
Security Onion Documentation

Release 2.4

Security Onion Solutions, LLC

Apr 11, 2024

TABLE OF CONTENTS

1	About	1
1.1	Security Onion	1
1.2	Security Onion Solutions, LLC	1
1.3	Documentation	1
2	Introduction	3
2.1	Network Visibility	5
2.2	Host Visibility	5
2.3	Analysis Tools	6
2.4	Workflow	9
2.5	Deployment Scenarios	10
2.6	Conclusion	10
3	License	11
4	First Time Users	13
5	Getting Started	35
5.1	Best Practices	36
5.2	Architecture	37
5.3	Hardware Requirements	47
5.4	Download	53
5.5	VMware	53
5.6	VirtualBox	55
5.7	Proxmox	56
5.8	Booting Issues	57
5.9	Airgap	58
5.10	Installation	59
5.11	Amazon Cloud Image	60
5.12	Azure Cloud Image	66
5.13	Google Cloud Image	70
5.14	Configuration	77
5.15	After Installation	78
6	Security Onion Console (SOC)	81
6.1	Alerts	83
6.2	Dashboards	90
6.3	Hunt	100
6.4	Cases	100
6.5	PCAP	110
6.6	Grid	115

6.7	Downloads	122
6.8	Administration	123
6.9	Kibana	127
6.10	Elastic Fleet	130
6.11	Osquery Manager	137
6.12	InfluxDB	138
6.13	CyberChef	139
6.14	Playbook	141
6.15	ATT&CK Navigator	145
7	Security Onion Desktop	149
7.1	Chromium	151
7.2	NetworkMiner	151
7.3	Wireshark	152
8	Network Visibility	153
8.1	AF-PACKET	154
8.2	Stenographer	154
8.3	Suricata	157
8.4	Zeek	164
8.5	Strelka	169
8.6	Intrusion Detection Honeypot	171
9	Host Visibility	175
9.1	Elastic Agent	175
9.2	Syslog	177
9.3	Sysmon	178
10	Logs	181
10.1	Ingest	181
10.2	Logstash	183
10.3	Redis	187
10.4	Elasticsearch	189
10.5	ElastAlert	195
10.6	Curator	198
10.7	Data Fields	199
10.8	Alert Data Fields	199
10.9	Elastalert Fields	200
10.10	Zeek Fields	201
10.11	Community ID	201
10.12	SOC Logs	202
11	Updating	205
11.1	soup	205
11.2	End Of Life	210
12	Accounts	211
12.1	Passwords	211
12.2	MFA	213
12.3	Adding Accounts	214
12.4	Listing Accounts	215
12.5	Disabling Accounts	216
12.6	Role-Based Access Control (RBAC)	217
12.7	Kratos	222
12.8	OpenID Connect (OIDC)	223

13 Services	237
14 Customizing for Your Environment	239
14.1 SOC Customization	239
14.2 nginx	242
14.3 Proxy	244
14.4 Firewall	246
14.5 Email	250
14.6 NTP	250
14.7 Console	251
14.8 SSH	252
14.9 Hostname	252
14.10 IP Address	252
14.11 Web Access URL	252
15 Tuning	255
15.1 BPF	255
15.2 Managing Rules	257
15.3 Adding Local Rules	260
15.4 Managing Alerts	261
15.5 High Performance Tuning	266
15.6 Salt	267
16 Tricks and Tips	269
16.1 Backup	269
16.2 Docker	270
16.3 Jupyter Notebook	274
16.4 Adding a new disk	277
16.5 Network Installation	278
16.6 PCAPs for Testing	281
16.7 Removing a Node	282
16.8 Syslog Output	283
16.9 UTC and Time Zones	283
16.10 pfSense	283
16.11 Endgame	285
17 Utilities	287
17.1 jq	287
17.2 so-allow	287
17.3 so-elastic-auth-password-reset	288
17.4 so-elasticsearch-query	288
17.5 so-import-pcap	289
17.6 so-import-evtx	291
17.7 so-monitor-add	291
17.8 so-status	291
17.9 so-test	292
17.10 so-user	293
18 Help	295
18.1 FAQ	295
18.2 Directory Structure	299
18.3 Tools	300
18.4 Support	301
18.5 Community Support	301
18.6 Help Wanted	302

19 Security	305
19.1 Vulnerability Disclosure	305
19.2 Product and Supply Chain Integrity	305
20 Release Notes	307
20.1 Known Issues	307
21 Appendix	319
22 Cheat Sheet	323

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://readthedocs.org/projects/securityonion/downloads/pdf/2.4/>.

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.

To submit a new page, you can submit a pull request (PR) to the 2.4 branch of the `securityonion-docs` repo at <https://github.com/Security-Onion-Solutions/securityonion-docs>.

Pages are written in RST format and you can find several RST guides on the Internet including https://thomas-cokelaer.info/tutorials/sphinx/rest_syntax.html.

1.3.5 Naming Convention

New documentation pages should use the following naming convention:


- all lowercase
- `.rst` file extension
- ideally, the name of the page should be one simple word (for example: `suricata.rst`)
- try to avoid symbols if possible
- if symbols are required, use hyphens (NOT underscores)

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 your choice of either *Zeek* or *Suricata*, full packet capture, 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*.

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.



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 recorded by *Stenographer*.

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 *Intrusion Detection Honeypot* 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 (SOC)

Security Onion Console (SOC) 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*.


The screenshot shows the Security Onion Alerts console. The interface includes a sidebar with navigation options like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, and Administration. The main area displays a table of alerts with the following columns: Count, rule.name, event.module, and event.severity_label. The table lists various alerts such as 'ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)' with a count of 59 and 'ET MALWARE Trickbot Checkin Response' with a count of 4. The interface also features search filters, a time range selector (2021/06/30 00:00:00 AM - 2021/07/01 00:00:00), a Refresh button, and a Fetch Limit of 50.

Count	rule.name	event.module	event.severity_label
59	ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)	suricata	low
4	ET MALWARE Trickbot Checkin Response	suricata	high
4	ET POLICY HTTP traffic on port 443 (POST)	suricata	medium
3	ET HUNTING GENERIC SUSPICIOUS POST to Dotted Quad with Fake Browser 1	suricata	medium
3	ET MALWARE VNCStartServer BOT Variant CnC Beacon	suricata	high
3	ET MALWARE VNCStartServer USR Variant CnC Beacon	suricata	high
3	ET POLICY PE EXE or DLL Windows file download HTTP	suricata	high
2	ET HUNTING curl User-Agent to Dotted Quad	suricata	medium
2	ET INFO Dotted Quad Host DLL Request	suricata	medium
2	ET MALWARE Win32/Trickbot Data Exfiltration M2	suricata	high
2	ET POLICY curl User-Agent Outbound	suricata	medium
1	ET DNS Query to a *.top domain - Likely Hostile	suricata	medium
1	ET EXPLOIT ETERNALBLUE Probe Vulnerable System Response MS17-010	suricata	high
1	ET EXPLOIT Possible ETERNALBLUE Probe MS17-010 (Generic Flags)	suricata	high
1	ET HUNTING Suspicious POST with Common Windows Process Names - Possible Process List Exfiltration	suricata	high
1	ET HUNTING Suspicious Windows Commands in POST Body (ipconfig)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net config)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net view)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (nltest)	suricata	medium

Security Onion Console (SOC) 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.

The screenshot shows the Security Onion console interface for the 'Cases' section. The top navigation bar includes 'Overview', 'Alerts', 'Dashboards', 'Hunt', 'Cases', 'PCAP', 'Grid', 'Downloads', and 'Administration'. The 'Cases' section is active, displaying a search bar with 'Open Cases' and a filter for 'Last 12 months'. Below the search bar, there are two filter buttons: 'NOT so_case.status:closed' and 'NOT so_case.category:template'. The main content area shows a table of cases:

Timestamp	so_case.title	so_case.status	so_case.severity	so_case.assigneeid	so_case.createTime
2022-04-18 19:08:32.859 +00:00	Attempts to exploit log4j vulnerability against public facing web servers in DMZ	in progress	high	jim@example.com	2022-01-19T21:14:17.101028031Z
2022-04-18 19:07:33.320 +00:00	John Doe in Accounting received phishing email	new	critical	bill@example.com	2022-01-19T21:12:04.652244498Z
2022-01-19 21:11:18.165 +00:00	SQL injection attempts against web servers in DMZ	in progress	medium	jim@example.com	2022-01-19T21:08:48.897561928Z

At the bottom of the table, it indicates 'Rows per page: 50' and '1-3 of 3'.

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

The screenshot shows the Security Onion console interface for PCAP retrieval. The top navigation bar is the same as in the previous screenshot. The 'PCAP' section is active, displaying a search bar with '# 1001', 'securityonion', and '176.10.125.8:80'. Below the search bar, there are filter results and a table of captured packets:

Num	Timestamp	Type	Source IP	Source Port	Destination IP	Destination Port	Flags	Length
0	2021-06-30 20:48:37.602 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	SYN	66
1	2021-06-30 20:48:37.760 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	SYN ACK	58
2	2021-06-30 20:48:37.760 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	ACK	54
3	2021-06-30 20:48:37.760 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	PSH ACK	148
4	2021-06-30 20:48:37.760 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	ACK	54
5	2021-06-30 20:48:37.927 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	ACK	1514
6	2021-06-30 20:48:37.927 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	PSH ACK	1358
7	2021-06-30 20:48:37.927 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	ACK	54
8	2021-06-30 20:48:37.930 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	ACK	1514
9	2021-06-30 20:48:37.930 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	PSH ACK	1370

At the bottom of the table, it indicates 'Rows per page: 10' and '1-10 of 24'. There is a 'LOAD MORE' button below the table.

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.

The screenshot shows the CyberChef interface with a recipe named 'From Hexdump'. The recipe steps are: 'From Hexdump', 'Strip HTTP headers', 'Strip HTTP headers', and 'Strings'. The 'Strings' step is currently active, showing a list of strings extracted from the input. The input field contains hex data, and the output field shows the decoded text, including 'This is a Windows NT windowed dynamic link library'.

2.3.3 Playbook

Playbook allows you to create a Detection Playbook, which itself consists of individual plays. These plays are fully self-contained and describe the different aspects around the particular detection strategy.

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.
- 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.
- Develop a play in *Playbook* that will automatically alert on observables moving forward and update your coverage in *ATT&CK Navigator*.
- 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* 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!

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>

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, download our ISO image as shown in the [Download](#) section.

Then install the ISO image and configure for IMPORT as shown below (also see the [Installation](#) and [Configuration](#) sections). This can be done in a minimal virtual machine with as little as 4GB RAM, 2 CPU cores, and 200GB of storage. For more information about virtualization, please see the [VMware](#), [VirtualBox](#), and [Proxmox](#) sections.

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, specify your username and password:

```
#####
##          ** W A R N I N G **          ##
##          _____                    ##
##  Installing the Security Onion ISO    ##
## on this device will DESTROY ALL DATA ##
##          and partitions!              ##
##          ** ALL DATA WILL BE LOST **  ##
#####
Do you wish to continue? (Type the entire word 'yes' to proceed.) yes

A new administrative user will be created. This user will be used for setting up and administering Security Onion.

Enter an administrative username: doug


Let's set a password for the doug user:

Enter a password:
Re-enter the password:
```


Once installation is complete, you are prompted to reboot:

```
Initial Install Complete. Press [Enter] 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 (SOC)*:




Set password for *Security Onion Console (SOC)*:




Confirm password for *Security Onion Console (SOC)*:




Select how to access *Security Onion Console (SOC)*:




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 (SOC)*:



After logging in, you will see the *Security Onion Console (SOC)* Overview page:

Getting Started

New to Security Onion 2? Click the menu in the upper-right corner and you'll find links for [Help](#) and a [Cheat Sheet](#) that will help you best utilize Security Onion to hunt for evil! In addition, check out our free Security Onion 2 Essentials online course, available on our [Training](#) website.

If you're ready to dive in, take a look at the [Alerts](#) interface to see what Security Onion has detected so far. Then go to the [Dashboards](#) interface for a general overview of all logs collected or go to the [Hunt](#) interface for more focused threat hunting. Once you've found something of interest, escalate it to [Cases](#) to then collect evidence and analyze observables as you work towards closing the case.

What's New

To see all the latest features and fixes in this version of Security Onion, click the upper-right menu and then click the [What's New](#) link.

Enterprise Appliances

Want the best hardware for your enterprise deployment? Check out our [enterprise appliances!](#)

Customize This Space

Make this area your own by customizing the content in the [Config](#) interface.

Check *Grid* to verify all services are running properly:

ID	Role	Address	Version	Model	EPS	Last Heard From	Age	Status
securityonion	Import	192.168.199.143	2.4.60	N/A	0	a few seconds ago	an hour	OK

Node Status

ID: securityonion
 Role: Import
 Address: 192.168.199.143
 Version: 2.4.60
 Model: N/A
 Date Created: 2024-03-12 14:28:21.385 +...
 Last Heard From: 2024-03-12 15:41:01.742 +...
 Age: an hour
 OS Uptime: an hour
 Last Synchronized: 12 minutes ago
 Process Status: OK
 Connection Status: OK
 Elasticsearch Status: OK
 RAID Status: Feature Unavailable
 Consumption EPS: 0
 Memory Usage: 72.0% of 4.0 GB
 Swap Usage: 21.3% of 8.6 GB
 CPU Usage: 3.6%
 I/O Wait: 0.8%
 Root Partition Usage: 19.1% of 87.0 GB
 NSM Partition Usage: 8.6% of 169.6 GB
 Elastic Storage Used: 0.3 GB
 InfluxDB Storage Used: 0.0 GB

Container Status

Container	Hunt	Status
so-dockerregistry	+	running
so-elastic-fleet	+	running
so-elastic-fleet-package-registry	+	running
so-elasticsearch	+	running
so-idstools	+	running
so-influxdb	+	running
so-kibana	+	running
so-kratos	+	running
so-nginx	+	running
so-sensoroni	+	running
so-soc	+	running
so-telegraf	+	running

Appliance Images

Appliance images are only displayed for official Security Onion Solutions appliances.

While on the *Grid* page, you can upload a PCAP or EVTX file:


The screenshot shows the Security Onion dashboard with a modal dialog titled "Upload a PCAP or EVT File". The dialog includes a file upload icon, a text input field, and a note: "Maximum upload size: 26,214,400 Bytes". Below the input field are "UPLOAD" and "CANCEL" buttons. The background shows system metrics for a Security Onion instance, including memory usage (70.7% of 4.0 GB), CPU usage (3.8%), and various service statuses.

Review alerts on the [Alerts](#) page:


The screenshot displays the "Alerts" page in the Security Onion interface. It features a search bar with "Custom" selected, a date range of "2021/06/30 00:00:00 AM - 2021/07/01 00:00:00", and a "REFRESH" button. Below the search bar, there are "Fetch Limit" (set to 50) and "Filter Results" options. The main content is a table of alerts:

Count	rule.name	event.module	event.severity_label
59	ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)	suricata	low
4	ET MALWARE Trickbot Checkin Response	suricata	high
4	ET POLICY HTTP traffic on port 443 (POST)	suricata	medium
3	ET HUNTING GENERIC SUSPICIOUS POST to Dotted Quad with Fake Browser 1	suricata	medium
3	ET MALWARE VNCStartServer BOT Variant CnC Beacon	suricata	high
3	ET MALWARE VNCStartServer USR Variant CnC Beacon	suricata	high
3	ET POLICY PE EXE or DLL Windows file download HTTP	suricata	high
2	ET HUNTING curl User-Agent to Dotted Quad	suricata	medium
2	ET INFO Dotted Quad Host DLL Request	suricata	medium
2	ET MALWARE Win32/Trickbot Data Exfiltration M2	suricata	high
2	ET POLICY curl User-Agent Outbound	suricata	medium
1	ET DNS Query to a *.top domain - Likely Hostile	suricata	medium
1	ET EXPLOIT ETERNALBLUE Probe Vulnerable System Response MS17-010	suricata	high
1	ET EXPLOIT Possible ETERNALBLUE Probe MS17-010 (Generic Flags)	suricata	high
1	ET HUNTING Suspicious POST with Common Windows Process Names - Possible Process List Exfiltration	suricata	high
1	ET HUNTING Suspicious Windows Commands in POST Body (ipconfig)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net config)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net view)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (nirtest)	suricata	medium

Review other logs on the [Dashboards](#) page:



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:

The screenshot shows the CyberChef interface with a recipe containing three operations: 'From Hexdump', 'Strip HTTP headers', and 'Strings'. The 'Strings' operation is configured with 'Single byte' encoding, a minimum length of 9, and 'Alphanumeric' matching. The 'Input' field contains a hex dump of a network packet, and the 'Output' field shows the extracted strings.

If you need to refer back to previous PCAP jobs, you can find them on the [PCAP](#) page:

The screenshot shows the Security Onion PCAP page. The table below lists the PCAP jobs:

ID	Owner	Date Queued	Date Completed	Sensor ID	Status	Actions
1001	doug@example.com	2024-03-12 15:22:10.335 +00:00	2024-03-12 15:22:11.191 +00:00	securityonion	Completed	
1002	doug@example.com	2024-03-12 15:38:42.108 +00:00	2024-03-12 15:38:43.150 +00:00	securityonion	Completed	

Rows per page: 10 1-2 of 2

IMPORTANT installations do not support remote agents, but if you were running another installation type you could download the Elastic Agent installer from [Downloads](#):

Security Onion

Downloads

Elastic Agent Installers

Evaluation installs and Import installs do not support remote elastic agents. The links below are shown for demonstration purposes only.

- [Windows x86_64 Installer](#)
- [Linux x86_64 Installer](#)
- [macOS x86_64 Installer](#)
- [macOS.arm64 Installer](#)

These [Elastic Agent](#) installers are customized for this specific [Elastic Fleet](#) installation. These files are not signed. If you need signed non-customized Elastic Agent installers, you can get them from [elastic.co](#).

Version: 2.4.60 © 2024 Security Onion Solutions, LLC License: ELv2

The *Administration* section allows to you manage user accounts:

Security Onion

Users

Options

Users Enabled: 1 / 1

Email Address	First Name	Last Name	Note	Role	Status
doug@example.com				superuser	⚠

Profile

First Name

Last Name

Note

UPDATE

Roles

- analyst
- auditor
- limited-analyst
- limited-auditor
- superuser

Access Control

New password

CHANGE PASSWORD LOCK USER

DELETE

Rows per page: 10 1-1 of 1

Version: 2.4.60 © 2024 Security Onion Solutions, LLC License: ELv2

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:

The screenshot shows the Security Onion web interface. The top navigation bar includes a menu icon, the Security Onion logo, and a user profile icon. The left sidebar contains a list of navigation items: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Users, Grid Members, Configuration, and License Key (which is highlighted). Below the sidebar is a 'Tools' section with links to Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Licensing' and has two tabs: 'LICENSE KEY' (selected) and 'LICENSE TERMS'. Under the 'LICENSE KEY' tab, there is a 'License Key' field with a pencil icon and a status indicator that says 'Status: Unprovisioned'.

All this in a minimal VM with only 4GB RAM!

```
top - 15:19:20 up 58 min, 1 user, load average: 0.16, 0.60, 1.08
Tasks: 238 total, 1 running, 235 sleeping, 0 stopped, 2 zombie
%Cpu(s): 1.3 us, 0.6 sy, 0.0 ni, 97.3 id, 0.5 wa, 0.2 hi, 0.1 si, 0.0 st
MiB Mem : 3849.6 total, 156.9 free, 2688.6 used, 1265.4 buff/cache
MiB Swap: 8192.0 total, 6088.7 free, 2103.2 used, 1169.0 avail Mem

  PID USER      PR  NI    VIRT    RES    SHR  S  %CPU  %MEM    TIME+  COMMAND
 87499 socore    20   0 1671212 128500 6888  S   3.7   3.3   1:06.65 sensoroni
 63171 elastic+ 20   0 4885232 991068 14996  S   2.0  25.1  20:33.77 java
 89728 kibana    20   0  11.8g 391520 19584  S   1.3   9.9   2:54.97 node
 59676 socore    20   0  1600    644    600  S   0.3   0.0   0:00.39 entrypoint.sh
128367 root       20   0 1893304 92232 21884  S   0.3   2.3   0:17.53 filebeat
 1 root      20   0  174208   9468  5776  S   0.0   0.2   0:04.24 systemd
 2 root      20   0    0      0     0  S   0.0   0.0   0:00.01 kthreadd
 3 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 rcu_gp
 4 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 rcu_par_gp
 5 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 slab_flushng
 6 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 netns
 8 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 kworker/0:0H-events_highpri
 9 root      20   0    0      0     0  I   0.0   0.0   0:01.13 kworker/u8:0-flush-252:0
10 root      0 -20   0      0     0  I   0.0   0.0   0:01.24 kworker/0:1H-kblockd
11 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 mm_percpu_wq
12 root      20   0    0      0     0  S   0.0   0.0   0:00.00 rcu_tasks_rude_
13 root      20   0    0      0     0  S   0.0   0.0   0:00.00 rcu_tasks_trace
14 root      20   0    0      0     0  S   0.0   0.0   0:00.27 ksoftirqd/0
15 root      20   0    0      0     0  I   0.0   0.0   0:05.65 rcu_sched
16 root      rt   0    0      0     0  S   0.0   0.0   0:00.00 migration/0
18 root      20   0    0      0     0  S   0.0   0.0   0:00.00 cpuhp/0
19 root      20   0    0      0     0  S   0.0   0.0   0:00.00 cpuhp/1
20 root      rt   0    0      0     0  S   0.0   0.0   0:00.05 migration/1
21 root      20   0    0      0     0  S   0.0   0.0   0:00.24 ksoftirqd/1
23 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 kworker/1:0H-events_highpri
24 root      20   0    0      0     0  S   0.0   0.0   0:00.00 cpuhp/2
25 root      rt   0    0      0     0  S   0.0   0.0   0:00.07 migration/2
26 root      20   0    0      0     0  S   0.0   0.0   0:00.25 ksoftirqd/2
28 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 kworker/2:0H-events_highpri
29 root      20   0    0      0     0  S   0.0   0.0   0:00.00 cpuhp/3
30 root      rt   0    0      0     0  S   0.0   0.0   0:00.06 migration/3
31 root      20   0    0      0     0  S   0.0   0.0   0:00.26 ksoftirqd/3
33 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 kworker/3:0H-events_highpri
38 root      20   0    0      0     0  S   0.0   0.0   0:00.00 kdevtmpfs
39 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 inet_frag_wq
40 root      20   0    0      0     0  S   0.0   0.0   0:00.02 kauditd
41 root      20   0    0      0     0  S   0.0   0.0   0:00.00 khungtaskd
42 root      20   0    0      0     0  S   0.0   0.0   0:00.00 oom_reaper
43 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 writeback
44 root      20   0    0      0     0  S   0.0   0.0   0:03.49 kcompactd0
45 root      25   5    0      0     0  S   0.0   0.0   0:00.00 ksmd
46 root      39  19   0      0     0  S   0.0   0.0   0:00.17 khugepaged
104 root      0 -20   0      0     0  I   0.0   0.0   0:00.00 kintegrityd
```

GETTING STARTED

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 *Architecture* section 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 Requirements* 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 Image*, *Azure Cloud Image*, and *Google Cloud Image* sections.

What if I have trouble booting the ISO image?

Check out the *Booting Issues* 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 *After Installation* section.

5.1 Best Practices

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

5.1.1 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 Requirements* 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.
- Prefer taps to span ports when possible.
- Make sure that any network firewalls have the proper firewall rules in place to allow ongoing operation and updates (see the *Firewall* section).

5.1.2 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.

5.1.3 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.

5.1.4 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.

5.2 Architecture

If you're going to deploy Security Onion, you should first decide on what type of deployment you want. 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 forward nodes. This section will discuss what those different deployment types look like from an architecture perspective.


5.2.1 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.



5.2.2 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.



5.2.3 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.




5.2.4 Distributed

A standard distributed deployment includes a **manager node**, one or more **forward 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 forward 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.



5.2.5 Node Types

Management

The **manager** node runs *Security Onion Console (SOC)* and *Kibana*. It has its own local instance of *Elasticsearch*, but that's mainly used for storing *Cases* data and central configuration. An analyst connects to the manager node from a client workstation (perhaps *Security Onion Desktop*) to execute queries and retrieve data. Please keep in mind that a dedicated manager node requires separate search nodes.

The manager node runs the following components:

- *Security Onion Console (SOC)*
- *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 (SOC)*
- *Elasticsearch*
- *Logstash*
- *Kibana*
- *ElastAlert*
- *Redis*

Forward Node

A `forward` 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 *Stenographer* remains on the forward node itself.

Forward nodes run the following components:

- *Zeek*
- *Suricata*
- *Stenographer*

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.



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 talk to `Logstash` on the manager and then would try to connect to the receiver node. Once connected they would send their logs to the receiver like nothing was wrong. The search nodes connect to both the manager and receiver nodes and pull events from the `Redis` queue. If the manager goes down, the 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 *Intrusion Detection Honeypot* node mimics common services such as HTTP, FTP, and SSH. Any interaction with these fake services will automatically result in an alert.




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*
- *Logstash*
- *Zeek*
- *Suricata*
- *Stenographer*

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.

5.3 Hardware Requirements

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.

5.3.1 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!

5.3.2 Minimum Specs

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.

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

5.3.3 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 (<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!

5.3.4 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.

5.3.5 NIC

You'll need at least one wired network interface dedicated to management (preferably connected to a dedicated management network). We recommend using static IP addresses where possible.

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.

5.3.6 UPS

Like most IT systems, Security Onion has databases and those databases don't like power outages or other ungraceful shutdowns. To avoid power outages and having to manually repair databases, please consider a UPS.

5.3.7 Elastic Stack

Please refer to the *Architecture* section for detailed deployment scenarios.

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

5.3.8 Standalone Deployments

In a standalone deployment, the manager components and the sensor components all run on a single box, therefore, your hardware requirements will reflect that. 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.

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 metadata, 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.

5.3.9 Manager node with local log storage and search

In an enterprise distributed deployment, a manager node will store logs from itself and forward 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.

5.3.10 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*. The amount of available RAM directly impacts the size of the *Redis* queue.
- Disk: Used for general OS purposes and storing *Kibana* dashboards.

Please refer to the *Architecture* section for detailed deployment scenarios.

5.3.11 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.

5.3.12 Forward Node (Sensor)

A forward 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.

5.3.13 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.

5.3.14 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*.

5.3.15 Intrusion Detection Honeypot (IDH) Node

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

5.3.16 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 *Stenographer* 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*, *Zeek*, and *Stenographer* (full packet capture) 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: $50\text{Mb/s} = 6.25\text{ MB/s} = 375\text{ MB/minute} = 22,500\text{ MB/hour} = 540,000\text{ MB/day}$. 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 Testing*. For a production deployment, you'll need a SPAN/monitor port on an existing switch or a dedicated TAP. We recommend dedicated TAPs 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
- Mikrotik
- 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>.

5.4 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/2.4/main/DOWNLOAD_AND_VERIFY_ISO.md.

Warning: If you download our ISO image and then scan it with antivirus software, it is possible that one or more of the files included in the ISO image may generate 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\`. This is part of *Strelka* and it is being incorrectly flagged as a backdoor when it is really just a Yara ruleset that looks for backdoors. In some cases, the alert may be for a sample EXE that is included in *Strelka* but again a false positive.

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/>.

5.5 VMware

5.5.1 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. If you don't already have VMware, you can download VMware Workstation Player from <https://www.vmware.com/products/player/playerpro-evaluation.html>.

Note: With the sniffing interface in bridged mode, you will 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.

5.5.2 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:

1. From the VMware main window, select File >> New Virtual Machine.
2. Select Typical installation >> Click Next.
3. Installer disc image file >> SO ISO file path >> Click Next.
4. Choose Linux, then choose the closest Linux distribution and click Next.
5. Specify virtual machine name and click Next.
6. Specify disk size (minimum 200GB), store as single file, click Next.
7. Customize hardware and increase Memory and Processors based on the *Hardware Requirements* section.
8. 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) – in this tutorial, this will be the management interface.
9. Add >> Network Adapter (Bridged) - this will be the sniffing (monitor) interface.
10. Click Close.
11. Click Finish.
12. Power on the virtual machine and then follow the installation steps for your desired installation type in the *Installation* section.

5.5.3 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:

1. From the VMware Fusion main window, click File and then click New.
2. Select the Installation Method appears. Click Install from disc or image and click Continue.
3. Create a New Virtual Machine appears. Click Use another disc or disc image..., select our ISO image, click Open, then click Continue.
4. Choose Operating System appears. Click Linux, choose the closest Linux distribution, then click Continue.
5. Choose Firmware Type appears. Click Legacy BIOS and then click Continue.
6. Finish screen appears. Click the Customize Settings button.
7. Save As screen appears. Give the VM a name and click the Save button.
8. Settings window appears. Click Processors & Memory.
9. Processors & Memory screen appears. Increase processors and memory based on the *Hardware Requirements* section. Click the Add Device... button.
10. Add Device screen appears. Click Network Adapter and click the Add... button.
11. Network Adapter 2 screen appears. This will be the sniffing (monitor) interface. Select your desired network adapter configuration. Click the Show All button.

12. Settings screen appears. Click **Hard Disk (SCSI)**.
13. **Hard Disk (SCSI)** screen appears. Increase the disk size to at least 200GB depending on your use case. Click the **Apply** button.
14. Close the **Settings** window.
15. At the window for your new VM, click the **Play** button to power on the virtual machine.
16. Follow the installation steps for your desired installation type in the *Installation* section.

5.5.4 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>.

5.5.5 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
```

5.6 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>.

