

---

# **Security Onion Documentation**

***Release 2.3***

**Security Onion Solutions, LLC**

**Apr 06, 2024**



# TABLE OF CONTENTS

<b>1</b>	<b>About</b>	<b>1</b>
1.1	Security Onion . . . . .	1
1.2	Security Onion Solutions, LLC . . . . .	1
1.3	Documentation . . . . .	2
<b>2</b>	<b>Introduction</b>	<b>5</b>
2.1	Network Security Monitoring . . . . .	7
2.2	Enterprise Security Monitoring . . . . .	8
2.3	Analysis Tools . . . . .	8
2.4	Workflow . . . . .	12
2.5	Deployment Scenarios . . . . .	13
2.6	Conclusion . . . . .	13
<b>3</b>	<b>License</b>	<b>15</b>
<b>4</b>	<b>First Time Users</b>	<b>17</b>
<b>5</b>	<b>Getting Started</b>	<b>45</b>
5.1	Best Practices . . . . .	46
5.2	Architecture . . . . .	47
5.3	Hardware Requirements . . . . .	55
5.4	Partitioning . . . . .	61
5.5	Download . . . . .	62
5.6	VMware . . . . .	63
5.7	VirtualBox . . . . .	65
5.8	Proxmox . . . . .	66
5.9	Bootng Issues . . . . .	66
5.10	Airgap . . . . .	67
5.11	Installation . . . . .	67
5.12	AWS Cloud AMI . . . . .	69
5.13	Azure Cloud Image . . . . .	75
5.14	Configuration . . . . .	80
5.15	After Installation . . . . .	81
<b>6</b>	<b>Security Onion Console (SOC)</b>	<b>83</b>
6.1	Alerts . . . . .	85
6.2	Dashboards . . . . .	92
6.3	Hunt . . . . .	103
6.4	Cases . . . . .	104
6.5	PCAP . . . . .	114
6.6	Grid . . . . .	117

6.7	Downloads	118
6.8	Administration	118
6.9	Kibana	119
6.10	Grafana	123
6.11	CyberChef	126
6.12	Playbook	130
6.13	FleetDM	134
6.14	ATT&CK Navigator	136
<b>7</b>	<b>Analyst VM</b>	<b>139</b>
7.1	NetworkMiner	140
7.2	Wireshark	143
<b>8</b>	<b>Network Visibility</b>	<b>149</b>
8.1	AF-PACKET	150
8.2	Stenographer	150
8.3	Suricata	153
8.4	Zeek	158
8.5	Strelka	168
8.6	Intrusion Detection Honeypot	174
<b>9</b>	<b>Host Visibility</b>	<b>181</b>
9.1	osquery	181
9.2	Beats	184
9.3	Wazuh	186
9.4	Syslog	192
9.5	Sysmon	193
9.6	Autoruns	194
<b>10</b>	<b>Logs</b>	<b>197</b>
10.1	Ingest	197
10.2	Filebeat	199
10.3	Logstash	217
10.4	Redis	222
10.5	Elasticsearch	223
10.6	ElastAlert	234
10.7	Curator	237
10.8	Data Fields	239
10.9	Alert Data Fields	240
10.10	Elastalert Fields	241
10.11	Zeek Fields	241
10.12	Community ID	242
10.13	Re-Indexing	242
10.14	SOC Logs	243
10.15	Other Supported Logs	244
10.16	Related Information	247
<b>11</b>	<b>Updating</b>	<b>253</b>
11.1	soup	253
11.2	End Of Life	258
<b>12</b>	<b>Accounts</b>	<b>261</b>
12.1	Passwords	261
12.2	MFA	262
12.3	Adding Accounts	263



12.4	Listing Accounts . . . . .	264
12.5	Disabling Accounts . . . . .	265
12.6	Role-Based Access Control (RBAC) . . . . .	265
<b>13</b>	<b>Services</b>	<b>271</b>
<b>14</b>	<b>Customizing for Your Environment</b>	<b>273</b>
14.1	SOC Customization . . . . .	273
14.2	Proxy Configuration . . . . .	278
14.3	Firewall . . . . .	279
14.4	Email Configuration . . . . .	284
14.5	NTP . . . . .	285
14.6	SSH . . . . .	285
14.7	Changing Hostname . . . . .	287
14.8	Changing IP Addresses . . . . .	287
14.9	Changing Web Access URL . . . . .	287
<b>15</b>	<b>Tuning</b>	<b>289</b>
15.1	Salt . . . . .	289
15.2	Homenet . . . . .	291
15.3	BPF . . . . .	293
15.4	Managing Rules . . . . .	295
15.5	Adding Local Rules . . . . .	298
15.6	Managing Alerts . . . . .	299
15.7	High Performance Tuning . . . . .	307
<b>16</b>	<b>Tricks and Tips</b>	<b>309</b>
16.1	Backups . . . . .	309
16.2	Docker . . . . .	310
16.3	DNS Anomaly Detection . . . . .	313
16.4	ICMP Anomaly Detection . . . . .	314
16.5	Jupyter Notebook . . . . .	314
16.6	Machine Learning . . . . .	317
16.7	Adding a new disk . . . . .	319
16.8	PCAPs for Testing . . . . .	320
16.9	Removing a Node . . . . .	321
16.10	Syslog Output . . . . .	322
16.11	UTC and Time Zones . . . . .	322
<b>17</b>	<b>Utilities</b>	<b>323</b>
17.1	jq . . . . .	323
17.2	so-allow . . . . .	323
17.3	so-elastic-auth . . . . .	325
17.4	so-elasticsearch-query . . . . .	325
17.5	so-import-pcap . . . . .	326
17.6	so-import-evtx . . . . .	327
17.7	so-monitor-add . . . . .	328
17.8	so-status . . . . .	329
17.9	so-test . . . . .	330
17.10	so-zeek-logs . . . . .	331
<b>18</b>	<b>Help</b>	<b>333</b>
18.1	FAQ . . . . .	333
18.2	Directory Structure . . . . .	337
18.3	Tools . . . . .	338

18.4	Support . . . . .	339
18.5	Community Support . . . . .	339
18.6	Help Wanted . . . . .	340
<b>19</b>	<b>Security</b>	<b>343</b>
19.1	Vulnerability Disclosure . . . . .	343
19.2	Product and Supply Chain Integrity . . . . .	343
<b>20</b>	<b>Appendix</b>	<b>345</b>
20.1	Appendix A: Security Onion 16.04 . . . . .	345
20.2	Appendix B: Ubuntu 18.04 . . . . .	348
<b>21</b>	<b>Release Notes</b>	<b>351</b>
21.1	2.3.300 [20240401] Changes . . . . .	351
21.2	2.3.290 [20240229] Changes . . . . .	351
21.3	2.3.280 [20231128] Changes . . . . .	351
21.4	2.3.270 [20231006] Changes . . . . .	352
21.5	2.3.260 [20230620] Changes . . . . .	352
21.6	2.3.250 [20230519] Changes . . . . .	352
21.7	2.3.240 [20230426] Changes . . . . .	353
21.8	2.3.230 [20230417] Changes . . . . .	353
21.9	2.3.220 Hotfix [20230301] Changes . . . . .	353
21.10	2.3.220 [20230224] Changes . . . . .	353
21.11	2.3.210 [20230202] Changes . . . . .	353
21.12	2.3.200 [20230113] Changes . . . . .	354
21.13	2.3.190 Hotfix [20221207] Changes . . . . .	354
21.14	2.3.190 [20221205] Changes . . . . .	354
21.15	2.3.182 [20221109] Changes . . . . .	355
21.16	2.3.181 [20221021] Changes . . . . .	355
21.17	2.3.180 [20221014] Changes . . . . .	355
21.18	2.3.170 [20220922] Changes . . . . .	355
21.19	2.3.160 [20220829] Changes . . . . .	355
21.20	2.3.150 [20220820] Changes . . . . .	356
21.21	2.3.140 Hotfix [20220812] Changes . . . . .	356
21.22	2.3.140 Hotfix [20220719] Changes . . . . .	356
21.23	2.3.140 Changes . . . . .	356
21.24	2.3.130 Changes . . . . .	357
21.25	2.3.120 Changes . . . . .	358
21.26	2.3.110 Hotfix [20220407] Changes . . . . .	359
21.27	2.3.110 Hotfix [20220407] Known Issues . . . . .	359
21.28	2.3.110 Hotfix [20220405] Changes . . . . .	359
21.29	2.3.110 Hotfix [20220401] Changes . . . . .	359
21.30	2.3.110 Changes . . . . .	359
21.31	2.3.100 Hotfix [20220301] Changes . . . . .	360
21.32	2.3.100 Hotfix [20220203] Changes . . . . .	360
21.33	2.3.100 Hotfix [20220202] Changes . . . . .	360
21.34	2.3.100 Changes . . . . .	360
21.35	2.3.91 Changes . . . . .	361
21.36	2.3.90 Hotfix [20211213] . . . . .	361
21.37	2.3.90 Hotfix [20211210] . . . . .	362
21.38	2.3.90 Hotfix [20211206] . . . . .	362
21.39	2.3.90 Hotfix [AIRGAPFIX] . . . . .	362
21.40	2.3.90 Hotfix [WAZUH] . . . . .	362
21.41	2.3.90 Changes . . . . .	362

21.42	2.3.80	Changes	364
21.43	2.3.70	Hotfix [WAZUH]	365
21.44	2.3.70	Hotfix [GRAFANA_DASH_ALLOW]	365
21.45	2.3.70	Hotfix [CURATOR]	365
21.46	2.3.70	Changes	365
21.47	2.3.61	Hotfix [STENO, MSEARCH]	366
21.48	2.3.61	Changes	366
21.49	2.3.60	Hotfix [ECSFIX, HEAVYNODE, FBPIPELINE, CURATORAUTH] Changes	366
21.50	2.3.60	Changes	367
21.51	2.3.52	Changes	368
21.52	2.3.51	Changes	368
21.53	2.3.50	Changes	368
21.54	2.3.50	Known Issues	370
21.55	2.3.40	Changes	370
21.56	2.3.40	Known Issues	371
21.57	2.3.30	Changes	372
21.58	2.3.30	Known Issues	373
21.59	2.3.21	Changes	373
21.60	2.3.10	Changes	375
21.61	2.3.10	Known Issues	377
21.62	2.3.2	Changes	377
21.63	2.3.1	Changes	377
21.64	2.3.1	Known Issues	377
21.65	2.3.0	Changes	378
21.66	2.2.0	Changes	379
21.67	2.1.0	Changes	380
21.68	2.0.3	Changes	380
21.69	2.0.2	Changes	381
21.70	2.0.1	Changes	381
21.71	2.0.0	Changes	381
<b>22</b>		<b>Cheat Sheet</b>	<b>383</b>



## 1.1 Security Onion

Security Onion is a free and open platform for threat hunting, enterprise security monitoring, and log management. It includes our own tools for *Alerts*, *Dashboards*, *Hunt*, *PCAP*, and *Cases* as well as other tools such as *Playbook*, *FleetDM*, *osquery*, *CyberChef*, *Elasticsearch*, *Logstash*, *Kibana*, *Suricata*, *Zeek*, and *Wazuh*. 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 an army of distributed sensors for your enterprise in minutes!

---

**Note:** Security Onion started in 2008 and was originally based on the Ubuntu Linux distribution. Throughout the years, the Security Onion version tracked the version of Ubuntu it was based on. For example, the last major version of Security Onion was based on Ubuntu 16.04 and so it was called Security Onion 16.04. Security Onion is now container based and thus no longer limited to just Ubuntu. To signify this change, Security Onion now has its own versioning scheme and this new platform is Security Onion 2.

---

## 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/2.3/>.

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 2.3 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](https://thomas-cokelaer.info/tutorials/sphinx/rest_syntax.html).

### 1.3.5 Naming Convention

Our goal is to allow you to easily guess and type the URL of the documentation you want to go to. For example, if you want to read more about Suricata, you can type the following into your browser: <https://securityonion.net/docs/suricata>

To achieve this goal, 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)

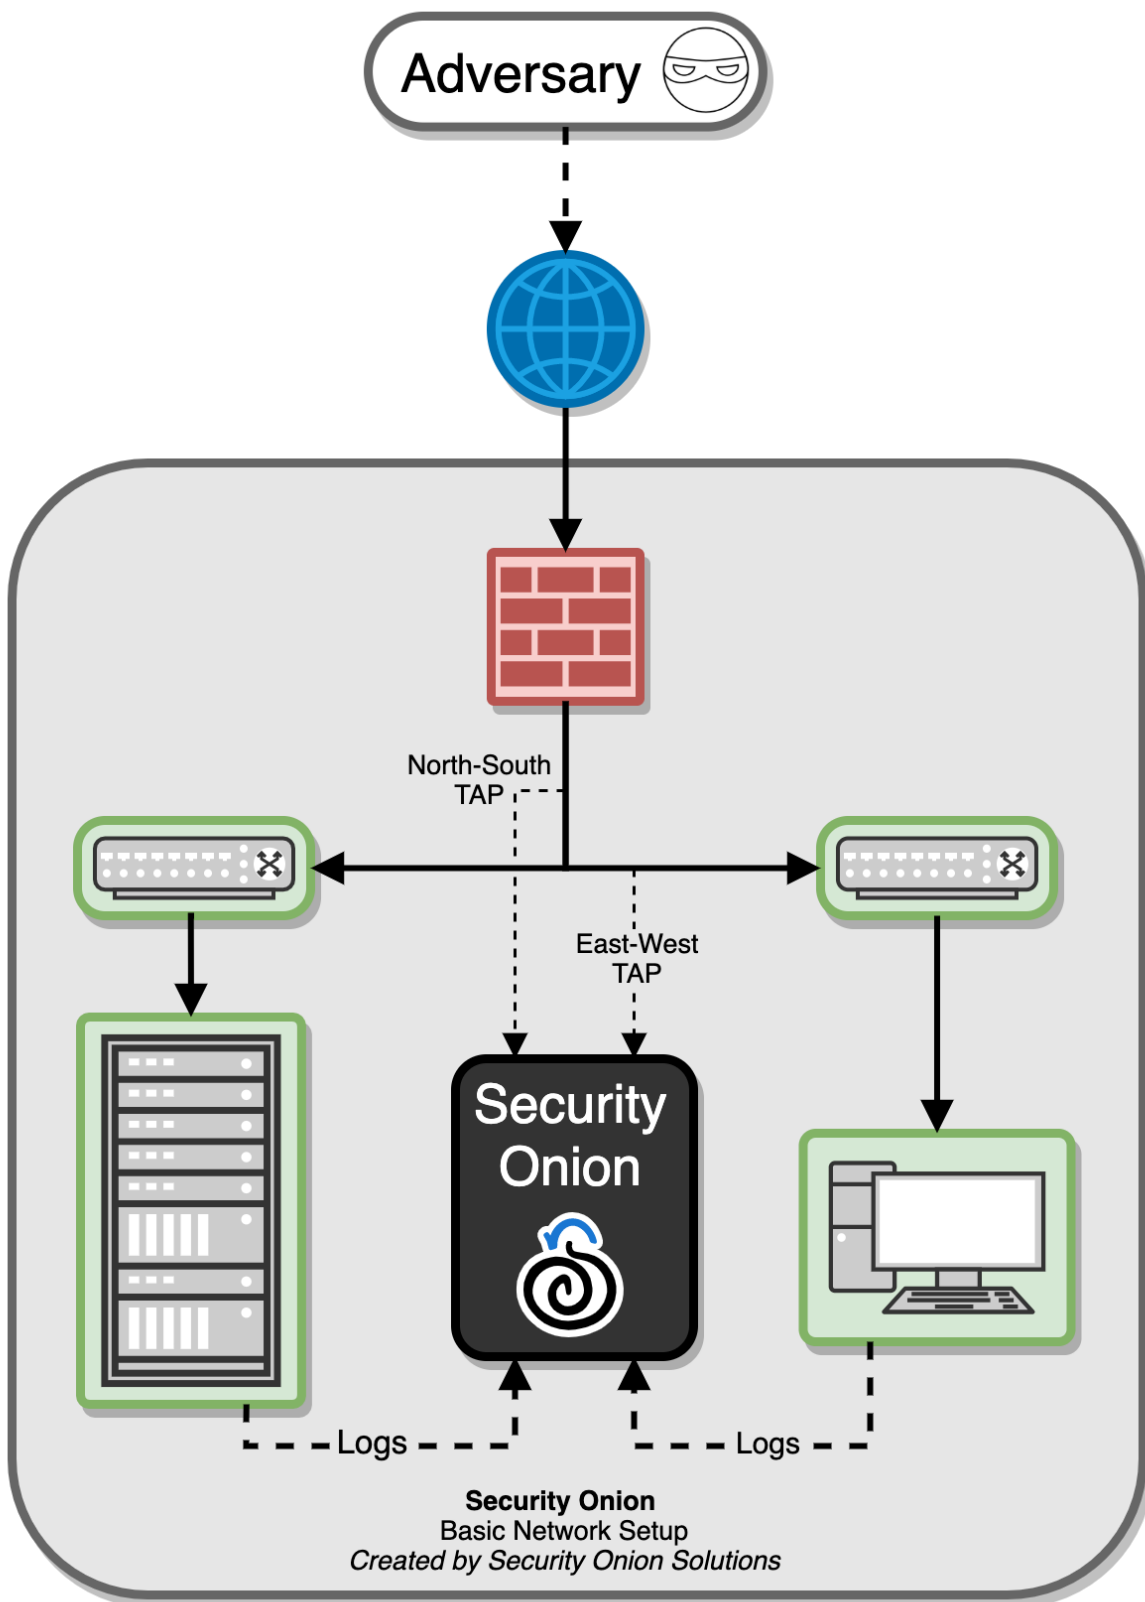




## **INTRODUCTION**

Security Onion is a free and open platform for Network Security Monitoring (NSM) and Enterprise Security Monitoring (ESM). NSM is, put simply, monitoring your network for security related events. It might be proactive, when used to identify vulnerabilities or expiring SSL certificates, or it might be reactive, such as in incident response and network forensics. Whether you're tracking an adversary or trying to keep malware at bay, NSM provides context, intelligence, and situational awareness of your network. ESM takes NSM to the next level and includes endpoint visibility and other telemetry from your enterprise.

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.



Many assume NSM is a solution they can buy to fill a gap; purchase and deploy solution XYZ and problem solved. The belief that you can buy an NSM denies the fact that the most important word in the NSM acronym is “M” for Monitoring. Data can be collected and analyzed, but not all malicious activity looks malicious at first glance. While automation and correlation can enhance intelligence and assist in the process of sorting through false positives and malicious indicators, there is no replacement for human intelligence and awareness. We don’t want to disillusion you. Security Onion isn’t a silver bullet that you can setup, walk away from and feel safe. Nothing is and if that’s what you’re looking for you’ll never find it. Security Onion will provide visibility into your network traffic and context around alerts and anomalous events, but it requires a commitment from you the defender to review alerts, monitor the network activity, and most importantly, have a willingness, passion, and desire to learn.

## 2.1 Network Security Monitoring

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)

Security Onion includes an *Intrusion Detection Honeypot* 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.

## 2.2 Enterprise Security Monitoring

In addition to network visibility, Security Onion provides endpoint visibility via agents like *Beats*, *osquery*, and *Wazuh*.

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)

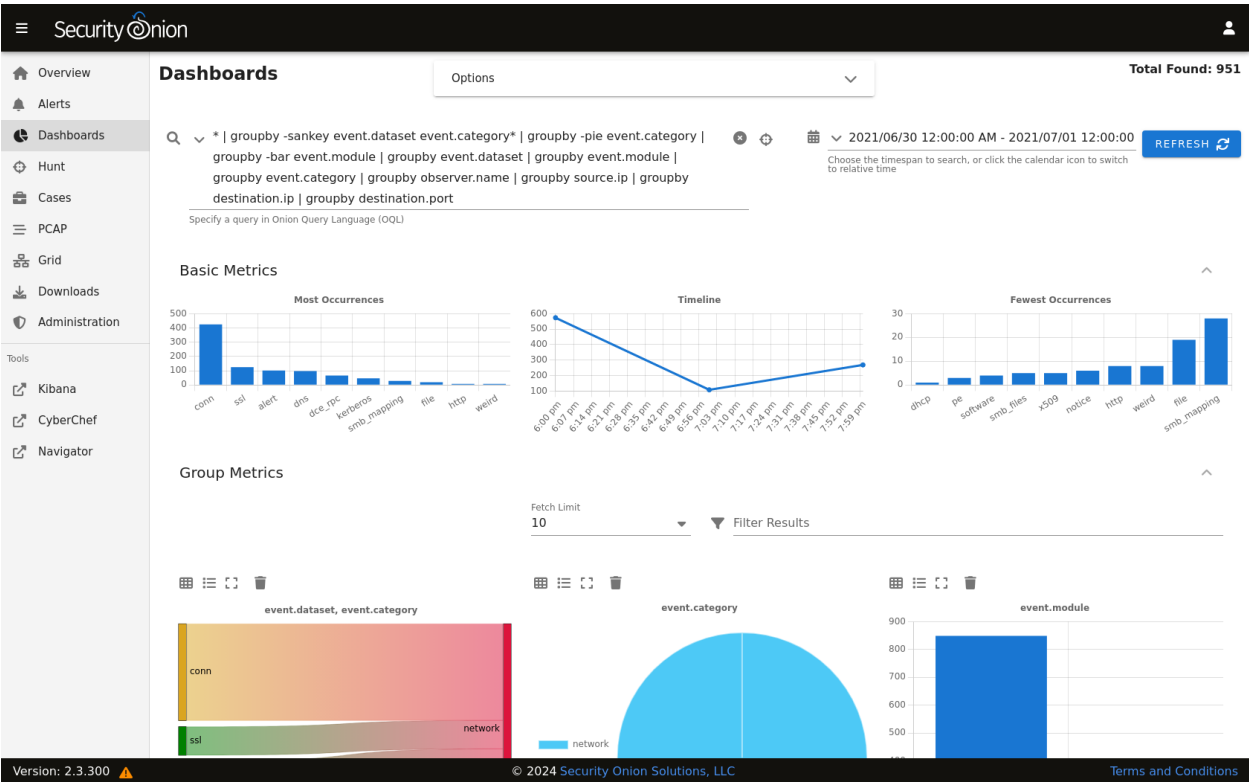
*Security Onion Console (SOC)* is the first thing you see when you log into Security Onion. It includes our *Alerts* interface which allows you to see all of your NIDS alerts from *Suricata* and HIDS alerts from *Wazuh*.

**Security Onion Console (SOC) Alerts Interface**

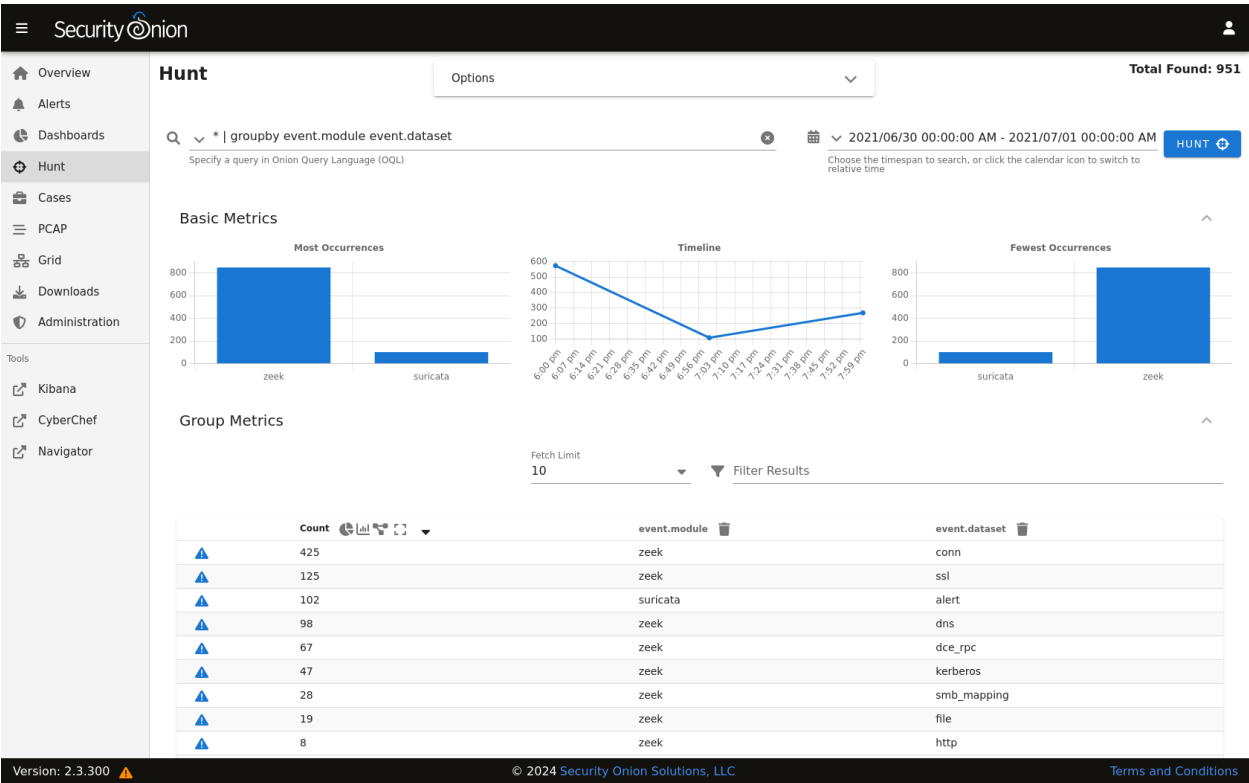
Version: 2.3.300 © 2024 Security Onion Solutions, LLC Terms and Conditions

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/HIDS 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.

SecurityOnion

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Cases

Options

Total Found: 3

+

Open Cases

Specify a query in Onion Query Language (OQL)

NOT so\_case.status:closed NOT so\_case.category:template

Last 12 months

REFRESH

Timestamp	so_case.title	so_case.status	so_case.severity	so_case.assigned	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

Rows per page: 50 1-3 of 3

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

SecurityOnion

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

# 1001

securityonion

176.10.125.8:80

172.16.3.130:49457

Filter Results

HEX

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

Rows per page: 10 1-10 of 24

LOAD MORE

Version: 2.3.300

© 2024 Security Onion Solutions, LLC

Terms and Conditions

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.

Download CyberChef

Last build: 8 months ago

OptionsAbout / Support

Operations

Search...

Favourites

To Base64

From Base64

To Hex

From Hex

To Hexdump

From Hexdump

URL Decode

Regular expression

Entropy

Fork

Magic

Data format

Encryption / Encoding

Public Key

Arithmetic / Logic

Networking

Language

Utils

Date / Time

Extractors

Compression

Hashing

Recipe

From Hexdump

Strip HTTP headers

Strip HTTP headers

Strings

Encoding: Single byte

Minimum length: 9

Match: Alphanumeri...

Display total

Sort

Unique

STEP

BAKE!

Auto Bake

Input

0000 47 45 54 20 2F 31 30 35 2E 64 6C 6C 20 48 54 54 GET /105.dll HTT

0016 50 2F 31 2E 31 0D 0A 43 6F 6E 6E 65 63 74 69 6F P/1.1..Connectio

0032 6E 3A 20 4B 65 65 70 2D 41 6C 69 76 65 0D 0A 55 n: Keep-Alive..U

0048 73 65 72 2D 41 67 65 6E 74 3A 20 63 75 72 6C 2F ser-Agent: curl/

0064 37 2E 37 34 2E 30 0D 0A 48 6F 73 74 3A 20 31 37 7.74.0..Host: 17

0080 36 2E 31 30 2E 31 32 35 2E 38 0D 0A 0D 0A 6.10.125.8...

0000 48 54 54 50 2F 31 2E 31 20 32 30 30 20 4F 4B 0D HTTP/1.1 200 OK.

0016 0A 44 61 74 65 3A 20 57 65 64 2C 20 33 30 20 4A .Date: Wed, 30 J

0032 75 6E 20 32 30 32 31 20 32 30 3A 34 38 3A 33 38 un 2021 20:48:38

0048 20 47 4D 54 0D 0A 53 65 72 76 65 72 3A 20 41 70 GMT..Server: Ap

0064 61 63 68 65 2F 32 2E 34 2E 32 35 20 28 44 65 62 ache/2.4.25 (Deb

0080 69 61 6E 29 0D 0A 4C 61 73 74 2D 4D 6F 64 69 66 ian)..Last-Modif

0096 69 65 64 3A 20 57 65 64 2C 20 33 30 20 4A 75 6E ied: Wed, 30 Jun

0112 20 32 30 32 31 20 31 38 3A 35 37 3A 33 34 20 47 2021 18:57:34 G

0128 4D 54 0D 0A 45 54 61 67 3A 20 22 32 65 30 30 2D MT..ETag: "2e00-

0144 35 63 36 30 30 34 62 61 30 38 63 63 64 22 0D 0A Sc6004ba08ccdf..

0160 41 63 63 65 70 74 2D 52 61 6E 67 65 73 3A 20 62 Accept-Ranges: b

0176 79 74 65 73 0D 0A 43 6F 6E 74 65 6E 74 2D 4C 65 ytes..Content-Le

0192 6E 67 74 68 3A 20 31 31 37 37 36 0D 0A 4B 65 65 ngth: 11776..Kee

0208 70 2D 41 6C 69 76 65 3A 20 74 69 6D 65 6F 75 74 p-Alive: timeout

0224 3D 35 2C 20 6D 61 78 3D 31 30 30 0D 0A 43 6F 6E =5, max=100..Con

0240 6E 65 63 74 69 6F 6E 3A 20 4B 65 65 70 2D 41 6C nection: Keep-Al

Output

!This is a Windows NT windowed dynamic link library

KERNEL32.DLL

VirtualAlloc

TstDll.dll

TstDll.dll

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.

Home
Logged in as analyst
My account
Sign out

DETECTION PLAYBOOKS
Search:
Detection Playbooks

Activity
Playbook
Sigma Editor

### Playbook

Filters
Status
open
Add filter

Options

Apply Clear Save

#	Status	Level	Playbook	Product	Title	Updated
623	Draft	medium	community	windows	Harvesting of Wifi Credentials Using netsh.exe	05/13/2020 02:07 PM
622	Draft	medium	community	windows	Advanced IP Scanner	05/13/2020 02:07 PM
621	Draft	high	imported	windows	Whoami Execution	05/13/2020 02:05 PM
620	Draft	medium	imported	osquery	New Sensitive Shared Resource	05/13/2020 01:30 PM
618	Inactive	medium	community	windows	XSL Script Processing	05/03/2020 10:00 AM
617	Draft	high	community	windows	Wreset UAC Bypass	05/01/2020 08:58 PM
616	Draft	high	community	windows	Microsoft Workflow Compiler	05/01/2020 08:57 PM
615	Draft	critical	community	windows	Wmpirve Spawning Process	05/01/2020 08:57 PM
614	Draft	high	community	windows	WMI Spawning Windows PowerShell	05/01/2020 08:57 PM
613	Draft	high	community	windows	WMI Persistence - Script Event Consumer	05/01/2020 08:57 PM
612	Draft	critical	community	windows	WMI Backdoor Exchange Transport Agent	05/01/2020 08:57 PM
611	Draft	high	community	windows	Windows 10 Scheduled Task SandboxEscaper 0-day	05/01/2020 08:57 PM
610	Draft	high	community	windows	Run Whoami as SYSTEM	05/01/2020 08:57 PM
609	Draft	high	community	windows	Shells Spawned by Web Servers	05/01/2020 08:57 PM
608	Draft	high	community	windows	Webshell Detection With Command Line Keywords	05/01/2020 08:57 PM
607	Draft	medium	community	windows	Java Running with Remote Debugging	05/01/2020 08:57 PM
606	Draft	high	community	windows	Possible Privilege Escalation via Weak Service Permissions	05/01/2020 08:57 PM
605	Draft	high	community	windows	Bypass UAC via WReset.exe	05/01/2020 08:57 PM
604	Draft	high	community	windows	Bypass UAC via Fodhelper.exe	05/01/2020 08:57 PM
603	Draft	high	community	windows	Bypass UAC via CMSTP	05/01/2020 08:57 PM
602	Draft	medium	community	windows	Domain Trust Discovery	05/01/2020 08:57 PM
601	Draft	high	community	windows	Terminal Service Process Spawn	05/01/2020 08:57 PM
600	Draft	high	community	windows	Tasks Folder Evasion	05/01/2020 08:57 PM
599	Draft	medium	community	windows	Tap Installer Execution	05/01/2020 08:57 PM
598	Draft	high	community	windows	System File Execution Location Anomaly	05/01/2020 08:57 PM

Previous 2 3 ... 13 Next
(1-25/310) Per page: 25, 75, 150

Also available in: Atom | CSV | PDF

Custom queries
All Plays
Disabled Plays
Draft Plays
Playbook - Community Sigma
Playbook - Internal

## 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](#).
- 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.

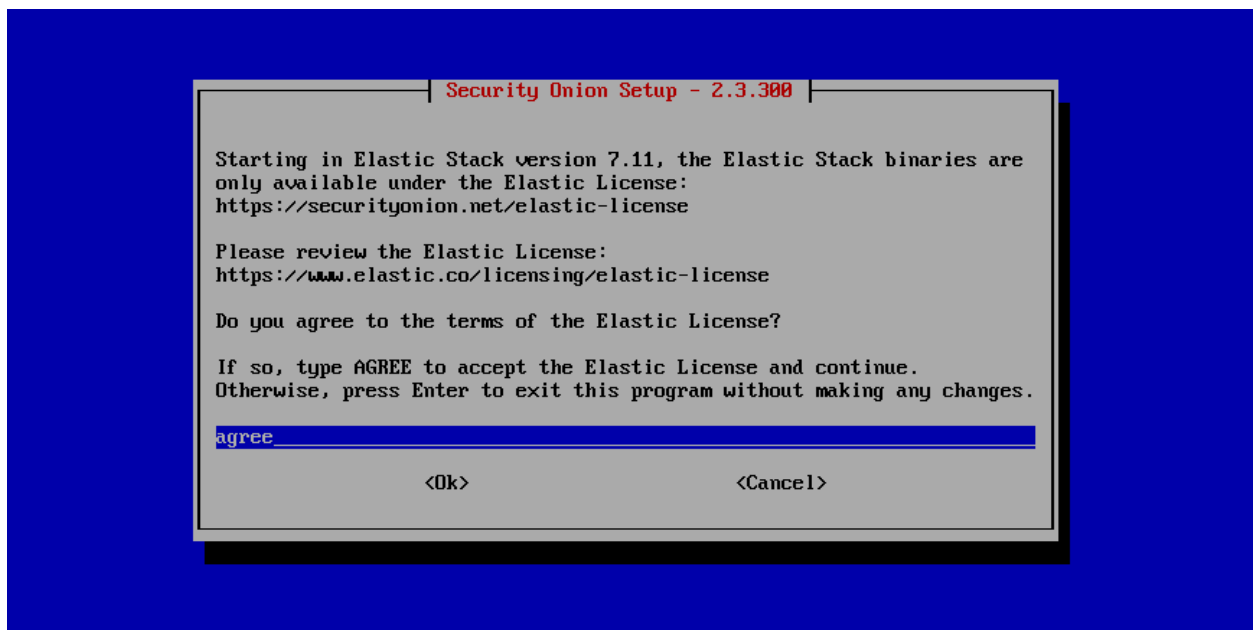
## 2.6 Conclusion

After you install Security Onion, you will have network and endpoint detection, comprehensive metadata, and full packet capture. 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. The vast majority of software included in Security Onion is licensed under OSI-approved open source licenses. However, starting in Elastic 7.11, the Elastic Stack is licensed under the Elastic License. During installation, you will be prompted to accept the Elastic License:



---

**Note:**

You can find the full text of the Elastic License at:  
<https://www.elastic.co/licensing/elastic-license>

For more information about the Elastic license change, please see:  
<https://securityonion.net/elastic-license>

---



## FIRST TIME USERS

If this is your first time using Security Onion 2, then we highly recommend that you start with a simple IMPORT installation using our Security Onion ISO image (see the [Download](#) section). This can be done in a minimal virtual machine (see the [VMware](#) and [VirtualBox](#) sections) with as little as 4GB RAM, 2 CPU cores, and 200GB of storage.

The following screenshots will walk you through:

- installing our Security Onion ISO image
- configuring for IMPORT
- optionally enabling the Analyst environment (see the [Analyst VM](#) section)
- running `so-import-pcap` and importing one or more pcap files

After following the screenshots, you can skip to the [Security Onion Console \(SOC\)](#) section.

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

# Security Onion

Security Onion 2.3.300

```
Install Security Onion 2.3.300
Install Security Onion 2.3.300 Analyst Workstation
Install Security Onion 2.3.300 IDH Node
Install Security Onion 2.3.300 in basic graphics mode
Install Security Onion 2.3.300 Analyst Workstation in basi
Install Security Onion 2.3.300 IDH Node in basic graphics
Test this media & install Security Onion 2.3.300
  Press Tab for full configuration options on menu items.
Troubleshooting >
```



```
#####
##          ** 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

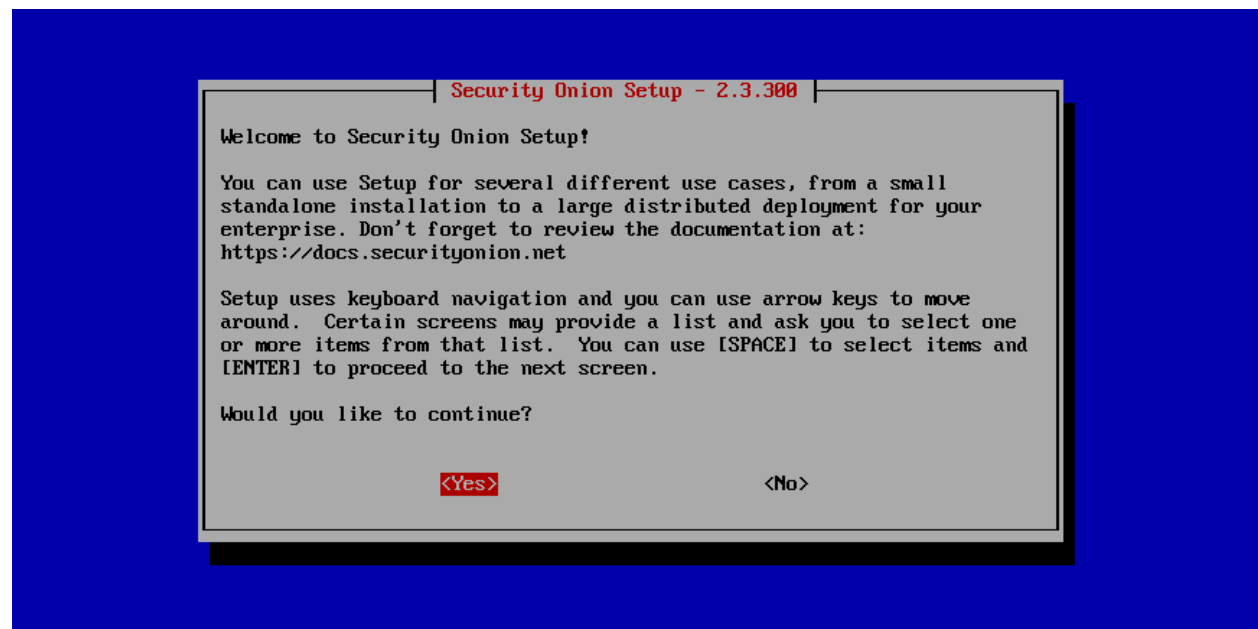
A new administrative user will be created. This user will be used for setting up and administering Security Onion.

Enter an administrative username: doug

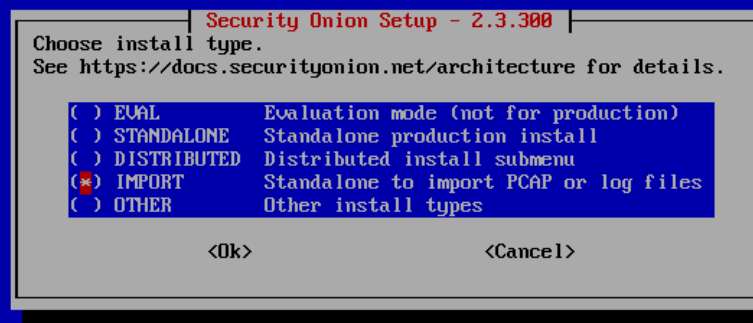

Let's set a password for the doug user:


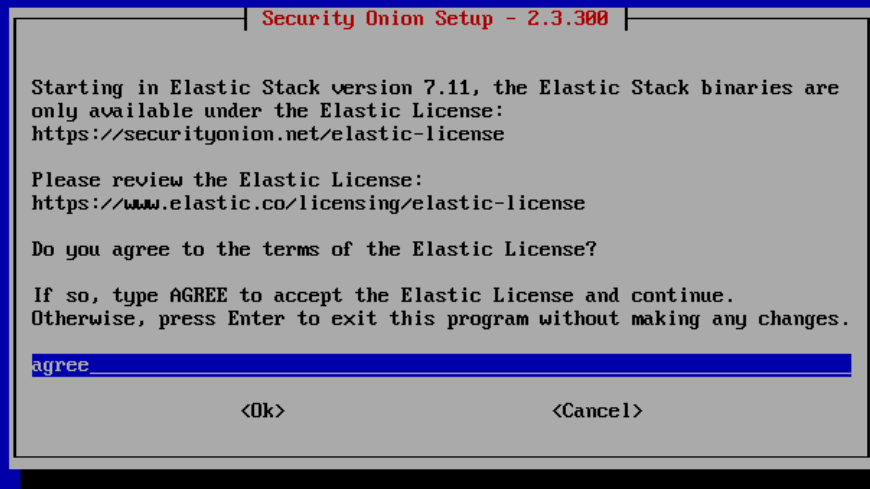
Enter a password:
Re-enter the password:
```



Initial Install Complete. Press [Enter] to reboot!


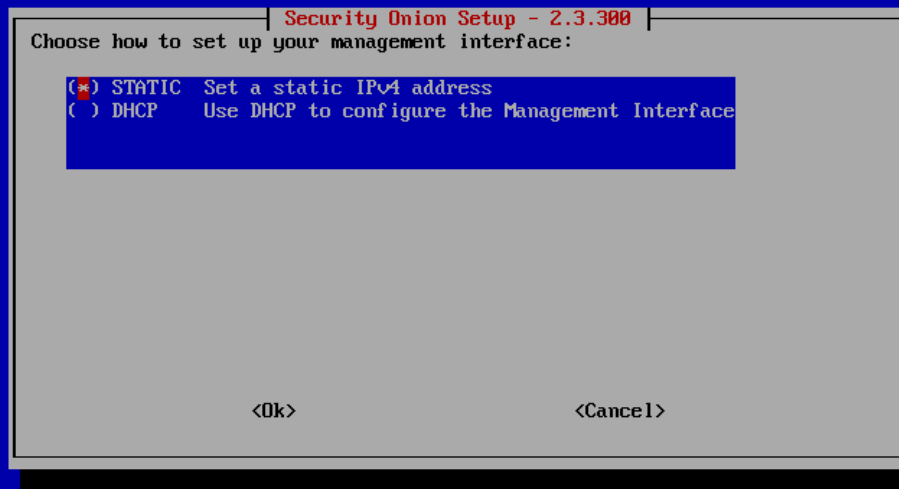























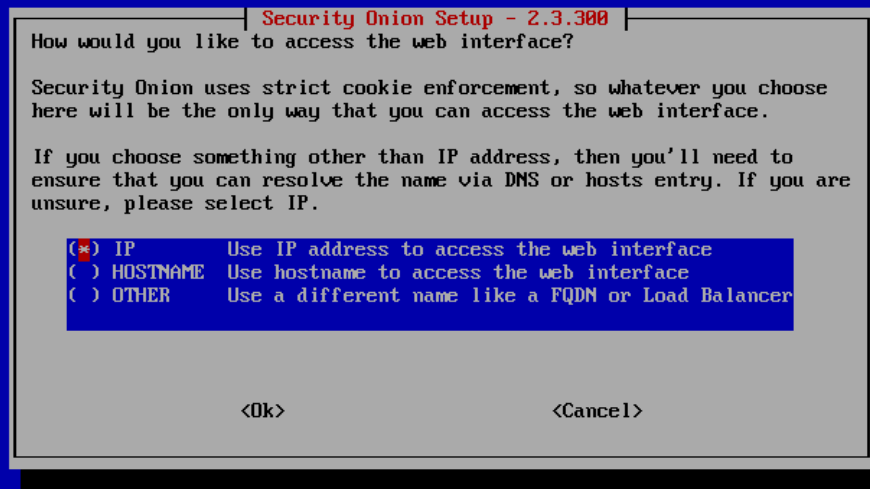
























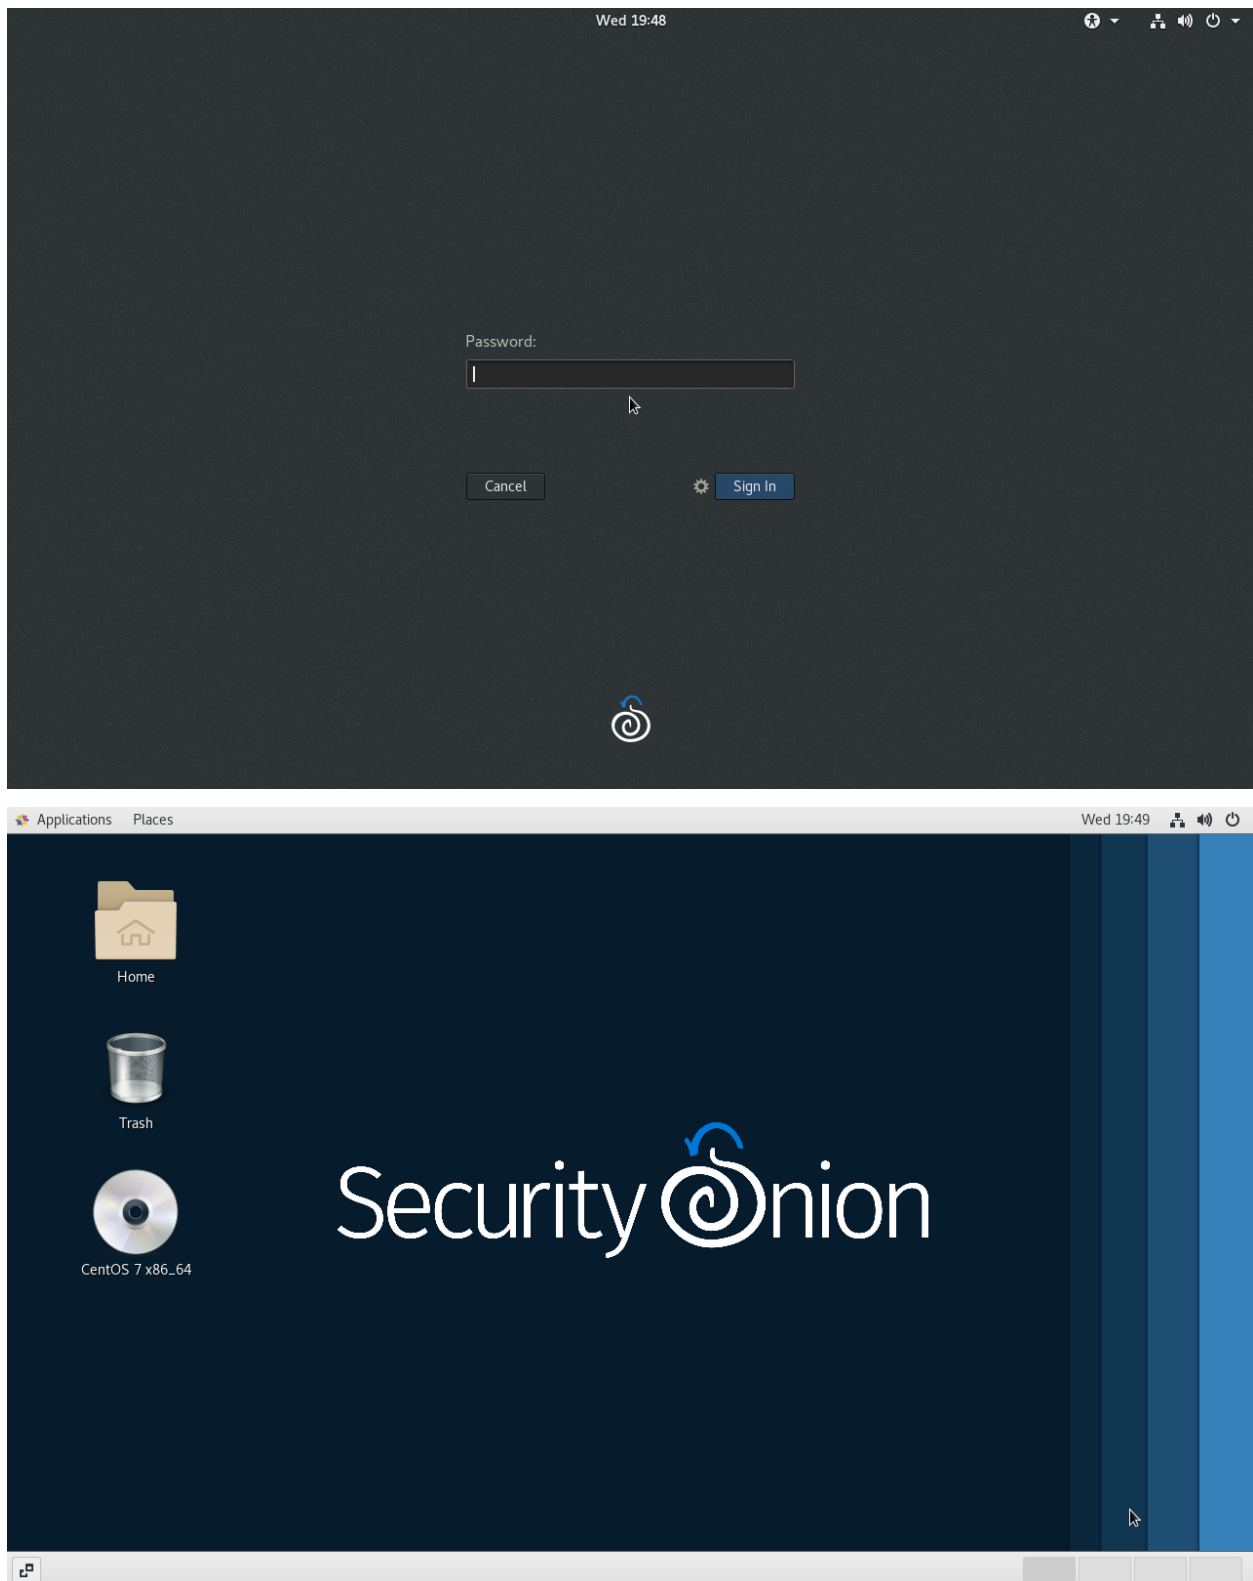
```
[doug@securityonion ~]$ sudo so-analyst-install_
```

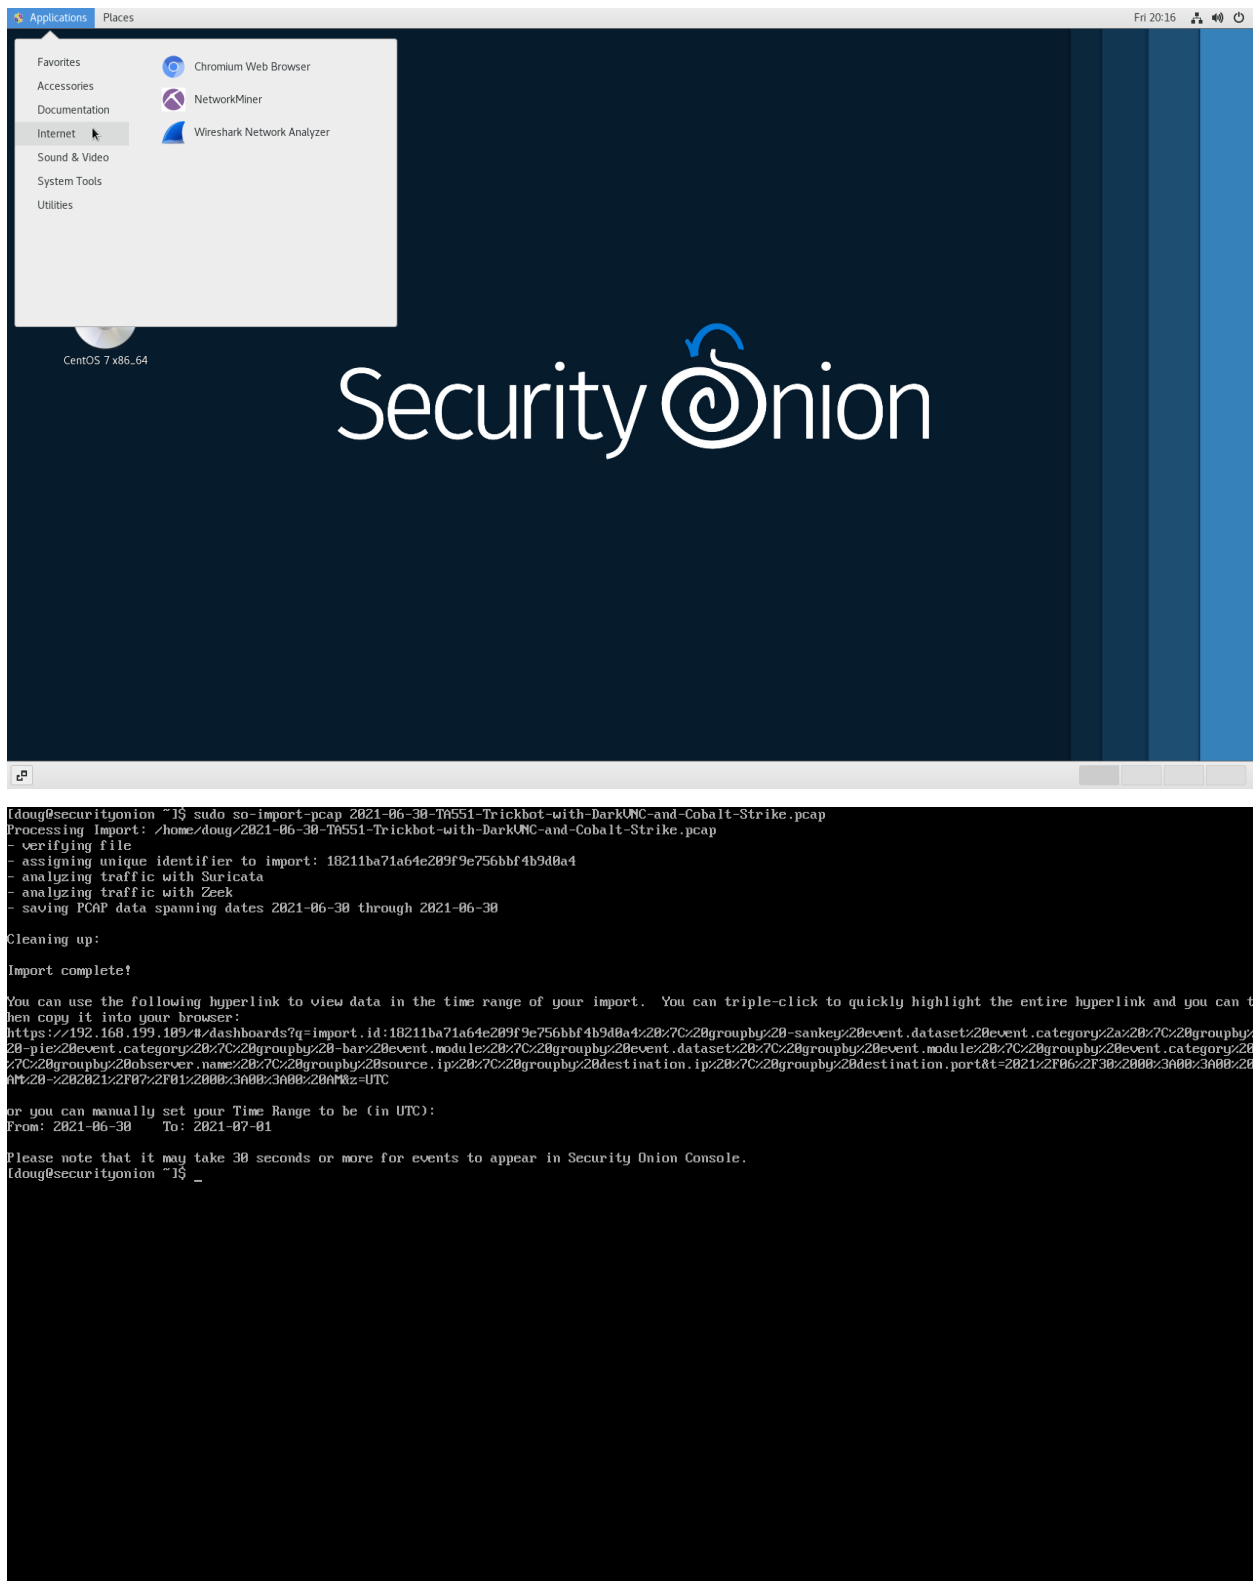
```
Function: file.symlink
  Name: /etc/systemd/system/default.target
  Result: True
  Comment: Created new symlink /etc/systemd/system/default.target -> /lib/systemd/system/graphical.target
  Started: 19:38:13.248554
  Duration: 12.767 ms
  Changes:
    -----
    new:
      -----
      /etc/systemd/system/default.target
    -----
ID: trusted_ca
Function: x509.pem_managed
  Name: /etc/pki/ca-trust/source/anchors/ca.crt
  Result: True
  Comment: File /etc/pki/ca-trust/source/anchors/ca.crt updated
  Started: 19:38:13.268888
  Duration: 13.142 ms
  Changes:
    -----
    diff:
      -----
      New file
    -----
ID: update_ca_certs
Function: cmd.run
  Name: update-ca-trust
  Result: True
  Comment: Command "update-ca-trust" run
  Started: 19:38:13.287288
  Duration: 884.996 ms
  Changes:
    -----
    pid:
      24163
    retcode:
      0
    stderr:
    stdout:

Summary for local
-----
Succeeded: 6 (changed=6)
Failed:    0
-----
Total states run:    6
Total run time: 162.163 s

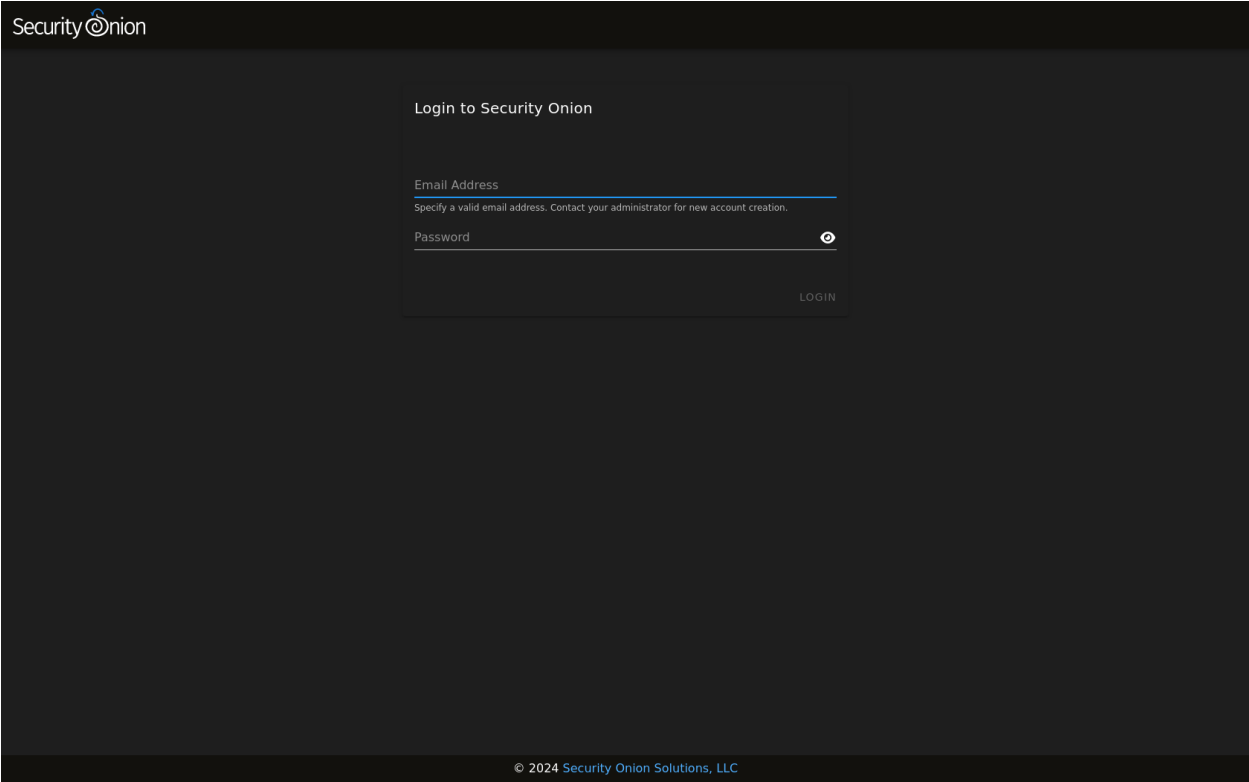
Analyst workstation has been installed!
Press ENTER to reboot or Ctrl-C to cancel.
```











Security Onion

⌵

⌵

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Overview

Getting Started

New to Security Onion 2? Click the menu in the upper-right corner and you'll find links for [Help](#) and a [Cheatsheet](#) 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.

Customize This Space

Make this area your own by customizing the content. The content is stored in the `.motd.md` file, which uses the common Markdown (.md) format. To learn more about the format, please see [markdownguide.org](#).

To customize this content, login to the manager via SSH and execute the following command:

```
sudo cp /opt/so/saltstack/default/salt/soc/files/soc/motd.md /opt/so/saltstack/local/salt/soc/files/soc/
```

Then edit the new file as desired using your favorite text editor.

