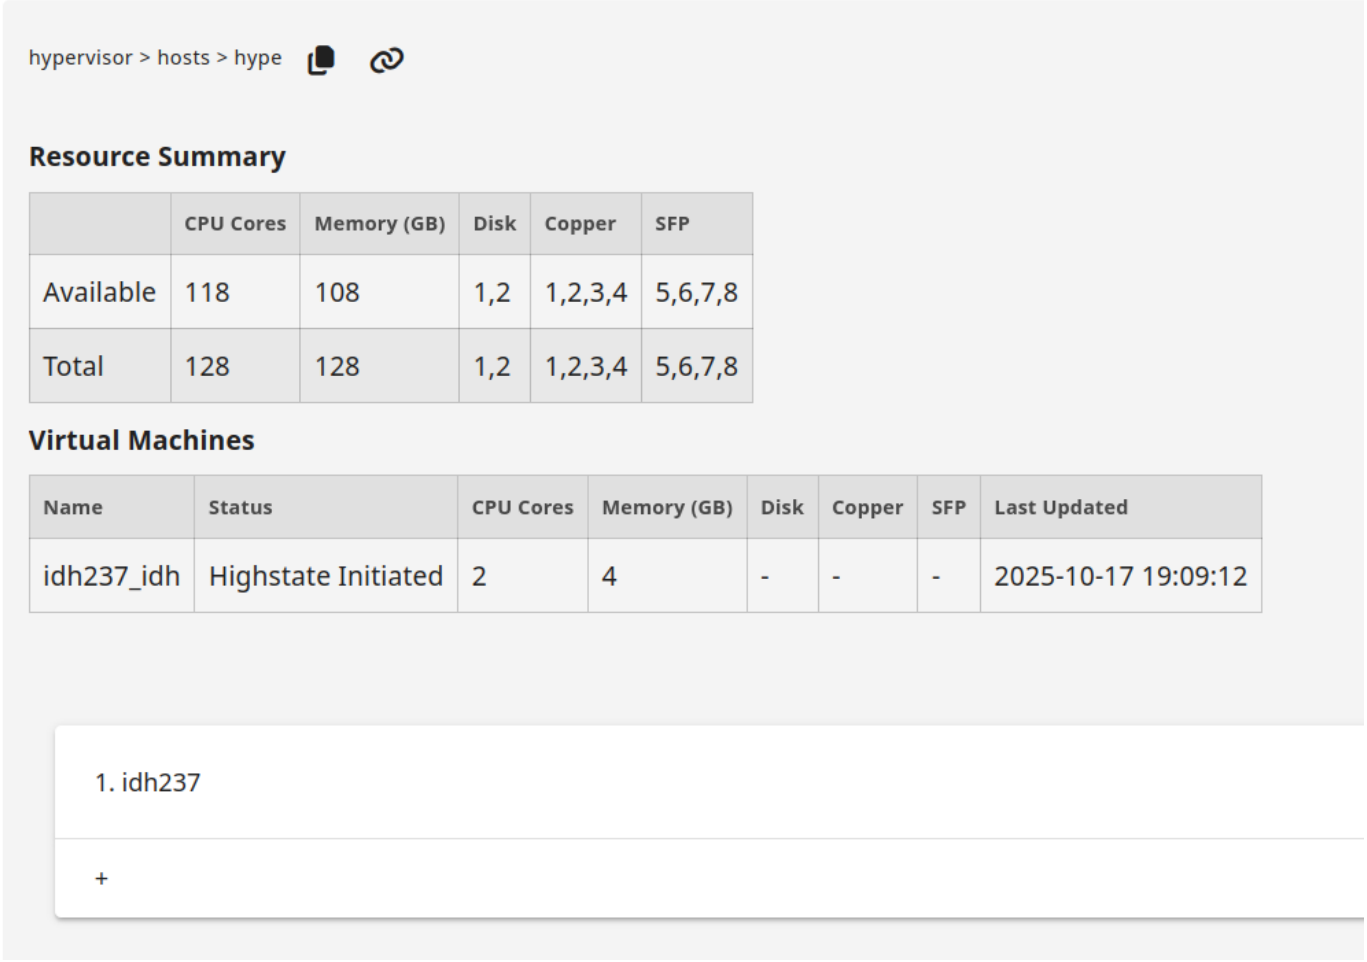

21.16.2 VM Creation Status

After a user has created a VM by filling out the form and clicking the green check mark, the web browser can be periodically refreshed to check the status of the VM creation. This is shown in the status column of the previous screenshot showing the highstate initiated.

Processing We detected that the user has requested to create a VM.

Hypervisor NSM Disk Full If the user requested a virtual disk for /nsm and there is not enough space on /nsm of the hypervisor, then this error will be seen. The user should delete this VM from the SOC Grid Configuration and either free up space on the hypervisor or create a new VM with a smaller /nsm.

IP Configuration Static or DHCP is being configured within the VM image before creation.

Starting Create The hypervisor has cloned the base image and has created the new VM image.

Executing Deploy Script The VM is being provisioned and the salt-minion is being bootstrapped.

Initialize Minion Pillars The VM has told the Security Onion manager to create its minion pillars.

Created Instance The VM creation process is complete.

Volume Creation If the user chose a virtual disk for /nsm, then this process creates it and fully allocates the space on the hypervisor. This step may take a while depending on the size of /nsm that was requested.

Volume Configuration The disk is being assigned to the VM.

Hardware Configuration The requested hardware has assigned to the VM.

Highstate Initiated The VM has been started and the first highstate is currently running. Subsequent highstates will not update the `Last Updated` column.

Destroyed Instance The instance has been destroyed. This VM will be removed from the `Virtual Machines` table after 48 hours.

21.16.3 Stopping or Starting a VM

If you need to stop or start a VM, you can use the `Execute VM power operations` menu at the bottom of the VM form:

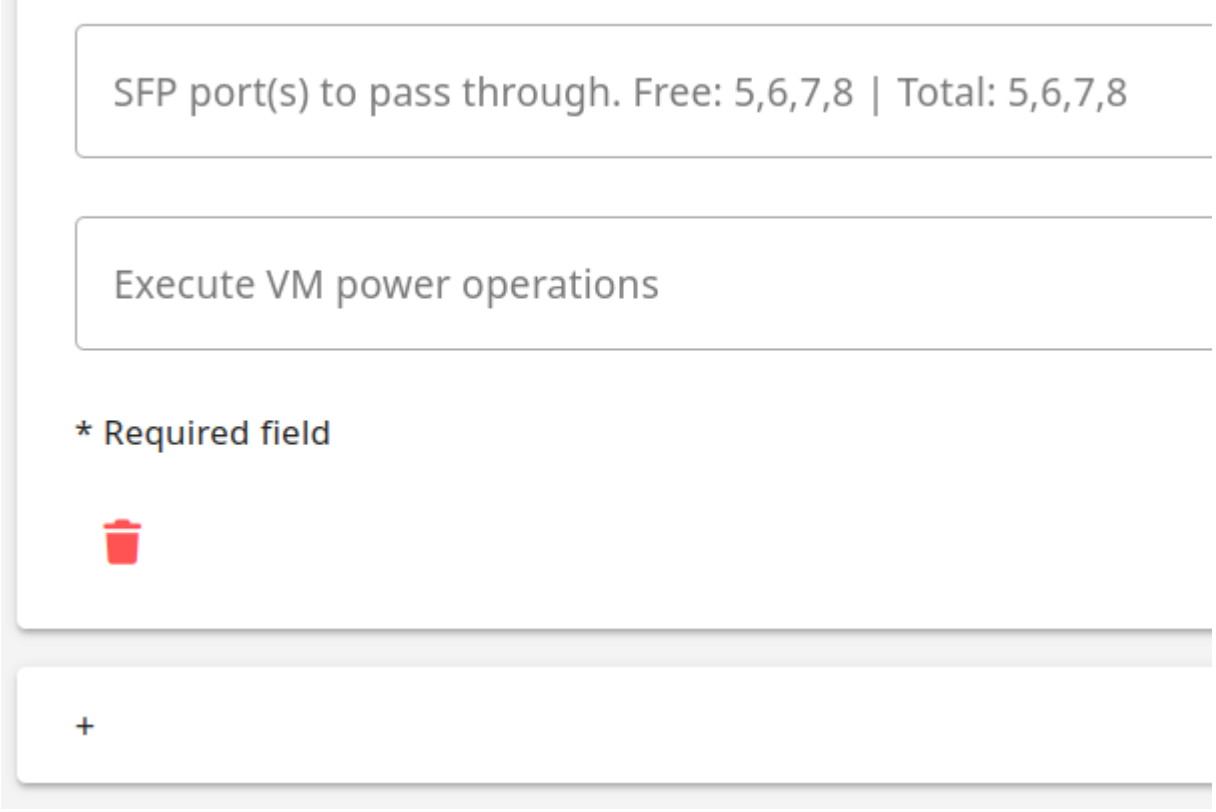

Here is the list of VM power operations and what they actually do:

- Reboot: gracefully reboot the VM (virt.reboot)
- Reset: forcefully reset the VM (virt.reset)
- Shutdown: gracefully shut down the VM (virt.shutdown)
- Start: start the VM (virt.start)
- Stop: forcefully stop the VM (virt.stop)

21.16.4 Deleting a VM

To delete a VM, click the trash icon:

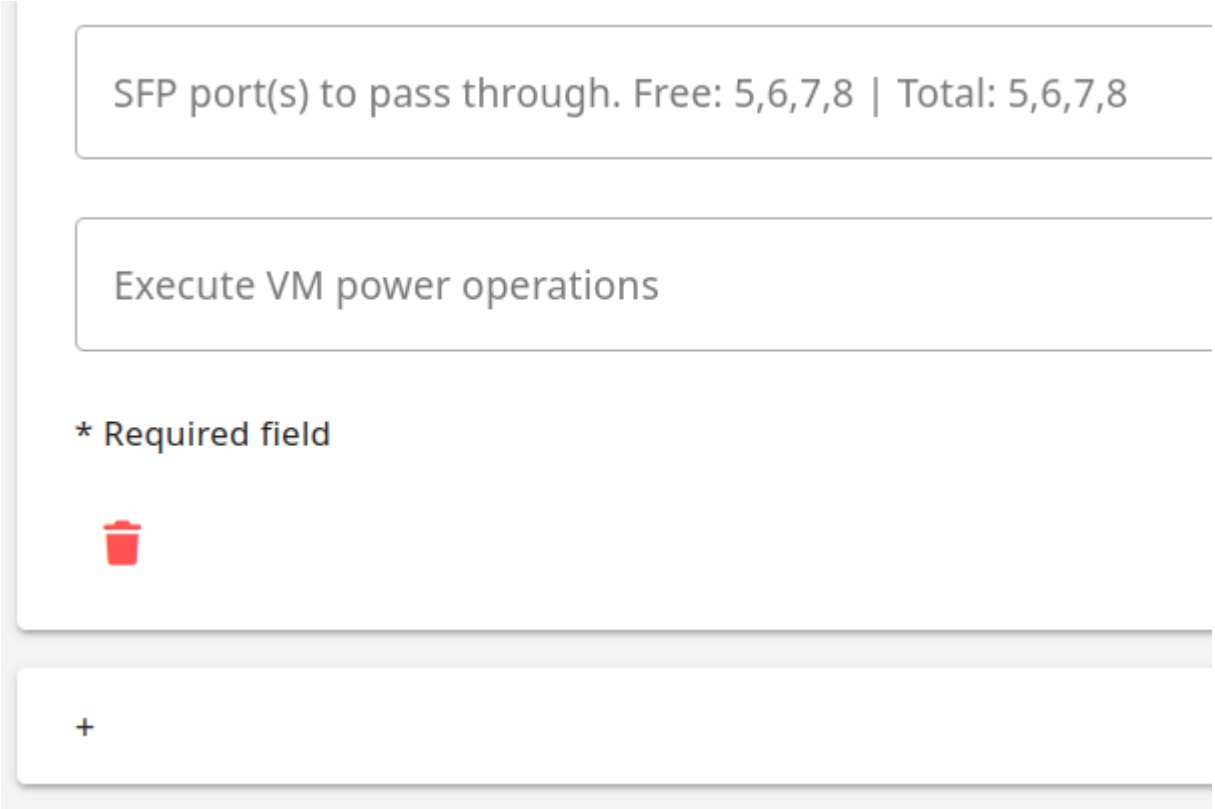

It will ask for confirmation:

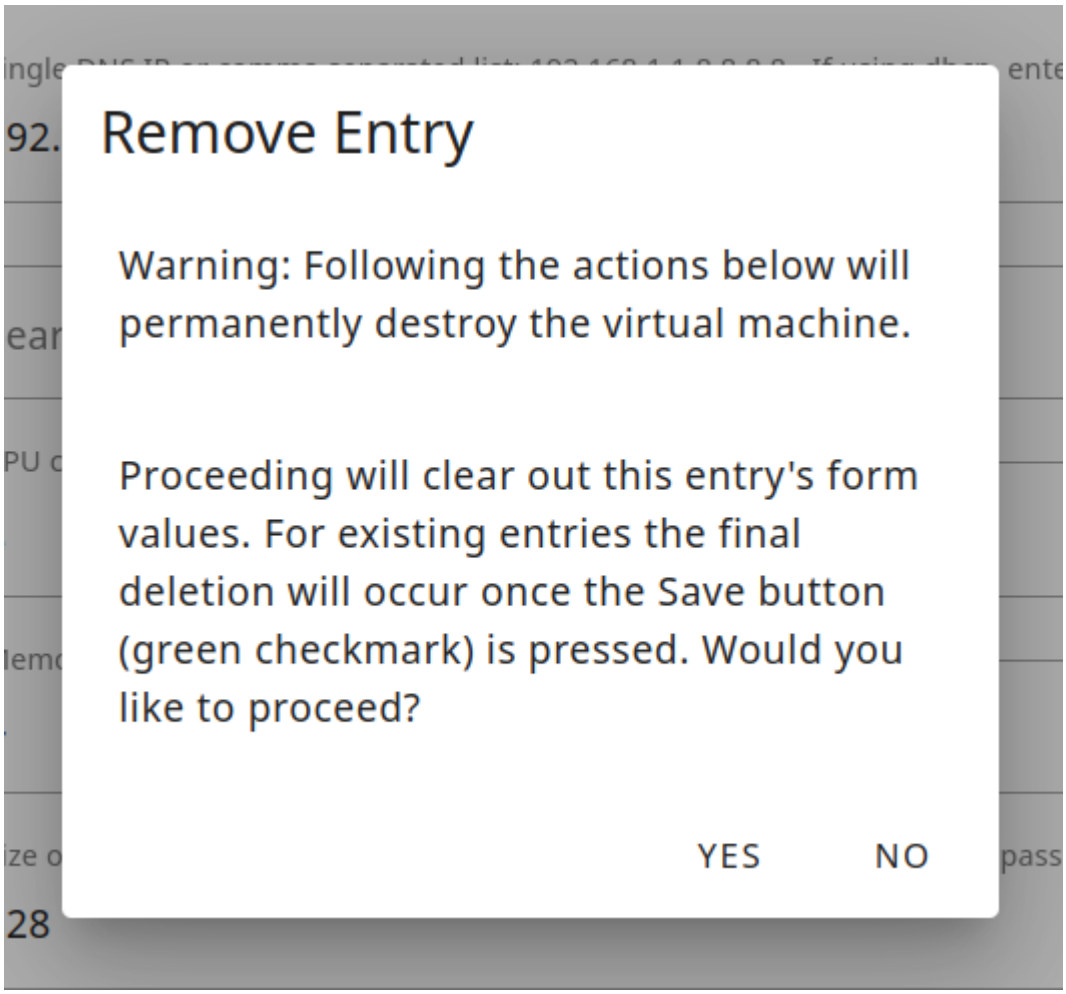

Once you've confirmed, it will show that it is pending deletion:

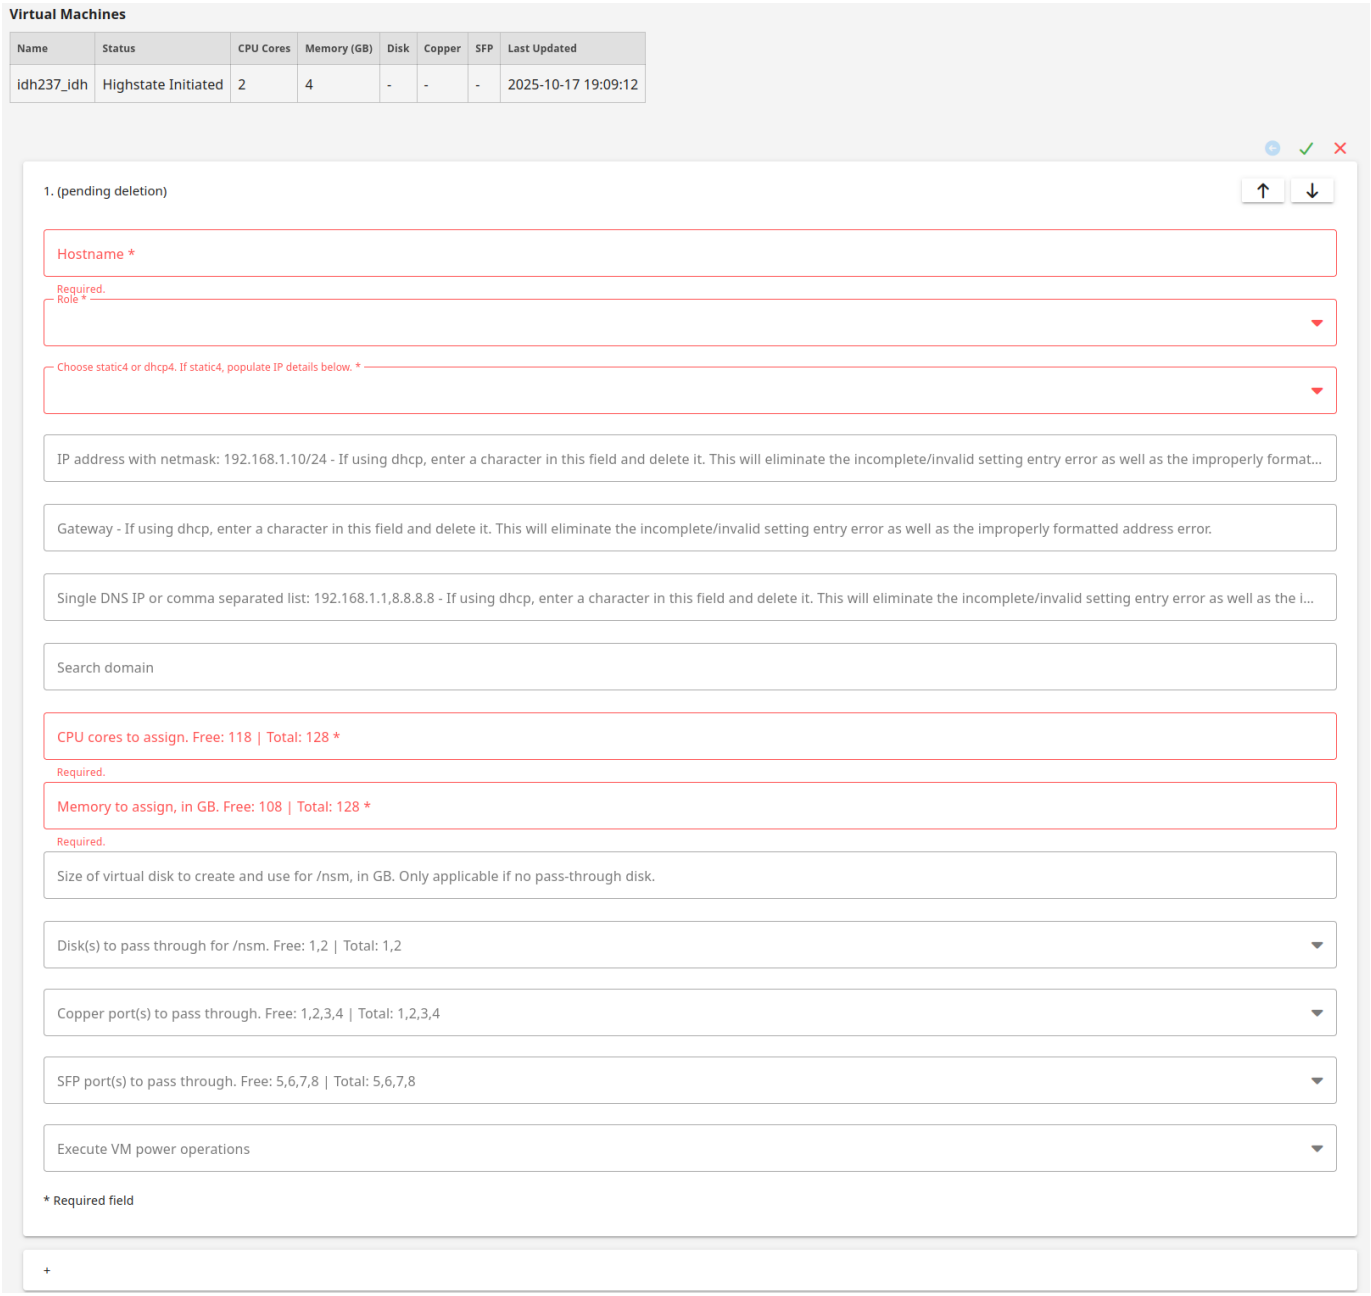

The backend then deletes the VM, releases the hardware used by the VM, and updates the hypervisor annotation used by SOC. This process should take less than 30 seconds.

You will then need to refresh your web browser to see that CPU, memory, and free hardware have been updated. The deleted VM should then be listed as `Destroyed Instance`.

> **Warning**
>
> If you delete a VM and attempt to immediately create a new VM prior to the backend releasing the hardware, then you will not be able to pass through the hardware that was previously used by the deleted VM.

21.17 Reports

Security Onion Pro customers can produce reports and data exports for analytical purposes. This is not a data migration or full data export feature but rather is intended for analysts to present relevant data to the broader team. Consequently, exports are limited in the amount of data that will be produced.

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license to enable this feature.

21.17.1 Reports

Reports are generated in Portable Document Format (PDF) in the background. Once a report is requested, a confirmation message will appear to notify the user that the report is being generated and will be available for download on the Reports screen.

Two types of reports are available: standard and custom.

Standard Reports

Standard reports are pre-built reports provided by the Security Onion team.

A *Productivity Report* provides general team productivity metrics. This includes:

- Ingested event metrics
- Alert metrics (total, acknowledged, escalated, by severity, etc)
- Case metrics (total, by status, by assignee, comments by user, etc)
- Time tracking (total case hours, case hours by user)

Productivity reports can be generated by navigating to the Reports interface and clicking the `Add Job` button near the top of the screen. A popup will appear allowing the user to select the Productivity report.

A *Case Report* provides detailed information about a specific case. This includes:

- ID, Title, description
- Current status, assignee, etc
- Hours logged against the case
- All case comments and the date and author
- All detections involved in the case
- All attachments (excluding actual file content)
- Observable data
- Related events
- Case audit history

Case reports can be generated by opening a case and clicking the Export icon above the Case title.

> **Note**
>
> Customer feedback will help the Security Onion team determine which additional reports will be provided in future releases. Provide feedback by reaching out to our Support team.