5.6.1 Creating VM

- Launch VirtualBox and click the **New** button.
- Provide a name for the virtual machine (**Security Onion 2.4** 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 Requirements* section and then click the **Next** button.
- Specify virtual hard disk size as needed per the *Hardware Requirements* 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.

5.6.2 Guest Additions

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

5.7 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>.

5.7.1 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.

5.7.2 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)`.

5.7.3 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.

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](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>.

5.8 Booting Issues

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


- Verify the ISO image using hashes or GPG key.
- Verify that your machine is x86-64 architecture (standard Intel or AMD 64-bit).
- If you're trying to run a 64-bit virtual machine, verify that your 64-bit processor supports virtualization and that virtualization is enabled in the BIOS.
- 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.

5.9 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, then you may want to consider the airgap option as everything will install via the ISO image.



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.
- *Suricata* NIDS rules from Emerging Threats (ET) are copied to `/nsm/rules/suricata`.
- Yara rules for *Strelka* are copied to `/nsm/rules/yara`.
- Sigma rules for *Playbook* 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.

5.9.1 Rule Updates

Our ISO image includes the Emerging Threats (ET) ruleset. When *soup* updates an airgap system via ISO, it automatically installs the latest ET rules as well. If you would like to switch to a different ruleset like Emerging Threats Pro (ETPRO), then you can manually copy the ETPRO rules to `/nsm/rules/suricata/emerging-all.rules` using a command like:

```
cat /path/to/ETPRO_rules/*.rules > /nsm/rules/suricata/emerging-all.rules
```

5.10 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 Image](#), [Azure Cloud Image](#), or [Google Cloud Image](#) sections.

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



1. Review the [Hardware Requirements](#) and [Release Notes](#) sections.
2. Download and verify our ISO image as shown in the [Download](#) section.
3. Boot the ISO in a machine that meets the minimum hardware specs.
4. Follow the prompts to complete the installation and reboot.
5. You may need to eject the ISO image or change the boot order of the machine to boot from the newly installed OS.
6. Login using the username and password you set in the installer.

7. 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
```

8. Proceed to the *Configuration* section.

5.11 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_230525

Warning: Existing 2.4 RC1 or newer Security Onion AMI installations should use the *soup* command to upgrade to newer versions of Security Onion. Attempting to switch to a newer AMI from the AWS Marketplace could cause loss of data and require full grid re-installation. Upgrading from Security Onion 2.3 or beta versions of 2.4 is unsupported.

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.

5.11.1 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.

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.xlarge
- Storage: 256GB EBS (Optimized) gp3

Evaluation

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

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.xlarge
- Storage: 300GB EBS (Optimized) gp3

Search Nodes

- Quantity: 2 or more
- Type: m5ad.xlarge
- 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

5.11.2 Create Monitoring Interface

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

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.

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**

5.11.3 Create Security Onion Instances

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

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
```

5.11.4 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 elastaalert*/_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.

5.11.5 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.

5.11.6 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.

5.11.7 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 Forward 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.

5.11.8 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>.

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**

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**

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**

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/
```

5.12 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 2.4 RC1 or newer Security Onion Azure Image installations should use the *soup* command to upgrade to newer versions of Security Onion. Attempting to switch to a newer image from the Azure Marketplace could cause loss of data and require full grid re-installation. Upgrading from Security Onion 2.3 or beta versions of 2.4 is unsupported.

Note: Azure has put on hold their Virtual TAP preview feature, which means in order to install a Security Onion sensor in the Azure cloud you will need to use a packet broker offering from the Azure Marketplace. 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.

5.12.1 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.

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_D4as_v4

- Storage: 256GB Premium SSD

Evaluation

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

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_D4as_v4
- Storage: 256GB Premium SSD

Search Nodes

- Quantity: 2 or more
- Type: Standard_D4as_v4
- Storage: 256GB Premium SSD

Sensor monitoring the VPN ingress

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

5.12.2 Create Monitoring Interface

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

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.

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`

5.12.3 Create Security Onion Instances

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 2 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
```

5.12.4 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.

5.12.5 Search Node Setup

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

5.12.6 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.

5.12.7 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.

Note: Azure has put on hold their Virtual TAP preview feature, which means in order to install a Security Onion sensor in the Azure cloud you will need to use a packet broker offering from the Azure Marketplace. See more information here: <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-tap-overview>

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/
```

5.13 Google Cloud Image

If you would like to deploy Security Onion in Google Cloud Platform (GCP), choose the Security Onion 2 image listed on the Google Marketplace: https://securityonion.net/google/?ref=_ptnr_soc_docs_230824

Warning: Existing 2.4 RC1 or newer Security Onion Google Image installations should use the *soup* command to upgrade to newer versions of Security Onion. Attempting to switch to a newer image from the Google Marketplace could cause loss of data and require full grid re-installation. Upgrading from Security Onion 2.3 or beta versions of 2.4 is unsupported.

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 deployment 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.

5.13.1 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 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.

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-4
- 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.

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-4
- Storage: 300GB Balanced Persistent Disk

Search Nodes

- Quantity: 2 or more
- Type: n2-standard-4
- Storage: 256GB Balanced Persistent Disk
- Storage: 375GB Local Disk (NVMe)

Sensor monitoring the VPN ingress

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

5.13.2 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.

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`.

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.

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.

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.

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.

5.13.3 Create Security Onion Instances

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 2 official VM image.
- 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. Specify the network tag `so-collector` for this VM.
- (Distributed “Manager” node or Single-Node grid only) If not using a VPN, enable the Allow HTTPS traffic from the Internet checkbox, and specify allowed source IP ranges. Under network tags, type `https-server` and press <ENTER>.
- 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) Under Disks, click Add Local SSD. Choose NVMe and select the desired disk capacity based on anticipated log/event retention.
- If requested, review GCP Marketplace Terms, and if acceptable click the corresponding checkbox.
- Select: Create

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:

```
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
```

5.13.4 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.

5.13.5 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.

5.13.6 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.

5.13.7 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: Google		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 Forward 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.

5.13.8 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 -nmi 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.

5.14 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.

5.14.1 Import

One of the easiest ways to get started with Security Onion is using it to forensically analyze pcap and log files. Just install Security Onion in Import mode and then import pcap files or Windows event logs in EVTX format using the [Grid](#) page.

5.14.2 Evaluation

Evaluation Mode is ideal for classroom or small lab environments. Evaluation is **not** designed for production usage. Choose EVAL, follow the prompts (see screenshots below), and then proceed to the *After Installation* section.

5.14.3 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 *After Installation* section.

5.14.4 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 forward nodes).

Build the manager by running Setup, selecting the DISTRIBUTED install submenu, 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 install submenu, 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 *After Installation* section.

5.15 After Installation

5.15.1 Services

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

The screenshot displays the Security Onion Grid management interface. The top navigation bar includes 'Overview', 'Alerts', 'Dashboards', 'Hunt', 'Cases', 'PCAP', 'Grid', 'Downloads', and 'Administration'. The 'Grid' section is active, showing a table of nodes with columns for ID, Role, Address, Version, Model, EPS, Last Heard From, Age, and Status. A single node is listed with ID 'securityonion', Role 'Import', Address '192.168.199.143', Version '2.4.60', Model 'N/A', EPS '0', Last Heard From 'a few seconds ago', and Age 'an hour'. The status is 'OK'.

Below the table, there are three panels:

- Node Status:** Shows details for the selected node, including ID, Role, Address, Version, Model, Date Created, Last Heard From, Age, OS Uptime, Last Synchronized, Process Status (OK), Connection Status (OK), Elasticsearch Status (OK), RAID Status (Feature Unavailable), and Consumption EPS (0). It also displays resource usage: Memory (72.0% of 4.0 GB), Swap (21.3% of 8.6 GB), CPU (3.6%), I/O Wait (0.8%), Root Partition (19.1% of 87.0 GB), NSM Partition (8.6% of 169.6 GB), Elastic Storage (0.3 GB), and InfluxDB Storage (0.0 GB).
- Container Status:** Lists running containers with columns for Container, Hunt, and Status. Containers include so-dockerregistry, so-elastic-fleet, so-elastic-fleet-package-registry, so-elasticsearch, so-idstools, so-influxdb, so-kibana, so-kratos, so-nginx, so-sensoroni, so-soc, and so-telegraf.
- Appliance Images:** A section for appliance images, noting that they are only displayed for official Security Onion Solutions appliances.

The footer of the interface shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

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
```

5.15.2 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.

The screenshot shows the Security Onion Administration console. The left sidebar contains a navigation menu with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users, Grid Members, Configuration (selected), and License Key. Below this is a 'Tools' section with links to Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and shows a list of hostgroups under the 'firewall' category. The 'analyst' hostgroup is selected and highlighted. A modal window is open, displaying 'List of IP or CIDR blocks to allow access to this hostgroup.' with a 'VIEW DEFAULT' button and a text input field containing '192.168.199.0/24'. Below the input field is a dropdown menu labeled 'Select a node to modify'.

5.15.3 SSH

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


5.15.4 Data Retention

- Review the *Elasticsearch* section to see if you need to change any of the default index retention settings.
- Review the *Stenographer* section to see if you need to change the Disk Free Percentage or Maximum Files settings.

5.15.5 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 *Security Onion Console (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 *Tuning* section.

SECURITY ONION CONSOLE (SOC)



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.



Login to Security Onion

Password Login

Email Address
Specify a valid email address. Contact your administrator for new account creation.

Password

LOGIN

Passwordless Login (Security Key / WebAuthn / PassKey)

Email Address

LOGIN

© 2024 Security Onion Solutions, LLC

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*, *PCAP*, *Kibana*,

CyberChef, *Playbook*, and *ATT&CK Navigator*. While *Alerts*, *Dashboards*, *Hunt*, *Cases*, 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.

6.1 Alerts

Security Onion Console (SOC) 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*.

The screenshot shows the Security Onion Alerts interface. At the top, there's a navigation bar with the Security Onion logo and a user profile icon. Below it, a sidebar on the left contains navigation options: Overview, Alerts (selected), Dashboards, Hunt, Cases, PCAP, Grid, Downloads, and Administration. The main content area is titled 'Alerts' and features a search bar with 'Custom' selected, a date range of '2021/06/30 00:00:00 AM - 2021/07/01 00:00:00', and a 'REFRESH' button. Below the search bar, there's a 'Fetch Limit' set to 50 and a 'Filter Results' button. The main area displays a table of alerts with the following columns: Count, rule.name, event.module, and event.severity_label. The table contains 20 rows of alert data.

Count	rule.name	event.module	event.severity_label
59	ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)	suricata	low
4	ET MALWARE Trickbot Checkin Response	suricata	high
4	ET POLICY HTTP traffic on port 443 (POST)	suricata	medium
3	ET HUNTING GENERIC SUSPICIOUS POST to Dotted Quad with Fake Browser 1	suricata	medium
3	ET MALWARE VNCStartServer BOT Variant CnC Beacon	suricata	high
3	ET MALWARE VNCStartServer USR Variant CnC Beacon	suricata	high
3	ET POLICY PE EXE or DLL Windows file download HTTP	suricata	high
2	ET HUNTING curl User-Agent to Dotted Quad	suricata	medium
2	ET INFO Dotted Quad Host DLL Request	suricata	medium
2	ET MALWARE Win32/trickbot Data Exfiltration M2	suricata	high
2	ET POLICY curl User-Agent Outbound	suricata	medium
1	ET DNS Query to a *.top domain - Likely Hostile	suricata	medium
1	ET EXPLOIT ETERNALBLUE Probe Vulnerable System Response MS17-010	suricata	high
1	ET EXPLOIT Possible ETERNALBLUE Probe MS17-010 (Generic Flags)	suricata	high
1	ET HUNTING Suspicious POST with Common Windows Process Names - Possible Process List Exfiltration	suricata	high
1	ET HUNTING Suspicious Windows Commands in POST Body (ipconfig)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net config)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net view)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (nltest)	suricata	medium

At the bottom of the interface, there's a footer with 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

6.1.1 Options

At the top of the page, there is an Options menu that allows you to set options such as Acknowledged/Escalated, Automatic Refresh Interval, and Time Zone.

Toggles

The first toggle is labeled `Temporarily enable advanced interface features`. If you enable this option, then the interface will show more advanced features similar to *Dashboards* and *Hunt*. These advanced features are only enabled temporarily so if you navigate away from the page and then return to the page, it will default back to its simplified view.

The Acknowledged and Escalated toggles control what alerts are displayed:

- 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*.

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.

6.1.2 Query Bar

The query bar defaults to `Group By Name, Module` which groups the alerts by `rule.name` and `event.module`. If you want to send your current Alerts query to *Hunt*, you can click the crosshair icon to the right of the query bar.



You can click the dropdown box to select other queries which will group by other fields.



6.1.3 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.

6.1.4 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.
- Clicking the bell icon acknowledges an 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.
- Clicking the blue exclamation icon 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).
- 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.

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.

	Count	rule.name	
 	130	ET JA3 Hash - Possible Malware - Various Trickbot/Kovter/Dridex	
 	Quick Drilldown	ET MALWARE Tordal/Hancitor/Chanitor Checkin	 Include
 	60	ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)	 Exclude
 	41	ET HUNTING Suspicious GET To gate.php with no Referer	 Only
 	41	ET INFO Suspicious Windows NT version 9 User-Agent	 Drilldown
 	41	ET MALWARE Generic gate .php GET with minimal headers	

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.

	Timestamp ▼	rule.name
>	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102
>	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PC RAT/Gh0st CnC traffic (OUTBOUND) 102

- Include
- Exclude
- Only
- Group By
- Clipboard ▼
- Actions ▲
- Hunt
- Correlate
- PCAP
- Google
- VirusTotal

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 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 or pivot to other pages.

	Timestamp ▼	rule.name
	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family
@timestamp	2021-04-29T16:37:38.577Z	
destination.geo.continent_name	Asia	
destination.geo.country_iso_code	HK	
destination.geo.country_name	Hong Kong	
destination.geo.ip	58.64.132.141	
destination.geo.location.lat	22.25	
destination.geo.location.lon	114.1667	
destination.geo.timezone	Asia/Hong_Ko	
destination.ip	58.64.132.141	
destination.port	80	
ecs.version	1.6.0	
event.category	network	
event.dataset	alert	
event.module	suricata	
event.severity	3	
event.severity_label	high	
host.name	securityonion	
ingest.timestamp	2021-04-29T16:37:39.366Z	

- Include
- Exclude
- Only
- Group By
- Clipboard ▼
- Actions ▲
- Hunt
- Correlate
- PCAP
- Google
- VirusTotal

6.1.5 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 value to your existing search to only show search results that include that value.

Exclude

Clicking the Exclude option will exclude the selected value from your existing search results.

Only

Clicking the Only option will start a new search for the selected value and retain any existing groupby terms.

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:

- 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.
- 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 for the selected process.
- Clicking the Process Ancestors option will show all parent processes for the selected process.

Please note that some of these actions will only display on the Actions menu if you click on a specific log type. For example, the Process Info and Process Ancestors options will only appear if you click on a log that contains the process.entity_id field.

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

6.2 Dashboards

Security Onion Console (SOC) 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!>

6.2.1 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 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 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.

6.2.2 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*, endpoint logs, and firewall logs.


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, 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.


6.2.4 Basic Metrics

The first section of output 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.



6.2.5 Group Metrics

The middle section of output is the Group Metrics section. It 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.

Count	    ▼	event.dataset 	event.category 
538		access	host
266		conn	network
215		file	network
116		ssl	network
83		x509	network
63		dns	network
60		alert	network
43		application	host
42		dce_rpc	network
24		weird	network

Rows per page: 10 ▼ 1-10 of 10 < >

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).



6.2.6 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.

Events

Fetch Limit
100

Filter Results

	Timestamp	source.ip	source.port	destination.ip	destination.port	network.transport	network.protocol
>	2024-01-11 16:24:10.997 -05:00	fe80::184f:67ff:fe10:33ac	133	ff02::2	133	icmp	
>	2024-01-11 16:24:10.787 -05:00	fe80::184f:67ff:fe10:33ac	143	ff02::16	0	icmp	
>	2024-01-11 16:24:09.168 -05:00	::	143	ff02::16	0	icmp	
>	2024-01-11 13:42:04.058 -05:00	192.168.3.65	1036	188.72.243.72	80	http	http
>	2024-01-11 13:42:03.202 -05:00	192.168.3.35	1032	188.124.5.107	80	http	http
>	2024-01-11 13:42:02.316 -05:00	207.46.26.185	1863	192.168.10.128	175	http	http
>	2024-01-11 13:42:01.523 -05:00	192.168.10.125	1368	212.58.253.68	80	http	http
>	2024-01-11 13:42:01.193 -05:00	192.168.10.128	1663	80.157.169.193	80	http	http
>	2024-01-11 13:42:01.193 -05:00	192.168.10.128	1664	80.157.169.193	80	http	http
>	2024-01-11 13:42:01.144 -05:00	192.168.10.128	1656	209.85.171.100	80	http	http

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.

The screenshot displays a table of network traffic records. The top row shows a record with the following values: Timestamp: 2024-01-11 16:24:10.997 -05:00, source.ip: fe80::184f:67ff:fe10:33ac, source.port: 133, and destination.ip: ff02::2. Below this, a list of fields and their corresponding values is shown. A context menu is open over the 'client.oui' field, which has the value 'unkno'. The menu options include: Include, Exclude, Only, Group By, New Group By, Clipboard, Actions, Hunt, Add to Case, Correlate, PCAP, CyberChef, Google, and VirusTotal.

Field	Value
@timestamp	2024-01-11T21:24:10.997Z
client.bytes	0
client.ip	fe80::184f:67ff:fe10:33ac
client.ip_bytes	336
client.oui	unkno
client.packets	7
client.port	133
connection.bytes.missed	0
connection.local.originator	true
connection.local.responder	false
connection.state	OTH
connection.state.description	No SY c (a 'partial connection' that was not later closed)
container.id	conn.
data_stream.dataset	zeek
data_stream.namespace	so
data_stream.type	logs
destination.ip	ff02::2
destination.port	134
ecs.version	8.0.0
elastic_agent.id	39fca
elastic_agent.snapshot	false
elastic_agent.version	8.10.4

6.2.7 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.

The backend data fetch took 0.035 seconds. The total round trip took 0.06 seconds.

6.2.8 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 value to your existing search to only show search results that include that value.

Exclude

Clicking the Exclude option will exclude the selected value from your existing search results.

Only

Clicking the Only option will start a new search for the selected value and 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:

- 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.

- 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 for the selected process.
- Clicking the `Process Ancestors` option will show all parent processes for the selected process.

Please note that some of these actions will only display on the Actions menu if you click on a specific log type. For example, the `Process Info` and `Process Ancestors` options will only appear if you click on a log that contains the `process.entity_id` field.

If you'd like to add your own custom actions, see the [SOC Customization](#) section.

6.2.9 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.

6.3 Hunt

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

Hunt Options Total Found: 949

Q * | groupby event.module event.dataset 2021/06/30 00:00:00 AM - 2021/07/01 00:00:00 AM HUNT

Specify a query in Onion Query Language (OQL) Choose the timespan to search, or click the calendar icon to switch to relative time

Basic Metrics

Most Occurrences

Module	Count
zeek	425
suricata	125

Timeline

Fewest Occurrences

Module	Count
suricata	8
zeek	19

Group Metrics

Fetch Limit: 10 Filter Results

Count	event.module	event.dataset
425	zeek	zeek.conn
125	zeek	zeek.ssl
102	suricata	suricata.alert
98	zeek	zeek.dns
67	zeek	zeek.dce_rpc
47	zeek	zeek.kerberos
28	zeek	zeek.smb_mapping
19	zeek	zeek.file
8	zeek	zeek.http

Version: 2.4.60 © 2024 Security Onion Solutions, LLC License: ELv2

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.

6.4 Cases

Security Onion Console (SOC) 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.

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

6.4.1 Installation


Cases is a part of *Security Onion Console (SOC)*. It's automatically enabled when doing an Import, Eval, Standalone, Manager, or ManagerSearch installation. If you want the quickest and easiest way to try out Cases, you can follow our *First Time Users* guide to install a minimal Import installation.

6.4.2 Creating a New Case

On a new deployment, Cases will be empty until you create 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.

	Count	source.ip
	1,892	10.7.9.102
	937	172.16.2.4
	702	10.7.9.102

Once you click the escalate button, you can choose to escalate to a new case or an existing case.


+ Escalate to new case

Attach event to a recently viewed case:

-  Attempts to exploit log4j vulnerability against public facing web servers in DMZ
-  SQL Injection attempts against web servers in DMZ
-  John Doe in Accounting received phishing email

6.4.3 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/>.



6.4.4 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.

The screenshot shows the Security Onion interface. The main content area displays the 'Attachments' tab for a case titled "Attempts to exploit log4j vulnerability against public facing web servers in DMZ". The interface includes a sidebar with navigation options, a main content area with a table of attachments, and a right-hand panel with summary and details.

Attachments Table:

Actions	Created	Updated	Filename
	Jan 19, 2022 9:24 PM	Jan 19, 2022 9:24 PM	vendor_statement.jpg

Attachment Details:

- Filename: vendor_statement.jpg (1,285,029 Bytes)
- SHA256: 1a5270ad07ef885de0bc4e106df937f5e9b27b55dc59b3a0b0ea26642dcf0e9
- SHA1: 234f783a2edc560a66628e9caec904f43932b0a7
- MD5: 2a09d90dc3b393e4650ecf6e0632622a
- Description: Vendor statement
- TLP: white
- Tags: log4j

Summary Panel:

- Assignee: jim@example.com
- Status: in progress
- Severity: high
- Priority: 1
- TLP: amber
- PAP: red
- Category: general
- Tags: log4j

6.4.5 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.

The screenshot shows the Security Onion interface. The main content area displays the 'Observables' tab for a case titled "Attempts to exploit log4j vulnerability against public facing web servers in DMZ". The interface includes a sidebar with navigation options, a main content area with a table of observables, and a right-hand panel with summary and details.

Observables Table:

Actions	Created	Updated	Type	Value
> ⚡	Jan 19, 2022 9:27 PM	Jan 19, 2022 9:27 PM	ip	195.54.160.149
> ⚡	Jan 19, 2022 9:27 PM	Jan 19, 2022 9:27 PM	ip	175.6.210.66
> ⚡	Jan 19, 2022 9:29 PM	Jan 19, 2022 9:29 PM	uri_path	/\${jndi:ldap://121.140.99.236:1389...}
> ⚡	Jan 19, 2022 9:31 PM	Jan 19, 2022 9:31 PM	user-agent	/\${jndi:ldap://121.140.99.236:1389...}

Summary Panel:

- Assignee: jim@example.com
- Status: in progress
- Severity: high
- Priority: 1
- TLP: amber
- PAP: red
- Category: general
- Tags: log4j

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.

6.4.6 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.

Attempts to exploit log4j vulnerability against public facing web servers in DMZ

Several vulnerabilities were recently announced in log4j:
<https://logging.apache.org/log4j/2.x/security.html>

COMMENTS ATTACHMENTS OBSERVABLES **EVENTS** HISTORY

Filter Results

Actions	Timestamp	ID	Category	Module	Dataset
> +	2022-01-01 12:32:03....	-JvFPh4B8vgkSQ33E1Og	network	zeek	http
> +	2022-01-01 12:32:03....	4pvFPh4B8vgkSQ33E...	network	zeek	notice
> +	2022-01-01 12:32:03....	_JvFPh4B8vgkSQ33E1...	network	zeek	http
> +	2022-01-01 12:32:03....	5JvFPh4B8vgkSQ33EVFu	network	zeek	notice
> +	2022-01-01 12:32:09....	_pvFPh4B8vgkSQ33E1...	network	zeek	http
> +	2022-01-01 12:32:09....	55vFPh4B8vgkSQ33E...	network	zeek	notice
> +	2022-01-01 19:55:04....	oJvFPh4B8vgkSQ33FV...	network	zeek	http
> +	2022-01-01 19:55:04....	PjvFPh4B8vgkSQ33EIPd	network	zeek	notice
> +	2022-01-01 19:55:04....	OjvFPh4B8vgkSQ33EIPd	network	zeek	notice

Summary

Assignee: jim@example.com

Status: in progress

Details

Severity: high

Priority: 1

TLP: amber

PAP: red

Category: general

Tags:

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.

Attempts to exploit log4j vulnerability against public facing web servers in DMZ

Several vulnerabilities were recently announced in log4j:
<https://logging.apache.org/log4j/2.x/security.html>

COMMENTS ATTACHMENTS OBSERVABLES **EVENTS** HISTORY

Filter Results

Actions	Timestamp	ID	Category	Module	Dataset
∨ +	2022-01-01 12:32:03....	-JvFPh4B8vgkSQ33E1Og	network	zeek	http

⊕ @timestamp: 2022-01-01T12:32:03.715Z

⊕ client.ip: 175.6.210.66

⊕ client.port: 55736

⊕ destination.geo.city_name: Ashburn

⊕ destination.geo.continent_name: North America

⊕ destination.geo.country_iso_code: US

⊕ destination.geo.country_name: United States

⊕ destination.geo.ip: 198.71.247.91

⊕ destination.geo.location.lat: 39.0469

⊕ destination.geo.location.lon: -77.4903

⊕ destination.geo.region_iso_code: US-VA

Summary

Assignee: jim@example.com

Status: in progress

Details

Severity: high

Priority: 1

TLP: amber

PAP: red

Category: general

Tags:

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.