Finally, restart SOC to make the changes take effect:

```
sudo so-soc-restart
```

Brought to you by:

Security Onion SOLUTIONS

Version: 2.3.300

© 2024 Security Onion Solutions, LLC

[Terms and Conditions](#)

Security Onion

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Alerts

Options

Total Found: 102

Q

Group By Name, Module

+

2021/06/30 00:00:00 AM - 2021/07/01 00:00:00

REFRESH

Fetch Limit

50

Filter Results

	Count	rule.name	event.module	event.severity_label
	59	ET POLICY OpenSSL Demo CA - Internet Wldgits 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

Version: 2.3.300
© 2024 Security Onion Solutions, LLC
Terms and Conditions

Security Onion

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Dashboards

Options

Total Found: 951

Q

\* | groupby -sankey event.dataset event.category\* | groupby -pie event.category | groupby -bar event.module | groupby event.dataset | groupby event.module | groupby event.category | groupby observer.name | groupby source.ip | groupby destination.ip | groupby destination.port

+

2021/06/30 12:00:00 AM - 2021/07/01 12:00:00

REFRESH

Fetch Limit

10

Filter Results

Specify a query in Onion Query Language (OQL)

Basic Metrics

Most Occurrences

Timeline

Fewest Occurrences

Group Metrics

event.dataset, event.category

event.category

event.module

Version: 2.3.300
© 2024 Security Onion Solutions, LLC
Terms and Conditions

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Hunt

Options

Total Found: 951

Q \* | groupby event.module event.dataset

Specify a query in Onion Query Language (OQL)

2021/06/30 00:00:00 AM - 2021/07/01 00:00:00 AM

HUNT

Choose the timespan to search, or click the calendar icon to switch to relative time

Basic Metrics

Most Occurrences

Timeline

Fewest Occurrences

Group Metrics

Fetch Limit 10

Filter Results

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

Version: 2.3.300

© 2024 Security Onion Solutions, LLC

Terms and Conditions

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

# 1001

securityonion

176.10.125.8:80

172.16.3.130:49457

Filter Results

HEX

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

Rows per page: 10 1-10 of 24

LOAD MORE

Version: 2.3.300

© 2024 Security Onion Solutions, LLC

Terms and Conditions

[illegible]

☰

SecurityOnion

👤

🏠 Overview

🔔 Alerts

📊 Dashboards

🔍 Hunt

📁 Cases

📄 PCAP

🗃️ Grid

📄 Downloads

🛡️ Administration

Tools

🔗 Kibana

🔗 CyberChef

🔗 Navigator

PCAP

+

ID ▲	Owner	Date Queued	Date Completed	Sensor ID	Status	Actions ▲
<a href="#">1001</a>	doug@example.com	2024-03-27 19:51:28.875 +00:00	2024-03-27 19:51:29.771 +00:00	securityonion	Completed	<div>👤 ⚙️</div>

Rows per page: 10 1-1 of 1 < >

☰

SecurityOnion

👤

🏠 Overview

🔔 Alerts

📊 Dashboards

🔍 Hunt

📁 Cases

📄 PCAP

🗃️ Grid

📄 Downloads

🛡️ Administration

Tools

🔗 Kibana

🔗 CyberChef

🔗 Navigator

Grid

▼ Filter Results

ID ▲	Role(s)	Address	Description	Version	Model	Date Updated	Earliest PCAP	Uptime ▲	Status
>	securityonion	Import	192.168.199.109						
				2.3.300	N/A	2024-03-27 19:52:23.855 +00:00	2024-03-27 19:52:23.855 +00:00	6 minutes	OK

Rows per page: 10 1-1 of 1 < >

SecurityOnion

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Downloads

When installing packages such as osquery or beats onto remote systems be sure to run `so-allow` on the Security Onion Manager node to allow network access through the firewall.

Elasticsearch Utilities (8.10.4)

Winlogbeat (Windows)

Filebeat DEB (Linux [amd64]: Debian, Ubuntu)

Filebeat RPM (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)

Metricbeat DEB (Linux [amd64]: Debian, Ubuntu)

Metricbeat RPM (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)

Auditbeat DEB (Linux [amd64]: Debian, Ubuntu)

Auditbeat RPM (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)

Wazuh Agents (3.13.1-1)

MSI (Windows)

DEB (Linux [amd64]: Debian, Ubuntu)

RPM (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)

PKG (MacOS)

osquery Packages and Configs

MSI (Windows)

DEB (Linux [amd64]: Debian, Ubuntu)

RPM (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)

PKG (MacOS)

RPM & DEB Config Flag File

MSI Config Flag File

These packages and configs are osquery files, customized for this specific Fleet install and will only be generated if Fleet has been installed. Due to macOS packaging constraints, the macOS PKG has not been customized for this Fleet install - osquery/Launcher will need to be configured post-deployment. These files are not signed. Signed, non-customized osquery packages can be obtained directly from [osquery.io](#). For further Fleet & osquery information, view our online help.

Version: 2.3.300

© 2024 Security Onion Solutions, LLC

Terms and Conditions

SecurityOnion

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Users

Email Address	First Name	Last Name	Note	Role(s)	Status	Actions
doug@example.com				superuser		

Rows per page: 10 1-1 of 1

Version: 2.3.300

© 2024 Security Onion Solutions, LLC

Terms and Conditions

42

Chapter 4. First Time Users

```

top - 19:34:53 up 9 min, 1 user, load average: 0.11, 0.33, 0.27
Tasks: 181 total, 1 running, 180 sleeping, 0 stopped, 0 zombie
%Cpu(s): 6.5 us, 2.5 sy, 0.0 ni, 90.9 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3879960 total, 483284 free, 2619564 used, 777112 buff/cache
KiB Swap: 8388684 total, 8231164 free, 157440 used, 983208 avail Mem

```

PID	USER	PR	NI	UIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
3567	root	20	0	1232940	70888	2920	S	10.3	1.8	0:06.09	salt-minion
8319	elastic+	20	0	3558856	999404	12716	S	2.7	25.8	1:02.52	java
9	root	20	0	0	0	0	S	0.3	0.0	0:00.64	rcu_sched
1551	root	20	0	1209280	28532	2600	S	0.3	0.7	0:00.57	salt-master
3490	root	20	0	1324872	67040	3460	S	0.3	1.7	0:01.17	salt-minion
6032	socore	20	0	748200	25196	2600	S	0.3	0.6	0:00.34	sensoroni
8587	kibana	20	0	11.2g	423072	12344	S	0.3	10.9	0:19.99	node
20130	doug	20	0	162120	2200	1520	R	0.3	0.1	0:00.01	top
1	root	20	0	193752	5476	2764	S	0.0	0.1	0:01.05	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd
3	root	20	0	0	0	0	S	0.0	0.0	0:00.13	kuorker/0:0
4	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kuorker/0:0H
6	root	20	0	0	0	0	S	0.0	0.0	0:00.06	ksoftirqd/0
7	root	rt	0	0	0	0	S	0.0	0.0	0:00.01	migration/0
8	root	20	0	0	0	0	S	0.0	0.0	0:00.00	rcu_bh
10	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	lru-add-drain
11	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	watchdog/0
12	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	watchdog/1
13	root	rt	0	0	0	0	S	0.0	0.0	0:00.01	migration/1
14	root	20	0	0	0	0	S	0.0	0.0	0:00.05	ksoftirqd/1
16	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kuorker/1:0H
18	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kdevtmpfs
19	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	netns
20	root	20	0	0	0	0	S	0.0	0.0	0:00.00	khungtaskd
21	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	writeback
22	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kintegrityd
23	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	bioaset
24	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	bioaset
25	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	bioaset
26	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kblockd
27	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	md
28	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	edac-poller
29	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	watchdogd
30	root	20	0	0	0	0	S	0.0	0.0	0:00.01	kuorker/0:1
31	root	20	0	0	0	0	S	0.0	0.0	0:00.03	kuorker/1:1
33	root	20	0	0	0	0	S	0.0	0.0	0:00.66	kswapd0
34	root	25	5	0	0	0	S	0.0	0.0	0:00.00	ksmd
35	root	39	19	0	0	0	S	0.0	0.0	0:00.00	khugepaged
36	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	crypto
44	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kthrotld
45	root	20	0	0	0	0	S	0.0	0.0	0:00.12	kuorker/u4:1
46	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kmpath_rdacd
47	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kaluad





## 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.

**Which ISO image should I download?**

You can download our Security Onion ISO image or a standard 64-bit CentOS 7 or Ubuntu 20.04 ISO image. We recommend our Security Onion ISO image for most use cases, but you should review the [Partitioning](#), [Release Notes](#), and [Download](#) sections for more information.

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

See the [VMware](#) and [VirtualBox](#) sections.

**How do I deploy Security Onion in the cloud?**

See the [AWS Cloud AMI](#) and [Azure Cloud Image](#) sections.

**What if I have trouble booting the ISO image?**

Check out the [Bootling 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?**

The [Installation](#) section has steps for our Security Onion ISO image and for standard CentOS 7 and Ubuntu 20.04 ISO images.

**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 our Security Onion ISO image for the quickest and easiest installation experience (see 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 OSs 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 [Grafana](#) 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 files using *so-import-pcap* or evtx files using *so-import-evtx*. You can then view those logs in *Security Onion Console (SOC)*.

### 5.2.2 Evaluation

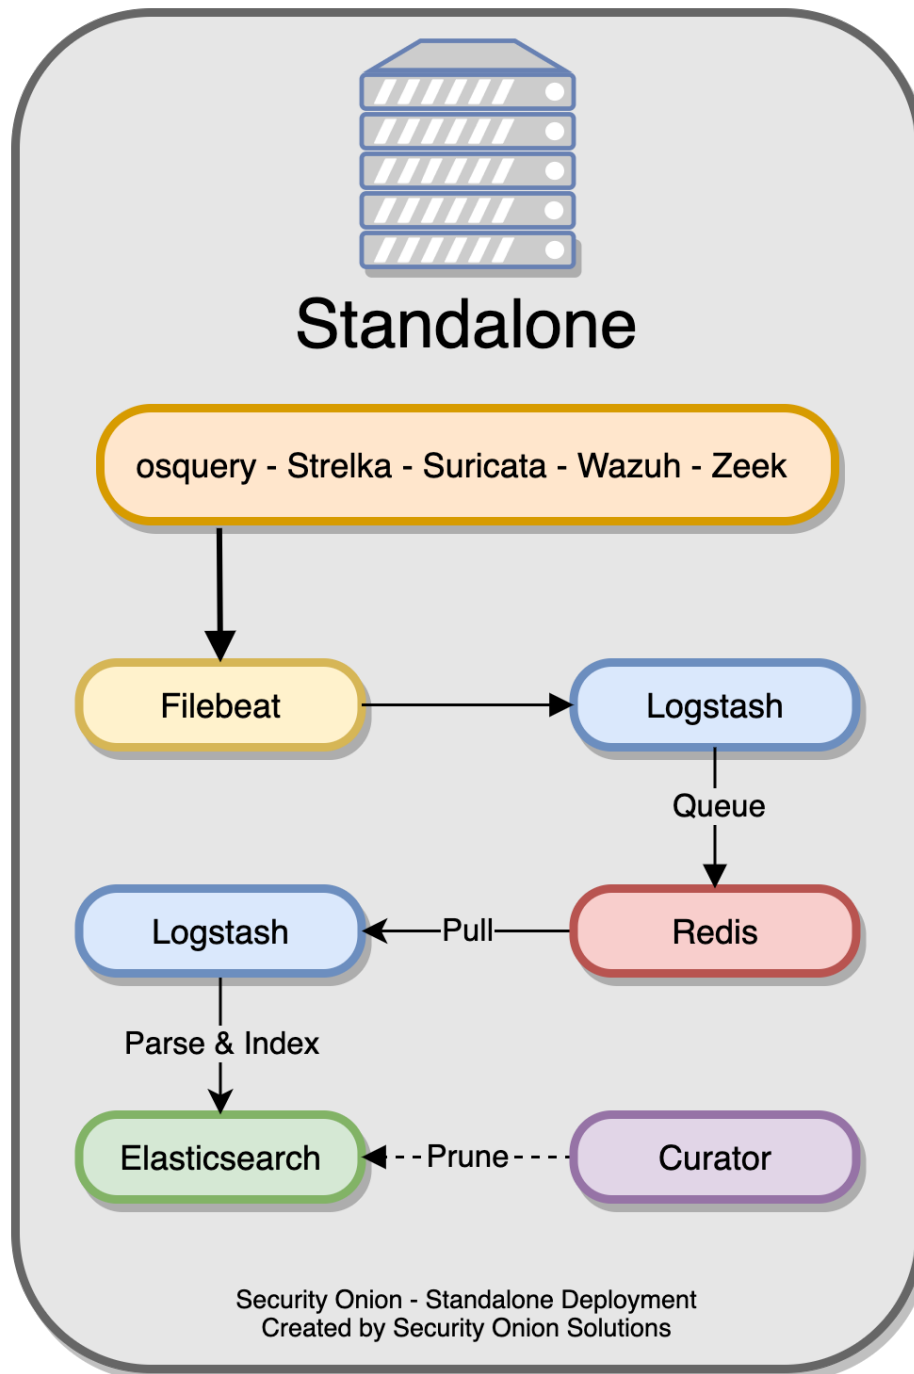
The next architecture is **Evaluation**. It's a little more complicated than **Import** because it has a network interface dedicated to sniffing live traffic from a TAP or span port. Processes monitor the traffic on that sniffing interface and generate logs. *Filebeat* collects those logs and sends them 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.



### 5.2.3 Standalone

Standalone is similar to Evaluation in that all components run on one box. However, instead of *Filebeat* 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

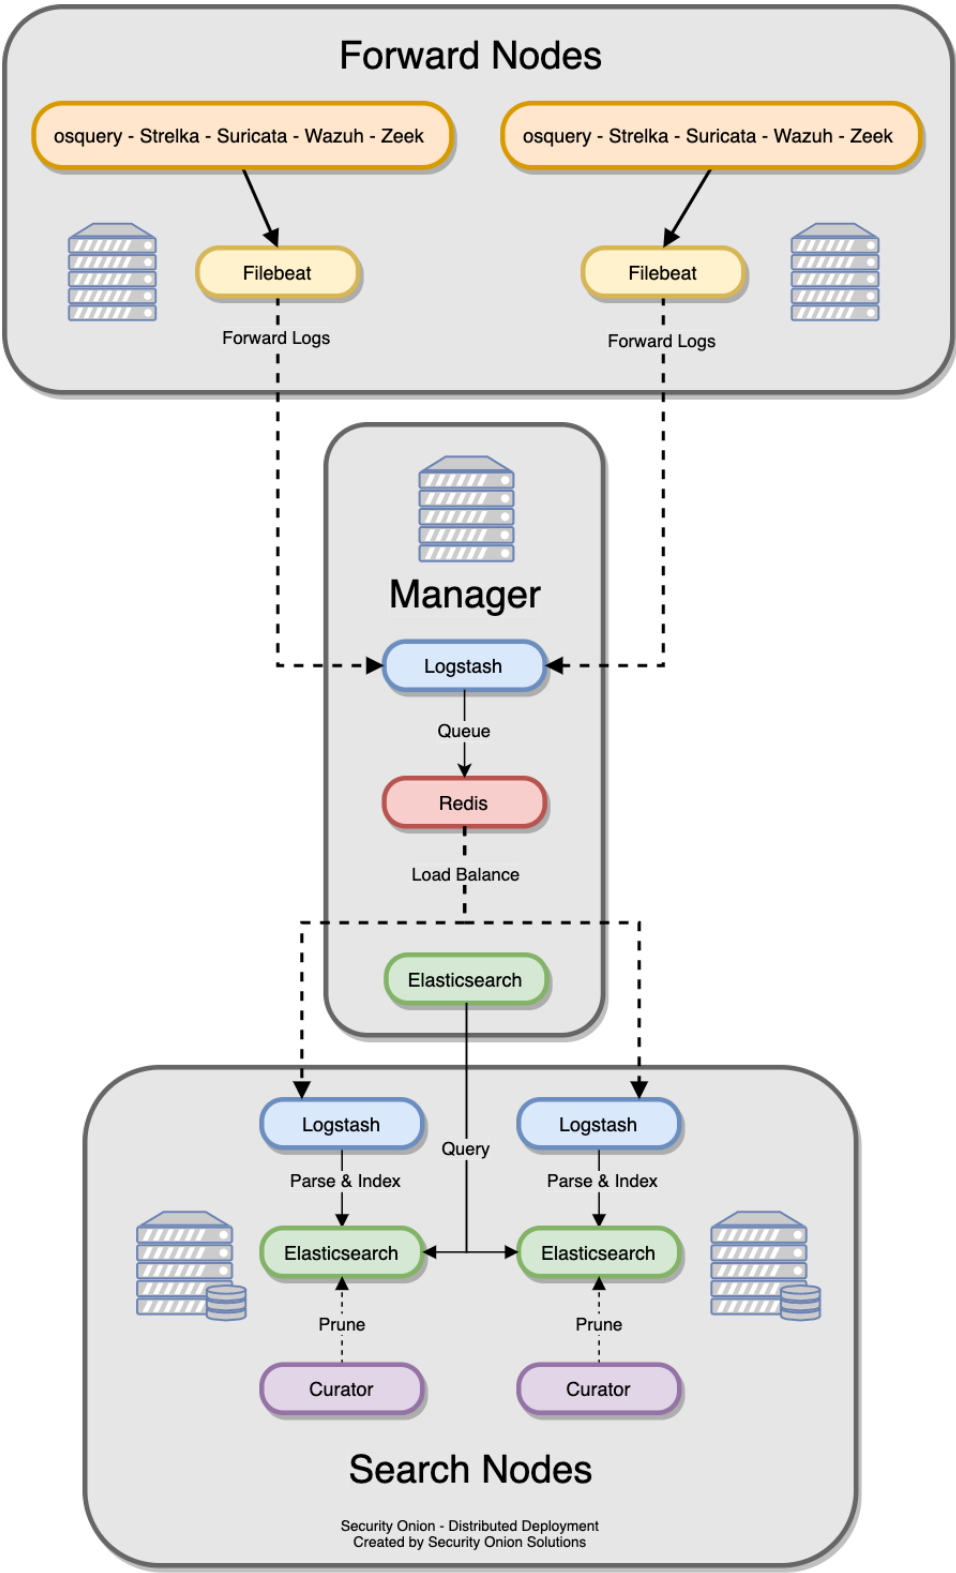
A standard distributed deployment includes a **manager node**, one or more **forward nodes** running network sensor components, and one or more **search nodes** running Elastic search components. This architecture may cost more upfront, but it provides for greater scalability and performance, as you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


---



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



**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.



## 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 an *Analyst VM*) 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*
- *Curator*
- *ElastAlert*
- *Redis*
- *Wazuh*

### 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.

Security Onion can add search nodes using *cross cluster search* where each search node is independent or using traditional Elastic clustering where all search nodes join to form a single cluster.

Search Nodes run the following components:

- *Elasticsearch*
- *Logstash*
- *Curator*
- *Wazuh*

### 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*
- *Curator*
- *ElastAlert*

- *Redis*
- *Wazuh*

### Forward Node

A **forward node** is a sensor that forwards all logs via *Filebeat* 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).

Forward nodes run the following components:

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

### Heavy Node

**Warning:** Heavy nodes are NOT recommended for most users.

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*
- *Curator*
- *Zeek*
- *Suricata*
- *Stenographer*
- *Wazuh*

### Fleet Standalone Node

A *FleetDM* Standalone Node is ideal when there are a large amount of osquery endpoints deployed. It reduces the amount of overhead on the manager node by transferring the workload associated with managing osquery endpoints to a dedicated system. It is also useful for off-network osquery endpoints that do not have remote access to the Manager node as it can be deployed to the DMZ and TCP/8090 made accessible to your off-network osquery endpoints.

If the Manager Node was originally setup with *FleetDM*, your grid will automatically switch over to using the *FleetDM* Standalone Node instead as a grid can only have one *FleetDM* instance active at a time.

*FleetDM* Standalone Nodes run the following components:

- *FleetDM*

## Receiver Node

Security Onion includes a Receiver Node option. The 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, *Filebeat* 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.

## Intrusion Detection Honeypot (IDH) Node

Security Onion includes an *Intrusion Detection Honeypot* 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.

# 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

If you just want to import a pcap using *so-import-pcap*, then you can configure Security Onion 2 as an Import Node with the following minimum specs:

- 4GB RAM
- 2 CPU cores
- 200GB storage

For all other configurations, the minimum specs for running Security Onion 2 are:

- 12GB RAM
- 4 CPU cores
- 200GB storage

---

**Note:** These minimum specs are for EVAL mode with minimal services running. These 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.

---

**Warning:** If using the beta tool logscan, a CPU/vCPU with AVX support will be required, with AVX2 support recommended. Most CPU's produced since 2011 will have AVX support, but be sure to check the specs of your particular processor. This will be more important to virtualized environments, where the default vCPU type will typically not support either instruction set.

### 5.3.3 Production Deployments

Security Onion 2 is a new platform with more features than previous versions of Security Onion. These additional features result in higher hardware requirements. For best results, we recommend purchasing new hardware to meet the new requirements.

---

**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 would most likely need swap space to avoid issues.

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 ES 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, 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 ES 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 ES 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 ES 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.

- CPU: Used to parse incoming events, index incoming events, 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 ES 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, 2x CPU, 1x 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 Partitioning

Now that you understand *Hardware Requirements*, we should next discuss disk partitioning. If you're installing Security Onion for a production deployment, you'll want to pay close attention to partitioning to make sure you don't fill up a partition at some point.

### 5.4.1 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 `/`.

### 5.4.2 ISO

If you use our Security Onion ISO image, it will automatically partition your disk for you. If you instead use CentOS 7 or Ubuntu 20.04, you will most likely need to manually modify their default partition layout.

### 5.4.3 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. Our Security Onion ISO image uses LVM by default.

### 5.4.4 `/boot`

You probably want a dedicated `/boot` partition of at least 500MB at the beginning of the drive.

### 5.4.5 `/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.

### 5.4.6 `/`

`/` (the root partition) currently contains `/var/lib/docker/` (more on that below) and thus you'll want at least 100GB.

### 5.4.7 Docker

Docker images are currently written to `/var/lib/docker/`. The current set of Docker images uses 27GB 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.

### 5.4.8 Other

If you install using a standard CentOS 7 or Ubuntu 20.04 ISO, then those installers 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.

### 5.4.9 Example

Here's an example of how our current Security Onion ISO image partitions a 1TB disk:

- 500MB `/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

## 5.5 Download

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

You can either download our Security Onion ISO image (based on CentOS 7) **or** download a standard 64-bit CentOS 7 or Ubuntu 20.04 ISO image and then add our Security Onion components. **Please keep in mind that we only support CentOS 7 and Ubuntu 20.04.**

---

**Tip:** For most use cases, we recommend using our Security Onion ISO image as it's the quickest and easiest method.

---

**Warning:** ALWAYS verify the checksum of ANY downloaded ISO image! Regardless of whether you're downloading our Security Onion ISO image or a standard CentOS or Ubuntu ISO image, you should ALWAYS verify the downloaded ISO image to ensure it hasn't been tampered with or corrupted during download.

- If downloading our Security Onion 2.3 ISO image, you can find the download link and verification instructions here: [https://github.com/Security-Onion-Solutions/securityonion/blob/master/VERIFY\\_ISO.md](https://github.com/Security-Onion-Solutions/securityonion/blob/master/VERIFY_ISO.md)
- If downloading an Ubuntu or CentOS ISO image, please verify that ISO image using whatever instructions they provide.

**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\strelka\yara\`. 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 one of the ISO images above, 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/>.

---

## 5.6 VMware

### 5.6.1 Overview

In this section, we'll cover creating a virtual machine (VM) for our Security Onion 2 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.

---

### 5.6.2 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>.

### 5.6.3 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 Security Onion 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, CentOS 7 64-Bit** 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.

## 5.6.4 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 Security Onion 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**, click **CentOS 7 64-bit**, 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.

## 5.6.5 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 CentOS:

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

## 5.7 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>.

### 5.7.1 Creating VM

First, launch VirtualBox and click the "New" button. Provide a name for the virtual machine ("Security Onion" for example) and specify the type ("Linux") and version (this could be CentOS/RedHat or Ubuntu depending on which version you're installing), then click "Continue." We'll next define how much memory we want to make available to our virtual machine based on the [Hardware Requirements](#) section.

Next, we'll create a virtual hard drive. Specify "Create a virtual hard drive now" then click "Create" to choose the hard drive file type "VDI (VirtualBox Disk Image)" and "Continue." For storage, we have the options of "Dynamically allocated" or "Fixed size." For a client virtual machine, "Dynamically allocated" is the best choice as it will grow the hard disk up to whatever we define as the maximum size on an as needed basis until full, at which point Security Onion's disk cleanup routines will work to keep disk space available. If you happen to be running a dedicated sensor in a virtual machine, I would suggest using "Fixed size," which will allocate all of the disk space you define up front and save you some disk performance early on. Once you've settled on the storage allocation, click "Continue" and provide a name from your hard disk image file and specify the location where you want the disk file to be created if other than the default location. For disk size, you'll want at least 200GB so you have enough capacity for retrieving/testing packet captures and downloading system updates. Click "Create" and your Security Onion VM will be created.

At this point, you can click "Settings" for your new virtual machine so we can get it configured. Mount the Security Onion ISO file so our VM can boot from it to install Linux. Click the "Storage" icon, then under "Controller: IDE" select the "Empty" CD icon. To the right, you'll see "CD/DVD Drive" with "IDE Secondary" specified with another CD icon. Click the icon, then select "Choose a virtual CD/DVD disk file" and browse to where you downloaded the Security Onion ISO file, select it then choose "Open." Next click "Network" then "Adapter 2." You'll need to click the checkbox to enable it then attach it to "Internal Network." Under the "Advanced" options, set "Promiscuous Mode" to "Allow All." Click "Ok" and we are ready to install the operating system.

Hit the "Start" button with your new virtual machine selected and after a few seconds the boot menu will load.

Follow the installation steps for your desired installation type in the [Installation](#) section.

---

**Tip:** You'll notice two icons on the top right in VirtualBox Manager when you select your virtual machine: Details and Snapshots. Click "Snapshots" then click the camera icon and give your snapshot a name and description. Once we have a snapshot, we'll be able to make changes to the system and revert those changes back to the state we are preserving.

---

## 5.7.2 Guest Additions

If you want to install VirtualBox Guest Additions, please see:

<https://wiki.centos.org/HowTos/Virtualization/VirtualBox/CentOSguest>

## 5.8 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>.

### 5.8.1 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.

### 5.8.2 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)`.

### 5.8.3 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 (not the Security Onion VM). One option is to enable passthrough and pass the sniffing NIC 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\(e\)\\_Passthrough](https://pve.proxmox.com/wiki/PCI(e)_Passthrough)

For other options, please see <https://github.com/Security-Onion-Solutions/securityonion/discussions/8245>.

## 5.9 Booting Issues

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

- Verify the downloaded 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 64-bit ISO image for CentOS 7 or Ubuntu 20.04. If they don't run, then you should double-check your 64-bit compatibility.

---

**Tip:** If all else fails but standard 64-bit CentOS 7 or Ubuntu 20.04 installs normally, then you can always install our components on top of them as described on the [Installation](#) page.

---

## 5.10 Airgap

Security Onion is committed to allowing users to run a full install on networks that do not have Internet access. You will need to use our Security Onion ISO image as it includes everything you need to run without Internet access. Setup will ask if you want to configure the installation for airgap and will then make the appropriate modifications to make this work properly. Please note that the airgap option is intended to be consistent across your deployment, so if you are on an airgap network you should choose the airgap option when installing the manager and all nodes (recent versions enforce this automatically).

### 5.10.1 Key Differences

There are a few differences between an airgap install and a normal install with Internet access:

- All CentOS repos are removed and replaced with a new repo that runs on the manager.
- During the install, all of the necessary RPMs are copied from the ISO to a new repo located in `/nsm/repo/`.
- All devices in the grid will now use this repo for updates to packages.
- The latest Emerging Threats (ET) Open rules are copied to `/nsm/repo/rules/` so that the manager can access them and users can use the standard Security Onion process for managing NIDS rules.
- Yara rules are copied to `/nsm/repo/rules/strelka/` for *Strelka* file analysis.
- 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.

### 5.10.2 Rule Updates

The Security Onion 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/repo/rules/emerging-all.rules` using a command like:

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

## 5.11 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 Amazon AWS using our AMI, you can skip to the [AWS Cloud AMI](#) section. If you want to deploy in Azure using our image, you can skip to the [Azure Cloud Image](#) section.

---

Having downloaded your desired ISO according to the [Download](#) section, it's now time to install! There are separate sections below to walk you through installing using our Security Onion ISO image (based on CentOS 7) **or** installing standard CentOS 7 or Ubuntu 20.04 and then installing our components on top.

**Warning:** For most use cases, we recommend using our Security Onion ISO image as it's the quickest, easiest, and most consistent method. If you're not going to use our Security Onion ISO image and you're building a distributed deployment, then we recommend keeping the base OS consistent across all nodes in the deployment.

### 5.11.1 Installation using Security Onion ISO Image

If you want to install Security Onion using our ISO image:

1. Review the [Hardware Requirements](#) and [Release Notes](#) sections.
2. [Download and verify our Security Onion ISO image](#).
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.

### 5.11.2 Installation on Ubuntu or CentOS

If you want to install Security Onion on CentOS 7 or Ubuntu 20.04 (**not** using our Security Onion ISO image), follow these steps:

1. Review the [Hardware Requirements](#) and [Release Notes](#) sections.
2. Download the ISO image for your preferred flavor of 64-bit CentOS 7 or Ubuntu 20.04. 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.
4. Reboot into your new installation.
5. Login using the username and password you specified during installation.
6. Install prerequisites. If you're using CentOS 7:

```
sudo yum -y install git
```

If you're using Ubuntu 20.04:

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



7. Download our repo and start the Setup process:

```
git clone https://github.com/Security-Onion-Solutions/securityonion
cd securityonion
sudo bash so-setup-network
```

8. Proceed to the *Configuration* section.

9. NOTE: If any interfaces intended to be used for monitoring were automatically configured via DHCP during Ubuntu installation, setup will ask you to remove them from other network management tools. The following steps will be required to ensure the devices are managed by `nmcli`:

- Remove monitor interface declarations from `/etc/netplan/00-installer-config.yaml` and then run:

```
sudo netplan apply
sudo touch /etc/NetworkManager/conf.d/10-globally-managed-devices.conf
sudo service network-manager restart
```

- Re-run setup.

## 5.12 AWS Cloud AMI

If you would like to deploy Security Onion 2.3 in AWS, we have an AMI that is already built for you. As this is an older version of Security Onion, the launch of the VM is a little different from installing the latest version on the marketplace. The details are described below.

**Warning:** Existing 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.

**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](https://docs.aws.amazon.com/directoryservice/latest/admin-guide/gsg_create_vpc.html).

### 5.12.1 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 losing *Elasticsearch* data. However, for temporary evaluation installations, where there is little concern for data loss, ephemeral instance storage should 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: 200GB EBS (Optimized) gp3

Evaluation

- Quantity: 1
- Type: t3a.2xlarge
- Storage: 100GB 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. Note that when using virtual machines with the minimum RAM requirements you may need to enable memory swapping.

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.xlarge
- Storage: 500GB EBS (Optimized) gp3

### 5.12.2 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.
- Select: **Create**

#### 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**

## 5.12.3 Create Security Onion Instances

### Instance Creation

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

- From the EC2 dashboard select: **Launch Instance**
- Browse for more AMIs. Change to the **Community AMIs** tab and search for **security onion official**. Choose the 2.3.260 version.
- 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 drop-down 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 drop-down 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 2 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`. 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
```

### 5.12.4 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. For distributed manager nodes using ephemeral storage, if you would like to use traditional *Elasticsearch* clustering, select Advanced and answer Yes. Continue instructions below for applicable nodes.

AWS provides a built-in NTP server at IP 169.254.169.123. This can be used when prompted for an NTP host.

---

**Note:** As Security Onion 2.3 approaches EOL, new cloud images are no longer being released. However, maintenance releases are available for cloud installation. After the manager completes setup it is important to run `sudo soup` to apply all newer maintenance releases and patches. This only is necessary on the manager node.

---

#### All Distributed Manager Nodes

For distributed manager nodes, if connecting sensors through the VPN instance then follow the instructions below to allow remote sensor connections:

Run `so-firewall includehost minion <inside interface of your VPN concentrator>`. Ex:

```
so-firewall includehost minion 10.99.1.10
```

Run `so-firewall includehost sensor <inside interface of your VPN concentrator>`. Ex:

```
so-firewall --apply includehost sensor 10.99.1.10
```

At this time your Manager is ready for remote minions to start connecting.

#### Distributed Manager Nodes using Traditional Elasticsearch Clustering

For distributed manager nodes using ephemeral storage that chose to use traditional *Elasticsearch* clustering, make the following changes in `/opt/so/saltstack/local/pillar/global.sls`:

```
replicas: 1
```

Then, restart *Logstash*:

```
sudo so-logstash-restart
```

Next, fix *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 elastalet*/_settings -X PUT -d '{"index" : { "number_of_replicas" : 1 } }'
```

### 5.12.5 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.

### 5.12.6 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.

### 5.12.7 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 address of the manager inside AWS when prompted for the manager IP.

The AWS internal VPN endpoint IP will need to be added to the `minion` and `sensor` arrays in `/opt/so/saltstack/local/salt/firewall/hostgroups.local.yaml` on the manager. This will open up the required ports for the remote sensor to communicate with the manager, since it is behind a NAT.

For instance, assuming the following architecture:

Remote Network (including Forward Node, 192.168.33.13) <--> Remote Network VPN Endpoint, 192.168.33.10 <--> Internet <--> AWS VPN Endpoint, 10.55.1.10 <--> AWS Security Onion Manager, 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 `minion` and `sensor` arrays in `/opt/so/saltstack/local/salt/firewall/hostgroups.local.yaml`.

### 5.12.8 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/
```

## 5.13 Azure Cloud Image

If you would like to deploy Security Onion 2.3 in Azure, we have an Azure compatible image that is already built for you. As this is an older Security Onion version, it must be accessed via the Azure CLI rather than the Azure Marketplace.

**Warning:** Existing Security Onion installations in Azure 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.

---

**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/>.

---

### 5.13.1 Requirements

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

While Azure has recently begun offering ephemeral storage, which potentially could offer increased disk performance for search nodes, the setup required for configuring them is out of scope of this documentation. We recommend 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: 200GB Premium SSD

Evaluation

- Quantity: 1
- Type: Standard\_D8as\_v4
- Storage: 200GB 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. Note that when using virtual machines with the minimum RAM requirements you may need to enable memory swapping.

### 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

## 5.13.2 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

## 5.13.3 Create Security Onion Instances

### Instance Creation

As this is an older Security Onion version, it must be accessed via the Azure CLI rather than the Azure Marketplace. To configure a Security Onion instance (repeat for each node in a distributed grid), follow these steps:

- Using the Azure CLI (available in the Azure web portal if not installed locally) run the following commands to launch a new Security Onion 2.3 instance, adjusting the parameters as needed:

```
az vm image terms show --urn securityonionsolutions:securityonion:so2:2.3.260
az vm image terms accept --urn securityonionsolutions:securityonion:so2:2.3.260
az group create --name mySOGroup --location eastus
az vm create \
  --resource-group mySOGroup \
  --name so-vm-manager \
  --admin-username onion \
  --generate-ssh-keys \
  --image securityonionsolutions:securityonion:so2:2.3.260 \
  --plan-name so2 \
  --plan-product securityonion \
  --plan-publisher securityonionsolutions \
  --size Standard_D8ds_v4
```

- Choose the appropriate Size based on the desired hardware requirements. For assistance on determining resource requirements please review the Requirements section above.
- Ensure you have access to the generated SSH private key, or copy the key locally if using the Azure CLI in the web portal
- Stop the new VM after deployment completes.
- Edit the VM and:

- Adjust the OS disk size to be at least 100GB in size.
- 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.
- If this VM is a single-node grid, or is sensor node, attach the monitoring network interface created earlier.
- Adjust public IP, virtual networks, etc as necessary.
- 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
```

### 5.13.4 Manager Setup

After SSH'ing into the node, setup will begin automatically. Follow the prompts, selecting the appropriate install options. Continue instructions below for applicable nodes.

**Note:** As Security Onion 2.3 approaches EOL, new cloud images are no longer being released. However, maintenance releases are available for cloud installation. After the manager completes setup it is important to run `sudo soup` to apply all newer maintenance releases and patches. This only is necessary on the manager node.

#### All Distributed Manager Nodes

For distributed manager nodes, if connecting sensors through the VPN instance, adjust the Security Onion firewall as shown in the below commands:

Run `so-firewall includehost minion <inside interface of your VPN concentrator>`. Ex:

```
so-firewall includehost minion 10.99.1.10
```

Run `so-firewall includehost sensor <inside interface of your VPN concentrator>`. Ex:

```
so-firewall --apply includehost sensor 10.99.1.10
```

At this time your Manager is ready for remote minions to start connecting.

### 5.13.5 Search Node Setup

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

### 5.13.6 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.

### 5.13.7 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/
```

## 5.14 Configuration

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

---

**Note:** Setup uses keyboard navigation and you can use arrow keys to move around. Certain screens may provide a list and ask you to select one or more items from that list. You can use the space bar to select items and the Enter key to proceed to the next screen.

---

**Warning:** If you use DHCP and your IP address changes, this can cause problems. If you want to use DHCP, make sure that you have a DHCP reservation so that your IP address does not change. Otherwise, use a static IP address to be safe.

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 run *so-import-pcap* to import pcap files or *so-import-evtx* to import Windows event logs in EVTX format.

### 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 have no sniffing interfaces of its own).

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.

Build the manager by running **Setup**, selecting the **DISTRIBUTED** install submenu, and choosing the **New Deployment** option. Save the **soremove** password so that you can join nodes to the manager.

Build search nodes by running **Setup**, selecting the **DISTRIBUTED** install submenu, and choosing the **Existing Deployment** option. When prompted, join them to the manager node using the password that you created for the **soremove** account.

Build forward nodes by running **Setup**, selecting the **DISTRIBUTED** install submenu, and choosing the **Existing Deployment** option. When prompted, join them to the manager node using the password that you created for the **soremove** account.

Proceed to the *After Installation* section.

## 5.15 After Installation

### 5.15.1 SSH Key Change

Depending on what kind of installation you did, you may have seen a warning at the end of **Setup** about SSH key changes.

NOTE: You will receive a warning upon SSH reconnect that the host key has changed.

This is expected due to hardening of the OpenSSH server config.

The host key algorithm will now be ED25519, follow the instructions given by your SSH client to remove the old key fingerprint then retry the connection.

<Ok>

For more information, see the [SSH](#) section.

### 5.15.2 Adjust firewall rules using `so-allow`

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 allow other IP addresses, you can manually run [so-allow](#).

### 5.15.3 Services

- Verify services are running with the [so-status](#) command:

```
sudo so-status
```

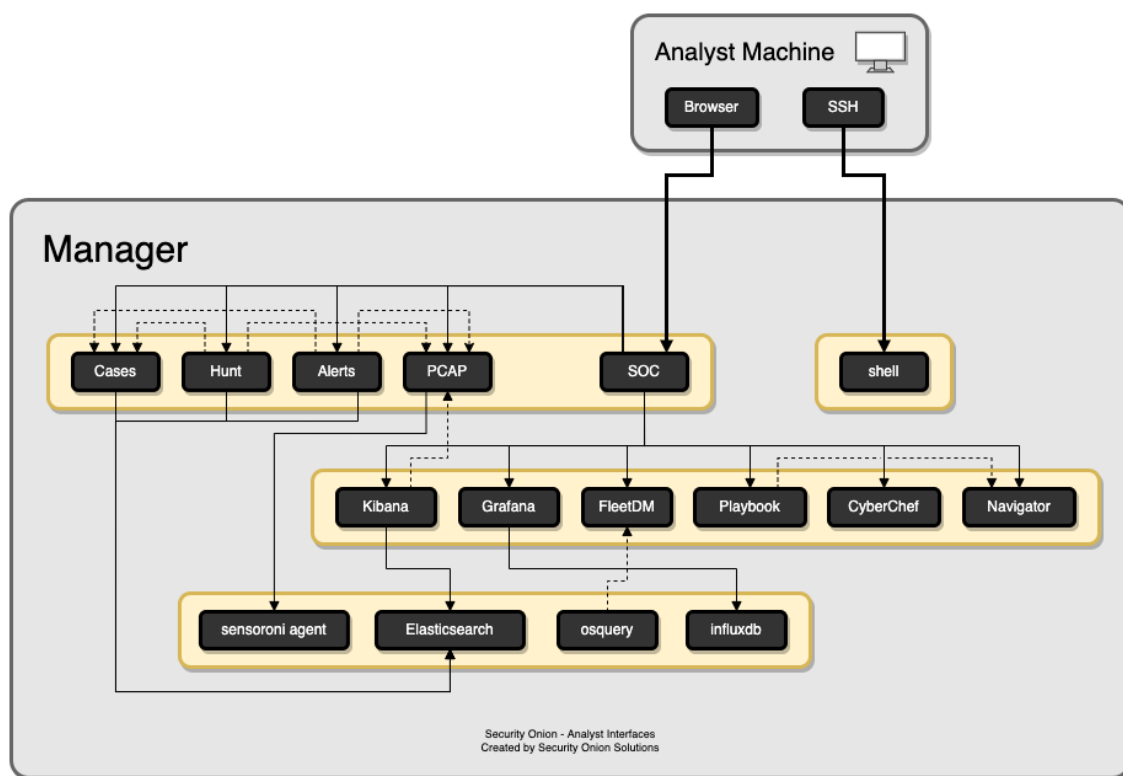
### 5.15.4 Data Retention

- Review the [Curator](#) and [Elasticsearch](#) sections to see if you need to change any of the default index retention settings.

### 5.15.5 Other

- Full-time analysts may want to connect using a dedicated [Analyst VM](#).
- Any IDS/NSM system needs to be tuned for the network it's monitoring. Please see the [Tuning](#) section.
- Configure the OS to use your preferred [NTP](#) server.

## SECURITY ONION CONSOLE (SOC)



Once you've run `so-allow` and allowed your IP address, you can then connect to Security Onion Console (SOC) with your web browser. We recommend chromium or 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

Email Address

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

Password



LOGIN

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.



Security Onion

Overview

**Getting Started**

New to Security Onion 2? Click the menu in the upper-right corner and you'll find links for [Help](#) and a [Cheatsheet](#) that will help you get started. In addition, check out our free Security Onion 2 Essentials online course, available on our [Trail](#).

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 identified a threat 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.

**Customize This Space**

Make this area your own by customizing the content. The content is stored in the `motd.md` file, which uses the common Markdown (.md) format. To learn more about the format, please see [markdownguide.org](#).

To customize this content, login to the manager via SSH and execute the following command:

```
sudo cp /opt/so/saltstack/default/salt/soc/files/soc/motd.md /opt/so/saltstack/local/salt/soc/files/soc/
```

Then edit the new file as desired using your favorite text editor.

Finally, restart SOC to make the changes take effect:

```
sudo so-soc-restart
```

Brought to you by:

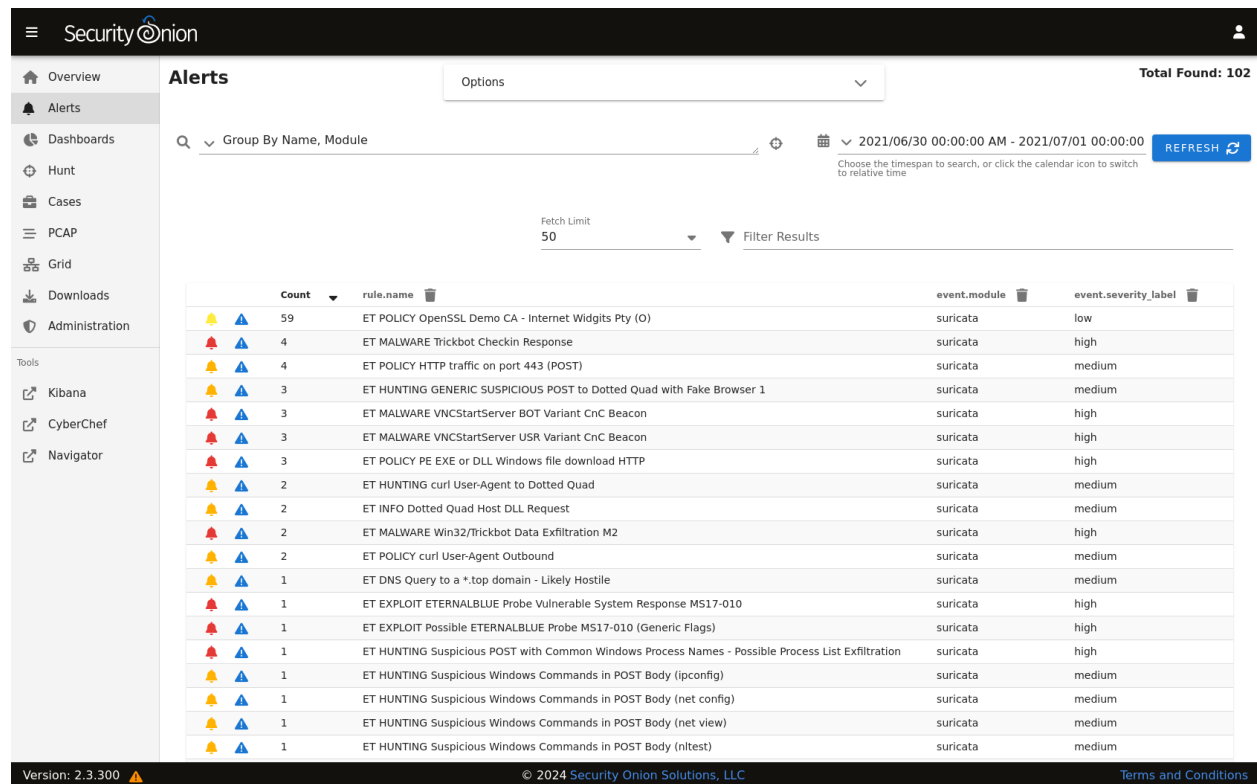
Security Onion SOLUTIONS

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)* gives you access to our Alerts interface. This interface gives you an overview of the alerts that Security Onion is generating and allows you to quickly drill down into details, pivot to [Hunt](#) or the [PCAP](#) interface, and escalate alerts to [Cases](#).



Security Onion Alerts interface showing a list of alerts. The interface includes a sidebar with navigation options (Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration) and a main area displaying a table of alerts. The table columns are Count, rule.name, event.module, and event.severity\_label. The alerts are filtered by time range (2021/06/30 00:00:00 AM - 2021/07/01 00:00:00) and a search query (Group By Name, Module). The total number of alerts found is 102. The interface also includes a 'REFRESH' button and a 'Filter Results' option.


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 ETHERNALBLUE Probe Vulnerable System Response MS17-010	suricata	high
1	ET EXPLOIT Possible ETHERNALBLUE 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

## 6.1.1 Options

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


### Toggles

There is 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.

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
 	<b>Quick Drilldown</b>	ET MALWARE Tordal/Hancitor/Chanitor Checkin
 	60	ET POLICY OpenSSL Demo CA - Internet Widgits Pty (O)
 	41	ET HUNTING Suspicious GET To gate.php with no Referer
 	41	ET INFO Suspicious Windows NT version 9 User-Agent
 	41	ET MALWARE Generic gate .php GET with minimal headers

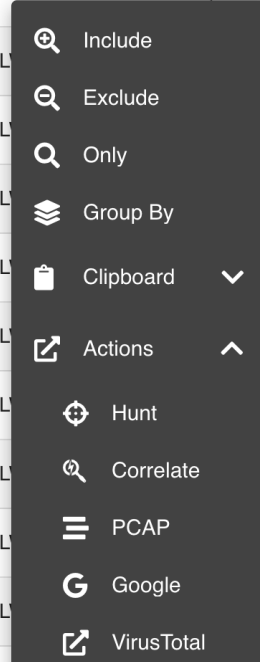
-  Include
-  Exclude
-  Only
-  Drilldown

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 PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.577 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.576 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102
>	 	2021-04-29 12:37:38.575 -04:00	ET MALWARE Backdoor family PCrat/Gh0st CnC traffic (OUTBOUND) 102



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
<div> </div>		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.

#### Clipboard

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

#### Actions

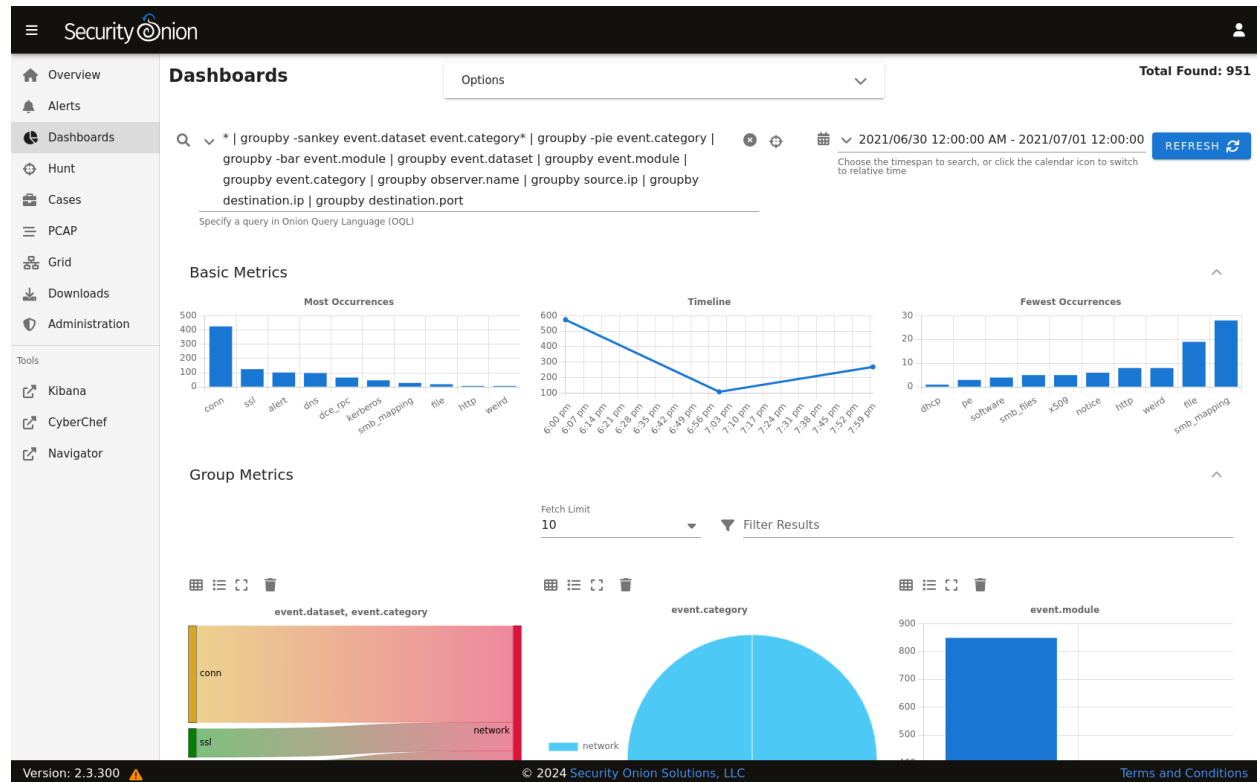
The Actions sub-menu has several different options:

- Clicking the Hunt option will start a new search for the selected value and will aggregate the results by `event.module` and `event.dataset` to give you a good overview of what types of data are available for that indicator.
- 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.

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

## 6.2 Dashboards

*Security Onion Console (SOC)* has its own Dashboards interface. This interface 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 drop-down menu that allows you to set options such as Auto Apply, Exclude case data, Automatic Refresh Interval, and Time Zone.

#### Auto Apply

The Auto Apply option defaults to enabled:




When enabled, Dashboards 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.

### Automatic Refresh Interval

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

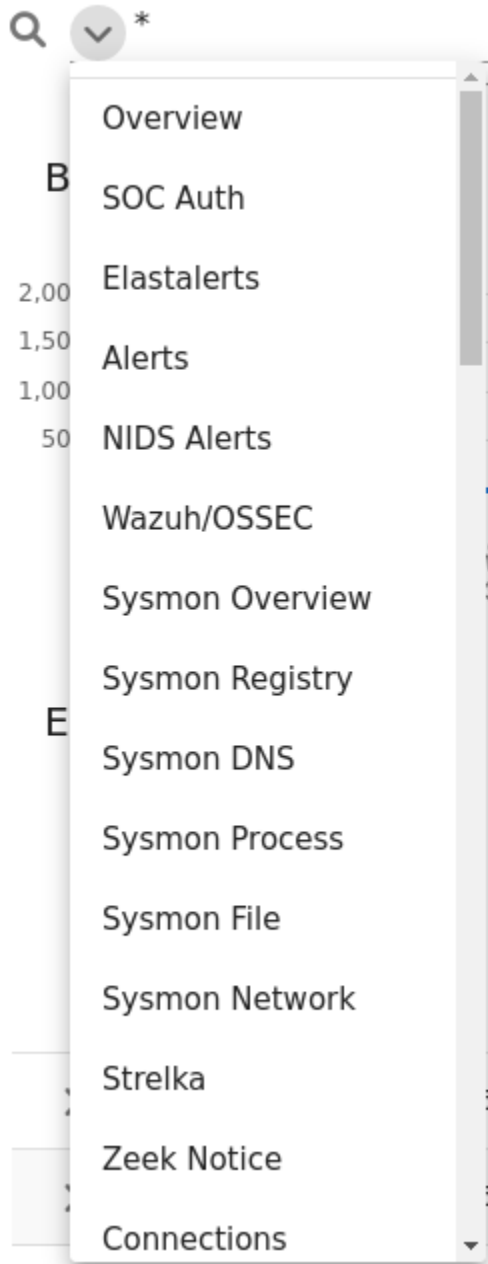


### Time Zone

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*, HIDS alerts from *Wazuh*, 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 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: `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	log.id.uuid	network.community_id
> ▲ 2022-08-08 21:27:00.681 +00:00	10.8.8.101	50461	104.238.220.131	80	http	http	CNUcGi4sjsLRs9Uw9	1:J5S04ymnxGdJn5x6a9bfZUX7H7s=
> ▲ 2022-08-08 21:27:00.681 +00:00	10.8.8.101	50462	104.238.220.131	80	http	http	Ce7wamDffkpu3bLc	1:CtuBQZH4FjhfttN4rVHqGIlC0c=
> ▲ 2022-08-08 21:26:43.995 +00:00	10.8.8.101	50460	146.190.25.131	44	ssl	ssl	CYnPS14RYINmCCXnc9	1:3kCJMwTEckQPXOWG0jPNCXMI0xY=
> ▲ 2022-08-08 21:26:43.936 +00:00	10.8.8.101	59531	10.8.8.1	53	dns	dns	CVnATp1SP2DMo3qNn5	1:wMFN3gurTH3Orec3Rb/QnsMW5g8=
> ▲ 2022-08-08 21:21:39.614 +00:00	10.8.8.101	50436	146.190.25.131	44	ssl	ssl	CX9po12HkI8kOCs3	1:/2vRFDiecPSbrOwvCUccZt1pZ8=
> ▲ 2022-08-08 21:21:39.562 +00:00	10.8.8.101	55126	10.8.8.1	53	dns	dns	CcxJUr1lJUn2j0mYz2	1:OKISDoL5cUWm3gcAaSt6d9XK80o=
> ▲ 2022-08-08 21:16:34.160 +00:00	10.8.8.101	50409	146.190.25.131	44	ssl	ssl	CdiaG52dujkblik3c4	1:tlytDaj5TIwfnlprf2GqTelFDR8=
> ▲ 2022-08-08 21:16:34.093 +00:00	10.8.8.101	55126	10.8.8.1	53	dns	dns	CPIEKVKaZaVZ6m3V3	1:OKISDoL5cUWm3gcAaSt6d9XK80o=
> ▲ 2022-08-08 21:11:15.698 +00:00	10.8.8.101	50385	146.190.25.131	44	ssl	ssl	CtonpC4losTuff6jm6	1:YrGg7N22UWVzjVixU4WNECzWF50=
> ▲ 2022-08-08 21:11:15.640 +00:00	10.8.8.101	55126	10.8.8.1	53	dns	dns	C3e4Qa2qld2G25FT6e	1:OKISDoL5cUWm3gcAaSt6d9XK80o=

Rows per page: 10 51-60 of 100

The backend data fetch took 0.013 seconds. The total round trip took 0.026 seconds.

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 ▼	source.ip	source.port	destination.ip
▼ 	2021-04-29 12:33:08.848 -04:00	10.10.10.70	1044	10.10.10.10
 @timestamp	2021-04-29T16:33:08.848Z			
 client.ip	10.10.10.70			
 client.ip_bytes	48			
 client.packets	1			
 client.port	1044			
 connection.bytes.missed	0			
 connection.history	Sr			
 connection.local.originator	true			
 connection.local.responder	true			
 connection.state	REJ			
 connection.state_description	Connection			
 destination.ip	10.10.10.			
 destination.port	4445			
 ecs.version	1.6.0			
 event.category	network			
 event.dataset	conn			
 event.module	zeek			
 ingest.timestamp	2021-04-2			
 log.file.path	/nsm/zeek			
 log.id.uid	Cuzlf73R			
 log.offset	4475233			

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

## 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.

### 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 aggregate the results by `event.module` and `event.dataset` to give you a good overview of what types of data are available for that indicator.
- 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.

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
```

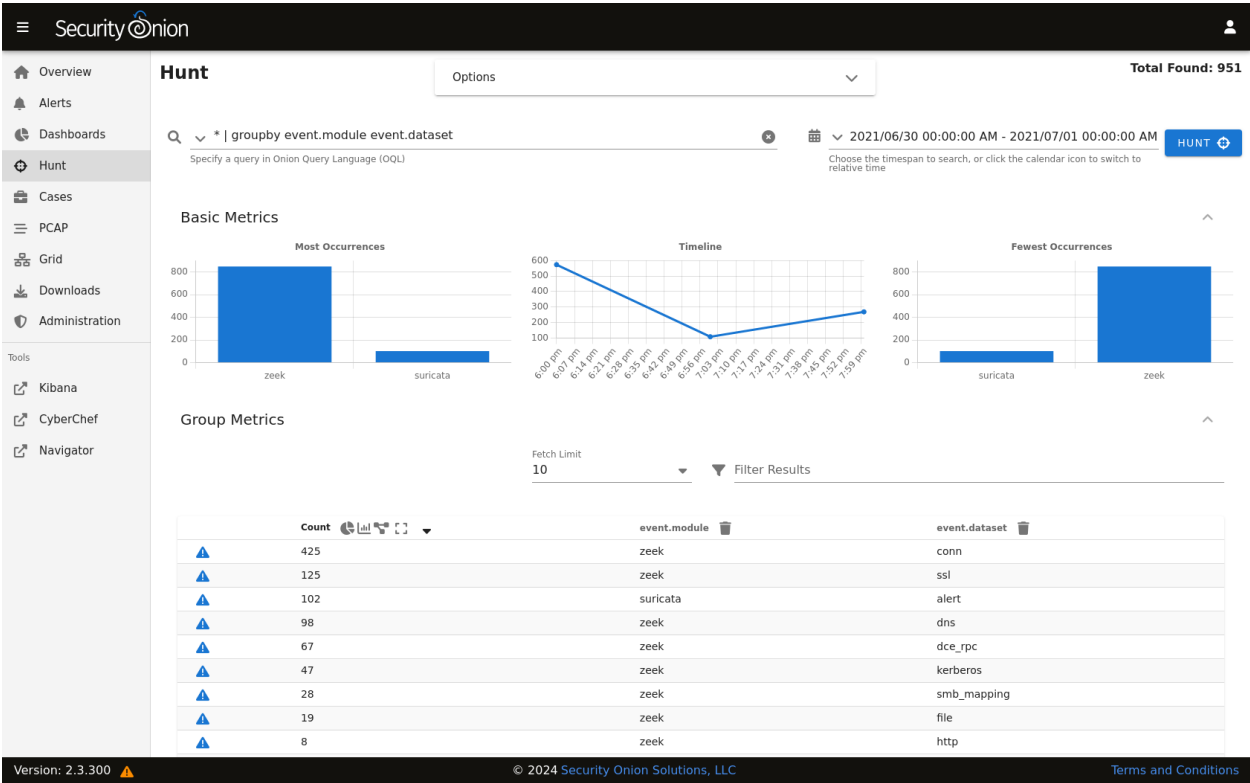
Please note that adding the asterisk to a non-string field may not work as expected. As an alternative, you may be able to use the asterisk with the equivalent `keyword` field if it is available. For example, `source.geo.ip*` may return 0 results, or a query failure error, but `source.geo.ip.keyword*`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

## 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

Hunt is similar to our *Dashboards* interface but is tuned more for threat hunting.



The main difference between Hunt and *Dashboards* is that Hunt's default queries are more focused than the overview queries in *Dashboards*. Also, 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.

The screenshot shows the Security Onion console interface. At the top, a search bar contains the query: `event.dataset:conn | groupby source.ip destination.ip network.protocol destination.port`. Below the search bar, a dropdown menu is open, displaying a list of saved queries. The queries are as follows:

- \* | groupby observer.name**  
Default Query — Show all events grouped by the origin host
- \* | groupby event.module event.dataset**  
Log Type — Show all events grouped by module and dataset
- \_type:elastalert | groupby rule.name**  
Elastalerts —
- event.dataset: alert | groupby event.module**  
Alerts — Show all alerts grouped by alert source
- event.category: network AND event.dataset: alert | groupby rule.category rule.gid rule.uuid rule.name**  
NIDS Alerts — Show all NIDS alerts grouped by alert
- event.module:ossec AND event.dataset:alert | groupby rule.category**  
Wazuh/OSSEC Alerts — Show all Wazuh alerts grouped by category
- event.module:ossec AND event.dataset:alert | groupby process.command\_line**  
Wazuh/OSSEC Commands — Show all Wazuh alerts grouped by command line
- event.module:ossec AND event.dataset:alert | groupby process.name.keyword**  
Wazuh/OSSEC Processes — Show all Wazuh alerts grouped by process name
- event.module:ossec AND event.dataset:alert | groupby user.escalated.keyword**  
Wazuh/OSSEC Users — Show all Wazuh alerts grouped by username
- event.module:sysmon | groupby event.dataset**  
Sysmon Events — Show all Sysmon logs grouped by event type
- event.module:sysmon | groupby event.dataset, user.name.keyword**  
Sysmon Usernames — Show all Sysmon logs grouped by username
- event.dataset:notice | groupby notice.note notice.message**  
Zeek Notice — Show notices from Zeek

## 6.4 Cases

Security Onion includes its own 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!](https://youtu.be/y_kr_hrtqVc!)

## 6.4.1 Installation

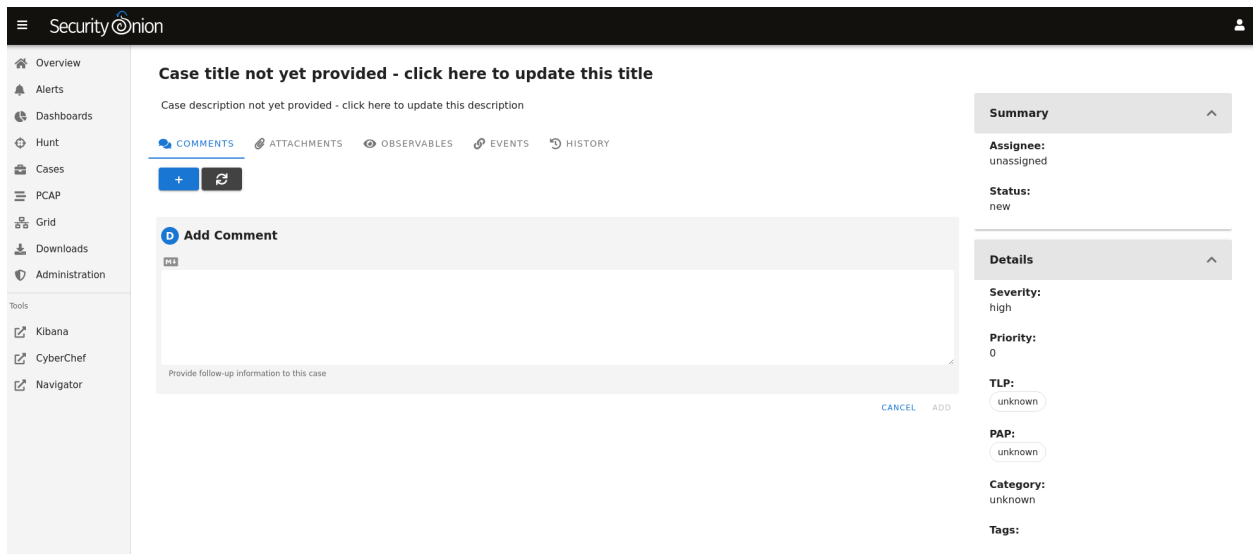
Cases is a part of *Security Onion Console (SOC)*. It's automatically enabled when doing an Import, Eval, Standalone, Manager, or ManagerSearch installation. If you want the quickest and easiest way to try out Cases, you can follow our *First Time Users* guide to install a minimal Import installation.