Standard reports can be lightly customized via the Configuration screen. Search for the report name, such as "Case report". Customizations include changing layout formatting, text content, and hiding or showing additional data. The ability to show additional data is limited to what the built-in report has already queried. For example, the built-in Case Report shows the TLP of each observable, but not the PAP. So if you want to see the PAP and not the TLP, then you could customize to do so. Remember to synchronize the grid after saving customizations.

Custom Reports

While the standard reports can be lightly customized, some customers may need more flexibility in customizing reports. For this, Security Onion Pro offers custom report types.

Custom reports are generated by navigating to the Reports interface and clicking the Add Job button. The popup will show a list of available custom reports.

Customers can create a limited number of custom reports by opening the Configuration screen and searching for `custom report`. Custom reports allow users to design their own report layout and choose which data is queried and then rendered on the report. Customizing these reports requires an understanding of markdown format (https://www.markdownguide.org/) and a general understanding of OQL syntax. For best results, report designers should build queries in the Hunt interface to ensure the data returned is what they want for their report. Then the report designer will copy that same OQL query into the custom report.

A sample report is provided in the standard Security Onion Pro installation. Here is the content of this report file followed by a brief review of each section.

```
{{- /* query.myDocEvents.Oql = metadata.type: _doc | groupby event.module, event.dataset | sortby @timestamp desc */ -}}
{{- /* query.myDocEvents.MetricLimit = 10 */ -}}
{{- /* query.myDocEvents.EventLimit = 100 */ -}}

Security Onion Custom Report
============================

{{ if .Error }}
**NOTE: This report encountered a problem extracting the relevant data and may not be complete.**

**Error:** {{.Error}}
{{ end }}

Records must have been created or updated during the following time frame in order to be reflected in this report.

**Report Start Date:** {{formatDateTime "Mon Jan 02 15:04:05 -0700 2006" .BeginDate}}

**Report End Date:** {{formatDateTime "Mon Jan 02 15:04:05 -0700 2006" .EndDate}}

## Sample Doc Events

**Total Events:** {{ formatNumber "%d" "en" .Results.myDocEvents.TotalEvents}}

### Event Counts By Module and Dataset

| Count | Proportion | Module | Dataset |
| ----- | ---------- | ------ | ------- |
{{ range sortMetrics "Value" "desc" .Results.myDocEvents.Metrics.groupby_0_event_module_event_dataset -}}
| {{ formatNumber "%.0f" "en" .Value}} | {{ formatNumber "%.1f" "en" .Percentage}}% | {{index .Keys 0}} | {{index .Keys 1}} |
{{end}}

### Individual Events (Limited to first {{.Results.myDocEvents.Criteria.EventLimit}})

| Event Time | Module | Dataset | Category |
| ---------- | ------ | ------- | -------- |
{{ range .Results.myDocEvents.Events -}}
| {{.Timestamp}} | {{.Payload.event_module}} | {{.Payload.event_dataset}} | {{.Payload.event_category}} |
{{end}}
```

Starting at the top: The first three lines are internal instructions to the report generator. They tell the report generator to run the OQL query `metadata.type: _doc | groupby event.module, event.dataset | sortby @timestamp desc` and store the results into the `.Results.myDocEvents` variable, for later reference in the report. The `MetricLimit = 10` parameter limits the groupby metrics to the top 10 counts, and the `EventLimit = 100` limits the raw event results to 100 at most. These are important parameters because without them you could end up with a PDF document hundreds of pages in length, and that much data is unlikely to be helpful when presenting security-related findings to others.

Pay close attention to the traditional template `{{ ... }}` syntax as well as the `{{- /* ... */ -}}` internal report instruction syntax. A missed character will cause the report generation to fail.

Next is the title of the report. Notice how the full length of the title is followed by an equally long series of equal `=` characters. This denotes that the previous line is the report title, and will be the name that appears in the Add Job popup list on the Reports interface. Refer to the Markdown guide reference earlier in this section for more help with defining titles, headers, sections, and other styling within your report document.

Following the title is a `{{ if .Error }} ... {{ end }}` block. This is useful for presenting errors encountered during the rendering of the report. If there is no error then this section will not be visible in the final generated report.

Next is the content of the report itself. You'll see report text content mixed with data outputs. For example, if 12 events matched the OQL query provided at the top of this report, then the line containing `**Total Events:**` will be rendered as:

Total Events: 12

Finally, there are two tabular outputs included. The first is showing "Event counts by Module and Dataset" and, according to markdown format, will be rendered in a table structure including the column headers "Count, Proportion, Module, and Dataset". Note that those `|` characters are used to separate the column data. The first column will output the Value of the first groupby Metric (referenced as `groupby_0_event_module_event_dataset`), which will be a "count" value. The `%.0f` is instructing the output to hide the decimal place values, and the `en` is instructing that the number be formatted in the English locale (1,220 vs 1.220). The second column will show the proportion of events that matched this row out of the full set of events, the third column will show the event module for that row, and the last column will show the event dataset for that row.

The second tabular output is showing the list of events (up to 100 as previously instructed at the top of the file). This is not a metric so there is no `.Value` count column on this table. Instead, the column will show the actual event field values. Notice how the first column is the event timestamp, which is useful for showing chronologically-ordered lists of events. The remaining columns can be any field that is associated with that event. Use Hunt to determine which fields are available for that query.

> **Note**
>
> Customizing reports can be intimidating to those not familiar with this level of detail. Pro customers with professional service hours can utilize the experience of the Security Onion support team to help get you started.

Remember to synchronize the grid after saving custom reports.

21.17.2 Tabular Exports

Also included with this feature is the ability to export data from several of the SOC screens, including:

- Alerts
- Dashboards
- Hunt
- Cases (listing)
- Detections (listing)

The exported data will be saved to a CSV file and include the headers as the first line of the CSV. The CSV will be generated in the background and will then be available for download in the Reports interface.

Group and Event data can be exported by clicking on the data export icon, typically found in the top-left corner of the Group metric table or graph, or the event table. While the export icon is visible on the graph mode of a group metric, it will still export the underlying data in CSV format. It does not attempt to export the graph visualization. However, users can use the browser's Print feature to send those graph visualizations and dashboards to a printer or PDF file.

CSV exports will export all underlying data, up to a configured max. Therefore, while the Dashboards interface could show 10 entries in a Group metric because of how the group limits are set, the CSV could contain more than the 10, if the backing data has more aggregated metrics available.

> **Warning**
>
> Exporting CSV data while in the relative time mode, such as 24 hours, will often result in exported CSV data that is different from what is on the screen at the time when the export was initiated. This is because the export is re-running the same query, using the same relative time but starting at the time of the export, versus when the dashboard was queried earlier. To export the identical data shown on the SOC interface first change the time range to absolute mode and define the start and end times prior to refreshing the results and clicking the export icon.

CSV exports are limited to 10,000 records by default.

> **Note**
>
> This CSV export feature is not intended for Elasticsearch data migration purposes.