6.4.7 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.

The screenshot shows the Security Onion interface with the History tab selected. The case title is "Attempts to exploit log4j vulnerability against public facing web servers in DMZ". Below the title, there is a link to a security advisory: <https://logging.apache.org/log4j/2.x/security.html>. The interface includes tabs for COMMENTS, ATTACHMENTS, OBSERVABLES, EVENTS, and HISTORY (which is active). A "Filter Results" section is present above a table of actions.

Actions	User	Time	Kind	Operation
>	doug@example.com	Jan 19, 2022 9:14 PM	Case	+ Create
>	doug@example.com	Jan 19, 2022 9:14 PM	Events	+ Create
>	doug@example.com	Jan 19, 2022 9:14 PM	Case	Update
>	doug@example.com	Jan 19, 2022 9:15 PM	Case	Update
>	tom@example.com	Jan 19, 2022 9:16 PM	Comments	+ Create
>	tom@example.com	Jan 19, 2022 9:16 PM	Case	Update
>	tom@example.com	Jan 19, 2022 9:16 PM	Case	Update
>	tom@example.com	Jan 19, 2022 9:16 PM	Case	Update
>	tom@example.com	Jan 19, 2022 9:16 PM	Case	Update
>	tom@example.com	Jan 19, 2022 9:16 PM	Case	Update
>	tom@example.com	Jan 19, 2022 9:16 PM	Case	Update

On the right side, there is a "Summary" section with the following details:

- Assignee:** jim@example.com
- Status:** in progress

Below the summary is a "Details" section with the following information:

- Severity:** high
- Priority:** 1
- TLP:** amber
- PAP:** red
- Category:** general
- Tags:** log4j

6.4.8 Overview Page

Once you have one or more cases, you can use the main Cases page to get an overview of all cases.

The screenshot shows the Security Onion interface with the Overview page for Cases. The page title is "Cases" and it shows "Total Found: 3". There is a search bar with the text "Open Cases" and a "REFRESH" button. Below the search bar, there are two filters: "NOT so_case.status:closed" and "NOT so_case.category:template". A table lists the cases with columns for Timestamp, so_case.title, so_case.status, so_case.severity, so_case.assigneeid, and so_case.createTime.

Timestamp	so_case.title	so_case.status	so_case.severity	so_case.assigneeid	so_case.createTime
2022-04-18 19:08:32.859 +00:00	Attempts to exploit log4j vulnerability against public facing web servers in DMZ	in progress	high	jim@example.com	2022-01-19T21:14:17.101028031Z
2022-04-18 19:07:33.320 +00:00	John Doe in Accounting received phishing email	new	critical	bill@example.com	2022-01-19T21:12:04.652244498Z
2022-01-19 21:11:18.165 +00:00	SQL injection attempts against web servers in DMZ	in progress	medium	jim@example.com	2022-01-19T21:08:48.897561928Z

At the bottom right, there is a "Rows per page" dropdown set to 50 and a "1-3 of 3" indicator.

6.4.9 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 `Temporarily enable advanced interface features`. If you enable this option, then the interface will show more advanced features similar to *Dashboards* and *Hunt*. These advanced features are only enabled temporarily so if you navigate away from the page and then return to the page, it will default back to its simplified view.

6.4.10 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 crosshair 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.

6.4.11 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.

6.4.12 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.

6.4.13 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.

6.4.14 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	URI	URL	User Agent
Alienvault OTX	✓		✓						✓
Echotrail						✓			
Elasticsearch	✓	✓	✓	✓	✓	✓	✓	✓	✓
EmailRep					✓				
Greynoise				✓					
LocalFile	✓		✓	✓		✓		✓	
Malwarebazaar			✓						
Malware Hash Registry			✓						
Pulsedive	✓		✓	✓			✓	✓	✓
Spamhaus				✓					
Sublime Platform		✓							
Threatfox	✓		✓	✓					
Urlhaus								✓	
Urlscan								✓	
Virustotal	✓		✓	✓				✓	
WhoisLookup	✓								

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:

May 18, 2022 1:12 PM May 18, 2022 1:12 PM hash 8a62d103168974fba9c61edab336038c

Value (hash, filename, etc.):
8a62d103168974fba9c61edab336038c

Description:

IOC:
No

TLP:
white

Tags:

Analyzer Results:

Job: 1054	May 18, 2022 1:13 PM	Analzers Processed: 5	
localfile	✔ No results found	Bytes: 62	▼
malwarehashregistry	⚠ Further investigation needed	Bytes: 161	▼
otx	ℹ Analysis complete	Bytes: 434	▼
pulsedive	✔ No results found	Bytes: 82	▼
virusotal	⚠ Malicious	Bytes: 31,959	▼

user@test.local • May 18, 2022 1:13 PM

user@test.local • May 18, 2022 1:12 PM

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:

Job: 1054	May 18, 2022 1:13 PM	Analzers Processed: 5	
localfile	✔ No results found	Bytes: 62	▼
malwarehashregistry	⚠ Further investigation needed	Bytes: 161	▼
otx	ℹ Analysis complete	Bytes: 434	▼
pulsedive	✔ No results found	Bytes: 82	▼

Expanding the collapsed row will reveal a more detailed view of the analysis:

virusotal ▲ Malicious Bytes: 31,959 ^

```
{
  "response": {
    "data": [
      {
        "attributes": {
          "authentihash": "2a659725ff96571699cfb3d9bfc55c00e596bbfeceada29f4647525ed477d10b6",
          "bytehero_info": "Trojan.Win32.Heur.089",
          "creation_date": 1216826506,
          "first_submission_date": 1611819074,
          "last_analysis_date": 1611819074,
          "last_analysis_results": {
            "ALYac": {
              "category": "malicious",
              "engine_name": "ALYac",
              "engine_update": "20210128",
              "engine_version": "1.1.3.1",
              "method": "blacklist",
              "result": "Trojan.Ranapama.AMY"
            },
            "APEX": {
              "category": "malicious",
              "engine_name": "APEX",
              "engine_update": "20210128",
              "engine_version": "6.125",
              "method": "blacklist",
              "result": "Malicious"
            }
          }
        }
      }
    ]
  }
}
```

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 all configurable settings, including 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/dev/salt/sensoroni/files/analyzers>.

6.5 PCAP

Security Onion Console (SOC) includes a PCAP interface which allows you to access your full packet capture that was written to disk by *Stenographer*.


In most cases, you'll pivot to PCAP from a particular event in *Alerts*, *Dashboards*, or *Hunt* by choosing the PCAP action on the action menu.

Events

Fetch Limit 100 Filter Results

	Timestamp	source.ip	source.port	destination.ip	destination.port	network.transport	network.protocol
>	2024-01-11 16:24:10.997 -05:00	fe80::184f:67ff:fe10:33ac	133	ff02::2	133	icmp	icmp
>	2024-01-11 16:24:10.787 -05:00	fe80::184f:67ff:fe10:33ac	143	ff02::16	0	icmp	icmp
>	2024-01-11 16:24:09.168 -05:00	::	143	ff02::16	0	icmp	icmp
>	2024-01-11 13:42:04.058 -05:00	192.168.3.65	1036	188.72.243.72	80	tcp	http
>	2024-01-11 13:42:03.202 -05:00	192.168.3.35	1032	188.124.5.107	80	tcp	http
>	2024-01-11 13:42:02.316 -05:00	207.46.26.185	1863	192.168.10.128	179	tcp	http
>	2024-01-11 13:42:01.523 -05:00	192.168.10.125	1368	212.58.253.68	80	tcp	http
>	2024-01-11 13:42:01.193 -05:00	192.168.10.128	1663	80.157.169.193	80	tcp	http
>	2024-01-11 13:42:01.193 -05:00	192.168.10.128	1664	80.157.169.193	80	tcp	http
>	2024-01-11 13:42:01.144 -05:00	192.168.10.128	1656	209.85.171.100	80	tcp	http

Alternatively, you can go directly to the PCAP interface, click the blue + button, and then put in your search criteria to search for a particular stream.



Security Onion will then locate the stream and render a high level overview of the packets.

The screenshot shows the Security Onion console interface. On the left is a navigation sidebar with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, and Tools. The main area displays a table of network traffic. At the top, there are search filters for ID (# 1001), host (securityonion), source IP (176.10.125.8:80), and destination IP (172.16.3.130:49457). Below the filters, there are options to 'Filter Results', 'HEX', and a download icon. The table has the following columns: Num, Timestamp, Type, Source IP, Source Port, Destination IP, Destination Port, Flags, and Length. The data rows are as follows:

Num	Timestamp	Type	Source IP	Source Port	Destination IP	Destination Port	Flags	Length
0	2021-06-30 20:48:37.602 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	SYN	66
1	2021-06-30 20:48:37.760 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	SYN ACK	58
2	2021-06-30 20:48:37.760 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	ACK	54
3	2021-06-30 20:48:37.760 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	PSH ACK	148
4	2021-06-30 20:48:37.760 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	ACK	54
5	2021-06-30 20:48:37.927 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	ACK	1514
6	2021-06-30 20:48:37.927 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	PSH ACK	1358
7	2021-06-30 20:48:37.927 +00:00	TCP	172.16.3.130	49457	176.10.125.8	80	ACK	54
8	2021-06-30 20:48:37.930 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	ACK	1514
9	2021-06-30 20:48:37.930 +00:00	TCP	176.10.125.8	80	172.16.3.130	49457	PSH ACK	1370

At the bottom right of the table, there is a 'Rows per page' dropdown set to 10, and a pagination indicator '1-10 of 24'. A 'LOAD MORE' button is located at the bottom left of the table area.


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. By disabling Show all packet data and HEX, we can 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.

You can send all of the visible packet data to *CyberChef* by clicking the CyberChef icon on the right side of the table header. 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.



Finally, you can download the full pcap file by clicking the download button on the far right side of the table header. If you are using *Security Onion Desktop*, then the pcap will automatically open in *NetworkMiner*. Alternatively, you could open the pcap in *Wireshark*.



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

A screenshot of the Security Onion web console interface. The top navigation bar shows "Security Onion" and a user profile icon. A left sidebar contains a menu with items like Overview, Alerts, Dashboards, Hunt, Cases, PCAP (selected), Grid, Downloads, Administration, and Tools. The main content area is titled "PCAP" and features a blue "+" button. Below it is a table with columns: ID, Owner, Date Queued, Date Completed, Sensor ID, Status, and Actions. Two rows of PCAP files are listed, both with a "Completed" status. At the bottom right of the table, there is a "Rows per page" dropdown set to 10 and a "1-2 of 2" pagination indicator. The footer contains "Version: 2.4.60", "© 2024 Security Onion Solutions, LLC", and "License: ELv2".

ID	Owner	Date Queued	Date Completed	Sensor ID	Status	Actions
1001	doug@example.com	2024-03-12 15:22:10.335 +00:00	2024-03-12 15:22:11.191 +00:00	securityonion	Completed	[User] [X]
1002	doug@example.com	2024-03-12 15:38:42.108 +00:00	2024-03-12 15:38:43.150 +00:00	securityonion	Completed	[User] [X]

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/`.

6.5.1 Troubleshooting

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

- Verify that *Stenographer* is enabled.
- 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.
- Check to see if you have any *BPF* configuration that may cause *Stenographer* to ignore the traffic.
- Make sure that there is plenty of free space on */nsm* so that *Stenographer* can carve the stream and write the output to disk.

6.6 Grid

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

The screenshot displays the Security Onion Grid interface. At the top, the 'Grid' title is followed by an 'Options' dropdown and 'Grid EPS: 0'. Below this is a table with columns: ID, Role, Address, Version, Model, EPS, Last Heard From, Age, and Status. A single node is listed with ID 'securityonion', Role 'Import', Address '192.168.199.143', Version '2.4.60', Model 'N/A', EPS '0', Last Heard From 'a few seconds ago', and Age 'an hour'. The Status column contains a green 'OK' button.

Below the table are three panels:

- Node Status:** Shows details for the selected node, including ID, Role, Address, Version, Model, Date Created, Last Heard From, Age, OS Uptime, Last Synchronized, Process Status (OK), Connection Status (OK), Elasticsearch Status (OK), RAID Status (Feature Unavailable), Consumption EPS (0), and various resource usage bars (Memory: 72.0% of 4.0 GB, Swap: 21.3% of 8.6 GB, CPU: 3.6%, I/O Wait: 0.8%, Root Partition: 19.1% of 87.0 GB, NSM Partition: 8.6% of 169.6 GB, Elastic Storage Used: 0.3 GB, InfluxDB Storage Used: 0.0 GB).
- Container Status:** A table listing containers and their status:

Container	Hunt	Status
so-dockerregistry	+	running
so-elastic-fleet	+	running
so-elastic-fleet-package-registry	+	running
so-elasticsearch	+	running
so-idstools	+	running
so-influxdb	+	running
so-kibana	+	running
so-kratos	+	running
so-nginx	+	running
so-sensoroni	+	running
so-soc	+	running
so-telegraf	+	running
- Appliance Images:** A note stating 'Appliance images are only displayed for official Security Onion Solutions appliances.'

At the bottom of the interface, the version '2.4.60' is shown on the left, '© 2024 Security Onion Solutions, LLC' in the center, and 'License: ELv2' on the right.

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.

Starting in Security Onion 2.4.40, there is a new checkbox in the options dropdown near the top of the page. This checkbox 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 SOC 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.

6.6.1 Node Status

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

Note: Starting in Security Onion 2.4.40, a significant number of new metrics are included in the `Node Status` section. Older versions will not have all of the metrics shown below.

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 SOC 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.

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.

Stenographer Loss

The Stenographer Loss field shows the percentage of dropped packets due to *Stenographer* being unable to keep up with the flow of network data. *Stenographer* is responsible for writing down all packets to disk, as well as indexing these packets.

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 SOC *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 this 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 screenshot displays the Security Onion web interface. On the left is a navigation sidebar with options like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, and Administration. The main content area shows the status of a node with ID 'securityonion'. It lists various system metrics such as Memory Usage (70.7% of 4.0 GB), Swap Usage (21.2% of 8.6 GB), CPU Usage (3.8%), and I/O Wait (1.0%). A table of containers is visible on the right, including 'so-dockerregistry', 'so-elastic-fleet', 'so-elastic-fleet-package-registry', 'so-elasticsearch', 'so-idstools', 'so-influxdb', and 'so-telegraf', all with a status of 'running'. An 'Upload a PCAP or EVTX File' dialog box is centered over the interface, featuring a blue header, a file selection icon, a text input field, and 'UPLOAD' and 'CANCEL' buttons. A note at the top right of the interface states: 'Appliance images are only displayed for official Security Onion Solutions appliances.'

- 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 SOC web interface to become temporarily unavailable.

- Clicking the question mark button takes you to this help document.

6.6.2 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.

6.6.3 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.

6.6.4 Other Grid Pages

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

6.7 Downloads

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

Security Onion

Overview
Alerts
Dashboards
Hunt
Cases
PCAP
Grid
Downloads
Administration

Tools
Kibana
Elastic Fleet
Osquery Manager
InfluxDB
CyberChef
Navigator

Downloads

Elastic Agent Installers

Evaluation installs and Import installs do not support remote elastic agents. The links below are shown for demonstration purposes only.

- [Windows x86_64 Installer](#)
- [Linux x86_64 Installer](#)
- [macOS x86_64 Installer](#)
- [macOS.arm64 Installer](#)

These [Elastic Agent](#) installers are customized for this specific [Elastic Fleet](#) installation. These files are not signed. If you need signed non-customized Elastic Agent installers, you can get them from [elastic.co](#).

Version: 2.4.60 © 2024 Security Onion Solutions, LLC License: ELv2

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*.

6.8 Administration

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

6.8.1 Users

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

The screenshot displays the Security Onion Users management page. The interface includes a sidebar with navigation options like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, and Tools. The main content area is titled 'Users' and shows a table of user accounts. Below the table are three panels: Profile, Roles, and Access Control. The Profile panel has fields for 'Locked', 'Analyst', and 'Account is locked'. The Roles panel lists roles like 'analyst', 'auditor', 'limited-analyst', 'limited-auditor', and 'superuser'. The Access Control panel has a 'New password' field and buttons for 'CHANGE PASSWORD', 'UNLOCK USER', and 'DELETE'.

Email Address	First Name	Last Name	Note	Role	Status
doug@example.com	Doug	Burks	TOTP MFA Enabled	superuser	
limited.analyst@example.com	Limited	Analyst	User has changed password	limited-analyst	
limited.auditor@example.com	Limited	Auditor	User hasn't changed password	limited-auditor	
locked.analyst@example.com	Locked	Analyst	Account is locked	analyst	

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 *Role-Based Access Control (RBAC)* section.

The Status column will show different icons depending on the status of the account. In the screenshot above:

- the first account is enabled and has TOTP *MFA* enabled
- the second account is enabled and has changed their password but does not have *MFA* enabled
- the third account is enabled but has not yet changed their password and does not have *MFA* enabled
- the fourth account is locked

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

6.8.2 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.

The screenshot shows the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user profile icon. The left sidebar lists various system components, with 'Grid Members' highlighted under the 'Administration' section. The main content area is titled 'Grid Members' and features a descriptive paragraph: 'A distributed grid is made up of member nodes. Member nodes will request to join the grid and remain in a pending state until an administrator has accepted the node. If a pending member node is not yet listed as pending, then it's possible that the wrong manager host was provided during setup or there could be a connectivity problem.' Below this text are three panels: 'Pending Members' (None), 'Denied Members' (None), and 'Rejected Members' (None). To the right, the 'Accepted Members' panel shows a single member 'securityonion_import' with a green checkmark and a 'REVIEW' button.

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.

6.8.3 Configuration

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

The screenshot shows the Security Onion Configuration interface. On the left is a navigation menu with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration, and License Key. Below this is a 'Tools' section with various management tools. The main content area is titled 'Grid Configuration' and features a tree view on the left with components like backup, bpf, elastalert, elasticfleet, elasticsearch, firewall, global, host, idh, idstools, influxdb, kibana, kratos, logstash, and manager. A filter field is at the top of this list. To the right, there are 'Grid Administration Quick Links' (including NTP, Firewall, and various connection settings) and 'Analyst Quick Links' (including Sigma, Suricata, Zeek, BPFs, and SOC). A blue banner at the top of the quick links section reads: 'Select a setting from the tree view on the left or the quick links on the right.' The top right corner shows 'Modified: 39 / 348' and the bottom footer contains version and license information.

The most common configuration options are shown in the quick links on the right side. On the left side, you can 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, you can click it to go to the documentation for that component.

If you're not sure of which component a particular setting may belong to, you can 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:

- 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: If you see a key that includes `_x_`, it is a placeholder value used to represent a period (.).


Some settings can be applied across the entire grid or to specific nodes. If you apply a setting to a specific node, it will override the grid setting.

By default, the Configuration page only shows the most widely used settings. If you want to see all settings, you can go to the Options bar at the top of the page and then click the toggle labeled `Show all configurable settings, including advanced settings`.

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

6.8.4 License Key

In the future, we will offer some new enterprise features for Security Onion. If you are interested in those features and purchase a license key, then this screen will allow you to enter your license key and then show the status of that license key.



The screenshot displays the Security Onion console interface. The top navigation bar shows the Security Onion logo and a user profile icon. The left sidebar contains a menu with categories: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration, License Key (highlighted), and Tools. The main content area is titled "Licensing" and contains two tabs: "LICENSE KEY" (active) and "LICENSE TERMS". Below the tabs is a large text input field labeled "License Key" with a pencil icon for editing. Below the input field, the status is displayed as "Status: Unprovisioned". The footer of the console shows "Version: 2.4.60", "© 2024 Security Onion Solutions, LLC", and "License: ELv2".

6.9 Kibana

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

6.9.1 Authentication

Log into Kibana using the same username and password that you use for *Security Onion Console (SOC)*.

You can add new user accounts to both Kibana and *Security Onion Console (SOC)* at the same time as shown in the *Adding Accounts* section. Please note that if you instead create accounts directly in Kibana, then those accounts will only have access to Kibana and not *Security Onion Console (SOC)*.

6.9.2 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 *Playbook* and *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>.

6.9.3 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.

6.9.4 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.

6.9.5 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.



6.9.6 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

6.9.7 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/master/xpack-spaces.html#spaces-control-feature-visibility>.

6.9.8 More Information

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

6.10 Elastic Fleet

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

6.10.1 Configuration

Elastic Fleet is pre-configured during Security Onion setup, however, centralized management of configuration is provided within its user interface inside of Kibana.

Configuration options for various components are detailed below.

6.10.2 Agents

This section displays registered Elastic agents (<https://docs.securityonion.net/en/2.4/elastic-agent.html>) and allows the user to add additional 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 `S0-Grid-Nodes` agent policy. We do not recommend removing policy settings for Security Onion grid node agents.

6.10.3 Adding Agents

To add a new agent to your deployment, see the following:

<https://docs.securityonion.net/en/2.4/elastic-agent.html#deployment>

6.10.4 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 `S0-Grid-Nodes` agent policy is comprised of the following integrations:

- `elasticsearch-logs` (Elasticsearch integration)
- `import-evt-x-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)

6.10.5 Agent Policies - endpoints-initial

Agent installers downloaded from SOC -> 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

- 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

6.10.6 Integrations

Note: Security Onion 2.4.10 supports the following Elastic integrations:

- aws
- azure
- cloudflare
- elasticsearch
- endpoint
- fleet_server
- fim
- github
- google_workspace
- log
- osquery_manager
- redis
- system
- tcp
- udp
- windows
- 1password

Security Onion 2.4.20 supports these additional Elastic integrations:

- apache
- auditd
- barracuda
- cisco_asa
- crowdstrike
- darktrace
- f5_bigip
- fortinet
- fortinet_fortigate
- gcp
- http_endpoint
- httpjson
- juniper
- juniper_srx
- kafka_log

- lastpass
- m365_defender
- microsoft_defender_endpoint
- microsoft_dhcp
- netflow
- o365
- okta
- panw
- pfsense
- sentinel_one
- sonicwall_firewall
- symantec_endpoint
- ti_abusech
- ti_misp
- ti_otx
- ti_recordedfuture
- zscaler_zia
- zscaler_zpa

Security Onion 2.4.30 supports these additional Elastic integrations:

- auth0
- carbonblack_edr
- checkpoint
- cisco_duo
- cisco_meraki
- cisco_umbrella
- fireeye
- mimecast
- pulse_connect_secure
- snyk
- sophos
- sophos_central
- tenable_sc
- vsphere

Security Onion 2.4.40 supports these additional Elastic integrations:

- cisco_ftd
- cisco_ios

- cisco_ise
- iis
- microsoft_sqlserver
- mysql
- proofpoint_tap
- snort
- ti_anomali
- ti_cybersixgill
- ti_threatq

Security Onion 2.4.50 supports these additional Elastic integrations:

- citrix_adc
- citrix_waf
- nginx
- winlog

Security Onion 2.4.60 supports these additional Elastic integrations:

- journald
 - ti_cybersixgill
-

You can read more about Elastic integrations at <https://docs.elastic.co/integrations>.

6.10.7 Adding an Integration

New integrations can be added to existing policies to provide increased visibility and more comprehensive monitoring.

To add an integration to an existing policy:

From Fleet -> Agent policies -> \$Policy name, click Add Integration and follow the steps for adding the integration.

When setting up 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), once again, consider adding it to the Fleet Server policy. If that is not feasible, and it will be added to the Grid Nodes policy, make sure to set the firewall rules correctly so that you are not opening ports on all of your nodes.

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.

6.10.8 Adding a Custom Integration

A custom integration can be added by adding an integration such as the Custom Logs integration. We can specify various settings relative to the data source and define additional actions to be performed.

6.10.9 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.

6.10.10 Data Streams

Data collected by Elastic Agent is sent to a data stream (<https://www.elastic.co/guide/en/fleet/current/fleet-overview.html#data-streams-intro>) 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.

6.10.11 Settings

The section provides details such as:

- Fleet server hosts in your deployment
- Configured outputs
 - specifies where data will be sent
 - this should include Elasticsearch 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.

6.10.12 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.

The screenshot shows the Security Onion console interface. The left sidebar contains navigation links for Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Users, Grid Members, Configuration, and License Key. The main content area is titled "Grid Configuration" and features a tree view of configuration options. The "excluded" option is selected, showing a list of Zeek logs: analyzer, broker, capture_loss, cluster, and conn-summary. A "VIEW DEFAULT" button is located above the list. Below the list is a dropdown menu labeled "Select a node to modify".

At the top of the page, click the Options menu and then enable the Show all configurable settings, including 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.

6.10.13 More Information

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

6.11 Osquery Manager

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

6.11.1 More Information

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

6.12 InfluxDB

Security Onion Console (SOC) 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.

6.12.1 Authentication

Log into InfluxDB using the same username and password that you use for *Security Onion Console (SOC)*.

If you need to reset your InfluxDB password, you can reset your *Security Onion Console (SOC)* password via the *Administration* interface which will also update your InfluxDB password.

6.12.2 Configuration

You can configure InfluxDB by going to *Administration* → Configuration → influxdb.

The screenshot displays the Security Onion Administration interface. The top navigation bar shows the Security Onion logo and a user profile icon. The left sidebar contains a menu with categories: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), License Key, Tools, Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled "Grid Configuration" and shows a list of configuration items under the "influxdb" category. The "influxdb" item is expanded, showing a list of sub-items: backup, bpf, elastalert, elasticfleet, elasticsearch, firewall, global, host, idh, idstools, buckets, config, and downsample. The "influxdb" item is highlighted in blue and labeled "enabled". To the right of the configuration list is a "Filter" input field. Below the configuration list is a "VIEW DEFAULT" button. The main content area also displays a configuration form for InfluxDB. The form has a title "You can enable or disable InfluxDB." and a "VIEW DEFAULT" button. The form contains two input fields: "Current Grid Value" with the value "false" and "securityonion_import" with the value "true". The footer of the page shows the version "2.4.60", copyright "© 2024 Security Onion Solutions, LLC", and license "License: ELv2".

You can configure Telegraf by going to *Administration* → Configuration → telegraf.

The screenshot shows the Security Onion Administration console. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), License Key, and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, Navigator). The main content area is titled 'Grid Configuration' and shows a list of grid members. The 'telegraf' member is expanded, showing a configuration table with the following values:

Key	Value
Current Grid Value	false
securityonion_import	true

At the top of the configuration area, there is a 'VIEW DEFAULT' button and a message: 'You can enable or disable Telegraf.' The footer of the console displays 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

6.12.3 More Information

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

6.13 CyberChef

Security Onion Console (SOC) 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.

- 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.
- The recipe area in the middle, where you can drag the operations that you want to use and specify arguments and options.

6.13.1 Screenshot



6.13.2 Accessing

To access CyberChef, log into *Security Onion Console (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.