## 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/>.

Security Onion

Overview
Alerts
Dashboards
Hunt
Cases
PCAP
Grid
Downloads
Administration

Tools
Kibana
CyberChef
Navigator

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

+

Please check all web servers in the DMZ and determine which ones have log4j and if they are vulnerable. Also check Security Onion and other monitoring systems.

10 web servers have log4j:

- For 6 of them, we were able to update to log4j 2.17.1.
- For 2 of them, we were not able to update log4j but were able to apply the mitigation of removing the vulnerable class.
- One is awaiting an official vendor update.
- One is a piece of software that is EOL and no longer supported. We need to determine if we can safely apply a workaround.

Security Onion release an [initial hotfix on 2021/12/10](#).  
Security Onion then released a [second hotfix on 2021/12/13](#).  
[Security Onion 2.3.91](#) updated to Elastic 7.16.2 which includes Log4j 2.17.0.  
Security Onion 2.3.100 updates log4j to version 2.17.1 to satisfy vulnerability scanners.

Summary

Assignee:  
jim@example.com

Status:  
in progress

Details

Severity:  
high

Priority:  
1

TLP:  
amber

PAP:  
red

Category:  
general

Tags:  
log4j

## 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.

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

vendor\_statement.jpg (1,285,029 Bytes)

**SHA256:**  
 1a5270ad07ef885de0bc4e106df937f5e9b27b55dc59b3a0b0ea26642dcf0e9

**SHA1:**  
 234f783a2edc560a66628e9caec904f43932b0a7

**MD5:**  
 2a09d90dc3b393e4650ecf6e0632622a

**Description:**  
 Vendor statement

**TLP:**  
 white

**Tags:**

Summary

**Assignee:**  
 jim@example.com

**Status:**  
 in progress

**Details**

**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.

**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	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	/S{jndiIdap:/121.140.99.236:138...
	Jan 19, 2022 9:31 PM	Jan 19, 2022 9:31 PM	user-agent	S{jndiIdap:/121.140.99.236:1389...

Summary

**Assignee:**  
 jim@example.com

**Status:**  
 in progress

**Details**

**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).

## 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....	-JvFPn4B8vgkSQ33E1Og	network	zeek	http
> +	2022-01-01 12:32:03....	4pvFPn4B8vgkSQ33E...	network	zeek	notice
> +	2022-01-01 12:32:03....	_JvFPn4B8vgkSQ33E1...	network	zeek	http
> +	2022-01-01 12:32:03....	5jvFPn4B8vgkSQ33EVfu	network	zeek	notice
> +	2022-01-01 12:32:09....	_pvFPn4B8vgkSQ33E1...	network	zeek	http
> +	2022-01-01 12:32:09....	55vFPn4B8vgkSQ33E...	network	zeek	notice
> +	2022-01-01 19:55:04....	oJvFPn4B8vgkSQ33FV...	network	zeek	http
> +	2022-01-01 19:55:04....	PjvFPn4B8vgkSQ33EIPd	network	zeek	notice
> +	2022-01-01 19:55:04....	OjvFPn4B8vgkSQ33EIPd	network	zeek	notice

**Summary**

**Assignee:**  
jim@example.com

**Status:**  
in progress

**Details**

**Severity:**  
high

**Priority:**  
1

**TLP:**  
amber

**PAP:**  
red  
Permissible Actions Protocol

**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....	-JvFPn4B8vgkSQ33E1Og	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  
Permissible Actions Protocol

**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 web interface. The left sidebar contains navigation links: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Tools, Kibana, CyberChef, and Navigator. The main content area is titled 'Attempts to exploit log4j vulnerability against public facing web servers in DMZ'. Below the title, there's a link to a log4j vulnerability announcement. The 'HISTORY' tab is selected, showing a table of actions. The table has columns: Actions, User, Time, Kind, and Operation. The actions are performed by 'doug@example.com' and 'tom@example.com'. On the right, there's a 'Summary' section with fields for Assignee (jim@example.com), Status (in progress), and Details (Severity: high, Priority: 1, TLP: amber, PAP: red, Category: general, Tags: log4j).

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

## 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 web interface with the 'Cases' tab selected. The top bar shows 'Total Found: 3'. Below the title, there's a search bar with 'Open Cases' and a query input field. The table below shows a list of cases with columns: Timestamp, so\_case.title, so\_case.status, so\_case.severity, so\_case.assigneeid, and so\_case.createTime. The cases are: 'Attempts to exploit log4j vulnerability against public facing web servers in DMZ' (in progress, high), 'John Doe in Accounting received phishing email' (new, critical), and 'SQL injection attempts against web servers in DMZ' (in progress, medium). At the bottom, there's a 'Rows per page' dropdown set to 50 and a '1-3 of 3' indicator.

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

### 6.4.9 Options

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

There is 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 dropdown 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 *Backups* section.

### 6.4.14 Analyzers

Security Onion includes 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	Hash	IP	Mail	Other	URI	URL	User Agent
Alienvault OTX	✓	✓						✓
EmailRep				✓				
Greynoise			✓					
LocalFile	✓	✓	✓		✓		✓	
Malware Hash Registry		✓						
Pulsedive	✓	✓	✓			✓	✓	✓
Spamhaus			✓					
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	Analysts 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	Analysts 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:

virustotal

 Malicious

Bytes: 31,959



```
{
  "response": {
    "data": [
      {
        "attributes": {
          "authentihash": "2a659725ff96571699cfb3d9bfc55c00e596bbfeceda29f4647525ed477d10b6",
          "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"
            }
          }
        }
      ]
    }
  }
}
```

## 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
- EmailRep
- GreyNoise
- LocalFile
- Pulsedive
- Urlscan
- VirusTotal

To configure an analyzer, consult its documentation via the provided README file. For example, the README for the VirusTotal analyzer can be found at <https://github.com/Security-Onion-Solutions/securityonion/tree/dev/salt/sensoroni/files/analyzers/virustotal>.

## 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)* gives you access to our PCAP interface. This interface allows you to access your full packet capture that was recorded 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	log.id.uid	network.community_id
2022-08-08 21:27:00.681 +00:00	10.8.8.101	50461	104.238.220.131	80	http	http	CNUcGi4sjsLSRs9Uw9	1J55O4ymnxGdjN5x6a9bfZUX7H7s=
2022-08-08 21:27:00.681 +00:00	10.8.8.101	50462	104.238.220.131	80	http	http	Ce7wamDffkpu3bLc	1CtuBQZH4FjhfttN4rVHgGllC0c=
2022-08-08 21:26:43.995 +00:00	10.8.8.101	50460	146.190.25.131	443	ssl	ssl	CYnPS14RYINmCCXnc9	1:3kCJMwTEckQPXOWG0PNCXMI0xY=
2022-08-08 21:26:43.936 +00:00	10.8.8.101	59531	10.8.8.1	53	dns	dns	CVnATp1SP2DMo3qNn5	1:wMFN3gurTH3Orec3Rb/QnsMW5g8=
2022-08-08 21:21:39.614 +00:00	10.8.8.101	50436	146.190.25.131	443	ssl	ssl	CX9po12hHkl8kOCs3	1:/2vRFDecPSbrOwvCuczcZt1pZ8=
2022-08-08 21:21:39.562 +00:00	10.8.8.101	55126	10.8.8.1	53	dns	dns	CcxJUr1lJUn2j0mYz2	1:OKISDoL5cUWm3gcAaSt6d9XK80o=
2022-08-08 21:16:34.160 +00:00	10.8.8.101	50409	146.190.25.131	443	ssl	ssl	CdiaG52dujkbk3c4	1:ttYtDaj5Tlwnlprf2GqTeIFDR8=
2022-08-08 21:16:34.093 +00:00	10.8.8.101	55126	10.8.8.1	53	dns	dns	CPiEKVKazaVZ6m3V3	1:OKISDoL5cUWm3gcAaSt6d9XK80o=
2022-08-08 21:11:15.698 +00:00	10.8.8.101	50385	146.190.25.131	443	ssl	ssl	CtonpC4IosTuf6jm6	1:YrGg7N22UWVzjVlxU4WNECzWFs0=
2022-08-08 21:11:15.640 +00:00	10.8.8.101	55126	10.8.8.1	53	dns	dns	C3e4Qa2Qld2G25FT6e	1:OKISDoL5cUWm3gcAaSt6d9XK80o=

Rows per page: 10 51-60 of 100

The backend data fetch took 0.013 seconds. The total round trip took 0.026 seconds.

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

Overview Alerts Dashboards Hunt Cases PCAP Grid Downloads Administration

Tools Kibana CyberChef Navigator

PCAP

+ ID Owner

1001	doug@example.com
1002	doug@example.com
1003	doug@example.com
1005	doug@example.com

Add job

Sensor ID

The sensor ID must match an actual sensor ID in order for this job to be processed.

Import ID

UUID value that is output from so-import-pcap. Only needed for imported PCAPs.

Source IP

Optional source IP address to include in this job filter

Source Port

Optional source TCP port to include in this job filter

Destination IP

Optional destination IP address to include in this job filter

Destination Port

Optional destination TCP port to include in this job filter

Filter Begin

Filter start time in RFC 3339 format (Ex: 2020-10-16 13:00:00.230-04:00). Unused for imported PCAPs.

Filter End

Filter end time in RFC 3339 format (Ex: 2020-10-16 15:30:00.230-04:00). Unused for imported PCAPs.

CLEAR CANCEL ADD

Sensor ID Status Actions

securityonion	Completed	
securityonion	Completed	
securityonion	Completed	
securityonion	Completed	

Rows per page: 10 1-4 of 4

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

Security Onion

# 1001 securityonion 176.10.125.8:80 172.16.3.130:49457

Filter Results

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

Rows per page: 10 1-10 of 24

LOAD MORE

Version: 2.3.300 © 2024 Security Onion Solutions, LLC Terms and Conditions

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.

Security Onion

# 1005 securityonion 192.168.222.177:56106 151.101.193.67:80

Filter Results

```
GET /Logon.php?LOB=RB0&pageLabel=page_ HTTP/1.1
User-Agent: curl/7.29.0
Host: cnn.com
Accept: */*

HTTP/1.1 301 Moved Permanently
Server: Varnish
Retry-After: 0
Content-Length: 0
Cache-Control: public, max-age=600
Location: http://www.cnn.com/Logon.php?LOB=RB0&pageLabel=page_
Accept-Ranges: bytes
Date: Thu, 20 Jan 2022 13:17:12 GMT
Via: 1.1 varnish
Connection: close
Set-Cookie: countryCode=US; Domain=.cnn.com; Path=/; SameSite=Lax
Set-Cookie: stateCode=GA; Domain=.cnn.com; Path=/; SameSite=Lax
Set-Cookie: geoData=grovetown[GA]30813[US]NA[-500]broadband[33.460]-82.210; Domain=.cnn.com; Path=/; SameSite=Lax
X-Served-By: cache-fty21368-FTY
X-Cache: HIT
X-Cache-Hits: 0
```

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.

Security Onion

# 1005 securityonion 192.168.222.177:56106


Filter Results

```
GET /Logon.php?LOB=RBG&_pageLabel=page_ HTTP/1.1
User-Agent: curl/7.29.0
Host: cnn.com
Accept: */*

HTTP/1.1 301 Moved Permanently
Server: Varnish
Retry-After: 0
Content-Length: 0
Cache-Control: public, max-age=600
Location: http://www.cnn.com/Logon.php?LOB=RBG&_pageLabel=page_
Accept-Ranges: bytes
Date: Thu, 20 Jan 2022 13:17:12 GMT
Via: 1.1 varnish
Connection: close
Set-Cookie: countryCode=US; Domain=.cnn.com; Path=/; SameSite=Lax
Set-Cookie: stateCode=GA; Domain=.cnn.com; Path=/; SameSite=Lax
Set-Cookie: geoData=grovetown|GA|30813|US|NA|-500|broadband|33.4
X-Served-By: cache-fty21368-FTY
X-Cache: HIT
X-Cache-Hits: 0
```

Copy to clipboard  
Hunt  
CyberChef  
Google  
VirusTotal

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 an [Analyst VM](#), 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.

Security Onion

PCAP

ID	Owner	Date Queued	Date Completed	Sensor ID	Status	Actions
1001	doug@example.com	2022-01-08 21:22:47.216 +00:00	2022-01-08 21:22:47.975 +00:00	securityonion	Completed	Hunt CyberChef VirusTotal
1002	doug@example.com	2022-01-08 21:41:33.384 +00:00	2022-01-08 21:41:33.629 +00:00	securityonion	Completed	Hunt CyberChef VirusTotal
1003	doug@example.com	2022-01-09 12:34:32.322 +00:00	2022-01-09 12:34:32.625 +00:00	securityonion	Completed	Hunt CyberChef VirusTotal
1005	doug@example.com	2022-05-25 20:04:50.982 +00:00	2022-05-25 20:04:51.644 +00:00	securityonion	Completed	Hunt CyberChef VirusTotal

Rows per page: 10 1-4 of 4



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/`.

## 6.6 Grid

*Security Onion Console (SOC)* gives you access to our Grid interface. This interface allows you to quickly check the status of all nodes in your grid. It also includes a few different EPS (events per second) measurements:

- EPS (also shown as Production EPS) is how much a node is producing. This is taken from the number of events out in *Filebeat*.
- Consumption EPS is how much a search node is consuming.
- Grid EPS in the upper right corner is the sum of all Consumption EPS measurements in the entire grid.

The screenshot shows the Security Onion Grid interface. On the left is a sidebar with navigation links: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid (selected), Downloads, and Administration. Below these are tools: Kibana, Grafana, CyberChef, Playbook, FleetDM, and Navigator. The main content area is titled 'Grid' and shows a table of nodes. Above the table is a 'Filter Results' button. The table has columns: ID, Role(s), Address, Description, Version, Model, EPS, Date Updated, Earliest PCAP, Uptime, and Status. There are three rows of nodes. The first row is expanded, showing a detailed view of the 'sos-manager' node, including images of the front and rear of the appliance and a status summary: Online Since: May 26, 2022 6:11 PM, Production EPS: 15, Consumption EPS: 0, Process Status: OK, Connection Status: OK, and Raid Status: OK. The other two rows are 'sos-search-01' and 'sos-sensor-01', both with status 'OK'. At the bottom right, it says 'Rows per page: 10' and '1-3 of 3'.

ID	Role(s)	Address	Description	Version	Model	EPS	Date Updated	Earliest PCAP	Uptime	Status
sos-manager	Manager	172.16.224.248	SOS Grid Manager - Rack B/2	2.3.130	SOSMN	15	2022-05-26 19:04:56.116 +00:00	2022-05-26 19:04:56.115 +00:00	an hour	OK
sos-search-01	Search	172.16.224.249	SOS Search - Rack B/2	2.3.130	N/A	0	2022-05-26 19:04:56.262 +00:00	2022-05-26 19:04:56.093 +00:00	34 minutes	OK
sos-sensor-01	Forward	172.16.224.250	SOS Forward Sensor - Augusta	2.3.130	N/A	11	2022-05-26 19:05:01.324 +00:00	2022-05-26 18:57:26.782 +00:00	16 minutes	OK

If you have purchased our official Security Onion Solutions appliances, then the grid page will show pictures of the front and rear of the appliances, useful for walking through connectivity discussions with personnel in the data center. If you are not using official Security Onion Solutions appliances, then it will simply display a message to that effect. This also applies to the Raid Status field in the expanded grid panel: If you are using an official Security Onion appliance with raid support then you will see the corresponding status appear in that field.

If any nodes show anything other than OK, then run `so-status` on the node and check logs in `/opt/so/log/`.

## 6.7 Downloads

The *Security Onion Console (SOC)* Downloads page gives you access to some files that you might need to download:

☰

SecurityOnion

👤

🏠 Overview

🔔 Alerts

📊 Dashboards

🔍 Hunt

📁 Cases

📄 PCAP

🗃️ Grid

📁 Downloads

🔧 Administration

Tools

🔗 Kibana

🔗 CyberChef

🔗 Navigator

Downloads

When installing packages such as osquery or beats onto remote systems be sure to run `so-allow` on the Security Onion Manager node to allow network access through the firewall.

Elasticsearch Utilities (8.10.4)

- [Winlogbeat](#) (Windows)
- [Filebeat DEB](#) (Linux [amd64]: Debian, Ubuntu)
- [Filebeat RPM](#) (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)
- [Metricbeat DEB](#) (Linux [amd64]: Debian, Ubuntu)
- [Metricbeat RPM](#) (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)
- [Auditbeat DEB](#) (Linux [amd64]: Debian, Ubuntu)
- [Auditbeat RPM](#) (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)

Wazuh Agents (3.13.1-1)

- [MSI](#) (Windows)
- [DEB](#) (Linux [amd64]: Debian, Ubuntu)
- [RPM](#) (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)
- [PKG](#) (MacOS)

osquery Packages and Configs

- [MSI](#) (Windows)
- [DEB](#) (Linux [amd64]: Debian, Ubuntu)
- [RPM](#) (Linux [x86\_64]: Amazon, CentOS, Rocky, Fedora, Oracle, SUSE)
- [PKG](#) (MacOS)
- [RPM & DEB Config Flag File](#)
- [MSI Config Flag File](#)

These packages and configs are osquery files, customized for this specific Fleet install and will only be generated if Fleet has been installed. Due to macOS packaging constraints, the macOS PKG has not been customized for this Fleet install - osquery/Launcher will need to be configured post-deployment. These files are not signed. Signed, non-customized osquery packages can be obtained directly from [osquery.io](#). For further Fleet & osquery information, view our [online help](#).

Version: 2.3.300 ⚠️

© 2024 Security Onion Solutions, LLC

[Terms and Conditions](#)

## 6.8 Administration

*Security Onion Console (SOC)* includes an Administration page which shows current users:

☰

SecurityOnion

👤

🏠 Overview

🔔 Alerts

📊 Dashboards

🔍 Hunt

📁 Cases

📄 PCAP

🗃️ Grid

📁 Downloads

🔧 Administration

👤 Users

Users

Email Address ▲	First Name	Last Name	Note	Role(s)	Status	Actions
disabled@example.com				analyst	🚫	
doug@example.com				superuser	👤	
tom@example.com				analyst	👤	

Rows per page: 10 ▼ 1-3 of 3 < >

The Role(s) column lists roles assigned to the user as defined in the *Role-Based Access Control (RBAC)* section.

The Status column will show a different icon depending on the status of the account. In the screenshot above, the first account is disabled, the second account is enabled and has *MFA* enabled, and the third account is enabled but does not

118

Chapter 6. Security Onion Console (SOC)

have *MFA* enabled. Hovering over the icon in the Status column will show you these details as well.

## 6.9 Kibana

From <https://www.elastic.co/kibana>:

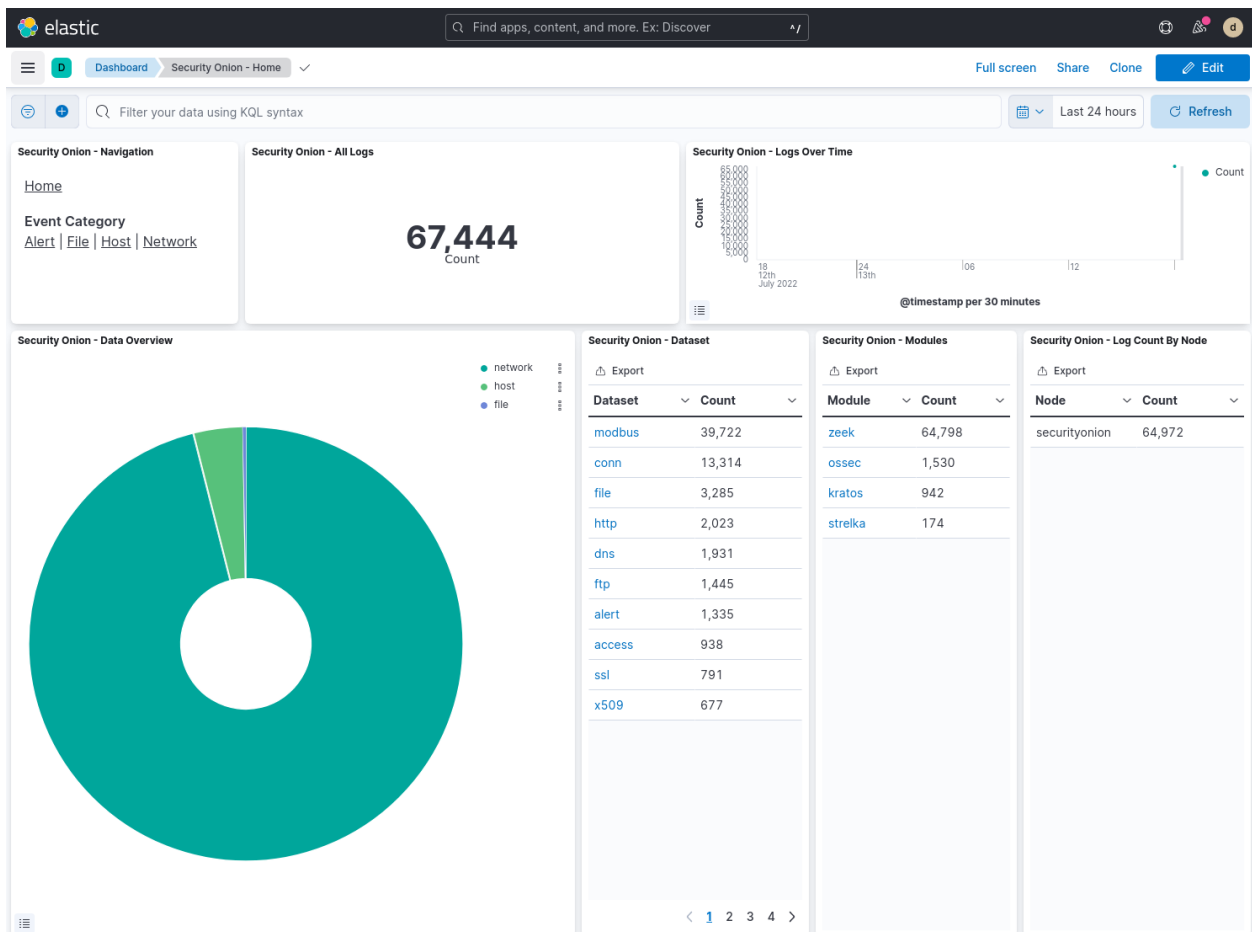
Kibana is a free and open user interface that lets you visualize your Elasticsearch data and navigate the Elastic Stack. Do anything from tracking query load to understanding the way requests flow through your apps.

### 6.9.1 Authentication

We support Elastic authentication via *so-elastic-auth*. Elastic authentication is automatically enabled for new installations, so you will need to log into Kibana using the same username and password that you use for *Security Onion Console (SOC)*.

### 6.9.2 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*, *Suricata*, and *Wazuh*. 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.

We've included the old 16.04 dashboards in case you performed an in-place upgrade and have any old 16.04 data. These dashboards are named with the `z16.04` prefix and will only show old 16.04 data from the old `*:logstash-*` indices. The new Security Onion 2 dashboards are all named with the Security Onion prefix and they should be used for any new data stored in the new `*:so-*` indices.

If you ever need to reload 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 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 Pivoting

Kibana uses multiple hyperlinked fields to accelerate investigations and decision-making:

#### PCAP/Cases

The `_id` field has a hyperlink which is labeled as *Hunt* and optionally pivot to PCAP/Cases. Clicking this hyperlink takes you to *Hunt* and searches for that particular record. From *Hunt*, you can then escalate the event to *Cases* or pivot to full packet capture via our *PCAP* interface (assuming it's a network event). You can usually find the `_id` field as the rightmost column in the log panels at the bottom of the dashboards.

24 documents	
network.community_id	_id
1:VpZ0Yp8+e0Z/blpdGIzTMR4E3DE=	<a href="#">Hunt and optionally pivot to PCAP/Cases</a>
1:VpZ0Yp8+e0Z/blpdGIzTMR4E3DE=	<a href="#">Hunt and optionally pivot to PCAP/Cases</a>

You can also find the hyperlinked `_id` field by drilling into a row in the log panel.

Security Onion - All Logs					
Time ↓	source.ip	source.port	destination.ip	destination.port	log.id.uid
Jan 12, 2022 @ 17:20:55.655	10.22.5.47	50513	204.79.197.212	80	Co38V12y8ZKjuHH0ld
Expanded document					
Table JSON					
Actions	Field		Value		
	_id		<a href="#">Hunt and optionally pivot to PCAP/Cases</a>		

## Indicator Dashboard

Several fields are hyperlinked to the Indicator dashboard to allow you to get all the information you can about a particular indicator. Here are just a few:

```
uid
source.ip
source.port
destination.ip
destination.port
```

### 6.9.4 Search Results

Search results in the dashboards and through Discover 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.

### 6.9.5 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.

### 6.9.6 Configuration

Kibana's configuration can be found in `/opt/so/conf/kibana/`. However, please keep in mind that most configuration is managed with [Salt](#), so if you manually make any modifications in `/opt/so/conf/kibana/`, they may be overwritten at the next salt update.

`/opt/so/conf/kibana/etc/kibana.yml` can be managed using the `kibana` pillar placed in the manager pillar file located under `/opt/so/saltstack/local/pillar/minions/`. The manager pillar file will end with either `*_manager.sls`, `*_managersearch.sls`, `*_standalone.sls`, or `*_eval.sls` depending on the manager type that was chosen during install.

- An example of a Kibana pillar may look as follows:

```
kibana:
  config:
    elasticsearch:
      requestTimeout: 120000
    data:
      autocomplete:
        valueSuggestions:
          timeout: 2000
          terminateAfter: 200000
    logging:
      root:
        level: warn
```

### 6.9.7 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
```

### 6.9.8 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>.

## 6.9.9 More Information

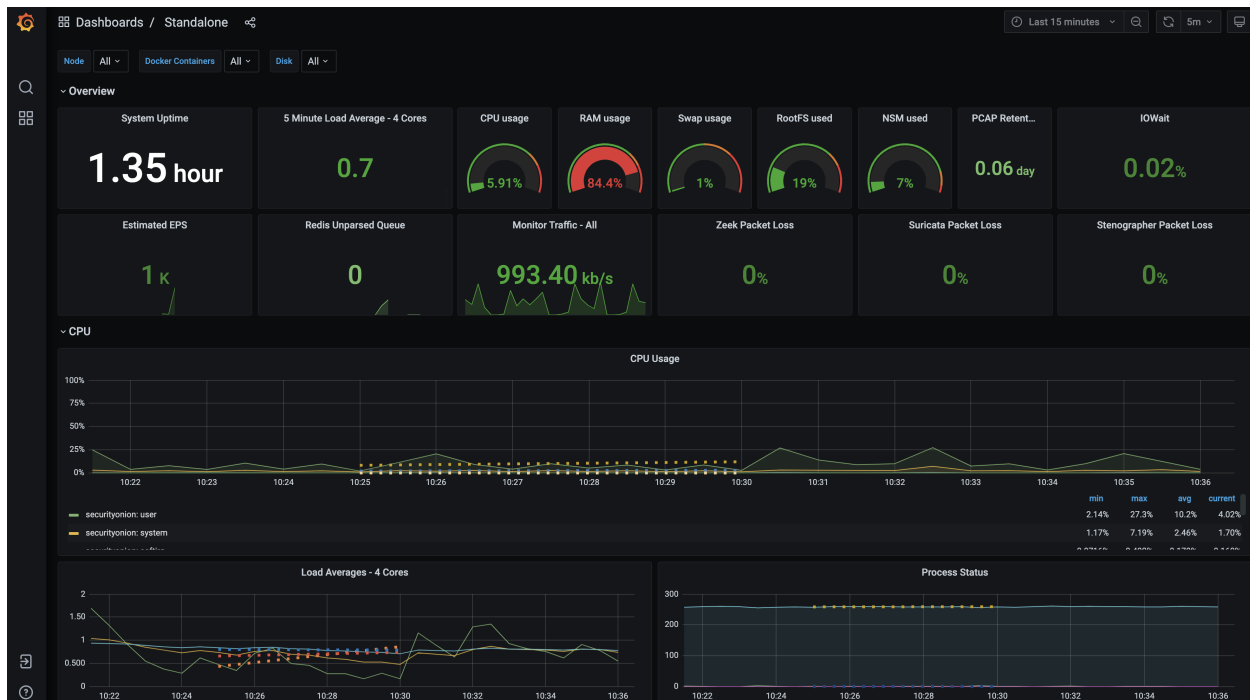
**Note:** For more information about Kibana, please see <https://www.elastic.co/kibana>.

## 6.10 Grafana

Once you've logged into *Security Onion Console (SOC)*, you can then click the Grafana link to see system health information.



You will start on the Security Onion Grid Overview dashboard. Depending on what kind of deployment you have, there will be at least one more dashboard available if you click the Dashboards icon on the left.



### 6.10.1 Data

Grafana has both high-resolution data and downsampled low-resolution data. Some Grafana graphs have dotted lines that show previous data that has been downsampled. High-resolution data will be purged after 30 days, leaving just the downsampled low-resolution data.

At some point, you may want to update your Grafana data as shown below. If you have a distributed deployment, you will run all commands on the manager.

If you want to remove some old data prior to downsampling, you can run `so-influxdb-clean`. This is optional and not required. `so-influxdb-clean` will ask how many days or weeks worth of data you want to retain and remove all data older than that.

If you want to downsample all data, run `so-influxdb-downsample`. This process could take a while depending on system resources and the amount of data that needs to be downsampled. For each measurement, the script will go day by day starting at 7/21/20 and downsample that day's data from the `autogen` retention policy into the `so_long_term` retention policy.

Once the downsampling is complete and you verify the downsampled data is available in Grafana (other than Processes, Disk I/O, Memory), you can remove the old data and free up space. The `so-influxdb-drop-autogen` script will remove the `autogen` retention policy and thus remove the old data that we previously downsampled.



## 6.10.2 Accounts

By default, you will be viewing Grafana as an anonymous user. If you want to make changes to the default Grafana dashboards, you will need to log into Grafana with username `admin` and the randomized password found via `sudo salt-call pillar.get secrets`.

## 6.10.3 Configuration

Grafana configuration can be found in `/opt/so/conf/grafana/etc/`. However, please keep in mind that most configuration is managed with *Salt*, so if you manually make any modifications in `/opt/so/conf/grafana/etc/`, they may be overwritten at the next salt update. The default configuration options can be seen in `/opt/so/saltstack/default/salt/grafana/defaults.yaml`. Any options not specified in here, will use the Grafana default.

## 6.10.4 Example

If you want to configure and enable SMTP for Grafana, place the following in the `global.sls` file. If you have files referenced in the config file, those can be placed in `/opt/so/saltstack/default/salt/grafana/etc/files/`. Those files will be then be placed in `/opt/so/conf/grafana/etc/files` on the minion and mapped to `/etc/grafana/config/files/` within the container.

```
grafana:
  config:
    smtp:
      enabled: true
      host: smtphost.mydomain:25
      user: myuser
      # If the password contains # or ; you have to wrap it with triple quotes wrapped
      ↪by single quotes. Ex """"#password;""""
      password: mypassword
#   cert_file: /etc/grafana/config/files/smtp_cert_file.crt
#   key_file: /etc/grafana/config/files/smtp_key_file.key
#   skip_verify: false
      from_address: admin@grafana.localhost
      from_name: Grafana
#   ehlo_identity: dashboard.example.com
```

## 6.10.5 More Information

---

**Note:** Check out our Grafana Alarms video at <https://youtu.be/8FmZ4MRe8Uk>.

For more information about Grafana, please see <https://grafana.com/>.

---

## 6.11 CyberChef

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

The Cyber Swiss Army Knife

CyberChef is a simple, intuitive web app for carrying out all manner of “cyber” operations within a web browser. These operations include simple encoding like XOR or Base64, more complex encryption like AES, DES and Blowfish, creating binary and hexdumps, compression and decompression of data, calculating hashes and checksums, IPv6 and X.509 parsing, changing character encodings, and much more.

The tool is designed to enable both technical and non-technical analysts to manipulate data in complex ways without having to deal with complex tools or algorithms.

There are four main areas in CyberChef:

1. The input box in the top right, where you can paste, type or drag the text or file you want to operate on.
2. The output box in the bottom right, where the outcome of your processing will be displayed.
3. The operations list on the far left, where you can find all the operations that CyberChef is capable of in categorised lists, or by searching.
4. The recipe area in the middle, where you can drag the operations that you want to use and specify arguments and options.

### 6.11.1 Screenshot

The screenshot shows the CyberChef web application interface. The top bar includes a download link, the last build date (8 months ago), and links for options, about, and support. The interface is divided into four main sections:

- Operations:** A sidebar on the left with a search bar and a list of operations categorized by function (Data format, Encryption / Encoding, Public Key, Arithmetic / Logic, Networking, Language, Utils, Date / Time, Extractors, Compression, Hashing).
- Recipe:** The central area where operations are added to a sequence. It shows a 'From Hexdump' operation followed by two 'Strip HTTP headers' operations. Below these are 'Strings' and 'Display total' options.
- Input:** A large text area on the top right where the user can paste or type data. It contains a sample HTTP GET request.
- Output:** A text area on the bottom right that displays the result of the operations. It shows a hex dump of the request body.

At the bottom of the interface, there is a 'STEP' button, a 'BAKE!' button with a chef icon, and an 'Auto Bake' checkbox.

## 6.11.2 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.

## 6.11.3 Example

Suppose you are looking at an interesting HTTP file download in *PCAP* and want to extract the file. 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.

Operations	Recipe	Input	Output
Search...	From Hexdump	length: 1448550 Lines: 19601	time: 101ms length: 311750 lines: 1137
Favourites		0000 47 45 54 20 2F 65 76 69 6C 2E 70 6E 67 20 48 54 GET /evil.png HT	
Data format		0016 54 50 2F 31 2E 31 0D 0A 48 6F 73 74 3A 20 31 39 TP/1.1..Host: 19	
Encryption / Encoding		0032 32 2E 31 36 38 2E 32 32 32 2E 32 0D 0A 55 73 2.168.222.22..Us	
Public Key		0048 65 72 2D 41 67 65 6E 74 3A 20 63 75 72 6C 2F 37 er-Agent: curl/7	
Arithmetic / Logic		0064 2E 37 36 2E 31 0D 0A 41 63 63 65 70 74 3A 20 2A .76.1..Accept: *	
Networking		0080 2F 2A 0D 0A 0D 0A HTTP/1.0 200 OK.	
Language		0000 48 54 54 50 2F 31 2E 30 20 32 30 30 20 4F 4B 0D .	
Utils		0016 0A Server: SimpleHT	
Date / Time		0000 53 65 72 76 65 72 3A 20 53 69 6D 70 6C 65 48 54 TP/0.6 Python/2.	
Extractors		0016 54 50 2F 30 2E 36 20 50 79 74 68 6F 6E 2F 32 2E 7.17..	
Compression		0032 37 2E 31 37 0D 0A Date: Sun, 22 Au	
Hashing		0000 44 61 74 65 3A 20 53 75 6E 2C 20 32 32 20 41 75 g 2021 10:25:31	
Code tidy		0016 67 20 32 30 32 31 20 31 30 3A 32 35 3A 33 31 20 GMT..	
Forensics		0032 47 4D 54 0D 0A Content-type: im	
Multimedia		0000 43 6F 6E 74 65 6E 74 2D 74 79 70 65 3A 20 69 6D age/png..	
Other		0016 61 67 65 2F 70 6E 67 0D 0A Content-Length:	
Flow control		0000 43 6F 6E 74 65 6E 74 2D 4C 65 6E 67 74 68 3A 20 311475..	
		0016 33 31 31 34 37 35 0D 0A Last-Modified: S	
		0000 4C 61 73 74 2D 4D 6F 64 69 66 69 65 64 3A 20 53 un, 22 Aug 2021	
		0016 75 6E 2C 20 32 32 20 41 75 67 20 32 30 32 31 20 10:25:26 GMT..	
		0032 31 30 3A 32 35 3A 32 36 20 47 4D 54 0D 0A ..	
		0000 0D 0A .PNG.....IHDR	
		0000 89 50 4E 47 0D 0A 1A 0A 00 00 00 0D 49 48 44 52 .....	
		0016 00 00 02 08 00 00 02 0C 08 06 00 00 00 90 6D 54 .....mT	
		0032 75 00 00 0A E0 69 43 43 50 49 43 43 20 50 72 6F u....iCCPICC Pro	
		0048 66 69 6C 65 00 00 48 89 95 97 07 54 53 59 1A 80 file..H....TSY..	
		0064 EF 7B E9 8D 00 09 11 90 12 7A 13 A4 13 40 4A E8 .{.....z...@J.	
		Output	
		GET /evil.png HTTP/1.1	
		Host: 192.168.222.22	
		User-Agent: curl/7.76.1	
		Accept: */*	
		HTTP/1.0 200 OK	
		Server: SimpleHTTP/0.6 Python/2.7.17	
		Date: Sun, 22 Aug 2021 10:25:31 GMT	
		Content-type: image/png	
		Content-Length: 311475	
		Last-Modified: Sun, 22 Aug 2021 10:25:26 GMT	
		.PNG	
		...	
		IHDR.....mTu..	
		àïCCPICC Profile..H....TSY..i{é.. ..z.ø.@Jè.Æ%\$....bBP°	
		+.#8*... 2.E"	
		.....±EP.Ü..	
		È:X*î.XÄiüÿÿ;+ÿiüiçpsp.%.+øÆÉ.d.³øáp^iQ_x&n.à.6	
		..çòd.vhhødzp%ç. ü¶Ä. .yp_E./..ñ....NÄEx..."ä.0"i..u.N.,E.Lò.érSÄ..'8e	
		.çLpðSÈ.'m"Ä%6..0ærW)....=3...ø!."l-æ.Ä."E0.'a0.Fä.9....<	
		(°)b/.BG..ö').ñ.øðiÄ!(x®@iÄü.d.tñÿ+5ÿ[20ä01..A.J.Ä...iB% 'i .....L³.?i?	
		ÉBy@040dpñ0Içü.)0 /.æd..Gä'..9I..oA4K3A.±.Wp1iæJgäE0e.zj.E0.+...æ\Q0.i.	
		WE.I0x+0Ry, ". .0Bk&0.	
		ç0.Ü.è.q.k³....Ü±3ü.Äi..²XEn . .i.M.Ä^,â%.%I.U0.ÿ.zYv.bm.r8g0.*00...:iÄ.ø.	
		'âa.('...	

You may want to apply an operation from the left column. One option is to use the **Extract Files** operation. If you choose this option, you may want to specify certain file types for extraction. In this case, let's instead remove the client

HTTP headers using the Strip HTTP headers operation.

The screenshot displays the CyberChef interface with the following components:

- Operations Panel (Left):** A list of operations including 'Strip', 'Strip HTML tags', 'Strip HTTP headers' (highlighted), 'Microsoft Script Decoder', 'Favourites', 'Data format', 'Encryption / Encoding', 'Public Key', 'Arithmetic / Logic', 'Networking', 'Language', 'Utils', 'Date / Time', 'Extractors', 'Compression', 'Hashing', 'Code tidy', 'Forensics', 'Multimedia', 'Other', and 'Flow control'.
- Recipe Panel (Top Left):** Shows a sequence of operations: 'From Hexdump' and 'Strip HTTP headers'.
- Input Panel (Top Right):** Displays the raw hex and ASCII data of an HTTP GET request for /evil.png. The ASCII part shows headers like 'Server: SimpleHTTP/0.6 Python/2.7.17', 'Date: Sun, 22 Aug 2021 10:25:31 GMT', 'Content-type: image/png', and 'Content-Length: 311475'.
- Output Panel (Bottom Right):** Shows the result after stripping the headers, displaying the raw PNG data in hex and ASCII format.
- Bottom Bar:** Includes a 'STEP' button, a 'BAKE!' button with a chef icon, and an 'Auto Bake' checkbox.

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. Here, CyberChef is automatically applying Strip HTTP headers again to remove the web server HTTP headers and then rendering the actual PNG image.

Operations

strip

Strip HTML tags

Strip HTTP headers

Microsoft Script Decoder

Favourites

Data format

Encryption / Encoding

Public Key

Arithmetic / Logic

Networking

Language

Utils

Date / Time

Extractors

Compression

Hashing

Code tidy

Forensics

Multimedia

Other

Flow control

Recipe

From Hexdump

Strip HTTP headers

Strip HTTP headers

Render Image

Input format

Raw

Input


length: 1448550  
lines: 19601

0000 47 45 54 20 2F 65 76 69 6C 2E 70 6E 67 20 48 54 GET /evil.png HT  
0016 54 50 2F 31 2E 31 0D 0A 48 6F 73 74 3A 20 31 39 TP/1.1..Host: 19  
0032 32 2E 31 36 38 2E 32 32 32 2E 32 32 0D 0A 55 73 2.168.222.22..Us  
0048 65 72 2D 41 67 65 6E 74 3A 20 63 75 72 6C 2F 37 er-Agent: curl/7  
0064 2E 37 36 2E 31 0D 0A 41 63 63 65 70 74 3A 20 2A .76.1..Accept: \*  
0080 2F 2A 0D 0A 0D 0A /\*....  
0000 48 54 54 50 2F 31 2E 30 20 32 30 30 20 4F 4B 0D HTTP/1.0 200 OK.  
0016 0A .  
0000 53 65 72 76 65 72 3A 20 53 69 6D 70 6C 65 48 54 Server: SimpleHT  
0016 54 50 2F 30 2E 36 20 50 79 74 68 6F 6E 2F 32 2E TP/0.6 Python/2.  
0032 37 2E 31 37 0D 0A 7.17..  
0000 44 61 74 65 3A 20 53 75 6E 2C 20 32 32 20 41 75 Date: Sun, 22 Au  
0016 67 20 32 30 32 31 20 31 30 3A 32 35 3A 33 31 20 g 2021 10:25:31  
0032 47 4D 54 0D 0A GMT..  
0000 43 6F 6E 74 65 6E 74 2D 74 79 70 65 3A 20 69 6D Content-type: im  
0016 61 67 65 2F 70 6E 67 0D 0A age/png..  
0000 43 6F 6E 74 65 6E 74 2D 4C 65 6E 67 74 68 3A 20 Content-Length:  
0016 33 31 31 34 37 35 0D 0A 311475..  
0000 4C 61 73 74 2D 4D 6F 64 69 66 69 65 64 3A 20 53 Last-Modified: S  
0016 75 6E 2C 20 32 32 20 41 75 67 20 32 30 32 31 20 un, 22 Aug 2021  
0032 31 30 3A 32 35 3A 32 36 20 47 4D 54 0D 0A 10:25:26 GMT..  
0000 0D 0A ..  
0000 89 50 4E 47 0D 0A 1A 0A 00 00 00 0D 49 48 44 52 .PNG.....IHDR  
0016 00 00 02 08 00 00 02 0C 08 06 00 00 00 90 6D 54 .....MT  
0032 75 00 00 0A E0 69 43 43 50 49 43 43 20 50 72 6F u.....iCCPICC Pro  
0048 66 69 6C 65 00 00 48 89 95 97 07 54 53 59 1A 80 file..H....TSY..  
0064 EF 7B E9 8D 00 09 11 90 12 7A 13 A4 13 40 4A E8 .{.....Z...@J.

Output

time: 202ms  
length: 311475  
lines: 1125

INSERT WITTY SECURITY ONION  
REFERENCE HERE



STEP

BAKE!

Auto Bake

6.11.4 More Information

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

6.11. CyberChef

129

## 6.12 Playbook

### 6.12.1 Overview

Playbook is a web application available for installation on Manager nodes. 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 a particular detection strategy.

The screenshot displays the 'DETECTION PLAYBOOKS' interface. At the top, there's a navigation bar with 'Home', 'Logged in as analyst', 'My account', and 'Sign out'. Below this is a search bar and a dropdown menu set to 'Detection Playbooks'. The main section is titled 'Playbook' and includes a 'Filters' section with 'Status' (set to 'open') and 'Options'. A table lists various plays, each with a checkbox, ID, status, level, playbook name, product, title, and update date. The table is paginated, showing items 1 through 13. On the right, a sidebar lists 'Custom queries' including 'All Plays', 'Disabled Plays', 'Draft Plays', 'Playbook - Community Sigma', and 'Playbook - Internal'.

#	Status	Level	Playbook	Product	Title	Updated
623	Draft	medium	community	windows	Harvesting of Wifi Credentials Using netsh.exe	05/13/2020 02:07 PM
622	Draft	medium	community	windows	Advanced IP Scanner	05/13/2020 02:07 PM
621	Draft	high	imported	windows	Whoami Execution	05/13/2020 02:05 PM
620	Draft	medium	imported	osquery	New Sensitive Shared Resource	05/13/2020 01:30 PM
618	Inactive	medium	community	windows	XSL Script Processing	05/03/2020 10:00 AM
617	Draft	high	community	windows	Wareset UAC Bypass	05/01/2020 08:58 PM
616	Draft	high	community	windows	Microsoft Workflow Compiler	05/01/2020 08:57 PM
615	Draft	critical	community	windows	Wimpvise Spawning Process	05/01/2020 08:57 PM
614	Draft	high	community	windows	WMI Spawning Windows PowerShell	05/01/2020 08:57 PM
613	Draft	high	community	windows	WMI Persistence - Script Event Consumer	05/01/2020 08:57 PM
612	Draft	critical	community	windows	WMI Backdoor Exchange Transport Agent	05/01/2020 08:57 PM
611	Draft	high	community	windows	Windows 10 Scheduled Task SandboxEscaper 0-day	05/01/2020 08:57 PM
610	Draft	high	community	windows	Run Whoami as SYSTEM	05/01/2020 08:57 PM
609	Draft	high	community	windows	Shells Spawned by Web Servers	05/01/2020 08:57 PM
608	Draft	high	community	windows	Webshell Detection With Command Line Keywords	05/01/2020 08:57 PM
607	Draft	medium	community	windows	Java Running with Remote Debugging	05/01/2020 08:57 PM
606	Draft	high	community	windows	Possible Privilege Escalation via Weak Service Permissions	05/01/2020 08:57 PM
605	Draft	high	community	windows	Bypass UAC via WSRset.exe	05/01/2020 08:57 PM
604	Draft	high	community	windows	Bypass UAC via Fodhelper.exe	05/01/2020 08:57 PM
603	Draft	high	community	windows	Bypass UAC via CMSTP	05/01/2020 08:57 PM
602	Draft	medium	community	windows	Domain Trust Discovery	05/01/2020 08:57 PM
601	Draft	high	community	windows	Terminal Service Process Spawn	05/01/2020 08:57 PM
600	Draft	high	community	windows	Tasks Folder Evasion	05/01/2020 08:57 PM
599	Draft	medium	community	windows	Tap Installer Execution	05/01/2020 08:57 PM
598	Draft	high	community	windows	System File Execution Location Anomaly	05/01/2020 08:57 PM

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

## 6.12.2 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>.

## 6.12.3 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.

## 6.12.4 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.



### 6.12.5 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.

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**.

### 6.12.6 Viewing Playbook Alerts

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

### 6.12.7 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>).

### 6.12.8 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.



### 6.12.9 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.

### 6.12.10 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.

### 6.12.11 Log Sources and Field Names

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

- *osquery*
- network (via *Zeek* logs)
- Windows Eventlogs and *Sysmon* (shipped with *osquery* or *winglobeat*)

The pre-loaded Plays depend on Sysmon and Windows Eventlogs shipped with *winlogbeat* or *osquery*.

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>

### 6.12.12 .Security subfield

Playbook uses the `.security` subfield that is generated by a special analyzer ([https://github.com/neu5ron/es\\_stk](https://github.com/neu5ron/es_stk)). This analyzer allows case insensitive wildcard searches and is designed specifically for security logs.

### 6.12.13 Adding Additional Rulesets

As previously mentioned, 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 pulled in by editing a pillar value under `/opt/so/saltstack/local/pillar/global.sls`:

```
soctopus:
  playbook:
    rulesets:
      - windows
```

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.

Next, restart SOctopus:

```
so-soctopus-restart
```

Finally, tell Playbook to pull in the new rules:

```
so-playbook-ruleupdate
```

This can take a few minutes to complete if pulling in a large amount of new rules.

### 6.12.14 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
```

### 6.12.15 More Information

---

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

---

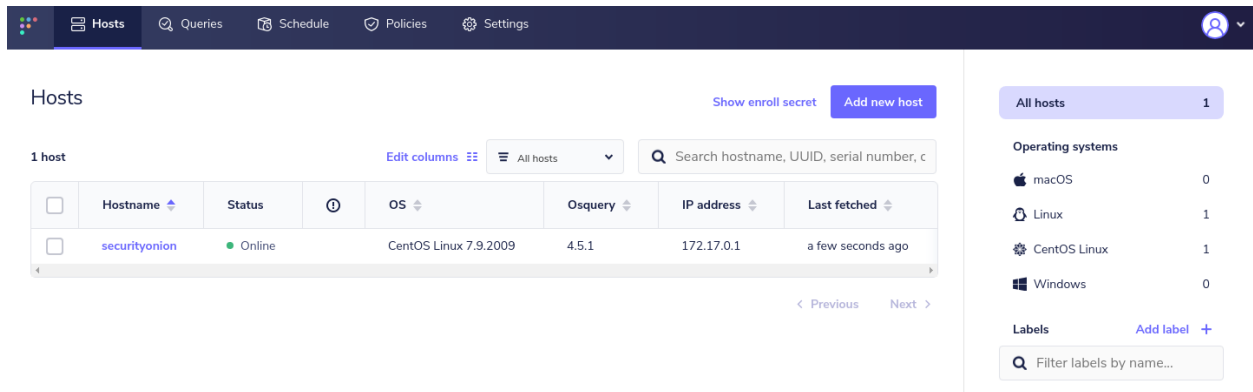
## 6.13 FleetDM

From <https://fleetdm.com/>:

Ask questions about your servers, containers, and laptops running Linux, Windows, and macOS. Quickly deploy osquery and scale your fleet to 50,000+ devices on top of a stable core technology.

### 6.13.1 Usage

If you enabled Fleet during setup, you can then login to Fleet using the email address and password that you entered during the installer. You can edit the password or add a new Fleet user within Fleet itself.



Custom *osquery* packages are generated during setup and you can find them on the [Downloads](#) page in *Security Onion Console (SOC)*. Before you install a package on an endpoint, use *so-allow* on your manager node to configure the firewall to allow inbound *osquery* connections.

## 6.13.2 Configuration

Fleet configuration can be found in `/opt/so/conf/fleet/`. However, please keep in mind that if you make any changes to this directory they may be overwritten since the configuration is managed with *Salt*.

## 6.13.3 Diagnostic Logging

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

```
sudo docker logs so-fleet
```

## 6.13.4 fleetctl

`fleetctl` is a command-line utility that allows you to manage your Fleet instance and run live queries from the cli.

If using `fleetctl` from the Manager and Fleet is enabled on the Manager, first set the `fleetctl` login configuration:

```
./fleetctl config set --address https://localhost:8080 --url-prefix fleet --tls-skip-verify
```

Then login using a valid username and password:

```
./fleetctl login
```

`fleetctl` documentation can be found at <https://fleetdm.com/docs/using-fleet/fleetctl-cli>.

### 6.13.5 Adding Query Packs

#### CLI

You can bulk add queries and packs to FleetDM using `fleetctl`. `/opt/so/conf/fleet/packs/` is mapped to the FleetDM container, so you can drop your query packs there and then reference them like this:

```
sudo docker exec -it so-fleet fleetctl apply -f /packs/<yourpack>.yaml
```

#### GUI

Single queries and packs can also be added via the GUI in FleetDM. Go to the **Queries** section, click the **Create New Query** button, and add the query. Test the live query, verify that it works correctly, and save it. You can then either add it to an existing pack or create a new pack. Do this by going to the **Schedule** section, choosing **Advanced**, and using the options there. Your query will not work if it isn't also scheduled, so set a short interval for testing, then update to a more sensible gap for production.

### 6.13.6 More Information

---

**Note:** For more information about `osquery`, please see the *osquery* section.

For more information about Fleet, please see <https://fleetdm.com/>.

---

## 6.14 ATT&CK Navigator

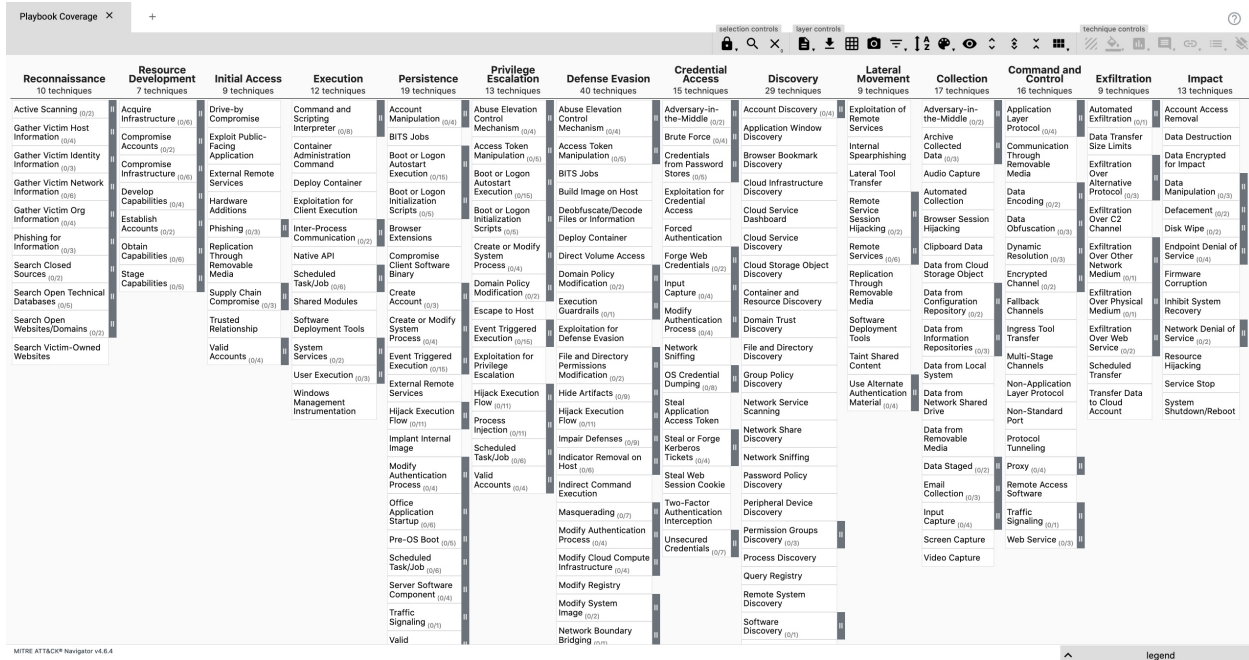
From <https://github.com/mitre-attack/attack-navigator>:

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.

## 6.14.1 Accessing

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



## 6.14.2 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.

## 6.14.3 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*.

## 6.14.4 More Information

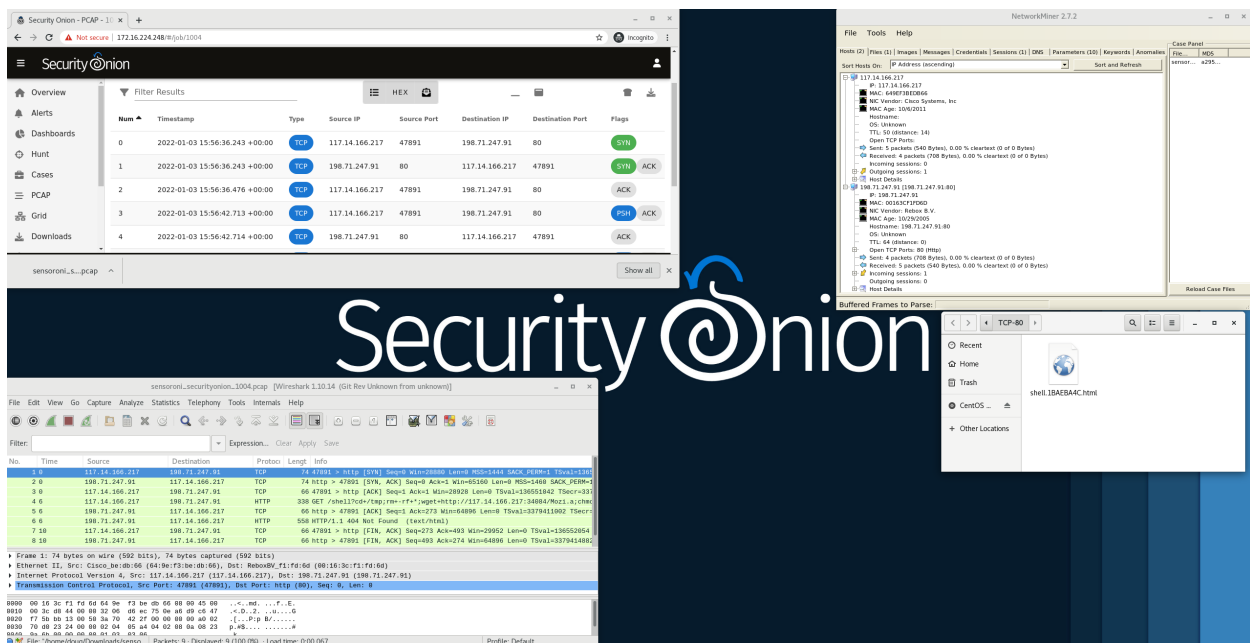
### Note:

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



## ANALYST VM

Full-time analysts may want to create a dedicated Analyst VM. This is a full GNOME desktop environment including Chromium web browser, *NetworkMiner*, *Wireshark*, and other analyst tools. This allows you to investigate pcaps and other potentially malicious artifacts without impacting your Security Onion deployment or your workstation.



**Note:** The Analyst desktop currently only supports CentOS, so you'll either need to use our Security Onion ISO image (recommended) or a manual installation of CentOS 7.

### Installation

There are a few different ways to install the Analyst desktop:

- Our Security Onion ISO image includes a boot menu option for Analyst installs that will partition your disk appropriately and immediately perform an Analyst installation.
- If you're doing a network installation on CentOS 7 (NOT using our Security Onion ISO image), then in our normal Setup wizard you can choose **OTHER** and then choose **ANALYST**.
- You can manually run `sudo so-analyst-install`. This script 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 VM itself.

### Joining to Grid

You can optionally join your Analyst 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 Analyst installation to connect. Please note that Analyst installations won't actually display on the *Grid* page.

If you choose not to join your Analyst installation to your grid, then you may need to run *so-allow* on the manager node and choose the *analyst* option to allow the traffic through the host-based *Firewall*.

### Disabling

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

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

## 7.1 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.

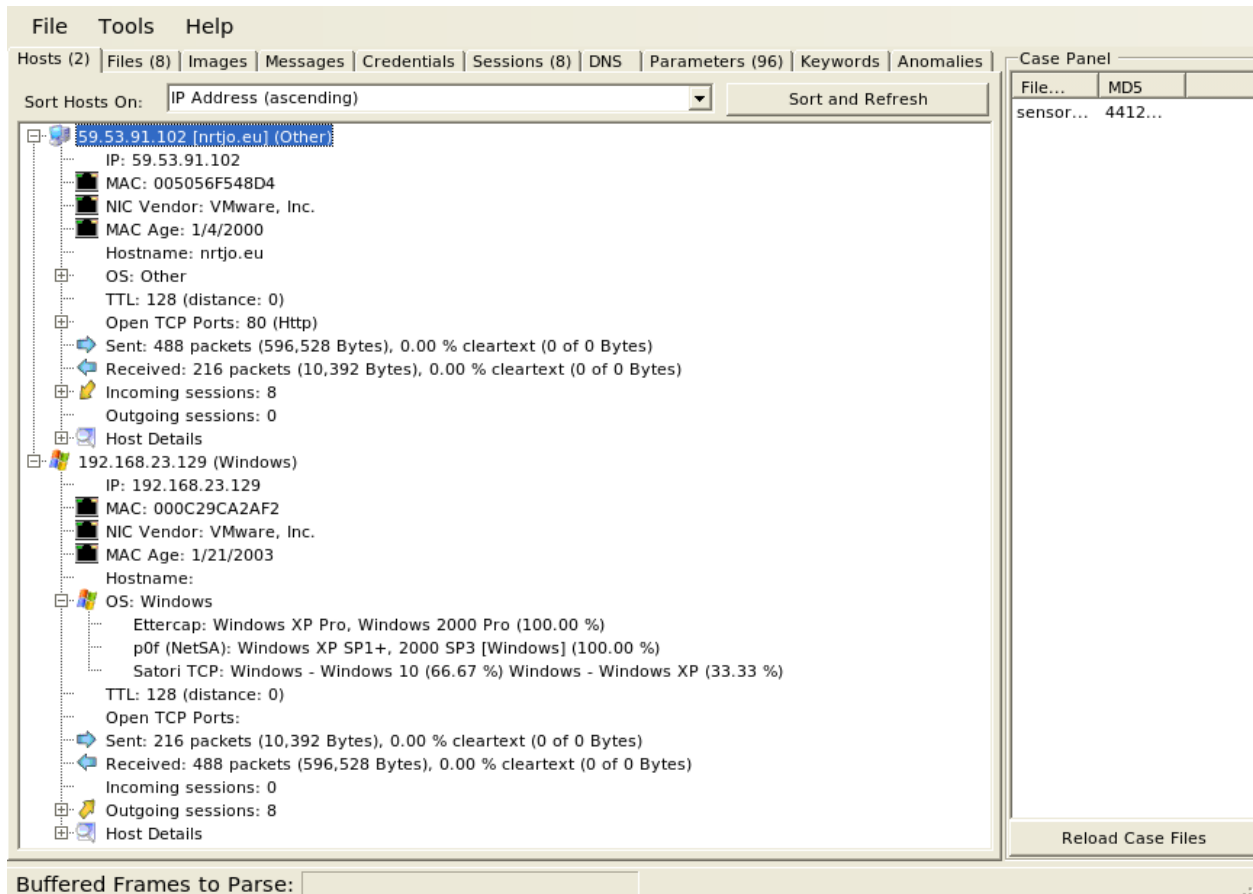
NetworkMiner makes it easy to perform advanced Network Traffic Analysis (NTA) by providing extracted artifacts in an intuitive user interface. The way data is presented not only makes the analysis simpler, it also saves valuable time for the analyst or forensic investigator.

### 7.1.1 Usage

NetworkMiner is a part of our *Analyst VM* installation. `so-analyst-install` automatically registers NetworkMiner as a pcap handler, so if you download a pcap file from the *PCAP* interface, you can simply click on it to open in NetworkMiner.




## 7.1.2 Screenshot




## 7.1.3 Example

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.