21.17.3 Report Customization Reference

PDF Render Instructions

- Generate a table of contents at the beginning of the exported report:

```
{{- /* pdf_param: -generate-toc */ -}}
```

- Disable TLS verification when reaching out to pull external images for inclusion in the report:

```
{{- /* pdf_param: -insecure */ -}}
```

- Add a page break whenever it encounters a horizontal rule (HR) tag in the report, such as `---`:

```
{{- /* pdf_param: -new-page-on-hr */ -}}
```

- Change the report output layout from *portrait* (default) to *landscape*:

```
{{- /* pdf_param: -orientation */ -}}
{{- /* pdf_param: landscape */ -}}
```

OQL Query Instructions

- Define a new OQL query that filters by `some.field` and `another.field`, and includes a single aggregated metric of counts by both `event.modules` and `event.dataset` with a max limit of no more than 10 aggregations per groupby and no more than 25 total event records returned:

```
{{- /* query.mySpecialQuery.Oql = some.field: 123 AND another.field: "Jason" | groupby event.module, event.dataset */ -}}
{{- /* query.mySpecialQuery.MetricLimit = 10 */ -}}
{{- /* query.mySpecialQuery.EventLimit = 25 */ -}}
```

Accessing Query Results

- Loop through a query's event results. Note that all fields containing a period will have their periods replaced with an underscore to avoid colliding with the template syntax:

```
{{ range .Results.mySpecialQuery.Events -}}
{{.Timestamp}} | {{.some_field}} | {{ .another_field }}
{{ end }}
```

- Loop through a query's aggregation results. Note that aggregations have special fields available:
  - `.Value` The calculated aggregation, typically a summed count
  - `.Percentage` The calculated ratio * 100 of this group's occurrences compared to the total number of groups for this aggregation
  - `.Keys[#index]` Where is the 0-based index of the groupby fields. In this example, `.Keys 0` will be the `event.module` value, and `.Keys 1` will be the `event.dataset` value for this group.

```
{{ range sortMetrics "Value" "desc" .Results.mySpecialQuery.Metrics.groupby_0_event_module_event_dataset -}}
{{ formatNumber "%.0f" "en" .Value }} | {{ formatNumber "%.1f" "en" .Percentage }}% | {{ index .Keys 0 }} | {{ index .Keys 1 }} |
{{ end }}
```

If there had been a second `groupby` clause then it would have been referenced as `groupby_1_fieldx`.

Formatters

- Format a date and time into a specific format of *Mon Jan 02 15:04:05 -0700 2006*. Always reuse the same date and time as shown below for the format template:

```
{{ formatDateTime "Mon Jan 02 15:04:05 -0700 2006" .SomeDataField }}
```

- Format a numeric value to an integer, decimal, etc output, following the given locale, in this case *English*:

```
{{ formatNumber "%.2f" "en" .SomeNumericField }}
```

Sorters

- Sort the given metrics by the .Value field in descending order:

```
{{ range sortMetrics .Value "desc" .Results.mySpecialQuery.Metrics.groupby_1_fieldx }}
```

- Sort the given case data by the .CreateTime field (or event's timestamp) in ascending order (applicable to Case Report):

```
{{ range sortComments .CreateTime "asc" .Comments }}
{{ range sortDetections .CreateTime "asc" .Detections }}
{{ range sortArtifacts .CreateTime "asc" .Attachments }}
{{ range sortArtifacts .CreateTime "asc" .Observables }}
{{ range sortRelatedEvents "fields:soc_timestamp" "asc" .RelatedEvents }}
{{ range sortHistory .CreateTime "asc" .History }}
```

Filters

- Convert the given user ID into the user's email address. Applies to any valid user ID event field or indexed metric field:

```
{{ .SomeField | getUserDetail "email" }}
```

- Convert the value to uppercase:

```
{{ .SomeField | upper }}
```

- Convert the value to lowercase:

```
{{ .SomeField | lower }}
```

- Join an array value into a string, separated by commas:

```
{{ .SomeField | join "," }}
```

21.18 Onion AI

> **Note**
>
> This is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at https://securityonion.com/pro for more information about purchasing a Security Onion Pro license.

> **Note**
>
> Onion AI is disabled by default. To enable it for your Security Onion Pro deployment, go to the Administration --> Configuration page and navigate to soc --> config --> server --> client --> assistant --> enabled.

The Onion AI Assistant is your personal AI helper designed to assist you with a variety of tasks and provide information on demand. We support accessing LLMs from a variety of sources. Several tools have been made available as you interact with the assistant so it can access up-to-date information and resources.

21.18.1 Adapters

There are a variety of adapters available to connect to different LLM providers. When configuring an adapter, you're instructing SOC how to connect to an LLM provider. You configure the models separately in Available Models and tell each model which adapter to use. This allows you to connect to multiple providers at the same time and choose which models you want to use from each provider. The available adapters are:

- **SOAI**: The Security Onion AI adapter connects to our own hosted gateway and provides access to a variety of models. This is the default adapter, and is hosted in the cloud with an allotted number of credits per Security Onion Pro licensed grids.
- **Gemini**: Google's Gemini models can be accessed through this adapter. You can connect through either the Gemini Developer API or the Vertex API. These models are hosted in Google's cloud and require the use of your own Google Cloud key, or Gemini API key.
- **OpenAI Responses**: This adapter can connect to the newest OpenAI compatible APIs that support the Responses protocol. This adapter allows Security Onion grids to connect to locally-hosted LLMs.
- **OpenAI Chat**: This adapter can connect to any OpenAI Chat compatible APIs, an older protocol still supported by many AI providers. This adapter allows Security Onion grids to connect to locally-hosted LLMs.

Most features are supported across all adapters, but capabilities may change depending on the model or provider you connect to. Credits are unique to SOAI, it's the only adapter that responds with a credit balance. Multiple adapters can be configured at the same time.

Configuration

A fresh install will come with the SOAI adapter pre-configured and ready to use. To connect to other providers or modify the SOAI adapter, go to Administration --> Configuration page, be sure "Show advanced settings" is turned on in the Options at the top of the page, and check under soc --> config --> server --> modules --> assistant --> adapters. Here you can add as many adapters as you need. Each adapter defines how to connect to a provider. Different adapters have different configuration requirements. Every adapter needs a unique name, the protocol it uses, and a Health Timeout measured in seconds. Below is a table of which fields are required for which adapters:

	Adapter Name	Protocol	API Url	API Key	Service Account JSON	Service Locatio
SOAI	R	R	R	--	--	--
Gemini	R	R	--	O	O	O
OpenAI Responses	R	R	R	O	--	--
OpenAI Chat	R	R	R	O	--	--

R = Required, O = Optional, -- = Not Used

SOAI

This adapter should not require any configuration, aside from the API Url, which will be different for Security Onion Pro customers located in different regions. To change the SOAI API Url, go to the Administration --> Configuration page and navigate to soc --> config --> server --> modules --> assistant --> adapters. Select the SOAI adapter and update the API Url field. If your license is setup to use an alternate cloud region, the correct API Url will be included with your Security Onion license key. If an alternate URL is not specified, use the default API Url.

Any additionally supplied fields are ignored. Generally at most one instance of this adapter should be configured at a time.

CREDITS

Using models over the SOAI adapter will consume credits from your Security Onion Pro license. Credits are only applicable to the SOAI adapter. If you use a local model or another external provider (such as Gemini or OpenAI), SOAI credits are not consumed.

The Security Onion Pro license includes an initial amount of credits to get started. For long term usage planning contact your Security Onion account representative. They will assist with estimating credit usage rates as well as the provisioning of additional credits. Credits are consumed based on the number of tokens used in the conversation including user input, assistant output, and tool usage. Your organization's balance can be viewed at the top right of the assistant page or on the management page under Administration --> AI Metrics.

Gemini

Security Onion supports connecting to Google's Gemini models through either the Gemini Developer API or the Vertex API.

To connect over the Gemini Developer API, you will need to generate an API Key in the Google Cloud Console or Google AI Studio, and provide it in the API Key field. To connect over the Vertex API, you will need to create a Service Account with the appropriate permissions for the models you would like to access, generate a JSON key for that account, paste the JSON into the Service Account JSON field, and you must also specify the Service Account Location. Some Gemini models may only be available in certain locations, such as `global`. The API Url field is not used for Gemini connections.

If an API Key and Service Account JSON + Service Account Location are both provided, SOC will attempt to connect using the Service Account and the API Key will be ignored.

This adapter will not use any credits from your Security Onion Pro license, but Google will charge you based on their pricing. Security Onion Solutions, LLC does not provide support for Google Gemini, aside from assisting with the adapter configuration referenced on this page.