6.13.3 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.

6.13.4 More Information

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

6.14 Playbook

Security Onion Console (SOC) includes a link on the sidebar that takes you to Playbook which allows you to create a **Detection Playbook**, which itself consists of individual **Plays**. These Plays are fully self-contained and describe the different aspects around a particular detection strategy.

The key components of a Play are:

1. Objective and context - what exactly are we trying to detect and why?
2. What are the follow-up actions required to validate and/or remediate when results are seen?
3. The actual query needed to implement the Play's objective. In our case, the *ElastAlert* / *Elasticsearch* configuration.

Any results from a Play (low, medium, high, critical severity) are available to view within *Dashboards*, *Hunt*, or *Kibana*. High or critical severity results from a Play will generate an Alert within the Security Onion Console *Alerts* interface.

The final piece to Playbook is automation. Once a Play is made active, the following happens:

- The required *ElastAlert* config is put into production
- *ATT&CK Navigator* layer is updated to reflect current coverage

6.14.1 Getting Started

You can access Playbook by logging into *Security Onion Console (SOC)* and clicking the Playbook link. You will see over 500 plays already created that have been imported from the Sigma Community repository of rules at <https://github.com/Neo23x0/sigma/tree/master/rules>.

6.14.2 Creating a new Play

Plays are based on Sigma rules - from <https://github.com/Neo23x0/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.

To create a new play, click on the **Sigma Editor** menu link. Either Load a sample Sigma rule or paste one into the Sigma field and click **Convert**. This will convert the Sigma into a query that you can use in *Dashboards*, *Hunt*, or *Kibana* to confirm that it will work for your target log.

Refer to the Log Sources and Field Names section below for details around what field names to use in the Sigma etc.

Once you are ready to create the Play, click **Create Play From Sigma**. If the Play creation is successful, you will be redirected to the newly created Play - it will have a status of **Draft**.

The lifecycle of a Play is as follows:

1. Draft (Initial state)
2. Active (In Production)
3. Inactive (Temporarily moved out of production)
4. Archived (Play has been superseded/retired)

A Play can also have the status of **Disabled**, which means that it is broken in some way and should not be made Active.

6.14.3 Editing a Play

Click on **Edit** to edit a Play. There will only be a few fields that you can modify - to make edits to the others (**Title**, **Description**, etc), you will need to edit the Sigma inside the Sigma field. Keep in mind that the Sigma is YAML formatted, so if you have major edits to make it is recommended to lint it and/or **Convert** it through the Sigma Editor to confirm that it is formatted correctly. Be sure to remove the prepended and postpended Playbook-specific syntax highlighting before linting/convert - `{{collapse(View Sigma) <pre><code class="yaml"> and </code></pre>}}`.

Once you save your changes, Playbook will update the rest of the fields to match your edits, including regenerating the Elastalert rule if needed.

6.14.4 Putting a Play into Production

When you are ready to start alerting on your Play, change the Status of the play to **Active**. This will create the *ElastAlert* config. Any edits made to the Play in Playbook will automatically update the *ElastAlert* configuration.

The Elastalert rules are located under `/opt/so/rules/elastalert/playbook/<PlayID>.yaml`. Elastalert rules created by Playbook will run every 3 minutes, with a `buffer_time` of 15 minutes.

Warning: Performance testing is still ongoing. We recommend avoiding the `Malicious Nishang PowerShell Commandlets` play as it can cause serious performance problems. You may also want to avoid others with a status of `experimental`.

6.14.5 Viewing Playbook Alerts

When results from your Plays are found (ie alerts), they are available to view within *Alerts*.

6.14.6 Tuning Plays

If you have a Play that is generating false positives, you can tune it by adding a *Custom Filter* to the Play.

For example, suppose you are seeing a large amount of `Suspicious Service Path Modification` alerts. Drilling down into the alerts, it appears to be a legitimate configuration change by one of the IT Ops Service Accounts. This can be tuned out by doing the following:

- Open the Play and click `Edit`
- Add the following filter in the `Custom Filter` field (YAML Formatting!):

```
sofilter:
  User: SA_ITOPS
```

The `sofilter` syntax is important - add as many top-level filter clauses as you need, but they should all start with `sofilter` - for example `sofilter1`, `sofilter2`

- Click `Submit` and Playbook will take care of the rest, which includes automatically adding the custom filter to the rule when it is converted.

Custom filters are applied right away (written out to the backend ElastAlert rule file), but ElastAlert could take a couple minutes to pick up on the change, as it runs rules every 3 minutes.

It is not recommended to edit the Sigma directly for Community rules, as if there is ever an update for that Sigma rule from the Sigma rules repo, your changes will get overwritten.

Finally, if you are seeing legitimate executions that are not unique to your environment, you might consider submitting a PR to the rule in the Sigma repo (<https://github.com/SigmaHQ/sigma/tree/master/rules>).

6.14.7 User Accounts

By default, once a user has authenticated through SOC they can access Playbook without having to login again to the app itself. This anonymous access has the permissions of the analyst role.

If you need administrator access to Playbook, you can login as `admin` with the randomized password found via `sudo salt-call pillar.get secrets`. However, the Playbook UI is designed to be used with a user that has an analyst role. Using an admin account will be very confusing to newcomers to Playbook, since many of the fields will now be shown/editable and it will look much more cluttered.

6.14.8 Disable Anonymous Access and Create User Accounts

If you need your team to login with individual user accounts, you can disable anonymous access and create new user accounts and add them to the analyst group which will give them all the relevant permissions.

To do this, login with a user that has administrative access, and navigate to `Administration -> Users -> New User`. Fill out the relevant fields. By default, Playbook is not connected to an email server so password resets via email will not work. Once the new user has been created, go back to `Administration -> Users` and select the newly created user. There will be a `Groups` tab, from which you can add the user to the Analyst group. This will give the user all the needed permissions.

To disable anonymous access, login with a user that has administrative access and navigate to `Administration -> Projects -> Detection Playbooks`. Unselect the `Public` checkbox.

6.14.9 Misc Notes

`so-playbook-sync` runs every 5 minutes. This script queries Playbook for all active plays and then checks to make sure that there is an *ElastAlert* config for each play. It also runs through the same process for inactive plays.

6.14.10 Log Sources and Field Names

Sigma support currently extends to the following log sources in Security Onion:

- Windows Eventlogs and *Sysmon* (via *Elastic Agent*)
- *osquery* (via *Elastic Agent*)
- *network* (via *Zeek* logs)

The pre-loaded Plays depend on Sysmon and Windows Eventlogs shipped with *Elastic Agent*.

For best compatibility, use the following Sigma Taxonomy:

- Process Creation: <https://github.com/Neo23x0/sigma/wiki/Taxonomy#process-creation-events>
- Network: <https://github.com/Neo23x0/sigma/wiki/Taxonomy#specific>

The current Security Onion Sigma field mappings can be found here: <https://github.com/Security-Onion-Solutions/securityonion-image/blob/master/so-soctopus/so-soctopus/playbook/securityonion-baseline.yml>

6.14.11 Adding Additional Rulesets

The pre-loaded Plays come from the community Sigma repository at <https://github.com/Neo23x0/sigma/tree/master/rules>. The default config is to only pull in the Windows rules. The rest of the rules from the community repository can be added via *Administration* → Configuration → *soctopus* as shown below.

The screenshot shows the Security Onion Administration console. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users, Grid Members, Configuration (selected), License Key, and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, Navigator). The main content area is titled 'Grid Configuration' and shows a list of log sources under the 'soctopus' category. The 'enabled' option is highlighted. A modal dialog is open, allowing the user to enable or disable SOctopus. The 'Current Grid Value' is set to 'false', and there is a dropdown to 'Select a node to modify'. The footer shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

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 all configurable settings, including advanced settings. option.
- On the left side, navigate to soctopus → playbook → rulesets.
- On the right side, add one or more of the following: application, category, cloud, compliance, linux, macos, network, web, windows. These are based on the top level directories from the Sigma community repository rule's folder.
- At the top of the page, click the SYNCHRONIZE GRID button under the Options menu.

6.14.12 Diagnostic Logging

Playbook logs can be found in `/opt/so/log/playbook/`. Depending on what you're looking for, you may also need to look at the *Docker* logs for the container:

```
sudo docker logs so-playbook
```

6.14.13 More Information

Note: Check out our Detecting Hashes video at <https://youtu.be/pK8mS60Sk5s!>

6.15 ATT&CK Navigator

Security Onion Console (SOC) 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.

6.15.1 Accessing

To access Navigator, log into *Security Onion Console (SOC)* and then click the Navigator hyperlink on the left side.



6.15.2 Default Layer - Playbook

The default layer is titled **Playbook** and is automatically updated when a Play from *Playbook* is made active/inactive. This allows you to see your Detection Playbook coverage across the ATT&CK framework.

Right-clicking any Technique and selecting **View Related Plays** will open Playbook with a pre-filtered view of any plays that are tagged with the selected Technique.

6.15.3 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*.

6.15.4 More Information


Note:

For more information about ATT&CK Navigator, please see:

<https://github.com/mitre-attack/attack-navigator>

SECURITY ONION DESKTOP

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. Starting with Security Onion 2.4.20, Desktop nodes will now 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 Desktop desktop environment, find the *workstation* setting in your *Salt* pillar and change `enabled: true` to `enabled: false`:

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

7.1 Chromium

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

7.1.1 More Information

Note:

For more information about Chromium, please see:
<https://www.chromium.org/chromium-projects/>

7.2 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.

7.2.1 Usage

NetworkMiner is a part of our *Security Onion Desktop* installation.

7.2.2 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.

7.2.3 More Information

Note:

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

7.3 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.

7.3.1 Usage

Wireshark 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 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.

7.3.3 More Information

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

NETWORK VISIBILITY

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



8.1 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:

- *Stenographer*
- *Suricata*
- *Zeek*

8.1.1 VLAN tags

Warning:

Please note that *Stenographer* should correctly log traffic on a VLAN but won't log the actual VLAN tags due to the way that *AF-PACKET* works:

<https://github.com/google/stenographer/issues/211>

8.1.2 More Information

Note:

For more information about AF-PACKET, please see:

https://www.kernel.org/doc/Documentation/networking/packet_mmap.txt

8.2 Stenographer

Security Onion uses Stenographer to write network traffic to disk. From <https://github.com/google/stenographer>:

Stenographer is a full-packet-capture utility for buffering packets to disk for intrusion detection and incident response purposes. It provides a high-performance implementation of NIC-to-disk packet writing, handles deleting those files as disk fills up, and provides methods for reading back specific sets of packets quickly and easily.

Stenographer uses *AF-PACKET* for packet acquisition. It's important to note that Stenographer is totally independent from *Suricata* and *Zeek*. This means that Stenographer has no impact on your NIDS alerts and protocol metadata.

8.2.1 Output

Stenographer writes full packet capture to `/nsm/pcap/`. It will automatically start purging old data once the partition reaches the `DiskFreePercentage` setting as shown below.

8.2.2 Analysis

You can access full packet capture via the *PCAP* interface:

Alerts, *Dashboards*, *Hunt*, and *Kibana* allow you to easily pivot to the *PCAP* interface.

8.2.3 Command Line

You can also access packet capture from the command line of the box where the pcap is stored using a Stenographer query as defined at <https://github.com/google/stenographer#querying>. In the following examples, replace “YourStenoQueryHere” with your actual Stenographer query.

The first option is using docker to run `steno-read`. If the query succeeds, you can then find the resulting pcap file in `/nsm/pcaptmp/` in the host filesystem:

```
sudo docker exec -it so-steno steno-read "YourStenoQueryHere" -w /tmp/new.pcap
```

We’ve included a wrapper script called `so-pcap-export` to make this a little easier. For example:

```
sudo so-pcap-export "YourStenoQueryHere" output
```

If the query succeeds, you can then find the resulting output `.pcap` file in `/nsm/pcapout/` in the host filesystem.

8.2.4 Configuration

You can configure Stenographer by going to *Administration* → Configuration → pcap.

The screenshot shows the Security Onion web interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), License Key, Tools, Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and shows a tree view of configuration options. The 'pcap' folder is expanded, and the 'config' sub-option is selected, showing a status of 'enabled'. The right pane displays configuration details for 'Current Grid Value' (false) and 'securityonion_import' (true). A footer bar shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

8.2.5 Disk Free Percentage

An important configuration item to be aware of is steno's DiskFreePercentage setting. From <https://github.com/google/stenographer/blob/master/INSTALL.md#threads>:

DiskFreePercentage: The amount of space to keep free in the packets directory. *stenographer* will delete files in this thread's packets directory when free disk space decreases below this percentage.

You can find this setting at *Administration* → Configuration → pcap → config → diskfreepercentage.

If you have a distributed deployment with dedicated forward nodes, then the default value of 10 should be reasonable since Stenographer should be the main consumer of disk space in the /nsm partition. However, if you have systems that run both Stenographer and *Elasticsearch* at the same time (like eval and standalone installations), then you'll want to make sure that this value is no lower than 21 so that you avoid *Elasticsearch* hitting its watermark setting at 80% disk usage. If you have an older standalone installation, then you may need to manually change this value to 21.

8.2.6 Maximum Files

By default, Stenographer limits the number of files in the pcap directory to 30000 to avoid limitations with the ext3 filesystem. However, if you're using the ext4 or xfs filesystems, then it is safe to increase this value. So if you have a large amount of storage and find that you only have 3 weeks worth of PCAP on disk while still having plenty of free space, then you may want to increase this default setting. To do so, you can go to *Administration* → Configuration → pcap → config → maxdirectoryfiles and set the value to something appropriate for your system.

8.2.7 Diagnostic Logging

Diagnostic logging for Stenographer can be found at `/opt/so/log/stenographer/`. Depending on what you're looking for, you may also need to look at the *Docker* logs for the container:

```
sudo docker logs so-steno
```

8.2.8 Disabling

Since Stenographer is totally independent from *Suricata* and *Zeek*, you can disable it without impacting your NIDS alerts or protocol metadata. If you decide to disable Stenographer, you can do so by going to *Administration* → Configuration → pcap → enabled.

8.2.9 VLAN Tags

Warning:

Please note that Stenographer should correctly record traffic on a VLAN but won't log the actual VLAN tags due to the way that *AF-PACKET* works:

<https://github.com/google/stenographer/issues/211>

8.2.10 More Information

Note: For more information about stenographer, please see <https://github.com/google/stenographer>.

8.3 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*:

The screenshot shows the Security Onion Alerts dashboard. The interface includes a sidebar with navigation options like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, and Administration. The main area displays a table of alerts with the following columns: Count, rule.name, event.module, and event.severity_label. The table lists various alerts such as 'ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)' and 'ET MALWARE Trickbot Checkin Response'. At the top right, it indicates 'Total Found: 102'. There is also a search bar, a date range selector, and a refresh button.

Count	rule.name	event.module	event.severity_label
59	ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)	suricata	low
4	ET MALWARE Trickbot Checkin Response	suricata	high
4	ET POLICY HTTP traffic on port 443 (POST)	suricata	medium
3	ET HUNTING GENERIC SUSPICIOUS POST to Dotted Quad with Fake Browser 1	suricata	medium
3	ET MALWARE VNCStartServer BOT Variant CnC Beacon	suricata	high
3	ET MALWARE VNCStartServer USR Variant CnC Beacon	suricata	high
3	ET POLICY PE EXE or DLL Windows file download HTTP	suricata	high
2	ET HUNTING curl User-Agent to Dotted Quad	suricata	medium
2	ET INFO Dotted Quad Host DLL Request	suricata	medium
2	ET MALWARE Win32/trickbot Data Exfiltration M2	suricata	high
2	ET POLICY curl User-Agent Outbound	suricata	medium
1	ET DNS Query to a *.top domain - Likely Hostile	suricata	medium
1	ET EXPLOIT ETERNALBLUE Probe Vulnerable System Response MS17-010	suricata	high
1	ET EXPLOIT Possible ETERNALBLUE Probe MS17-010 (Generic Flags)	suricata	high
1	ET HUNTING Suspicious POST with Common Windows Process Names - Possible Process List Exfiltration	suricata	high
1	ET HUNTING Suspicious Windows Commands in POST Body (ipconfig)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net config)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (net view)	suricata	medium
1	ET HUNTING Suspicious Windows Commands in POST Body (nltest)	suricata	medium

If enabled, Suricata metadata (protocol logs) can be found in *Dashboards*, *Hunt*, and *Kibana*.

8.3.1 Community ID

Security Onion enables Suricata’s built-in support for *Community ID*.

8.3.2 Configuration

You can configure Suricata by going to *Administration* → *Configuration* → *suricata*.

The screenshot shows the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user profile icon. The left sidebar lists various system components and tools. The main content area is titled 'Grid Configuration' and shows a tree view of configuration options. The 'enabled' option under 'suricata' is selected. A modal window is open, displaying configuration values for 'Current Grid Value' (false) and 'securityonion_import' (true).

If you would like to configure/manage IDS rules, please see the [Managing Rules](#) and [Managing Alerts](#) sections.

8.3.3 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.

8.3.4 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.

8.3.5 Performance

If Suricata is experiencing packet loss, then you may need to do one or more of the following: tune the ruleset (see the [Managing Alerts](#) section), apply a *BPF*, adjust `max-pending-packets` in the Suricata configuration, or 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

8.3.6 Thresholding

To edit the thresholding configuration, please see the *Managing Alerts* section.

8.3.7 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 filter out the unnecessary metadata by writing rules. We have included some examples at <https://github.com/Security-Onion-Solutions/securityonion/blob/dev/salt/idstools/sorules/filters.rules>.

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:

The screenshot shows the Security Onion web interface. The left sidebar contains a navigation menu with 'Configuration' selected. The main panel displays the 'Grid Configuration' page, which includes a search filter and a list of engine options. The 'mdengine' option is highlighted. A modal dialog is open, prompting the user to select an engine for meta data generation, with 'ZEEK' currently selected in the input field.

8.3.8 File Extraction

If you choose Suricata for metadata, it will extract files from network traffic and *Strelka* will then analyze those extracted files. If you would like to extract additional file types, then you can add file types as shown at <https://github.com/Security-Onion-Solutions/securityonion/blob/dev/salt/idstools/sorules/extraction.rules>.

8.3.9 PCAP

Starting in Security Onion 2.4.60, you now have the option of switching full packet capture from *Stenographer* to Suricata.

Warning: This Suricata PCAP feature is in BETA! We recommend that you test this feature thoroughly in a test environment.

If you would like to experiment with Suricata PCAP, then you can go to *Administration* → *Configuration* → *Global* and select the *pcapengine* setting. That setting should default to *STENO* but you can change it to either *TRANSITION* or *SURICATA*. If you don't need your old *Stenographer* PCAP at all, then you can immediately set *pcapengine* to *SURICATA* and manually delete the contents of the *Stenographer* PCAP and index directories. However, most folks will probably want to use the *TRANSITION* option as it will keep *Stenographer* running but not capturing traffic so that you can retrieve older *Stenographer* PCAP as well as new Suricata PCAP. *Stenographer* will then start purging its old PCAP as Suricata uses more space. Once your old *Stenographer* PCAP has fully aged off, you can change the *pcapengine* setting to *SURICATA* to fully disable *Stenographer*.

Differences between Suricata and Stenographer for PCAP

- *Stenographer* indexes PCAP which allows instant retrieval of PCAP sessions from disk. When a Suricata PCAP is requested, a process searches the PCAP files and retrieves the appropriate packets for the flow.
- Since *Stenographer* indexes PCAP, it stores the PCAP in a special format. Suricata writes standard PCAP files which can be copied off to another system and then opened with any standard libpcap tool.
- Suricata can compress PCAP using lz4 compression.
- Suricata supports conditional PCAP if you only want to write PCAP when certain conditions are met.
- Suricata has the ability to stop capturing PCAP once a flow reaches a specific stream depth. Security Onion sets this stream depth to 1MB by default. This means that once the PCAP flow reaches 1MB, Suricata will stop recording packets for that flow.
- Currently, there is NO SUPPORT for a PCAP specific *BPF* for Suricata. If you apply a *BPF* to Suricata, it will apply to not only PCAP but also standard NIDS alerts and metadata if enabled.

Conditional PCAP

If you switch to Suricata PCAP, it will write all traffic to PCAP by default. If you would like to limit Suricata to only writing PCAP when certain conditions are met, you can go to *Administration* -> Configuration -> Suricata -> pcap -> conditional and change it 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 -> pcap. Some settings are considered advanced settings so you will only see them if you enable the `Show all configurable settings, including advanced settings.` option.

- `compression`: Set to `none` to disable compression. Set to `lz4` to enable lz4 compression but note that this requires more CPU cycles.
- `lz4-level`: Specify the level of lz4 compression. `0` for no compression. `16` for maximum compression.
- `maxsize`: Max size in GB to use for PCAP stored on the sensor.
- `filesize`: File size for the PCAP files that get written.
- `use-stream-depth`: Set to `no` to ignore the stream depth and capture the entire flow. Set this to `yes` to truncate the flow based on the stream depth.

8.3.10 Disabling

Suricata can be disabled by going to *Administration* → Configuration → suricata → enabled.

8.3.11 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
```

8.3.12 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.
- 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 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.

8.3.13 Stats

For detailed Suricata statistics, check `/opt/so/log/suricata/stats.log`.

8.3.14 Testing Rules

To test a new rule, use the following utility on a node that runs Suricata (ie Forward 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.

8.3.15 VLAN Tags

If your network traffic has VLAN tags, then Suricata will log them. *Dashboards* has a VLAN dashboard which will show this data.

8.3.16 More Information


Note: For more information about Suricata, please see <https://suricata.io>.

8.4 Zeek

Security Onion includes Zeek for network protocol analysis and metadata. From <https://www.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*:



8.4.1 Community ID

Security Onion enables Zeek's built-in support for *Community ID*.

8.4.2 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.

8.4.3 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.

The screenshot shows the Security Onion web interface. The left sidebar contains navigation links for Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Users, Grid Members, Configuration, and License Key. The main content area is titled "Grid Configuration" and shows a tree view of configuration options. The "zeek" folder is expanded, and the "excluded" option is selected. A modal window displays the "Current Grid Value" for "excluded", listing: analyzer, broker, capture_loss, cluster, and conn-summary. A dropdown menu below the modal says "Select a node to modify".

8.4.4 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`.

8.4.5 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.

8.4.6 Disabling

Zeek can be disabled by going to *Administration* → Configuration → zeek → enabled.

8.4.7 Syslog

To forward Zeek logs to an external syslog collector, please see the *Syslog Output* section.

8.4.8 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 creates logs, such as:

conn.log

- TCP/UDP/ICMP connections
- 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>

ftp.log

- FTP activity
- For more information, see:

<https://docs.zeek.org/en/latest/scripts/base/protocols/ftp/info.zeek.html#type-FTP::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 and several different ICS/SCADA protocols. These analyzers are enabled by default.

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*.

8.4.9 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.

8.4.10 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 forward 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 forward nodes, it may take up to 15 minutes for `intel` to sync to the forward 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>

<https://zeek.org/2014/01/23/intelligence-data-and-bro/>

8.4.11 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
```

8.4.12 More Information

Note: For more information about Zeek, please see <https://www.zeek.org/>.

8.5 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.

8.5.1 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 read more about YARA rules in the *Adding Local Rules* section.

8.5.2 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*.

8.5.3 Configuration

You can configure Strelka by going to *Administration* → *Configuration* → *strelka*.

The screenshot displays the Security Onion Administration web interface. On the left is a navigation sidebar with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), and License Key. Below this is a 'Tools' section with various monitoring tools. The main content area is titled 'Grid Configuration' and shows a tree view of configuration options. Under the 'strelka' category, the 'backend' option is expanded and highlighted in blue, showing the value 'enabled'. To the right of the tree is a configuration form for the selected option, which includes a 'VIEW DEFAULT' button, a 'Current Grid Value' input field containing 'false', and a 'Select a node to modify' dropdown menu. The top of the interface shows the Security Onion logo and a user profile icon. The bottom of the interface displays the version '2.4.60', copyright '© 2024 Security Onion Solutions, LLC', and license 'License: ELv2'.

8.5.4 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
```

8.5.5 More Information

Note: For more information about Strelka, please see <https://github.com/target/strelka>.

8.6 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=PLljFITO9rB17mESq7Z9OeFKvVh39vJW34&index=5>.

8.6.1 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!

8.6.2 Configuration

- Run Setup and select the DISTRIBUTED install 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*.

8.6.3 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 Plays within *Playbook* look for certain logs emitted by OpenCanary to generate alerts, which can be viewed in the *Alerts* interface.

8.6.4 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).

8.6.5 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.

8.6.6 Custom Configuration

You can configure IDH by going to *Administration* → Configuration → idh.

The screenshot shows the Security Onion Administration interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users, Grid Members, Configuration (selected), License Key, and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, Navigator). The main content area is titled 'Grid Configuration' and shows a list of services under the 'idh' category. The 'enabled' checkbox is checked. The 'Current Grid Value' is 'false'. A dropdown menu is visible for selecting a node to modify. The footer shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

8.6.7 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 all configurable settings, including 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.

8.6.8 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-ConfiguringtheSystemsNetwork.html#ol-netconf-config-tui>.

HOST VISIBILITY

Security Onion can consume many kinds of host logs. You can send 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*.

9.1 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.

9.1.1 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, forward 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 *Security Onion Console (SOC) Downloads* page and download the proper Elastic agent for the operating system of that endpoint. Don't forget to allow the agent to connect through the firewall by going to *Administration* -> Configuration -> firewall -> hostgroups.

The screenshot shows the Security Onion web interface. On the left, a navigation sidebar includes 'Overview', 'Alerts', 'Dashboards', 'Hunt', 'Cases', 'PCAP', 'Grid', 'Downloads', 'Administration', 'Users', 'Grid Members', 'Configuration', and 'License Key'. The 'Configuration' section is expanded to show 'hostgroups', with 'analyst' selected. The main content area is titled 'Grid Configuration' and shows a list of hostgroups: backup, bpf, elastalert, elasticfleet, elasticsearch, firewall, and hostgroups. Under 'hostgroups', several options are listed: analyst (selected), beats_endpoint, beats_endpoint_ssl, desktop, elastic_agent_endpoint, eval, fleet, and heavynode. On the right, a modal window displays 'List of IP or CIDR blocks to allow access to this hostgroup.' with a 'VIEW DEFAULT' button, a 'Current Grid Value' field containing '192.168.199.0/24', and a 'Select a node to modify' dropdown menu.

Once there, select the `elastic_agent_endpoint` option.

Note: If you'd like to see this in action, check out our Youtube video at <https://youtu.be/cGmQMsFuAvw>.

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
```

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/8.9/elastic-endpoint-deploy-reqs.html>.

9.1.2 Logs

Once the agent starts sending logs, you should be able to find them in *Dashboards*, *Hunt*, or *Kibana*.

9.1.3 Management

You can manage your agents using *Elastic Fleet*.

9.1.4 Live Queries

You can query your agents in realtime using *Osquery Manager*.

9.1.5 Integrations

You can read more about integrations in the *Elastic Fleet* section and at <https://docs.elastic.co/integrations>.

9.1.6 More Information

Note: For more information about the Elastic Agent, please see <https://www.elastic.co/guide/en/fleet/current/fleet-overview.html>.