### 7.1.4 More Information

**Note:**

For more information about NetworkMiner, please see:

<https://www.netresec.com/?page=networkminer>

## 7.2 Wireshark

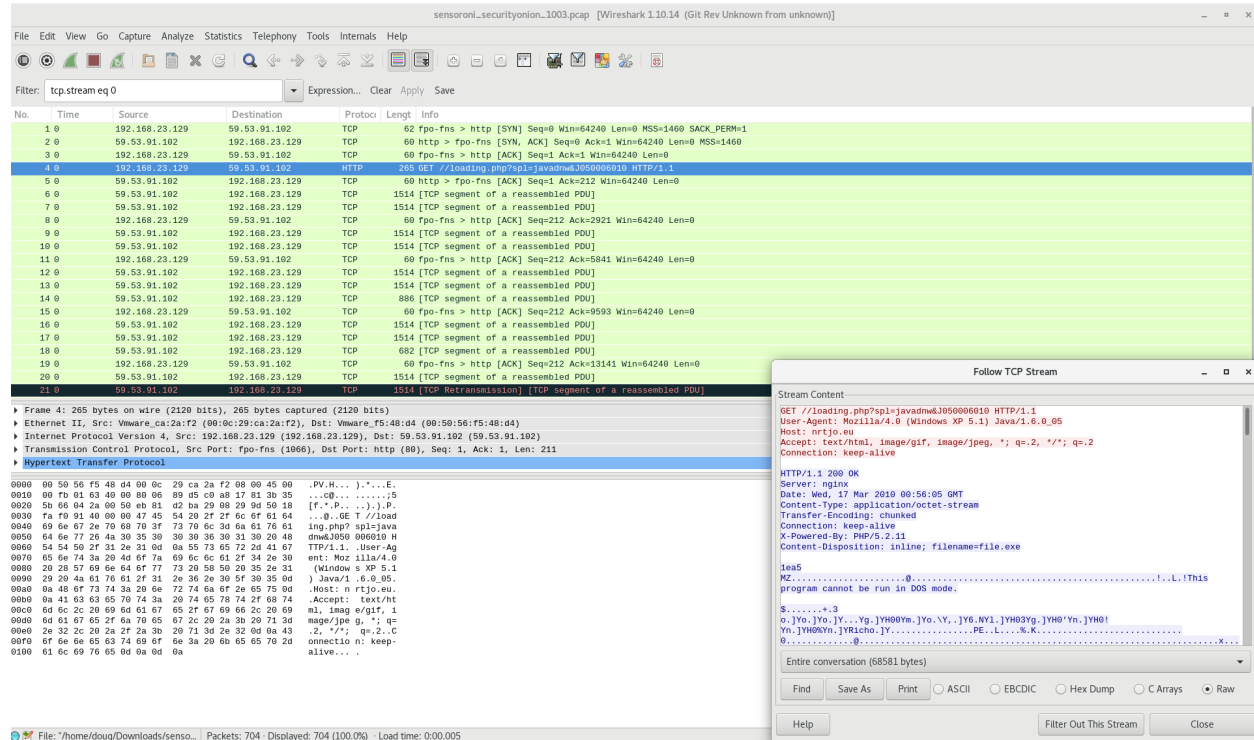
From <https://www.wireshark.org/>:

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.

## 7.2.1 Usage

Wireshark is a part of our *Analyst VM* installation.

## 7.2.2 Screenshot



## 7.2.3 Example

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.

File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help

Filter: Expression... Clear Apply Save

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	192.168.222.1	192.168.222.22	TCP	74	50560 → http [SYN] Seq=0 Win=65535 Len=0 MSS=1460 WS=128 SACK_PERM=1 TSval=1291891406 TSecr=0
2	0.000000	192.168.222.22	192.168.222.1	TCP	74	http → 50560 [SYN, ACK] Seq=0 Ack=1 Win=65535 Len=0 MSS=1460 SACK_PERM=1 TSval=151780193 TSecr=1291891406 WS=128
3	0.000000	192.168.222.1	192.168.222.22	TCP	66	50560 → http [ACK] Seq=1 Ack=1 Win=65792 Len=0 TSval=1291891407 TSecr=151780193
4	0.000000	192.168.222.1	192.168.222.22	HTTP	152	GET /evil.png HTTP/1.1
5	0.000000	192.168.222.22	192.168.222.1	TCP	66	http → 50560 [ACK] Seq=1 Ack=87 Win=65152 Len=0 TSval=151780194 TSecr=1291891407
6	0.000000	192.168.222.22	192.168.222.1	TCP	83	[TCP segment of a reassembled PDU]
7	0.000000	192.168.222.1	192.168.222.22	TCP	66	50560 → http [ACK] Seq=87 Ack=18 Win=65792 Len=0 TSval=1291891410 TSecr=151780196
8	0.000000	192.168.222.22	192.168.222.1	TCP	104	[TCP segment of a reassembled PDU]
9	0.000000	192.168.222.1	192.168.222.22	TCP	66	50560 → http [ACK] Seq=87 Ack=56 Win=65792 Len=0 TSval=1291891410 TSecr=151780196
10	0.000000	192.168.222.22	192.168.222.1	TCP	103	[TCP segment of a reassembled PDU]
11	0.000000	192.168.222.1	192.168.222.22	TCP	66	50560 → http [ACK] Seq=87 Ack=93 Win=65792 Len=0 TSval=1291891410 TSecr=151780196
12	0.000000	192.168.222.22	192.168.222.1	TCP	91	[TCP segment of a reassembled PDU]
13	0.000000	192.168.222.1	192.168.222.22	TCP	66	50560 → http [ACK] Seq=87 Ack=118 Win=65792 Len=0 TSval=1291891410 TSecr=151780196
14	0.000000	192.168.222.22	192.168.222.1	TCP	90	[TCP segment of a reassembled PDU]
15	0.000000	192.168.222.1	192.168.222.22	TCP	66	50560 → http [ACK] Seq=87 Ack=142 Win=65792 Len=0 TSval=1291891410 TSecr=151780197
16	0.000000	192.168.222.22	192.168.222.1	TCP	112	[TCP segment of a reassembled PDU]
17	0.000000	192.168.222.1	192.168.222.22	TCP	66	50560 → http [ACK] Seq=87 Ack=188 Win=65792 Len=0 TSval=1291891411 TSecr=151780197

▶ Frame 4: 152 bytes on wire (1216 bits), 152 bytes captured (1216 bits)

▶ Ethernet II, Src: CompuLab\_16:f6:ed (00:01:c0:16:f6:ed), Dst: AeonTec\_4c:94:b8 (00:07:32:4c:94:b8)

▶ Internet Protocol Version 4, Src: 192.168.222.1 (192.168.222.1), Dst: 192.168.222.22 (192.168.222.22)

▶ Transmission Control Protocol, Src Port: 50560 (50560), Dst Port: http (80), Seq: 1, Ack: 1, Len: 86

▶ Hypertext Transfer Protocol

```


0000  00 07 32 4c 94 b8 00 01 c0 16 f6 ed 08 00 45 00  ..2L.....E.
0010  00 8a 00 00 40 00 00 06 fd 04 c0 a8 de 01 c0 a8  ....@.....
0020  de 16 c5 80 00 50 18 31 5d 05 34 93 f7 0b 80 18  ....P.1].4....
0030  02 02 90 3d 00 00 01 01 08 0a 4d 00 b2 cf 09 0b  ...=.N.....
0040  fb 61 47 45 54 20 2f 05 76 09 6c 2e 70 0e 07 20  .aGET /e vil.png
0050  48 54 54 50 2f 31 2e 31 0d 0a 48 6f 73 74 3a 20  HTTP/1.1 .Host:
0060  31 39 32 2e 31 36 38 2e 32 32 32 2e 32 32 0d 0a  192.168. 222.22..
0070  55 73 65 72 2d 41 67 65 6e 74 3a 20 63 75 72 6c  User-Agent: curl
0080  2f 37 2e 37 36 2e 31 0d 0a 41 63 63 65 70 74 3a  /7.76.1. .Accept:
0090  20 2a 2f 2a 0d 0a 0d 0a  ../*.....
  
```

To extract files from HTTP traffic, click File - Export Objects - HTTP.

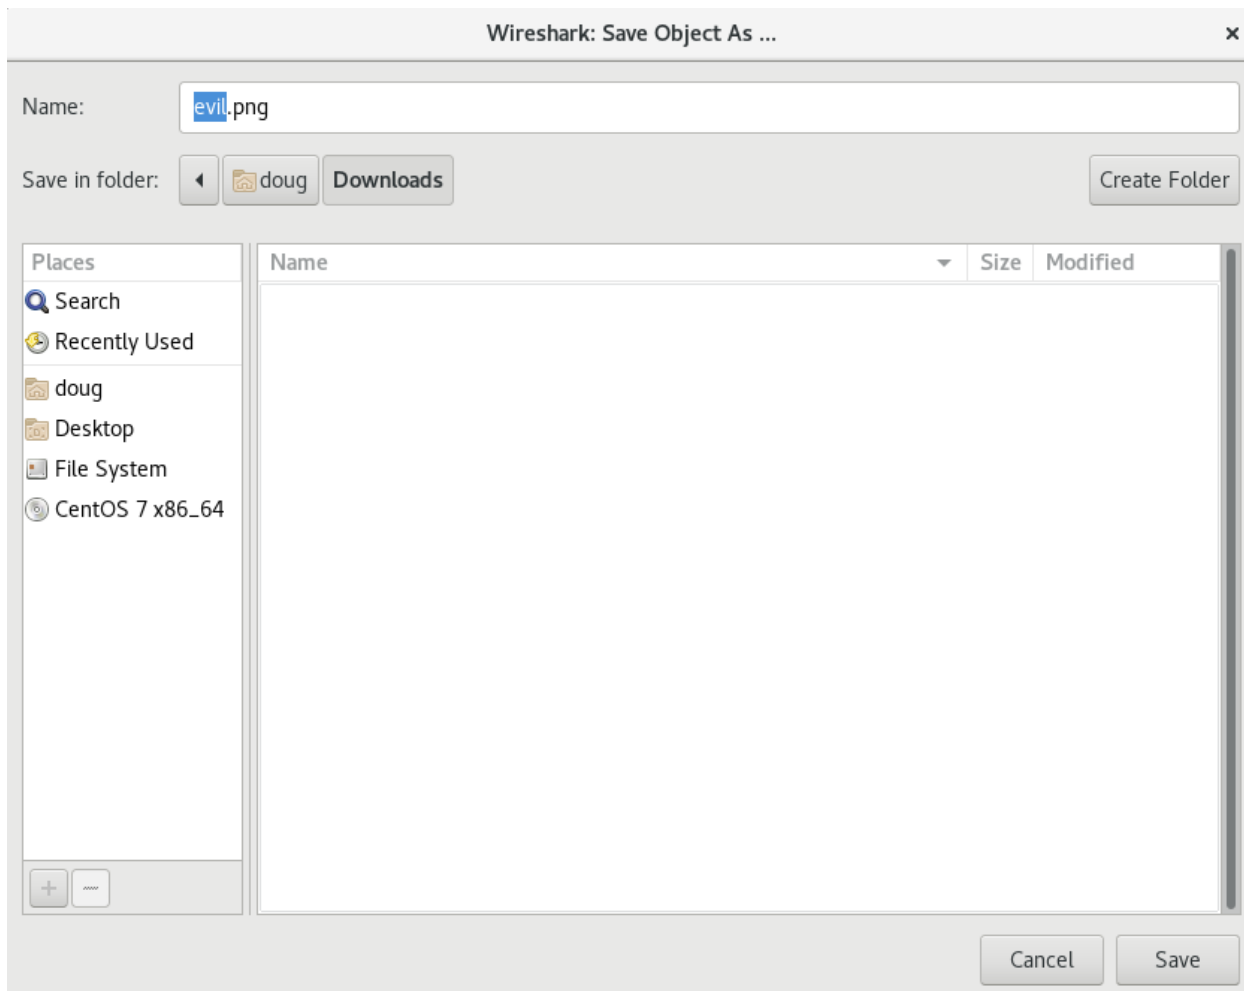
File Edit View Go Capture Analyze Statistics Telephony Tools

- Open... Ctrl+O
- Open Recent
- Merge...
- Import from Hex Dump...
- Close Ctrl+W
- Save Ctrl+S
- Save As... Shift+Ctrl+S
- File Set
- Export Specified Packets...
- Export Packet Dissections
- Export Selected Packet Bytes... Ctrl+H
- Export SSL Session Keys...
- Export Objects**
  - HTTP**
  - DICOM
  - SMB/SMB2
- Print... Ctrl+P
- Quit Ctrl+Q


Select the file(s) to save.



Specify where to save them.



Review the extracted file(s).



### 7.2.4 More Information

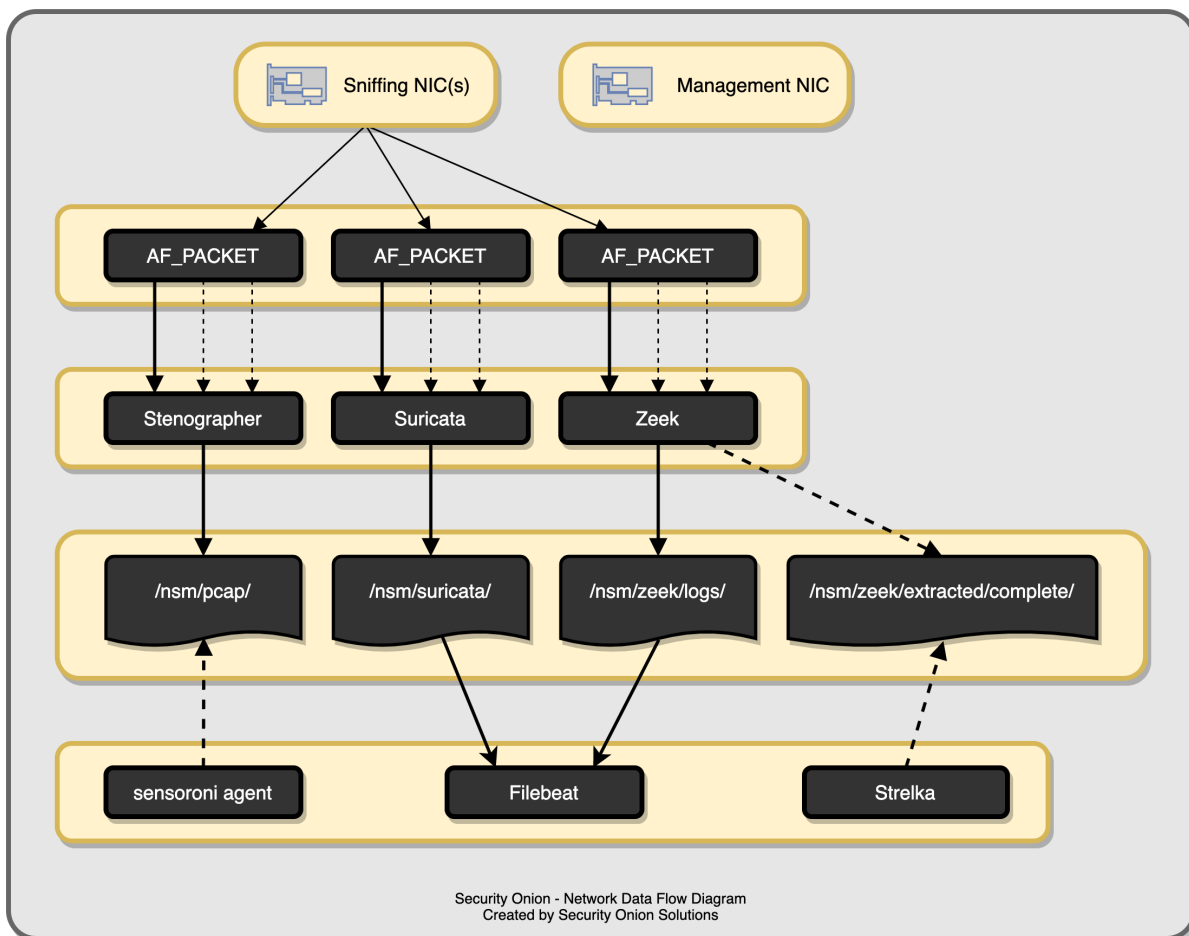
**Note:** For more information about Wireshark, please see <https://www.wireshark.org/>.





## 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

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 packet acquisition:

- *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](https://www.kernel.org/doc/Documentation/networking/packet_mmap.txt)

---

## 8.2 Stenographer

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.

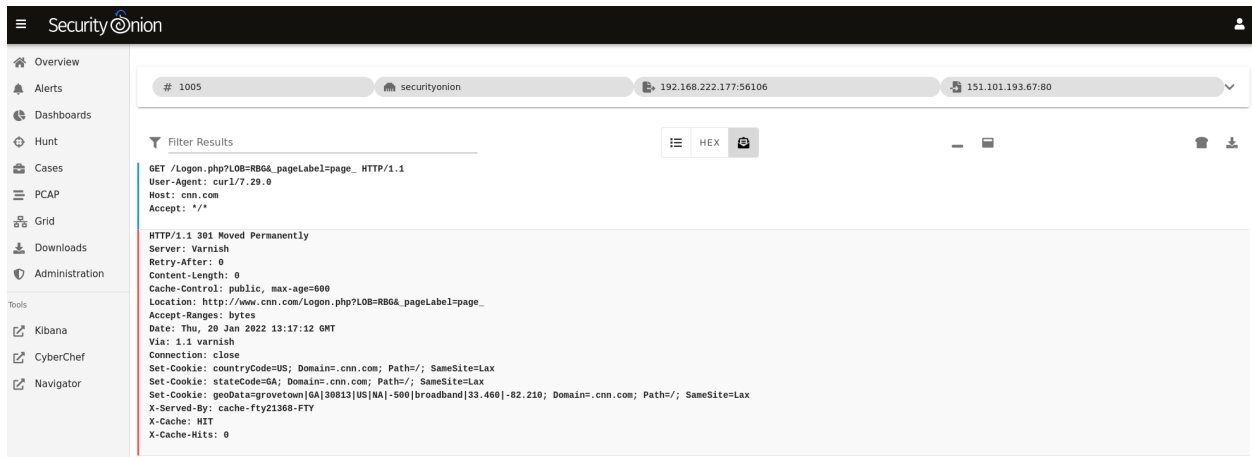
Stenographer uses *AF-PACKET* for packet acquisition.

## 8.2.1 Output

Stenographer writes full packet capture to `/nsm/pcap/`. It will automatically start purging old data once the partition reaches 90%. This value is configurable as shown in the Configuration section 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 steno query as defined at <https://github.com/google/stenographer#querying>. In the following examples, replace “YourStenoQueryHere” with your actual steno 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

Stenographer reads its configuration from `/opt/so/conf/steno/`. However, please keep in mind that if you make any changes to this directory they may be overwritten since the configuration is managed with *Salt*.

For example, suppose you want to change the default value for purging old pcap. You could add the `diskfreepercentage` option in the `steno` section of the *Salt* pillar and set the value to something appropriate for your system. For example:

```
steno:
  diskfreepercentage: 20
```

### 8.2.5 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 add the `maxfiles` option in the `steno` section of the *Salt* pillar and set the value to something appropriate for your system. For example:

```
steno:
  maxfiles: 120000
```

### 8.2.6 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
```

### 8.2.7 Disabling

If you need to disable Stenographer, you can do so in two different ways. If you just want to disable it on a single sensor, then you can edit that sensor's `minion.sls` file. If the file doesn't already have a `steno` section, then add the following to the end of the file:

```
steno:
  enabled: false
```

If you want to disable Stenographer globally across all your sensors, then you can add that entry to your `global.sls` file.

### 8.2.8 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>

8.2.9 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*:

Security Onion

Overview

Alerts

Dashboards

Hunt

Cases

PCAP

Grid

Downloads

Administration

Tools

Kibana

CyberChef

Navigator

Alerts

Options

Total Found: 102

Q

Group By Name, Module

2021/06/30 00:00:00 AM - 2021/07/01 00:00:00

REFRESH

Choose the timespan to search, or click the calendar icon to switch to relative time

Fetch Limit 50

Filter Results

	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

Version: 2.3.300

© 2024 Security Onion Solutions, LLC

Terms and Conditions

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 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.

To change the number of workers:

- Stop sensor processes:

```
sudo so-suricata-stop
```

- Edit `/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls` and change the `suriprocs` variable to the desired number of workers.
- Start sensor processes:

```
sudo so-suricata-start
```

---

**Note:**

For other tuning considerations, please see:

<https://suricata.readthedocs.io/en/latest/performance/tuning-considerations.html>

---

For best performance, Suricata should be pinned to specific CPUs. In most cases, you'll 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](https://github.com/brokenscripts/cpu_pinning)

---

To pin Suricata workers to specific CPUs:

- Stop sensor processes:

```
sudo so-suricata-stop
```

- Edit `/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls` and add the following under `sensor`:

```
suripins:  
- <cpu_1>  
- <cpu_2>  
- <cpu_3>
```

- Note: To avoid inconsistent Suricata workers being allocated, ensure `suriprocs` is removed from under `sensor`: or is equivalent to the number of cpu cores being pinned.
- Start sensor processes:

```
sudo so-suricata-start
```

### 8.3.3 HOME\_NET

To configure HOME\_NET, please see the *Homenet* section.

### 8.3.4 EXTERNAL\_NET

By default, EXTERNAL\_NET is set to any (which includes HOME\_NET) to detect lateral movement inside your environment.

### 8.3.5 Configuration

You can configure Suricata's `suricata.yaml` using *Salt*. The defaults for this have been defined in <https://github.com/Security-Onion-Solutions/securityonion/blob/master/salt/suricata/defaults.yaml>. Under `suricata:config`, the pillar structure follows the same YAML structure of the `suricata.yaml` file.

For example, suppose you want to change Suricata's EXTERNAL\_NET setting from the default of any to !\$HOME\_NET. You could add the following to the global pillar file (`/opt/so/saltstack/local/pillar/global.sls`) or minion pillar file (`/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls`) on the manager:

```
suricata:
  config:
    vars:
      address-groups:
        EXTERNAL_NET: "!$HOME_NET"
```

From the manager, then run:

```
sudo salt $SENSORNAME_$ROLE state.highstate
```

---

**Note:** Check out our Suricata Variables video at <https://youtu.be/8TfnocUuQnM!>

---

Some of the settings normally found in `suricata.yaml` can be found in the sensor pillar instead of the Suricata pillar. These options are: HOMENET, default-packet-size, and the CPU affinity settings for pinning the processes to CPU cores or how many processes to run.

If you would like to configure/manage IDS rules, please see the *Managing Rules* and *Managing Alerts* sections.

### 8.3.6 Thresholding

To enable thresholding for SIDS, reference the example pillar at <https://github.com/Security-Onion-Solutions/securityonion/blob/master/pillar/thresholding/pillar.example>.

To view the acceptable syntax, view the file located at <https://github.com/Security-Onion-Solutions/securityonion/blob/master/pillar/thresholding/pillar.usage>.

This pillar can be added to *Salt* in either the global pillar file (`/opt/so/saltstack/local/pillar/global.sls`) or minion pillar file (`/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls`).

**Warning:** Salt sls files are in YAML format. When editing these files, please be very careful to respect YAML syntax, especially whitespace. For more information, please see [https://docs.saltproject.io/en/latest/topics/troubleshooting/yaml\\_idiosyncrasies.html](https://docs.saltproject.io/en/latest/topics/troubleshooting/yaml_idiosyncrasies.html).

### 8.3.7 Metadata

Depending on what options you choose in Setup, it may ask if you want to use *Zeek* or *Suricata* for metadata. If you choose *Suricata* and later find that some 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>.

The global pillar on your manager node controls the metadata engine on each sensor. Only one metadata engine at a time is supported.

To change your grid's metadata engine from Zeek to Suricata:

- On the manager, edit the global pillar file (`/opt/so/saltstack/local/pillar/global.sls`) and change the `mdengine` variable from `ZEEK` to `SURICATA`.

- Stop Zeek on all nodes:

```
sudo salt '*' cmd.run 'so-zeek-stop'
```

- Update all nodes:

```
sudo salt '*' state.highstate
```

- You may need to remove `so-zeek` from `/opt/so/conf/so-status/so-status.conf` on each sensor node.

### 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 Disabling

Suricata can be disabled by setting `enabled: false` in the `suricata` *Salt* pillar.

If you just want to disable Suricata on a single sensor, then you can edit that sensor's `minion.sls` file. If the file doesn't already have a `suricata` section, then add the following to the end of the file:

```
suricata:
  enabled: false
```

If you want to disable Suricata globally across all your sensors, then you could add that entry to your `global.sls` file.



### 8.3.10 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
```

### 8.3.11 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 *Homenet* 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 *Filebeat*, *Logstash*, and *Elasticsearch* logs for any pipeline issues that may be preventing the alerts from being written to the `so-ids` indices.
- 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.

### 8.3.12 Stats

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

### 8.3.13 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.

8.3.14 More Information

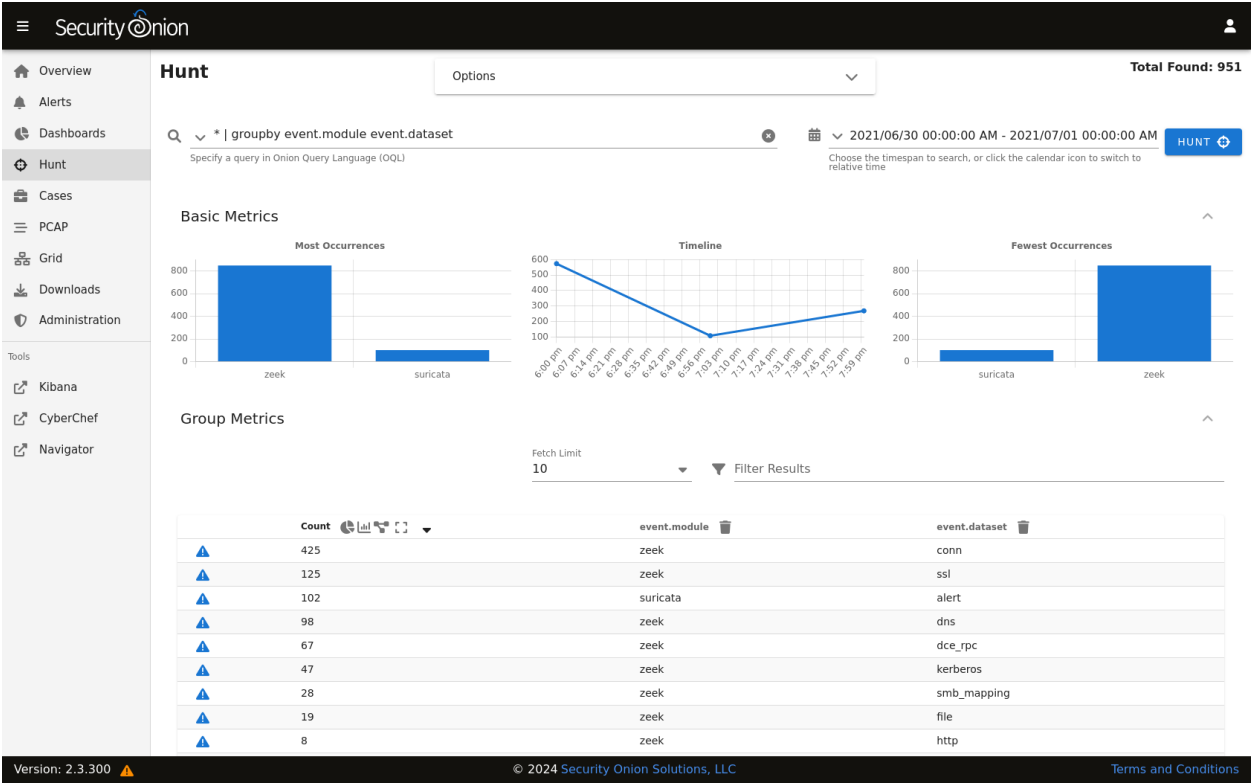
**Note:** For more information about Suricata, please see <https://suricata.io>.

8.4 Zeek

Zeek is formerly known as Bro. 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 conn (connection) 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 [Grafana](#). 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 Performance

Zeek uses [AF-PACKET](#) so that you can spin up multiple Zeek workers to handle more traffic.

To change the number of AF-PACKET workers for [Zeek](#):

- Stop Zeek:

```
sudo so-zeek-stop
```

- Edit `/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls` and change the `zeek_lbprocs` variable to the desired number of cores.

- Start Zeek:

```
sudo so-zeek-start
```

For best performance, Zeek should be pinned to specific CPUs. In most cases, you'll 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](https://github.com/brokenscripts/cpu_pinning).

---

To pin Zeek workers to specific CPUs:

- Stop sensor processes:

```
sudo so-zeek-stop
```

- Edit `/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls` and add the following under `sensor`:

```
zeek_pins:
  - <cpu_1>
  - <cpu_2>
  - <cpu_3>
```

- Note: To avoid inconsistent Zeek workers being allocated, ensure `zeek_lbprocs` is removed from under `sensor`: or is equivalent to the number of cpu cores being pinned.
- Start sensor processes:

```
sudo so-zeek-start
```

### 8.4.4 Syslog

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

### 8.4.5 Logs

Zeek logs are stored in `/nsm/zeek/logs`. They are collected by *Filebeat*, 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*.

If you want to specify what Zeek logs are ingested, you can use *so-zeek-logs*.

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.zeeb.org/en/latest/scripts/base/protocols/ssl/main.zeeb.html#type-SSL::Info>

## notice.log

- Zeek notices
- For more information, see:

<https://docs.zeeb.org/en/latest/scripts/base/frameworks/notice/main.zeeb.html#type-Notice::Info>

## Other Zeek logs

Zeek also provides other logs by default and you can read more about them at <https://docs.zeeb.org/en/latest/script-reference/log-files.html>.

Security Onion includes protocol analyzers for STUN, TDS, and Wireguard traffic and several different ICS/SCADA protocols. By default, these analyzers are enabled and will log to the filesystem. For new installations, these logs are automatically ingested into Elasticsearch. If you upgrade an existing deployment and would like to ingest these logs, you can enable them via *so-zeek-logs*.

Security Onion also includes 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 add `bzar` to your `global.sls` as shown in the Custom Scripts section below. 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*.

## 8.4.6 VLAN tags

If you want to include VLAN tags in your `conn.log`, then you can enable the `protocols/conn/vlan-logging` script (already included in Zeek but disabled by default) in your Zeek configuration as shown in the Custom Scripts section below.

## 8.4.7 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/>

<https://github.com/weslambert/securityonion-misp>

## 8.4.8 Custom Scripts

Custom scripts can be added to `/opt/so/saltstack/local/salt/zeek/policy/custom/<$custom-module>` on the manager. The custom folder is mapped to Zeek through Docker on the minions. Once the script module is created, the configuration for `local.zeek` will need to be updated. In Security Onion 2, this configuration is abstracted into a *Salt* pillar. For example, we would copy the following lines from `/opt/so/saltstack/default/pillar/zeek/init.sls` into the `global.sls` file, replacing `$module-name` on the last line with the actual module name:

```
zeek:
  local:
    '@load':
      - misc/loaded-scripts
      - tuning/defaults
      - misc/capture-loss
      - misc/stats
      - frameworks/software/vulnerable
      - frameworks/software/version-changes
      - protocols/ftp/software
      - protocols/smtp/software
      - protocols/ssh/software
      - protocols/http/software
      - protocols/dns/detect-external-names
      - protocols/ftp/detect
      - protocols/conn/known-hosts
      - protocols/conn/known-services
      - protocols/ssl/known-certs
      - protocols/ssl/validate-certs
      - protocols/ssl/log-hostcerts-only
      - protocols/ssh/geo-data
      - protocols/ssh/detect-bruteforcing
      - protocols/ssh/interesting-hostnames
      - protocols/http/detect-sqli
      - frameworks/files/hash-all-files
      - frameworks/files/detect-MHR
```

(continues on next page)

(continued from previous page)

```

- policy/frameworks/notice/extend-email/hostnames
- ja3
- hassh
- intel
- cve-2020-0601
- securityonion/bpfconf
- securityonion/communityid
- securityonion/file-extraction
- oui-logging
- icsnpp-modbus
- icsnpp-dnp3
- icsnpp-bacnet
- icsnpp-ethercat
- icsnpp-enip
- icsnpp-opcua-binary
- icsnpp-bsap
- icsnpp-s7comm
- zeek-plugin-tds
- zeek-plugin-profinet
- zeek-spicy-wireguard
- zeek-spicy-stun
- custom/$module-name

```

Once the configuration has been updated, Zeek can be restarted with `sudo so-zeek-restart` on applicable nodes to pick up the changes. Finally, `/nsm/zeek/logs/current/loaded_scripts.log` can be checked to ensure the new module has been loaded. For example:

```
grep mynewmodule /nsm/zeek/logs/current/loaded_scripts.log
```

### 8.4.9 Custom Script Example: log4j

Corelight has developed a Zeek package to detect log4j exploitation attempts and it can be found at <https://github.com/corelight/cve-2021-44228>. This package contains Zeek scripts which can easily be loaded into your Security Onion deployment.

First, we need to make sure that the custom directory exists on the manager:

```
sudo mkdir -p /opt/so/saltstack/local/salt/zeek/policy/custom/
```

Next, download the Zeek package to a temporary location:

```
git clone https://github.com/corelight/cve-2021-44228.git
```

Now we need to move the Zeek scripts to the Zeek custom directory:

```
sudo mv cve-2021-44228/scripts /opt/so/saltstack/local/salt/zeek/policy/custom/cve-2021-44228
```

Next, we need to configure Zeek to load the new scripts. If `/opt/so/saltstack/local/pillar/global.sls` does not already contain a `zeek:` section, then copy and paste the following at the end of the file (be careful when pasting to respect yaml indentation):

```

zeek:
  local:
    '@load':
      - misc/loaded-scripts
      - tuning/defaults
      - misc/capture-loss
      - misc/stats
      - frameworks/software/vulnerable
      - frameworks/software/version-changes
      - protocols/ftp/software
      - protocols/smtp/software
      - protocols/ssh/software
      - protocols/http/software
      - protocols/dns/detect-external-names
      - protocols/ftp/detect
      - protocols/conn/known-hosts
      - protocols/conn/known-services
      - protocols/ssl/known-certs
      - protocols/ssl/validate-certs
      - protocols/ssl/log-hostcerts-only
      - protocols/ssh/geo-data
      - protocols/ssh/detect-bruteforcing
      - protocols/ssh/interesting-hostnames
      - protocols/http/detect-sqli
      - frameworks/files/hash-all-files
      - frameworks/files/detect-MHR
      - policy/frameworks/notice/extend-email/hostnames
      - ja3
      - hassh
      - intel
      - cve-2020-0601
      - securityonion/bpfconf
      - securityonion/communityid
      - securityonion/file-extraction
      - custom/cve-2021-44228

```

Within 15 minutes, *Salt* should automatically restart Zeek where necessary. If you don't want to wait, you can manually restart Zeek with the following command. If you have a distributed deployment, you could run this command on each sensor manually or use *Salt* to run the command across all sensors at once:

```
sudo so-zeek-restart
```

#### 8.4.10 Modifying base protocol scripts

If you need to modify base protocol scripts, you can do so as follows. In this example, we are modifying the default ports that Zeek considers for the MySQL analyzer:

```

const ports = { 11434/tcp, 13306/tcp };

event zeek_init() &priority=5
{

```

(continues on next page)



(continued from previous page)

```
Analyzer::register_for_ports(Analyzer::ANALYZER_MYSQL, ports);
}
```

### 8.4.11 Configuration

You can use [Salt](#) to manage Zeek's `local.zeek`, `node.cfg` and `zeekctl.cfg`:

`local.zeek`: The allowed options for this file are `@load`, `@load-sigs` and `redef`. An example of configuring this pillar can be seen below.

`node.cfg`: The pillar items to modify this file are located under the `sensor` pillar in the minion pillar file. The options that can be customized in the file include: `interface`, `lb_procs`, `pin_cpus`, and `af_packet_buffer_size`.

`zeekctl.cfg`: An example of customizing this can be seen below. The allowed options can be seen in <https://github.com/Security-Onion-Solutions/securityonion/blob/master/salt/zeek/files/zeekctl.cfg.jinja>.

Here is an example of how we would modify `local.zeek`. We can see the default pillar assignments used for `local.zeek` in `/opt/so/saltstack/default/pillar/zeek/init.sls`. This file should never be modified as it could be updated in the future and any modification made would be overwritten. The global or minion pillar files should be used for making changes as they are stored in `/opt/so/saltstack/local/`, and that directory isn't overwritten during a Security Onion code update.

```
zeek:
  zeekctl:
    MailTo: root@localhost
    MailConnectionSummary: 1
    MinDiskSpace: 5
    MailHostUpDown: 1
    LogRotationInterval: 3600
    LogExpireInterval: 0
    StatsLogEnable: 1
    StatsLogExpireInterval: 0
    StatusCmdShowAll: 0
    CrashExpireInterval: 0
    SitePolicyScripts: local.zeek
    LogDir: /nsm/zeek/logs
    SpoolDir: /nsm/zeek/spool
    CfgDir: /opt/zeek/etc
    CompressLogs: 1
  local:
    '@load':
      - misc/loaded-scripts
      - tuning/defaults
      - misc/capture-loss
      - misc/stats
      - frameworks/software/vulnerable
      - frameworks/software/version-changes
      - protocols/ftp/software
      - protocols/smtp/software
      - protocols/ssh/software
      - protocols/http/software
      - protocols/dns/detect-external-names
      - protocols/ftp/detect
```

(continues on next page)

(continued from previous page)

```

- protocols/conn/known-hosts
- protocols/conn/known-services
- protocols/ssl/known-certs
- protocols/ssl/validate-certs
- protocols/ssl/log-hostcerts-only
- protocols/ssh/geo-data
- protocols/ssh/detect-bruteforcing
- protocols/ssh/interesting-hostnames
- protocols/http/detect-sqli
- frameworks/files/hash-all-files
- frameworks/files/detect-MHR
- policy/frameworks/notice/extend-email/hostnames
- ja3
- hassh
- intel
- cve-2020-0601
- securityonion/bpfconf
- securityonion/communityid
- securityonion/file-extraction
'@load-sigs':
- frameworks/signatures/detect-windows-shells
redef:
- LogAscii::use_json = T;
- LogAscii::json_timestamps = JSON::TS_ISO8601;

```

In this file, there are two keys under `zeek`, `zeekctl` and `local`. We will be using `zeek:local` for this example since we are modifying the `zeek.local` file. We will address `zeek:zeekctl` in another example where we modify the `zeekctl.cfg` file.

Under `zeek:local`, there are three keys: `@load`, `@load-sigs`, and `redef`. In the pillar definition, `@load` and `@load-sigs` are wrapped in quotes due to the `@` character. Under each of the keys, there is a list of items that will be added to the `local.zeek` file with the appropriate directive of either `@load`, `@load-sigs` or `redef`. In order to modify either of the lists, the entire list must be redefined in either the global or minion pillar file.

If we have a node where `protocols/ssh/detect-bruteforcing` is generating a lot of noise and we want to tell Zeek to stop loading that script, we would do the following. Since we just want to turn it off for that specific node, we would open `/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls`. At the bottom, we would append the following:

```

zeek:
  local:
    '@load':
      - misc/loaded-scripts
      - tuning/defaults
      - misc/capture-loss
      - misc/stats
      - frameworks/software/vulnerable
      - frameworks/software/version-changes
      - protocols/ftp/software
      - protocols/smtp/software
      - protocols/ssh/software
      - protocols/http/software
      - protocols/dns/detect-external-names

```

(continues on next page)

(continued from previous page)

```

- protocols/ftp/detect
- protocols/conn/known-hosts
- protocols/conn/known-services
- protocols/ssl/known-certs
- protocols/ssl/validate-certs
- protocols/ssl/log-hostcerts-only
- protocols/ssh/geo-data
- protocols/ssh/interesting-hostnames
- protocols/http/detect-sqli
- frameworks/files/hash-all-files
- frameworks/files/detect-MHR
- policy/frameworks/notice/extend-email/hostnames
- ja3
- hassh
- intel
- cve-2020-0601
- securityonion/bpfconf
- securityonion/communityid
- securityonion/file-extraction

```

We redefined the @load list in the minion pillar file, but we left out the `protocols/ssh/detect-bruteforcing`. This will override the value defined in the `/opt/so/saltstack/default/pillar/zeek/init.sls` and the global pillar file if it is defined there, and prevent the script from being added to the `local.zeek` file. If we wanted to add a script to be loaded, then we would add our script to the list. Since we aren't changing @load-sigs or redef, then we do not need to add them here. Once the file is saved, and the node checks in with manager, the `local.zeek` file will be updated and the `so-zeek` docker container will be restarted.

Let's see an example of how we would modify the `zeekctl.cfg` file. From the example above, we know that the default pillar values are set for zeek in `/opt/so/saltstack/default/pillar/zeek/init.sls`. The default pillar values for `zeekctl.cfg` are as follows:

```

zeek:
  zeekctl:
    MailTo: root@localhost
    MailConnectionSummary: 1
    MinDiskSpace: 5
    MailHostUpDown: 1
    LogRotationInterval: 3600
    LogExpireInterval: 0
    StatsLogEnable: 1
    StatsLogExpireInterval: 0
    StatusCmdShowAll: 0
    CrashExpireInterval: 0
    SitePolicyScripts: local.zeek
    LogDir: /nsm/zeek/logs
    SpoolDir: /nsm/zeek/spool
    CfgDir: /opt/zeek/etc
    CompressLogs: 1

```

For anything not defined here, Zeek will use its own defaults. The options that are allowed to be managed with the pillar can be found at <https://github.com/Security-Onion-Solutions/securityonion/blob/master/salt/zeek/files/zeekctl.cfg.jinja>.

In order to add or modify an option in `zeekctl`, we will need to modify either the global or minion pillar file. For

example, if we wanted to turn log compression off and change the timeout for Broker communication events to 20 seconds globally, we would add the following to the global pillar file.

```
zeek:
  zeekctl:
    compresslogs: 0
    commtimeout: 20
```

Since `zeek:zeekctl` is a dictionary with dictionary values, we do not need to redefine the entire pillar here like we did for `zeek:local` above. Once the pillar file is saved and the node checks in with the manager, the `zeekctl.cfg` file will be updated and the `so-zeek` container will be restarted.

### 8.4.12 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
```

### 8.4.13 Disabling

Zeek can be disabled by setting `enabled: false` in the `zeek` *Salt* pillar.

If you just want to disable Zeek on a single sensor, then you can edit that sensor's `minion.sls` file. If the file doesn't already have a `zeek` section, then add the following to the end of the file:

```
zeek:
  enabled: false
```

If you want to disable Zeek globally across all your sensors, then you could add that entry to your `global.sls` file.

### 8.4.14 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.

Depending on what options you choose in Setup, it may ask if you want to use *Zeek* or *Suricata* for metadata. Whichever engine you choose for metadata will then extract files from 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](#), [Hunt](#), or [Kibana](#). Here is an example of Strelka detecting Poison Ivy RAT:
















The screenshot shows the Security Onion Alerts page. The left sidebar contains navigation links: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, and Administration. The main content area is titled 'Alerts' and includes a search bar with the placeholder 'Specify a query in Onion Query Language (OQL)'. Below the search bar, there are controls for 'Group By Name, Module', a 'Fetch Limit' of 500, and a 'Filter Results' button. A table displays the alert results with columns: Count, rule.name, event.module, and event.severity\_label. The table shows one alert with a count of 1, rule.name 'RAT\_PoisonIvy', event.module 'strelka', and event.severity\_label 'high'. The bottom right of the table shows 'Rows per page: 50' and '1-1 of 1'.

Count	rule.name	event.module	event.severity_label
1	RAT_PoisonIvy	strelka	high

Drilling into that alert, we find more information about the file and the YARA rule:

 event.dataset	alert
 event.module	strelka
 event.severity	3
 event.severity_label	high
 file.accessed	1630321382
 file.ctime	1630321382
 file.depth	0
 file.flavors.yara	[ "mz_file" ]
 file.mime_type	[ "application/x-dosexec" ]
 file.mtime	1630321382
 file.name	tmpsvz2s7yk
 file.permissions	rw-----
 file.scanners	[ "ScanEntropy", "ScanExiftool", "ScanHash", "ScanHeader", "ScanPe", "ScanYara" ]
 file.size	8192
 file.source	/dev/shm/tmpsvz2s7yk
 file.tree.node	f0ac2c0b-a1d8-401e-98d6-d03a1aef8644
 file.tree.root	f0ac2c0b-a1d8-401e-98d6-d03a1aef8644
 hash.elapsed	0.020976
 hash.md5	e2c33fa7a3802289d46a7c3e4e1df342
 hash.sha1	d8fd563fbbdea43c78841ccca49e8c5a3fe47cbc
 hash.sha256	35c35bc56ce3064f6236db4432fDCF578d098353076d3fBE1e600fa926bc6227
 hash.ssdeep	192:JJGc1ZI2+VAfNxI1THs6xgzgVGjPIROInQAkhFo2A:JJGcMJxDTHfRmoc

 rule.author	Kevin Breen <kevin@technarchy.net>
 rule.date	01.04.2014
 rule.description	Detects PoisonIvy RAT
 rule.filetype	exe
 rule.maltype	Remote Access Trojan
 rule.name	RAT_PoisonIvy
 rule.reference	<a href="http://malwareconfig.com/stats/PoisonIvy">http://malwareconfig.com/stats/PoisonIvy</a>
 scan.entropy.elapsed	0.000277
 scan.entropy.entropy	6.030109054353968
 scan.exiftool	[ "SourceFile=/dev/shm/tmpsvz2s7yk", "ExifToolVersion=11.88", "FileName=tmpsvz2s7yk", "Directory=/dev/shm", "FileSize=8.0 kB", "FileModifyDate=1630321382", "FileAccessDate=1630321382", "FileInodeChangeDate=1630321382", "FilePermissions=rw-----", "FileType=Win32 EXE", "FileTypeExtension=exe", "MIMEType=application/octet-stream", "MachineType=Intel 386 or later, and compatibles", "TimeStamp=1199631091", "ImageFileCharacteristics=No relocs, Executable, No line numbers, No symbols, 32-bit", "PEType=PE32", "LinkerVersion=5.12", "CodeSize=512", "InitializedDataSize=7168", "UninitializedDataSize=0", "EntryPoint=520", "OSVersion=4", "ImageVersion=4", "SubsystemVersion=4", "Subsystem=Windows GUI" ]

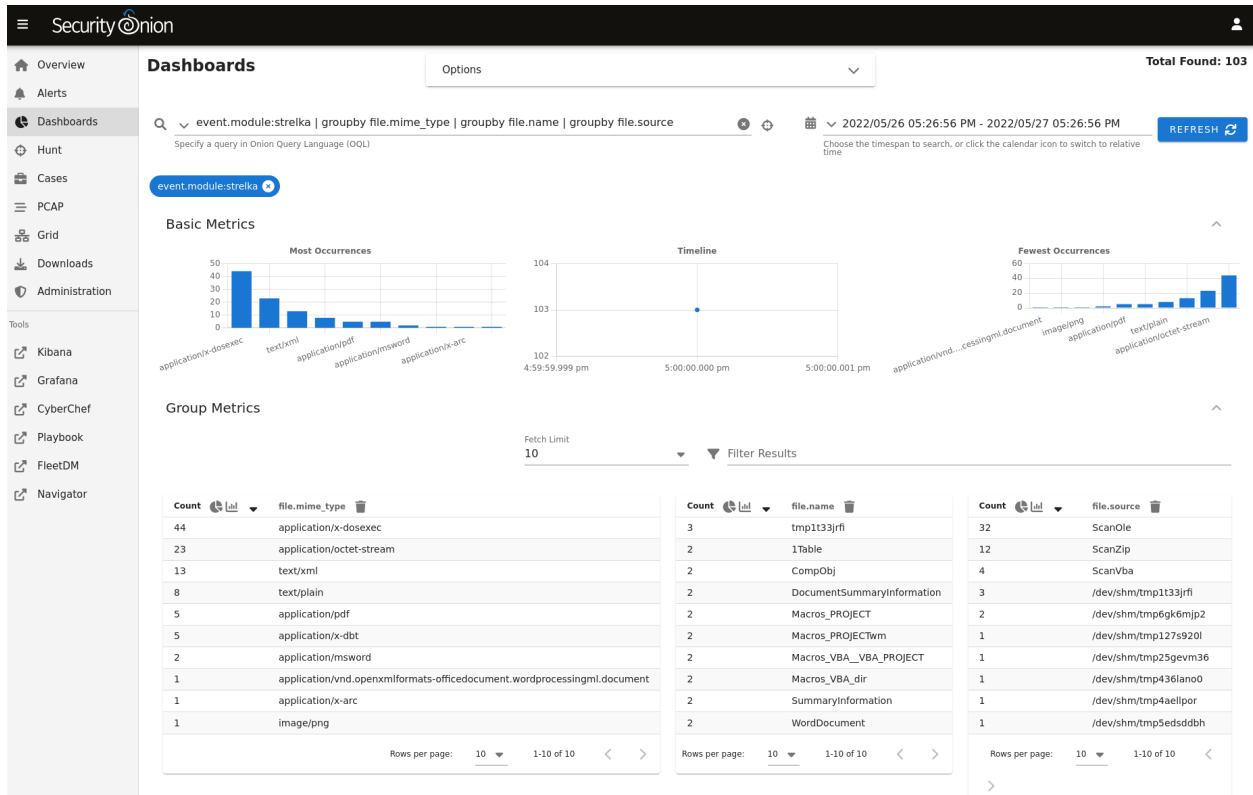
 scan.pe.size_of_stack_reserve	1048576
 scan.pe.size_of_uninitialized_data	0
 scan.pe.subsystem_version	4
 scan.pe.summary.section_md5	[ "e2b705b27f9783b8d25641bd5f34eb27", "9e5912d9f35aa91102fcd5f4740ef0a" ]
 scan.pe.summary.section_sha1	[ "6f3f04bde3817992a3fcb1fa9e22f1a472c0005b", "9a084593794e4a38007068ccbecf78b7a3d5adb7" ]
 scan.pe.summary.section_sha256	[ "6b85472307dee17ba961a68d9791a529fee61b33ba2d727475c1d349c98e3641", "5a51af1e3c8e448fad09bdd76bcc6642e49701d180208fb86befee821a23ef25" ]
 scan.pe.symbols.imported	[ "ExitProcess" ]
 scan.pe.symbols.libraries	[ "kernel32.dll" ]
 scan.pe.symbols.table	[ { "library": "kernel32.dll", "symbols": [ "ExitProcess" ], "type": "import" } ]
 scan.pe.total.libraries	1
 scan.pe.total.resources	0
 scan.pe.total.sections	2
 scan.pe.total.symbols	1
 scan.yara.elapsed	0.917085
 scan.yara.matches	[ "RAT_PoisonIvy" ]

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](#), [Hunt](#), and [Kibana](#). Here's an example of the default Strelka dashboard in [Dashboards](#):





## 8.5.3 Configuration

Strelka reads its configuration from `/opt/so/conf/strelka/`. However, please keep in mind that if you make any changes to this directory they may be overwritten since the configuration is managed with *Salt*.

## 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>.

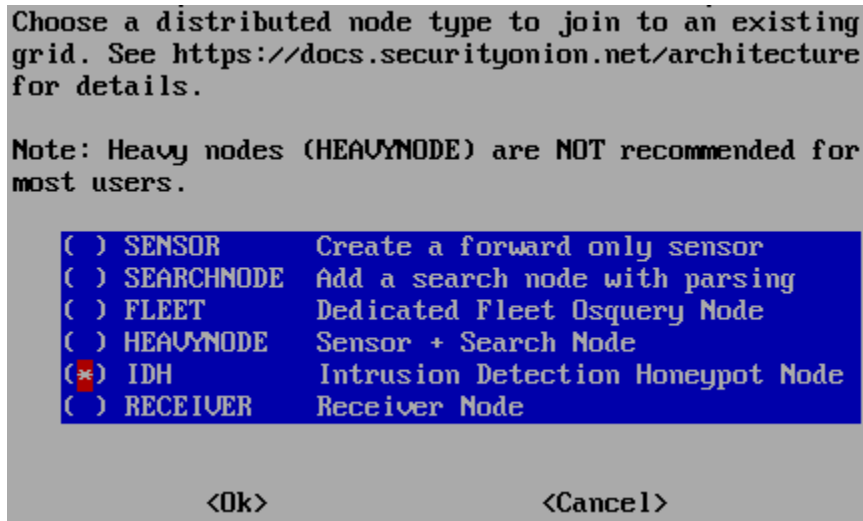
### 8.6.1 Installation

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.

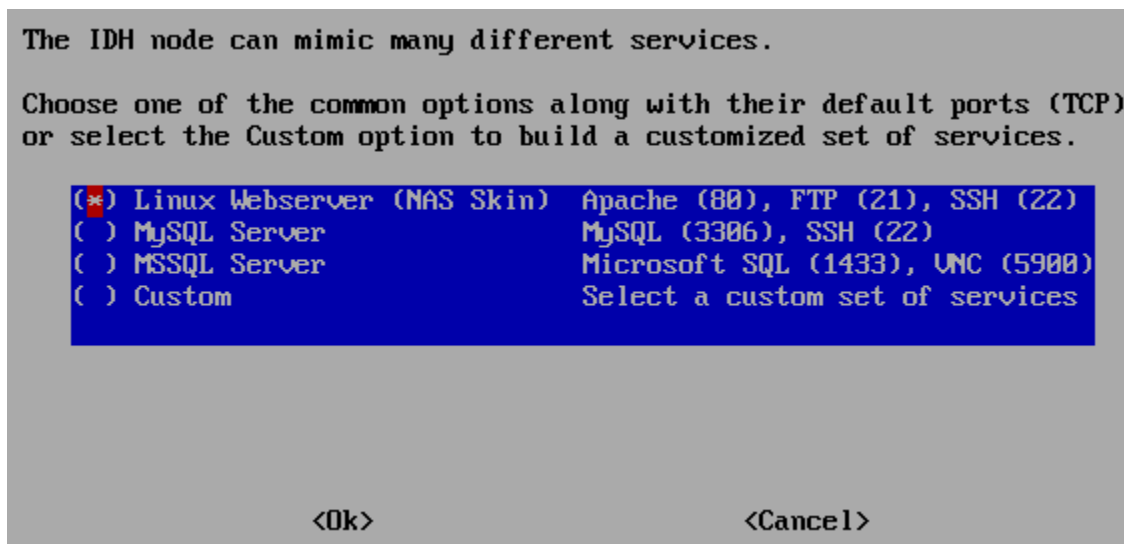
**Warning:** The IDH node is designed to be placed *inside* your network perimeter! It should not be accessible from the Internet!

## 8.6.2 Screenshots

Run Setup, select the DISTRIBUTED install submenu, choose the Existing Deployment option, and then select the IDH option.



Then choose from one of the common honeypot options or choose the Custom option to customize.



You can optionally prevent the IDH services from listening on the management interface.

By default, the IDH services selected in the previous screen will be bound to all interfaces and IP addresses on this system.

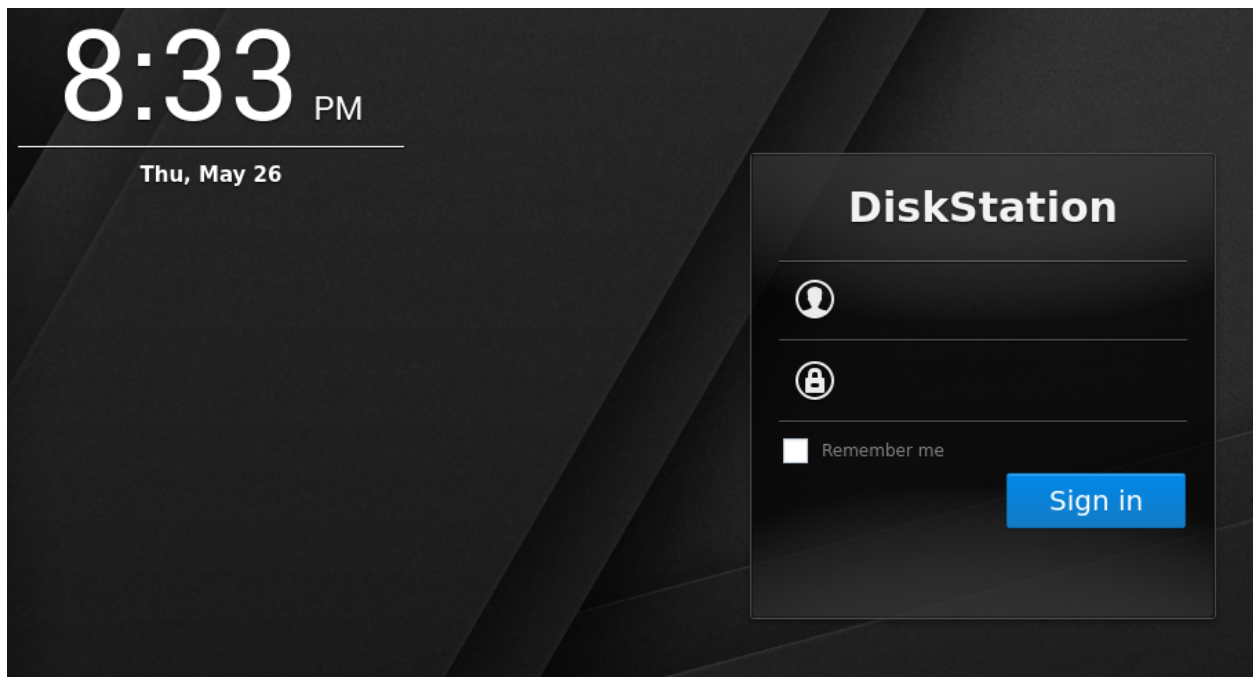
If you would like to prevent IDH services from being published on this system's management IP, you can select the option below.

☒ 192.168.230.231 Disable IDH services on this management IP

<Ok>

<Cancel>

If you chose the Linux Webserver option above, anybody who connects to the webserver will see a fake login page.



Connections to honeypot services will result in SO IDH alerts that can be seen in [Alerts](#).

**Alerts** Options Total Found: 15

Search: Group By Name, Module  
Specify a query in Onion Query Language (OQL)

Fetch Limit: 500 Filter Results

Count	rule.name	event.module	event.severity_label
12	SO IDH - SSH Accessed	playbook	critical
1	SO IDH - FTP Login Attempt	playbook	critical
1	SO IDH - HTTP Accessed	playbook	critical
1	SO IDH - HTTP Login Attempt	playbook	critical

### 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 SSH

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

Services can be customized in two ways:

- 1) Changing the default port.

Note: The previously-configured port is not automatically closed on the IDH node’s firewall, so it will need to be closed manually. This can be done by running the following command on the IDH node, substituting the protocol and port as needed:

```
iptables -D INPUT -p tcp -m tcp --dport 80 -j ACCEPT
```

- 2) Service-specific config, like SSH version string. All of the defaults for these types of configuration can be found in the following files on the manager:

`/opt/so/saltstack/default/salt/idh/defaults/$Service.defaults.yaml`

Both of these configurations can be implemented by editing the node’s minion pillar, which is stored on the manager in `/opt/so/saltstack/local/pillar/minions/$IDH-Hostname_idh.sls`.

### 8.6.7 Custom Configuration - Example: Port Change

For example, suppose that we already have the HTTP service running but we want to change the default port from 80 to 8080.

**Warning:** The following configuration files are YAML, and so tabs are not permitted. Please only use spaces and pay special attention to the number of spaces!

First, we need to copy the default configuration for the HTTP service. This can be found on the manager in `/opt/so/saltstack/default/salt/idh/defaults/http.defaults.yaml`. Out of the defaults defined there, we just need the following:

```

idh:
  opencanary:
    config:
      http.port: 80

```

Next, we will edit the minion sls file (`/opt/so/saltstack/local/pillar/minions/$IDH-Hostname_idh.sls`) and add the previously copied config, but swap `80` for `8080`. The minion sls file should look something like this:

```

idh:
  services:
    - http
    - ftp
    - ssh
  opencanary:
    config:
      http.port: 8080

```

With this configuration changed, we can now make it active on the IDH node by using *Salt* to apply the `idh` and `firewall` states. Run the following from the manager (replacing `$IDH-Hostname` with your actual IDH hostname):

```
sudo salt '$IDH-Hostname*' state.apply idh,firewall
```

You should now be able to browse to the HTTP server on the IDH node on TCP/8080. You should be able to override any other configuration in the `http.defaults.yaml` in a similar manner.

### 8.6.8 Custom Configuration - Example: Custom HTTP Skin

For example, suppose that we already have the HTTP service running but we want to change the default skin to a custom skin.

**Warning:** The following configuration files are YAML, and so tabs are not permitted. Please only use spaces and pay special attention to the number of spaces!

First off, copy the custom skin folder to the following directory on the Manager:

```
/opt/so/saltstack/local/salt/idh/skins/http/custom/
```

It should look like this:

```
/opt/so/saltstack/local/salt/idh/skins/http/custom/CustomSkin01/
```

Next, edit the minion sls file (`/opt/so/saltstack/local/pillar/minions/$IDH-Hostname_idh.sls`) and add an override for the custom skin. The minion sls file should look something like this:

```

idh:
  services:
    - http
    - ftp
    - ssh
  opencanary:
    config:
      http.skin: CustomSkin01

```

With this configuration changed, we can now make it active on the IDH node by using *Salt* to apply the `idh` states. Run the following from the manager (replacing `$IDH-Hostname` with your actual IDH hostname):

```
sudo salt '$IDH-Hostname*' state.apply idh
```

You should now be able to browse to the HTTP server on the IDH node and see your custom skin.



## HOST VISIBILITY

When you logged into *Security Onion Console (SOC)*, you may have seen some host logs from *Wazuh*. Security Onion can also consume many other kinds of host logs as well. You can send logs to Security Onion via your choice of either *osquery*, *Beats*, *Wazuh*, or *Syslog*:

- Choose *osquery* if you want some live response actions and maybe light log transport. A good example here is a roaming laptop where log volume is low and you might want to send its logs to a dedicated *FleetDM* node in the DMZ.
- Choose *Wazuh* if you want HIDS functionality and log transport. However, please be aware that Wazuh will not be included in the upcoming Security Onion 2.4.
- Choose *Beats* for dedicated log transport. Examples would be high volume domain controllers or Windows Event Collectors.
- 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* and/or *Autoruns*. Those additional logs can then be transported by whatever mechanism you chose above.

### 9.1 osquery

From <https://osquery.io/>:

Osquery uses basic SQL commands to leverage a relational data-model to describe a device.

#### 9.1.1 FleetDM

Security Onion includes *FleetDM* to manage your osquery deployment. For more information, please see the *FleetDM* section.

#### 9.1.2 Agents - Deployment

To deploy an osquery agent to an endpoint, go to the *Security Onion Console (SOC) Downloads* page and download the proper osquery agent for the operating system of that endpoint. Use *so-allow* to allow the osquery agent to connect to port 8090 on the manager. Then install the osquery agent and it should check into the manager and start showing up in *FleetDM*.

Osquery will attempt to connect to the manager via the manager's IP or Hostname - whichever was selected during the manager setup. If the hostname is used, the endpoints need to be able to resolve that hostname to the manager's IP. See this value by running the following command on the manager: `sudo salt-call pillar.get global:url_base`.

If this value ever changes, the osquery packages under the *Security Onion Console (SOC) Downloads* page will need to be regenerated.

All the packages (except for the macOS PKG) are customized for the specific Security Onion grid they were downloaded from, and include all the necessary configuration to connect to that grid. The macOS package is a stock Launcher package, and will require additional configuration once it has been deployed.

For macOS deployments, install the package and then configure the following:

- Update `/etc/so-launcher/secret` with the *FleetDM* enroll secret. This can be found by running the following on the manager:

```
sudo salt-call pillar.get secrets:fleet_enroll-secret
```

- Update `/etc/so-launcher/launcher.flags`, changing the hostname to your manager hostname and changing the port from 443 to 8090
- Update `/etc/so-launcher/roots.pem` with the contents from the following file (on your manager): `/etc/ssl/certs/intca.crt`

At this point, osquery should connect up to *FleetDM* within a few minutes. If not, try manually restarting the osquery agent on the macOS endpoint:

```
sudo launchctl kickstart -k system/com.so-launcher.launcher
```

If required, you can review logs for issues:

```
sudo tail /var/log/so-launcher/launcher-stderr.log
```

### 9.1.3 Agents - Updating

Security Onion uses Launcher as a management wrapper around osquery. This allows for a simpler configuration as well as auto-updates of Launcher and osquery. Launcher will check every hour to see if an update is available and, if so, will download and install it. This is the default configuration, but can be changed within the osquery Flags file.

In an airgap environment where the endpoints do not have Internet access, updated osquery packages can be downloaded from the Security Onion Console and used to update the endpoints. Osquery packages are periodically updated on the manager as new versions of osquery are released.

### 9.1.4 Agents - Troubleshooting

Agent logs on Windows endpoints can be found under the Application channel in the Windows Eventlog and the source is Launcher.

### 9.1.5 Agents - Regenerating Install Packages

To regenerate packages, run the following on the manager (it may take up to 5 minutes to rebuild the packages):

```
sudo salt-call state.apply fleet.event_gen-packages
```

### 9.1.6 Hunt or Kibana

All osquery logs can be found by using the following query:

```
event.module: osquery
```

#### Kibana Dashboard: Host Data → Modules/Osquery

This dashboard gives an overview of the osquery logs in the system. As long as the default osquery configuration is used, this dashboard should work out of the box regardless of how you schedule or name your queries and packs.

### 9.1.7 Shipping Windows Eventlogs

Windows Eventlogs from the local Windows system can be shipped with osquery to Security Onion. Current parsing support extends to core Windows Eventlog channels ( `Security` , `Application` , `System` ) as well as Sysmon under the default channel location. These logs will show up in Security Onion as `event.dataset: windows_eventlog` or `event.dataset: sysmon`.

- Confirm that you can successfully live query the logs: `SELECT * FROM windows_events limit 10;`
- Save a new query: Query -> Manage Queries -> Create New Query `SELECT * FROM windows_events;` -> Save
- Add the new query to a query pack that targets a Windows host - how often it should run depends on log volume on the local host; start off with 180 seconds, differential logging: Packs -> Manage Packs -> Select + Edit Pack (Modify Targets for Windows only if needed, Modify Logging options as needed)
- Save pack + Enable pack, if needed.

Please refer to the osquery documentation for further information on osquery Evented tables: <https://osquery.readthedocs.io/en/stable/development/pubsub-framework/#the-pub-sub-evented-data-framework-of-osquery>

### 9.1.8 Community ID

We sponsored the development of *Community ID* support for osquery to allow for quicker and easier log correlation from different data types.

### 9.1.9 More Information

---

**Note:** For more information about osquery, please see <https://osquery.readthedocs.io/en/stable/>.

---

## 9.2 Beats

We can use Elastic Beats to facilitate the shipping of endpoint logs to Security Onion's Elastic Stack. Currently, testing has only been performed with Filebeat (multiple log types) and Winlogbeat (Windows Event logs).

---

**Note:** In order to receive logs from Beats, 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. Please note that, 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.

---

### 9.2.1 so-allow

Run *so-allow* and select the *b* option to allow your Beats agents to send their logs to *Logstash* port 5044/tcp.

### 9.2.2 Version

When downloading a Beats agent, make sure the version number matches the version of Elastic running on your Security Onion deployment.

### 9.2.3 Winlogbeat

Navigate to the Downloads page in *Security Onion Console (SOC)* and download the linked Winlogbeat agent. This will ensure that you get the correct version of Winlogbeat for your Elastic version.

Install Winlogbeat and copy `winlogbeat.example.yml` to `winlogbeat.yml` if necessary. Then configure `winlogbeat.yml` as follows:

- Make sure that the `setup.dashboards.enabled` setting is commented out or disabled.
- Disable the `output.elasticsearch` output.
- Enable the `output.logstash` output and configure it to send logs to port 5044 on your management node.
- If you are shipping *Sysmon* logs, confirm that your Winlogbeat configuration simply collects the *Sysmon* logs and does NOT use the Elastic Sysmon processors section as Security Onion will do all the necessary parsing.

Once `winlogbeat.yml` is configured properly, start the Winlogbeat service.

---

**Note:** Check out our Winlogbeat video at <https://youtu.be/Xz-7oDrZdQY!>

---

### 9.2.4 Installation

To install a Beat, follow the instructions provided for the respective Beat, with the exception of loading the index template, as Security Onion uses its own template file to manage Beats fields.

Filebeat

<https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-installation.html>

Winlogbeat

<https://www.elastic.co/guide/en/beats/winlogbeat/current/winlogbeat-installation.html>

If installing Filebeat on a Linux distribution, you will want to ensure that the service is started after a reboot. We can ensure this by running the following commands after install:

```
sudo update-rc.d filebeat defaults
sudo update-rc.d filebeat enable
```

## 9.2.5 Encryption

**Warning:** Beats communication with *Logstash* is not encrypted by default. If you require encryption, you will need to manually configure it.

### Configuring Encryption for Beats

There are a few considerations when enabling encryption for Beats. If you enable it on the default port then all connections on 5044 will be required to use encryption. The other option is to create a custom port for encryption and send only encrypted beats to that port.

### Using the Beats default port 5044 with encryption

Copy 0009\_input\_beats.conf to the local directory:

```
cp /opt/so/saltstack/default/salt/logstash/pipelines/config/so/0009_input_beats.conf /
  ↪opt/so/saltstack/local/salt/logstash/pipelines/config/so/0009_input_beats.conf
```

Copy your certificates to the proper directory on the manager. You will need a cert from the ca that you are signing the cert from, as well as the cert and key.