OpenAI Responses

The *OpenAI Responses* adapter can connect to the newest OpenAI compatible APIs that support the Responses protocol. To connect, you will need to provide an API Url. The API Url should be the base url for the provider you are connecting to, for example `https://api.openai.com/v1/`. An API Key may or may not be required depending on the provider you are connecting to. If you're unsure if your provider supports the Responses protocol, check with their support or try connecting with this adapter. If it doesn't work, you can try the *OpenAI Chat* adapter which supports an older protocol that is still widely used.

OpenAI Chat

The *OpenAI Chat* adapter can connect to any OpenAI compatible API that supports the Chat protocol. To connect, you will need to provide an API Url. The API Url should be the base url for the provider you are connecting to, for example `https://api.openai.com/v1/`. An API Key may or may not be required depending on the provider you are connecting to.

21.18.2 Available Models

In order to control what models are available inside SOC, individual models must be configured. To configure a model, go to the Administration --> Configuration page, be sure "Show advanced settings" is turned on in the Options at the top of the page, and check under soc --> config --> server --> client --> assistant --> availableModels. Here you can add as many models as you need. Each model needs a unique name, the adapter it uses, and a model identifier which tells the adapter which model to connect to on the provider.

> **Note**
>
> If a provider offers models that aren't listed in availableModels, they will not be accessible within SOC.

21.18.3 Local Model Considerations

Security Onion now supports local models through any OpenAI-compatible endpoint. Because the assistant relies on large context windows, the minimum recommended context length is 128k tokens. The following open-source models have been tested with OnionAI:

- **GPT OSS 120B** -- (US) Fast inference but limited to a 128k context window. Accuracy is fair.
- **Kimi 2.5** -- (China) A capable model with average accuracy and a 256k context window sufficient for most tasks. Note that this model requires significant VRAM to maintain performance.
- **GLM 5** -- (China) Average accuracy with a 200k context window.
- **Qwen 3.5** -- (China) Average accuracy with a 200k context window.

> **Note**
>
> Local models will not match the performance of proprietary foundational models such as those from Anthropic, Google, or OpenAI.

21.18.4 Hosting Local Models

Hosting your own models requires powerful and expensive hardware. For beginners we recommend using a tool such as LM Studio. **You need at least 96GB of VRAM** to host your own models locally. The speed and accuracy of OnionAI when hosted locally is based on the hardware that you are using. For the most accurate results we recommend using credits with OnionAI.

21.18.5 Available Tools

The assistant currently runs in the cloud. In order to reference local information, SOC makes the following tools available to the assistant:

- **query_events**: This read-only tool allows the assistant to query security events from your local Security Onion instance similar to how you would use the Hunt page.
- **query_cases**: This tool allows the assistant to query cases from your local Security Onion instance similar to how you would use the Cases page.
- **query_detections**: This tool allows the assistant to query detections from your local Security Onion instance similar to how you would use the Detections page.
- **get_playbooks**: When the assistant uses this read-only tool, SOC will gather the playbooks, execute their queries, and return all the data ready for analysis. This tool is read-only.
- **ack_alerts**: The assistant can use this tool to query and acknowledge Alerts in your local Security Onion instance.
- **escalate_alerts**: Similar to the ack tool, this tool allows the assistant to query for alerts and escalate them to a new case.
- **create_detection**: The assistant will use this tool to create new detections in your local Security Onion instance.
- **add_overrides**: This tool allows the assistant to tune existing detections by adding overrides to them.
- **toggle_detections**: The assistant can use this tool to enable or disable existing detections in your local Security Onion instance.
- **update_detection_content**: If a detection needs its content updated, the assistant can use this tool to update the content and keep the same metadata, overrides, and enabled status.
- **update_overrides**: This tool allows the assistant to update existing overrides for detections.

Permissions

When the assistant requests any of these tools, the user will be prompted to approve or reject the usage. Approval will have the tool run and the result entered into the conversation followed by a response from the AI. Rejection will notify the assistant and allow it to respond.

Auto Approve Read-Only Tools

The following read-only tools can be used by the assistant without requiring explicit user approval:

- query_events
- query_cases
- query_detections
- get_playbooks

By default, permission must be granted for each tool request. However, users can enable auto-approval for read-only tools by toggling the slider in the Options dropdown at the top of the assistant page. With the option enabled, read-only tools will auto approve, while tools that modify data will still request explicit user approval.

> **Note**
>
> Conversation data is stored locally and not used to train any models. However, be cautious about sharing sensitive or personal information during your interactions as others in your organization may have access to the conversation history.

21.18.6 Customizing the System Prompt

The system prompt sets the behavior and context for the AI assistant. Users can customize this prompt to tailor the assistant's responses to their specific needs, such as describing network resources, environments, or workflows. The prompt is included in every session and can be modified at any time on the Administration --> Configuration page under SOC --> config --> server --> modules --> assistant --> systemPromptAddendum.

Your system prompt addendum will be added after Security Onion's default system prompt.

> **Note**
>
> Be cautious when customizing the system prompt, as it can significantly influence the assistant's behavior and responses. A longer prompt will also use more credits.

21.18.7 Metrics

Superusers can review token usage and conversation history for all users by going to Administration --> AI Metrics. This page provides usage statistics for a given date range. The page starts with a table of usage by user. Clicking a user's binoculars icon on the right hand side will show any sessions the user interacted with during the selected date range, even deleted sessions. Clicking on a session's binoculars icon will show the full conversation. Administrators can adjust who has permissions via RBAC roles.

To provide an accurate history, deleted sessions are retained on the metrics page even after being deleted by the user.

22. Telemetry

22.1 SOC Telemetry

SOC can optionally send telemetry data to Google Analytics. This telemetry helps the Security Onion development team improve the product. Specifically, by knowing which user-interface features are being used, and how the user interacts with the SOC user interface, the development team can better prioritize new features, improvements to the existing user interface, and begin deprecating features that are rarely used. Deprecating unused features will help developers avoid spending their time and effort maintaining and upgrading areas of the product that aren't widely used. This allows more time to be spent on new features and bug fixes, directly benefiting Security Onion users.

22.1.1 Configuring

During setup, or during non-automated soup invocations, the user must choose to enable or disable SOC telemetry collection.

After installation, Grid administrators can enable or disable SOC Telemetry via the configuration interface. Search for `SOC Telemetry` in the Configuration screen.

After changing the `SOC Telemetry` configuration setting, the grid must be resynchronized. This happens automatically once every 15 minutes, or manually if the grid administrator clicks the Synchronize Grid option, at the top of the Configuration screen. Grid synchronization can take several minutes to complete.

Also, the browser will cache the previous configuration setting. Therefore, to ensure the browser is using the new setting value, make sure Grid synchronization completes, and then perform a hard browser refresh on the SOC UI. This can be performed via CTRL+SHIFT+R or CMD+SHIFT+R.

22.1.2 Included Data

The primary objective of the SOC Telemetry data collection is to understand what features are being used in SOC. Specifically, the data being collected relates to user interface navigation flows. Additional context data, such as the version of the software, the viewed page path, etc. may also be included.

Grid-specific data is intentionally excluded from the data collection, where possible.

Due to the nature of how internet web requests work, the originating IP address and User-Agent information of the web browser as well as host referrer information is known at the time of the data collection.

See the Retention Period section below to learn more about how long this data is retained.

22.1.3 Network Parameters

The SOC Telemetry data originates from the SOC browser running on the analyst workstation. The domains being accessed from the SOC browser are:

- `www.googletagmanager.com`
- `www.google-analytics.com`

The telemetry data is sent using TLS encryption.

22.1.4 Retention Period

Data stored in Google Analytics is configured to be automatically removed after two months. This is the shortest interval that Google Analytics provides.

22.2 Operating System Updates Telemetry

Security Onion periodically checks for package updates to ensure the operating system (OS) and related applications are kept patched and updated. These updates are critical to protecting the OS and installed packages from recently exposed vulnerabilities. This is different from Security Onion product upgrades via soup, which is manually invoked. When the OS package updates begin, the request to the Security Onion repository server(s) can include a limited set of telemetry data.