9.2 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.

The screenshot shows the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user profile icon. The left sidebar contains a navigation menu with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users, Grid Members, Configuration (selected), and License Key. Below the sidebar is a 'Tools' section with links to Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and shows a list of hostgroups under the 'firewall' category. The 'analyst' hostgroup is selected, and a modal window is open showing the 'Current Grid Value' as '192.168.199.0/24' and a dropdown menu to 'Select a node to modify'.

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.

9.3 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.

9.3.1 Integration

Josh Brower wrote a great paper on integrating sysmon into Security Onion:

<https://www.sans.org/reading-room/whitepapers/forensics/sysmon-enrich-security-onion-039-s-host-level-capabilities-35837>

Please note that the paper is a few years old and was therefore written for an older version of Security Onion.

9.3.2 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 config 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>

9.3.3 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.

9.3.4 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.

9.3.5 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>

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.

10.1 Ingest

Here's an overview of how logs are ingested in various deployment types.

10.1.1 Import

Core Pipeline: Elastic Agent [IMPORT Node] → Elasticsearch Ingest [IMPORT Node]
Logs: Zeek, Suricata

10.1.2 Eval

Core Pipeline: Elastic Agent [EVAL Node] → Elasticsearch Ingest [EVAL Node]
Logs: Zeek, Suricata, Osquery/Fleet

Osquery Shipper Pipeline: Osquery [Endpoint] → Fleet [EVAL Node] → Elasticsearch Ingest via Core Pipeline
Logs: WEL, Osquery, syslog

10.1.3 Standalone

Core Pipeline: Elastic Agent [SA Node] → Logstash [SA Node] → Redis [SA Node] ↔ Logstash [SA Node] →
Elasticsearch Ingest [SA Node]
Logs: Zeek, Suricata, Osquery/Fleet, syslog

WinLogbeat: Winlogbeat [Windows Endpoint] → Logstash [SA Node] → Redis [SA Node] ↔ Logstash [SA Node]
→ Elasticsearch Ingest [SA Node]
Logs: WEL, Sysmon

10.1.4 Fleet Standalone

Pipeline: Elastic Agent [Fleet Node] -> Logstash [M | MS] -> Elasticsearch Ingest [S | MS]

Logs: Osquery

10.1.5 Manager (separate search nodes)

Core Pipeline: Elastic Agent [Fleet | Forward] -> Logstash [Manager] -> Redis [Manager]

Logs: Zeek, Suricata, Osquery/Fleet, syslog

WinLogbeat: Winlogbeat [Windows Endpoint]-> Logstash [Manager] -> Redis [Manager]

Logs: WEL

10.1.6 Manager Search

Core Pipeline: Elastic Agent [Fleet | Forward] -> Logstash [MS] -> Redis [MS] <-> Logstash [MS] -> Elasticsearch Ingest [MS]

Logs: Zeek, Suricata, Osquery/Fleet, syslog

Pipeline: Elastic Agent [MS] -> Logstash [MS] -> Elasticsearch Ingest [MS]

Logs: Local Osquery/Fleet

WinLogbeat: Winlogbeat [Windows Endpoint]-> Logstash [MS] -> Elasticsearch Ingest [MS]

Logs: WEL

10.1.7 Heavy

Pipeline: Elastic Agent [Heavy Node] -> Logstash [Heavy] -> Redis [Heavy] <-> Logstash [Heavy] -> Elasticsearch Ingest [Heavy]

Logs: Zeek, Suricata, Osquery/Fleet, syslog

10.1.8 Search

Pipeline: Redis [Manager] → Logstash [Search] → Elasticsearch Ingest [Search]
 Logs: Zeek, Suricata, Osquery/Fleet, syslog

10.1.9 Forward

Pipeline: Elastic Agent [Forward] → Logstash [M | MS] → Elasticsearch Ingest [S | MS]
 Logs: Zeek, Suricata, syslog

10.2 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.

10.2.1 Configuration

You can configure Logstash by going to *Administration* → Configuration → logstash.

The screenshot shows the Security Onion Administration interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (Users, Grid Members, Configuration, License Key), and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, Navigator). The main content area is titled 'Grid Configuration' and shows a tree view of configuration options under 'logstash', including 'config', 'enabled', 'settings', 'manager', 'nginx', 'ntp', 'patch', 'pcap', 'playbook', and 'redis'. The 'enabled' checkbox is checked. A modal window is open with the text 'You can enable or disable Logstash.' and a 'VIEW DEFAULT' button. Below this, there is a text input field for 'Current Grid Value' containing 'false', and a dropdown menu labeled 'Select a node to modify'.

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.

10.2.2 Parsing

Logstash does not parse logs in Security Onion, so modifying existing parsers or adding new parsers should be done via *Elasticsearch*.

10.2.3 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>.

10.2.4 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"
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
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 *Security Onion Console (SOC)*, navigate to *Administration* -> Configuration. At the top of the page, click the Options menu and then enable the Show all configurable settings, including 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
```

10.2.5 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.

10.2.6 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.

10.2.7 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
```

10.2.8 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 }
```

10.2.9 More Information

Note: For more information about Logstash, please see <https://www.elastic.co/products/logstash>.

10.3 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.

10.3.1 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).

10.3.2 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.



The screenshot shows the Security Onion Administration web interface. The left sidebar contains a navigation menu with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), License Key, and Tools. The main content area is titled "Grid Configuration" and shows a list of services including idstools, influxdb, kibana, kratos, logstash, manager, nginx, ntp, patch, pcap, playbook, redis (expanded to show config and enabled), and sensor. The "enabled" option for redis is highlighted. On the right, a configuration panel for Redis is visible, showing a "VIEW DEFAULT" button, a "Current Grid Value" field set to "false", and a "Select a node to modify" dropdown. The footer of the interface displays "Version: 2.4.60", "© 2024 Security Onion Solutions, LLC", and "License: ELv2".

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>

10.3.3 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
```

10.3.4 More Information

Note: For more information about Redis, please see <https://redis.io/>.

10.4 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.

10.4.1 Data

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-2022-03-08.0001/_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>.

Schema

Security Onion tries to adhere to the Elastic Common Schema wherever possible. Otherwise, additional fields or slight modifications to native Elastic field mappings may be found within the data.

Management

Elasticsearch indices are managed by the `so-elasticsearch-indices-delete` utility and ILM (<https://www.elastic.co/guide/en/elasticsearch/reference/current/index-lifecycle-management.html>).

`so-elasticsearch-indices-delete` handles size-based index deletion and ILM handles the following:

- size-based index rollover
- time-based index rollover
- time-based content tiers
- time-based index deletion

Default ILM policies are preconfigured and associated with various data streams and index templates in `/opt/so/saltstack/default/salt/elasticsearch/defaults.yaml`.

10.4.2 Querying

You can query Elasticsearch using web interfaces like *Alerts*, *Dashboards*, *Hunt*, and *Kibana*. You can also query Elasticsearch from the command line using a tool like `curl`. You can also use `so-elasticsearch-query`.

10.4.3 Authentication

You can authenticate to Elasticsearch using the same username and password that you use for *Security Onion Console (SOC)*.

You can add new user accounts to both Elasticsearch and *Security Onion Console (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 *Security Onion Console (SOC)*.

10.4.4 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
```

10.4.5 Storage

All of the data Elasticsearch collects is stored under `/nsm/elasticsearch/`.

10.4.6 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:

For more about Elasticsearch ingest parsing, please see:

<https://www.elastic.co/guide/en/elasticsearch/reference/current/ingest.html>

10.4.7 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.

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`.

10.4.8 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>

10.4.9 Configuration

You can configure Elasticsearch by going to *Administration* → Configuration → elasticsearch.

The screenshot shows the Security Onion Administration interface. The main content area is titled "Grid Configuration" and shows a tree view of configuration options. The "elasticsearch" folder is expanded, and the "esheap" option is selected. The "esheap" configuration is set to "600m". The interface also shows a "Filter" field and a "Current Grid Value" field. The footer of the interface displays "Version: 2.4.60", "© 2024 Security Onion Solutions, LLC", and "License: ELv2".

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.

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.

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>

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.

10.4.10 Deleting Indices

Size-based Index Deletion

Size-based deletion of Elasticsearch indices occurs based on the value of cluster-wide `elasticsearch.retention.retention_pct`, which is derived from the total disk space available for `/nsm/elasticsearch` across all nodes in the Elasticsearch cluster. The default value for this setting is 50 percent.

To modify this value, first navigate to *Administration* -> Configuration. At the top of the page, click the Options menu and then enable the Show all configurable settings, including advanced settings option. Then navigate to elasticsearch -> retention -> retention_pct. The change 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.

If your indices are using more than `retention_pct`, then `so-elasticsearch-indices-delete` will delete old indices until disk space is back under `retention_pct`.

Time-based Index Deletion

Time-based deletion occurs through the use of the `$data_stream.policy.phases.delete.min_age` setting within the lifecycle policy tied to each index and is controlled by ILM. It is important to note that size-based deletion takes priority over time-based deletion, as disk may reach `retention_pct` and indices will be deleted before the `min_age` value is reached.

Policies can be edited within the SOC administration interface by navigating to *Administration* -> Configuration -> elasticsearch -> \$index -> policy -> phases -> delete -> min_age. Changes will take effect when a new index is created.

10.4.11 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>

10.4.12 Clearing

If you want to clear all Elasticsearch data including documents and indices, you can run the `so-elastic-clear` command.

10.4.13 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 this, add the following to your Elasticsearch *Salt* config:

```
config:
  ingest:
    geoip:
      downloader:
        enabled: true
```

10.4.14 More Information

Note:

For more information about Elasticsearch, please see:
<https://www.elastic.co/products/elasticsearch>

10.5 ElastAlert

From <https://elastalert2.readthedocs.io/en/latest/elastalert.html#overview>:

ElastAlert 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.

10.5.1 Configuration

You can modify ElastAlert configuration by going to *Administration* -> Configuration -> elastalert.

The screenshot shows the Security Onion web interface. The left sidebar has a navigation menu with items like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Users, Grid Members, Configuration, and License Key. The main content area is titled 'Grid Configuration' and shows a list of configuration items under the 'elastalert' category, with 'enabled' selected. A modal window is open, displaying the 'Current Grid Value' as 'false' and a dropdown menu to 'Select a node to modify'.

10.5.2 ElastAlert Rules

ElastAlert rules are stored in `/opt/so/rules/elastalert/`.

By default, ElastAlert rules are configured with an output type of `debug`, which simply outputs to a log file found in `/opt/so/log/elastalert/`.

ElastAlert logs to *Elasticsearch* indices. You can search those indices in *Dashboards*, *Hunt*, or *Kibana*. *Security Onion Console (SOC)* does not automatically search the `elastalert` indices by default so if you want to use *Dashboards* or *Hunt* to search ElastAlert logs, then you'll need to adjust the appropriate configuration setting. Find it in the Administration -> Configuration screen by filtering for `elastic.index` and selecting Options (at the top) and toggle on "Show all configurable settings". Add `*:elastalert*` to the `index` setting. The new setting value should resemble the following:

```
*:so-*,*:endgame-*,*:logs-*,*:elastalert*
```

Slack

To have ElastAlert send alerts to something like Slack, we can simply change the alert type and details for a rule like so:

```
alert:
- "slack":
  slack_webhook_url: "https://hooks.slack.com/services/YOUR_WEBHOOK_URI"
```

Email - Internal

To have ElastAlert send to email, we could do something like the following:

```
alert:
- "email"
email:
- "youremail@yourcompany.com"
smtp_host: "your_company_smtp_server"
smtp_port: 25
from_addr: "elastalert@yourcompany.com"
```

Email - External

If we need to use an external email provider like Gmail, we can add something like the following:

```
alert:
- "email"
email:
- "youremail@gmail.com"
smtp_host: "smtp.gmail.com"
smtp_port: 465
smtp_ssl: true
from_addr: "youremail@gmail.com"
smtp_auth_file: '/opt/elastalert/rules/smtp_auth_file.txt'
```

Then create a new file called `/opt/so/rules/elastalert/smtp_auth_file.txt` and add the following:

```
user: youremail@gmail.com
password: yourpassword
```

so-elastalert-create

`so-elastalert-create` is a tool created by [Bryant Treacle](#) that 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 originally written by [Bryant Treacle](#) 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`.

Defaults

With Security Onion's example rules, Elastalert is configured by default to only count the number of hits for a particular match, and will not return the actual log entry for which an alert was generated.

This is governed by the use of `use_count_query: true` in each rule file.

If you would like to view the data for the match, you can simply remark this line in the rule file(s). This may impact performance negatively, so testing the change in a single file at a time may be the best approach.

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
```

<https://elastalert2.readthedocs.io/en/latest/ruletypes.html#buffer-time>

<https://github.com/Yelp/elastalert/issues/805>

10.5.3 Diagnostic Logging

Elastalert diagnostic logs are in `/opt/so/log/elastalert/`. Depending on what you're looking for, you may also need to look at the *Docker* logs for the container:

```
sudo docker logs so-elastalert
```

10.5.4 More Information

Note: For more information about ElastAlert, please see <https://elastalert2.readthedocs.io/en/latest/>.

10.6 Curator

From <https://www.elastic.co/guide/en/elasticsearch/client/curator/current/about.html#about>:

Elasticsearch Curator helps you curate, or manage, your Elasticsearch indices and snapshots by:

1. Obtaining the full list of indices (or snapshots) from the cluster, as the actionable list
2. Iterate through a list of user-defined filters to progressively remove indices (or snapshots) from this actionable list as needed.
3. Perform various actions on the items which remain in the actionable list.

Warning: Starting in Security Onion 2.4.40, Curator is no longer included in Security Onion. To learn more about index maintenance, please see the *Elasticsearch* section.

10.6.1 More Information

Note:

For more information about Curator, please see:

<https://www.elastic.co/guide/en/elasticsearch/client/curator/current/about.html#about>

10.7 Data Fields

This page references the various types of data fields utilized by the Elastic Stack in Security Onion.

10.7.1 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>

10.7.2 Fields

Alert Data Fields

Elastalert Fields

Zeek Fields

10.7.3 Template files

Fields are mapped to their proper type using template files found in `/opt/so/conf/elasticsearch/templates/`.

10.8 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/2.4/main/salt/elasticsearch/files/ingest/suricata.alert>

<https://github.com/Security-Onion-Solutions/securityonion/blob/2.4/main/salt/elasticsearch/files/ingest/common.nids>

<https://github.com/Security-Onion-Solutions/securityonion/blob/2.4/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
```

10.9 Elastalert Fields

The following lists 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
```

10.10 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/2.4/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/2.4/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/2.4/main/salt/elasticsearch/files/ingest/zeek.common>

which in turn calls the `common` pipeline (`/opt/so/conf/elasticsearch/ingest/common`):

<https://github.com/Security-Onion-Solutions/securityonion/blob/2.4/main/salt/elasticsearch/files/ingest-dynamic/common>

10.11 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.

10.11.1 More Information

Note:

For more information about Community ID, please see:

<https://github.com/corelight/community-id-spec>

10.12 SOC Logs

Standard *Security Onion Console (SOC)* logs can be found at `/opt/so/log/soc/`.

10.12.1 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/`. To look for successful SOC logins, you can run the following:

```
sudo zgrep "Identity authenticated successfully and was issued an Ory Kratos Session Cookie" /opt/so/log/kratos/*
```

Those logs should be ingested into *Elasticsearch* and available for searching in *Dashboards*, *Hunt*, and *Kibana*. Both *Dashboards* and *Hunt* have pre-defined queries for SOC auth logs.

The screenshot displays the Security Onion Dashboards interface. The search query is: `event.module:kratos AND event.dataset:audit AND msg:authenticated | groupby http_request.headers.x-real-ip | groupby identity_id | groupby http_request.headers.user-agent`. The results are summarized in three charts: 'Most Occurrences' (bar chart), 'Timeline' (line chart), and 'Fewest Occurrences' (bar chart). Below the charts, there are three tables showing group metrics for `http_request.headers.x-real-ip`, `identity_id`, and `http_request.headers.user-agent`.

Count	http_request.headers.x-real-ip
2	172.16.224.248
1	172.16.224.1

Count	identity_id
2	c1ee0de7-9ae0-4d3c-958f-6b21565e207d
1	1da175f8-3d2d-4fea-8dd9-304245006891

Count	http_request.headers.user-agent
2	Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/102.0.5005.115 Safari/537.36
1	Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.0.0 Safari/537.36

identity_id

Once you see the auth logs, you will notice that the login is logged as `identity_id`. You can find your desired `identity_id` as follows, replacing `USERNAME@DOMAIN.COM` with your desired SOC username:

```
echo "select * from identities;" | sudo sqlite3 /nsm/kratos/db/db.sqlite | grep   
↳ USERNAME@DOMAIN.COM | cut -d\| -f1
```


UPDATING

In this section, we'll review how to keep Security Onion up-to-date.

11.1 soup

soup stands for Security Onion UPdater. To install updates, run the `soup` command:

```
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 when running `so-status`. This may be due to the initial 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 salt-call state.highstate
```

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.

11.1.1 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 is 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.

11.1.2 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.

11.1.3 OS Updates

soup checks for missing OS updates and ask if you want to install them.

You can configure automatic OS updates by going to *Administration* → Configuration → patch.

The screenshot shows the Security Onion Administration web interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users, Grid Members, Configuration (selected), License Key, and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, Navigator). The main content area is titled 'Grid Configuration' and shows a tree view of configuration options. The 'os' option is expanded, showing 'enabled' (selected), 'schedule_to_run', and 'schedules'. A right-hand panel titled 'Enable OS updates' shows a 'VIEW DEFAULT' button, a 'Current Grid Value' field set to 'true', and a 'Select a node to modify' dropdown menu. The footer indicates 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

11.1.4 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.

11.1.5 Log

If soup displays any errors, you can check `/root/soup.log` for additional clues.

11.1.6 ssh

If you run soup via ssh and the ssh session terminates, then any processes running in that session would terminate. You should avoid leaving soup unattended especially if the machine you are ssh'ing 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.

11.1.7 Airgap

When you run soup on an *Airgap* install, it will ask for the location of the upgrade media. You can do one of the following:

- burn the latest ISO image to a DVD and insert it in the DVD drive
- flash the ISO image to a USB drive and insert that USB drive
- simply copy the ISO file itself to the airgapped manager

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 image):

```
sudo soup -y -f /home/YourUser/securityonion-2.4.XYZ-YYYYMMDD.iso
```

11.1.8 Agents

If soup updated to a new version of the Elastic stack, then you might need to update your Elastic Agents via *Elastic Fleet*.

11.1.9 log_size_limit

soup will check your *Elasticsearch* `log_size_limit` values to see if they are over the recommended values. If so, it will ask you to update the values in `/opt/so/saltstack/local/pillar/minions/`. When updating these files, please update any and all instances of `log_size_limit` as it may exist as `elasticsearch:log_size_limit` or `manager:log_size_limit`.

11.1.10 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
```

11.1.11 Automation

soup can be automated as follows (assuming you've previously accepted the Elastic license):

```
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
```

11.1.12 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`.

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:2.4.0 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 salt-call state.highstate` and wait for it to complete before running `soup` again.

11.1.13 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/heavy nodes in your deployment are updated. This is especially important if you are using true clustering for *Elasticsearch*.

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 salt-call state.highstate` from a search node and make sure there are no errors. Typically if it works on one node it will work on the rest. Forward nodes are less complex and will update as they check in so you can monitor those from the *Grid* section of *Security Onion Console (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.

11.2 End Of Life

This page lists End Of Life (EOL) dates for older versions of Security Onion and older components.

Security Onion 2.3 reaches EOL on April 6, 2024 (please migrate to Security Onion 2.4):

<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>

ACCOUNTS

In Security Onion, there are two main types of accounts:

- operating system (OS) accounts
- application accounts used when authenticating to *Security Onion Console (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.

12.1 Passwords

12.1.1 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
```

12.1.2 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
```

12.1.3 Password Logins to SOC

Log into *Security Onion Console (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 *Security Onion Console (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.

12.1.4 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 *Security Onion Console (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.

12.2 MFA

You can enable Multi-Factor Authentication (MFA) to further protect your account. This can be enabled in *Security Onion Console (SOC)* by clicking the user icon in the upper right corner, clicking **Settings**, and then going to the **Security** tab.

12.2.1 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.

If you have a user account on multiple Security Onion deployments with TOTP activated, they may be listed identically in your authenticator app. If so, you should be able to edit the listing in your authenticator app so that you can distinguish between them.

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.

12.2.2 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.

12.3 Adding Accounts

12.3.1 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
```

We recommend creating usernames in lower case for consistency.

For more information, please see the `adduser` manual by typing `man adduser`.

12.3.2 SOC

If you need to add a new account to *Security Onion Console (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.

12.4 Listing Accounts

12.4.1 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
```

12.4.2 SOC

You can get a list of users in *Security Onion Console (SOC)* by navigating to the *Administration* interface and then clicking Users:

The screenshot shows the Security Onion Administration interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users (selected), Grid Members, Configuration, License Key, and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef). The main content area is titled 'Users' and shows 'Users Enabled: 3 / 4'. A table lists the following users:

Email Address	First Name	Last Name	Note	Role	Status
doug@example.com	Doug	Burks	TOTP MFA Enabled	superuser	Active
limited.analyst@example.com	Limited	Analyst	User has changed password	limited-analyst	Active
limited.auditor@example.com	Limited	Auditor	User hasn't changed password	limited-auditor	Warning
locked.analyst@example.com	Locked	Analyst	Account is locked	analyst	Locked

Below the table are three panels for editing the selected user (locked.analyst@example.com):

- Profile:** Fields for 'Locked', 'Analyst', and 'Account is locked'. An 'UPDATE' button is at the bottom.
- Roles:** A list of roles with checkboxes: analyst, auditor, limited-analyst, limited-auditor, and superuser.
- Access Control:** A 'New password' field with a 'CHANGE PASSWORD' button, and 'UNLOCK USER' and 'DELETE' buttons.

At the bottom right, there is a pagination control showing 'Rows per page: 10' and '1-4 of 4'.

For more information about the Users page, please see the *Administration* section.

12.5 Disabling Accounts

12.5.1 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`.

12.5.2 SOC

If you need to disable an account in *Security Onion Console (SOC)*, you can go to the *Administration* interface, expand the user account, and click the LOCK USER button.

After disabling a user account, the *Administration* page will show the disabled user account with a disabled icon in the Status column:

The screenshot shows the Security Onion Administration interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users (selected), Grid Members, Configuration, License Key, and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef). The main content area is titled 'Users' and shows a table of users. The table has columns for Email Address, First Name, Last Name, Note, Role, and Status. The users listed are:

Email Address	First Name	Last Name	Note	Role	Status
> doug@example.com	Doug	Burks	TOTP MFA Enabled	superuser	
> limited.analyst@example.com	Limited	Analyst	User has changed password	limited-analyst	
> limited.auditor@example.com	Limited	Auditor	User hasn't changed password	limited-auditor	
∨ locked.analyst@example.com	Locked	Analyst	Account is locked	analyst	

Below the table, there is a detailed view for the selected user 'locked.analyst@example.com'. This view is divided into three sections:

- Profile:** Shows the user's status as 'Locked' and 'Analyst'. There is an 'UPDATE' button.
- Roles:** A list of roles with checkboxes: analyst, auditor, limited-analyst, limited-auditor, and superuser.
- Access Control:** A section for password management with a 'New password' field, a 'CHANGE PASSWORD' button, and 'UNLOCK USER' and 'DELETE' buttons.

At the bottom right, there is a pagination control showing 'Rows per page: 10' and '1-4 of 4'.

For more information about the Users page, please see the *Administration* section.

12.6 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*.

12.6.1 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 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

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.

12.6.2 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.

12.6.3 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.

12.6.4 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.

12.6.5 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",
```

(continues on next page)

(continued from previous page)

```

    "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: The format of the json in this file must match the request body outlined in the Elas-

tic 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 '. |<br>↳map_values(keys)'
```

3. Run a salt highstate from the manager:

```
sudo salt-call state.highstate
```

12.6.6 Defining Security Onion Roles

There are two ways to define a custom Security Onion role:

- 1) Building it from scratch using the built-in permissions and default roles available as outlined later in this document, or
- 2) 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:

```
case-admin: eastcoast-analyst<br>event-admin: eastcoast-analyst<br>node-monitor: eastcoast-analyst<br>user-monitor: eastcoast-analyst<br>job-user: eastcoast-analyst
```

2. Alternatively, the `eastcoast-analyst` role could be created by inheriting the permissions of the `analyst` role:

```
limited-analyst: eastcoast-analyst
```

Security Onion Privileges and Default Roles

The available low-level Security Onion privileges are listed in the table below:

<i>cases/read</i>	Read all case-related information for all cases
<i>cases/write</i>	Create and update cases, and escalate events to cases
<i>events/read</i>	Read from Elasticsearch
<i>events/write</i>	Write to Elasticsearch
<i>events/ack</i>	Acknowledge alerts
<i>jobs/read</i>	View all PCAP jobs
<i>jobs/pivot</i>	Pivot to PCAP job from event
<i>jobs/write</i>	Request arbitrary PCAP jobs
<i>jobs/delete</i>	Delete any PCAP job
<i>jobs/process</i>	Update, read, and attach packets to all pending PCAP jobs †
<i>nodes/read</i>	View all nodes in grid
<i>nodes/write</i>	Update node information †
<i>roles/read</i>	View all users' roles
<i>roles/write</i>	Change any user's role
<i>users/read</i>	View all users
<i>users/write</i>	Change any user's password
<i>users/delete</i>	Delete any user

These discrete privileges are then collected into privilege groups as defined below:

case-admin	<i>cases/write</i>
case-monitor	<i>cases/read</i>
event-admin	<i>events/read, events/write, events/ack</i>
event-monitor	<i>events/read</i>
job-admin	<i>jobs/read, jobs/pivot, jobs/write, jobs/delete</i>
job-monitor	<i>jobs/read</i>
job-user	<i>jobs/pivot</i>
job-processor	<i>jobs/process †</i>
node-admin	<i>nodes/read, nodes/write</i>
node-monitor	<i>nodes/read</i>
user-admin	<i>roles/read, roles/write, users/read, users/write, users/delete</i>
user-monitor	<i>roles/read, users/read</i>

† intended for use by Sensoroni agents only

12.7 Kratos

Security Onion Console (SOC) authentication is handled by Kratos. You can read more about Kratos at <https://github.com/ory/kratos>.

12.7.1 Configuration

You can configure Kratos by going to *Administration* → Configuration → kratos.

The screenshot shows the Security Onion Administration web interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), and License Key. Under Administration, there are sub-items: Users, Grid Members, Configuration, and License Key. Under Configuration, there are sub-items: idstools, influxdb, kibana, kratos (expanded), config, log, selfservice, and session. Under session, there are sub-items: lifespan (highlighted), oidc, logstash, manager, nginx, ntp, and patch. The main content area is titled 'Grid Configuration' and shows a configuration card for 'kratos'. The card has a title 'Defines the length of a login session.' and a 'VIEW DEFAULT' button. Below the button is a text input field labeled 'Current Grid Value' with the value '24h'. The footer of the interface shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

12.7.2 More Information

Note: For more information about Kratos, please see <https://github.com/ory/kratos>.