```
cp myca.crt /opt/so/saltstack/local/salt/logstash/etc/certs/
cp mybeats.crt /opt/so/saltstack/local/salt/logstash/etc/certs/
cp mybeats.key /opt/so/saltstack/local/salt/logstash/etc/certs/
```

Next make your config look like the one below. Note that the paths are not the same due to docker.

```
input {
  beats {
    port => "5044"
    ssl => true
    ssl_certificate_authorities => ["/usr/share/logstash/certs/myca.crt"]
    ssl_certificate => "/usr/share/logstash/certs/mybeats.crt"
    ssl_key => "/usr/share/logstash/certs/mybeats.key"
    tags => [ "beat-ext" ]
  }
}
```

## 9.2.6 Log files

### Filebeat

Windows: C:\Program Files\Filebeat\filebeat.log

Linux: /var/log/filebeat/filebeat

### Winlogbeat

C:\Program Files\Winlogbeat\winlogbeat.log

Default fields: <https://www.elastic.co/guide/en/beats/winlogbeat/master/exported-fields-eventlog.html>

## 9.2.7 Data

In *Dashboards* and *Hunt*, you can find Beats data by searching for `_index:"*:so-beats-*`.

In *Kibana*, you can find Beats data on the Host dashboard or by searching for `_index:"*:so-beats-*` in Discover.

## 9.3 Wazuh

### 9.3.1 Description

From <https://wazuh.com/>:

Wazuh is a free, open source and enterprise-ready security monitoring solution for threat detection, integrity monitoring, incident response and compliance.

### 9.3.2 Usage

Security Onion utilizes Wazuh as a Host Intrusion Detection System (HIDS) on each of the Security Onion nodes.

The Wazuh components include:

**manager** - runs inside of `so-wazuh` Docker container and performs overall management of agents

**API** - runs inside of `so-wazuh` Docker container and allows for remote management of agents, querying, etc.

**agent** - runs directly on each host and monitors logs/activity and reports to **manager**

The Wazuh API runs at TCP port 55000 locally, and currently uses the default credentials of `user:foo` and `password:bar` for authentication. Keep in mind, the API port is not exposed externally by default. Therefore, firewall rules need to be in place to reach the API from another location other than the Security Onion node on which the targeted Wazuh manager is running.

Since the manager runs inside a Docker container, many of the Wazuh binaries that you might want to run will need to be run inside the Docker container. For example, to run `agent_upgrade`:

```
sudo so-wazuh-agent-upgrade
```

### 9.3.3 Configuration

The main configuration file for Wazuh is `/opt/so/conf/wazuh/ossec.conf`.

### 9.3.4 Syslog

If you want to send Wazuh logs to an external syslog collector, please see the *Syslog Output* section.

### 9.3.5 Active Response

Sometimes, Wazuh may recognize legitimate activity as potentially malicious and engage in Active Response to block a connection. This may result in unintended consequences such as blocking of trusted IPs. To prevent this from occurring, you can add your IP address to a safe list and change other settings in `/opt/so/conf/wazuh/ossec.conf` in the `<!-- Active response -->` section. *so-allow* does this for you automatically when you allow analyst connections.

### 9.3.6 Email

If you want *Wazuh* to send email, you can modify `/opt/so/conf/wazuh/ossec.conf` as follows:

```
<global>
<email_notification>yes</email_notification>
<email_to>YourUsername@YourDomain.com</email_to>
<smtp_server>YourMailRelay.YourDomain.com</smtp_server>
<email_from>ossec@YourDomain.com</email_from>
<email_maxperhour>100</email_maxperhour>
</global>
```

Then restart *Wazuh*:

```
sudo so-wazuh-restart
```

You can specify the severity of an event for which *Wazuh* will send email alerts by specifying an appropriate value for `email_alert_level` in `/opt/so/conf/wazuh/ossec.conf`. If you notice `email_alert_level` is not being respected for a certain rule, it may be that the option is overridden by `<options>alert_by_email</options>` being set for a rule. You can modify this behavior in `/opt/so/conf/wazuh/rules/local_rules.xml`.

You can also find an explanation of the alert levels at <https://www.ossec.net/docs/manual/rules-decoders/rule-levels.html>.

### 9.3.7 Tuning Rules

#### New Rules

You can add new rules in `/opt/so/rules/hids/local_rules.xml`.

## Modify Existing Rules

You can modify existing rules by copying the rule to `/opt/so/rules/hids/local_rules.xml`, making your changes, and adding `overwrite="yes"` as shown at <https://documentation.wazuh.com/current/user-manual/ruleset/custom.html#changing-an-existing-rule>. To suppress a Wazuh alert, you can add the rule and include `noalert="1"` in the rule section.

The overall process would be as follows:

1. First, find the existing rule in `/opt/so/rules/hids/ruleset/rules/`.
2. Copy the rule to `/opt/so/rules/hids/local_rules.xml`.
3. Put the rule inside `<group>` `</group>` tags and give it a name.
4. Update the `<rule>` section to include `noalert="1"` along with `overwrite="yes"`.
5. Finally, restart Wazuh with `sudo so-wazuh-restart`.

Here is an example to suppress “Windows Logon Success” and “Windows User Logoff” alerts:

```
<group name="customrules,">
  <rule id="60106" level="3" noalert="1" overwrite="yes">
    <if_sid>60103</if_sid>
    <field name="win.system.eventID">^528$|^540$|^673$|^4624$|^4769$</field>
    <description>Windows Logon Success</description>
    <options>no_full_log</options>
    <mitre>
      <id>T1078</id>
    </mitre>
    <group>authentication_success,pci_dss_10.2.5,gpg13_7.1,gpg13_7.2,gdpr_IV_
↪32.2,hipaa_164.312.b,nist_800_53_AU.14,nist_800_53_AC.7,tsc_CC6.8,tsc_CC7.2,
↪tsc_CC7.3,</group>
  </rule>

  <rule id="60137" level="3" noalert="1" overwrite="yes">
    <if_sid>60103</if_sid>
    <field name="win.system.eventID">^538$|^551$|^4634$|^4647$</field>
    <description>Windows User Logoff</description>
    <options>no_full_log</options>
    <group>pci_dss_10.2.5,gdpr_IV_32.2,hipaa_164.312.b,nist_800_53_AU.14,nist_
↪800_53_AC.7,tsc_CC6.8,tsc_CC7.2,tsc_CC7.3,</group>
  </rule>
</group>
```

---

**Note:** This will not remove existing alerts that were generated before applying the new rule. Also note that this only suppresses the alert and not the underlying log.

---



## Child Rules

In addition to overwriting rules, another option is to add child rules using `if_sid`. In this example, suppose you are receiving Wazuh alerts for PAM: Login session closed and want to stop receiving those alerts for a particular user account.

Let's start by using `ossec-logtest` with a default configuration and pasting in a sample log:

```
[doug@securityonion ~]$ sudo docker exec -it so-wazuh /var/ossec/bin/ossec-logtest
2022/02/24 17:52:49 ossec-testrule: INFO: Started (pid: 2298).
ossec-testrule: Type one log per line.

Feb 24 17:46:19 securityonion sshd[37140]: pam_unix(sshd:session): session closed for
↪user doug

**Phase 1: Completed pre-decoding.
    full event: 'Feb 24 17:46:19 securityonion sshd[37140]: pam_unix(sshd:session):
↪session closed for user doug'
    timestamp: 'Feb 24 17:46:19'
    hostname: 'securityonion'
    program_name: 'sshd'
    log: 'pam_unix(sshd:session): session closed for user doug'

**Phase 2: Completed decoding.
    decoder: 'pam'
    dstuser: 'doug'

**Phase 3: Completed filtering (rules).
    Rule id: '5502'
    Level: '3'
    Description: 'PAM: Login session closed.'
**Alert to be generated.
```

This shows us the rule that would fire and its Rule id of 5502. Now let's add the following rule to `/opt/so/rules/hids/local_rules.xml`:

```
<rule id="100002" level="1">
  <if_sid>5502</if_sid>
  <match>doug</match>
  <description>ignore logins from doug</description>
</rule>
```

Finally, let's re-run `ossec-logtest`:

```
[doug@securityonion ~]$ sudo docker exec -it so-wazuh /var/ossec/bin/ossec-logtest
2022/02/24 17:54:26 ossec-testrule: INFO: Started (pid: 2305).
ossec-testrule: Type one log per line.

Feb 24 17:46:19 securityonion sshd[37140]: pam_unix(sshd:session): session closed for
↪user doug

**Phase 1: Completed pre-decoding.
```

(continues on next page)

(continued from previous page)

```

    full event: 'Feb 24 17:46:19 securityonion sshd[37140]: pam_unix(sshd:session):
↪session closed for user doug'
    timestamp: 'Feb 24 17:46:19'
    hostname: 'securityonion'
    program_name: 'sshd'
    log: 'pam_unix(sshd:session): session closed for user doug'

**Phase 2: Completed decoding.
    decoder: 'pam'
    dstuser: 'doug'

**Phase 3: Completed filtering (rules).
    Rule id: '100002'
    Level: '1'
    Description: 'ignore logins from doug'
**Alert to be generated.
```

This shows us that Wazuh no longer fires rule 5502 but now fires our new alert. Once your changes are complete, restart Wazuh with `sudo so-wazuh-restart`.

### 9.3.8 Apparmor DENIED Alerts

If you're running on Ubuntu, then you most likely get Wazuh HIDS alerts for Apparmor DENIED. In most cases, this is due to telegraf (part of *Grafana*) trying to use `ptrace`. To silence these alerts, you might want to add a child rule as shown in the section above. Apparmor DENIED is Wazuh sid 52002 so that would go in the `if_sid` section. We could then use a regular expression to look for `operation="ptrace"` and `comm="telegraf"` but allow any digit value in the pid field:

```

<rule id="100002" level="0">
  <if_sid>52002</if_sid>
  <description>ignore apparmor denied messages from telegraf</description>
  <regex>apparmor="DENIED" operation="ptrace" profile="docker-default" pid=\d+ comm=
↪"telegraf"</regex>
</rule>
```

We could add this to `/opt/so/rules/hids/local_rules.xml` and then restart Wazuh with `sudo so-wazuh-restart`.

### 9.3.9 Adding Agents

**Warning:** Before deploying lots of Wazuh agents to your enterprise, please be aware that Wazuh will not be included in the upcoming Security Onion 2.4. For more information, please see <https://blog.securityonion.net/2023/02/security-onion-in-2022-and-2023.html>.

Navigate to the Downloads page in *Security Onion Console (SOC)* and download the appropriate Wazuh agent for your endpoint. This will ensure that you get the correct version of Wazuh. If your endpoint is not listed there, you can check the Wazuh website at <https://documentation.wazuh.com/3.13/installation-guide/packages-list/index.html>.

**Warning:** It is important to ensure that you download the agent that matches the version of your Wazuh server. For example, if your Wazuh server is version 3.13.1, then you will want to deploy Wazuh agent version 3.13.1.

You can verify the version of your current Wazuh server using the following command:

```
sudo docker exec -it so-wazuh dpkg -l |grep wazuh
```

Once you've installed the Wazuh agent on the host(s) to be monitored, then perform the steps defined here:

<https://documentation.wazuh.com/3.13/user-manual/registering/command-line-registration.html>

Please keep in mind that when you run `manage_agents` you will need to do so inside the `so-wazuh` container like this:

```
sudo so-wazuh-agent-manage
```

You also may need to run `so-allow` to allow traffic from the IP address of your Wazuh agent(s).

**Warning:** The Windows Wazuh agent installs with incorrect permissions on `ossec.conf` which could allow users to escalate privileges. However, most users configure that Wazuh agent using Wazuh Agent Manager which then sets the permissions correctly. If you don't use the Wazuh Agent Manager for configuration, then you may need to manually fix the permissions on `ossec.conf`. For more information, please see <https://github.com/Security-Onion-Solutions/securityonion/discussions/9390>.

### 9.3.10 Maximum Number of Agents

Security Onion is configured to support a maximum number of 14000 Wazuh agents reporting to a single Wazuh manager.

### 9.3.11 Automated Deployment

If you would like to automate the deployment of Wazuh agents, the Wazuh server includes `ossec-authd`. You can read more about `ossec-authd` at <https://documentation.wazuh.com/3.13/user-manual/reference/daemons/ossec-authd.html>.

When using `ossec-authd`, be sure to add a firewall exception for agents to access port 1515/tcp on the Wazuh manager node by running `so-allow` and choosing the `r` option.

### 9.3.12 API

The Wazuh API runs on port 55000 and requires a user to be created for access. To add a new user, run `so-wazuh-user-add` as follows (replacing `newuser` with the actual username you'd like to create):

```
sudo so-wazuh-user-add newuser
```

When prompted, provide a password for the new user. Once the user has been added, then restart Wazuh:

```
sudo so-wazuh-restart
```

Once restarted, try accessing the API locally from the node using the newly created user and password:

```
curl -k -u newuser:password https://localhost:55000
```

You should receive a message similar to the following indicating success:

```
{"error":0,"data":{"msg":"Welcome to Wazuh HIDS API","api_version":"v3.13.1","hostname":  
  ↪ "securityonion-is-the-coolest","timestamp":"Wed Feb 02 2022 13:09:03    GMT+0000 (UTC)  
  ↪ "}}
```

If you receive a 401 (Unauthorized) error message, double-check the credentials or try running `sudo so-wazuh-user-passwd` if necessary. You can also check the user file inside the Docker container:

```
sudo docker exec -it so-wazuh cat /var/ossec/api/configuration/auth/user
```

### 9.3.13 Diagnostic Logging

Wazuh diagnostic logs are in `/nsm/wazuh/logs/`. Depending on what you're looking for, you may also need to look at the *Docker* logs for the container:

```
sudo docker logs so-wazuh
```

### 9.3.14 More Information

---

**Note:** For more information about Wazuh, please see <https://documentation.wazuh.com/3.13/>.

---

## 9.4 Syslog

If you want to send syslog from other devices, you should check to see if the device has an existing *Filebeat* module at <https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-modules.html>. If so, using the *Filebeat* module should provide some parsing by default.

If your device does not have an existing *Filebeat* module, you can still collect standard syslog by running *so-allow* on the manager and then choosing 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.

## 9.5 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.

---

**Note:** Check out our Sysmon video at <https://youtu.be/Xz-7oDrZdQY!>

---

### 9.5.1 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](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.

### 9.5.2 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>

### 9.5.3 Transport

Sysmon logs can be collected and transported using *Beats*, *osquery*, or *Wazuh*.

### 9.5.4 Winlogbeat

If you are shipping Sysmon logs via Winlogbeat (see the *Beats* section), confirm that your Winlogbeat configuration does NOT use the Elastic Sysmon module. Security Onion will do all the necessary parsing.

## 9.5.5 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.

## 9.5.6 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>

---

## 9.6 Autoruns

From <https://docs.microsoft.com/en-us/sysinternals/downloads/autoruns>:

This utility, which has the most comprehensive knowledge of auto-starting locations of any startup monitor, shows you what programs are configured to run during system bootup or login, and when you start various built-in Windows applications like Internet Explorer, Explorer and media players. These programs and drivers include ones in your startup folder, Run, RunOnce, and other Registry keys. Autoruns reports Explorer shell extensions, toolbars, browser helper objects, Winlogon notifications, auto-start services, and much more. Autoruns goes way beyond other autostart utilities.

### 9.6.1 Integration

#### Pertinax

Josh Brower developed a great project called Pertinax to normalize autoruns data and integrate it into Security Onion:

<https://github.com/defensivedepth/Pertinax/wiki/Introduction>

Execute autoruns and ar-normalize.ps1 as shown here:

<https://github.com/defensivedepth/Pertinax/wiki/Reference%20Architecture>

## AutorunsToWinEventLog

Another method for integrating Autoruns into your logging infrastructure is AutorunsToWinEventLog:  
<https://github.com/palantir/windows-event-forwarding/tree/master/AutorunsToWinEventLog>

### 9.6.2 Downloads

Download Autoruns here:

<https://download.sysinternals.com/files/Autoruns.zip>

Download ar-normalize.ps1 here:

<https://raw.githubusercontent.com/defensivedepth/Pertinax/master/normalize/ar-normalize.ps1>

### 9.6.3 More Information

---

#### Note:

For more information about Autoruns, please see:

<https://docs.microsoft.com/en-us/sysinternals/downloads/autoruns>

---





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: Filebeat [IMPORT Node] → ES Ingest [IMPORT Node]  
Logs: Zeek, Suricata

### 10.1.2 Eval

Core Pipeline: Filebeat [EVAL Node] → ES Ingest [EVAL Node]  
Logs: Zeek, Suricata, Wazuh, Osquery/Fleet

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

### 10.1.3 Standalone

Core Pipeline: Filebeat [SA Node] → Logstash [SA Node] → Redis [SA Node] ↔ Logstash [SA Node] → ES Ingest [SA Node]  
Logs: Zeek, Suricata, Wazuh, Osquery/Fleet, syslog

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

### 10.1.4 Fleet Standalone

Pipeline: Filebeat [Fleet Node] → Logstash [M | MS] → ES Ingest [S | MS]

Logs: Osquery

### 10.1.5 Manager (separate search nodes)

Core Pipeline: Filebeat [Fleet | Forward] → Logstash [Manager] → Redis [Manager]

Logs: Zeek, Suricata, Wazuh, Osquery/Fleet, syslog

WinLogbeat: Winlogbeat [Windows Endpoint] → Logstash [Manager] → Redis [Manager]

Logs: WEL

### 10.1.6 Manager Search

Core Pipeline: Filebeat [Fleet | Forward] → Logstash [MS] → Redis [MS] ↔ Logstash [MS] → ES Ingest [MS]

Logs: Zeek, Suricata, Wazuh, Osquery/Fleet, syslog

Pipeline: Filebeat [MS] → Logstash [MS] → ES Ingest [MS]

Logs: Local Wazuh, Osquery/Fleet

WinLogbeat: Winlogbeat [Windows Endpoint] → Logstash [MS] → ES Ingest [MS]

Logs: WEL

### 10.1.7 Heavy

Pipeline: Filebeat [Heavy Node] → Logstash [Heavy] → Redis [Heavy] ↔ Logstash [Heavy] → ES Ingest [Heavy]

Logs: Zeek, Suricata, Wazuh, Osquery/Fleet, syslog

### 10.1.8 Search

Pipeline: Redis [Manager] → Logstash [Search] → ES Ingest [Search]

Logs: Zeek, Suricata, Wazuh, Osquery/Fleet, syslog

### 10.1.9 Forward

Pipeline: Filebeat [Forward] → Logstash [M | MS] → ES Ingest [S | MS]

Logs: Zeek, Suricata, Wazuh, syslog

## 10.2 Filebeat

From <https://www.elastic.co/beats/filebeat>:

Filebeat helps you keep the simple things simple by offering a lightweight way to forward and centralize logs and files.

On an Evaluation installation, Filebeat sends logs directly to *Elasticsearch*. For other installation types, Filebeat sends to *Logstash*.

### 10.2.1 Configuration

You can configure Filebeat inputs and output using *Salt*. An example of the filebeat pillar can be seen at <https://github.com/Security-Onion-Solutions/securityonion/blob/master/salt/filebeat/pillar.example>

Any inputs that are added to the pillar definition will be in addition to the default defined inputs. In order to prevent a *Zeek* log from being used as input, the `zeeklogs:enabled` pillar will need to be modified. The easiest way to do this is via *so-zeek-logs*.

### 10.2.2 Diagnostic Logging

Filebeat's log can be found in `/opt/so/log/filebeat/`.

To debug Filebeat, copy `/opt/so/saltstack/default/salt/filebeat/etc/filebeat.yml` to `/opt/so/saltstack/local/salt/filebeat/etc/filebeat.yml`, then change the `logging.level` value to `debug`. Next, restart Filebeat with `so-filebeat-restart`. Be sure to remove the local file after debugging.

Depending on what you're looking for, you may also need to look at the *Docker* logs for the container:

```
sudo docker logs so-filebeat
```

### 10.2.3 Modules

We support official Filebeat modules and you can learn more at <https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-modules.html>.

#### Example 1: AWS Cloudtrail Logs

If you would like to parse AWS Cloudtrail logs using the Filebeat `cloudtrail` module, you can enable the Filebeat module on any nodes that require it. Depending on your deployment, you might add the following configuration to the global pillar in `global.sls`, the manager's minion pillar in `/opt/so/saltstack/local/pillar/minions/$managename_manager.sls`, and/or the search node pillars in `/opt/so/saltstack/local/pillar/minions/`. If you have a distributed deployment using cross cluster search, then you will need to enable it for the manager and each search node. If you have a distributed deployment using Elastic clustering, then it only needs to be enabled for the manager.

Here's the configuration:

```
filebeat:
  third_party_filebeat:
    modules:
      aws:
```

(continues on next page)

(continued from previous page)

```
cloudtrail:
  enabled: true
  var.queue_url: https://sqs.$REGION.amazonaws.com/$ACCOUNTID/$QUEUENAME
  var.access_key_id: ABCD1234
  var.secret_access_key: ABCD1234ABCD1234
```

Access key details can be found within the AWS console by navigating to My Security Credentials -> Access Keys.

## Example 2: Fortinet Logs

If you would like to parse Fortinet logs using the Filebeat `fortinet` module, you can enable the Filebeat module on any nodes that require it. Depending on your deployment, you might add the following configuration to the global pillar in `global.sls`, the manager's minion pillar in `/opt/so/saltstack/local/pillar/minions/$managename_manager.sls`, and/or the search node pillars in `/opt/so/saltstack/local/pillar/minions/`. If you have a distributed deployment using cross cluster search, then you will need to enable it for the manager and each search node. If you have a distributed deployment using Elastic clustering, then it only needs to be enabled for the manager.

Here's the configuration:

```
filebeat:
  third_party_filebeat:
    modules:
      fortinet:
        firewall:
          enabled: true
          var.input: udp
          var.syslog_host: 0.0.0.0
          var.syslog_port: 9004
```

(Please note that *Firewall* ports still need to be opened on the minion to accept the Fortinet logs.)

## Walkthrough: AWS Cloudtrail Logs

In this brief walkthrough, we'll use the `aws` module for Filebeat to ingest `cloudtrail` logs from Amazon Web Services into Security Onion.

Credit goes to Kaiyan Sheng and Elastic for having an excellent starting point on which to base this walkthrough: <https://www.elastic.co/blog/getting-aws-logs-from-s3-using-filebeat-and-the-elastic-stack>.

Please follow the steps below to get started.

The official Elastic documentation for the Google Workspace module can be found here:

<https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-aws.html>

NOTE: This module requires that the user have a valid AWS service account, and credentials/permissions to access to the SQS queue we will be configuring.

### AWS Cloudtrail Configuration

*Create an SQS queue:*

Navigate to Amazon SQS -> Queues, and click `Create queue`.

Specify queue details, choosing to use a `Standard` queue, and providing a name:

## Details

### Type

Choose the queue type for your application or cloud infrastructure.

 You can't change the queue type after you create a queue.

### ☒ Standard [Info](#)

At-least-once delivery, message ordering isn't preserved

- At-least once delivery
- Best-effort ordering

### Name

demo-queue

A queue name is case-sensitive and can have up to 80 characters. You can use alphanumeric characters, hyphens (-),

Specify an Advanced policy and add policy configuration (changing to suit your environment, as needed):

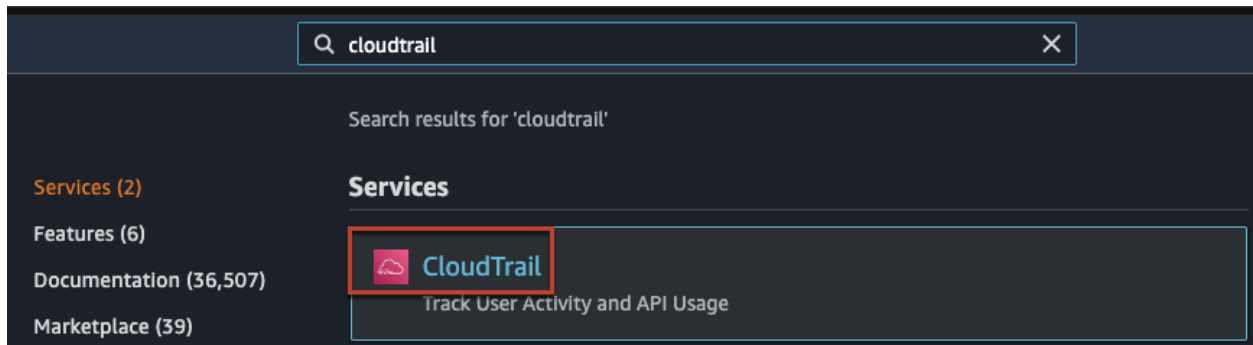
```
{
  "Version": "2012-10-17",
  "Id": "example-ID",
  "Statement": [
    {
      "Sid": "example-statement-ID",
      "Effect": "Allow",
      "Principal": {
        "Service": "s3.amazonaws.com"
      },
      "Action": [
        "SQS:SendMessage"
      ],
      "Resource": "arn:aws:sqs:<region>:<account-id>:<queue-name>",
      "Condition": {
        "StringEquals": { "aws:SourceAccount": "<account-id>" }
      }
    }
  ]
}
```

After the queue has been created, you will be redirected to a summary screen.

From here, copy the provided URL value. This value will be used to populate the queue URL in Security Onion's Filebeat configuration.

### Create a Trail:

We'll create a trail using the AWS Cloudtrail console. To get to the Cloudtrail console, search for `cloudtrail` in the AWS search bar at the top of the screen within the main console, and select CloudTrail:



From the main page of the Cloudtrail console, we can create our trail by clicking `Create a trail`:

## Create a trail with AWS CloudTrail

Get started with AWS CloudTrail by creating a trail to log your AWS account activity.

**Create a trail**

Next, we'll configure some basic details, and choose to use a new s3 bucket with our trail:

CloudTrail > Dashboard > Create trail

Step 1  
**Choose trail attributes**

---


Step 2  
Choose log events

---

Step 3  
Review and create

## Choose trail attributes

### General details

A trail created in the console is a multi-region trail. [Learn more](#) 

### Trail name

Enter a display name for your trail.

demo-trail

3-128 characters. Only letters, numbers, periods, underscores, and dashes

☐ Enable for all accounts in my organization

To review accounts in your organization, open AWS Organizations. [See all](#)

### Storage location [Info](#)

☒ Create new S3 bucket

Create a bucket to store logs for the trail.

We'll also need to specify an alias for a KMS key:

#### Customer managed AWS KMS key

- ☒ New
- ☐ Existing

#### AWS KMS alias

demo-trail-key

KMS key and S3 bucket must be in the same region.

Scroll down, and click Next.

From here, we'll select the type of log events we want to include with our trail:

CloudTrail > Dashboard > Create trail

Step 1  
[Choose trail attributes](#)

Step 2  
**Choose log events**

Step 3  
[Review and create](#)


## Choose log events

**Events** [Info](#)  
Record API activity for individual resources, or for all current and future resources in AWS account. [Additional charges apply](#)

**Event type**  
Choose the type of events that you want to log.

<input checked="" type="checkbox"/> <b>Management events</b> Capture management operations performed on your AWS resources.	<input type="checkbox"/> <b>Data events</b> Log the resource operations performed on or within a resource.	<input type="checkbox"/> <b>Insights events</b> Identify unusual activity, error behavior in your account.
--	---	---

**Management events** [Info](#)  
Management events show information about management operations performed on resources in your AWS account.

 Charges apply to log management events on this trail because you are logging at least one other copy of management events in your account.

**API activity**  
Choose the activities you want to log.

<input checked="" type="checkbox"/> <b>Read</b>	<input checked="" type="checkbox"/> <b>Write</b>
<input type="checkbox"/> Exclude AWS KMS events	
<input type="checkbox"/> Exclude Amazon RDS Data API events	

We'll then review our changes and click **Create Trail**:



CloudTrail > Dashboard > Create trail

Step 1

Choose trail attributes

---

Step 2

Choose log events

---

Step 3

**Review and create**

## Review and create

### Step 1: Choose trail attributes

#### General details

Trail name

demo-trail

Multi-region trail

Yes

Apply trail to my organization

Not enabled

The trail should now be created and viewable in Cloudtrail -> Trails. The Status column should display as Logging. Because we chose to create a new bucket when creating the trail, an s3 bucket should already be created.

We'll need to ensure our bucket is configured correctly by modifying the event notification properties. To do this, we'll navigate to Amazon S3 -> \$BucketName -> Properties -> Event notifications -> Create event notification:

**Event notifications (1)**


Send a notification when specific events occur in your bucket. [Learn more](#)

Edit Delete **Create event notification**

<input type="checkbox"/>	Name	Event types	Filters	Destination type	Destination
<input type="checkbox"/>	test-for-so	All object create events	-	SQS queue	<a href="#">demo-queue</a>

Under Event Types, we can select the type of events for which we would like to receive notifications to our SQS queue:

## Event types

Specify at least one type of event for which you want to receive notifications. [Learn more](#) 

- ☒ All object create events
  - s3:ObjectCreated:\*
  - ☒ Put
    - s3:ObjectCreated:Put
  - ☒ Post
    - s3:ObjectCreated:Post
  - ☒ Copy
    - s3:ObjectCreated:Copy
  - ☒ Multipart upload completed
    - s3:ObjectCreated:CompleteMultipartUpload

We'll also need to select the queue where events will be published:

### Specify SQS queue

- ☒ Choose from your SQS queues
- ☐ Enter SQS queue ARN

SQS queue

demo-queue ▼

If we would like to log bucket access events, we can enable `Server Access Logging` (within the bucket `Properties` section):

## Server access logging

Log requests for access to your bucket.

### Server access logging

- ☐ Disable
- ☒ Enable

### Security Onion Configuration

Now that we've configured our Cloudtrail trail and SQS queue, we need to place our credential information into our


Filebeat module configuration within Security Onion. In this example, we'll edit the minion pillar for the node we want to pull in the AWS Cloudtrail logs – in this case, a standalone node. In a distributed environment, this would likely be the manager node.

Edit `/opt/so/saltstack/local/pillar/minions/$minion_standalone.sls`, adding the following configuration (if you are already using other modules, simply append the module specific configuration without adding the `filebeat.third_party_filebeat.modules` portion):

```
filebeat:
  third_party_filebeat:
    modules:
      aws:
        cloudtrail:
          enabled: true
          var.queue_url: https://sqs.us-east-2.amazonaws.com/$youraccountid/demo-queue
          var.access_key_id: ABCDE1234
          var.secret_access_key: AbCdeFG...
```

Next, restart Filebeat on the node, with `so-filebeat-restart`.

After a few minutes, assuming there are logs to be gathered, Filebeat should pull in those logs from AWS, and an Elasticsearch index named `so-aws-$DATE` should be created. This can be verified by navigating to Hunt or Kibana, searching for `event.module:aws`:

 event.action	GetTrailStatus
 event.created	2021-07-16T14:08:53.770Z
 event.dataset	aws.cloudtrail
 event.id	65e55ad3-5919-46ee-b260-71e7c3c2d10d
 event.ingested	2021-07-16T14:08:55.507497011Z
 event.kind	event
 event.module	aws

We can also run the `so-elasticsearch-query` command, like so:

`so-elasticsearch-query _cat/indices | grep aws`

```
[root@securityonion-stdl-testing ~]# so-elasticsearch-query _cat/indices | grep aws
green open so-aws-2021.07.16          I2FStBwSSJi1_PZSxQRK6Q 1 0      672    0
```

Congratulations! You've ingested AWS Cloudtrail logs into Security Onion!

## Walkthrough: Google Workspace Audit Logs

In this brief walkthrough, we'll use the `google_workspace` module for Filebeat to ingest `admin` and `user_accounts` logs from Google Workspace into Security Onion.

Please follow the steps below to get started.

The official Elastic documentation for the Google Workspace module can be found here:

[https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-google\\_workspace.html](https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-google_workspace.html)

NOTE: This module requires that the user have a valid Google Workspace administrator account. You'll also need to set up a project within Google Cloud if that has not already been done (will set up as needed during the walkthrough).

### Google Cloud/Workspace Configuration

Google provides documentation for setting up a service account here:

<https://support.google.com/workspacemigrate/answer/9222993?hl=en>

In this example, we'll choose the automated method of service account creation (using a script and the Cloud Shell).

We can enter the Cloud Shell by clicking the Cloud Shell icon (right-hand side of screen) from **console.cloud.google.com** (signed in as our Google Workspaces Super Administrator):



Once opened, we will run the following command:

```
python3 <(curl -s -S -L https://git.io/gwm-create-service-account)
```

```
Welcome to Cloud Shell! Type "help" to get started.
Your Cloud Platform project in this session is set to dfirlab.
Use "gcloud config set project [PROJECT_ID]" to change to a different project.
wes@cloudshell:~ (dfirlab)$ python3 <(curl -s -S -L https://git.io/gwm-create-service-account)
```

After running the command, we will be provided a menu (press Enter to continue):

```
Welcome! This script will create and authorize the resources that are necessary to use Google Workspace Migrate. The following steps will be performed on your behalf:
1. Create a Google Cloud Platform project
2. Enable APIs
3. Create a service account
4. Authorize the service account
5. Create a service account key

In the end, you will be prompted to download the service account key. This key can then be used for GWM.

If you would like to perform these steps manually, then you can follow the instructions at <https://support.google.com/workspacemigrate/answer/10839762>.

Press Enter to continue or 'n' to exit:
```

The script will proceed through the steps until the first phase of setup is complete:

```
Creating project...
gwm-1626103894017 successfully created ✓
Verifying acceptance of Terms of service...
Terms of service acceptance verified ✓
Enabling APIs...
APIs successfully enabled ✓
Creating service account...
gwm-service-account successfully created ✓
```

After the first phase of setup, you will be provided a URL to visit and authorize the changes. When authorizing changes, make sure to add the following OAuth scope to the client:

<https://www.googleapis.com/auth/admin.reports.audit.readonly>

Add a new client ID

https://www.googleapis.com/auth/admin.directory.␣

OAuth scopes (comma-delimited)
✕

https://www.googleapis.com/auth/admin.directory.␣

OAuth scopes (comma-delimited)
✕

.googleapis.com/auth/admin.reports.audit.readonly␣

OAuth scopes (comma-delimited)

CANCEL AUTHORIZE

Navigate back to the Cloud Shell and press Enter to proceed through the rest of the setup:

```

Creating service account key...
Service account key successfully created ✓
Verifying service account authorization...
Service account successfully authorized ✓
Verifying API access...
API access verified ✓
Done! ✓

If you have already downloaded the file, then you may close this page. Please remember that this file is highly sensitive.
access. You should treat it just like you would a password.

```

You will be prompted to download a file containing the service account credentials:

## Download File

Please confirm that you wish to download the following files:

/tmp/gwm-service-account-key-2021-07-12-15-30-26.json

Download

Cancel

Ensure this file is kept safe. We will provide it to Filebeat in the Security Onion Filebeat module configuration.

## Security Onion Configuration

Now that we've set up a service account and obtained a credentials file, we need to place it into our Filebeat module configuration within Security Onion. In this example, we'll edit the minion pillar for the node we want to pull in the Google Workspace logs – in this case, a standalone node. In a distributed environment, this would likely be the manager node.

Copy the credentials file to `/opt/so/conf/filebeat/modules/` as `credentials_file.json`.

Edit `/opt/so/saltstack/local/pillar/minions/$minion_standalone.sls`, adding the following configuration (if you are already using other modules, simply append the module specific configuration without adding the `filebeat.third_party_filebeat.modules` portion):

```
filebeat:
  third_party_filebeat:
    modules:
      google_workspace:
        admin:
          enabled: true
          var.jwt_file: "/usr/share/filebeat/modules.d/credentials_file.json"
          var.delegated_account: "adminuser@yourdomain.com"
        user_accounts:
          enabled: true
          var.jwt_file: "/usr/share/filebeat/modules.d/credentials_file.json"
          var.delegated_account: "adminuser@yourdomain.com"
```

Next, restart Filebeat on the node, with `so-filebeat-restart`.

After a few minutes, assuming there are logs to be gathered, Filebeat should pull in those logs from Google Workspace, and an Elasticsearch index named `so-google_workspace-$DATE` should be created. This can be verified by navigating to Hunt or Kibana, searching for `event.module:google_workspace`:

event.created	2021-07-12T17:00:27.732Z
event.dataset	google_workspace.admin
event.id	-1381061208029214585
event.ingested	2021-07-12T17:00:29.523312099Z
event.module	google_workspace
event.original	{       "actor": {         "email": "wes@dfirlab.com",         "client_access": "CLIENT_ACCESS",         "parameters": {           "mailsettings/2.0/",           "https://www.googleapis.com/auth/admin.directory.orgunit",           "https://www.googleapis.com/auth/admin.directory.rolemanager",           "https://www.googleapis.com/auth/admin.directory.rolemanager"         }       },       "time": "2021-07-12T16:56:43.115Z",       "uniqueQ"     }
event.provider	admin
event.type	[       "change"     ]
fileset.name	admin

We can also run the `so-elasticsearch-query` command, like so:

```
so-elasticsearch-query _cat/indices | grep google_workspace
```

```
[root@securityonion-stdl-testing modules]# so-elasticsearch-query _cat/indices | grep google_workspace
green open so-google_workspace-2021.07.12 K4TftUGuRDdW-dBekcmW8Q 1 0 2 0 68.2kb 68.2kb
```

Congratulations! You've ingested Google Workspace logs into Security Onion!

## Walkthrough: Okta System Logs

In this brief walkthrough, we'll use the `okta` module for Filebeat to ingest system logs from Okta into Security Onion. Please follow the steps below to get started.

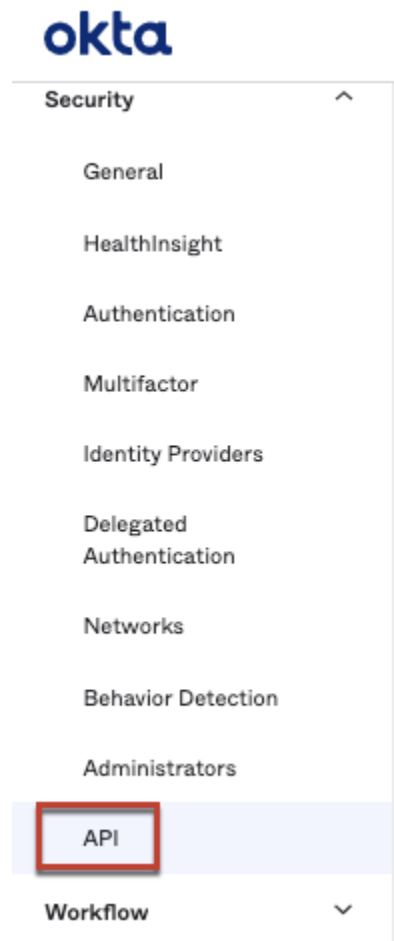
The official Elastic documentation for the Okta module can be found here:

<https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-okta.html>

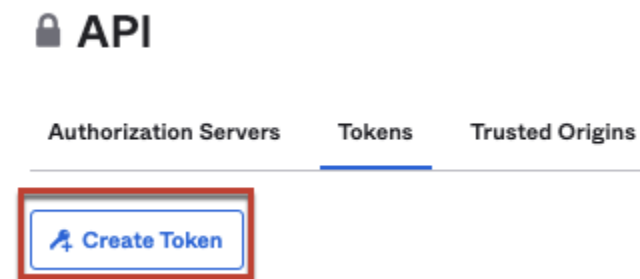
NOTE: This module requires that the user have a valid API token for access to their Okta instance.

### Okta Configuration

Within the Okta administrative console, from the pane on the left-hand side of the screen, navigate to `Security->API`.



Next, navigate to Tokens, and click Create Token:



Enter a name for the token, then click Create Token:



## Create Token

✕

What do you want your token to be named?

The token name is used for tracking API calls.



A confirmation message like the following should appear:

## Create Token

✕

**Token created successfully!**

Please make a note of this token as it will be the only time that you will be able to view it. After this, it will be stored as a hash for your protection.

Ensure the token provided below the message is saved and stored securely.

### Security Onion Configuration

Now that we've got our token, we need to place it into our Filebeat module configuration within Security Onion. In this example, we'll edit the minion pillar for the node we want to pull in the Okta logs – in this case, a standalone node. In a distributed environment, this would likely be the manager node.

Edit `/opt/so/saltstack/local/pillar/minions/$minion_standalone.sls`, adding the following configuration (if you are already using other modules, simply append the module specific configuration without adding the `filebeat.third_party_filebeat.modules` portion):

```
filebeat:
  third_party_filebeat:
    modules:
      okta:
        system:
          enabled: true
          var.url: https://$yourdomain/api/v1/logs
          var.api_key: '$yourtoken'
```

Next, restart Filebeat on the node, with `so-filebeat-restart`.

After a few minutes, assuming there are logs to be gathered, Filebeat should pull in those logs from Okta, and an Elasticsearch index named `so-okta-$DATE` should be created. This can be verified by navigating to Hunt or Kibana, searching for `event.module:okta`:

 okta.client.user_agent.browser	CHROME
 okta.client.user_agent.os	Mac OS X
 okta.client.user_agent.raw_user_agent	Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7)
 okta.client.zone	null
 okta.debug_context.debug_data.request_id	YOxINfD9AgaS9EyC4WGwkQAACeQ
 okta.debug_context.debug_data.request_uri	/api/internal/tokens
 okta.debug_context.debug_data.url	/api/internal/tokens?expand=user
 okta.display_message	Create API token
 okta.event_type	system.api_token.create
 okta.outcome.result	SUCCESS
 okta.security_context.as.number	12083
 okta.security_context.as.organization.name	wideopenwest finance
 okta.security_context.domain	knology.net
 okta.security_context.is_proxy	false
 okta.security_context.isp	wideopenwest finance
 okta.target	[ { "alternate_id": "unknown", "display_name": "test-token2",

We can also run the `so-elasticsearch-query` command, like so:

```
so-elasticsearch-query _cat/indices | grep okta
```

```
[root@securityonion-stdl-testing pillar]# so-elasticsearch-query _cat/indices | grep okta
green open so-okta-2021.07.12      C3tjZq9uS0mk1GHgSpao5w 1 0      10      0      208b      208b
```

Congratulations! You've ingested Okta logs into Security Onion!

## Walkthrough: Netflow Logs

In this brief walkthrough, we'll use the `net.flow` module for Filebeat to ingest Netflow logs into Security Onion.

---

**Note:** Check out our Netflow video at <https://youtu.be/ew5gtVjAs7g!>

---

The official Elastic documentation for the Netflow module can be found here:

<https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-module-netflow.html>

Overview of steps:

- enable third party module
- update docker config
- update firewall config
- build logstash pipeline

### Enable third party module

If you would like to ingest Netflow logs using the Filebeat `netflow` module, you can enable the Filebeat module on any nodes that require it. Depending on your deployment, you might add the following configuration to the global pillar in `global.sls`, the manager's minion pillar in `/opt/so/saltstack/local/pillar/minions/$managename_manager.sls`, and/or the search node pillars in `/opt/so/saltstack/local/pillar/minions/`. If you have a distributed deployment using cross cluster search, then you will need to enable it for the manager and each search node. If you have a distributed deployment using Elastic clustering, then it only needs to be enabled for the manager.

Here's the configuration:

```
filebeat:
  third_party_filebeat:
    modules:
      netflow:
        log:
          enabled: true
          var.netflow_host: 0.0.0.0
          var.netflow_port: 2055
```

### Update docker config

Next, we need to add an extra listening port to the Filebeat container. We'll start by making a local copy the `filebeat init.sls` file.

```
sudo cp /opt/so/saltstack/default/salt/filebeat/init.sls /opt/so/saltstack/local/salt/
↪ filebeat/init.sls
```

Next, set permissions on the file:

```
sudo chown socore:socore /opt/so/saltstack/local/salt/filebeat/init.sls
```

Edit `/opt/so/saltstack/local/salt/filebeat/init.sls` and add port 2055 to the `port_bindings` section of the `so-filebeat` config:

```
- port_bindings:
  - 0.0.0.0:514:514/udp
  - 0.0.0.0:514:514/tcp
  - 0.0.0.0:2055:2055/udp
  - 0.0.0.0:5066:5066/tcp
```

Save the file and run `sudo salt-call state.apply filebeat` to allow *Salt* to recreate the container. You can check that the config has applied by running `sudo docker ps | grep so-filebeat`. You should see `0.0.0.0:2055->2055/udp` among the other existing listening ports.

### Update firewall config

The next step is to add a host group and port group for Netflow traffic to allow it through the firewall. Replace `172.30.0.0/16` with whatever is appropriate for your network.

```
sudo so-firewall addhostgroup netflow
sudo so-firewall addportgroup netflow
sudo so-firewall includehost netflow 172.30.0.0/16
sudo so-firewall addport netflow udp 2055
```

Edit `/opt/so/saltstack/local/pillar/minions/<manager.sls>` to add iptables rules to allow the new netflow groups:

```
firewall:
  assigned_hostgroups:
    chain:
      DOCKER-USER:
        hostgroups:
          netflow:
            portgroups:
              - portgroups.netflow
      INPUT:
        hostgroups:
          netflow:
            portgroups:
              - portgroups.netflow
```

Save the file and then run `sudo salt-call state.apply firewall` to enable the new firewall rules.

### Build logstash pipeline

Now the module is enabled, the container is listening on the right port, and the firewall is allowing traffic to get to the container. Next is to ensure that the Netflow pipeline is enabled, or the data will not be saved to the ES database.

Note: If you have a distributed setup, you need to run the following command on the search nodes as well:

```
sudo docker exec -i so-filebeat filebeat setup modules --pipelines --modules netflow -M
↪ "netflow.log.enabled=true" -c /usr/share/filebeat/module-setup.yml
```

You should see `Loaded Ingest pipelines`. Once that is complete, run `sudo so-filebeat-restart`.

Assuming you have Netflow sources sending data, you should now start to see data in [Dashboards](#) or [Hunt](#). Group by `event.dataset` and you should now have `netflow.log` entries appearing.

## 10.2.4 More Information

---

**Note:** For more information about Filebeat, please see <https://www.elastic.co/beats/filebeat>.

---

## 10.3 Logstash

From <https://www.elastic.co/products/logstash> :

Logstash is a free and open server-side data processing pipeline that ingests data from a multitude of sources, transforms it, and then sends it to your favorite “stash.”

When Security Onion 2 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.

### 10.3.1 Configuration

You can configure Logstash using *Salt*. Here are a few of the settings which you may need to tune in `/opt/so/saltstack/local/pillar/minions/$MINION_ROLE.sls` under `logstash_settings`.

#### `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](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.

### 10.3.2 Parsing

Since Logstash no longer parses logs in Security Onion 2, modifying existing parsers or adding new parsers should be done via [Elasticsearch](#).

### 10.3.3 Adding New Logs

If you want to add a new log to the list of logs that are sent to [Elasticsearch](#) for parsing, you can update the logstash pipeline configurations by adding to `/opt/so/saltstack/local/salt/logstash/pipelines/config/custom/`.

If you are modifying or adding a new manager pipeline, then first copy `/opt/so/saltstack/default/pillar/logstash/manager.sls` to `/opt/so/saltstack/local/pillar/logstash/`, then add the following to the `manager.sls` file under the local directory:

```
logstash:
  pipelines:
    manager:
      config:
        - so/0009_input_beats.conf
        - so/0010_input_hhbeats.conf
        - so/9999_output_redis.conf.jinja
        - custom/9999_output_custom.conf.jinja
```

If you are modifying or adding a new search pipeline for all search nodes, then first copy `/opt/so/saltstack/default/pillar/logstash/search.sls` to `/opt/so/saltstack/local/pillar/logstash/`, then add the following to the `search.sls` file under the local directory:

```
logstash:
  pipelines:
    search:
      config:
        - so/0900_input_redis.conf.jinja
        - so/9000_output_zeek.conf.jinja
        - so/9002_output_import.conf.jinja
        - so/9034_output_syslog.conf.jinja
        - so/9100_output_osquery.conf.jinja
        - so/9400_output_suricata.conf.jinja
        - so/9500_output_beats.conf.jinja
        - so/9600_output_ossec.conf.jinja
        - so/9700_output_strelka.conf.jinja
        - custom/9701_output_custom.conf.jinja
```

If you only want to modify the search pipeline for a single search node, then the process is similar to the previous example. However, instead of placing `logstash:pipelines:search:config` in `/opt/so/saltstack/local/pillar/logstash/search.sls`, it would be placed in `/opt/so/saltstack/local/pillar/minions/$hostname_searchnode.sls`.

### 10.3.4 Logstash Parsing

If you want to add a legacy Logstash parser (not recommended) then you can copy the file to `local`. Once the file is in `local`, then depending on which nodes you want it to apply to, you can add the proper value to either `/opt/so/saltstack/local/pillar/logstash/manager.sls`, `/opt/so/saltstack/local/pillar/logstash/search.sls`, or `/opt/so/saltstack/local/pillar/minions/$hostname_searchnode.sls` as in the previous examples.

### 10.3.5 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>.

### 10.3.6 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 [module] =~ "zeek" and [dataset] =~ "dns" {
    udp {
      id => "cloned_events_out"
      host => "192.168.x.x"
      port => 1001
      codec => "json_lines"
    }
  }
}
```

**Warning:** When using the `tcp` output plugin, if the destination host/port is down, it will cause the Logstash pipeline to be blocked. To avoid this behavior, 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, Suricata's `dataset` value will still be set to `common`, as the events have not yet been processed by the Ingest Node configuration.

Copy `/opt/so/saltstack/default/pillar/logstash/manager.sls` to `/opt/so/saltstack/local/pillar/logstash/manager.sls`, and append your newly created file to the list of config files used for the manager pipeline:

- `custom/myfile.conf`

Restart Logstash on the manager with `so-logstash-restart`.

Monitor events flowing through the output with `curl -s localhost:9600/_node/stats | jq .pipelines.manager`.

### 10.3.7 Modified Event Forwarding

To forward events to an external destination AFTER they have traversed the Logstash pipelines (NOT ingest node pipelines) used by Security Onion, perform the same steps as above, but instead of adding the reference for your Logstash output to `manager.sls`, add it to `search.sls` instead, and then restart services on the search nodes with something like:

```
sudo salt "*_search*" cmd.run "so-logstash-restart"
```

Monitor events flowing through the output with `curl -s localhost:9600/_node/stats | jq .pipelines.search` on the search nodes.

Please keep in mind that events will be forwarded from all applicable search nodes, as opposed to just the manager.

### 10.3.8 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.

### 10.3.9 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
```

### 10.3.10 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 }'
```

### 10.3.11 More Information

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

## 10.4 Redis

From <https://redis.io/>:

Redis is an open source (BSD licensed), in-memory data structure store, used as a database, cache and message broker. It supports data structures such as strings, hashes, lists, sets, sorted sets with range queries, bitmaps, hyperloglogs and geospatial indexes with radius queries.

On Standalone (non-Eval) installations and distributed deployments, *Logstash* on the manager node outputs to Redis. Search nodes can then consume from Redis.

### 10.4.1 Queue

To see how many logs are in the Redis queue:

```
sudo so-redis-count
```

If the queue is backed up and doesn't seem to be draining, try stopping *Logstash* on the manager node:

```
sudo so-logstash-stop
```

Then monitor the queue to see if it drains:

```
watch 'sudo so-redis-count'
```

If the Redis queue looks okay, but you are still having issues with logs getting indexed into *Elasticsearch*, you will want to check the *Logstash* statistics on the search node(s).

### 10.4.2 Tuning

We configure Redis to use 812MB of your total system memory. If you have sufficient RAM available, you may want to increase the `redis_maxmemory` setting in `/opt/so/saltstack/local/pillar/global.sls`. This value is in Megabytes so to set it to use 8 gigs of ram you would set the value to 8192.

*Logstash* on the manager node is configured to send to Redis. For best performance, you may want to ensure that `batch` is set to `true` and then tune the `ls_pipeline_batch_size` variable 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` 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>

### 10.4.3 Diagnostic Logging

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

```
sudo docker logs so-redis
```

### 10.4.4 More Information

---

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

---

## 10.5 Elasticsearch

From <https://www.elastic.co/products/elasticsearch>:

Elasticsearch is a distributed, RESTful search and analytics engine capable of addressing a growing number of use cases. As the heart of the Elastic Stack, it centrally stores your data for lightning fast search, fine-tuned relevancy, and powerful analytics that scale with ease.

### 10.5.1 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*.

### 10.5.2 Authentication

We support Elastic authentication via *so-elastic-auth*.

### 10.5.3 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
```

## 10.5.4 Storage

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

## 10.5.5 Parsing

In Security Onion 2, Elasticsearch receives unparsed logs from *Logstash* or *Filebeat*. 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`.

---

### Note:

For more about Elasticsearch ingest parsing, please see:

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

---

## 10.5.6 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`, and these references can be modified in the `elasticsearch` *Salt* pillar:

## Custom Templates

To add a custom index template, ensure the custom or modified component templates are copied to `/opt/so/saltstack/local/salt/elasticsearch/templates/component/so/`.

Next, copy `/opt/so/saltstack/default/pillar/elasticsearch/index_templates.sls` to `/opt/so/saltstack/local/pillar/elasticsearch/`.

Edit the file similar to the following, adding your custom index template details and the references to the component templates you wish to associate to the index template:

`/opt/so/saltstack/local/pillar/elasticsearch/index_templates.sls`

```
elasticsearch:
  index_settings:
    so-custom:
      index_sorting: False
      index_template:
        index_patterns:
          - so-custom*
        template:
          mappings:
            dynamic_templates:
              - strings_as_keyword:
                  mapping:
                    ignore_above: 1024
                    type: keyword
                  match_mapping_type: string
            date_detection: false
          settings:
            index:
              mapping:
                total_fields:
                  limit: 1500
              sort:
                field: "@timestamp"
                order: desc
              refresh_interval: 30s
              number_of_shards: 1
              number_of_replicas: 0
            composed_of:
              - custom-mappings
              - custom-settings
          priority: 500
```

Next, apply the Elasticsearch state for the relevant nodes (or wait for the next highstate):

```
sudo salt-call state.apply elasticsearch
```

Upon successful application, the resultant index template will be created in `/opt/so/conf/elasticsearch/templates/index` with a filename that consists of the custom index key value (`so-custom` in this case) and a static `-template.json` suffix. We can check to see if the file exists and check the contents of the file with the following command:

```
cat /opt/so/conf/elasticsearch/templates/index/so-custom-template.json
```

We can also check to ensure that both the associated component templates and the index template itself were loaded into Elasticsearch:

```
so-elasticsearch-component-templates-list | grep custom
so-elasticsearch-index-templates-list | grep custom
```

## 10.5.7 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>

## 10.5.8 Configuration

### Pillar Files

All configuration changes take place in *Salt* pillar files. There are two places that hold pillar settings for Elasticsearch. The pillars are:

/opt/so/saltstack/local/pillar/minions/\$minion.sls

```
elasticsearch:
  mainip: 10.66.166.22
  mainint: eth0
  esheap: 4066m
  esclustername: {{ grains.host }}
  node_type: search
  es_port: 9200
  log_size_limit: 3198
  node_route_type: hot
```

/opt/so/saltstack/local/pillar/global.sls

```
elasticsearch:
  true_cluster: False
  replicas: 0
  discovery_nodes: 1
  hot_warm_enabled: False
  cluster_routing_allocation_disk.threshold_enabled: true
  cluster_routing_allocation_disk.watermark_low: '95%'
  cluster_routing_allocation_disk.watermark_high: '98%'
  cluster_routing_allocation_disk.watermark_flood_stage: '98%'
  script.painless.regex.enabled: true
  index_settings:
    so-beats:
      index_template:
        template:
          settings:
            index:
              number_of_shards: 1
      warm: 7
      close: 30
```

(continues on next page)

(continued from previous page)

```

    delete: 365
so-endgame:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7
  close: 30
  delete: 365
so-firewall:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7
  close: 30
  delete: 365
so-flow:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
          close: 45
          delete: 365
so-ids:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7
  close: 30
  delete: 365
so-import:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7
  close: 73000
  delete: 73001
so-osquery:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7

```

(continues on next page)

(continued from previous page)

```
close: 30
delete: 365
so-ossec:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7
  close: 30
  delete: 365
so-strelka:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7
  close: 30
  delete: 365
so-syslog:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 1
  warm: 7
  close: 30
  delete: 365
so-zeek:
  index_template:
    template:
      settings:
        index:
          number_of_shards: 2
```

## Customization

You can completely customize your Elasticsearch configuration via [Salt](#) pillars. This allows `elasticsearch.yml` customizations to be retained when doing upgrades of Security Onion. Depending on your customization goal, you can specify settings in either the global pillar or the minion pillar. Create the `config` sub-section if it does not already exist in your pillar and then place your configuration options under that sub-section. For example, to change the `node_concurrent_recoveries` setting:

```
elasticsearch:
  config:
    routing:
      allocation:
        node_concurrent_recoveries: 4
```



**Warning:** Please be very careful when adding items under the `config` sub-section to avoid typos and other errors that would interfere with Elasticsearch. After making changes, keep a close eye on Elasticsearch to make sure the change is working as intended.

### 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 customize the `indices.query.bool.max_clause_count` value for any boxes running Elasticsearch in your deployment.