22.2.1 Configuring

Automatic package updates can be enabled or disabled via the configuration interface. Search for `patch` in the Configuration screen.

After changing this configuration setting, the grid must be resynchronized. This happens automatically once every 15 minutes, or manually if the grid administrator clicks the Synchronize Grid option, at the top of the Configuration screen. Grid synchronization can take several minutes to complete.

22.2.2 Included Data

The primary objective of the included telemetry data is to understand which versions of Security Onion are deployed and on which platforms. This information helps the Security Onion development team determine how to prioritize support for older versions, and whether there is justification to start testing or stop testing on various operating systems and architectures.

Additional data may be included to provide the development team with information about which features have been enabled. This data can change from release to release as it often relates to new development work.

Also, grids with license keys installed will include the license key identifier. Grids using the standard unprovisioned license do not have a license key identifier.

Due to the nature of how internet web requests work, the originating IP address and User-Agent information of the web browser as well as host referrer information is known at the time of the data collection.

22.2.3 Network Parameters

The OS Updates Telemetry data originates from the manager node. The domains being accessed from the manager node are:

- `sigs.securityonion.net`
- `repo.securityonion.net`
- `repo-alt.securityonion.net`

The telemetry data is sent using TLS encryption.

22.3 Airgap

Grids installed within airgapped environments will automatically disable telemetry. In this scenario, the `SOC Telemetry` configuration setting will have no effect and the automatic package updates will be disabled. See the Airgap section for more information about environments detached from the internet.

> **Note**
>
> If a grid is switched from airgap to non-airgap, and if the SOC Telemetry is not explicitly disabled in the grid by an administrator, the SOC app running in the browser will send telemetry.

23. Security

23.1 Vulnerability Disclosure

If you have any security concerns regarding Security Onion or believe you have uncovered a vulnerability, please send an email to security@securityonion.net per the following guidelines:

- Include a description of the issue and steps to reproduce
- Use plain text format in the email (no Word documents or PDF files)

Please do NOT disclose publicly until we have had sufficient time to resolve the issue.

> **Note**
>
> This security address should be used only for undisclosed vulnerabilities. Dealing with fixed issues or general questions on how to use Security Onion should be handled via the normal Support channels.

23.2 Beg Bounties

We do not participate in "beg bounties":

https://www.troyhunt.com/beg-bounties/

23.3 Product and Supply Chain Integrity

Security Onion is based on free and open software. Third-party components, as well as the software that the Security Onion team develops, is built from source code that is readily available for the public to review. Community contributors, or anyone for that matter, can request to have notifications pushed to them when a change is accepted into the public repositories. This is very different from closed source software since those closed source code bases are only visible to a very small group of developers. Further, if a closed source code base does not have formal code review procedures in place, or lacks infrastructure around the code base to make others aware of new changes, this further restricts visibility and review of changes. These deficiencies allow attackers that gain unauthorized access to a closed source code base to make changes without others detecting it.

When upstream, third-party components are updated in Security Onion, the change requires multiple checks before it can be merged into the master (released) branch. First, all commits must be signed using cryptography before being allowed into the master branch. Second, code reviews and approvals from multiple team members are required before the pull requests can be merged. Both of these restrictions are enforced by the source code repository itself, which eliminates risk of a human mistake allowing the process to be bypassed. Further, changes to the Security Onion source code repositories cause notifications to be delivered to the Security Onion development team, as well as anyone in the public who choose to be notified of such changes. On top of this, Security Onion developers are required (enforced by the repository itself) to use multi-factor authentication in order to approve changes.

Additionally, Security Onion's build infrastructure runs both unit level tests and fully automated end-to-end tests on every release, to ensure the Security Onion platform, and its components, continue to operate as expected. When a change is merged into Security Onion, whether it's to upgrade an upstream component or a modification to the source code maintained by the Security Onion developers, which breaks the automated tests, we are notified and take action to review the failure and root cause. Often this results in our developers chasing down upstream code commits to find out why something changed, and if it was intended or not. Fortunately, these investigations are typically bug related, rather than malicious, and our team will either contribute a pull request to fix the upstream project, or file an issue to raise awareness to the project maintainers.

There is no guarantee that any software, open or closed source, will always be free from attacks. However, our commitment to open software, and our investments into repeatable processes and software automation and testing technologies improves Security Onion's posture when it comes to safe guarding the product and its user-base.

24. Software Bill of Materials

The following table lists the major software projects integrated into the current version of Security Onion.

Product	Version	Author	Project URL	License	Description
Alpine Linux	3.23.3	Alpine Linux Development Team	https://alpinelinux.org/	GNU GPL Version 3	Alpine Linux is a security-oriented, lightweight Linux distribution based on musl libc and busybox.
ATT&CK Navigator	5.3.0	MITRE	https://github.com/mitre-attack/attack-navigator	Apache License 2	The ATT&CK Navigator is designed to provide basic navigation and annotation of ATT&CK matrices, something that people are already doing today in tools like Excel.
CyberChef	10.22.1	GCHQ	https://github.com/gchq/CyberChef	Apache License 2	The "Cyber Swiss Army Knife" - a web app for encryption, encoding, compression and data analysis.
Docker	29.2.1	Docker	https://github.com/docker	Apache License 2	Docker is a set of platform as a service products that use OS-level virtualization to deliver software in packages called containers.
ElastAlert	2.28.0	Jason Ertel	https://github.com/jertel/elastalert2	Apache License 2	ElastAlert is a simple framework for alerting on anomalies, spikes, or other patterns of interest from data in Elasticsearch.
Elastic Agent	9.0.8	Elastic	https://github.com/elastic/elastic-agent	Elastic License Version 2	Beats is a free and open platform for single-purpose data shippers. They send data from hundreds or thousands of machines and

Product	Version	Author	Project URL	License	Description
					systems to Logstash or Elasticsearch.
Elasticsearch	9.0.8	Elastic	https://github.com/elastic/elasticsearch	Elastic License Version 2	Elasticsearch is a distributed, open source search and analytics engine for all types of data, including textual, numerical, geospatial, structured, and unstructured.
evtx	0.11.0	Omer Benamram	https://github.com/omerbenamram/evtx	MIT License	A cross-platform parser for the Windows XML EventLog format.
evtx2es	1.8.0	Shinta Nakano	https://github.com/Security-Onion-Solutions/evtx2es	MIT License	A fast library for parsing and importing Windows Event Logs into Elasticsearch.
ExifTool	12.60	Phil Harvey	https://github.com/exiftool/exiftool	GNU GPL Version 3	ExifTool is a platform-independent Perl library plus a command-line application for reading, writing and editing meta information in a wide variety of files.
Hydra	25.4.0	Ory	https://github.com/ory/hydra	Apache License 2	Ory Hydra is a hardened, OpenID Certified OAuth 2.0 Server and OpenID Connect Provider optimized for low-latency, high throughput, and low resource consumption
InfluxDB	2.7.12	InfluxData	https://github.com/influxdata/influxdb/tree/main-2.x	MIT License	InfluxDB is an open source time series platform. This includes APIs for storing and querying data, processing it in the

Product	Version	Author	Project URL	License	Description
					background for ETL or monitoring and alerting purposes, user dashboards, and visualizing and exploring the data and more.
Kibana	9.0.8	Elastic	https://github.com/elastic/kibana	Elastic License Version 2	Kibana is an open source frontend application that sits on top of the Elastic Stack, providing search and data visualization capabilities for data indexed in Elasticsearch.
Kratos	25.4.0	Ory	https://github.com/ory/kratos	Apache License 2	Ory Kratos is the developer-friendly, security-hardened and battle-tested Identity, User Management and Authentication system for the Cloud.
Logstash	9.0.8	Elastic	https://github.com/elastic/logstash	Elastic License Version 2	Logstash is a server-side data processing pipeline that ingests data from a multitude of sources simultaneously, transforms it, and then sends it to your favorite "stash."
NetworkMiner	2.8.1	Netresec	https://www.netresec.com/?page=NetworkMiner	GNU GPL Version 2	NetworkMiner is an open source Network Forensic Analysis Tool (NFAT) and can be used as a passive network sniffer/packet capturing tool in order to detect operating systems,