12.8 OpenID Connect (OIDC)

Starting with Security Onion version 2.4.30, 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: The OIDC feature is an enterprise-level feature of Security Onion. Contact Security Onion Solutions, LLC via our website at <https://securityonionsolutions.com> for more information about purchasing a license to enable this feature.

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.

12.8.1 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 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`. Click *Register*, and on the resulting screen find the application ID for this new app registration. It will also resemble a UUID.

[Home](#) > [App registrations](#) >

Register an application ...

* Name

The user-facing display name for this application (this can be changed later).

 ✓

Supported account types

Who can use this application or access this API?

- Accounts in this organizational directory only (SO Dev only - Single tenant)
- Accounts in any organizational directory (Any Microsoft Entra ID tenant - Multitenant)
- Accounts in any organizational directory (Any Microsoft Entra ID tenant - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)
- Personal Microsoft accounts only

[Help me choose...](#)

Redirect URI (optional)

We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

 ✓ ✓


Register an app you're working on here. Integrate gallery apps and other apps from outside your organization by adding from [Enterprise applications](#).

By proceeding, you agree to the [Microsoft Platform Policies](#) 

[Register](#)

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.

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.

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`.

API APIs & Services

- Enabled APIs & services
- Library
- Credentials
- OAuth consent screen
- Page usage agreements

← Create OAuth client ID

A client ID is used to identify a single app to Google's OAuth servers. If your app runs on multiple platforms, each will need its own client ID. See [Setting up OAuth 2.0](#) for more information. [Learn more](#) about OAuth client types.

Application type *
Web application

Name *
Security Onion

The name of your OAuth 2.0 client. This name is only used to identify the client in the console and will not be shown to end users.

i The domains of the URIs you add below will be automatically added to your [OAuth consent screen](#) as [authorized domains](#).

Authorized JavaScript origins ?

For use with requests from a browser

+ ADD URI

Authorized redirect URIs ?

For use with requests from a web server

URIs 1 *
https://invalid.com/auth/self-service/methods/oidc/callback/SSO

+ ADD URI

Note: It may take 5 minutes to a few hours for settings to take effect

CREATE

CANCEL

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*.

Settings / Developer Settings

Register a new OAuth application

Application name *
Security Onion
Something users will recognize and trust.

Homepage URL *
https://invalid.com
The full URL to your application homepage.

Application description
Login to the Security Onion Console
This is displayed to all users of your application.

Authorization callback URL *
https://invalid.com/auth/self-service/methods/oidc/callback/SSO
Your application's callback URL. Read our [OAuth documentation](#) for more information.

Enable Device Flow
Allow this OAuth App to authorize users via the Device Flow.
Read the [Device Flow documentation](#) for more information.

Register application Cancel

© 2023 GitHub, Inc. [Terms](#) [Privacy](#) [Security](#) [Status](#) [Docs](#) [Contact GitHub](#) [Pricing](#) [API](#) [Training](#) [Blog](#) [About](#)

Once the app is created a new screen will show the newly create OAuth application settings, including the generated client ID and secret.

Settings / Developer settings / Security Onion

General

Optional features

Advanced

Security Onion

jertel owns this application.

Transfer ownership

You can list your application in the [GitHub Marketplace](#) so that other users can discover it.

List this application in the Marketplace

0 users

Revoke all user tokens

Client ID

096e48744d26a9fff230

Client secrets Generate a new client secret

Make sure to copy your new client secret now. You won't be able to see it again.

✓ 7a946e57cd293e796aa7b39570c7311688af517f [copy](#)

Added now by jertel

Never used

You cannot delete the only client secret. Generate a new client secret first.

Delete

Client secret

Application logo

Upload new logo

You can also drag and drop a picture from your computer.

Badge background color

#ffffff

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*.

The screenshot displays the Auth0 application configuration interface. At the top, the Security Onion logo is shown. Below it, a text input field for the logo URL contains `https://invalid.com/images/sos-logo.svg`. A descriptive note states: "The URL of the logo to display for the application, if none is set the default badge for this type of application will be shown. Recommended size is 150x150 pixels." Below this is the "Application Type" dropdown menu, currently set to "Single Page Application". A note explains: "The type of application will determine which settings you can configure from the dashboard." The "Application URIs" section is visible, containing the "Application Login URI" field with `https://invalid.com` and the "Allowed Callback URLs" field with `https://invalid.com/auth/self-service/methods/oidc/callback/SSO`. A note for the callback URLs states: "After the user authenticates we will only call back to any of these URLs. You can specify multiple valid URLs by comma-separating them (typically to handle different environments like QA or testing). Make sure to specify the protocol (https://) otherwise the callback may fail in some cases. With the exception of custom URI schemes for native clients, all callbacks should use protocol https:// . You can use Organization URL parameters in these URLs."

Scroll back to the top of the Auth0 Settings page where the Client ID and Secret are shown.

[← Back to Applications](#)

Security Onion

Single Page Application Client ID HzAmVrsEej79CXAx1QcG419Hh4rITAr0

[Quickstart](#) [Settings](#) [Addons](#) [Connections](#) [Organizations](#)

Basic Information

Name *

Domain

Client ID

Client Secret

The Client Secret is not base64 encoded.

Description

A free text description of the application. Max character count is 140.

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.

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.

12.8.2 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 <SSO>* the second time will take the newly linked user to the SOC interface. This additional login click is only required once.

12.8.3 Roles

When a new OIDC user logs into SOC, that user will not be assigned any roles. 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.

An administrator will need to login to SOC and assign roles to OIDC users via the Administration -> Users screen. This is a one time operation, per user.


12.8.4 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.

12.8.5 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.



Administrators may choose to disable password logins when using SSO, to ensure all logins must go through the external OIDC provider. On the SOC Configuration screen, enter `password.enabled` into the filter to locate that Advanced setting (ensure the *Show all configurable settings* toggle is enabled).

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:



Security Onion

Overview
Alerts
Dashboards
Hunt
Cases
PCAP
Grid
Downloads
Administration

Tools
Kibana
Elastic Fleet
Osquery Manager
InfluxDB
CyberChef
Playbook
Navigator

User Settings

Settings Profile **Security**

You may be prompted to login again when updating your security settings. If submitting a new password, you will need to verify your identity first with the old password. This is a security measure to protect your account.

Open ID Connect (OIDC)

Single Sign-On via an external identity provider has been enabled for SOC. Authentication settings, such as password changes, should be performed in the external identity system unless the Security Onion administrators have enabled local password logins concurrently with SSO.

12.8.6 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.

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
```


CUSTOMIZING FOR YOUR ENVIRONMENT

This section covers how to customize Security Onion for your environment.

14.1 SOC Customization

You can customize *Security Onion Console (SOC)* by going to *Administration* → Configuration → soc.

The screenshot shows the Security Onion Administration interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (selected), Users, Grid Members, Configuration (selected), License Key, and Tools (Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, Navigator). The main content area is titled 'Grid Configuration' and shows a list of configuration items under the 'soc' category. The 'actions' item is selected and highlighted. A modal window is open, displaying the 'Current Grid Value' for the 'actions' configuration. The value is a JSON object defining a list of actions. The modal also includes a 'VIEW DEFAULT' button and a text description: 'A list of actions a user can take from the SOC UI against a hunt, alert, and other records. The action must be defined in JSON object format, and contain a "name" key and "links" key. The links is a list of URLs, where the most suitable URL in the list will be the selected URL when the user clicks the action.'

```
Current Grid Value: [{"description": "actionHuntHelp", "icon": "fa-crosshairs", "links": [{"!/#/hunt?q=\"{value|escape}\" | groupby event.module* | groupby -sankey event.module* event.dataset | groupby event.dataset | groupby source.ip source.port destination.ip destination.port | groupby network.protocol | groupby source_geo.organization_name source_geo.country_name | groupby destination_geo.organization_name destination_geo.country_name | groupby rule.name rule.category event.severity_label | groupby
```

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.

14.1.1 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>.

14.1.2 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.

14.1.3 Links

You can also customize the links on the left side. To do so, go to *Administration* → Configuration → soc → server → client → tools.

14.1.4 Reverse DNS Lookups

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 → client → enableReverseLookup.

14.1.5 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.

14.1.6 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 → sessiontimeout.

14.1.7 Custom Queries

If you'd like to add your own custom queries to *Alerts*, *Dashboards*, or *Hunt*, you can go to *Administration* → Configuration → soc → server → client and then select the specific app you'd like to modify.

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 → 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.

14.1.8 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. For example, suppose you want to add AbuseIPDB with URL `https://www.abuseipdb.com/check/{value}`. Insert the following as the next to last line:

```
,{ "name": "AbuseIPDB", "description": "Search for this value at AbuseIPDB", "icon": "fa-external-link-alt", "target": "_blank", "links": [ "https://www.abuseipdb.com/check/{value}" ] }
```

You can also create background actions that don't necessarily result in the user being taken to a new page or tab. For example, if you want to have a new action submit a case to JIRA, you would define it as a background POST action. When it completes the POST, it will show an auto-fading message in SOC telling you that the action completed. Alternatively, instead of the auto-fading message you can have it pop a new tab (or redirect SOC tab) to JIRA. Because of CORS restrictions, SOC can't expect to have visibility into the result of the background POST so there is no attempt to parse the response of any background action, other than the status code/text from the request's response.

Here is an example of a background action that submits a javascript fetch to a remote resource and then optionally shows the user a second URL:

```
{
  "name": "My Background Action",
  "description": "Something wonderful!",
  "icon": "fa-star",
  "target": "_blank",
  "links": [
    "http://somewhere.invalid/?somefield={:client.ip|base64}"
  ],
  "background": true,
  "method": "POST",
  "options": {
    "mode": "no-cors",
    "headers": {
      "header1": "header1value",
      "header2": "header2value"
    }
  },
  "body": "something={value|base64}",
  "backgroundSuccessLink": "https://securityonion.net?code={responseCode}&text={responseStatus}",
  "backgroundFailureLink": "https://google.com?q={error}"
}
```

Note that the above JSON block cannot be pasted as-is into the SOC configuration screen, for the action field. Each custom action must be formatted onto a single line, as was shown in the earlier example. The immediate example above is formatted on multiple lines to make it easier to explain in the documentation below.

The `options` object is the same options object that will be passed into the Javascript `fetch()` method. You can read more about that at https://developer.mozilla.org/en-US/docs/Web/API/Fetch_API/Using_Fetch.

There may come a time where you are not sure what fields to target for the request body, or you may want to forward

events of different types that contain different field names. This is ideal if you would like to send the event to a case management system, a SOAR platform, or something similar. In this case, the `eventJson` variable can be used to pass the entire event as a JSON string.

To use this variable, construct the body of the request within the action configuration, like so:

```
"body": "{eventJson}"
```

NOTE: You may run into issues using the `eventJson` variable, depending on the size of the event and the amount of data being passed in the request.

14.1.9 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).
- `elasticsearch` - 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.

14.2 nginx

nginx is the main web server for Security Onion.

14.2.1 Configuration

You can modify nginx configuration by going to *Administration* → Configuration → nginx.

The screenshot shows the Security Onion Grid Configuration interface. The left sidebar contains a navigation menu with the following items: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), and License Key. Below the main menu is a 'Tools' section listing Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and shows a list of settings under the 'config' section. The 'throttle_login_burst' setting is selected and highlighted in blue. The current value is 12. A tooltip explains that this value represents the number of login requests that can burst without triggering request throttling. The interface includes a search filter, a list of tools, and a footer with version and license information.

14.2.2 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!


1. At the top of the page, click the Options menu and then enable the Show all configurable settings, including advanced settings. option.
2. On the left side, go to nginx, expand ssl, and then select the Replace Default Cert setting.
3. On the right side, change the setting to true and then click the checkmark to save the value.
4. On the left side, select the SSL/TLS Cert File setting.
5. On the right side, paste your new cert file and then click the checkmark to save it.
6. On the left side, select the SSL/TLS Key File setting.
7. On the right side, paste your new key file and then click the checkmark to save it.

14.2.3 More Information

Note: For more information about nginx, please see <https://nginx.org/>.

14.3 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.

14.3.1 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.

14.3.2 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 forward 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>
```

14.3.3 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-rule-update
```

Warning: Using `sudo su -` will ignore `/etc/environment`, instead use `sudo su` if you need to operate as root.

14.4 Firewall

This section will cover both network firewalls outside of Security Onion and the host-based firewall built into Security Onion.

14.4.1 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)

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 are not using our ISO image and are instead performing a *Network Installation*, you will also need access to the following:

- update repo for whatever base OS you're installing on (OS packages)

- download.docker.com (*Docker* packages)
- repo.saltstack.com (*Salt* packages)

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 below.

All nodes to Manager:

- TCP/443 - Sensoroni
- TCP/5000 - Docker registry
- TCP/8086 - influxdb
- TCP/4505 - Salt
- TCP/4506 - Salt

Elastic Agent:

- TCP/8220 (All nodes to Manager, Fleet nodes) - Elastic Agent management
- TCP/8443 (All nodes to Manager) - Elastic Agent binary updates
- TCP/5055 (All nodes to Manager, Fleet nodes, Receiver nodes) - Elastic Agent data

Search nodes from/to Manager:

- TCP/9300 - Node-to-node for *Elasticsearch*
- TCP/9696 - *Redis*

Elastic Fleet nodes to Manager:

- TCP/9200 - Node-to-node for *Elasticsearch*
- TCP/5056 - Logstash-to-Logstash for Elastic Agent data ingest

Elastic Fleet nodes to Receiver nodes:


- TCP/5056 - Logstash-to-Logstash for Elastic Agent data ingest

14.4.2 Host Firewall

The remainder of this section will cover the host firewall built into Security Onion.

14.4.3 Configuration

You can configure the firewall by going to *Administration* → Configuration → firewall → hostgroups.



If for some reason you can't access *Security Onion Console (SOC)*, you can use the `so-firewall` command to allow your IP address to connect (replacing `<IP ADDRESS>` with your actual IP address):

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

14.4.4 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.

14.4.5 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.

14.4.6 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.

14.4.7 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 all configurable settings, including 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.

1. At the top of the page, click the Options menu and then enable the Show all configurable settings, including advanced settings. option.
2. On the left side, go to firewall, select portgroups, locate the nginx portgroup, and then select tcp.
3. On the right side, select the manager node, specify your custom port to be added, and then click the checkmark to save the value.
4. 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.

1. At the top of the page, click the Options menu and then enable the Show all configurable settings, including advanced settings. option.
2. On the left side, go to firewall, select hostgroups, and then select customhostgroup0.
3. 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.
4. On the left side, go to firewall, select portgroups, select customportgroup0, and then select the appropriate protocol.
5. 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.
6. On the left side, go to firewall, role, and then select idh, chain, DOCKER-USER, hostgroups, customhostgroup0, portgroups.
7. 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.
8. The next time the IDH node checks in, it should get the appropriate firewall rules.

14.5 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.

14.5.1 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.

14.5.2 Elastalert

Follow the steps in the *ElastAlert* section.

14.6 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.

The screenshot shows the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user profile icon. The left sidebar contains a navigation menu with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Users, Grid Members, Configuration (selected), and License Key. Below this is a 'Tools' section with various management tools. The main content area is titled 'Grid Configuration' and shows a tree view of configuration options. Under 'ntp', the 'config' sub-option is expanded to show 'NTP Servers'. The 'NTP Servers' page has a 'Filter' box, a 'VIEW DEFAULT' button, and a text area showing the current grid value as '0.pool.ntp.org' and '1.pool.ntp.org'. Below this is a dropdown menu labeled 'Select a node to modify'.

14.6.1 IDS Alerts

Anybody can join the NTP Pool Project and provide NTP service. Occasionally, somebody provides NTP service from a residential DHCP address that at some point in time may have also been used for Tor. This results in IDS alerts for Tor nodes where the port is 123 (NTP). This is another good reason to modify the NTP configuration to pull time updates from your preferred NTP provider.

14.7 Console

When you log into the local bash console (tty1), 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 bash 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>.

14.8 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.

14.9 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.

14.10 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.

14.11 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.

The screenshot shows the Security Onion web interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), and License Key. Below these are tools like Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and shows a list of configuration items: backup, bpf, elastalet, elasticfleet, elasticsearch, firewall, global (expanded), airgap, mdengine (selected), pcapengine, soversion, url_base, host, idh, and idstools. A modal dialog is open over the 'mdengine' item, with the text 'Which engine to use for meta data generation. Options are ZEEK and SURICATA.' and a text input field containing 'ZEEK'.

Then select the `url_base` option.

To get the best performance out of Security Onion, you'll want to tune it for your environment. Start by creating Berkeley Packet Filters (BPFs) to ignore any traffic that you don't want your network sensors to process. Then tune your IDS rulesets. There may be entire categories of rules that you want to disable first and then look at the remaining enabled rules to see if there are individual rules that can be disabled. Once your rules and alerts are under control, then check to see if you have packet loss. If so, then tune the number of AF-PACKET workers for sniffing processes. If you are on a large network, you may need to do additional tuning like pinning processes to CPU cores. More information on each of these topics can be found in this section.

15.1 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.

15.1.1 Configuration

You can modify your BPF configuration by going to *Administration* -> Configuration -> bpf. You can apply BPF configuration to *Stenographer*, *Suricata*, or *Zeek*.

The screenshot displays the Security Onion web interface for Grid Configuration. On the left, a navigation menu includes sections for Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Users, Grid Members, Configuration, and License Key. The main area is titled 'Grid Configuration' and features a search filter and a list of BPF filters. The 'pcap' filter is highlighted. To the right, a panel titled 'List of BPF filters to apply to PCAP' contains a 'VIEW DEFAULT' button, a text input field for 'Current Grid Value', and a dropdown menu for 'Select a node to modify'. The footer shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

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)))
```

Warning:

Please note that *Stenographer* should correctly record traffic on a VLAN but won't log the actual VLAN tags due to the way that *AF-PACKET* works:

<https://github.com/google/stenographer/issues/211>

Adding Comments

Starting in Security Onion 2.4.30, comments can be added to the filters via the SOC UI. 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>

15.1.2 More Information

Note:

For more information about BPF, please see:

https://en.wikipedia.org/wiki/Berkeley_Packet_Filter

<https://biot.com/capstats/bpf.html>

15.2 Managing Rules

15.2.1 Updating Rules

Assuming you have Internet access, Security Onion will automatically update your NIDS rules on a daily basis. If you need to manually update your rules, you can run the following on your manager node:

```
sudo so-rule-update
```

If you have a distributed deployment and you update the rules on your manager node, then those rules will automatically replicate from the manager node to your sensors within 15 minutes. If you don't want to wait 15 minutes, you can force the sensors to update immediately by running the following command on your manager node:

```
sudo salt \* state.highstate
```

15.2.2 Configuration

You can modify your rule configuration by going to *Administration* → Configuration → idstools.

The screenshot shows the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user profile icon. The left sidebar contains a menu with categories: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), License Key, and Tools. The main content area is titled 'Grid Configuration' and features a filter input field. Below the filter is a tree view of grid configurations. Under the 'idstools' category, the 'enabled' option is selected and highlighted in blue. The right-hand panel displays the configuration details for the selected 'enabled' grid. It includes a 'VIEW DEFAULT' button and two input fields: 'Current Grid Value' with the value 'false' and 'securityonion_import' with the value 'true'. The interface also shows a 'Modified: 39 / 348' indicator in the top right corner.

15.2.3 Rulesets

Security Onion offers the following choices for rulesets to be used by *Suricata*.

15.2.4 ET Open

- optimized for *Suricata*
- **free**

For more information, see:

<https://rules.emergingthreats.net/open/>

15.2.5 ET Pro (Proofpoint)

- optimized for *Suricata*
- rules retrievable as released
- license fee per sensor (you are responsible for purchasing enough licenses for your entire deployment)

For more information, see:

<https://www.proofpoint.com/us/threat-insight/et-pro-ruleset>

15.2.6 Snort Community

- NOT optimized for *Suricata*
- community-contributed rules
- **free**

For more information, see:

<https://www.snort.org/downloads/#rule-downloads>

<https://www.snort.org/faq/what-are-community-rules>

15.2.7 Snort Registered

- NOT optimized for *Suricata*
- Snort SO (Shared Object) rules do NOT work with *Suricata*
- same rules as Snort Subscriber ruleset, except rules only retrievable after 30 days past release
- **free**

Since Shared Object rules won't work with *Suricata*, you may want to disable them using a regex like 're:soid [0-9]+' as described in the *Managing Alerts* section.

For more information, see:

<https://www.snort.org/downloads/#rule-downloads>

<https://snort.org/documents/registered-vs-subscriber>

15.2.8 Snort Subscriber (Talos)

- NOT optimized for *Suricata*
- Snort SO (Shared Object) rules do NOT work with *Suricata*
- rules retrievable as released
- license fee per sensor (you are responsible for purchasing enough licenses for your entire deployment)

Since Shared Object rules won't work with *Suricata*, you may want to disable them using a regex like 're:soid [0-9]+' as described in the *Managing Alerts* section.

For more information, see:

<https://www.snort.org/downloads/#rule-downloads>

<https://snort.org/documents/registered-vs-subscriber>


15.2.9 Other

- not officially managed/supported by Security Onion
- license fee may or may not apply

15.3 Adding Local Rules

15.3.1 NIDS

You can add local NIDS rules by going to *Administration* → Configuration → idstools.



The screenshot displays the Security Onion Administration web interface. The left sidebar shows the navigation menu with 'Configuration' selected. The main content area is titled 'Grid Configuration' and includes a filter input field. A tree view on the left lists various configuration categories, with 'idstools' expanded to show 'config' as 'enabled'. The right-hand panel, titled 'Options', contains a 'VIEW DEFAULT' button and two configuration fields: 'Current Grid Value' (set to 'false') and 'securityonion_import' (set to 'true'). The footer of the interface shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

At the top of the page, click the Options menu and then enable the Show all configurable settings, including advanced settings. option. Then navigate to idstools → rules → Local Rules. Add your new rule(s) and click the checkmark to save them. 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.

15.3.2 YARA

Default YARA rules are provided from Florian Roth's *signature-base* Github repo at <https://github.com/Neo23x0/signature-base>.

Local YARA Rules

To add local YARA rules, create a directory in `/opt/so/saltstack/local/salt/strelka/rules`, for example `localrules`. Inside of `/opt/so/saltstack/local/salt/strelka/rules/localrules`, add your YARA rules.

After adding your rules, update the configuration by running `so-strelka-restart` on all nodes running Strelka.

Alternatively, run `salt -G 'role:so-sensor' cmd.run "so-strelka-restart"` to restart Strelka on all sensors at once.

Remote YARA Rules

If you have Internet access and want to have `so-yara-update` pull YARA rules from a remote Github repo, copy `/opt/so/saltstack/local/salt/strelka/rules/`, and modify `repos.txt` to include the repo URL (one per line).

Next, run `so-yara-update` to pull down the rules. Finally, run `so-strelka-restart` to allow Strelka to pull in the new rules.

15.4 Managing Alerts

Network Security Monitoring, as a practice, is not a solution you can plug into your network, make sure you see blinking lights and tell people you are “secure.” It requires active intervention from an analyst to qualify the quantity of information presented. One of those regular interventions is to ensure that you are tuning properly and proactively attempting to reach an acceptable level of signal to noise.

15.4.1 Alerting Engines & Severity

There are two alerting engines within Security Onion: *Suricata* and *Playbook* (Sigma). Though each engine uses its own severity level system, Security Onion converts that to a standardized alert severity:

```
event.severity: 4 ==> event.severity_label: critical
```

```
event.severity: 3 ==> event.severity_label: high
```

```
event.severity: 2 ==> event.severity_label: medium
```

```
event.severity: 1 ==> event.severity_label: low
```

All alerts are viewable in *Alerts*, *Dashboards*, *Hunt*, and *Kibana*.

15.4.2 NIDS Testing

The easiest way to test that our NIDS is working as expected might be to simply access <http://testmynids.org/uid/index.html> from a machine that is being monitored by Security Onion. You can do so via the command line using `curl`:

```
curl testmynids.org/uid/index.html
```

Alternatively, you could also test for additional hits with a utility called `tmNIDS`, running the tool in interactive mode:

```
curl -sSL https://raw.githubusercontent.com/0xtf/testmynids.org/master/tmNIDS -  
↵o /tmp/tmNIDS && chmod +x /tmp/tmNIDS && /tmp/tmNIDS
```

If everything is working correctly, you should see a corresponding alert (GPL ATTACK_RESPONSE id check returned root) in *Alerts*, *Dashboards*, *Hunt*, or *Kibana*. If you do not see this alert, try checking to see if the rule is enabled in `/opt/so/rules/nids/suri/all.rules`:

```
grep 2100498 /opt/so/rules/nids/suri/all.rules
```

You can also test using `so-test`.

15.4.3 Identifying rule categories

Rulesets come with a large number of rules enabled (over 20,000 by default). You should only run the rules necessary for your environment, so you may want to disable entire categories of rules that don't apply to you. Run the following command to get a listing of categories and the number of enabled rules in each:

```
grep -v "^#" /opt/so/rules/nids/suri/all.rules | cut -d\" -f2 | awk '{print $1, $2}' |  
↵grep -v "^$" | sort | uniq -c |sort -nr
```

15.4.4 So what's next?

In tuning your sensor, you must first understand whether or not taking corrective actions on this signature will lower your overall security stance. For some alerts, your understanding of your own network and the business being transacted across it will be the deciding factor. For example, if you don't care that users are accessing Facebook, then you can silence the policy-based signatures for Facebook access.

Another consideration is whether or not the traffic is being generated by a misconfigured piece of equipment. If it is, then the most expedient measure may be to resolve the misconfiguration and then reinvestigate tuning.

There are multiple ways to handle overly productive signatures and we'll try to cover as many as we can without producing a full novel on the subject. After making one of the changes described below, your ruleset will need to be updated as shown in the *Managing Rules* section.

You can disable, modify, or threshold alerts by going to *Administration* → Configuration → idstools.

The screenshot displays the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user profile icon. The left sidebar contains a navigation menu with categories like Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (selected), and License Key. Below this is a 'Tools' section with links to Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and shows a tree view of configuration options. The 'idstools' category is expanded, and the 'enabled' option is selected. A modal window is open, showing the current grid value for 'securityonion_import' set to 'true'. The modal also includes a 'VIEW DEFAULT' button and a 'Current Grid Value' field.

15.4.5 Disable the alert

You can disable an alert by going to *Administration* → Configuration → idstools → sids → disabled.

If you want to disable multiple alerts at one time, you can use regular expressions. For example, to disable all alerts that contain heartbleed:

```
re:heartbleed
```

15.4.6 Modify the alert

You can modify an alert by going to *Administration* → Configuration → idstools → sids → modify.

To include an escaped \$ character in the regex pattern you'll need to make sure it's properly escaped. For example, if you want to modify SID 2009582 and change \$EXTERNAL_NET to \$HOME_NET:

```
2009582 "\\$EXTERNAL_NET" "$HOME_NET"
```

The first string is a regex pattern, while the second is just a raw value. You'll need to ensure the first of the two properly escapes any characters that would be interpreted by regex. The second only needs the \$ character escaped to prevent bash from treating that as a variable.