```
elasticsearch:
  config:
    indices.query.bool.max_clause_count: 4000
```

### 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-query _cat/indices
```

If you want to view the detail for each of those shards:

```
sudo so-elasticsearch-query _cat/shards
```

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 is defined in `/opt/so/saltstack/local/pillar/global.sls`. You can adjust shard counts for each index individually to meet your needs. The next time the node checks in, it will apply the settings automatically.

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. This can be modified in `/opt/so/saltstack/local/pillar/minions/$minion.sls`.

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/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 with one of the following options.

### Temporary

If you only need to increase the field limit temporarily, you can do something like:

```
curl -k -XPUT -H'Content-Type: application/json' https://localhost:9200/logstash-syslog-
  ↪*/_settings -d'{ "index.mapping.total_fields.limit": 6000 }'
```

The above command would increase the field limit for the `logstash-syslog-*` indice(s) to 6000. Keep in mind, this setting only applies to the current index, so when the index rolls over and a new one is created, your new settings will not apply.

### Persistent

If you need this change to be persistent, you can modify the `settings` stanza for the matched indices in the template:

```
"settings" : {
  "number_of_replicas": 0,
  "number_of_shards": 1,
  "index.refresh_interval" : "5s",
  "index.mapping.total_fields.limit": 6000
},
```

Then restart Logstash:

```
sudo so-logstash-restart
```

Please note that the change to the field limit will not occur immediately, only on index creation. Therefore, it is recommended to run the previously mentioned temporary command and modify the template file.

### 10.5.9 Closing Indices

Elasticsearch indices are closed based on the `close` setting shown in the global pillar above. This setting configures *Curator* to close any index older than the value given. The more indices are open, the more heap is required. Having too many open indices can lead to performance issues. There are many factors that determine the number of days you can have in an open state, so this is a good setting to adjust specific to your environment.

### 10.5.10 Deleting Indices

---

**Note:** This section describes how Elasticsearch indices are deleted in standalone deployments and distributed deployments using our default deployment method of cross cluster search. Index deletion is different for deployments using Elastic clustering and that is described in the Elastic clustering section later.

---

For standalone deployments and distributed deployments using cross cluster search, Elasticsearch indices are deleted based on the `log_size_limit` value in the minion pillar. If your open indices are using more than `log_size_limit` gigabytes, then *Curator* will delete old open indices until disk space is back under `log_size_limit`. If your total Elastic disk usage (both open and closed indices) is above `log_size_limit`, then `so-curator-closed-delete` will delete old closed indices until disk space is back under `log_size_limit`. `so-curator-closed-delete` does not use *Curator* because *Curator* cannot calculate disk space used by closed indices. For more information, see [https://www.elastic.co/guide/en/elasticsearch/client/curator/current/filtertype\\_space.html](https://www.elastic.co/guide/en/elasticsearch/client/curator/current/filtertype_space.html).

*Curator* and `so-curator-closed-delete` run on the same schedule. This might seem like there is a potential to delete open indices before deleting closed indices. However, keep in mind that *Curator*'s `delete.yml` is only going to see disk space used by open indices and not closed indices. So if we have both open and closed indices, we may be at `log_size_limit` but *Curator*'s `delete.yml` is going to see disk space at a value lower than `log_size_limit` and so it shouldn't delete any open indices.

For example, suppose our `log_size_limit` is 1TB and we have 30 days of open indices and 300 days of closed indices. We reach `log_size_limit` and both *Curator* and `so-curator-closed-delete` execute at the same time. *Curator*'s `delete.yml` will check disk space used but it will see that disk space is at maybe 100GB so it thinks we haven't reached `log_size_limit` and does not delete anything. `so-curator-closed-delete` gets a more accurate view of disk space used, sees that we have indeed reached `log_size_limit`, and so it deletes closed indices until we get lower than `log_size_limit`. In most cases, *Curator* deletion should really only happen if we have open indices without any closed indices.

### 10.5.11 Distributed Deployments

For distributed deployments, Security Onion 2 supports two different configurations for deploying Elasticsearch: cross cluster search and Elastic clustering.

#### Cross Cluster Search

Our traditional and default configuration for distributed Elasticsearch instances is `cross cluster search`. This means that each Elasticsearch instance is totally independent and the manager queries all Elasticsearch instances via cross cluster search. This lowers the amount of maintenance required and the required knowledge of Elasticsearch internals. This configuration is recommended for most users.

The `manager` node runs its own local copy of Elasticsearch, which manages cross-cluster search configuration for the deployment. This includes configuration for `search` nodes and `heavy` nodes (where applicable). This does not include `forward` nodes since they do not run Elastic Stack components.

Search nodes extend the storage and processing capabilities of the manager node, and run *Elasticsearch*, *Logstash*, and *Curator*. Search nodes are added to the manager node's cluster search configuration, so the data that resides on the nodes can be queried from the manager node.

Heavy nodes run sensor services and store their own logs in a local Elasticsearch instance. Heavy nodes are added to the manager node's cluster search configuration, so the data that resides on the nodes can be queried from the manager node. Heavy nodes are not recommended for most use cases.

When using a forward node, Elastic Stack components are not enabled. *Filebeat* forwards all logs 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 search nodes). From there, the data can be queried through the use of cross-cluster search.

## Elastic Clustering

For advanced users that require advanced features like shard replicas and hot/warm indices, Security Onion 2 also supports Elastic clustering. In this configuration, Elasticsearch instances join together to create a single cluster. However, please keep in mind that this requires more maintenance, more knowledge of Elasticsearch internals, and more traffic between nodes in the cluster.

**Warning:** Due to the increased complexity, we only recommend this option if you absolutely need cluster features.

Choose which type of manager to install:

```
( ) BASIC      Install manager with recommended settings
(*) ADVANCED   Do additional configuration to the manager
```


<Ok>

<Cancel>

Do you want to set up a traditional ES cluster for using replicas and/or Hot-Warm indices? Recommended only for those who have experience with ES clustering!

<Yes>

<No>



When using Elastic clustering, index deletion is based on the `delete` settings shown in the global pillar above. The `delete` settings in the global pillar configure *Curator* to delete indices older than the value given. For each index, please ensure that the `close` setting is set to a smaller value than the `delete` setting.

Let's discuss the process for determining appropriate `delete` settings. First, check your indices using *so-elasticsearch-query* to query `_cat/indices`. For example:

```
sudo so-elasticsearch-query _cat/indices | grep 2021.08.26
```

green	open	so-zeek-2021.08.26	rEtb1ERqQcyr7bfbnR95zQ	5	0	2514236	0	
→	2.4gb	2.4gb						
green	open	so-ids-2021.08.26	d3ySLbRHSJGRQ2oiS4pmMg	1	0	1385	147	
→	3.3mb	3.3mb						
green	open	so-ossec-2021.08.26	qYf1HWGUSn6fIOlOgFgJOQ	1	0	125333	61	
→	267.1mb	267.1mb						
green	open	so-elasticsearch-2021.08.26	JH8tOgr3QjaQ-EX080GEXw	1	0	61170	0	
→	32.7mb	32.7mb						
green	open	so-firewall-2021.08.26	Qx6_ZQS3QL6VGwIXIQ8mfQ	1	0	508799	0	
→	297.4mb	297.4mb						
green	open	so-syslog-2021.08.26	3HiYP3fgSPmoV-Nbs3dlDw	1	0	181207	0	
→	27mb	27mb						
green	open	so-kibana-2021.08.26	C6v6sazHSYiwqq5HxfokQg	1	0	745	0	
→	809.5kb	809.5kb						

Adding all the index sizes together plus a little padding results in 3.5GB per day. We will use this as our baseline.

If we look at our total `/nsm` size for our search nodes (data nodes in Elastic nomenclature), we can calculate how many days open or closed that we can store. The equation shown below determines the proper delete timeframe. Note that total usable space depends on replica counts. In the example below we have 2 search nodes with 140GB for 280GB total of `/nsm` storage. Since we have a single replica we need to take that into account. The formula for that is:

1 replica = 2 x Daily Index Size 2 replicas = 3 x Daily Index Size 3 replicas = 4 x Daily Index Size

Let's use 1 replica:

Total Space / copies of data = Usable Space

280 / 2 = 140

Suppose we want a little cushion so let's make Usable Space = 130

Usable NSM space / Daily Index Size = Days

For our example above lets fill in the proper values:

130GB / 3.5GB = 37.1428571 days rounded down to 37 days

Therefore, we can set all of our `delete` values to 37 in the `global.sls`.

### 10.5.12 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>

### 10.5.13 Clearing

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

### 10.5.14 GeoIP

Elasticsearch 8 no longer includes GeoIP databases by default, but we add GeoIP databases during the build process. 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
```

### 10.5.15 More Information

---

**Note:**

For more information about Elasticsearch, please see:

<https://www.elastic.co/products/elasticsearch>

---

## 10.6 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.

## 10.6.1 Configuration

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 `es_index_patterns` setting in your *Salt* pillar to include `*:elastalert*`:

```
es_index_patterns: '*:so-*,*:endgame-*,*: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). Keep in mind, this may impact performance negatively, so testing the change in a single file at a time may be the best approach.

### Timeframe

Keep in mind, 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>



## 10.6.2 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
```

## 10.6.3 More Information

---

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

---

## 10.7 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.

### 10.7.1 Configuration

Curator actions are stored in `/opt/so/conf/curator/action/`. These actions are run by cron jobs managed by *Salt*.

Curator defaults to closing indices older than 30 days. To modify this, edit `/opt/so/saltstack/local/pillar/global.sls` and change the close setting for each index under the `elasticsearch:index_settings` section.

---

**Note:** `/opt/so/saltstack/local/pillar/global.sls` only lists default indices. If you add any new indices, you will need to manually add them to that file.

---

Curator also deletes old indices to prevent your disk from filling up. The mechanism for this depends on how *Elasticsearch* is configured, so you can learn more in the *Elasticsearch* section.

For more information about the Curator close and delete settings, please see the *Elasticsearch* section.

## 10.7.2 Creating Actions

If you would like to create a custom Curator action, you will need to create a Curator action file and corresponding script file and then update Curator's state file.

### Curator action file

You can add the action file to `/opt/so/saltstack/local/salt/curator/files/action/`. Make sure the file ownership is `curator:socore`. The file will automatically get copied into `/opt/so/conf/curator/action/`.

### Script file

The script file is what actually executes Curator and specifies the action file. This script file must be placed in `/opt/so/saltstack/local/salt/curator/files/bin/`. See `/opt/so/saltstack/default/salt/curator/files/bin/` for examples of script files and copy one over to modify if needed.

### State file

Next, Curator's state file (`init.sls`) must be modified. This will be located at `/opt/so/saltstack/local/salt/curator/` and will copy files and create the cron job.

If `/opt/so/saltstack/local/salt/curator/init.sls` does not already exist, you can copy `/opt/so/saltstack/default/salt/curator/init.sls` to `/opt/so/saltstack/local/salt/curator/init.sls` and modify as shown below.

If `/opt/so/saltstack/local/salt/curator/init.sls` does already exist, create a backup of the file by copying it to a safe directory. Then, copy the default file located at `/opt/so/saltstack/default/salt/curator/init.sls` to the location of the current file and run a diff against the two `init.sls` files. Inside this file that was just copied over, the “cur” and “cron” sections must be added for your Curator job along with anything included in the diff output. Do not edit the original file in the directory.

To add the new Curator job, copy and modify one of the existing sections or use these examples:

For the “cur” section:

```
cur<custom-name>:
  file.managed:
    - name: /usr/sbin/so-curator-<custom-name>
    - source: salt://curator/files/bin/<your_script_file_name>
    - user: 934
    - group: 939
    - mode: 755
```

For the “cron” section:

```
so-curator<custom-name>:
  cron.present:
    - name: /usr/sbin/so-curator-<custom-name> > /opt/so/log/curator/<your_logfile>.log 2>
    ↪&1
    - user: root
    - minute: '*'
    - hour: '*'
    - daymonth: '*'
```

(continues on next page)

(continued from previous page)

```
- month: '* *  
- dayweek: '* *
```

This particular cron section will run the task every minute. After this, restart Curator with `sudo so-curator-restart` and note any errors (changes are not errors).

To confirm that the job was added correctly, run `crontab -l` and look for the new task's cron job.

If the task's cron job does not show up, then there may have been errors during the restart process. You must fix those errors for the cron job to be created.

### 10.7.3 Diagnostic Logging

When Curator completes an action, it logs its activity in a log file found in `/opt/so/log/curator/`. Depending on what you're looking for, you may also need to look at the *Docker* logs for the container:

```
sudo docker logs so-curator
```

### 10.7.4 Curator vs Index Lifecycle Management (ILM)

The goal of Security Onion is to allow you to concentrate on finding evil rather than spending time managing infrastructure. The default mode that Security Onion deploys allows each node to be independent and removes the complexity of shard migration across multiple nodes. These nodes will use Curator to trim indices as needed ensuring that they never run out of disk space. This is especially important when running in standalone mode. Finally, it should also be noted that ILM does not support deletion based on disk space.

### 10.7.5 More Information

---

**Note:**

For more information about Curator, please see:

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

---

## 10.8 Data Fields

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

### 10.8.1 ECS

We've begun transitioning to Elastic Common Schema (ECS). This is a work-in-progress and will continue as time goes on.

For more information about ECS, please see:

<https://www.elastic.co/guide/en/ecs/current/ecs-reference.html>

## 10.8.2 Fields

*Alert Data Fields*

*Elastalert Fields*

*Zeek Fields*

## 10.8.3 Template files

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

## 10.9 Alert Data Fields

*Elasticsearch* receives NIDS alerts from *Suricata* via *Filebeat* 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/master/salt/elasticsearch/files/ingest/suricata.alert>

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

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

## 10.10 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`

## 10.11 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/master/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/master/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/master/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/master/salt/elasticsearch/files/ingest-dynamic/common>

## 10.12 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*.

We also sponsored the development of Community ID support in *osquery*.

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>

### 10.12.1 More Information

---

#### Note:

For more information about Community ID, please see:

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

---

## 10.13 Re-Indexing

When changing mappings or index settings, we may need to re-index the existing indices to ensure there are no mapping conflicts.

One way to do this by using the following **experimental** example script:

<https://raw.githubusercontent.com/weslambert/securityonion-elastic-misc/master/so-elastic-reindex>

Pull down the script to your Security Onion box:

```
wget https://raw.githubusercontent.com/weslambert/securityonion-elastic-misc/master/so-elastic-reindex
```

Make the script executable:

```
sudo chmod +x so-elastic-reindex
```

Re-index all indices matching `logstash-*`, pulling the appropriate `refresh_interval` from the template named `logstash` in Elasticsearch:

```
sudo ./so-elastic-reindex -i "logstash-*" -t "logstash"
```

The script should then progress to re-index the matching indices, and inform you when it has completed.

**Warning:** Abnormal execution of this script may result in data loss– there are **NO GUARANTEES** this process will work perfectly for you.

## 10.14 SOC Logs

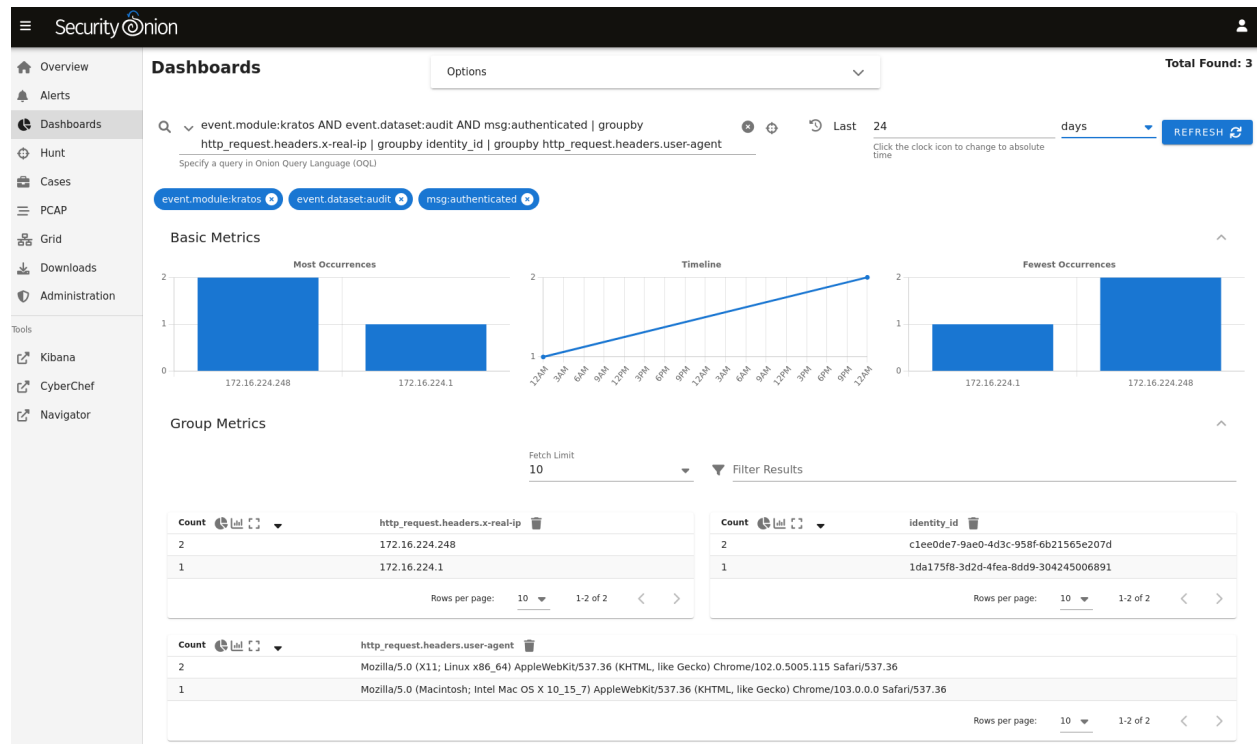
Standard *Security Onion Console (SOC)* logs can be found at `/opt/so/log/soc/`.

### 10.14.1 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.



## 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
```

## 10.15 Other Supported Logs

We include *Elasticsearch* ingest parsers for several log types that don't have *Filebeat* modules.

### 10.15.1 Example: pfSense

Security Onion includes *Elasticsearch* ingest parsers for pfSense firewall logs. Simply run *so-allow* as described in the *Syslog* section and then 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). You should then be able to see your firewall logs using the Firewall query in *Dashboards* or *Hunt*.

### 10.15.2 Example: RITA

Security Onion includes *Elasticsearch* ingest parsers for *RITA* logs. To enable this support, add the following in the relevant *Salt* minion pillar and then restart *Filebeat* on the minion(s):

```
rita:
  enabled: True
```

This will enable the following *Filebeat* inputs:

```
/nsm/rita/beacons.csv
/nsm/rita/explored-dns.csv
/nsm/rita/long-connections.csv
/nsm/rita/open-connections.csv
```

If you are installing *Filebeat* on a non-Security Onion node or your filenames differ, you will need to copy the *Filebeat* configuration from `/opt/so/saltstack/default/salt/filebeat/etc/filebeat.yml` to `/opt/so/saltstack/local/salt/filebeat/etc/filebeat.yml` (or modify on the non-Security Onion node in the normal Filebeat configuration file) and emulate the path/filename accordingly.

Once ingested into Security Onion, you should be able to search for *RITA* logs in *Dashboards* or *Hunt* using `event.module:rita | groupby event.dataset`.



Summary

Count	event.module	event.dataset
2,615	rita	beacon
1,000	rita	dns
29	rita	connection
2	rita	alert

Connections



















event.module	event.dataset	source.ip	destination.ip	event.duration ▼	connection.state
rita	connection	192.168.88.2	216.229.4.69	385	closed
rita	connection	10.55.100.100	65.52.108.225	385	closed
rita	connection	192.168.88.2	66.228.58.20	384	closed
rita	connection	192.168.88.2	216.218.220.101	384	closed
rita	connection	10.55.182.100	96.45.33.73	384	closed
rita	connection	10.55.182.100	173.243.138.200	384	closed
rita	connection	10.55.182.100	173.243.138.201	384	closed
rita	connection	10.55.100.107	111.221.29.113	384	closed
rita	connection	192.168.88.2	129.250.35.250	383	closed
rita	connection	192.168.88.2	162.210.111.4	383	closed

DNS

event.module	event.dataset	dns.question.name	dns.question.subdomain_count	dns.question.count ▼
rita	dns	3lift.com	10	707
rita	dns	2-01-2f1f-0001.cdx.cedexis.net	1	396
rita	dns	2mdn.net	2	379
rita	dns	20753598p.rlihub.com	1	234
rita	dns	2-01-2c3e-0010.cdx.cedexis.net	1	233
rita	dns	1605158521.rsc.cdn77.org	1	227
rita	dns	2-01-2c3e-003c.cdx.cedexis.net	1	220
rita	dns	2-01-37d2-0006.cdx.cedexis.net	1	195
rita	dns	2-01-3073-0012.cdx.cedexis.net	1	178
rita	dns	4-c-0003.c-msedge.net	2	16

Rows per page: 10 1-10 of 10

Beacon

 beacon.interval.dispersion	0
 beacon.interval.range	1
 beacon.interval.skew	0
 beacon.interval.top	1
 beacon.interval.top_count	297
 beacon.score	1
 beacon.size.dispersion	0
 beacon.size.range	104
 beacon.size.skew	0
 beacon.size.top	52
 beacon.size.top_count	1897
 destination.ip	10.233.233.5
 ecs.version	1.12.0
 event.category	network
 event.dataset	alert
 event.module	rita
 event.severity	3
 event.severity_label	high



Alert

If the value for beacon.score in a beacon record equals 1, an alert will be generated and viewable in [Alerts](#).

Group: rule.name ✕

Group: event.module ✕

Group: event.severity\_label ✕

	Count	rule.name
 	2	Potential C2 Beacon Activity

## 10.16 Related Information

### 10.16.1 RITA

From: <https://github.com/activecm/rita>

RITA is an open source framework for network traffic analysis.

The framework ingests Zeek Logs, and currently supports the following analysis features:

Beaconing: Search for signs of beaconing behavior in and out of your network

DNS Tunneling Search for signs of DNS based covert channels

Blacklisted: Query blacklists to search for suspicious domains and hosts

We can add RITA to Security Onion to enhance its current capabilities and leverage the great work from the folks at [Active Countermeasures](#). They've done a fantastic job of allowing RITA to be easy to integrate with Security Onion.

**Warning:** Please keep in mind we do not officially support RITA, so installation is at your own risk.

### Installation

To install RITA on Security Onion:

Download the install script:

```
wget https://raw.githubusercontent.com/activecm/rita/master/install.sh
```

Run the installer:

```
sudo bash ./install.sh --disable-zeek
```

### Configuration

You should set your own values for `InternalSubnets` in `/etc/rita/config.yaml` before importing data into RITA. See <https://github.com/activecm/rita#configuration-file> for more information.

### Usage

You can then import logs with:

```
rita import /nsm/zeek/logs/2019-09-04 dataset1
```

To see long connections, type:

```
rita show-long-connections dataset1
```

To see beacons, type:

```
rita show-beacons dataset1
```

Finally, you can issue an HTML report (viewable in browser) by typing:

```
rita html-report
```

See other available commands with:

```
rita --help
```

For additional information, see:

<https://github.com/activecm/rita>

## 10.16.2 Endgame

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.

---

### Configuration

#### During Security Onion Setup

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.

---

#### Post-Installation Setup

To configure Endgame ingestion on an existing Security Onion installation, perform the following steps.

Add the SMP to the firewall exceptions for the Security Onion node:

```
sudo so-firewall includehost endgame $smpip
```

Add the following to the soc pillar entry the manager's sls file in `/opt/so/saltstack/local/pillar/minions` to configure the pivot from SOC to Endgame (based on `agent.id`):

```
soc:
  endgamehost: $smpip
```

#### Configure Event Streaming in Endgame SMP

Once one of the two requirements above have been completed, the following must be configured in the Endgame web console:

Select **Administration** -> **Streaming** -> **Event Streaming** -> **Add**. **Requires Admin user role**

Ensure Logstash is selected under the **Destination** options.

Next, copy the contents of `/etc/ssl/certs/intca.crt` (on the Security Onion manager node) to the **Certificate** section.

*Endgame will attempt to verify the X.509 certificate attributes match the destination server, so you will also need to ensure the SMP can resolve the hostname of the Security Onion node (to match the certificate). This may require a hosts file entry on the SMP.*

Ensure the SMP is pointed to `https://$securityonion:3765` and save the configuration.

Navigate to **Administration** -> **Policy** -> **YOUR POLICY** -> **Settings** -> **Elastic Streaming** and enable **Event Streaming** if not already enabled.

Once events are batched and published from the Endgame SMP, you can search for them in [Dashboards](#) or [Hunt](#) using a query like `event.module:endgame`.

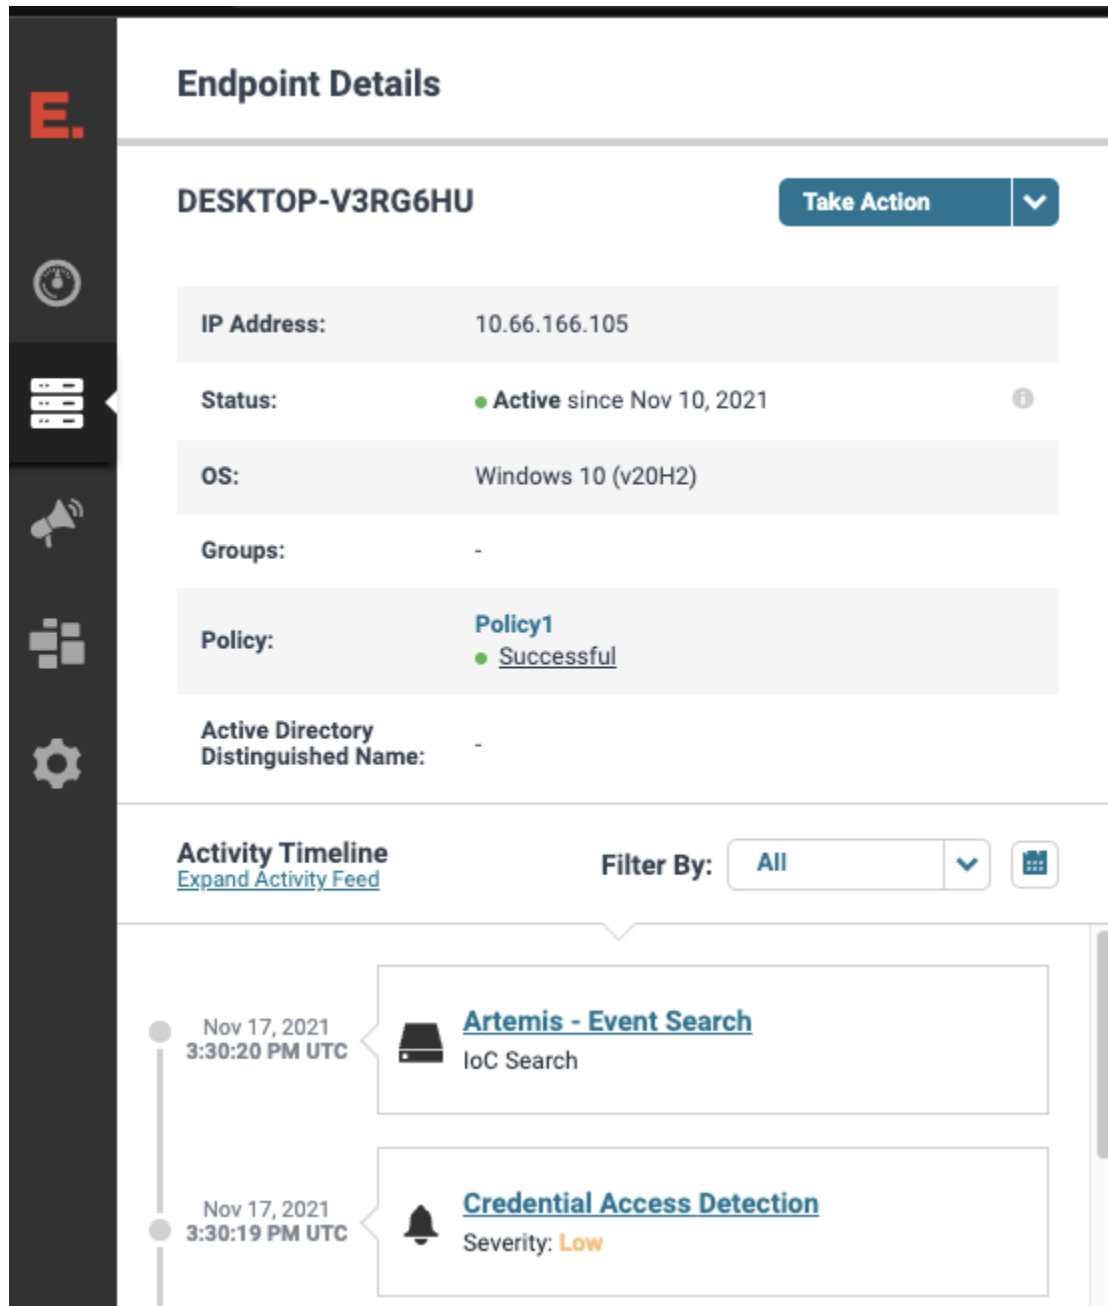
## Example Endgame Data

Count ▲	event.module	host.name	user.name	process.command_line
2	endgame	DESKTOP-V3RG6HU	bassmaster	"BackgroundTransferHost.exe" -ServerName:BackgroundTransferHost.1
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Program Files (x86)\Microsoft\Edge\Application\msedge.exe" --type=utility --utility-sub-type=i
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Program Files (x86)\Microsoft\Edge\Application\msedge.exe" --type=utility --utility-sub-type=i
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Program Files\WindowsApps\Microsoft.XboxGamingOverlay_5.721.10202.0_x64__8wekyb3
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Program Files\WindowsApps\Microsoft.XboxGamingOverlay_5.721.10202.0_x64__8wekyb3
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Users\bassmaster\AppData\Local\Microsoft\OneDrive\21.220.1024.0005\Microsoft.SharePoi
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Windows\system32\BackgroundTaskHost.exe" -ServerName:BackgroundTaskHost.WebAcco
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Windows\system32\backgroundTaskHost.exe" -ServerName:App.AppXe9cvj1thv1hmcw0csS
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Windows\system32\backgroundTaskHost.exe" -ServerName:App.AppXemn3t55sepg7q92mv
2	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\Windows\system32\backgroundTaskHost.exe" -ServerName:App.AppXggkaqzf6p31g37n0m
2	endgame	DESKTOP-V3RG6HU	bassmaster	C:\Users\bassmaster\AppData\Local\Microsoft\OneDrive\OneDriveStandaloneUpdater.exe
2	endgame	DESKTOP-V3RG6HU	bassmaster	C:\WINDOWS\System32\slui.exe -Embedding
2	endgame	DESKTOP-V3RG6HU	bassmaster	C:\WINDOWS\system32\svchost.exe -k BcastDVRUserService -s BcastDVRUserService
2	endgame	DESKTOP-V3RG6HU	bassmaster	C:\Windows\System32\smartscreen.exe -Embedding
2	endgame	DESKTOP-V3RG6HU	bassmaster	taskhostw.exe
2	endgame	DESKTOP-V3RG6HU	bassmaster	taskhostw.exe Install \$(Arg0)
2	endgame	DESKTOP-V3RG6HU	bassmaster	taskhostw.exe SyncFromCloud
4	endgame	DESKTOP-V3RG6HU	bassmaster	"C:\WINDOWS\system32\findstr.exe" password \passwords.txt

Pivot to Endgame Console


If Endgame support is enabled, a default *Endgame* pivot will be populated within SOC, based on the `agent.id` field:

agent.id	f48b5101-3ff3-493d-b124-102289139be7	<div><div>Include</div><div>Exclude</div><div>Only</div><div>Group By</div><div>Clipboard</div><div>Actions</div><div>Hunt</div><div>PCAP</div><div>CyberChef</div><div>Google</div><div>VirusTotal</div><div>Endgame</div></div>
agent.name	Endgame	
agent.type	Endgame	
agent.version	3.59.1	
ecs.version	1.9.0	
endgame.command_line	"C:\WINDOWS\system32\findstr.exe" pe	
endgame.event_subtype_full	termination_event	
endgame.event_type_full	process_event	
endgame.exit_code	1	
endgame.md5	804a6ae28e88689e0cf1946a6cb3fee5	
endgame.opcode	2	
endgame.original_file_name	FINDSTR.EXE	
endgame.parent_command_line	"C:\Windows\System32\WindowsPower	
endgame.parent_process_name	powershell.exe	
endgame.parent_process_path	C:\Windows\System32\WindowsPowerS	
endgame.pid	324	
endgame.ppid	9532	
endgame.process_name	findstr.exe	




The screenshot shows the 'Endpoint Details' page for a host named 'DESKTOP-V3RG6HU'. The left sidebar contains navigation icons: a red 'E.', a circular icon, a server rack icon, a megaphone icon, a grid icon, and a gear icon. The main content area has a 'Take Action' button with a dropdown arrow. Below this, several fields provide details about the endpoint: IP Address (10.66.166.105), Status (Active since Nov 10, 2021), OS (Windows 10 (v20H2)), Groups (-), Policy (Policy1, Successful), and Active Directory Distinguished Name (-). An 'Activity Timeline' section follows, with a 'Filter By: All' dropdown and a calendar icon. The timeline shows two events: 'Artemis - Event Search' (IoC Search) at 3:30:20 PM UTC and 'Credential Access Detection' (Severity: Low) at 3:30:19 PM UTC.

## Endpoint Details

**DESKTOP-V3RG6HU** [Take Action](#) 

**IP Address:** 10.66.166.105

**Status:** ● **Active** since Nov 10, 2021 



**OS:** Windows 10 (v20H2)


**Groups:** -


**Policy:** **Policy1**  
● [Successful](#)

**Active Directory Distinguished Name:** -

### Activity Timeline

[Expand Activity Feed](#) **Filter By:** [All](#)  

Nov 17, 2021 3:30:20 PM UTC  [Artemis - Event Search](#)  
IoC Search

Nov 17, 2021 3:30:19 PM UTC  [Credential Access Detection](#)  
Severity: **Low**





## 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
```

---

**Note:** Upgrading from Security Onion 2.3 to Security Onion 2.4 is not supported via `soup`, as the new version is a major change, including a new operating system. More information on the upgrade process can be found in the Security Onion 2.4 documentation: <https://docs.securityonion.net/en/2.4/appendix.html>

---

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

There is an option during *Configuration* to automatically install OS updates.

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

### 11.1.4 Elastic 8

When you update to Security Onion 2.3.140 or higher, the Elastic components will undergo a major version upgrade to version 8. Please review and follow the steps below. Failure to do so could result in loss of access to all data stored inside of Elastic and a non-functioning Security Onion installation.

- You should test the update in a test environment that closely matches your production environment. Do not upgrade your production environment without testing first!
- Elastic's requirement for upgrading to version 8 is to be on Elastic 7.17.0 or later (<https://www.elastic.co/guide/en/elasticsearch/reference/current/setup-upgrade.html>). This means that if you are currently using a version of Security Onion prior to 2.3.110, then you will need to update to Security Onion 2.3.130 before updating to 2.3.140 or higher. *soup* should check for applicable versions and advise on the proper method to upgrade appropriately.
- Once you are on Security Onion 2.3.110 or higher, you may want to log into *Kibana* and run the Upgrade Assistant to check for any unexpected issues before upgrading to 2.3.140 or higher. Upgrade Assistant may display issues for *ElastAlert* indices, but they will be deleted automatically when you update. You may need to investigate any non-ElastAlert issues.
- If your deployment was originally installed with a version lower than 2.3.60 and you never enabled Elastic Auth, then you will need to enable Elastic Auth via *so-elastic-auth* before updating.
- Before upgrading, you should review any local modifications that you've applied in the past. In particular, some folks have experienced issues with local Filebeat modifications. For example, if you have the file `/opt/so/saltstack/local/salt/filebeat/init.sls`, then you should move or rename it before beginning the upgrade. After the upgrade, you can copy the new `/opt/so/saltstack/default/salt/filebeat/init.sls` to `/opt/so/saltstack/local/salt/filebeat/init.sls` and then re-apply your local modification.
- When you run *soup* to update to 2.3.140 or higher, it will check for old Elasticsearch indices created with Elasticsearch 6. If found, you will need to log into *Kibana* and use the Upgrade Assistant to update those indices.
- As *soup* is updating to Security Onion 2.3.140 or higher, *ElastAlert* indices will be automatically deleted to ensure compatibility with Elastic version 8. You don't have to worry about loss of alerts or alert data generated by *Playbook* as they should remain in their normal HIDS/NIDS/Playbook indices. However, if you currently

leverage the *ElastAlert* indices for other purposes, you may want to back up the indices to manually review at a later date as needed. This process is not supported or documented, and is left up to the user to perform as needed.

- *Elasticsearch* 8 no longer includes GeoIP databases by default. Read more about this in the *Elasticsearch* section.

### 11.1.5 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.

### 11.1.6 Log

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

### 11.1.7 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.

### 11.1.8 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-2.3.XYZ-YYYYMMDD.iso
```

### 11.1.9 Agents

If you've previously added any external agents (*Wazuh*, *Beats*, etc.), be sure to upgrade them to match the version of your upgraded components.

### 11.1.10 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`.

### 11.1.11 Kibana

After soup completes, if *Kibana* says `Kibana server is not ready yet` even after waiting a few minutes for it to fully initialize, then take a look at the Diagnostic Logging section of the *Kibana* page.

If Kibana loads but the dashboards display errors that they didn't before the upgrade, first shift-reload your browser to make sure there are no cache issues. If that doesn't resolve the issue, then you may need to reload the dashboards on your manager:

```
sudo rm /opt/so/state/kibana_*.txt
sudo salt-call state.apply kibana.so_savedobjects_defaults -l info queue=True
```

### 11.1.12 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
```

### 11.1.13 Errors

#### 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

```
Status: Downloaded newer image for quay.io/securityonion/so-acng:2.3.30
quay.io/securityonion/so-acng:2.3.30
% Total      % Received % Xferd  Average Speed   Time    Time     Time  Current
             Dload  Upload   Total     Spent    Left     Speed
100    543    100    543     0     0   1412     0  --:--:-- --:--:-- --:--:--   1414
There is a problem downloading the so-acng:2.3.30 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.

### 11.1.14 Distributed deployments

If you have a distributed deployment with a manager node and separate sensor nodes and/or search nodes, you **only** need to run soup on the manager. Once soup has completed, other nodes should update themselves at the next *Salt* highstate (typically within 15 minutes).

**Warning:** Just because the update completed on the manager does NOT mean the upgrade is complete on other nodes in the grid. Do not manually restart anything until you know that all the search/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.
- If *FleetDM* is enabled, then it generates new *osquery* packages.
- 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 reached EOL on April 6th, 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>

ELSA reached EOL on October 9, 2018:

<https://blog.securityonion.net/2018/04/6-month-eol-notice-for-elsa.html>

Xplico reached EOL on June 5, 2018:

<https://blog.securityonion.net/2017/12/6-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 controlled by the `so-user` command.

### 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

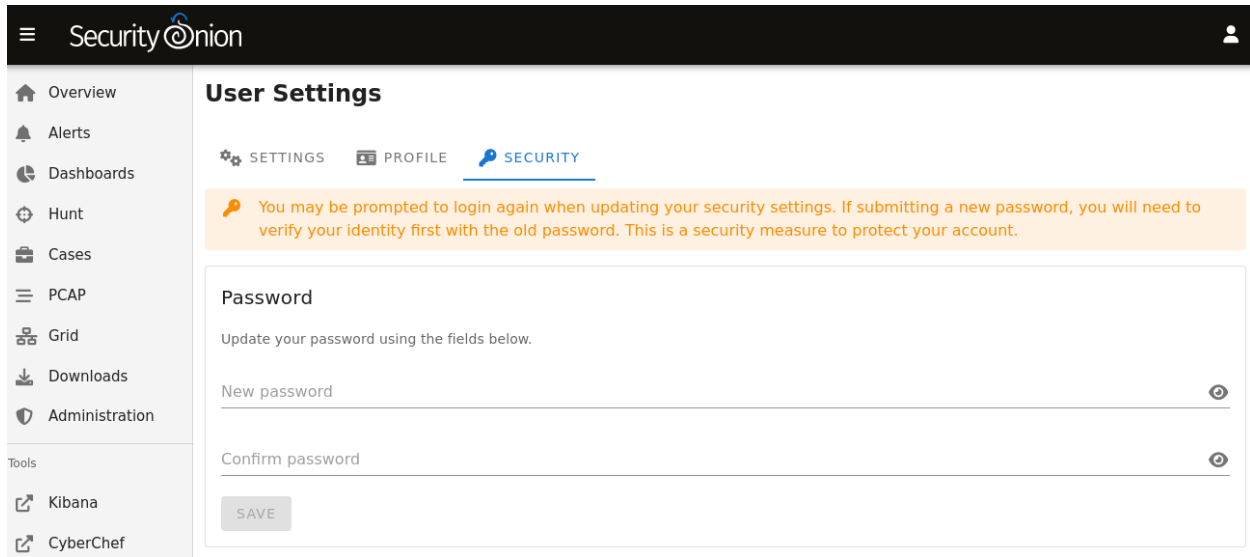
Your default user account should have `sudo` permissions. Command-line utilities that require administrative access can be prefixed with `sudo`. For example, the `so-status` command requires administrative access so you can run it with `sudo` as follows:

```
sudo so-status
```

#### 12.1.3 Security Onion Console (SOC)

Log into *Security Onion Console (SOC)* using the username and password you created in the Setup wizard.

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:



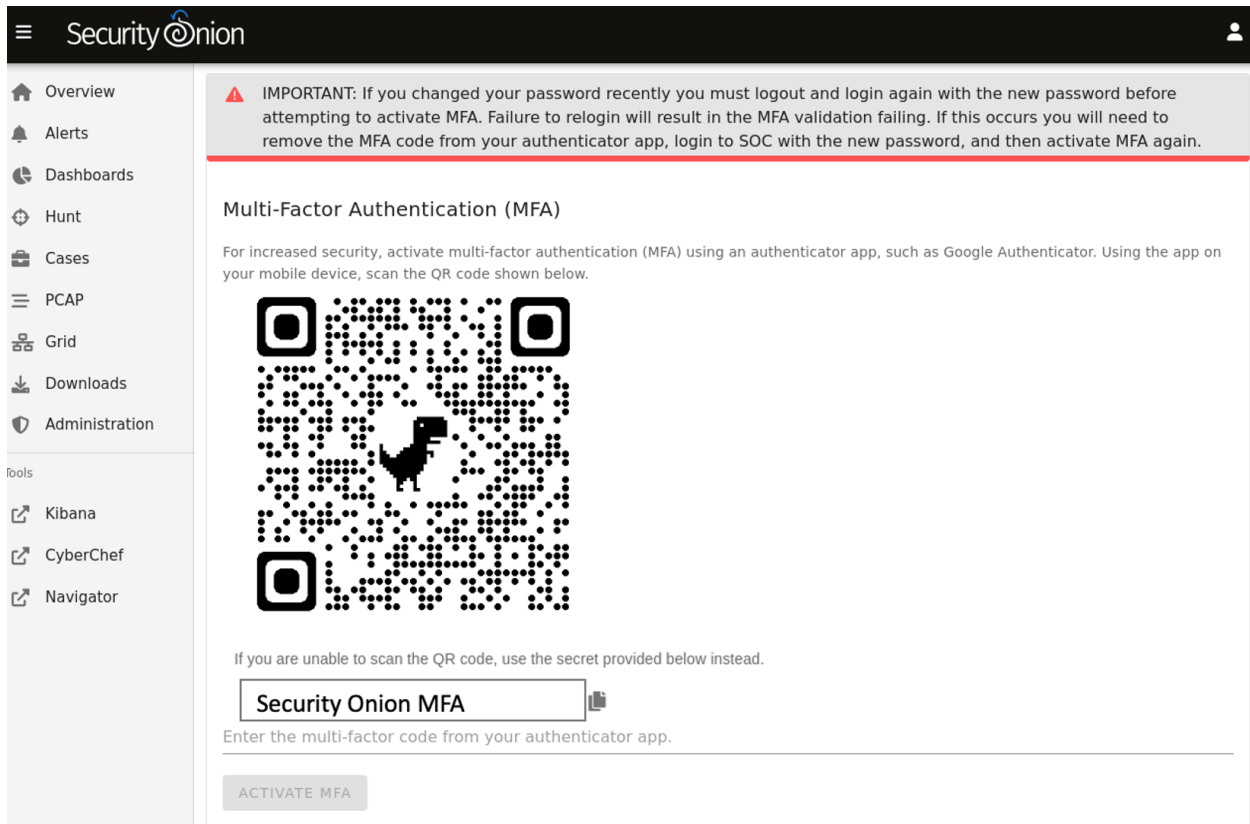
The screenshot shows the Security Onion web interface. The top navigation bar includes the Security Onion logo and a user icon. The left sidebar lists various tools: Overview, Alerts, Dashboards, Hunt, Cases, PCAP, Grid, Downloads, Administration, Kibana, and CyberChef. The main content area is titled 'User Settings' and has three tabs: SETTINGS, PROFILE, and SECURITY (which is active). A yellow warning box states: '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.' Below this, the 'Password' section prompts the user to 'Update your password using the fields below.' It contains two input fields: 'New password' and 'Confirm password', each with a toggle icon to the right. A 'SAVE' button is located below the 'Confirm password' field.

If you've forgotten your SOC password, you can reset it using the `so-user` command:

```
sudo so-user update username@example.com
```

## 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. Activating MFA 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.



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.

If you have multiple Security Onion deployments with MFA enabled, 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.

Note that if a user loses access to the authenticator app or otherwise needs to disable MFA, a Security Onion administrator can SSH into the manager node and run the following to change that user's password and temporarily disable MFA:

```
sudo so-user update username@example.com
```

## 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
```

For more information, please see `man adduser`.

## 12.3.2 SOC

If you need to add a new account to *Security Onion Console (SOC)* and *FleetDM*, you can use the `so-user-add` command and specify the user's email address. For example, to add a new account for `tom@example.com`:

```
sudo so-user-add tom@example.com
```

## 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

To list all *Security Onion Console (SOC)* accounts, you can use the `so-user-list` command:

```
sudo so-user-list
```

Alternatively, you can get a list of users in *Security Onion Console (SOC)* by clicking **Administration** and then **Users**:

Email Address ▲	First Name	Last Name	Note	Role(s)	Status	Actions
disabled@example.com				analyst		
doug@example.com				superuser		
tom@example.com				analyst		

Rows per page: 10 1-3 of 3 < >

The Status column will show a different icon depending on the status of the account. In the screenshot above, the first account is disabled, the second account is enabled and has *MFA* enabled, and the third account is enabled but does not have *MFA* enabled. Hovering over the icon in the Status column will show you these details as well.

## 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
```

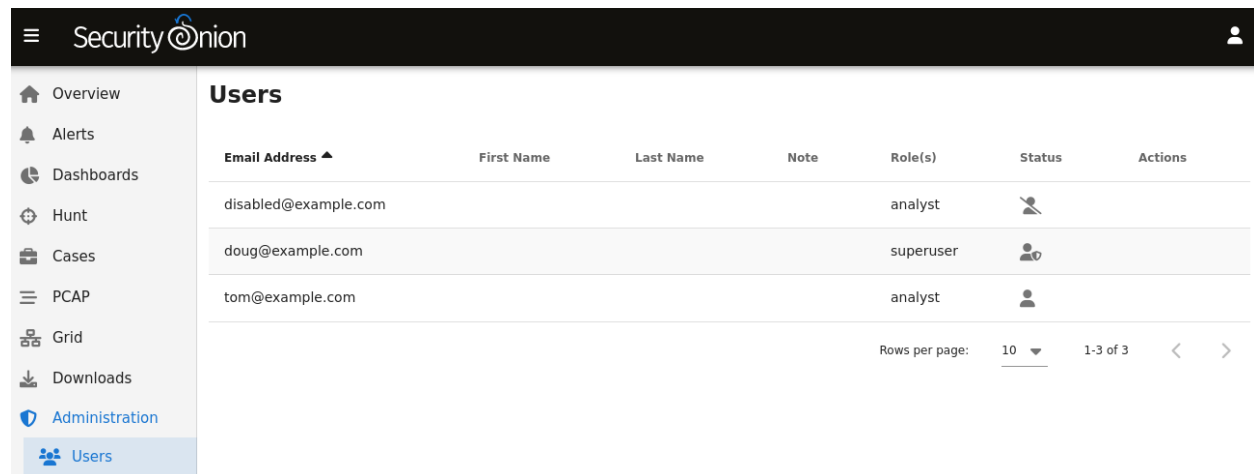
For more information, please see `man passwd` and `man usermod`.




### 12.5.2 SOC

If you need to disable an account in *Security Onion Console (SOC)* and *FleetDM*, you can use the `so-user-disable` command and specify the user's email address. For example, to disable the account for `tom@example.com`:

```
sudo so-user-disable tom@example.com
```

After disabling a user account, the *Security Onion Console (SOC) Administration* page will show the disabled user account with a disabled icon in the Status column:



Security Onion							
Users							
Email Address ▲	First Name	Last Name	Note	Role(s)	Status	Actions	
disabled@example.com				analyst			
doug@example.com				superuser			
tom@example.com				analyst			

Rows per page: 10 1-3 of 3

## 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 Elastic Auth Requirement

RBAC requires Elastic authentication which is enabled by default in version 2.3.60 and later. If you upgrade an older release, you may need to manually enable Elastic auth via *so-elastic-auth* before using RBAC features.

## 12.6.2 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.

---

### 12.6.3 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.

Upgrades of existing Security Onion grids to 2.3.80 or later will result in all existing users being assigned to the `superuser` role. Continue reading this document to learn how to change those role assignments, if necessary.

### 12.6.4 Adding a User With a Specific Role

To add a user with a role other than the default `analyst` role, use the `so-user add` command with an additional parameter for the role. For example, to add the user `tom@example.com` with the `auditor` role:

```
sudo so-user add "tom@example.com" auditor
```

### 12.6.5 Modifying User Roles

To add a role to an existing user you can use the `so-user addrole` command. For example to add the `analyst` role to the user `tom@example.com`:

```
sudo so-user addrole "tom@example.com" analyst
```

Conversely, to remove a role from a user use the `so-user delrole` command. To remove the `analyst` role from `tom@example.com` run:

```
sudo so-user delrole "tom@example.com" analyst
```

### 12.6.6 Creating Custom Roles

**Warning:** The creation of custom RBAC roles is an advanced feature that is recommended only for experienced administrators.

These steps will guide you through an example where we wish to introduce a new role called `eastcoast-analyst`, which will inherit the same Security Onion permissions as a `limited-analyst`, but will be restricted to only view a subset of documents in the Elastic stack. We base this role on the `limited-analyst` instead of the `analyst` role so that the user does not have the ability to create arbitrary PCAPs on any sensor.

1. For the Security Onion role: Follow the instructions in the next section entitled “Defining Security Onion Roles” to create a new role named `eastcoast-analyst`.
2. For the Elastic stack role: Create a new json role file named `eastcoast-analyst.json` under `/opt/so/saltstack/local/salt/elasticsearch/roles`. In this example we will define the new role that only allows access to documents from sensors on the east coast of the United States. Specifically, the role will include a query filter that limits search results to only include documents originating from sensors having a name prefixed with `nyc` (New York City) or `atl` (Atlanta).

`eastcoast-analyst.json`:

```
{
  "cluster": [
    "cancel_task",
```

(continues on next page)

(continued from previous page)

```

    "create_snapshot",
    "monitor",
    "monitor_data_frame_transforms",
    "monitor_ml",
    "monitor_rollup",
    "monitor_snapshot",
    "monitor_text_structure",
    "monitor_transform",
    "monitor_watcher",
    "read_ccr",
    "read_ilm",
    "read_pipeline",
    "read_slm"
  ],
  "indices": [
    {
      "names": [
        "so-*"
      ],
      "privileges": [
        "index",
        "maintenance",
        "monitor",
        "read",
        "read_cross_cluster",
        "view_index_metadata"
      ],
      "query": "{ \"bool\": { \"should\": [ { \"prefix\": { \"observer.name\": \"nyc\" } }, { \"prefix\": { \"observer.name\": \"atl\" } } ] } }"
    }
  ],
  "applications": [
    {
      "application": "kibana-.kibana",
      "privileges": [
        "feature_discover.all",
        "feature_dashboard.all",
        "feature_canvas.all",
        "feature_maps.all",
        "feature_ml.all",
        "feature_logs.read",
        "feature_visualize.all",
        "feature_infrastructure.read",
        "feature_apm.read",
        "feature_uptime.read",
        "feature_siem.read",
        "feature_dev_tools.read",
        "feature_advancedSettings.read",
        "feature_indexPatterns.read",
        "feature_savedObjectsManagement.read",
        "feature_savedObjectsTagging.read",
        "feature_fleet.all",

```

(continues on next page)



(continued from previous page)

```

        "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 Elastic docs here: <https://www.elastic.co/guide/en/elasticsearch/reference/current/security-api-put-role.html#security-api-put-role-request-body>.

The available cluster and indices permissions are explained in the Elastic docs here: <https://www.elastic.co/guide/en/elasticsearch/reference/current/security-privileges.html>.

The available kibana permissions can be obtained by running the following command on the manager node:

```

sudo so-elasticsearch-query _security/privilege/kibana-.kibana | jq '. |
  ↪map_values(keys)'

```

3. Run a salt highstate from the manager:

```

sudo salt-call state.highstate

```

## 12.6.7 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
event-admin: eastcoast-analyst
node-monitor: eastcoast-analyst

```

(continues on next page)

(continued from previous page)

```
user-monitor: eastcoast-analyst
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

## 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

Below are some ways in which you can customize *Security Onion Console (SOC)*. Once all customizations are complete, you can then restart SOC to make the changes take effect:

```
sudo so-soc-restart
```

#### 14.1.1 Login Page

You can customize the SOC login page with a login banner. The content of this banner is stored in the `banner.md` file, which uses the common Markdown (.md) format. You can learn more about Markdown format at <https://markdownguide.org>. To customize the login banner, copy `banner.md` as follows and then edit `/opt/so/saltstack/local/salt/soc/files/soc/banner.md` using your favorite text editor:

```
sudo cp /opt/so/saltstack/default/salt/soc/files/soc/banner.md /opt/so/saltstack/local/  
↪salt/soc/files/soc/
```

#### 14.1.2 Overview Page

After logging into SOC, you'll start on the main SOC Overview page which can be customized as well. The content of this page is stored in the `motd.md` file, which uses Markdown format as mentioned above. To customize the Overview page content, copy `motd.md` as follows and then edit `/opt/so/saltstack/local/salt/soc/files/soc/motd.md` using your favorite text editor:

```
sudo cp /opt/so/saltstack/default/salt/soc/files/soc/motd.md /opt/so/saltstack/local/  
↪salt/soc/files/soc/
```

### 14.1.3 Links

You can also customize the links on the left side. To do so, copy `tools.json` as follows and then edit `/opt/so/saltstack/local/salt/soc/files/soc/tools.json` using your favorite text editor:

```
sudo cp /opt/so/saltstack/default/salt/soc/files/soc/tools.json /opt/so/saltstack/local/
↪ salt/soc/files/soc/
```

### 14.1.4 Session Timeout

Another possible SOC customization is the 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. This can be adjusted by adding a pillar value to the manager node's pillar sls. For example, on an eval node, edit `/opt/so/saltstack/local/pillar/minions/eval_eval.sls` and add a new `kratos.sessiontimeout` value:

```
kratos:
  kratoskey: 'abcdef1234567890'
  sessiontimeout: 720h
```

### 14.1.5 Advanced Interface

The interface for *Alerts* and *Cases* is a simplified version of the interface seen in *Dashboards* and *Hunt*. *Alerts* and *Cases* have a new toggle under Options labeled `Temporarily enable advanced interface features`.

 `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.

In versions older than 2.3.160, you can temporarily enable this by opening your browser console and then pasting the following:

```
document.getElementById('hunt').__vue__.$parent._data.advanced=true
```

**Warning:** Making this change permanent requires changes to `soc.json` which changes at nearly every release. So if you proceed with this, please be prepared to update to the new `soc.json` every time you update.

If you would like to make this change permanent, make a copy of `soc.json`:

```
cp /opt/so/saltstack/default/salt/soc/files/soc/soc.json /opt/so/saltstack/local/salt/
↪ soc/files/soc/
```

Then edit `/opt/so/saltstack/local/salt/soc/files/soc/` using your favorite text editor, find the `alerts` or `cases` section, and set `advanced` to `true`. Don't forget that you will need to manually update `soc.json` every time you update to the latest version.

### 14.1.6 Custom Queries

If you'd like to add your own custom queries to *Alerts*, *Dashboards*, or *Hunt*, you can copy the relevant queries.json file from `/opt/so/saltstack/default/salt/soc/files/soc/` to `/opt/so/saltstack/local/salt/soc/files/soc/` and then add new entries. For example, if you want to add a new dashboard, you can copy `dashboards.queries.json` and then add your new dashboard there.

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 first copy the `alerts.queries.json`:

```
sudo cp -n /opt/so/saltstack/default/salt/soc/files/soc/alerts.queries.json /opt/so/
↳saltstack/local/salt/soc/files/soc/alerts.queries.json
```

Next edit the `/opt/so/saltstack/local/salt/soc/files/soc/alerts.queries.json` using your favorite text editor 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.

### 14.1.7 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 copy `/opt/so/saltstack/default/salt/soc/files/soc/menu.actions.json` to `/opt/so/saltstack/local/salt/soc/files/soc/menu.actions.json` and then add new entries. (Previous to Security Onion 2.3.60, this would be done in `alerts.actions.json` or `hunt.actions.json`.)

For example, suppose you want to add AbuseIPDB with URL `https://www.abuseipdb.com/check/{value}`. First, copy `/opt/so/saltstack/default/salt/soc/files/soc/menu.actions.json` to `/opt/so/saltstack/local/salt/soc/files/soc/menu.actions.json`:

```
sudo cp -n /opt/so/saltstack/default/salt/soc/files/soc/menu.actions.json /opt/so/
↳saltstack/local/salt/soc/files/soc/menu.actions.json
```

Next, edit `/opt/so/saltstack/local/salt/soc/files/soc/menu.actions.json` using your favorite text editor and insert the following as the next to last line of the file:

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

So once you've restarted SOC to make the change take effect:

- AbuseIPDB will display on the Actions menu.
- When you hover over that AbuseIPDB option, the description Search for this value at AbuseIPDB will appear.
- When you click the AbuseIPDB option, the browser will open a new tab and go to `https://www.abuseipdb.com/check/{value}` (replacing `{value}` with the original value that you clicked on that spawned the Action menu).

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}"
},
```

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](https://developer.mozilla.org/en-US/docs/Web/API/Fetch_API/Using_Fetch).

### 14.1.8 Cases

*Cases* comes with presets for things like category, severity, TLP, PAP, tags, and status. You can modify these presets by copying the appropriate presets file from `/opt/so/saltstack/default/salt/soc/files/soc/` to `/opt/so/saltstack/local/salt/soc/files/soc/`, making changes there, and then restarting SOC.

### 14.1.9 Escalation

*Alerts*, *Dashboards*, and *Hunt* display logs with a blue triangle that allows you to escalate the event which defaults to our own *Cases* interface. If for some reason you want to escalate to a different case management system, you can change this setting. To do so, locate the soc *Salt* pillar and then set `case_module` to one of the following values:

- `soc` - Enables the new built-in Case Management, with the new Escalation menu (default).
- `thehive` - Enables escalation directly to TheHive v3 instance running in the Security Onion cluster (only applicable to existing installations that upgrade to 2.3.100). Escalations will always open a new case; there will not be an advanced escalation menu popup. Note that Security Onion support for TheHive has ended, and TheHive



will no longer be included in future Security Onion releases. Therefore this option should only be considered for short-term, temporary usage.

- **elasticcases** - Enables escalation to the [Elastic Cases](#) tool. Escalations will always open a new case; there will not be an advanced escalation menu popup. This module will use the same user/pass that SOC uses to talk to Elastic. Note, however, that Elastic cases is actually a Kibana feature, therefore, when this setting is used, SOC will be communicating with the local Kibana service (via its API) for case escalations.
- **httpcase** - Enables escalation directly to an arbitrary web URL. Escalations will always open a new case; there will not be an advanced escalation menu popup. To use this module, you will need to add a second pillar value, for the pillar `httpcase_config`. The value can include some, or all, of the following settings:

```
"hostUrl": "http://some.external.host/some/api",
"headers": [
  "Authorization: basic Fa3Fa01mDmCC09dA",
  "x-some-key: 1122"
],
"verifyCert": true,
"createPath": "/some/url/path/to/create/a/case",
"createMethod": "PUT",
"createBody": "{ \"myid\": \"{{ '{{ .Id }}' }}\", \"title\": \"{{ '{{ .Title }}' }}\", \
↪ \"desc\": \"{{ '{{ .Description | js }}' }}\" }",
"createContentType": "application/json",
"createSuccessCode": 200
```

Example of a customized SOC pillar file located in `/opt/so/saltstack/local/pillar/minions/import_import.sls` (your file path will vary depending on your installation choices):

```
soc:
  es_index_patterns: '*:so-*,*:endgame-*'
  case_module: httpcase
  httpcase_config: |
    "hostUrl": "http://172.17.0.1/some/api",
    "headers": [
      "Authorization: basic Fa3Fa01mDmCC09dA",
      "x-some-key: 1122"
    ],
    "verifyCert": true,
    "createPath": "/some/url/path/to/create/a/case",
    "createMethod": "PUT",
    "createBody": "{ \"myid\": \"{{ '{{ .Id }}' }}\", \"title\": \"{{ '{{ .Title }}' }}\", \
↪ \"desc\": \"{{ '{{ .Description | js }}' }}\" }",
    "createContentType": "application/json",
    "createSuccessCode": 200
```

## 14.2 Proxy Configuration

Setup will ask if you want to connect through a proxy server and, if so, it will automatically configure the system for you. Otherwise, if you need to configure manually, please continue reading.

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.

---

### 14.2.1 Salt

In addition to the above, Security Onion also makes use of pillar values in the file `/opt/so/saltstack/local/pillar/minions/<HOSTNAME>_<ROLE>.sls` on managers. Edit that file as below, following the same substitutions from above:

```
...
manager:
  ...
  proxy: '<proxy_url>'
  no_proxy: 'localhost, 127.0.0.1, <management_ip>, <hostname>'
  ...
```

---

**Note:** The above snippet is truncated, ellipses ( `...` ) should be treated as one or more lines in the file.

---

### 14.2.2 Docker

To configure Docker proxy settings, please see <https://docs.docker.com/network/proxy/>.