Product	Version	Author	Project URL	License	Description
					sessions, hostnames, open ports etc. without putting any traffic on the network. NetworkMiner can also parse PCAP files for off-line analysis and to regenerate/ reassemble transmitted files and certificates from PCAP files. Only included in Security Onion Desktop.
Nginx	1.29.6	NGINX	https://github.com/nginxinc/docker-nginx	2-Clause BSD License	Nginx is a web server that can also be used as a reverse proxy, load balancer, mail proxy and HTTP cache.
OpenCanary	0.9.7	Thinkst Canary	https://github.com/thinkst/opencanary	3-Clause BSD License	OpenCanary is a multi-protocol network honeypot. It's primary use-case is to catch hackers after they've breached non-public networks.
Oracle Linux 9	9.7	Oracle Linux	https://www.oracle.com/linux/	GNU GPL Version 2	Oracle Linux is a Linux distribution packaged and freely distributed by Oracle, available partially under the GNU General Public License since late 2006. It is compiled from Red Hat Enterprise Linux source code, replacing Red Hat branding with Oracle's.
pcapfix	1.1.7	Robert Krause	https://github.com/Rup0rt/pcapfix/	GNU GPL Version 3	Pcapfix is a tool to repair your damaged or

Product	Version	Author	Project URL	License	Description
					corrupted pcap and pcapng files.
Python	3.14.3	Python Software Foundation	https://github.com/python/	Python Software Foundation License Version 2	Python is a programming language that lets you work quickly and integrate systems more effectively.
Redis	7.2.13	Redis	https://github.com/redis/redis	3-Clause BSD License	Redis is an open source in-memory data structure store used as a database, cache and message broker.
Salt	3006.19	Salt Project	https://github.com/saltstack/salt	Apache License 2	Salt is a distributed remote execution system used to execute commands and query data. It was developed in order to bring the best solutions found in the world of remote execution together and make them better, faster and more malleable. Salt accomplishes this via its ability to handle larger loads of information, and not just dozens, but hundreds or even thousands of individual servers, handle them quickly and through a simple and manageable interface.
Security Onion Console	3.0.0	Security Onion Solutions	https://github.com/Security-Onion-Solutions/securityonion-soc	Elastic License Version 2	Security Onion Console is a web interface to viewing alerts,

Product	Version	Author	Project URL	License	Description
					hunting, and PCAP analysis.
Strelka	1.0.1	Target	https://github.com/target/strelka	Apache License 2	Strelka is a real-time file scanning system used for threat hunting, threat detection, and incident response. Strelka's purpose is to perform file extraction and metadata collection at huge scale.
Suricata	8.0.4	OISF	https://github.com/OISF/suricata	GNU GPL Version 2	Suricata is a free and open source, mature, fast and robust network threat detection engine. Suricata inspects the network traffic using a powerful and extensive rules and signature language.
Telegraf	1.38.0	InfluxData	https://github.com/influxdata/telegraf	MIT License	Telegraf is a server-based agent for collecting and sending all metrics and events from databases, systems, and IoT sensors. Telegraf is written in Go and compiles into a single binary with no external dependencies, and requires a very minimal memory footprint.
Ubuntu	22.04, 26.04	Canonical	https://github.com/canonical	Canonical's IPRights Policy	The number 1 open source operating system powers millions of PCs and

Product	Version	Author	Project URL	License	Description
					laptops around the world.
Yara	4.3.1	VirusTotal	https://github.com/virustotal/yara	3-Clause BSD License	YARA is a tool aimed at (but not limited to) helping malware researchers to identify and classify malware samples.
Wireshark	3.4.10	Wireshark	https://github.com/wireshark/wireshark	GNU GPL Version 2	Wireshark is the world's foremost and widely-used network protocol analyzer. It lets you see what's happening on your network at a microscopic level and is the de facto (and often de jure) standard across many commercial and non-profit enterprises, government agencies, and educational institutions. Only included in Security Onion Desktop.
Zeek	8.0.6	Zeek	https://github.com/zeek/zeek/	3-Clause BSD License	A powerful framework for network traffic analysis and security monitoring.

25. Release Notes

25.0.1 Known Issues

For all known issues, please see https://github.com/Security-Onion-Solutions/securityonion/issues.

25.0.2 Release History

25.1 3.0.0 [20260331] Changes

- FEATURE: Configurable Elasticsearch vm.max_map_count setting
- FEATURE: Dynamically load Zeek plugins on zeek startup #15546
- FEATURE: Enable JA4+ License Acceptance #15560
- FEATURE: Parsing for Zeek websockets logs #15657
- FEATURE: Refresh login page with updated look
- FEATURE: Refresh SOC UI with updated look
- FEATURE: Support additional alt names in web cert
- FEATURE: Support docker ulimit customization #15581
- FEATURE: Suricata PCAP replacing Stenographer
- FIX: API 401 errors will no longer redirect #15611
- FIX: Cleanup file.absent and cron.absent
- FIX: Detections - Intermittent "error closing scroll" #14216
- FIX: Duplicated user roles when refreshing frontend at Administration > Users #15688
- FIX: Enabled / Disabled Buttons for SOC Grid Configuration Options #15649
- FIX: Fix rule validators in SOC #15533
- FIX: Global override configs should not apply to certain indices #15601
- FIX: Network Transport for suricata alerts should be lowercase #15668
- FIX: Sensors are not checking in while processing long jobs #15650
- FIX: so-suricata-testrule script #15396
- FIX: STIG V1R3
- FIX: Suricata address-groups vars allow negation #15664
- FIX: Unable to create detections via Connect API #15673
- UPGRADE: All frontend 3rd party deps
- UPGRADE: ATTACK Navigator to 5.3.0 #15680
- UPGRADE: CyberChef to 10.22.1 #15681
- UPGRADE: ElastAlert2 to 2.28.0 #15685
- UPGRADE: Golang 3rd party deps #15647
- UPGRADE: Golang to 1.26.1 #15580
- UPGRADE: Hydra to 25.4.0 #15678
- UPGRADE: Kafka to 3.9.2 #15684
- UPGRADE: Kratos to 25.4.0 #15677
- UPGRADE: Nginx to 1.29.6 #15686
- UPGRADE: OpenCanary to 0.9.7 #15679

- UPGRADE: Redis to 7.2.13 #15682
- UPGRADE: Suricata to 8.0.4 #15625
- UPGRADE: Telegraf to 1.38.0 #15683
- UPGRADE: Update Docker base images

26. Appendix

This appendix provides an overview of the process of upgrading from the old Security Onion 2.4 to the new Security Onion 3.

Tip

If you are a current Security Onion Solutions customer with Professional Services or Appliance coverage, contact SOS support and we can help you through this process.

Warning

Security Onion 3 only supports Oracle Linux 9. If you are running Security Onion 2.4 on some other unsupported distro, then you will need to perform a fresh installation of Security Onion 3.

Warning

We recommend trying this process in a test environment before attempting in your production environment.

If you have reviewed all of the warnings above and are ready to upgrade, you should be able to do the following.

First, make sure that your 2.4 installation is fully updated via soup:

```
sudo soup
```

Next, make sure there is a backup in /nsm/backup:

```
sudo ls -alh /nsm/backup
```

Once you have confirmed your backup, you may upgrade to Security Onion 3.

To begin, you will need to upgrade to the new version of soup:

```
sudo soupto3
```

Once you have the new version of soup, then you will need to run it to perform the upgrade:

```
sudo soup
```

27. Cheat Sheet

If you are viewing the online version of this documentation, you can click here for our Security Onion Cheat Sheet.

This was based on a cheat sheet originally created by Chris Sanders which can be found here:

https://chrissanders.org/2017/06/security-onion-cheat-sheet/