15.4.7 Rewrite the alert

In some cases, you may not want to use the modify option above, but instead create a copy of the rule and disable the original. You can add local rules as shown in the [Adding Local Rules](#) section. After pasting the rule, you may want to bump the SID into the 90,000,000 range and set the revision to 1. Then make any other changes to the rule. Now that we have a signature that will generate alerts a little more selectively, we need to disable the original SID as shown above.

15.4.8 Threshold

Thresholds, rate filters, and suppressions allow you to make finer grained decisions about certain alerts without having to rewrite them. The most common is a suppression which allows you to suppress alerts by specifying the SID, whether you want to track by source/destination/either, and the IP address or subnet. This way, you can still have certain alerts enabled, but the situations in which they alert are limited. It's important to note that with this functionality, care should be given to the thresholds being written to make sure they do not suppress legitimate alerts. You can learn more about Suricata thresholds at <https://docs.suricata.io/en/suricata-6.0.0/configuration/global-thresholds.html>.

You can manage threshold entries for *Suricata* by going to *Administration* → Configuration → suricata → thresholding → SIDS.

Usage:

```
<signature id>:
- threshold:
  gen_id: <generator id>
  type: <threshold | limit | both>
  track: <by_src | by_dst>
  count: <count>
  seconds: <seconds>
- rate_filter:
  gen_id: <generator id>
  track: <by_src | by_dst | by_rule | by_both>
  count: <count>
  seconds: <seconds>
  new_action: <alert | pass>
  timeout: <seconds>
- suppress:
  gen_id: <generator id>
  track: <by_src | by_dst | by_either>
  ip: <ip | subnet>
```

Please note that *Suricata* 6 has a 64-character limitation on the IP field in a threshold. You can read more about this at <https://redmine.openinfosecfoundation.org/issues/4377>.

Suppress

For example, suppose you want to suppress SID 2013030 where the source IP address is in the 10.10.3.0/24 subnet:

```
2013030:
- suppress:
  gen_id: 1
  track: by_src
  ip: 10.10.3.0/24
```

If you want to suppress a SID for multiple IP addresses, you can separate the IP addresses with a comma. For example, suppose you want to suppress SID 2013030 for the 10.10.3.0/24 subnet and also IP address 10.0.0.5:

```
2013030:
- suppress:
  gen_id: 1
  track: by_src
  ip: 10.10.3.0/24,10.0.0.5
```

15.4.9 Flowbits

idstools may seem like it is ignoring your disabled rules request if you try to disable a rule that has flowbits set.

Note: For a quick primer on flowbits, see <https://blog.snort.org/2011/05/resolving-flowbit-dependencies.html>.

For example, consider the following rules that reference the ET.MSSQL flowbit.

First rule:

```
alert tcp $HOME_NET any -> $EXTERNAL_NET !1433 (msg:"ET POLICY Outbound MSSQL Connection
↳ to Non-Standard Port - Likely Malware"; flow:to_server,established; content:"|12 01 00|
↳ "; depth:3; content:"|00 00 00 00 00 00 15 00 06 01 00 1b 00 01 02 00 1c 00|";
↳ distance:1; within:18; content:"|03 00|"; distance:1; within:2; content:"|00 04 ff 08
↳ 00 01 55 00 00 00|"; distance:1; within:10; flowbits:set,ET.MSSQL; classtype:bad-
↳ unknown; sid:2013409; rev:3;)
```

Second rule:

```
alert tcp $HOME_NET any -> $EXTERNAL_NET 1433 (msg:"ET POLICY Outbound MSSQL Connection
↳ to Standard port (1433)"; flow:to_server,established; content:"|12 01 00|"; depth:3;
↳ content:"|00 00 00 00 00 00 15 00 06 01 00 1b 00 01 02 00 1c 00|"; distance:1;
↳ within:18; content:"|03 00|"; distance:1; within:2; content:"|00 04 ff 08 00 01 55 00
↳ 00 00|"; distance:1; within:10; flowbits:set,ET.MSSQL; classtype:bad-unknown;
↳ sid:2013410; rev:4;)
```

Third rule:

```
alert tcp $HOME_NET any -> $EXTERNAL_NET !1433 (msg:"ET TROJAN Bancos.DV MSSQL CnC
↳ Connection Outbound"; flow:to_server,established; flowbits:isset,ET.MSSQL; content:
↳ "|49 00 B4 00 4D 00 20 00 54 00 48 00 45 00 20 00 4D 00 41 00 53 00 54 00 45 00 52 00|
↳ "; classtype:trojan-activity; sid:2013411; rev:1;)
```

If you try to disable the first two rules without disabling the third rule (which has `flowbits:isset,ET.MSSQL`) the third rule could never fire due to one of the first two rules needing to fire first. `idstools` helpfully resolves all of your flowbit dependencies, and in this case, is “re-enabling” that rule for you on the fly. Disabling all three of those rules by adding the following to `disablesid.conf` has the obvious negative effect of disabling all three of the rules:

```
1:2013409
1:2013410
1:2013411
```

When you run `sudo so-rule-update`, watch the “Setting Flowbit State...” section and you can see that if you disable all three (or however many rules share that flowbit) that the “Enabled XX flowbits” line is decremented and all three rules should then be disabled in your `all.rules`.

15.5 High Performance Tuning

15.5.1 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.

15.5.2 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>

15.5.3 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>

15.5.4 Disk/Memory

If you have plenty of RAM, disable swap altogether.

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/

15.5.5 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*.

15.6 Salt

From <https://docs.saltstack.com/en/latest/>:

Salt is a new approach to infrastructure management built on a dynamic communication bus. Salt can be used for data-driven orchestration, remote execution for any infrastructure, configuration management for any app stack, and much more.

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.

15.6.1 Firewall Requirements

Salt minions must be able to connect to the manager node on ports 4505/tcp and 4506/tcp:

<https://docs.saltproject.io/en/getstarted/system/communication.html>

15.6.2 Checking Status

You can use salt's `test.ping` to verify that all your nodes are up:

```
sudo salt \* test.ping
```

15.6.3 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'
```

15.6.4 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
```

15.6.5 Configuration

Many of the options that are configurable in Security Onion are done by going to *Administration* and then Configuration.

15.6.6 Salt Minion Startup Options

Currently, the salt-minion service startup is delayed by 30 seconds. This was implemented to avoid some issues that we have seen regarding Salt states that used the `ip_interfaces` grain to grab the management interface IP.

15.6.7 Diagnostic Logs

Diagnostic logs can be found in `/opt/so/log/salt/`.

15.6.8 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.

15.6.9 More Information

Note: For more information about Salt, please see <https://docs.saltstack.com/en/latest/>.

TRICKS AND TIPS

This section is a collection of miscellaneous tricks and tips for Security Onion.

16.1 Backup

Security Onion performs a daily backup of some critical files so that you can recover your grid from a catastrophic 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:

- `/etc/pki/` - All of the certs including the CA are backed up. Restoring this would allow you to communicate with your salt minions again.
- `/opt/so/saltstack/local/` - This includes all customizations done via *Administration* → Configuration.

You can configure backups by going to *Administration* → Configuration → backup.

The screenshot displays the Security Onion web interface. On the left, a navigation menu includes sections for Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Users, Grid Members, Configuration (highlighted), and License Key. Below this are various tools like Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main panel is titled 'Grid Configuration' and features a search filter and a tree view of configuration categories. The 'destination' category is selected, revealing a list of sub-categories including 'elasticsearch'. The 'elasticsearch' category is expanded, showing a configuration field for the directory to store backups, currently set to '/nsm/backup'.

16.1.1 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/`.

16.2 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.

16.2.1 Download

Our ISO image includes the Docker engine and all of our Docker images.

16.2.2 Security

To prevent tampering, our Docker images are signed using GPG keys. *soup* verifies GPG signatures any time Docker images are updated.

16.2.3 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.

16.2.4 Images

After installation, you can see all Docker images with the following command:

```
sudo docker images
```

16.2.5 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
```


16.2.6 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.

16.2.7 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.



16.2.8 Containers

Our Docker containers all belong to a common Docker bridge network, called `so-elastic-net`. 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.

16.2.9 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.

16.2.10 Dependencies

Playbook

so-playbook - REQ - Playbook Web App
so-navigator - OPT - Navigator Web App
so-soctopus - REQ - Automation

SOCTopus

so-soctopus - REQ - SOCTopus App
so-elasticsearch - OPT - Automation

Suricata

so-suricata - REQ - Suricata app

Kibana

so-kibana - REQ - Kibana Web App
so-elasticsearch - REQ -

Zeek

so-zeek - REQ - Zeek app

16.2.11 More Information

Note: For more information about Docker, please see <https://www.docker.com/what-docker>.

16.3 Jupyter Notebook


16.3.1 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>.

16.3.2 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:



The screenshot shows the Kibana user management interface for a user named 'jupyter'. The user's profile is displayed with a blue circular avatar containing the letter 'j'. The 'Profile' section includes fields for 'Username' (jupyter), 'Full name', and 'Email address'. The 'Privileges' section shows the user is assigned the 'viewer' role. A blue 'Update user' button and a 'Cancel' link are at the bottom.

Profile
Provide personal details.

Username
jupyter
Username can't be changed once created.

Full name

Email address

Privileges
Assign roles to manage access and permissions.

Roles
viewer ×

[Learn what privileges individual roles grant.](#)

Update user Cancel

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.

The screenshot displays the Security Onion web interface. On the left, a navigation sidebar includes 'Overview', 'Alerts', 'Dashboards', 'Hunt', 'Cases', 'PCAP', 'Grid', 'Downloads', 'Administration', 'Users', 'Grid Members', 'Configuration' (highlighted), and 'License Key'. Below this is a 'Tools' section with links to Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and features an 'Options' dropdown menu. A filter input field is present above a list of hostgroups: backup, bpf, elastalert, elasticfleet, elasticsearch, firewall, and hostgroups. Under 'hostgroups', the 'analyst' hostgroup is selected and highlighted in blue. To the right, a configuration panel for the 'analyst' hostgroup is shown, containing a 'VIEW DEFAULT' button, a 'Current Grid Value' input field with the value '192.168.199.0/24', and a 'Select a node to modify' dropdown menu. The footer of the interface shows 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

At the top of the page, click the Options menu and enable the Show all configurable settings, including 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.

16.3.3 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 warning like the following due to the fact that we are not performing verification for certificates:

```
/opt/conda/lib/python3.9/site-packages/elasticsearch/connection/http_urllib3.py:209: UserWarning: Connecting to https://192.168.6.100:9200 using SSL with verify_certs=False is insecure.
warnings.warn(
```

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:

	process	winlog	tags	@timestamp	file	@version	event	user
0	{'pid': 3956, 'entity_id': 'EBE732EE-504F-61A5...	{'execution': {'ThreadID': 4400, 'ProcessID': ...	velociraptor	2021-11-30T04:22:45.357Z	{'target': 'C:\Program Files\WindowsApps\Micro...	1	{'code': '11', 'module': 'sysmon', 'category': ...	NaN
1	{'pid': 3956, 'entity_id': 'EBE732EE-504F-61A5...	{'execution': {'ThreadID': 4400, 'ProcessID': ...	velociraptor	2021-11-30T04:22:45.357Z	{'target': 'C:\Program Files\WindowsApps\Micro...	1	{'code': '11', 'module': 'sysmon', 'category': ...	NaN
2	{'pid': 3956, 'entity_id': 'EBE732EE-504F-61A5...	{'execution': {'ThreadID': 4400, 'ProcessID': ...	velociraptor	2021-11-30T04:22:45.357Z	{'target': 'C:\Program Files\WindowsApps\Micro...	1	{'code': '11', 'module': 'sysmon', 'category': ...	NaN
3	{'pid': 3956, 'entity_id': 'EBE732EE-504F-61A5...	{'execution': {'ThreadID': 4400, 'ProcessID': ...	velociraptor	2021-11-30T04:22:45.357Z	{'target': 'C:\Program Files\WindowsApps\Micro...	1	{'code': '11', 'module': 'sysmon', 'category': ...	NaN
4	{'pid': 3956, 'entity_id': 'EBE732EE-504F-61A5...	{'execution': {'ThreadID': 4400, 'ProcessID': ...	velociraptor	2021-11-30T04:22:45.357Z	{'target': 'C:\Program Files\WindowsApps\Micro...	1	{'code': '11', 'module': 'sysmon', 'category': ...	NaN
...
3190	{'pid': 3224, 'entity_id': 'EBE732EE-6D06-61A5...	{'execution': {'ThreadID': 4400, 'ProcessID': ...	velociraptor	2021-11-30T01:00:55.162Z	{'target': 'C:\Windows\SoftwareDistribution\Do...	1	{'code': '', 'module': 'sysmon', 'category': ...	NaN
3191	{'parent': {'entity_id': 'EBE732EE-511E-61A5-9...	{'execution': {'ThreadID': 4400, 'ProcessID': ...	velociraptor	2021-11-30T01:00:55.162Z	NaN	1	{'code': '', 'module': 'sysmon', 'category': ...	{'name': 'NT AUTHORITY\SYSTEM'}

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):

	Dataset	Executable	Target
0	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.AspNetCore.SignalR.Protocols.Json.dll
1	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.Bcl.AsyncInterfaces.dll
2	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.AspNetCore.SignalR.Protocols.MessagePack.dll
3	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.Extensions.Caching.Abstractions.dll
4	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.AspNetCore.SignalR.Client.dll
5	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.AspNetCore.SignalR.Common.dll
6	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.Extensions.Caching.Memory.dll
7	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.AspNetCore.SignalR.Client.Core.dll
8	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Microsoft.AspNetCore.Http.Features.dll
9	file_create	C:\WINDOWS\system32\svchost.exe	C:\Program Files\WindowsApps\Microsoft.YourPhone_1.21102.134.0_x64__8wekyb3d8bbwe\YourPhoneServer\Humanizer.dll

Obviously, there is much more we can do with this data other than just running the above example code. Happy hunting!

16.4 Adding a new disk

If you ever need to add a new disk to expand your `/nsm` partition, there are at least 3 different ways to do this.

Warning: Before doing this in production, make sure you practice this on a non-production system!

16.4.1 Method 1: LVM (Logical Volume Management)

If you installed using LVM, then you should be able to use LVM to add new disk space to your LVM partitions.

16.4.2 Method 2: Mount a separate drive to /nsm

If you aren't using LVM, you can mount a drive directly to /nsm. If doing this after installation, you will need to stop services, move the data, and then restart services as shown below.

Stop services:

```
sudo systemctl disable salt-minion
sudo reboot
```

That should prevent most things from starting. If performing this on a manager you will need to do `sudo service docker stop` after the reboot.

Move the data:

```
sudo mv /nsm /nsm.old
sudo mkdir /nsm
# add your new file system to mount to /nsm in /etc/fstab
sudo mount -a
# make sure it's mounted correctly before continuing!
sudo mv /nsm.old/* /nsm/
sudo rm -rf /nsm.old
```

Restart services:

```
sudo systemctl enable salt-minion
sudo reboot
```

16.4.3 Method 3: Make /nsm a symlink to the new logging location

A variation on Method 2 is to make /nsm a symbolic link to the new logging location. Certain services like AppArmor may need special configuration to handle the symlink.

16.5 Network Installation

For most use cases, we **HIGHLY** recommend using our official Security Onion images as shown in the *Installation* section. Our official images are the only fully 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 full support.

If **NONE** of the above apply to you, you may be able to install one of the following operating systems and then perform a network installation:

- Oracle Linux 9
- Rocky Linux 9
- Alma Linux 9
- CentOS Stream 9
- RHEL 9
- Ubuntu 22.04
- Debian 12

For the least amount of issues, choose Oracle Linux 9 since it's used for our official images. Rocky Linux 9, CentOS Stream 9, and Alma Linux 9 should also work but they are not fully tested. Another option might be RHEL 9 itself although that is a paid option.

If you really want to run Ubuntu 22.04 or Debian 12, then please note that these distros may work but they get even less testing and therefore you will be more likely to run into issues. If you choose Ubuntu 22.04, we recommend the Ubuntu 22.04 Server ISO image and selecting the `Ubuntu Server` installation option as there are known issues when choosing the `Ubuntu Server (minimized)` option.

16.5.1 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 Requirements* 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

16.5.2 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.

1. Review the *Hardware Requirements* and *Release Notes* sections.
2. Download the ISO image for your desired x86-64 operating system. Verify the ISO image and then boot from it.
3. 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.
4. Reboot into your new installation.
5. Login using the username and password you specified during installation.
6. Install prerequisites. If you're using a RHEL flavor like Oracle Linux 9:

```
sudo dnf -y install git
```

If you're using a Debian flavor like Ubuntu:

```
sudo apt -y install git curl ethtool
```

7. Download our repo and start the Setup process:

```
git clone -b 2.4/main https://github.com/Security-Onion-Solutions/securityonion
cd securityonion
sudo bash so-setup-network
```

8. Proceed to the *Configuration* section.

16.6 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*.

16.6.1 Replay

You can use *tcpreplay* to replay any standard pcap to the sniffing interface of your Security Onion sensor.

16.6.2 Import

A drawback to using `tcpdump` is that it's replaying the pcap as new traffic and thus the timestamps that you see in *Security Onion Console (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.

16.7 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.

16.7.1 Removing from Salt

You can remove a node from *Salt* by going to *Administration* → *Grid Members*.

Security Onion

Grid Members

A distributed grid is made up of member nodes. Member nodes will request to join the grid and remain in a pending state until an administrator has accepted the node. If a pending member node is not yet listed as pending, then it's possible that the wrong manager host was provided during setup or there could be a connectivity problem.

Pending Members	Accepted Members
None	✓ securityonion_import REVIEW
Denied Members	
None	
Rejected Members	
None	

Version: 2.4.60 © 2024 Security Onion Solutions, LLC License: ELv2

Find the Grid Member you would like to remove, click the **REVIEW** button, and then click the **DELETE** button.

16.7.2 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
```

16.7.3 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.

16.8 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.

16.9 UTC and Time Zones

When you run Security Onion Setup, it sets the operating system timezone 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 timezone of your web browser and then render those UTC timestamps in local time. *Alerts*, *Dashboards*, and *Hunt* allow you to manually set your timezone under Options.

16.10 pfSense

pfSense is a free and open firewall that can be found at <https://www.pfsense.org/>. 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 Youtube video at <https://www.youtube.com/watch?v=aoH8qZwAxek>.

16.10.1 Simple Parser

The first option is to use our simple parser for pfSense firewall logs. 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.

The screenshot shows the Security Onion web interface. The left sidebar contains navigation options: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration (highlighted), Users, Grid Members, Configuration (highlighted), and License Key. Below the sidebar are various tools like Kibana, Elastic Fleet, Osquery Manager, InfluxDB, CyberChef, and Navigator. The main content area is titled 'Grid Configuration' and shows a list of hostgroups under the 'firewall' category. The 'analyst' hostgroup is selected. A modal window is open, showing a 'Current Grid Value' of '192.168.199.0/24' and a dropdown menu to 'Select a node to modify'. The footer of the interface displays 'Version: 2.4.60', '© 2024 Security Onion Solutions, LLC', and 'License: ELv2'.

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.

16.10.2 Elastic Integration for pfSense

The second option 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, we need to allow the traffic from the pfSense firewall to port 9001:

1. Navigate to *Administration* → Configuration.
2. At the top of the page, click the **Options** menu and then enable the **Show all configurable settings, including 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 dashboard to see your firewall logs.

16.11 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.

16.11.1 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.

16.11.2 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.

This section covers some of the utilities in Security Onion.

17.1 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.

17.1.1 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.

17.1.2 More Information

Note: For more information about jq, please see <https://stedolan.github.io/jq/>.

17.2 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.

The screenshot displays the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user profile icon. The left sidebar lists various system components, with 'Configuration' highlighted. The main panel shows the 'Grid Configuration' section, specifically the 'firewall' hostgroups. The 'analyst' hostgroup is selected, and a modal window is open, showing the 'Current Grid Value' as '192.168.199.0/24' and a dropdown menu to 'Select a node to modify'.

17.3 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.

17.4 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.

17.4.1 Usage

```
so-elasticsearch-query <PATH> [ARGS,...]
```

Where:

- `PATH` represents the elastic function being requested.
- `ARGS` is used to specify additional, optional curl parameters.

17.4.2 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
```

If you want to delete an old index, you can do that using the `-XDELETE` option. For example, to delete the *Zeek* index for 2022/05/07:

```
sudo so-elasticsearch-query so-zeek-2022.05.07 -XDELETE
```

17.5 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 *Security Onion Console (SOC)* can find them
- provide a hyperlink for you to view all alerts and logs in *Security Onion Console (SOC)*

In addition to viewing alerts and logs in *Security Onion Console (SOC)*, you can also find logs in *Kibana*.

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.

17.5.1 Screenshot

```
[doug@localhost ~]$ sudo so-import-pcap 2021-06-30-70551-Trickbot-with-DarkUNC-and-Cobalt-Strike.pcap
Processing Import: /home/doug/2021-06-30-70551-Trickbot-with-DarkUNC-and-Cobalt-Strike.pcap
- verifying file
- assigning unique identifier to import: 18211ba71a64e289f9e756bbf4b9d8a4
- analyzing traffic with Suricata
- analyzing traffic with Zeek
- found PCAP data spanning dates 2021-06-30 through 2021-06-30

Import complete!

Use the following hyperlink to view the imported data. Triple-click to quickly highlight the entire hyperlink and then copy it into a browser:
https://192.168.199.143/#/dashboards?q=import.id:18211ba71a64e289f9e756bbf4b9d8a4&20%7C%20groupby%20-sankey%20event.dataset%20event.category%2a%20%7C%20groupby%20-pie%20event.category%20%7C%20groupby%20-bar%20event.module%20%7C%20groupby%20event.dataset%20%7C%20groupby%20event.module%20%7C%20groupby%20event.category%20%7C%20groupby%20observer.name%20%7C%20groupby%20source.ip%20%7C%20groupby%20destination.ip%20%7C%20groupby%20destination.port&t=2021%2F06%2F30%2000%3A00%3A00%20AM%20-%202021%2F07%2F01%2000%3A00%3A00%20AM&z=UTC

or, manually set the Time Range to be (in UTC):
From: 2021-06-30 To: 2021-07-01

Note: It can take 30 seconds or more for events to appear in Security Onion Console.
[doug@localhost ~]$ _
```

17.5.2 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.

17.5.3 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.

17.5.4 Examples

If you don't already have some pcap files to import, see *PCAPs for Testing* 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!

17.6 so-import-evtx

`so-import-evtx` will import one or more evtx files into Security Onion.

17.6.1 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 *Security Onion Console (SOC)*. You can also find logs in *Kibana*.

17.7 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.

17.8 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.

17.8.1 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
```

17.9 so-test

so-test will run so-tcpreplay to replay some pcap samples to your sniffing interface.

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'
```

(continues on next page)

(continued from previous page)

```

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.

17.10 so-user

Security Onion Console (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
```

17.10.1 Listing SOC Users

To see a list of all SOC users, use the `list` operation:

```
sudo so-user list
```

17.10.2 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.

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: *Tools*
- Check log files in `/opt/so/log/` or other locations for any errors or possible clues:
 - Setup `/root/so/setup.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*.

18.1 FAQ

Install / Update / Upgrade

Users / Passwords

Support / Help

IDS engines

Security Onion internals

Tuning

Miscellaneous

18.1.1 Install / Update / Upgrade

Why won't the ISO image boot on my machine?

Please see the *Booting Issues* section.

What's the recommended procedure for installing Security Onion?

Please see the *Installation* 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 Requirements* section.
back to top

18.1.2 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

18.1.3 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

18.1.4 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

18.1.5 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 Structure* section.

Why does Security Onion use UTC?

Please see the *UTC and Time Zones* section.

Why are the timestamps in Kibana not in UTC?

Please see the *UTC and Time Zones* section.

Why is my disk filling up?

Security Onion records full packet capture to disk via *Stenographer*.

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. Endpoint agent traffic is encrypted where supported. This includes the *Elastic Agent* which supports encryption with additional configuration. 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

18.1.6 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 *Adding a new disk* section.

back to top

18.1.7 Miscellaneous

Where can I find interesting pcaps to replay?

Please see the *PCAPs for Testing* 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.

How can I add and test local rules?

Please see the *Adding Local Rules* section.

Can I connect Security Onion to Active Directory or LDAP?

We understand the appeal of integrating with directory services like Active Directory and LDAP, but we typically recommend against joining any security infrastructure (including Security Onion) to directory services. The reason is that when you get an adversary inside your network, one of their first goals is going to be gaining access to that directory. If they get access to the directory, then they get access to everything connected to the directory. For that reason, we recommend that all security infrastructure (including Security Onion) be totally separate from directory services.

back to top

18.2 Directory Structure

18.2.1 /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.

18.2.2 /opt/so/log

Debug logs are stored in /opt/so/log/.

18.2.3 /opt/so/rules

ElastAlert and *Suricata* rules are stored in /opt/so/rules/.

18.2.4 /opt/so/saltstack/local

Custom *Salt* settings can be added to /opt/so/saltstack/local/.

18.2.5 /nsm

The vast majority of data is stored in /nsm/.

18.2.6 /nsm/zeek

Zeek writes its protocol logs to /nsm/zeek/.

18.2.7 /nsm/elasticsearch

Elasticsearch stores its data in /nsm/elasticsearch/.

18.2.8 /nsm/pcap

Stenographer stores full packet capture in /nsm/pcap/.

18.3 Tools

Security Onion would like to thank the following projects for their contribution to our community!

(listed alphabetically)

- *ATT&CK Navigator*
- *CyberChef*
- *Docker*
- *ElastAlert*
- *Elasticsearch*
- *Elastic Agent*
- *InfluxDB*
- *Kibana*
- *Logstash*
- *Redis*

- *Salt*
- *Stenographer*
- *Strelka*
- *Suricata*
- *Zeek*

18.4 Support

18.4.1 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!

18.4.2 Community Support

If you need free support, you can reach out to our *Community Support*.

18.5 Community Support

18.5.1 Check Documentation First

First, check to see if your question has already been answered in the *Help* or *FAQ* sections.

18.5.2 Forum Guidelines

Before posting, please review the forum guidelines at <https://github.com/Security-Onion-Solutions/securityonion/discussions/1720>.

18.5.3 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>.

18.6 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.

18.6.1 Marketing Team

We need more folks to help spread the word about Security Onion by blogging, tweeting, and other social media.

18.6.2 Support Team

If you'd like help out other Security Onion users, please join the forum and start answering questions!

<https://securityonion.net/discuss>

18.6.3 Documentation Team

If you find that some information in our Documentation is incorrect or lacking, please feel free to submit Pull Requests via GitHub!

<https://github.com/Security-Onion-Solutions/securityonion-docs>

18.6.4 Core Development

Most of our code is on GitHub. Please feel free to submit pull requests!

<https://github.com/Security-Onion-Solutions>

18.6.5 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
- Scott Runnels
- Jon Schipp
- Brad Shoop
- Bryant Treacle
- William Wernert

19.1 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.

19.2 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.

RELEASE NOTES

20.1 Known Issues

If you notice an Elasticsearch status of `Pending` in the Grid interface, you can view affected indices 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.

20.1.1 2.4.60 [20240320] Changes

- FEATURE: Add Suricata classification.config for editing #12391
- FEATURE: Add Suricata support for full PCAP #12571
- FEATURE: Add default columns for endpoint.events datasets #12425
- FEATURE: Add new SOC action for Process Info #12421
- FEATURE: Add new endpoint dashboards #12428
- FEATURE: Additional Supported Integrations #5
- FEATURE: Improve Grid page Reboot indicators #12546
- FEATURE: Initial implementation of the new Detections system (currently disabled)
- FIX: Accept Uppercase emails #12559
- FIX: Change the default setting for steno diskfreepercentage on standalone installations to 21 #12541
- FIX: Download only newest packages for network installs
- FIX: EA packages are not downloadable once STIGs have been applied
- FIX: Endpoint diagnostic template pattern #12433
- FIX: Exclude templates from global overrides when necessary #12382

- FIX: Improve the accuracy of the stenoloss script #12477
- FIX: Receiver node Redis queue fills up using Managersearch without a Searchnode #12535
- FIX: Support Oinkcode values containing leading 0's #12506
- FIX: Update SOC annotations for Stenographer PCAP #12539
- FIX: Update correlate quick action with new icon #12387
- FIX: Update ks.cfg for appliances
- FIX: error.message mapping for system.syslog #12518
- FIX: so-saltstack-update should use the proper repo in 2.4 #12570
- UPGRADE: CyberChef 10.8.2 #12454
- UPGRADE: Kratos to 1.1.0 #12479
- UPGRADE: Suricata 7.0.4 #12609

20.1.2 2.4.50 [20240220] Changes

- FEATURE: Add Suricata PCAP module to Sensoroni (currently disabled) #12255
- FEATURE: Add new SOC action to show process ancestry #12345
- FEATURE: Add new dashboards for community_id and firewall auth #12323
- FEATURE: Additional Supported Integrations #4
- FEATURE: Allow user to create custom elastic search pipelines without copying them over via ssh
- FEATURE: Allow user to create custom logstash pipelines without copying them over via ssh
- FEATURE: Dedicated Fleet node should have an nginx entry and cert that works for /artifacts #11346
- FEATURE: Determine if Elastic is on its own mount point if so adjust size for watermark #12364
- FEATURE: Improve Correlate and Hunt actions on SOC Actions menu #12315
- FEATURE: RITA Logs #12226
- FEATURE: Support PCAP pivots for ICMP packets in SOC
- FIX: suricata.ike ingest pipeline does not exist #12174
- FIX: Add stenographer logging #12282
- FIX: Change field groupby button to new groupby #12228
- FIX: Correct SOC error messages related to malformed queries #12269
- FIX: Endpoint diagnostic collection index created with replicas #12256
- FIX: Expose node Reboot status as its own state; other grid/feature improvements
- FIX: Network Transport for suricata alerts should be lowercase #12217
- FIX: Strelka scan.pe.flags mapping #12251
- FIX: Sync the event.dataset values between the Windows Sysmon and ElasticAgent defend logs
- FIX: Syntax error running elastic fleet scripts during highstate
- FIX: User count logic providing inconsistent results #12258
- UPGRADE: CyberChef 10.6.0 #12310

- UPGRADE: Salt 3006.6 #12304
- UPGRADE: Strelka 0.24.01.18 #12229
- UPGRADE: Suricata 7.0.3 #12327
- UPGRADE: Zeek 6.0.3 #12225

20.1.3 2.4.40 [20240116] Changes

- FEATURE: Add geoip support to Suricata #11901
- FEATURE: Additional Supported Integrations #2 #11958
- FEATURE: Additional Supported Integrations #3 #12056
- FEATURE: Add server reboot notification to SOC #11852
- FEATURE: Allow an easy way to disable incoming events to a manager #12033
- FEATURE: Carve out the cert_chain_fps value from SSL traffic #11806
- FEATURE: Echotrail, Elasticsearch, MalwareBazaar, and ThreatFox Analyzers #12014
- FEATURE: Grid page status/metric enhancements #11971
- FEATURE: Manipulate event table columns #12145
- FEATURE: Sublime Platform Analyzer #11883
- FIX: Add force option to integrations #12017
- FIX: Adding extra_hosts for SOC, Elasticsearch and Logstash Docker containers fails #12015
- FIX: Begin kickstart consolidation
- FIX: Corrupt job files should not cause SOC to exit during startup #12082
- FIX: Disable Elastic Agent Downloads for Import and Eval mode
- FIX: Docker service sometimes not started or enabled on remote nodes during setup #12101
- FIX: Documentation links under SOC - Administration - Configuration need updating #11828
- FIX: FIM Integration #11847
- FIX: Ignore Zeek analyzer log #11892
- FIX: Improve salt-relay reponse integrity
- FIX: ISO image should default to 1GB /boot partition #12002
- FIX: Logstash pipeline to point to self instead of manager #12038
- FIX: Make sure optional integration pillar values are merged with defaults #12163
- FIX: Playbook Navigator Layer #11380
- FIX: Remove Curator
- FIX: Remove sudo entry for so-setup after setup completes
- FIX: Rerunning setup should uninstall local Elastic Agent #12030
- FIX: Show more readable column names for default Case list screen #12162
- FIX: SOC Hunt HTTP EXE query #11784
- FIX: so-elastic-fleet-reset non-destructive #12142

- FIX: so-playbook-reset #11790
- FIX: Update clear scripts #11991
- FIX: Update dashboard and hunt query for firewall logs #12021
- FIX: Update NIDS rule.reference in common.nids pipeline #11846
- UPGRADE: Salt 3006.5 #12143
- UPGRADE: SOC dependencies to latest versions #12041
- UPGRADE: Strelka 0.23.12.01 #11770

20.1.4 2.4.30 Hotfix [20231228] Changes

- FIX: Appliance kickstart files are not copying Elastic Agent tarballs #12081

20.1.5 2.4.30 Hotfix [20231219] Changes

- FIX: Update appliance kickstart scripts to fix issue with package copy #12044

20.1.6 2.4.30 Hotfix [20231204] Changes

- FIX: Choosing Desktop or IDH from ISO GRUB menu results in failure #11865
- FIX: Ensure airgap rule updates are being copied to the proper location #11932
- FIX: outdated import-evtx-logs pipeline versions #11889
- FIX: x509.pem_managed errors

20.1.7 2.4.30 Hotfix [20231121] Changes

- FIX: Salt minion service disabled highstate in upgrade to 2.4.30 #11851

20.1.8 2.4.30 Hotfix [20231117] Changes

- FIX: Elastic Defend Integration Policy Downgrade #11810
- FIX: Update SSL cert to avoid Google Chrome error (2.4) #11824

20.1.9 2.4.30 [20231113] Changes

- FEATURE: Additional Supported Integrations #11513
- FEATURE: Allow for BPF comments in SOC #11738
- FEATURE: OpenID Connect (OIDC) support
- FEATURE: so-elastic-fleet-reset #11697
- FEATURE: Sublime Platform Integration #11579
- FIX: Add -watch to soctopus saltstate for file SOCTopus.conf. Makes container restart @ highstate if file is updated. #11700

- FIX: Allow ICMP to allow a node to respond to ping #11495
- FIX: Allow standalone install type to work with 16GB of ram #11699
- FIX: Allow the setting up of data_warm to the nodes list in ES
- FIX: Data not returned from mine for network.ip_addrs #11502
- FIX: Delete all obsolete scripts and unused code (also check so-setup, so-functions)
- FIX: Fail so-setup if Elastic Fleet Setup encounters an error #11696
- FIX: Global BPF prevents new sensor from applying highstate #11610
- FIX: Improve error handling of Elasticsearch pipeline and template load scripts #11728
- FIX: Logs not parsed correctly when shipped from Fleet Node #11698
- FIX: Only heavy nodes should be treated as remote Elastic clusters in SOC #11553
- FIX: Reduce ISO size #11510
- FIX: Set days for warm for all so-* indices
- FIX: Show container download status during soup #11550
- FIX: Sigma DNS mapping #11498
- FIX: Suricata 7 pkt_src field needs to be parsed #11566
- FIX: The values for specific nodes in zeek.config.local.load are being populated incorrectly #11472
- UPGRADE: NetworkMiner 2.8.1 #11457
- UPGRADE: Salt 3006.3 #11529
- UPGRADE: SOC dependency Axios to 1.6.1 #11763
- UPGRADE: Sophos Integration #11548
- UPGRADE: Upgrade Elastic to 8.10.4
- UPGRADE: Upgrade InfluxDB to 2.7.1 and Telegraf to 1.28.2
- UPGRADE: Upgrade Suricata to 7.0.2
- UPGRADE: Zeek 6.0.2

20.1.10 2.4.20 Hotfix [20231012] Changes

- FIX: Elastic Defend Integration Policy Corrupted #11527

20.1.11 2.4.20 [20231006] Changes

- FEATURE: Add ingest parser for pfSense OpenVPN logs #7656
- FEATURE: Add new so-log-check tool to scan SO logging for anomalies
- FEATURE: Enable Analyzers to be managed through SOC #11211
- FEATURE: Grid screen improvements; support for desktop nodes
- FEATURE: Provide global replica value for index templates #10998
- FEATURE: SOC Grid Members should prompt for confirmation before actually deleting #11223
- FIX: Adding custom action to SOC causes the Endgame action to be replicated #11210

- FIX: Add Transform Role #11309
- FIX: CentOS stream 9 installation #11168
- FIX: Clean component template directory #11331
- FIX: Desktop via network install fails #10975
- FIX: Disable conn stats from being generated by default #11410
- FIX: Docker custom_bind_mounts not working for some containers #11122
- FIX: Duplicate cronjobs for filecheck #11400
- FIX: Elastic Agent - Installation “Not Accessible” Message #11191
- FIX: Elastic Fleet key and cert errors on heavynode #11026
- FIX: Exclude Zeek console log ingestion #11082
- FIX: Features pillar not showing all enabled features #11130
- FIX: Fleet plugin logs ERROR during kibana restart #10955
- FIX: Force nginx to run as user nobody #11402
- FIX: Heavy nodes are missing ElasticFleet integration policies #11189
- FIX: Heavy Nodes are not properly added to the soc.json #11192
- FIX: Improve consistency in cert storage across OS families #11162
- FIX: Improve default settings to avoid Elasticsearch hitting watermark #11305
- FIX: Kibana Elastic Agent Dashboard 404 #11018
- FIX: Maintain minion log in INFO level, add logrotate #10921
- FIX: Make sure a data stream is created for syslog #11209
- FIX: Make sure Elastic packages are loaded when changed #11428
- FIX: Minimum system requirements checks during setup #11324
- FIX: Minion log appears to show timezone bouncing #10922
- FIX: osquery not working on macOS
- FIX: Pre-load Integration Templates #11146
- FIX: Prevent repeated creation of unused Docker volumes #9941
- FIX: Remove default component templates to prevent conflicts #11260
- FIX: Remove OSSEC and add Playbook mappings for the SOC Alerts Event Table #11015
- FIX: Remove telegraf beats EPS script #11412
- FIX: Rename some SOC log fields to more unique field names #11429
- FIX: Reposync and yara rules shot not run in airgap #11427
- FIX: SOC Config pcap doc links should point to steno docs #11302
- FIX: SOC Config sensoroni doc links should point to correct docs #11362
- FIX: SOC doesn’t return user to login page after session expires #11438
- FIX: SOC fails to parse incomplete Elastic error response #11435
- FIX: SOC Grid Import inconsistency with larger files #11143

- FIX: Some packages are installed/removed and upgraded/downgraded every 15min #11458
- FIX: so-import-evtx incorrect dates #11332
- FIX: so-salt-minion-check not rendering as jinja #11390
- FIX: Stop zeek from trying to email reports #11407
- FIX: Strelka ingest pipeline should properly index entropy 0 values and float values in the same field
- FIX: Suricata filter and extraction rules are not properly updated #11229
- FIX: Update firewall docs for custom port and host groups #11053
- FIX: Update IDH Opencanary Modules to indicate they only apply to IDH nodes #10170
- UPGRADE: Kratos to v1.0.0
- UPGRADE: Suricata 6.0.14 #11319
- UPGRADE: Zeek 5.0.10 #11301

20.1.12 2.4.10 Hotfix [20230821] Changes

- FIX: Component templates not updated when packages are updated #11065
- FIX: Importing both PCAP and EVTX files fails #11030
- FIX: Logstash container missing on distributed receiver #11099
- FIX: pipeline with id logs-system.syslog-1.6.4 does not exist #11038
- FIX: Suricata permissions on Heavy Nodes are incorrect #11031

20.1.13 2.4.10 [20230815] Changes

- FEATURE: Auto-Upgrade Node Agents #10949
- FEATURE: Customize desktop environment #10957
- FIX: Custom actions, queries, tools can cause SOC restart to fail #11022
- FIX: Elastic Agents won't upgrade without Internet connection #10981
- FIX: Elastic Integrations not upgrading during SOUP #10984
- FIX: Elastic index settings annotations need synchronized with those specified in defaults #10999
- FIX: File extraction not working after switching from Zeek metadata to Suricata metadata #10973
- FIX: Fleet - url_base not working in cert CN #11003
- FIX: Improve wording for Firewall entries under Grid Administration Quick Links #10990
- FIX: Influx reporting No Results for Zeek Capture Loss #10956
- FIX: Suricata should not assume the interface will always be bond0 #10954
- FIX: Sysmon Events Table Field Rendering #10985
- FIX: so-desktop-install needs to change from Rocky to Oracle #10962
- FIX: soup may fail while trying to query Fleet server #10974

20.1.14 2.4.5 RC2 [20230807] Changes

- FEATURE: Add NetworkMiner to Security Onion Desktop #10865
- FEATURE: Add value from record in Hunt, etc as an observable to an existing or new case #7992
- FEATURE: Enable CommunityID for Elastic Defend Logs #10811
- FEATURE: Heavy Node Support #10671
- FEATURE: so-import-evt - timeshift #10743
- FEATURE: soup should rotate its log file #10951
- FIX: Dashboards with multiple groupby charts always filter by the first chart's, first groupby field #10856
- FIX: Disable offload on monitor NICs #10900
- FIX: EQL Field Mappings #10783
- FIX: Elastic Fleet Improvements #10846
- FIX: Firewall state custom host group assignments for single portgroup entry #10917
- FIX: IDH node #10882
- FIX: IPTables Persistence #10884
- FIX: Install Error: so-yara-download failed #10880
- FIX: Install screen - Firewall #10945
- FIX: List settings updated with blank values should be stored as empty lists #10936
- FIX: Login page shows error banner briefly on initial page load #10911
- FIX: RAID status on Grid page #10935
- FIX: SOC Auth dashboard #10878
- FIX: Security Onion Desktop state should default to Gnome Classic #10958
- FIX: sensor MTU setting in SOC Config should be read only #10883
- FIX: so-status taking several seconds to complete #10909
- FIX: soup #10902
- FIX: syslog not working #10896
- FIX: verbiage and links in soc_sensor.yaml #10906
- UPGRADE: Elastic 8.8.2 #10864

20.1.15 2.4.4 RC1 [20230728] Changes

- FEATURE: Add DNS lookup action to SOC #8655
- FEATURE: Add Oracle Linux Support #10844
- FEATURE: Add pivots for relational operators on numbers #8024
- FEATURE: Add relative Timeframe and Refresh Interval as URL Parameters to Hunt #3352
- FEATURE: Cases - Add ability to enable dynamic observable extraction #7972
- FEATURE: Oracle Linux ISO #10845

- FEATURE: Security Onion Desktop #10862
- FIX: Add retry to Elastic Agent installer #10488
- FIX: Case status code 404 error #10759
- FIX: Intermittent pcap retrieval #10750
- FIX: Navigator Errors #10742
- FIX: Remove .security subfield #10745
- UPGRADE: CyberChef 10.5.2 #10781
- UPGRADE: so-registry docker image #10727

20.1.16 2.4.3 Beta 4 [20230711] Changes

- FEATURE: Add link to Downloads page for convenient access to firewall settings #10702
- FEATURE: Add more SOC Config quick links #10563
- FEATURE: Add time zone selection to Grid page #8629
- FEATURE: Add webauthn support to SOC #10608
- FEATURE: Allow import of PCAP and EVTX via SOC UI #10413
- FEATURE: Elastic Fleet - Automatically Update Logstash Outputs #10746
- FEATURE: Elastic Fleet Server URL - Custom Domain #10744
- FEATURE: Supported Integrations #10590
- FEATURE: so-import-evtx #10673
- FIX: Strelka rule path #10715
- FIX: 2.4 ISO image won't install on Virtualbox #10534
- FIX: Account for Suricata XFF function in parsing and ingestion #8643
- FIX: Add more Zeek logs to excluded list #10569
- FIX: Analyzer requests and whoisit updates #10524
- FIX: Change Playbook index to data stream and update event.severity_label #10523
- FIX: Cleanup log-rotate.conf #10545
- FIX: Curator should ignore empty list #10512
- FIX: Don't override default integration ingest node pipelines #10542
- FIX: Ensure operations on records with "Missing" fields use correct search #8025
- FIX: Ensure packages aren't installed from default Rocky repos #10630
- FIX: Exclude System logs from Hunt/Dashboard Queries. #10122
- FIX: Finish SSL cert integration into SOC config UI #10533
- FIX: Improve SOC login error message for disabled users #8908
- FIX: Increase net.core.wmem_default value #10602
- FIX: InfluxDB NSM Disk Usage visualization #10520
- FIX: Integration logs not parsed correctly #10672

- FIX: Logstash soc.fields.query warning #10528
- FIX: Node description config setting should only apply at the node level #10562
- FIX: Remove default excluded rules from YARA repo #10718
- FIX: Review Kibana Dashboards #10664
- FIX: Rework dataset name and add tags based on suffix #10526
- FIX: Rework field to account for missing classifiers #10420
- FIX: SOC Config NTP quick link #10519
- FIX: Scheduled jobs trying to run during setup #10468
- FIX: Set Elastic Fleet certs to use url_base #10510
- FIX: Setup re-runs when SSH'ing into a successfully installed minion node #10498
- FIX: Strelka rule exclusions #10716
- FIX: Suricata DHCP logs not ingesting #10565
- FIX: Suricata dataset values for certain types of metadata #10551
- FIX: Update README.md #10554
- FIX: Update cheat sheet for 2.4 #10532
- UPGRADE: CyberChef 10.4.0 #10581
- UPGRADE: Suricata 6.0.13 #10594

20.1.17 2.4.2 Beta 3 [20230531] Changes

- FEATURE: Add additional alerts for Influxdb #10388
- FEATURE: Add link to SOC error messages that takes user to hunt and auto-searches for recent SOC-related errors. #10283
- FEATURE: Add Protected checkbox on Attachment upload form #10203
- FEATURE: Add support for Apple Silicon Elastic Agent Installer #10473
- FEATURE: Add support for EQL to Playbook #10471
- FEATURE: Allow for any docker container to have extra hosts and custom binds #10301
- FEATURE: Allow users to switch between airgap and non airgap. #10470
- FEATURE: Dedicated Elastic Fleet Node #10474
- FEATURE: Enable Elastic Defend Integration on Endpoints Policy #10475
- FEATURE: Integrate Elastic Artifact Repo #10053
- FEATURE: Integrate Elastic Package Registry #10472
- FEATURE: ISO image #10476
- FEATURE: Link the Grid Interface with Docker container log files #10149
- FEATURE: Prompt user to verify the manager nodes IP address if a DNS record is found during setup. #10334
- FEATURE: Quicklinks to common configs #10395
- FEATURE: SOC config UI should process each line individually with regex when multiline: True is set #10243

- FEATURE: Support authentication rate limiting #10308
- FIX: AWS Instances with forced IMDSv2 enabled fail to detect running in AWS #10205
- FIX: Cluster delete script should use different disk space logic when /nsm is shared among services #10418
- FIX: Correct SOC Annotations for idstools in Grid Configuration. #10208
- FIX: Correct SOC Annotations of Zeek in Grid Configuration. #10211
- FIX: Hunt Quick Drilldown #10377
- FIX: If mdengine is changed to Suricata, Zeek is still shown in so-status #10232
- FIX: Improve SOC configuration handling of lists #10219
- FIX: Improve soup's local file modification logic #8972
- FIX: In distributed deployment, Dashboards/Kibana only show data from the first sensor added. #10231
- FIX: Influxdb Elasticsearch cells showing duplicate data. #10336
- FIX: Kibana: Ensure _id fields beginning with a hyphen work properly when pivoting to SOC from Kibana #10305
- FIX: Logstash WARN logstash.outputs.elasticsearch on searchnode #10291
- FIX: Prepare SOUP for 2.4 #10056
- FIX: Prevent duplicate observables from being automatically created when attaching events to a case. #10123
- FIX: Review 2.4 file permissions and other local security changes #9110
- FIX: Setting CPU affinity or number of threads for Suricata not being applied. #10240
- FIX: Simplify cloud detection #10261
- FIX: Some SOC Config settings are only visible when Advanced is enabled #10429
- FIX: Strelka YARA Compilation #10271
- FIX: Suricata ignores the threads and always is set to 1 #10230
- FIX: Unable to disable PCAP via web configuration #10229
- FIX: Use pillar values to allow Zeek log ingestion selection from the UI #10322
- FIX: Zeek local policies are not being updated when changed in Current Grid value. #10209
- FIX: Zeek not ignoring lb_procs when Zeek pins configured #10215
- UPGRADE: Elastic 8.7.1 #10269
- UPGRADE: Kratos to 0.13.0 #10309
- UPGRADE: SOC external dependencies #10268
- UPGRADE: Suricata 6.0.12 #10311
- UPGRADE: Zeek 5.0.9 #10374

20.1.18 2.4.1 Beta 2 [20230424] Changes

- FIX: Add Dedicated Fleet Node #10054
- FIX: Don't create curl.config on Forward Nodes #10057
- FIX: Force case attachments to be downloaded #10186
- FIX: Improve Elasticsearch index deletion - so-elastic-clear #10109
- FIX: Improve Elasticsearch index deletion - so-elastic-cluster-delete-delete #10110
- FIX: Make sure Setup image downloads populate the screen and the log #10052
- FIX: Overview Customization link #10173
- FIX: Prevent Jinja syntax from being entered into config values via UI/API #10187
- FIX: Prevent Zeek from using a large amount of memory #10190
- FIX: Remove legacy Kibana dashboards #8555
- FIX: Remove template load from search nodes in distrib #10060
- FIX: SOC only displaying data for users assigned the superuser role #10068
- FIX: Sort grid members lists #10185
- FIX: Suricata DNS A and CNAME parsing #10117
- FIX: Using SOC Configuration to change mdengine from ZEEK to SURICATA fails #10189
- FIX: Zeek @local and @local-sigs need to strip the @ for config but replace in local.zeek #10050
- FIX: Zeek is not honoring lbprocs #10062
- UPGRADE: Elastic 8.7.0 #10059
- UPGRADE: Suricata 6.0.11 #10067
- UPGRADE: Zeek 5.0.8 #10107

20.1.19 2.4.0 Beta 1 [20230328] Changes

<https://blog.securityonion.net/2023/03/security-onion-24-beta-release-now.html>

APPENDIX

This appendix provides an overview of the process of migrating from the old Security Onion 2.3 to the new Security Onion 2.4.

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 2.4 is a MAJOR change, so please note the following:

- Security Onion 2.4 has higher hardware requirements, so you should check that your hardware meets those requirements.
- The /nsm partition must be on a separate disk.
- InfluxDB data is not migrated.
- If you have a distributed deployment, please note that 2.3 search nodes defaulted to cross cluster search whereas 2.4 defaults to full Elastic clustering. This means that you may need to rename or delete some Elasticsearch indices.
- We do not provide any guarantees that the upgrade process will work! If the upgrade fails, be prepared to perform a fresh installation of Security Onion 2.4.

For the reasons listed above, we recommend that most users procure new hardware and perform a fresh installation of Security Onion 2.4.

Tip: If you're planning to purchase new hardware, please consider official Security Onion appliances from Security Onion Solutions (<https://securityonionsolutions.com>). Our custom appliances have already been designed for certain roles and traffic levels and have Security Onion 2 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!

Warning: We recommend trying this process in a test environment before attempting in your production environment.

Warning: Please ensure that you have local access to the machine being upgraded via console, DRAC, IPMI, etc. Failure to do so could result in an unsuccessful upgrade, requiring a clean installation of Security Onion 2.4.

If you have reviewed all of the warnings above and still want to attempt migration, you should be able to do the following.

Note: If you have a distributed deployment, you will need to perform the steps on the manager first and then on each of the remaining nodes.

First, make sure that your 2.3 installation is fully updated via *soup*:

```
sudo soup
```

Next, make sure there is a backup in `/nsm/backup`:

```
sudo ls -alh /nsm/backup
```

Disable services and reboot:

```
sudo systemctl disable salt-minion
sudo reboot
```

Make sure docker containers are stopped:

```
sudo su -c "systemctl stop docker docker.socket"
sudo docker ps
```

If there are any remaining docker processes, stop them (replacing `$CONT_ID` with the actual ID):

```
sudo docker stop $CONT_ID
```

This can also be done in one command:

```
sudo docker ps | awk '!/CONTAINER/ { system("sudo docker stop " $1 ) }'
```

Unmount `/nsm`:

```
sudo umount /nsm
sudo vgchange -an /dev/mapper/nsm
sudo vgexport /dev/mapper/nsm
```

Boot the Security Onion 2.4 ISO image and go through the initial OS installation as shown in the *Installation* section.

Warning: During the OS installation, do NOT select the old NSM disk! It will be re-imported and used AFTER the OS install with LVM.

After installation and reboot, log in as the user you created during installation. Setup will automatically start, but you'll need to cancel setup and change partitioning (replacing `/home/user/` with your desired temporary location):

```
sudo cp -av /nsm/* /home/user/
sudo umount /nsm
sudo lvremove /dev/system/nsm

sudo lvresize -L +XG /dev/system/root
sudo xfs_growfs /dev/system/root
```

(continues on next page)

(continued from previous page)

```
sudo vgimport /dev/mapper/nsm
sudo vgchange -ay /dev/mapper/nsm
```

Add entry into `/etc/fstab` and then mount:

```
sudo mount -a
sudo systemctl daemon-reload
```

Remove `/nsm/repo` and `/nsm/docker-registry` from the old 2.3 `/nsm`.

Copy the `/nsm` contents of `/home/user/` (or wherever they were copied to) back to `/nsm` (`repo`, `docker-registry`, and `elastic-fleet`)

Run through setup as described in the [Configuration](#) section.

After setup, get the secrets pillar from `/nsm/backup` (replacing `2023_08_30` with the date of your most recent backup):

```
tar -xvf /nsm/backup/so-config-backup-2023_08_30.tar opt/so/saltstack/local/pillar/
↪secrets.sls
```

Replace the `mysql` secret in `secrets.sls` with the backed-up value:

```
docker exec -it so-mysql mysql -u root -p
# when prompted, enter the password from the 2.3 secrets.sls
```

At the `mysql` prompt, run the following query:

```
SELECT User, Host from mysql.user;
```

If you get the error `mysql error 1130: 172.17.1.1' is not allowed to connect to this mysql server`, then run the following:

```
UPDATE mysql.user SET Host = '172.17.1.1' WHERE User = 'root' AND Host = 'localhost';
```

Exit the `mysql` shell and restart the `so-mysql` container.

Run a full checkin:

```
sudo so-checkin
```


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/>