### 14.2.3 Git

To configure git to use a proxy for all users, add the following to `/etc/gitconfig`:

```
[http]
proxy = <proxy_url>
```

### 14.2.4 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.

## 14.3 Firewall

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

### 14.3.1 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 to the following:

- `repo.securityonion.net` (CentOS Updates)
- `raw.githubusercontent.com` (Security Onion public key)
- `pkg-containers.githubusercontent.com`
- `sigs.securityonion.net` (Signature files for Security Onion containers)
- `ghcr.io` (Container downloads)
- `rules.emergingthreatspro.com` (Emerging Threats IDS rules)
- `rules.emergingthreats.net` (Emerging Threats IDS open rules)
- `www.snort.org` (Paid Snort Talos ruleset)
- `github.com` (Strelka and Sigma rules updates)
- `geoip.elastic.co` (GeoIP updates for Elasticsearch)
- `storage.googleapis.com` (GeoIP updates for Elasticsearch)
- `notary.kolide.co` (osquery agent update)

- Ubuntu PPAs (OS Updates - Ubuntu only)
- download.docker.com (Docker packages - Ubuntu only)
- repo.saltstack.com (Salt packages - Ubuntu only)
- packages.wazuh.com (Wazuh packages - Ubuntu only)

In the case of a distributed deployment, you can configure your nodes to pull everything from the manager so that only the manager requires Internet access. Nodes will be configured to pull from `repocache.securityonion.net` but this URL does not actually exist on the Internet, it is just a special address for the manager proxy.

### 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:

- 22 (only needed for initial setup)
- 3142 (Apt-cacher-ng) (if manager proxy enabled, this is `repocache.securityonion.net` as mentioned above)
- 5000 (Docker registry)
- 8086 (influxdb)
- 4505 (Salt)
- 4506 (Salt)
- 5644 (Filebeat)
- 443 (Sensoroni)
- 8080 (Osquery, if enabled)

Search nodes from/to manager:

- 9300 (Node-to-node for Elasticsearch)
- 9696 (Redis)

### 14.3.2 Host Firewall

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

### 14.3.3 Port Groups

Port groups are a way of grouping together ports similar to a firewall port/service alias. For example, if you had a web server you could include 80 and 443 tcp into an alias or in this case a port group.

### 14.3.4 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.

### 14.3.5 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 using either *so-allow*, *so-firewall* or manually editing the yaml files. When setup is run on a new node, it will SSH to the manager using the *soremove* account and 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.

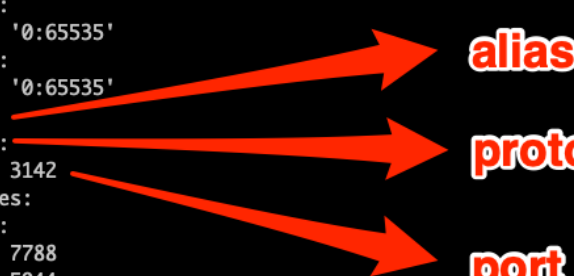
There are two directories that contain the yaml files for the firewall configuration.

`/opt/so/saltstack/default/salt/firewall/`

This directory contains the default firewall rules. The files in this directory should not be modified as they could possibly be overwritten during a *soup* update in the event we update those files.

`/opt/so/saltstack/default/salt/firewall/portgroups.yaml` is where the default port groups are defined.

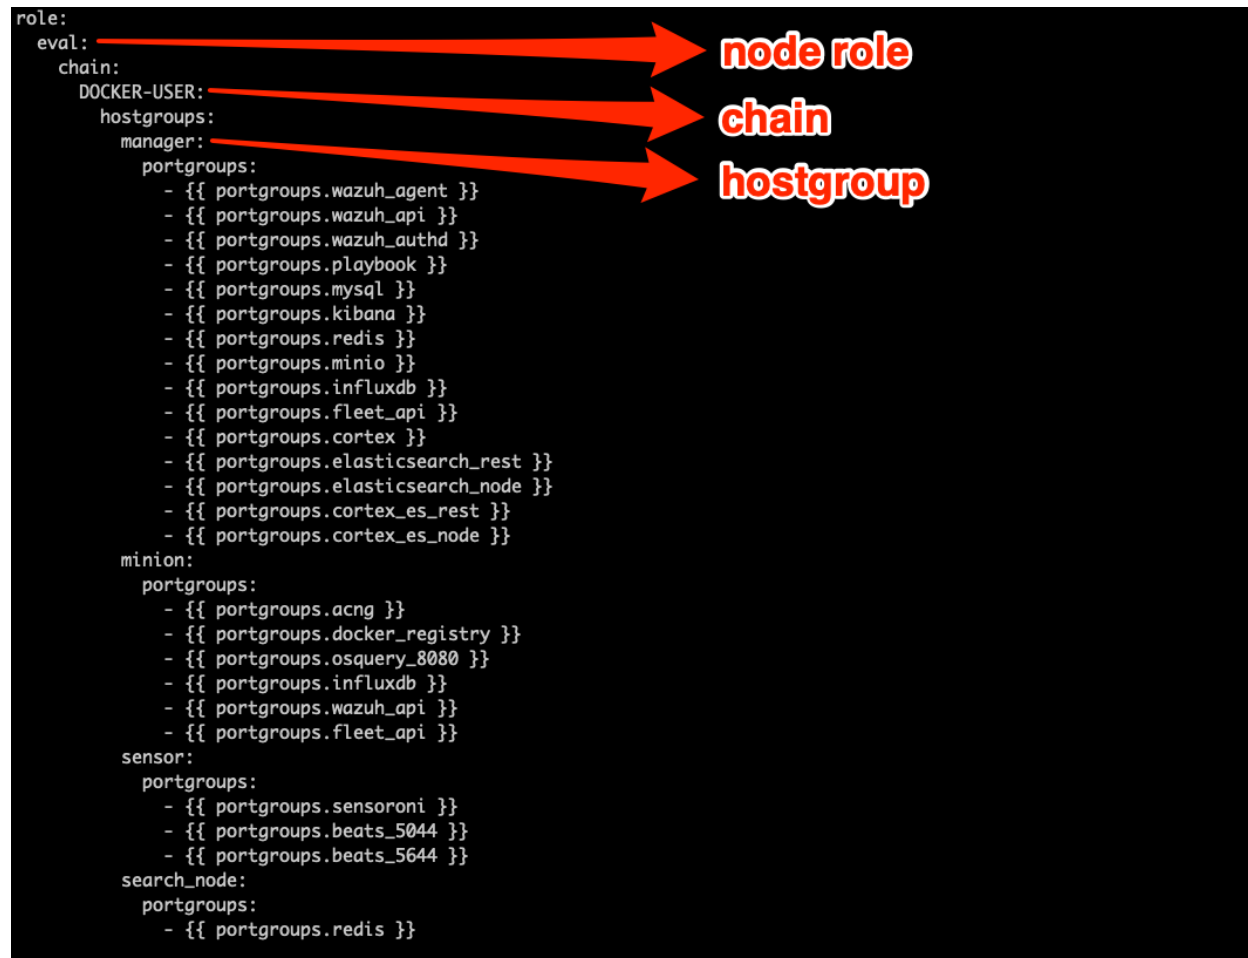
```
firewall:
  aliases:
    ports:
      all:
        tcp:
          - '0:65535'
        udp:
          - '0:65535'
      acng:
        tcp:
          - 3142
      agrules:
        tcp:
          - 7788
      beats_5044:
        tcp:
          - 5044
      beats_5644:
        tcp:
          - 5644
      cortex:
        tcp:
          - 9001
      cortex_es_node:
        tcp:
          - 9500
      cortex_es_rest:
        tcp:
```



The diagram illustrates the components of the firewall rule configuration. Three red arrows point from the 'acng' entry in the YAML file to the labels 'alias', 'proto', and 'port'. The 'alias' label points to the 'acng' key, the 'proto' label points to the 'tcp' key under 'acng', and the 'port' label points to the '3142' value under 'acng'.

`/opt/so/saltstack/default/salt/firewall/hostgroups.yaml` is where the default hostgroups are defined. There isn't much in here other than anywhere, dockernet, localhost and self.

`/opt/so/saltstack/default/salt/firewall/assigned_hostgroups.map.yaml` is where the default allow rules come together and pair hostgroups and portgroups and assign that pairing to a node based on its role in the grid. In the image below, we can see how we define some rules for an eval node.



### `/opt/so/saltstack/local/salt/firewall/`

This directory stores the firewall rules specific to your grid.

`/opt/so/saltstack/local/salt/firewall/portgroups.local.yaml` defines custom port groups.

`/opt/so/saltstack/local/salt/firewall/hostgroups.local.yaml` is where many default named hostgroups get populated with IPs that are specific to your environment. When you run `so-allow` or `so-firewall`, it modifies this file to include the IP provided in the proper hostgroup. Some node types get their IP assigned to multiple host groups.

`/opt/so/saltstack/local/salt/firewall/assigned_hostgroups.local.map.yaml` is where host group and port group associations would be made to create custom host group and port group assignments that would apply to all nodes of a certain role type in the grid.

### 14.3.6 Managing

Managing firewall rules for all devices should be done from the manager node using either *so-allow*, *so-firewall* or, for advanced cases, manually editing the yaml files.

### 14.3.7 Examples

#### Removing a host or network

If you previously added a host or network to your firewall configuration and now need to remove them, you can use *so-firewall* with the *excludehost* option. For example:

```
sudo so-firewall excludehost analyst 192.168.1.255
```

#### Allow hosts to send syslog to a sensor node

By default, if you use *so-allow* to add a host to the syslog hostgroup, that host will only be allowed to connect to the manager node. If we want to allow a host or group of hosts to send syslog to a sensor, then we can do the following:

1. Create a new host group that will contain the IPs of the hosts that you want to allow to connect to the sensor. This will add the host group to `/opt/so/saltstack/local/salt/firewall/hostgroups.local.yaml`. If the host group already exists, you can skip to step 2. Run the following on the manager:

```
sudo so-firewall addhostgroup <GROUP_NAME>
```

2. Add the desired IPs to the host group. This will add the IPs to the host group in `/opt/so/saltstack/local/salt/firewall/hostgroups.local.yaml`.

```
sudo so-firewall includehost <GROUP_NAME> <IP>
```

3. Since we reused the syslog port group that is already defined, we don't need to create a new port group. Now we have to build the association between the host group and the syslog port group and assign that to our sensor node. Add the following to the sensor minion pillar file located at `/opt/so/saltstack/local/pillar/minions/<HOSTNAME>_<ROLE>.sls`:

```
firewall:
  assigned_hostgroups:
    chain:
      DOCKER-USER:
        hostgroups:
          syslogtosensor1:
            portgroups:
              - portgroups.syslog
```

4. Now that the configuration is in place, you can either wait for the sensor to sync with Salt running on the manager, or you can force it to update its firewall by running the following from the manager:

```
sudo salt <HOSTNAME>_<ROLE> state.apply firewall
```

## Modify a default port group

In this example, we will be extending the default nginx port group to include port 8086 for a standalone node. By default, only the analyst hostgroup is allowed access to the nginx ports. At the end of this example IPs in the analyst host group, will be able to connect to 80, 443 and 8086 on our standalone node.

All the following will need to be run from the manager.

1. Add the custom nginx port group:

```
sudo so-firewall addportgroup nginx
```

2. Add the required ports to the port group. In this step we are redefining the nginx port group, so be sure to include the default ports as well if you want to keep them:

```
sudo so-firewall addport nginx tcp 80
sudo so-firewall addport nginx tcp 443
sudo so-firewall addport nginx tcp 8086
```

3. Associate this port group redefinition to a node. Add the following to the minion's sls file located at `/opt/so/saltstack/local/pillar/minions/<HOSTNAME>_<ROLE>.sls`:

```
firewall:
  assigned_hostgroups:
    chain:
      DOCKER-USER:
        hostgroups:
          analyst:
            portgroups:
              - portgroups.nginx
```

4. Apply the firewall state to the node, or wait for the highstate to run for the changes to happen automatically:

```
sudo salt-call state.apply firewall
```

**Warning:** Please review the [Salt](#) section to understand pillars and templates. Modifying these values outside of `so-allow` or `so-firewall` could lead to problems accessing your existing hosts. This is an advanced case and you most likely won't need to modify these files.

## 14.4 Email Configuration

Some applications rely on having a mail server in the OS itself and other applications (like [Wazuh](#)) have their own mail configuration and so they don't rely on a mail server in the OS itself.



### 14.4.1 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`.

### 14.4.2 Elastalert


Follow the steps on the [Elastalert](#) page.

### 14.4.3 Wazuh

Follow the steps in the [Wazuh](#) section.

## 14.5 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 use your preferred NTP provider. Setup will ask if you want to configure NTP:



If you're using our Security Onion ISO image, this can be set manually in `/etc/chrony.conf`.

### 14.5.1 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.

## 14.6 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.

### 14.6.1 Distributed Deployments

For distributed deployments, nodes only connect to the manager via SSH when they initially join the grid. That initial connection is done using the `soremove` account. If you enable 2FA for SSH, you will need to disable 2FA for the `soremove` account. The `soremove` account can be disabled when you are not adding any nodes to the grid.

### 14.6.2 Hardening

Some organizations require the removal of certain ciphers and algorithms from `sshd`. Setup will automatically do this for you by running `so-ssh-harden`:

**NOTE:** You will receive a warning upon SSH reconnect that the host key has changed.

This is expected due to hardening of the OpenSSH server config.

The host key algorithm will now be ED25519, follow the instructions given by your SSH client to remove the old key fingerprint then retry the connection.

<Ok>

Alternatively, you can manually run `sudo so-ssh-harden` or manually modify your `sshd_config` as follows:

```
sshd -T | grep "^ciphers" | sed -e "s/(3des-cbc|aes128-cbc|aes192-cbc|aes256-cbc|
↪arcfour|arcfour128|arcfour256|blowfish-cbc|cast128-cbc|rijndael-cbc@lysator.liu.
↪se)\,|/?//g" >> /etc/ssh/sshd_config
sshd -T | grep "^kexalgorithms" | sed -e "s/(diffie-hellman-group14-sha1|ecdh-sha2-
↪nistp256|diffie-hellman-group-exchange-sha256|diffie-hellman-group1-sha1|diffie-
↪hellman-group-exchange-sha1|ecdh-sha2-nistp521|ecdh-sha2-nistp384)\,|/?//g" >> /etc/
↪ssh/sshd_config
sshd -T | grep "^macs" | sed -e "s/(hmac-sha2-512,|umac-128@openssh.com,|hmac-sha2-
↪256,|umac-64@openssh.com,|hmac-sha1,|hmac-sha1-etm@openssh.com,|umac-64-
↪etm@openssh.com,|hmac-sha1)\,|/?//g" >> /etc/ssh/sshd_config
sshd -T | grep "^hostkeyalgorithms" | sed "s|ecdsa-sha2-nistp256,||g" | sed "s|ssh-rsa,
↪||g" >> /etc/ssh/sshd_config
```

**Warning:** Any time you modify `sshd_config`, there is a possibility of a syntax error preventing ssh from starting correctly which would then prevent you from accessing remotely. Please exercise caution in editing the file and have a backup method of accessing the box just in case.

### 14.6.3 Banner

Some organizations require a warning banner for SSH connections. Security Onion automatically configures SSH to display a warning banner.

```
#####
#####
###                                ###
###  UNAUTHORIZED ACCESS PROHIBITED  ###
###                                ###
#####
#####
```

You can customize this default Security Onion warning banner by editing `/etc/ssh/sshd-banner` with your favorite text editor.

## 14.7 Changing 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.

## 14.8 Changing IP Addresses

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.

## 14.9 Changing Web Access URL

If you need to change the URL for web access to Security Onion (for example, from IP to FQDN), do the following:

- Change the value for `url_base` in `/opt/so/saltstack/local/pillar/global.sls` (on the manager, if a distributed deployment).
- Run `sudo salt-call state.highstate` on all nodes (can first run on the manager, then `sudo salt \* state.highstate` from the manager for the remaining nodes).



## TUNING

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 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 2 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.

---

#### 15.1.1 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>

#### 15.1.2 Checking Status

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

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

### 15.1.3 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'
```

### 15.1.4 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
```

### 15.1.5 Configuration

Many of the options that are configurable in Security Onion 2 are done via pillar assignments in either the global or minion pillar files. Pillars are a Saltstack concept, formatted typically in YAML, that can be used to parameterize states via templating. Saltstack states are used to ensure the state of objects on a minion. In many of the use cases below, we are providing the ability to modify a configuration file by editing either the global or minion pillar file.

**Global pillar file:** This is the pillar file that can be used to make global pillar assignments to the nodes. It is located at `/opt/so/saltstack/local/pillar/global.sls`.

**Minion pillar file:** This is the minion specific pillar file that contains pillar definitions for that node. Any definitions made here will override anything defined in other pillar files, including global. This is located at `/opt/so/saltstack/local/pillar/minions/<minionid>.sls`.

**Default pillar file:** This is the pillar file located under `/opt/so/saltstack/default/pillar/`. Files here should not be modified as changes would be lost during a code update.

**Local pillar file:** This is the pillar file under `/opt/so/saltstack/local/pillar/`. These are the files that will need to be changed in order to customize nodes.

**Warning:** Salt sls files are in YAML format. When editing these files, please be very careful to respect YAML syntax, especially whitespace. For more information, please see [https://docs.saltproject.io/en/latest/topics/troubleshooting/yaml\\_idiosyncrasies.html](https://docs.saltproject.io/en/latest/topics/troubleshooting/yaml_idiosyncrasies.html).

Here are some of the items that can be customized with pillar settings:

- *Filebeat*
- *Firewall*
- *Managing Alerts*
- *Suricata*
- *Zeek*

### 15.1.6 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.

If you need to increase this delay, it can be done using the `salt:minion:service_start_delay` pillar. This can be done in the minion pillar file if you want the delay for just that minion, or it can be done in the `global.sls` file if it should be applied to all minions.

```
salt:
  minion:
    service_start_delay: 60 # in seconds.
```

Please keep this value below 90 seconds otherwise systemd will reach timeout and terminate the service.

### 15.1.7 Diagnostic Logs

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

### 15.1.8 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.

### 15.1.9 More Information

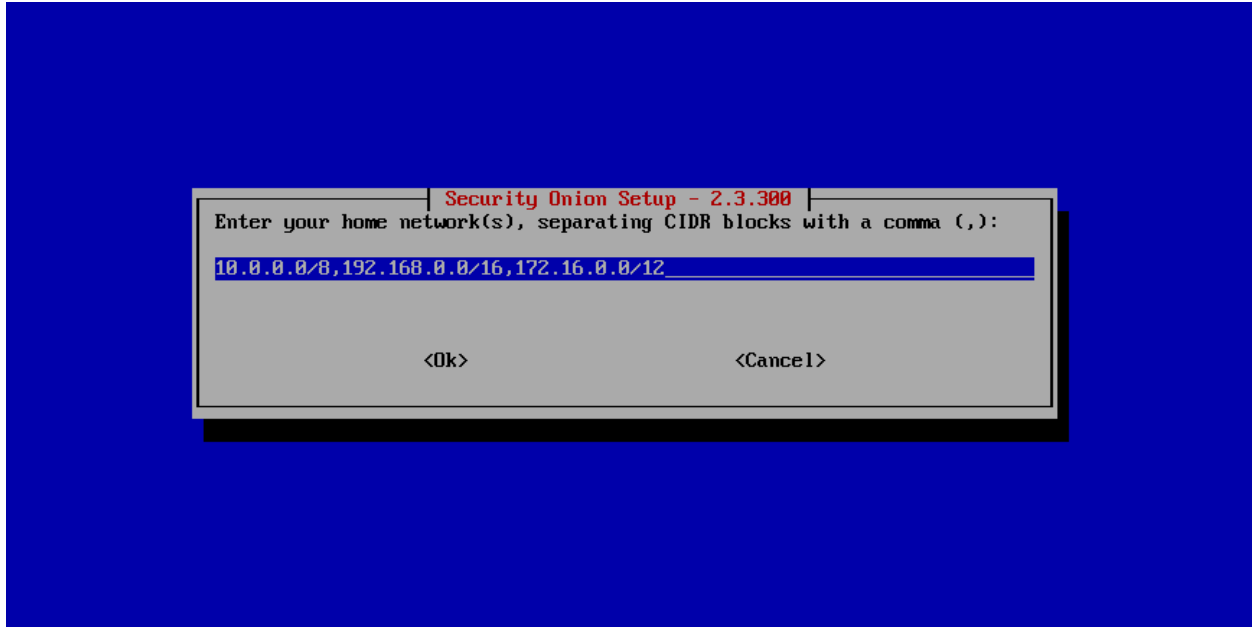
---

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

---

## 15.2 Homenet

The `homenet` variable defines the networks that are considered home networks (those networks that you are monitoring and defending). It is used for both *Suricata* and *Zeek*. The default value is RFC1918 private address space:



### 15.2.1 Configuration

A node can be assigned either the global homenet or its own homenet.

By default, a node will use the global homenet pillar value if it is defined in the global pillar file (`/opt/so/saltstack/local/pillar/global.sls`) under `global:hnmanager`.

```
global:
  soversion: '2.3.0'
  hnmanager: '10.0.0.0/8,192.168.0.0/16,172.16.0.0/12'
```

In order to define a per node homenet, it can be defined in the minion pillar file (`/opt/so/saltstack/local/pillar/minions/$SENSORNAME_$ROLE.sls`) under `sensor:hnsensor`.

```
sensor:
  interface: 'bond0'
  mainip: '172.16.106.112'
  mainint: 'eth0'
  zeek_lbprocs: 5
  suriprocs: 2
  manager: 'somanager1'
  mtu: 1500
  uniqueid: 1602623674
  hnsensor: 10.0.0.0/8
```

In order to sync the configuration change with the node, we can either wait for the node to automatically highstate on the predefined interval or we can force it. Since this homenet applies to *Suricata* and *Zeek*, we can apply the *suricata* and *zeek* states to the node.

- From the manager:

```
sudo salt $SENSORNAME_$ROLE state.apply suricata,zeek
```

or



- From the node:

```
sudo salt-call state.apply suricata,zeek
```

## 15.2.2 More Information

**Note:** For more information about Suricata and Zeek, please see the *Suricata* and *Zeek* sections.

## 15.3 BPF

BPF stands for Berkeley Packet Filter. From [https://en.wikipedia.org/wiki/Berkeley\\_Packet\\_Filter](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.

### 15.3.1 Configuration

#### Global BPF

You can specify your BPF in the global pillar on your manager node (`/opt/so/saltstack/local/pillar/global.sls`) and it will apply to all interfaces in your entire deployment by default. If there is no BPF configuration already in the file, you can append it to the bottom of the file.

If you have separate sensors reporting to that manager node, they will pull down the relevant BPF as part of the Salt update that runs every 15 minutes and then restart *Suricata*/*Stenographer*/*Zeek* so that the BPF change will take effect.

*Stenographer* example:

```
steno:
  bpf:
    - "Your BPF Here"
```

*Suricata* example:

```
nids:
  bpf:
    - "Your BPF Here"
```

*Zeek* example:

```
zeek:
  bpf:
    - "Your BPF Here"
```

## Node-Specific BPF

If you don't want your sensors to inherit BPF from the manager node, you can edit the minion sls file (`/opt/so/saltstack/local/pillar/minions/$Hostname.sls`), which will override any global BPF settings set from the global pillar.

### Simple Example

Most organizations want to configure *Stenographer* to not record full packet capture for port 443 since it should be mostly encrypted traffic anyway:

```
steno:
  bpf:
    - not port 443
```

### Quoting

YAML rules apply and so if you want to use a reserved YAML character such as `[] {} > | * & ! % # ` @ ,`, then you may need to enclose the entire line in double quotes. For example:

```
steno:
  bpf:
    - "(port 443)"
```

### 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`:

```
nids:
  bpf:
    - not host 192.168.1.2 and not host 192.168.1.3
```

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

```
nids:
  bpf:
    - 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>

## 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>

## 15.3.2 More Information

---

**Note:** Check out our BPF video at <https://youtu.be/uamNOjtUi4Y!>

For more information about BPF, please see:

[https://en.wikipedia.org/wiki/Berkeley\\_Packet\\_Filter](https://en.wikipedia.org/wiki/Berkeley_Packet_Filter)

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

---

## 15.4 Managing Rules

### 15.4.1 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
```

## 15.4.2 Rulesets

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

## 15.4.3 ET Open

- optimized for *Suricata*, but available for Snort as well
- **free**

For more information, see:

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

## 15.4.4 ET Pro (Proofpoint)

- optimized for *Suricata*, but available for Snort as well
- rules retrievable as released
- license fee per sensor (users are responsible for purchasing enough licenses for their entire deployment)

To enable the ET Pro ruleset in an already installed grid, modify the `/opt/so/saltstack/local/pillar/minions/<manager.sls>` file as follows:

```
idstools:
  config:
    ruleset: 'ETPRO'
    oinkcode: 'MYOINKCODE'
```

For more information, see:

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

## 15.4.5 Snort Community

- optimized for Snort
- community-contributed rules
- **free**

For more information, see:

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

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

### 15.4.6 Snort Registered

- optimized for Snort
- Snort SO (Shared Object) rules only work with Snort not *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>

### 15.4.7 Snort Subscriber (Talos)

- optimized for Snort
- Snort SO (Shared Object) rules only work with Snort not *Suricata*
- rules retrievable as released
- license fee per sensor (users are responsible for purchasing enough licenses for their 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.

To enable the Talos Subscriber ruleset in an already installed grid, modify the `/opt/so/saltstack/local/pillar/minions/<manager.sls>` file as follows:

```
idstools:
  config:
    ruleset: 'TALOS'
    oinkcode: 'MYOINKCODE'
```

For more information, see:

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

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

### 15.4.8 Other

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

To add other remotely-accessible rulesets, add an entry under `urls` for the ruleset URL in `/opt/so/saltstack/local/pillar/minions/<manager.sls>`:

```
idstools:
  config:
    ...primary ruleset...
    ...primary ruleset oinkcode...
  urls:
    - https://ruleseturlhere
```

## 15.5 Adding Local Rules

### 15.5.1 NIDS

You can add NIDS rules in `/opt/so/saltstack/local/salt/idstools/local.rules` on your manager. Within 15 minutes, *Salt* should then copy those rules into `/opt/so/rules/nids/local.rules`. The next run of *idstools* should then merge `/opt/so/rules/nids/local.rules` into `/opt/so/rules/nids/all.rules` which is what *Suricata* reads from.

If you don't want to wait for these automatic processes, you can run them manually from the manager (replacing `$SENSORNAME_$ROLE` as necessary):

```
sudo salt-call state.highstate
sudo so-rule-update
sudo salt $SENSORNAME_$ROLE state.apply suricata
```

For example:

- Let's add a simple rule to `/opt/so/saltstack/local/salt/idstools/local.rules` that's really just a copy of the traditional `id check returned root` rule:

```
alert ip any any -> any any (msg:"GPL ATTACK_RESPONSE id check returned root 2";
↪content:"uid=0|28|root|29|"; classtype:bad-unknown; sid:7000000; rev:1;)
```

- From the manager, tell *Salt* to update:

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

- Update rules:

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

- Restart *Suricata* (replacing `$SENSORNAME_$ROLE` as necessary):

```
sudo salt $SENSORNAME_$ROLE state.apply suricata
```

- If you built the rule correctly, then *Suricata* should be back up and running.
- You can then run `curl http://testmynids.org/uid/index.html` on the node to generate traffic which should cause this rule to alert (and the original rule that it was copied from, if it is enabled).

## 15.5.2 HIDS

You can add *Wazuh* HIDS rules in `/opt/so/rules/hids/local_rules.xml`.

## 15.5.3 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, modify `/opt/so/saltstack/default/salt/strelka/rules/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.

## 15.6 Managing Alerts

Security Onion generates a lot of valuable information for you the second you plug it into a TAP or SPAN port. Between *Zeek* logs, alert data from *Suricata*, and full packet capture from *Stenographer*, you have enough information to begin identifying areas of interest and making positive changes to your security stance.

---

**Note:** 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.

---

### 15.6.1 Alerting Engines & Severity

There are three alerting engines within Security Onion: *Suricata*, *Wazuh* 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*.

## 15.6.2 Wazuh HIDS Alerts

If you want to tune Wazuh HIDS alerts, please see the [Wazuh](#) section.

## 15.6.3 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/all.rules`:

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

You can also test using `so-test`.

## 15.6.4 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 rules in each:

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

## 15.6.5 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.

---

**Note:** Check out our NIDS tuning video at <https://youtu.be/1jEkFIEUCuI>!

---



### 15.6.6 so-rule

so-rule allows you to disable, enable, or modify NIDS rules. Run so-rule without any options to see the help output:

```
so-rule
usage: so-rule [-h] ...

optional arguments:
  -h, --help  show this help message and exit

commands:
  disabled    Manage and list disabled rules (add, remove, list)
  enabled     Manage and list enabled rules (add, remove, list)
  modify      Manage and list modified rules (add, remove, list)
```

### 15.6.7 Disable the SID

We can use so-rule to modify an existing NIDS rule. For example, suppose we want to disable SID 2100498. We can start by listing any currently disabled rules:

```
sudo so-rule disabled list
No rules disabled.
```

Next, let's disable SID 2100498:

```
sudo so-rule disabled add 2100498
Configuration updated. Would you like to apply your changes now? (y/N) y
Applying idstools state...
```

Once that completes, we can then verify that 2100498 is now disabled with so-rule disabled list:

```
sudo so-rule disabled list
Disabled rules:
- 2100498
```

Finally, we can check that 2100498 is commented out in /opt/so/rules/nids/all.rules:

```
grep 2100498 /opt/so/rules/nids/all.rules
# alert ip any any -> any any (msg:"GPL ATTACK_RESPONSE id check returned root"; content:
↪ "uid=0|28|root|29|"; classtype:bad-unknown; sid:2100498; rev:7; metadata:created_at_
↪ 2010_09_23, updated_at 2010_09_23;)
```

If you can't run so-rule, then you can modify configuration manually. Security Onion uses idstools to download new signatures every night and process them against a set list of user generated configurations. To enable or disable SIDs for *Suricata*, the *Salt* idstools pillar can be used in the minion pillar file (/opt/so/saltstack/local/pillar/minions/<managename>\_<role>.sls). In a distributed Security Onion environment, you only need to change the configuration in the manager pillar and then all other nodes will get the updated rules automatically.

If SID 4321 is noisy, you can disable it as follows:

```
idstools:
  sids:
    disabled:
      - 4321
```

From the manager, run the following to update the config:

```
sudo salt $SENSORNAME_$ROLE state.apply idstools
```

If you want to disable multiple rules at one time, you can use a regular expression, but make sure you enclose the full entry in single quotes like this:

```
idstools:
  sids:
    disabled:
      - 're:heartbleed'
```

## 15.6.8 Modify the SID

We can use `so-rule` to modify an existing NIDS rule. For example, suppose that we want to modify SID 2100498 and replace any instances of “returned root” with “returned root test”. We can start by listing any rules that are currently modified:

```
sudo so-rule modify list
No rules currently modified.
```

Let’s first check the syntax for the add option:

```
sudo so-rule modify add -h
usage: so-rule modify add [-h] [--apply] SID|REGEX SEARCH_TERM REPLACE_TERM

positional arguments:
  SID|REGEX      A valid SID (ex: "4321") or regular expression pattern (ex:
                  "re:heartbleed|spectre")
  SEARCH_TERM    A quoted regex search term (ex: "\\$EXTERNAL_NET")
  REPLACE_TERM   The text to replace the search term with

optional arguments:
  -h, --help      show this help message and exit
  --apply         After updating rule configuration, apply the idstools state.
```

Now that we understand the syntax, let’s add our modification:

```
sudo so-rule modify add 2100498 "returned root" "returned root test"
Configuration updated. Would you like to apply your changes now? (y/N) y
Applying idstools state...
```

Once the command completes, we can verify that our modification has been added:

```
sudo so-rule modify list
Modified rules + modifications:
  - 2100498 "returned root" "returned root test"
```

Finally, we can check the modified rule in `/opt/so/rules/nids/all.rules`:

```
grep 2100498 /opt/so/rules/nids/all.rules
alert ip any any -> any any (msg:"GPL ATTACK_RESPONSE id check returned root test";
↪content:"uid=0|28|root|29|"; classtype:bad-unknown; sid:2100498; rev:7;
↪metadata:created_at 2010_09_23, updated_at 2010_09_23;)
```

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:

```
sudo so-rule modify add 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.

If you can't run `so-rule`, you can modify the configuration manually in the manager pillar file at `/opt/so/saltstack/local/pillar/minions/<managername>_<role>.sls` (where `<role>` is `manager`, `managersearch`, `standalone`, or `eval` depending on the manager type that was chosen during install). In this file, the `idstools` section has a `modify` sub-section where you can add your modifications. For example:

```
idstools:
  sids:
    modify:
      - '2019401 "seconds \d+" "seconds 3600"'
```

From the manager, run the following to update the config:

```
sudo salt $SENSORNAME_$ROLE state.apply idstools
```

If you need to modify a part of a rule that contains a special character, such as a \$ in variable names, the special character needs to be escaped in the search part of the modify string. For example:

```
idstools:
  sids:
    modify:
      - '2826931 "\\$EXTERNAL_NET" "!$HOME_NET"'
```

## 15.6.9 Rewrite the signature

In some cases, you may not want to use the modify option above, but instead create a local copy of the rule as shown in the [Adding Local Rules](#) section and then disable the original rule.

- Edit the `/opt/so/saltstack/local/salt/idstools/local.rules` file using `vi` or your favorite text editor:

```
sudo vi /opt/so/saltstack/local/salt/idstools/local.rules
```

- Paste the rule. You may want to bump the SID into the 90,000,000 range and set the revision to 1.
- Now that we have a signature that will generate alerts a little more selectively, we need to disable the original signature. As shown above, we edit the minion pillar and add the SID to the `idstools - sids - disabled` section.
- Finally, from the manager, update the config on the remote node:

```
salt $SENSORNAME_$ROLE state.highstate
```

### 15.6.10 Threshold

You can manage threshold entries for *Suricata* using *Salt* pillars. The format of the pillar file can be seen below, as well as in `/opt/so/saltstack/default/pillar/thresholding/pillar.usage` and `/opt/so/saltstack/default/pillar/thresholding/pillar.example`

---

**Note:** The signature id (SID) must be unique. If you have multiple entries for the same SID, it will cause an error in salt resulting in all of the nodes in your grid to error out when checking in.

---

Usage:

```
thresholding:
  sids:
    <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>
```

Example:

```
thresholding:
  sids:
    8675309:
      - threshold:
          gen_id: 1
          type: threshold
          track: by_src
          count: 10
          seconds: 10
      - threshold:
          gen_id: 1
          type: limit
          track: by_dst
          count: 100
          seconds: 30
      - rate_filter:
          gen_id: 1
          track: by_rule
          count: 50
```

(continues on next page)

(continued from previous page)

```

    seconds: 30
    new_action: alert
    timeout: 30
  - suppress:
    gen_id: 1
    track: by_either
    ip: 10.10.3.7
11223344:
  - threshold:
    gen_id: 1
    type: limit
    track: by_dst
    count: 10
    seconds: 10
  - rate_filter:
    gen_id: 1
    track: by_src
    count: 50
    seconds: 20
    new_action: pass
    timeout: 60
  - suppress:
    gen_id: 1
    track: by_src
    ip: 10.10.3.0/24

```

In order to apply the threshold to all nodes, place the pillar in `/opt/so/saltstack/local/pillar/global.sls`. If you want to apply the threshold to a single node, place the pillar in `/opt/so/saltstack/local/pillar/minions/<MINION_ID>.sls`

**Warning:**

Salt sls files are in YAML format. When editing these files, please be very careful to respect YAML syntax, especially whitespace. For more information, please see:

[https://docs.saltproject.io/en/latest/topics/troubleshooting/yaml\\_idiosyncrasies.html](https://docs.saltproject.io/en/latest/topics/troubleshooting/yaml_idiosyncrasies.html)

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>.

For example, the following threshold IP exceeds the 64-character limit:

```

thresholding:
  sids:
    2012454:
      - suppress:
        gen_id: 1
        track: by_dst
        ip: 1.1.1.1,2.2.2.2,3.3.3.3,4.4.4.4,5.5.5.5,6.6.6.6,7.7.7.7,8.8.8.8,9.9.9.9,10.
↪10.10.10,11.11.11.11

```

This results in the following error in the *Suricata* log:

```
<Error> - [ERRCODE: SC_ERR_PCRE_COPY_SUBSTRING(325)] - pcre_copy_substring failed
```

The solution is to break the ip field into multiple entries like this:

```
thresholding:
  sids:
    2012454:
      - suppress:
          gen_id: 1
          track: by_dst
          ip: 1.1.1.1,2.2.2.2,3.3.3.3,4.4.4.4,5.5.5.5,6.6.6.6,7.7.7.7,8.8.8.8
      - suppress:
          gen_id: 1
          track: by_dst
          ip: 9.9.9.9,10.10.10.10,11.11.11.11
```

### 15.6.11 Suppressions

A suppression rule allows you to make some finer grained decisions about certain rules without the onus of rewriting them. With this functionality we can suppress rules based on their signature, the source or destination address and even the IP or full CIDR network block. This way, you still have the basic ruleset, but the situations in which they fire are altered. It's important to note that with this functionality, care should be given to the suppressions being written to make sure they do not suppress legitimate alerts. See above for `suppress` examples.

### 15.6.12 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`.

## 15.7 High Performance Tuning

### 15.7.1 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.

### 15.7.2 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>

### 15.7.3 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>

### 15.7.4 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/](https://lonesysadmin.net/2013/12/22/better-linux-disk-caching-performance-vm-dirty_ratio/)

### 15.7.5 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 *Filebeat*, *Logstash*, *Redis*, and *Elasticsearch*.



## TRICKS AND TIPS

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

### 16.1 Backups

Security Onion 2 performs daily backups 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. 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 minion sls files and customizations.

#### 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. To backup this data, there are a few options.

The first option is to enable snapshots with *Curator* to snapshot data to an external storage device such as a NAS.

The second 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. To do this, you can update the *elasticsearch* *Salt* pillar using something like the following and then restart *Elasticsearch*:

```
elasticsearch:
  config:
    path.repo: '/usr/share/elasticsearch/data/repo'
```

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/`.

A third option is 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.

## 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

If you download our Security Onion ISO image, the Docker engine and these Docker images are baked right into the ISO image.

If you instead use another ISO image, our installer will download Docker images from [ghcr.io](https://ghcr.io) as necessary.

### 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 bridge with an IP of 172.17.0.1.

<https://docs.docker.com/engine/userguide/networking/#default-networks>

For many folks this is fine, but what if we actually use the 172.17.0.0/16 range within our internal network(s)? This results in a **conflict** when trying to assign IP addresses to interfaces and trying to route outside of the host.

It is currently possible to change this at install time. Once you change this default docker network you **MUST** configure all nodes in the grid to use this range:

- During setup choose change docker network range.
- Enter your desired address range. You do not need the /24 at the end.

### 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

### Fleet

so-fleet - REQ - Fleet Web App  
so-mysql - REQ - Fleet state data  
so-redis - REQ - Required for live querying

### 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 DNS Anomaly Detection

Dr. Johannes Ullrich of the SANS Internet Storm Center posted a great DNS Anomaly Detection script based on the query logs coming from his DNS server. We can do the same thing with *Zeek*'s `dns.log` (where *Zeek* captures all the DNS queries it sees on the network).

**Note:** Please note that the following script is only intended for standalone machines and will not work properly on distributed deployments. Another option which might work better is *ElastAlert* and its `new_term` rule.

Thanks to senatorhotchkiss on our mailing list for updating the original script to replace `bro-cut` with `jq`:

```
#!/bin/bash

ZEEK_LOGS="/nsm/zeek/logs"
TODAY=`date +%Y-%m-%d`
YESTERDAY=`date -d yesterday +%Y-%m-%d`
OLD_DIRS=`ls $ZEEK_LOGS | grep "20*-*" | egrep -v "current|stats|$TODAY|$YESTERDAY"`
TMPDIR=/tmp
OLDLOG=$TMPDIR/oldlog
NEWLOG=$TMPDIR/newlog
SUSPECTS=$TMPDIR/suspects

for DIR in $OLD_DIRS; do zcat $ZEEK_LOGS/$DIR/dns* | jq '{"id.resp_p"},"{"query"}' ; done
↪ | grep -v "^5353" | awk '{print $2}' | sort | uniq -c | sort -k2 > $OLDLOG
zcat $ZEEK_LOGS/$YESTERDAY/dns* | jq '{"id.resp_p"},"{"query"}' | grep -v "^5353" | awk '
↪ {print $2}' | sort | uniq -c | sort -k2 > $NEWLOG
join -1 2 -2 2 -a 2 $OLDLOG $NEWLOG | egrep -v '.* [0-9]+ [0-9]+$' | sort -nr -k2 |
↪ head -50 > $SUSPECTS

echo
echo "=====
echo "Top 50 First Time Seen DNS queries:"
echo "=====
cat $SUSPECTS
```

## 16.4 ICMP Anomaly Detection

At Security Onion Conference 2016, Eric Conrad shared some IDS rules for detecting unusual ICMP echo requests/replies and identifying C2 channels that may utilize ICMP tunneling for covert communication.

### 16.4.1 Usage

We can add the rules to `/opt/so/rules/nids/local.rules` and the variables to `suricata.yaml` so that we can gain better insight into ICMP echoes or replies over a certain size, containing particularly suspicious content, etc.

### 16.4.2 Presentation

You can find Eric's presentation here:

<https://www.ericconrad.com/2016/09/c2-phone-home-leveraging-securityonion.html>

### 16.4.3 Download

You can download the rules here:

<https://drive.google.com/file/d/0ByeHgv6rpa3gUDNuMUdobFBCNkk>

## 16.5 Jupyter Notebook

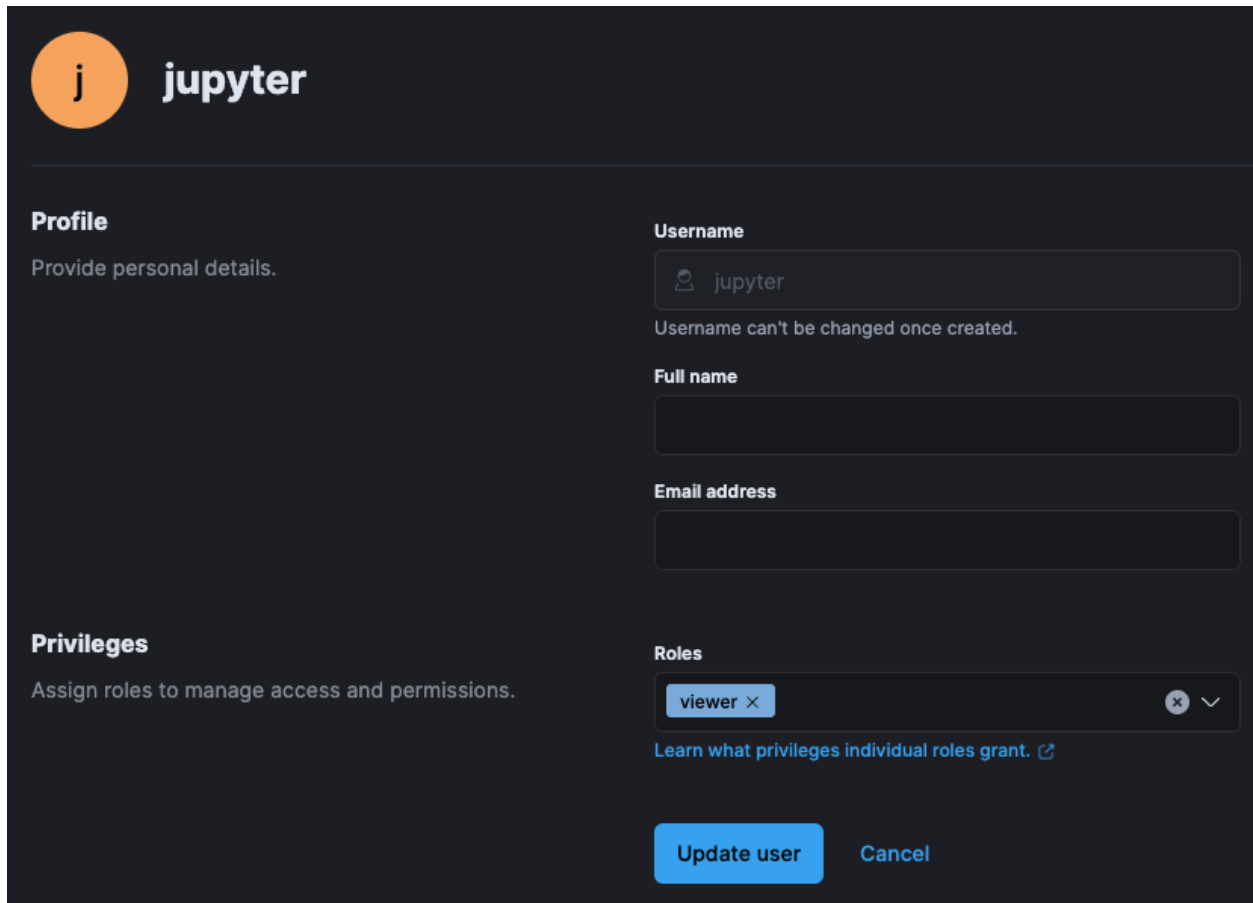
### 16.5.1 Overview

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>.

### 16.5.2 Security Onion Setup

#### Create Jupyter User

As a best practice, we'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 image shows a Jupyter user management interface. At the top left is the Jupyter logo (an orange circle with a white 'j') and the word 'jupyter' in white. Below this is a horizontal line. The interface is divided into two main sections: 'Profile' and 'Privileges'. The 'Profile' section has the heading 'Profile' and the subtext 'Provide personal details.' It contains three input fields: 'Username' (with a user icon and the text 'jupyter', and a note 'Username can't be changed once created.'), 'Full name', and 'Email address'. The 'Privileges' section has the heading 'Privileges' and the subtext 'Assign roles to manage access and permissions.' It contains a 'Roles' section with a dropdown menu showing 'viewer' and a link 'Learn what privileges individual roles grant.' with an external link icon. At the bottom right are two buttons: 'Update user' (blue) and 'Cancel' (grey).

### Security Onion Firewall

In order to allow network-based access to *Elasticsearch*, we'll need to allow the traffic through the host-based firewall using *so-allow*:

```
sudo so-allow
```

We'll choose the *e* option for the *Elasticsearch* REST API and provide our IP address(es) for which we could like to add an exception. Once complete, we should be able to reach the *Elasticsearch* instance. We can confirm connectivity using tools like *curl*, or Powershell's *Test-NetConnection*.

### 16.5.3 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!

## 16.6 Machine Learning

We have recently added a new tool (currently in beta) called logscan which utilizes machine learning models to detect anomalies in logs generated by Security Onion components.

**Warning:** Current and future ML components have dependencies that require special consideration to be made in regards to hardware or VM configurations prior to installation. Namely, a CPU/vCPU with AVX support is required, with AVX2 support recommended for better performance.

## 16.6.1 Listing components

To list all available ML components:

```
sudo so-learn list
```

---

**Note:** Currently logscan is the only ML component available. (Initially unavailable on air gapped installations. See warning below for more info.)

---

## 16.6.2 Enabling components

To enable an ML component:

```
sudo so-learn enable <component> # --apply to immediately apply your changes
```

## 16.6.3 Disabling components

To disable an ML component:

```
sudo so-learn disable <component> # --apply to immediately apply your changes
```

## 16.6.4 Logscan

**Warning:** Logscan will initially be unavailable on air gapped installations, therefore a networked installation is required to make use of the tool during this beta stage.

Logscan is log agnostic, but in its current implementation only scans logs from the built-in auth provider Kratos.

### Important Files and Directories

- App log: /opt/so/log/logscan/app.log
- Alerts log: /opt/so/log/logscan/alerts.log
- Data: /nsm/logscan/data

### Models

Logscan uses the following models to detect anomalous login activity on Security Onion Console:

- **K1:** Searches for high numbers of login attempts from single IPs in a 1 minute window
- **K5:** Searches for high ratios of login failures from single IPs in a 5 minute window
- **K60:** Searches for abnormal patterns of login failures from all IPs seen within a 1 hour window

## 16.7 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!

### 16.7.1 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.

### 16.7.2 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
```

### 16.7.3 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.

## 16.8 PCAPs for Testing

The easiest way to download pcaps 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://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 pcaps from the link above using a standard web browser or from the command line using a tool like *wget* or *curl*. Here are some examples.

To download the pcap from <https://www.malware-traffic-analysis.net/2020/09/16/index.html> using *wget*:

```
wget https://www.malware-traffic-analysis.net/2020/09/16/2020-09-16-Qakbot-infection-traffic.pcap.zip
```

To download a pcap from <https://www.netresec.com/?page=MACCDC>:

```
wget https://download.netresec.com/pcap/maccdc-2012/maccdc2012_000000.pcap.gz
```

### 16.8.1 tcpreplay

You can use *tcpreplay* to replay any standard pcap to the sniffing interface of your Security Onion sensor.

### 16.8.2 so-import-pcap

A drawback to using *tcpreplay* is that it's replaying the pcap as new traffic and thus the timestamps that you see in *Kibana* and other interfaces do not reflect the original timestamps from the pcap. To avoid this, a new tool was developed called *so-import-pcap*.

## 16.9 Removing a Node

There may come a time when you need to remove a node from your distributed deployment. To do this, you'll need to remove the node's configuration from a few different components.

### 16.9.1 Salt

First, log into your manager and list all *Salt* keys:

```
sudo salt-key
```

Then remove the node by deleting its key from *Salt* (replacing *nodename* with the actual node name):

```
sudo salt-key -d nodename
```

Remove the node from any `.sls` files in `/opt/so/saltstack/local/pillar/data/`.

### 16.9.2 Grafana

If necessary, remove the node's `json` file from the appropriate subdirectory under `/opt/so/conf/grafana/grafana_dashboards/` on the manager and then restart Grafana with:

```
sudo so-grafana-restart
```

You may also want to purge old Grafana data using `so-influxdb-clean` as described in the *Grafana* section.

### 16.9.3 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
```

### 16.9.4 FleetDM

You can delete the node from FleetDM via the *FleetDM* web interface.

### 16.9.5 Cross Cluster Search

If you are removing a search node, you will want to remove it from cross cluster search. To do so, you'll need to update that search node's settings in `_cluster/settings` and make sure that any settings are set to `null`. So you might want to start by doing the following query via *so-elasticsearch-query*:

```
sudo so-elasticsearch-query _cluster/settings
```

Then based on that output, update `_cluster/settings` by sending that node section back but with all settings set to `null`. You could use *so-elasticsearch-query* again or use *Kibana's Dev Tools* and paste something like the following text into the window (replacing *nodename* with the actual node name and adding any other settings as necessary):

```
PUT _cluster/settings
{
  "persistent": {
    "cluster": {
      "remote": {
        "nodename": {
          "skip_unavailable": null,
          "seeds": null
        }
      }
    }
  }
}
```

---

**Note:**

For more information, please see:

<https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-remote-clusters.html#configure-remote-clusters-dynamic>

---

## 16.10 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.

## 16.11 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.

## UTILITIES

This section covers some of the main 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

Security Onion locks down the *Firewall* by default. Depending on what kind of installation you do, Setup may walk you through allowing your analyst IP address(es). If you need to add other analyst IP addresses or open firewall ports for agents or syslog devices, you can run `sudo so-allow` and it will walk you through this process.

```
This program allows you to add a firewall rule to allow connections from a new ↵
↵ IP address.
```

```
Choose the role for the IP or Range you would like to add
```

(continues on next page)

(continued from previous page)

```
[a] - Analyst - ports 80/tcp and 443/tcp
[b] - Logstash Beat - port 5044/tcp
[e] - Elasticsearch REST API - port 9200/tcp
[f] - Strelka frontend - port 57314/tcp
[o] - Osquery endpoint - port 8090/tcp
[s] - Syslog device - 514/tcp/udp
[w] - Wazuh agent - port 1514/tcp/udp
[p] - Wazuh API - port 55000/tcp
[r] - Wazuh registration service - 1515/tcp
```

Please enter your selection:

## 17.2.1 Wazuh

If you choose the analyst option, `so-allow` will also add the analyst IP address to the *Wazuh* safe list. This will prevent *Wazuh* Active Response from blocking the analyst IP address.

## 17.2.2 Automation

In addition to the interactive menu shown above, you can pass desired options as command line arguments:

```
so-allow -h
```

```
Usage: /usr/sbin/so-allow [-abefhoprsrw] [ -i IP ]
```

This program allows you to add a firewall rule to allow connections **from** a new IP\_↵  
↵address **or** CIDR range.

If you run this program **with** no arguments, it will present a menu **for** you to choose your\_↵  
↵options.

If you want to automate **and** skip the menu, you can **pass** the desired options **as** command\_↵  
↵line arguments.

### EXAMPLES

To add 10.1.2.3 to the analyst role:

```
so-allow -a -i 10.1.2.3
```

To add 10.1.2.0/24 to the osquery role:

```
so-allow -o -i 10.1.2.0/24
```



## 17.3 so-elastic-auth

Security Onion supports Elastic authentication. This means that you will authenticate to *Elasticsearch* and *Kibana* using the same username and password that you use for *Security Onion Console (SOC)*.

Please note that if Elastic auth is enabled and you add a new user directly in *Kibana* via the Kibana Users page, then that new user will only have access to *Kibana* and no other apps. If you want the user to have access to all apps, make sure you add the user as shown in the *Adding Accounts* section.

### 17.3.1 New Installations

New installations will automatically enable Elastic auth.

### 17.3.2 Existing Installations

If you have an older installation (pre-2.3.60) that you've upgraded to a recent version and would like to enable Elastic auth, you can do so as follows:

```
sudo so-elastic-auth true
```

After manually enabling Elastic auth, each user will need to reset their password inside of *Security Onion Console (SOC)* as shown in the *Passwords* section and this will update their username and password in Elastic.

### 17.3.3 Service Accounts

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*.

### 17.5.1 Configuration

`so-import-pcap` requires you to run through Setup and choose a configuration that supports `so-import-pcap`. This includes Import Node and other nodes that include sensor services like Eval and Standalone. The quickest and easiest option is to choose Import Node which gives you the minimal services necessary to import a pcap.

If you're running a Security Onion Import node and run `so-status`, then it will show `so-steno`, `so-suricata`, and `so-zeek` as DISABLED since they are not sniffing live traffic. *Suricata* and *Zeek* will still analyze pcaps normally when running `so-import-pcap`. *Stenographer* is not used at all in Import mode.

### 17.5.2 Usage

Once Setup completes, you can then run `sudo so-import-pcap` and supply the full path to at least one pcap file. For example, to import a single pcap named `import.pcap`:

```
sudo so-import-pcap /full/path/to/import.pcap
```

To import multiple pcaps:

```
sudo so-import-pcap /full/path/to/import1.pcap /full/path/to/import2.pcap
```

Please note that if you import multiple pcaps at one time, `so-import-pcap` currently only provides a hyperlink for the last pcap in the list. If you need a hyperlink for each pcap, then you can run one pcap file per `so-import-pcap` and use a for-loop to iterate over your collection of pcap files.

`so-import-pcap` calculates the MD5 hash of the imported pcap and creates a directory in `/nsm/import/` for that hash. This is where `so-import-pcap` stores the alerts and logs generated by the traffic in the pcap. If you try to import that same pcap again, it will tell you that it has already imported that pcap. If for some reason you really do need to import that pcap again, you can remove that pcap's directory in `/nsm/import/` and then try again.

### 17.5.3 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

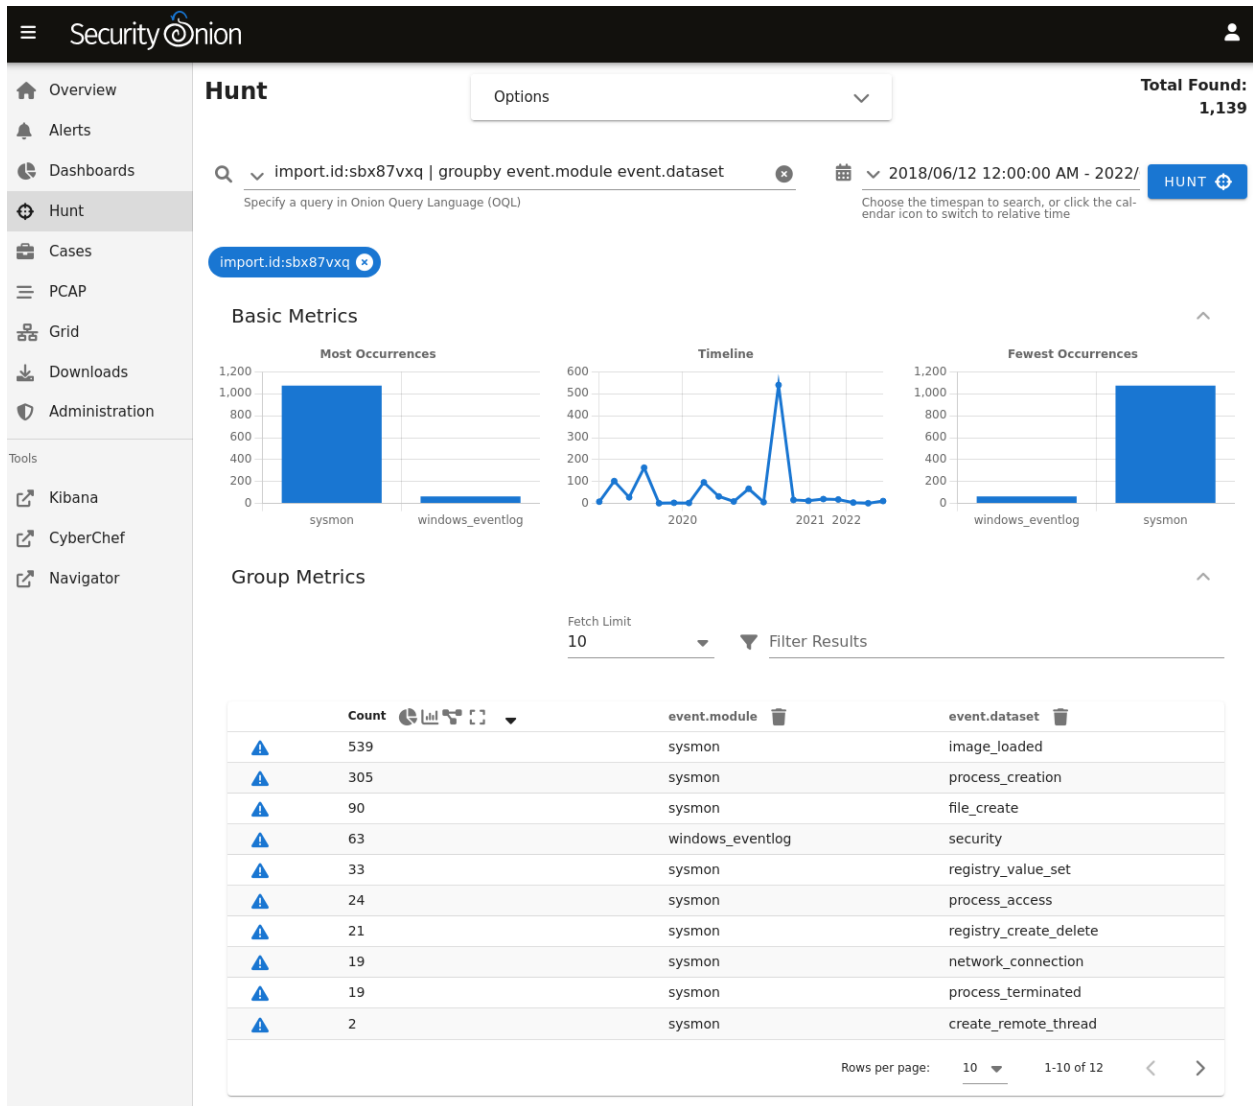
Run `sudo so-import-evtx` and supply the full path to at least one evtx file. For example, to import a single evtx file named `import.evtx`:

```
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, run `sudo so-status`:

```
Checking Docker status
```

```
Docker ----- [ OK ]
```

```
Checking container statuses
```

```
so-aptcacherng ----- [ OK ]
so-curator ----- [ OK ]
so-dockerregistry ----- [ OK ]
so-elastalert ----- [ OK ]
so-elasticsearch ----- [ OK ]
so-filebeat ----- [ OK ]
so-fleet ----- [ OK ]
so-grafana ----- [ OK ]
so-idstools ----- [ OK ]
so-influxdb ----- [ OK ]
so-kibana ----- [ OK ]
so-kratos ----- [ OK ]
so-logstash ----- [ OK ]
so-mysql ----- [ OK ]
so-nginx ----- [ OK ]
so-playbook ----- [ OK ]
so-redis ----- [ OK ]
so-sensoroni ----- [ OK ]
so-soc ----- [ OK ]
so-soctopus ----- [ OK ]
so-steno ----- [ OK ]
so-strelka-backend ----- [ OK ]
so-strelka-coordinator ----- [ OK ]
so-strelka-filestream ----- [ OK ]
so-strelka-frontend ----- [ OK ]
so-strelka-gatekeeper ----- [ OK ]
so-strelka-manager ----- [ OK ]
so-suricata ----- [ OK ]
so-telegraf ----- [ OK ]
so-wazuh ----- [ OK ]
```

`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:

```
sudo so-status -h

/usr/sbin/so-status  [-h] [-q|--quiet]

-h                Show this message.
-q|--quiet        Suppress the output and only return a
                  single status code for overall status
0:Ok, 1:Error, 2:Starting/Pending, 99:Installing SO

sudo so-status -q
echo $?
0
```

## 17.8.2 Import Node

If you're running a Security Onion Import node, then so-status will show so-steno, so-suricata, and so-zeek as DISABLED since they are not sniffing live traffic. *Suricata* and *Zeek* will still analyze pcaps normally when running *so-import-pcap*. *Stenographer* is not used at all in Import mode.

## 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'
  Started: 15:55:48.390107
```

(continues on next page)

(continued from previous page)

```

Duration: 1460.452 ms
Changes:
-----
container_id:
-----
added:
      f035103cd8bf43134b56d4b19d77a0ae9e7c09fcb54ef6da67cf89bef5fc4019
state:
-----
new:
    running
old:
    None

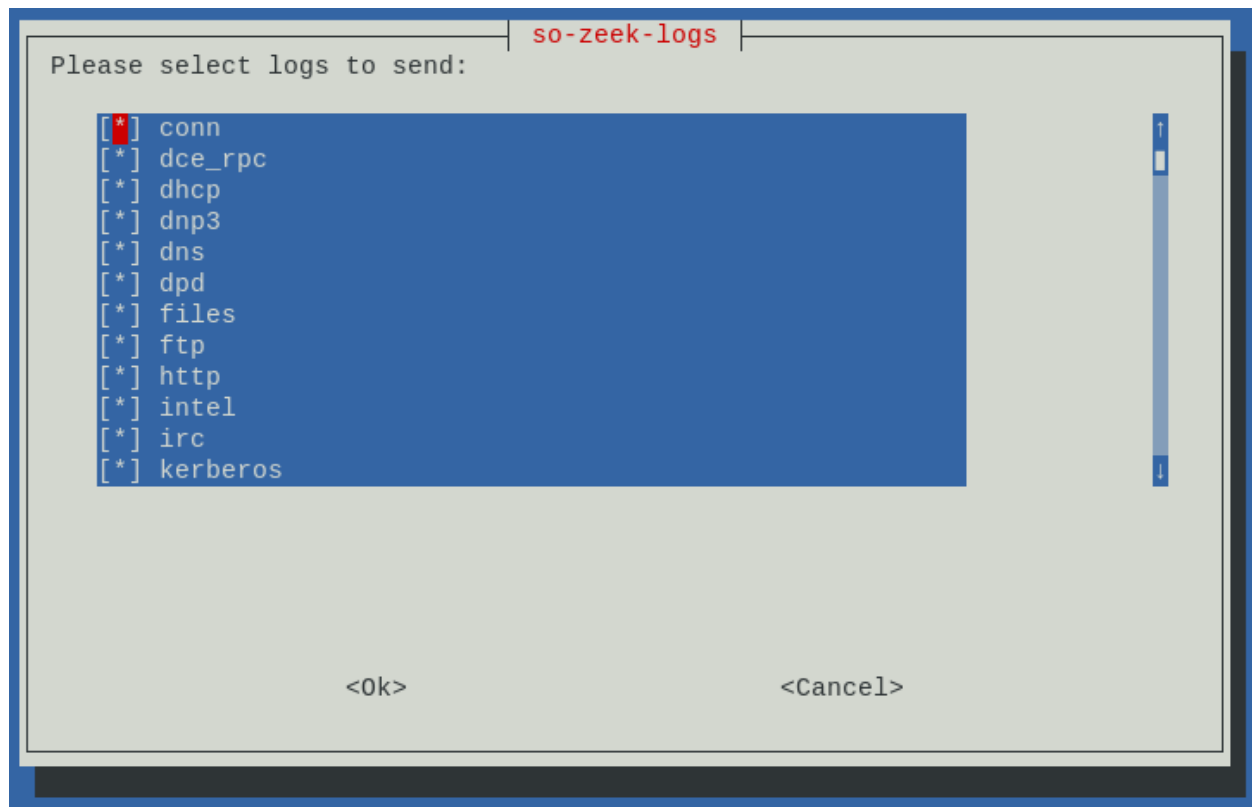
Summary for local
-----
Succeeded: 1 (changed=1)
Failed:    0
-----
Total states run:      1
Total run time: 1.460 s
Replaying PCAP(s) at 10 Mbps on interface bond0...
Actual: 111557 packets (12981286 bytes) sent in 10.38 seconds
Rated: 1249997.6 Bps, 9.99 Mbps, 10742.07 pps
Flows: 4102 flows, 394.99 fps, 2074477 flow packets, 45106 non-flow
Statistics for network device: bond0
  Successful packets:      55304
  Failed packets:         444
  Truncated packets:      0
  Retried packets (ENOBUFFS): 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-zeek-logs

If you want to specify what [Zeek](#) logs are ingested, you can use `so-zeek-logs`. It will show you a list of all [Zeek](#) logs and you can specify which of those logs are ingested. Once you've made your selection, it will modify the [Filebeat](#) configuration for you.





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/sosetup.log`
  - Suricata `/opt/so/log/suricata/suricata.log`
  - Zeek `/nsm/zeek/logs/current/`
  - Elasticsearch `/opt/so/log/elasticsearch/<hostname>.log`
  - Kibana `/opt/so/log/kibana/kibana.log`
  - Logstash `/opt/so/log/logstash/logstash.log`
  - Elastalert `/opt/so/log/elastalert/elastalert_stderr.log`
- Are you able to duplicate the problem on a fresh Security Onion installation?
- Check the [Known Issues](#) to see if this is a known issue that we are working on.
- If all else fails, please feel free to reach out for [Support](#).

## 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 Configuration* section.

#### Can I run Security Onion on Raspberry Pi or some other non-x86 box?

No, we only support 64-bit Intel/AMD 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 *Wazuh* agents and *osquery* agents. Elastic *Beats* agents support 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 Configuration* 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?

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), backups would be prohibitively expensive. Most organizations don’t do any 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/.

### 18.2.9 /nsm/wazuh

All *Wazuh* files are stored in /nsm/wazuh/. For convenience, we have placed symlinks for *Wazuh* config at /opt/so/conf/wazuh/ (linked to /nsm/wazuh/etc) and *Wazuh* rules at /opt/so/rules/hids/ (local\_rules.xml links to /nsm/wazuh/etc/rules/local\_rules.xml and ruleset links to /nsm/wazuh/ruleset).

## 18.3 Tools

Security Onion would like to thank the following projects for their contribution to our community!

(listed alphabetically)

- *ATT&CK Navigator*
- *Curator*
- *CyberChef*
- *Docker*
- *ElastAlert*
- *Elasticsearch*
- *Filebeat*
- *FleetDM*

- *Grafana*
- *Kibana*
- *Logstash*
- *osquery*
- *Redis*
- *Salt*
- *Stenographer*
- *Strelka*
- *Suricata*
- *Wazuh*
- *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
- 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](mailto: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.

## APPENDIX

This section covers upgrade scenarios.

### 20.1 Appendix A: Security Onion 16.04

This appendix covers the process of upgrading from the old Security Onion 16.04 to the new Security Onion 2.

**Warning:** Security Onion 2 is a MAJOR architectural change, so please note the following:

- Security Onion 2 has higher hardware requirements, so you should check that your hardware meets those requirements.
- Once you've upgraded from Ubuntu 16.04 to Ubuntu 18.04, you will essentially do a new installation of Security Onion 2 on top of Ubuntu 18.04. Very little data will be retained during the upgrade!
- There will be no way to migrate application accounts from Security Onion 16.04 to Security Onion 2.
- There will be no way to migrate sguild data from Security Onion 16.04 to Security Onion 2.
- You may need to purge pcap to make free space for the upgrade process. Any pcap remaining after the upgrade can only be accessed via tcpdump.
- 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.
- Ubuntu 18.04 reaches the end of Ubuntu standard support in April 2023 (<https://wiki.ubuntu.com/Releases>). So if you upgrade from Ubuntu 16.04 to 18.04 then you will need to immediately upgrade from 18.04 to 20.04 as shown in *Appendix B: Ubuntu 18.04*.

For the reasons listed above, we recommend that most users procure new hardware and perform a fresh installation of Security Onion 2.

---

**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!

---

If you have reviewed all of the warnings above and still want to attempt an in-place upgrade, you should be able to do the following.

**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.

First, make sure that Security Onion 16.04 is fully up-to-date:

```
sudo soup
```

Reboot:

```
sudo reboot
```

Copy and paste the following into a terminal to prepare for the upgrade process:

```
sudo rm /etc/apt/sources.list.d/securityonion-ubuntu-stable-xenial.list && \
sudo so-stop && \
sudo service syslog-ng stop && \
sudo service mysql stop && \
sudo service salt-minion stop ; \
sudo docker system prune -a -f && \
sudo sed -i 's|PREV="pre-.*$|PREV="pre-upgrade-to-18.04"|g' /var/lib/dpkg/info/
↪securityonion-bro.preinst && \
sudo /var/lib/dpkg/info/securityonion-bro.preinst install && \
sudo apt install update-manager-core -y && \
sudo sed -i 's|Prompt=never|Prompt=1ts|g' /etc/update-manager/release-upgrades && \
sudo pkill xscreensaver
```

Initiate the upgrade from Ubuntu 16.04 to Ubuntu 18.04:

```
sudo do-release-upgrade
```

You may be interactively prompted to provide an answer to the following questions or similar during the upgrade:

```
Non-superusers capture PCAP -> No
login.defs -> Install package maintainer's version
grub -> Choose to keep local version
sshd_config -> Choose to keep local version
syslog-ng.conf -> Choose to keep local version
```

At the end of the Ubuntu 18.04 upgrade process, you will be prompted to reboot. Do NOT reboot yet, as you will most likely need to re-install openssh-server:

```
sudo apt install openssh-server
```

Reboot:

```
sudo reboot
```

After rebooting, copy and paste the following:

```
sudo service apache2 stop && \
sudo systemctl disable apache2.service && \
sudo service mysql stop && \
sudo systemctl disable mysql.service && \
```

(continues on next page)

(continued from previous page)

```

sudo ntpdate -u time.nist.gov && \
sudo apt autoremove -f -y && \
for i in $(dpkg -l | grep securityonion | awk '{print $2}'); do sudo apt remove $i -y -f \
--purge; done && \
sudo mv /etc/salt/ /etc/salt_pre_upgrade && \
sudo mv /var/ossec /var/ossec_pre_upgrade && \
sudo apt purge salt-* -y && \
sudo apt install netplan.io -y && \
sudo apt purge -y ifupdown && \
sudo rm /etc/network/interfaces* && \
sudo mv /nsm/zeek/spool/ /nsm/zeek/old_spool && \
sudo mv /nsm/zeek/logs/stats/ /nsm/zeek/logs/old_stats && \
sudo sed -i 's/^*/#/' /etc/cron.d/salt-update

```

If you are upgrading a distributed deployment, do the following on the manager:

```

sudo systemctl stop redis.service && \
sudo systemctl disable redis.service && \
sudo apt purge redis -y

```

Remove all left-over unneeded packages:

```

sudo apt autoremove -y

```

Apply netplan for the management interface in `/etc/netplan/netplan.yaml` (create the file and ensure that the extension is `.yaml`). In the following examples, make sure to replace `ens18` with your actual management interface and replace all IP address information with your actual addresses.

If using DHCP (NOT recommended):

```

network:
  version: 2
  renderer: networkd
  ethernets:
    ens18:
      dhcp4: true

```

If using static IP:

```

network:
  version: 2
  renderer: networkd
  ethernets:
    ens18:
      addresses:
        - 10.10.10.2/24
      gateway4: 10.10.10.1
      nameservers:
        search: [mydomain]
        addresses: [10.10.10.1, 1.1.1.1]

```

For more netplan examples, please see: <https://netplan.io/examples/>

Apply the netplan configuration (may disconnect after this command, so ensure local access is available):

```
sudo netplan apply
```

Reboot:

```
sudo reboot
```

Delete “Wired connection 1” for later use as bond interface:

```
sudo nmcli con delete "Wired connection 1"
```

**Warning:** Don’t reboot yet!

Remove an old Docker configuration option:

```
sudo rm /etc/profile.d/securityonion-docker.sh
```

You may need to enable rsyslog:

```
sudo systemctl enable rsyslog
```

Download the Security Onion 2 repo:

```
git clone https://github.com/Security-Onion-Solutions/securityonion
cd securityonion
sudo bash so-setup-network
```

Follow the steps in the *Configuration* section.

Post-Installation:

While the files will still reside on disk, config files and settings will NOT be migrated to the appropriate format/locations for Security Onion 2.

Example configuration may include:

- IDS Rule Oinkcode/Thresholds/Disablements (/etc/nsm/rules/threshold.conf, /etc/nsm/pulledpork)
- Custom Logstash config (/etc/logstash/custom)
- Custom Zeek scripts or BPFs (/opt/zeek/share/zeek/policy, /etc/nsm/rules/bpf.conf)

Now that you’ve upgraded from Ubuntu 16.04 to 18.04, you should immediately upgrade from 18.04 to 20.04 as shown in *Appendix B: Ubuntu 18.04*.

## 20.2 Appendix B: Ubuntu 18.04

This appendix covers the process of upgrading a Security Onion 2.3 deployment from Ubuntu 18.04.

**Warning:**

- Ubuntu 18.04 reaches the end of Ubuntu standard support in April 2023 (<https://wiki.ubuntu.com/Releases>).



- We do not provide any guarantees that the Ubuntu in-place upgrade process will work! If the upgrade fails, be prepared to perform a fresh installation.
- 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 fresh installation.
- Plan on at least one hour per node to complete the upgrade process.
- If you have a distributed deployment with nodes configured to pull updates from the manager, you may need to comment out the proxy config in `/etc/apt/apt.conf.d/00Proxy` and make sure they can reach the Internet directly.
- Please keep in mind that, in Security Onion 2.4, we will begin phasing out support for Ubuntu (<https://blog.securityonion.net/2023/02/security-onion-in-2022-and-2023.html>).

If you have reviewed all of the warnings above and still want to attempt an in-place upgrade from Ubuntu 18.04 to Ubuntu 20.04, you should be able to follow the steps below.

**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, verify that the system is currently on Ubuntu 18.04:

```
grep VERSION_ID /etc/os-release
```

Next, make sure that all Ubuntu packages are fully up-to-date and then reboot:

```
sudo apt update && sudo apt upgrade -y && sudo reboot
```

After rebooting, copy and paste the following multi-line command to initiate the upgrade process:

```
sudo apt-mark unhold salt-* && \
sudo apt-mark unhold docker* && \
sudo sed -i 's|^deb \[arch=amd64\] https://download.docker.com/# deb \[arch=amd64\] https://\
↳download.docker.com|g' /etc/apt/sources.list && \
sudo sed -i 's|^deb|# deb|' /etc/apt/sources.list.d/wazuh.list && \
sudo sed -i 's|^deb|# deb|' /etc/apt/sources.list.d/saltstack.list && \
sudo do-release-upgrade
```

When prompted about configuration options, accept the default option. When the upgrade is complete, it will prompt you to reboot.

After rebooting, copy and paste the following multi-line command:

```
sudo sed -i 's|# deb \[arch=amd64\] https://download.docker.com/linux/ubuntu bionic|deb_\
↳[arch=amd64] https://download.docker.com/linux/ubuntu focal|g' /etc/apt/sources.list && \
↳ \
sudo sed -i 's|bionic|focal|' /etc/apt/sources.list && \
sudo sed -i 's|# deb|deb|' /etc/apt/sources.list.d/wazuh.list && \
sudo sed -i 's|# deb|deb|' /etc/apt/sources.list.d/saltstack.list && \
sudo sed -i 's|bionic|focal|' /etc/apt/sources.list.d/saltstack.list && \
sudo sed -i 's|18|20|' /etc/apt/sources.list.d/saltstack.list && \
sudo apt update && \
sudo apt upgrade -y && \
sudo apt install salt-common salt-master salt-minion -y
```

If this node has any sniffing interfaces, then you may need to update NetworkManager:

```
sudo touch /etc/NetworkManager/conf.d/10-globally-managed-devices.conf && \  
sudo systemctl restart NetworkManager.service
```

If this node is a manager, then you may need to remove some influxdb files:

```
sudo rm /opt/so/state/influxdb*
```

Run all salt states and review the output for any errors:

```
sudo salt-call state.highstate -linfo queue=True
```

Finally, make sure all services started properly:

```
sudo so-status
```

## RELEASE NOTES

### 21.1 2.3.300 [20240401] Changes

- FIX: Ignore YARA rules [#12678](#)
- UPGRADE: Strelka - 0.24.01.18 [#12676](#)
- UPGRADE: Suricata 6.0.17 [#12613](#)

### 21.2 2.3.290 [20240229] Changes

- UPGRADE: Suricata 6.0.16 [#12328](#)
- UPGRADE: Zeek 6.0.3 [#12236](#)

### 21.3 2.3.280 [20231128] Changes

- FEATURE: Add 2.3 EOL reminder icon to SOC UI
- FEATURE: Add note to soup about 2.3 EOL and 2.4 migration
- FIX: 2.3 Field Mapping for CodeIntegrity windows event logs [#11807](#)
- FIX: Update SSL cert to avoid Google Chrome error (2.3) [#11825](#)
- UPGRADE: 2.3 Grafana to Latest
- UPGRADE: 2.3 Strelka backend to 0.23.09.12
- UPGRADE: 2.3 Suricata to 6.0.15
- UPGRADE: 2.3 Update Zeek to 6.0
- UPGRADE: 2.3 Upgrade Elastic to 8.10.4
- UPGRADE: SOC dependency Axios to 1.6.1 [#11764](#)

## 21.4 2.3.270 [20231006] Changes

- FIX: Prevent repeated creation of unused Docker volumes #11354
- FIX: Support filtering against multiple groupby charts #10856
- UPGRADE: CyberChef 10.5.2 #10781
- UPGRADE: Elastic 8.8.2 #10913
- UPGRADE: Grafana 9.2.20 #10666
- UPGRADE: Suricata 6.0.14 #11318
- UPGRADE: Zeek 5.0.10 #11300

## 21.5 2.3.260 [20230620] Changes

- FIX: Update Analyzer Dependencies #10603
- UPGRADE: CyberChef 10.4.0 #10581
- UPGRADE: Grafana 9.2.19 #10583
- UPGRADE: Suricata 6.0.13 #10594

## 21.6 2.3.250 [20230519] Changes

- FIX: Bump SOCTopus Flask Version to 2.3.2 #10272
- FIX: Improve soup's local file modification logic #8972
- FIX: Kibana: Ensure \_id fields beginning with a hyphen work properly when pivoting to SOC from Kibana #10305
- FIX: Simplify cloud detection #10261
- FIX: Strelka YARA Compilation #10271
- UPGRADE: Elastic 8.7.1 #10269
- UPGRADE: FleetDM 4.31.1 #10379
- UPGRADE: Grafana 9.2.17 #10262
- UPGRADE: Kratos to 0.13.0 #10309
- UPGRADE: SOC external dependencies #10268
- UPGRADE: Suricata 6.0.12 #10311
- UPGRADE: Zeek 5.0.9 #10374

## 21.7 2.3.240 [20230426] Changes

- FIX: AWS Instances with forced IMDSv2 enabled fail to detect running in AWS #10205
- FIX: Case attachments should download to disk rather than open inline #10179

## 21.8 2.3.230 [20230417] Changes

- FIX: Suricata DNS A and CNAME parsing #10117
- UPGRADE: CyberChef 10.2.0 #9992
- UPGRADE: Grafana 9.2.15 #9889
- UPGRADE: Redis 6.2.11 #9980
- UPGRADE: Suricata 6.0.11 #10067
- UPGRADE: Zeek 5.0.8 #10107
- UPGRADE: nginx 1.22.1 #9742

## 21.9 2.3.220 Hotfix [20230301] Changes

- FIX: Curator configuration to align with requirements for Curator 8.0.x #9871

## 21.10 2.3.220 [20230224] Changes

- FIX: Trim rule.name to fit the so\_case.title field #9781
- UPGRADE: Elastic 8.6.2 #9804
- UPGRADE: FleetDM 4.27.1 #9750
- UPGRADE: Golang to 1.20 #9839
- UPGRADE: Grafana 9.2.10 #9808
- UPGRADE: Zeek 5.0.7 #9831

## 21.11 2.3.210 [20230202] Changes

- FIX: Add configured\_vulns\_ext\_vars.yar to Strelka YARA exclusion list #9701
- FIX: Improve Suricata DHCP parsing and dashboard #9678
- UPGRADE: CyberChef 9.55.0 #9673
- UPGRADE: Elastic 8.6.1 #9594
- UPGRADE: MySQL 5.7.24 to 5.7.41 #9616
- UPGRADE: Strelka 0.23.01.07 #9687
- UPGRADE: Suricata 6.0.10 #9403

- UPGRADE: Zeek 5.0.6 #9695

## 21.12 2.3.200 [20230113] Changes

- FEATURE: Add Consumption EPS to default Grid view #9359
- FEATURE: Allow custom IDH HTTP Skins #9045
- FEATURE: Improve SOC Dashboards #9450
- FEATURE: Update Zeek oui.dat during container build #9542
- FIX: EPS sorting in SOC Grid #9363
- FIX: Ensure file path is ascertainable by localfile.py for localfile analyzer #9342
- FIX: RITA number\_format\_exception “unable to convert to integer” #9503
- FIX: Superuser permission check for so-elasticsearch-indices-list #9461
- FIX: Sysmon logs are missing event.category and event.dataset #8194
- FIX: config.map.jinja updates for highlander setting #9321
- FIX: so-import utilities should hyperlink to dashboards #9373
- FIX: so-status runs some code before checking for root privileges #9270
- UPGRADE: Elastic 8.5.3 #9356

## 21.13 2.3.190 Hotfix [20221207] Changes

- FIX: Improve support for Suricata file extraction into Strelka #9318

## 21.14 2.3.190 [20221205] Changes

- FEATURE: Add Additional ICS Zeek Packages #9149
- FEATURE: Check hashes for duplicates before sending to strelka #9034
- FEATURE: Improve local copy of docs in SOC #9097
- FEATURE: so-pcap-export can run without needing to be attached to a TTY #8994
- FIX: Avoid deprecation warning in Zeek file extraction script #9123
- FIX: Change PyYAML .whl file name to comply with Joliet’s 240-character limit/threshold #8995
- FIX: Remove JA3er Analyzer #8984
- FIX: prevent ISO image from triggering common AV false positives #7297
- UPGRADE: CyberChef 9.49.0 #9180
- UPGRADE: Grafana 9.2.7 #8950
- UPGRADE: Strelka #8996
- UPGRADE: Suricata 6.0.9 #9229
- UPGRADE: Zeek 5.0.4 #9167

## 21.15 2.3.182 [20221109] Changes

- UPGRADE: Zeek 5.0.3 #9100

## 21.16 2.3.181 [20221021] Changes

- FIX: Add find to the zeek docker container #8922
- FIX: Implement workaround for case events query #8965

## 21.17 2.3.180 [20221014] Changes

- FEATURE: Add Additional Zeek Plugins to the Zeek Docker Container #8877
- FEATURE: Add new Sysmon dashboards #8870
- UPGRADE: Elastic 8.4.3 #8861
- UPGRADE: Suricata 6.0.8 #8820
- UPGRADE: Zeek 5.0.2 #8881

## 21.18 2.3.170 [20220922] Changes

- FEATURE: Events table(s) for Windows Events matching default view #8591
- FEATURE: Split the winlog.event\_data.Hashes field for Windows sysmon process creation events. #8593
- FIX: Mapping error when trying to index Strelka logs generated from ELF files. #8592
- UPGRADE: Elastic 8.4.1 #8794
- UPGRADE: Zeek 4.0.9 #8774

## 21.19 2.3.160 [20220829] Changes

- FEATURE: Add Advanced toggle to Alerts and Cases pages #7559
- FIX: Add gen\_webshells.yar to Strelka ignore list #8624
- UPGRADE: Zeek 4.0.8 #8610

## 21.20 2.3.150 [20220820] Changes

- FIX: Allow Filebeat to be disabled for sensors, idh, and fleet nodes #8404
- FIX: Display PCAP menu action on Dashboards page #8343
- FIX: Elasticsearch geoip lookups fail for some users in Elastic 8 #8373
- FIX: Fix TLP options in Cases to align with TLP 2.0 #8469
- FIX: Remove Elastic Pipeline Time graph from Grafana #8369
- FIX: Update TALOS rules snapshot version to 29200 #8551
- FIX: Use systemd drop-in file for salt-minion.service override #8441
- FIX: soup should check for local configurations before modifying anything #8423
- FIX: soup should only delete elastalert indices when upgrading from Elastic versions older than 8 #8536
- UPGRADE: CyberChef 9.46.0 #8299
- UPGRADE: Elastic 8.3.3 #8398

## 21.21 2.3.140 Hotfix [20220812] Changes

- FIX: so-curator-closed-delete-delete needs to reference new Elasticsearch directory #8529

## 21.22 2.3.140 Hotfix [20220719] Changes

- FIX: Revise Elastalert index check/deletion logic
- FIX: Ensure Elastalert is enabled before trying to run 'so-elastalert-stop'. Also suppress error output for when so-elastalert container is missing.

## 21.23 2.3.140 Changes

- FEATURE: Provide ability to maximize view of a groupby chart or table #8176
- FEATURE: Remember the state of the left menu (visible or hidden) between SOC refreshes #8186
- FEATURE: Remove disabled accounts from Case Assignee list #8184
- FEATURE: SOUP should not let you update to 2.3.140 or above unless you have updated to at least 2.3.110 #8239
- FEATURE: Support bulk observable data entry #8210
- FIX: Add ID to Filestream Inputs #8006
- FIX: Add event.category field to pfsense firewall logs #8112
- FIX: Add jinja to localfile.yaml #8196
- FIX: IDH ISO Disk partitions #8144
- FIX: Improve default dashboards #8136
- FIX: Strip whitespace after analyzer input strings (observable values) #8208



- FIX: Support group-by sorting memory on first group only #8133
- FIX: Using so-firewall to list default port groups #8264
- FIX: pam.d lastlog module breaks BPF capability #8188
- FIX: Ensure so-kibana indices can be cleaned up on search nodes #8262
- UPGRADE: Kratos to 0.10.1 #8227
- UPGRADE: Salt 3004.2 #8166
- UPGRADE: Suricata 6.0.6 #8279
- UPGRADE: Elastic 8.3.2 #7563
- UPGRADE: Redmine to 4.2.7 #8308

## 21.24 2.3.130 Changes

- FEATURE: Add “observable” button next to hash for case attachments #7222
- FEATURE: Add set of default analyzers #7945
- FEATURE: Make classification.config user-configurable #7918
- FEATURE: Native analyzer infrastructure #7944
- FEATURE: Playbook False Positive Tuning #8059
- FEATURE: SOC Dashboards #1211
- FIX: Allow quick actions on a field value with the number 0 #8023
- FIX: Elastalert query in Hunt #8049
- FIX: Ensure failed elastic queries show an error on the SOC UI #7846
- FIX: Firefox OQL edits should release focus after pressing ENTER #8063
- UPGRADE: ElastAlert 2 from 2.2.2 to 2.5.0 #8008
- UPGRADE: Elastic 7.17.4 #8002
- UPGRADE: FleetDM 4.14.0 #8012
- UPGRADE: Kratos 0.8.2-alpha.1 to 0.9.0-alpha.3 #7943
- UPGRADE: TensorFlow from 2.5 to 2.9.1 #8009
- UPGRADE: attack-navigator v4.6.4 #7977
- UPGRADE: Zeek 4.0.6 to 4.0.7 #8067

## 21.25 2.3.120 Changes

- FEATURE: Add ISO option to the installer for analyst workstation #7502
- FEATURE: Add new Hunt query for SOC logins #7327
- FEATURE: Add strelka-fileshot and strelka-oneshot binaries to analyst workstation #7670
- FEATURE: Expose Case user info (email address) in SOC Alert/Hunt/Cases results instead of user ID #7548
- FEATURE: Have Observables inherit their case's TLP by default #7642
- FEATURE: IDH - Separate MGT & IDH NIC #7589
- FEATURE: Remove thehive and cortex dockers #7501
- FEATURE: Stop hive related services in soup to 120 #7599
- FIX: Create .keyword shim for additional process fields #7633
- FIX: Elasticsearch & Logstash logs not compressed or cleaned #6932
- FIX: Failure of influxdb state if default shell is zsh #7730
- FIX: Hunt OR queries should work without parentheses #7540
- FIX: Improve Hunt query when pivoting from Cases Observables #7405
- FIX: Improve Zeek file extraction #7829
- FIX: Management IP is sometimes null at the end of setup #7113
- FIX: Prevent multiple instances of so-sensor-clean and so-playbook-sync #6622
- FIX: Prevent users from running *so-setup iso* on Ubuntu #7601
- FIX: Remove TheHive deps from Playbook #7483
- FIX: Run telegraf as non-root #7468
- FIX: Salt error during setup - [ERROR ] Unable to connect pusher: Stream is closed #7203
- FIX: Update syslog ingest pipeline per #5251 #6912
- FIX: remove soremote access after analyst install joins grid #7639
- FIX: soup should ensure salt-master service is running prior to running #7763
- FIX: surilogcompress not working correctly on some systems #7133
- UPGRADE: CyberChef 9.37.3 #7817
- UPGRADE: Elastic 7.17.3 #7807
- UPGRADE: FleetDM 4.12.1 #7725
- UPGRADE: Suricata 6.0.5 #7836
- UPGRADE: Zeek 4.0.6 #7839
- UPGRADE: nginx 1.20.2 #7808

## 21.26 2.3.110 Hotfix [20220407] Changes

- FIX: Previously failed Ubuntu minions will now be able to get the proper repo for install
- FIX: Fixed a regression in AIRGAP that was preventing salt from upgrading

## 21.27 2.3.110 Hotfix [20220407] Known Issues

- If you had a previous failed soup please ensure that the salt-master service is running before you run soup again.

## 21.28 2.3.110 Hotfix [20220405] Changes

- FIX: Change the salt bootstrap script to pull from the proper location for Ubuntu

## 21.29 2.3.110 Hotfix [20220401] Changes

- FIX: Updated Saltstack to 3004.1 to address CVE-2022-22934 #7701

## 21.30 2.3.110 Changes

- FEATURE: Full ECS data type compliance #6747
- FEATURE: Intrusion Detection Honeypot Node #7138
- FEATURE: Multi-Factor Authentication (MFA) for Security Onion #7316
- FEATURE: Populate Zeek's networks.cfg with \$HOME\_NET #6854
- FEATURE: SOC authentication logs will now be ingested into Elasticsearch #7354
- FEATURE: sort indices list alphabetically by index name #6969
- FIX: ACNG should clear the cache on restart #7114
- FIX: Abort so-user sync if Kratos database is locked #7459
- FIX: Add Endgame Index settings to the global.sls on new installs #7293
- FIX: Allow downgrades during docker\_install #7228
- FIX: Avoid telegraf apparmor issues #2560
- FIX: Composable Templates #4644
- FIX: Increase minimum password length from 6 to 8 characters #7352
- FIX: Navigator should ship with all needed files #1162
- FIX: Prevent Elasticsearch deprecation notices from causing installation failures #7353
- FIX: Random passwords generated at setup contain character combinations that cause problems with some containers #7233
- FIX: curator should exclude so-case\* indices #7270
- FIX: so-ip-update needs to update Kibana dashboards #7237

- FIX: so-status TTY improvements [#7355](#)
- UPGRADE: Elastic 7.17.1 [#7137](#)
- UPGRADE: FleetDM to 4.10.0 [#7245](#)
- UPGRADE: Grafana 8.4.1 [#7281](#)
- UPGRADE: Kratos 0.8.2-alpha.1 [#7351](#)

## 21.31 2.3.100 Hotfix [20220301] Changes

- FIX: Prevent curator from pruning case indices [#7270](#)

## 21.32 2.3.100 Hotfix [20220203] Changes

- FIX: SSLError for Logstash connecting to Redis if manager hostname contains uppercase [#7103](#)
- FIX: Add mixed case hostnames to automated testing

## 21.33 2.3.100 Hotfix [20220202] Changes

- FIX: Add new salt URL to the ACNG config for SSL passthrough
- FIX: Managers with capitals in the hostname will now properly pull from the salt mine [#7081](#)

## 21.34 2.3.100 Changes

- FEATURE: Add verbiage to soup to denote which branch is being used [#6763](#)
- FEATURE: Allow for an easy way to add a local repo directory for Elastic snapshots [#7034](#)
- FEATURE: Install Elasticsearch plugin - repository-s3 [#6139](#)
- FEATURE: Introduce new Cases module for native case management [#7019](#)
- FEATURE: Introduce new Receiver node type [#6469](#)
- FEATURE: Open event from Kibana in hunt [#6748](#)
- FEATURE: SOC error messages should show regardless of how far down the user has scrolled [#6977](#)
- FEATURE: Support sort order in Elasticsearch queries [#2577](#)
- FIX: Reinstall on Ubuntu 18.04 fails on docker install [#6467](#)
- FIX: Cleanup Invalid Kolide messages in nginx logs [#3989](#)
- FIX: Disable Wazuh on sensors if it is disabled globally [#7016](#)
- FIX: During a reinstall, remove existing certs and keys generated by the ssl and ca states [#7010](#)
- FIX: Enable SANs for all certificates [#6381](#)
- FIX: Fleet broken when default Docker IP range changed [#6603](#)
- FIX: Generate .security subfield for *message* field [#5106](#)

- FIX: Improve support for grouping by fields with spaces #6724
- FIX: Logstash inputs beats deprication #5194
- FIX: Playbook Field Mappings #3660
- FIX: Prevent the .security keyword from being added to the rule.uuid field in Playbook #6276
- FIX: Reduce excessive Elasticsearch log growth #5190
- FIX: Reinstall should not try to patch python3-influxdb modules if already patched. #6765
- FIX: Remove manager from /etc/hosts during install prompts #6492
- FIX: Remove xml header from ossec.conf #6658
- FIX: SOUP should check that en\_US.UTF-8 is available before switching to it #6599
- FIX: Salt does not generate a fleet.crt file with CUSTOM\_FLEET\_HOSTNAME #4319
- FIX: Typo in so-image-common output #6563
- FIX: Wazuh WEL Parsing #6829
- FIX: \_id fielddata deprecated message #6703
- FIX: elastic\_curl\_config depends on elastic\_curl\_config\_distributed #6811
- FIX: prevent the need for adding roles in a specific order when using so-user #6505
- FIX: so-preflight tries to run curl before it is installed #6899
- FIX: so-user update should automatically sync #6659
- UPGRADE: CyberChef 9.32.3 #6434
- UPGRADE: Elastic components to 7.16.3 #6860
- UPGRADE: FleetDM 4.8.0 #6828
- UPGRADE: Grafana 8.3.2 #6321
- UPGRADE: Salt to 3004 #6810
- UPGRADE: Zeek to 4.0.5 #6983

## 21.35 2.3.91 Changes

- UPGRADE: Elastic to 7.16.2 for log4j vulnerability mitigation

## 21.36 2.3.90 Hotfix [20211213]

- FIX: Remove JndiLookup class from Elasticsearch and Logstash jar files to address additional log4j attack vectors

## 21.37 2.3.90 Hotfix [20211210]

- FIX: Mitigate vulnerability in log4j

## 21.38 2.3.90 Hotfix [20211206]

- FIX: soup should now properly update 2.3.90 installs that had an issue with xml headers in the ossec.conf
- FIX: soup now has more logging
- FIX: soup now checks for the existence of the endgame group before trying to apply it on a re-soup
- FIX: so-elasticsearch-pipelines now uses the proper value for applying the pipelines

## 21.39 2.3.90 Hotfix [AIRGAPFIX]

- FIX: Airgap repo was created on distributed iso nodes even in non-airgap installs #6415

## 21.40 2.3.90 Hotfix [WAZUH]

- FIX: so-allow should not be modifying ossec.conf when Wazuh isn't installed #6317
- FIX: so-allow should not be writing an XML header to the ossec.conf file #6325
- FIX: Correct “exisiting” typo on whiptail prompt
- FIX: Soup will no longer attempt to validate a successful salt upgrade if salt wasn't upgraded on this soup run

## 21.41 2.3.90 Changes

- FEATURE: Add ASN annotation for GeoIP #5068
- FEATURE: Add Endgame Support for Security Onion #6166
- FEATURE: Add TI Module #5916
- FEATURE: Add additional flags to stenographer config #5851
- FEATURE: Add filebeat, auditbeat, and metricbeat downloads to SOC Download screen #5849
- FEATURE: Add logstash and redis input plugins to telegraf #5960
- FEATURE: Add so-deny script for removing access from firewall and other apps #4621
- FEATURE: Add support for escalation to Elastic Cases #6048
- FEATURE: Allow for Kibana customizations via pillar #3933
- FEATURE: Allow users to set their profile information #5846
- FEATURE: Allow vlan tagged NICs to be used as management interface #3687
- FEATURE: Create Pipeline Overview Dashboard for Grafana #6177
- FEATURE: Create script to reset elastic auth passwords #6206

- FEATURE: Enable Kibana Settings for encryption #6146
- FEATURE: Expose new user profile field for specifying a custom note about a user #5847
- FEATURE: HTTP module for SOC event escalation #5791
- FEATURE: Increase password lengths, provide a way to change existing passwords #6043
- FEATURE: Indicate that setup has completed at the very end of sosetup.log #5032
- FEATURE: Prevent SOUP from running if there is an issue with the manager pillar #5809
- FEATURE: Provide quick-select date ranges from Hunt/Alerts date range picker #5953
- FEATURE: SOC Hunt Timeline/Charts should be collapsible #5114
- FEATURE: Support Ubuntu 20.04 #601
- FEATURE: setup should run so-preflight #3497
- FIX: ACNG sometimes returns 503 errors when updating Ubuntu through the manager #6151
- FIX: Add details to Setup for Install Type menus #6105
- FIX: Adjust timeout in check\_salt\_minion\_status in so-functions #5818
- FIX: All templates should honor replica settings #6005
- FIX: Clear holds on Ubuntu installs #5588
- FIX: Consider making the airgap option only settable on the manager #5914
- FIX: Docker containers should not start unless file events are completed #5955
- FIX: Ensure soc\_users\_roles file is cleaned up if incorrectly mounted by Docker #5952
- FIX: Favor non-aggregatable data type when a cache field has multiple conflicting data types #5962
- FIX: Firefox tooltips stuck on Hunt and Alerts screens #6010
- FIX: Grafana sensor graphs only show interface graphs when selected individually #6007
- FIX: Kibana saved objects #5193
- FIX: Modify Steno packet loss calculation to show point in time packet loss #6060
- FIX: Remove CURCLOSEDAYS prompt in Setup since it is no longer used #6084
- FIX: Remove references to xenial (Ubuntu 16.04) from setup #4292
- FIX: Remove unnecessary screens from Analyst Setup #5615
- FIX: SOC docker should not start until file managed state runs #5954
- FIX: SOC unable to acknowledge alerts when not grouped by rule.name #5221
- FIX: Setup should ask if new or existing distributed deployment #6115
- FIX: Setup should prevent invalid characters in Node Description field #5937
- FIX: Support non-WEL Beats #6063
- FIX: Unnecessary Port Binding for so-steno #5981
- FIX: Use yaml.safe\_load() in so-firewall (thanks to @clairmont32) #5750
- FIX: Zeek state max depth not working #5558
- FIX: *so-ip-update* should grant mysql root user access on new IP #4811
- FIX: docker group can be given gid used by salt created groups #6071

- FIX: packetloss.sh gives an error every 10 min though ZEEK is disabled [#5759](#)
- FIX: so-import-evtx elastic creds & logging [#6065](#)
- FIX: so-user delete function causes re-migration of user roles [#5897](#)
- FIX: wazuh-register-agent times out after 15 minutes lower to 5 minutes [#5794](#)
- FIX: yum pkg.clean\_metadata occasionally fails during setup [#6113](#)
- UPGRADE: ElastAlert to 2.2.2 [#5751](#)
- UPGRADE: Elastic to 7.15.2 [#5752](#)
- UPGRADE: FleetDM to 4.5 [#6188](#)
- UPGRADE: Grafana to 8.2.3 [#5852](#)
- UPGRADE: Kratos to 0.7.6-alpha.1 [#5848](#)
- UPGRADE: Redis to 6.2.6 [#6140](#)
- UPGRADE: Suricata to 6.0.4 [#6274](#)
- UPGRADE: Telegraf to 1.20.3 [#6075](#)

## 21.42 2.3.80 Changes

- FEATURE: Ability to disable Zeek, Suricata [#4429](#)
- FEATURE: Add docs link to Setup [#5459](#)
- FEATURE: Add evtx support in Import Node [#2206](#)
- FEATURE: Consolidate whiptail screens when selecting optional components [#5456](#)
- FEATURE: Distinguish between Zeek generated syslog and normal syslog in hunt for event fields [#5403](#)
- FEATURE: Enable index sorting to increase search speed [#5287](#)
- FEATURE: Expose options for elasticsearch.yml via Salt pillar [#1257](#)
- FEATURE: Role-based access control (RBAC) [#5614](#)
- FEATURE: soup -y for automation [#5043](#)
- FIX: Add new default filebeat module indices to the global pillar. [#5526](#)
- FIX: all.rules file can become empty on non-airgap deployments if manager does not have access to the internet. [#3619](#)
- FIX: Curator cron should run less often [#5189](#)
- FIX: Improve unit test maintainability by refactoring to use Golang assertion library [#5604](#)
- FIX: Invalid password message should also mention dollar signs are not allowed [#5381](#)
- FIX: Max files for steno should use a pillar value for easy tuning. [#5393](#)
- FIX: Remove raid check for official cloud appliances [#5449](#)
- FIX: Remove watermark settings from global pillar. [#5520](#)
- FIX: SOC Username case sensitivity [#5154](#)
- FIX: so-user tool should validate password before adding user to SOC [#5606](#)
- FIX: Switch to new Curator auth params [#5273](#)



- UPGRADE: Curator to 5.8.4 #5272
- UPGRADE: CyberChef to 9.32.2 #5158
- UPGRADE: SOC UI 3rd Party dependencies to latest versions #5603
- UPGRADE: Zeek to 4.0.4 #5630

## 21.43 2.3.70 Hotfix [WAZUH]

- FIX: wazuh-agent is updated during setup on ISO, which causes service to fail to start #5354

## 21.44 2.3.70 Hotfix [GRAFANA\_DASH\_ALLOW]

- FIX: Grafana state trying to create undefined dashboards #5270

## 21.45 2.3.70 Hotfix [CURATOR]

- FIX: Rolled back curator change for true clustering deployments (will be fixed in next release) #5226
- FIX: Resolved benign error repeatedly logged to telegraf log file #5195

## 21.46 2.3.70 Changes

- FEATURE: Add sha.256 to suricata.fileinfo pipeline #4224
- FEATURE: Allow for adjustment of Kibana sampleSize setting in Discover dashboard #4969
- FEATURE: Allow for adjustment to automatic patch schedule #4985
- FEATURE: Require SOC login before allowing users to access playbook and soctopus #4623
- FEATURE: Scan kratos logs for anomalous login attempts #4710
- FEATURE: Send PCAP session transcript to CyberChef #5010
- FEATURE: Show model numbers of cloud-deployed nodes #4898
- FEATURE: Show warning when a user attempts to use a hostname or web domain entry that is not all lowercase #4791
- FEATURE: Simplify Grafana dashboard management and redesign dashboards #4674
- FEATURE: so-firewall needs an option to run apply by itself #4765
- FEATURE: so-pcap-export #4210
- FEATURE: SOUP - Prompt user when local modifications are detected #3860
- FIX: Add mapping to extracted file directory #4622
- FIX: Clarify missing appliance images message on SOC grid #5118
- FIX: Curator should only run on manager when set to use true clustering. #2806
- FIX: Disabled user still shows as active in GUI #5055

- FIX: Disallow blank passwords during ISO first stage setup (kickstart) #4947
- FIX: Disallow ctrl-c during the first stage of ISO setup #4948
- FIX: Improve raid failure detection on SOS Appliances #5064
- FIX: Improve verbiage for initial IPv4 prompt and so-allow prompt #5138
- FIX: Jinja the stream.reassembly.depth value in the Suricata defaults.yaml file #4293
- FIX: Remove so-elastic-features. #4542
- FIX: SOC login page missing the hide/show password icons #5087
- FIX: Wazuh data ingest error: data.port #3988

## 21.47 2.3.61 Hotfix [STENO, MSEARCH]

- FIX: Some browsers refuse to load SOC UI due to CSP blocking wss: protocol #4938
- FIX: Disabling steno raises errors when applying state.highstate / running soup update #4922
- FIX: Manager Search does not come up properly with true clustering enabled #4971

## 21.48 2.3.61 Changes

- FIX: Airgap link to Release Notes #4685
- FIX: CyberChef unable to load due to recent Content Security Policy restrictions #4885
- FIX: Suricata dns.response.code needs to be renamed to dns.response.code\_name #4770
- UPGRADE: alpine 3.12.1 to latest for Fleet image #4823
- UPGRADE: Elastic 7.13.4 #4730
- UPGRADE: Zeek 4.0.3 #4716

## 21.49 2.3.60 Hotfix [ECSFIX, HEAVYNODE, FBPIPELINE, CURATORAUTH] Changes

- FIX: Curator's authentication to Elasticsearch was incorrectly configured for the version currently in use.
- FIX: Some logs from Filebeat were not being properly routed to the correct pipeline causing the log to fill up the disk.
- FEATURE: All hotfixes going forward will have an ISO so that airgap users can follow the standard soup process as they would for normal releases.
- FIX: Hotfix to revert Strelka and Wazuh Elastic Common Schema (ECS) changes that weren't intended for 2.3.60.
- FIX: Correct SSL certificate common name (CN) to match heavy node hostnames. Only applicable to grids with heavy nodes. May require manual restart of Redis, Elasticsearch, Filebeat, and Logstash containers (in that order), once the heavy nodes have succeeded in applying highstate. For more information see the related blog post at <https://blog.securityonion.net/2021/07/security-onion-2360-heavy-node-hotfix.html>

## 21.50 2.3.60 Changes

- FEATURE: Ability to change default SOC timezone instead of using browser's timezone #4261
- FEATURE: Add SOC database to the backups #3748
- FEATURE: Add so-elasticsearch-query tool #4437
- FEATURE: Create a new Quick Drilldown option in SOC #4469
- FEATURE: Display Security Onion version number in so-setup #3348
- FEATURE: Elastic Auth #1423
- FEATURE: Implement retention policy for InfluxDB #3264
- FEATURE: New Grafana dashboards for InfluxDB RPs #4609
- FEATURE: Pillarize Filebeat Modules #3859
- FEATURE: Pivot from Alerts/Hunt to CyberChef #4081
- FEATURE: Pivot from SOC PCAP to CyberChef #1596
- FEATURE: Support adjustable SOC session timeout #4586
- FIX: Add a prompt when soup requires the path or cdrom device to be input #3551
- FIX: Add event\_data to Elasticsearch template(s) #4012
- FIX: Allow for spaces in password on kickstart script (ISO) #1079
- FIX: Change Acknowledge, Escalate, and expandEvent buttons from title to tooltip #4497
- FIX: Disallow so-suricata-start from running on the manager node #2977
- FIX: Ensure fixed PCAP files are readable by Suricata during so-import-pcap execution #4636
- FIX: Fail curl requests if the remote server responds with a failing status code #4266
- FIX: Implement error handling for soup #3220
- FIX: Improve PCAP job lookup performance by providing a tighter time range #4320
- FIX: Improve administrative username password prompt to prevent backspacing into text (ISO) #3099
- FIX: Improve soup for older installs #4617
- FIX: Include secure HTTP headers in nginx responses #4267
- FIX: Increase default search and proxy timeouts to 5 minutes #4321
- FIX: OS passwords including special characters like \$ and ! #4249
- FIX: Prevent highstate failure during soup #3559
- FIX: Prevent so-thehive-cortex from continuing to build if an issue is encountered installing Python packages #4032
- FIX: Setup should not prompt for node description when running import or eval #4004
- FIX: Trying to delete old pcap job results in error #4528
- FIX: Websocket session cleanup overly aggressive #4598
- FIX: so-user should support spaces in passwords for Fleet and TheHive users #4460
- FIX: zeek leaving post-terminate crash logs on every shutdown #4461
- UPGRADE: Elastic to 7.13 #4313

- UPGRADE: Kratos to 0.6.3-alpha.1 [#4282](#)
- UPGRADE: Redmine 4.2 (For Playbook) [#4159](#)
- UPGRADE: Suricata 6.0.3 [#4661](#)

## 21.51 2.3.52 Changes

- FIX: packetloss.sh can cause Zeek to segfault [#4398](#)
- FIX: soup now generates repo tarball with correct folder structure [#4368](#)
- UPGRADE: Zeek 4.0.2 [#4395](#)

## 21.52 2.3.51 Changes

- FIX: Mixed case sensor hostnames lead to incomplete PCAP jobs [#4220](#)
- FIX: Reconcile InfluxDB/Grafana containers in certain setup modes [#4207](#)
- FIX: Turn down log level for Salt States and Zeek [#4231](#)
- FIX: Correct downloaded PCAP filename [#4234](#)
- FIX: Truncate /root/wait\_for\_web\_response.log before each wait invocation [#4247](#)

## 21.53 2.3.50 Changes

- FEATURE: Add EPS Stats for Filebeat [#3872](#)
- FEATURE: Add copy-to-clipboard quick action menu option for copying a single field and value as 'field:value' [#3937](#)
- FEATURE: Add raid and so-status monitoring to SOC grid page [#3584](#)
- FEATURE: Add so-status to telegraf script executions and return a value [#3582](#)
- FEATURE: Add zeekctl wrapper script [#3441](#)
- FEATURE: Allow users to set an optional description for the node during setup [#2404](#)
- FEATURE: Initial implementation of enhanced websocket management [#3691](#)
- FEATURE: Combine proxy + package update questions into one menu [#3807](#)
- FEATURE: Configure NTP in Setup [#3053](#)
- FEATURE: Logstash pipeline stats wrapper [#3531](#)
- FEATURE: Need a way to have Hunt/Alerts perform groupbys that can optionally include event's that don't have a match for a group [#2347](#)
- FEATURE: Osquery WEL - Differentiate between Event & Ingest Timestamp [#3858](#)
- FEATURE: Provide customizable Login page banner content using markdown format [#3659](#)
- FEATURE: Provide customizable Overview tab content using markdown format [#3601](#)
- FEATURE: Redirect expired login form back to login page instead of showing error [#3690](#)

- FEATURE: Redirect to login when session expires [#3222](#)
- FEATURE: Show final selected options menu at the end of install [#3197](#)
- FEATURE: Show node and overall grid EPS on Grid Page [#3823](#)
- FEATURE: Telegraf should check for additional metrics if it is running on an appliance [#2716](#)
- FEATURE: VIM YAML Syntax Highlighting [#3966](#)
- FEATURE: allow for salt-minion start to be delayed on system start [#3543](#)
- FEATURE: check manager services (salt-master, so-status) during setup on a node [#1978](#)
- FEATURE: soup should check for OS updates [#3489](#)
- FIX: Alerts Total Found value should update when acknowledging or escalating [#2494](#)
- FIX: Alerts severity sort order [#1741](#)
- FIX: Change bro packet loss to be once per 2 minutes vs 30s [#3583](#)
- FIX: Check Zeek index close and delete settings for existing deployments [#3575](#)
- FIX: Firewall rules added via pillar only applies last hostgroup of the defined chain [#3709](#)
- FIX: Hunt not properly escaping special characters in Windows sysmon logs. [#3648](#)
- FIX: Hunt query for HTTP EXE downloads should work for both Zeek and Suricata [#3753](#)
- FIX: Incorrect retry syntax in CA and SSL states [#3948](#)
- FIX: Playbook Alert/Hunt showing incorrect timestamp [#2071](#)
- FIX: Properly handle unauthorized responses during API requests from SOC app [#2908](#)
- FIX: Reformat date/time on Grid and PCAP pages to enable sorting [#2686](#)
- FIX: Rename Fleet link in SOC to FleetDM [#3569](#)
- FIX: Suricata compress script should send it's output to /dev/null [#3917](#)
- FIX: Suricata cpu-affinity not being set if suriprocs is defined in minion pillar file. [#3926](#)
- FIX: TheHive Case Creation from Kibana Failure [#3870](#)
- FIX: WEL Shipping via Wazuh broken [#3857](#)
- FIX: Zeek Intel not working [#3850](#)
- FIX: ingest.timestamp should be date type [#3629](#)
- FIX: nmcli error during setup on Ubuntu + AMI [#3598](#)
- FIX: salt upgrade failure with versionlock [#3501](#)
- FIX: setup tries to connect to url used for proxy test even if the user chooses not to set one up [#3784](#)
- FIX: so-playbook-sync should only have one instance running [#3568](#)
- FIX: so-ssh-harden needs improvement [#3600](#)
- FIX: soup does not update /etc/soversion on distributed nodes [#3602](#)
- UPGRADE: Elastalert to 0.2.4-alt3 [#3947](#)
- UPGRADE: Salt 3003 [#3854](#)
- UPGRADE: Upgrade Grafana to 7.5.4 [#3916](#)
- UPGRADE: Upgrade external dependencies used by SOC [#3545](#)

## 21.54 2.3.50 Known Issues

- If you had previously enabled Elastic Features and then upgrade to Security Onion 2.3.50 or higher, you may notice some features missing in Kibana. You can enable or disable features as necessary 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>.
- If you have node names in mixed case (rather than all lower case), the Grid page may show the nodes as being in the Fault state. This is a cosmetic issue and has been resolved with a hotfix: <https://blog.securityonion.net/2021/05/security-onion-2350-hotfix-available.html>

## 21.55 2.3.40 Changes

- FEATURE: Add option for HTTP Method Specification/POST to Hunt/Alerts Actions #2904
- FEATURE: Add option to configure proxy for various tools used during setup + persist the proxy configuration #529
- FEATURE: Alerts/Hunt - Provide method for base64-encoding pivot value #1749
- FEATURE: Allow users to customize links in SOC #1248
- FEATURE: Display user who requested PCAP in SOC #2775
- FEATURE: Make SOC browser app connection timeouts adjustable #2408
- FEATURE: Move to FleetDM #3483
- FEATURE: Reduce field cache expiration from 1d to 5m, and expose value as a salt pillar #3537
- FEATURE: Refactor docker\_clean salt state to use loop w/ inspection instead of hardcoded image list #3113
- FEATURE: Run so-ssh-harden during setup #1932
- FEATURE: SOC should only display links to tools that are enabled #1643
- FEATURE: Update Sigmac Osquery Field Mappings #3137
- FEATURE: User must accept the Elastic licence during setup #3233
- FEATURE: soup should output more guidance for distributed deployments at the end #3340
- FEATURE: soup should provide some initial information and then prompt the user to continue #3486
- FIX: Add cronjob for so-suricata-eve-clean script #3515
- FIX: Change Elasticsearch heap formula #1686
- FIX: Create a post install version loop in soup #3102
- FIX: Custom Kibana settings are not being applied properly on upgrades #3254
- FIX: Hunt query issues with quotes #3320
- FIX: IP Addresses don't work with .security #3327
- FIX: Improve DHCP leases query in Hunt #3395
- FIX: Improve Setup verbiage #3422
- FIX: Improve Suricata DHCP logging and parsing #3397
- FIX: Keep RELATED,ESTABLISHED rules at the top of iptables chains #3288

- FIX: Populate `http.status_message` field #3408
- FIX: Remove “types removal” deprecation messages from elastic log. #3345
- FIX: Reword + fix formatting on ES data storage prompt #3205
- FIX: SMTP should read SNMP on Kibana SNMP view #3413
- FIX: Sensors can temporarily show offline while processing large PCAP jobs #3279
- FIX: Soup should log to the screen as well as to a file #3467
- FIX: Strelka port 57314 not immediately relinquished upon restart #3457
- FIX: Switch SOC to pull from fieldcaps API due to field caching changes in Kibana 7.11 #3502
- FIX: Syntax error in `/etc/sysctl.d/99-reserved-ports.conf` #3308
- FIX: Telegraf hardcoded to use https and is not aware of elasticsearch features #2061
- FIX: Zeek Index Close and Delete Count for curator #3274
- FIX: `so-cortex-user-add` and `so-cortex-user-enable` use wrong pillar value for api key #3388
- FIX: `so-rule` does not completely apply change #3289
- FIX: soup should recheck disk space after it tries to clean up. #3235
- UPGRADE: Elastic 7.11.2 #3389
- UPGRADE: Suricata 6.0.2 #3217
- UPGRADE: Zeek 4 #3216
- UPGRADE: Zeek container to use Python 3 #1113
- UPGRADE: docker-ce to latest #3493

## 21.56 2.3.40 Known Issues

- There was a typo in the Zeek index close and delete settings. We’ve fixed this for new installs in <https://github.com/Security-Onion-Solutions/securityonion/issues/3274>. If your deployment has more than 45 days of open Zeek indices, you may want to review these settings in `/opt/so/saltstack/local/pillar/global.sls` and modify them as necessary. This is being tracked in <https://github.com/Security-Onion-Solutions/securityonion/issues/3575>.
- If you had previously enabled Elastic Features and then upgrade to Security Onion 2.3.40 or higher, you may notice some features missing in Kibana. You can enable or disable features as necessary 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>.
- If you upgrade to 2.3.40 and then *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* section.

## 21.57 2.3.30 Changes

- Zeek is now at version 3.0.13.
- CyberChef is now at version 9.27.2.
- Elastic components are now at version 7.10.2. This is the last version that uses the Apache license.
- Suricata is now at version 6.0.1.
- Salt is now at version 3002.5.
- Suricata metadata parsing is now vastly improved.
- If you choose Suricata for metadata parsing, it will now extract files from the network and send them to Strelka. You can add additional mime types here: <https://github.com/Security-Onion-Solutions/securityonion/blob/dev/salt/idstools/sorules/extraction.rules>
- It is now possible to filter Suricata events from being written to the logs. This is a new Suricata 6 feature. We have included some examples here: <https://github.com/Security-Onion-Solutions/securityonion/blob/dev/salt/idstools/sorules/filters.rules>
- The Kratos docker container will now perform DNS lookups locally before reaching out to the network DNS provider.
- Network configuration is now more compatible with manually configured OpenVPN or Wireguard VPN interfaces.
- so-sensor-clean will no longer spawn multiple instances.
- Suricata eve.json logs will now be cleaned up after 7 days. This can be changed via the pillar setting.
- Fixed a security issue where the backup directory had improper file permissions.
- The automated backup script on the manager now backs up all keys along with the salt configurations. Backup retention is now set to 7 days.
- Strelka logs are now being rotated properly.
- Elastalert can now be customized via a pillar.
- Introduced new script so-monitor-add that allows the user to easily add interfaces to the bond for monitoring.
- Setup now validates all user input fields to give up-front feedback if an entered value is invalid.
- There have been several changes to improve install reliability. Many install steps have had their validation processes reworked to ensure that required tasks have been completed before moving on to the next step of the install.
- Users are now warned if they try to set “securityonion” as their hostname.
- The ISO should now identify xvda and nvme devices as install targets.
- At the end of the first stage of the ISO setup, the ISO device should properly unmount and eject.
- The text selection of choosing Suricata vs Zeek for metadata is now more descriptive.
- The logic for properly setting the LOG\_SIZE\_LIMIT variable has been improved.
- When installing on Ubuntu, Setup will now wait for cloud init to complete before trying to start the install of packages.
- The firewall state runs considerably faster now.
- ICMP timestamps are now disabled.
- Copyright dates on all Security Onion specific files have been updated.



- *so-tcpreplay* (and indirectly *so-test*) should now work properly.
- The Zeek packet loss script is now more accurate.
- Grafana now includes an estimated EPS graph for events ingested on the manager.
- Updated Elastalert to release *0.2.4-alt2* based on the <https://github.com/jertel/elastalert> alt branch.
- Pivots from Alerts/Hunts to action links will properly URI encode values.
- Hunt timeline graph will properly scale the data point interval based on the search date range.
- Grid interface will properly show “Search” as the node type instead of “so-node”.
- Import node now supports airgap environments.
- The *so-mysql* container will now show “healthy” when viewing the *docker ps* output.
- The Soctopus configuration now uses private IPs instead of public IPs, allowing network communications to succeed within the grid.
- The Correlate action in Hunt now groups the OR filters together to ensure subsequent user-added filters are correctly ANDed to the entire OR group.
- Add support to *so-firewall* script to display existing port groups and host groups.
- Hive init during Setup will now properly check for a running ES instance and will retry connectivity checks to TheHive before proceeding.
- Changes to the *.security analyzer* yields more accurate query results when using Playbook.
- Several Hunt queries have been updated.
- The pfSense firewall log parser has been updated to improve compatibility.
- Kibana dashboard hyperlinks have been updated for faster navigation.
- Added a new *so-rule* script to make it easier to disable, enable, and modify SIDs.
- ISO now gives the option to just configure the network during setup.

## 21.58 2.3.30 Known Issues

- Heavy Nodes are currently not compatible with Elastic true clustering: <https://github.com/Security-Onion-Solutions/securityonion/issues/3226>
- Custom Kibana settings are not being applied properly on upgrades: <https://github.com/Security-Onion-Solutions/securityonion/issues/3254>

## 21.59 2.3.21 Changes

- *soup* has been refactored. You will need to run it a few times to get all the changes properly. We are working on making this even easier for future releases.
- *soup* now has awareness of Elastic Features and now downloads the appropriate Docker containers.
- The Sensors interface has been renamed to Grid. This interface now includes all Security Onion nodes.
- Grid interface now includes the status of the node. The status currently shows either Online (blue) or Offline (orange). If a node does not check-in on time then it will be marked as Offline.
- Grid interface now includes the IP and Role of each node in the grid.

- Grid interface includes a new Filter search input to filter the visible list of grid nodes to a desired subset. As an example, typing in “sensor” will hide all nodes except those that behave as a sensor.
- The Grid description field can now be customized via the local minion pillar file for each node.
- SOC will now draw attention to an unhealthy situation within the grid or with the connection between the user’s browser and the manager node. For example, when the Grid has at least one Offline node the SOC interface will show an exclamation mark in front of the browser tab’s title and an exclamation mark next to the Grid menu option in SOC. Additionally, the favicon will show an orange marker in the top-right corner (dynamic favicons not supported in Safari). Additionally, if the user’s web browser is unable to communicate with the manager the unhealth indicators appear along with a message at the top of SOC that states there is a connection problem.
- Docker has been upgraded to the latest version.
- Docker should be more reliable now as Salt is now managing daemon.json.
- You can now install Elastic in a traditional cluster. When setting up the manager select Advanced and follow the prompts. Replicas are controlled in global.sls.
- You can now use Hot and Warm routing with Elastic in a traditional cluster. You can change the box.type in the minion’s sls file. You will need to create a curator job to re-tag the indexes based on your criteria.
- Telegraf has been updated to version 1.16.3.
- Grafana has been updated to 7.3.4 to resolve some XSS vulnerabilities.
- Grafana graphs have been changed to graphs vs guages so alerting can be set up.
- Grafana is now completely pillarized, allowing users to customize alerts and making it customizable for email, Slack, etc. See the docs here: <https://securityonion.net/docs/grafana>
- Yara rules now should properly install on non-airgap installs. Previously, users had to wait for an automated job to place them in the correct location.
- Strelka backend will not stop itself any more. Previously, its behavior was to shut itself down after fifteen minutes and wait for Salt to restart it to look for work before shutting down again.
- Strelka daily rule updates are now logged to */nsm/strelka/log/yara-update.log*
- Several changes to the setup script to improve install reliability.
- Airgap now supports the import node type.
- Custom Zeek file extraction values in the pillar now work properly.
- TheHive has been updated to support Elastic 7.
- Cortex image now includes whois package to correct an issue with the CERTatPassiveDNS analyzer.
- Hunt and Alert quick action menu has been refactored into submenus.
- New clipboard quick actions now allow for copying fields or entire events to the clipboard.
- PCAP Add Job form now retains previous job details for quickly adding additional jobs. A new Clear button now exists at the bottom of this form to clear out these fields and forget the previous job details.
- PCAP Add Job form now allows users to perform arbitrary PCAP lookups of imported PCAP data (data imported via the *so-import-pcap* script).
- Downloads page now allows direct download of Wazuh agents for Linux, Mac, and Windows from the manager, and shows the version of Wazuh and Elastic installed with Security Onion.
- PCAP job interface now shows additional job filter criteria when expanding the job filter details.
- Upgraded authentication backend to Kratos 0.5.5.
- SOC tables with the “Rows per Page” dropdown no longer show truncated page counts.

- Several Hunt errors are now more descriptive, particularly those around malformed queries.
- SOC Error banner has been improved to avoid showing raw HTML syntax, making connection and server-side errors more readable.
- Hunt and Alerts interfaces will now allow pivoting to PCAP from a group of results if the grouped results contain a `network.community_id` field.
- New “Correlate” quick action will pivot to a new Hunt search for all events that can be correlated by at least one of various event IDs.
- Fixed bug that caused some Hunt queries to not group correctly without a `.keyword` suffix. This has been corrected so that the `.keyword` suffix is no longer necessary on those groupby terms.
- Fixed issue where PCAP interface loses formatting and color coding when opening multiple PCAP tabs.
- Alerts interface now has a Refresh button that allows users to refresh the current alerts view without refreshing the entire SOC application.
- Hunt and Alerts interfaces now have an auto-refresh dropdown that will automatically refresh the current view at the selected frequency.
- The `so-elastalert-test` script has been refactored to work with Security Onion 2.3.
- The included Logstash image now includes Kafka plugins.
- Wazuh agent registration process has been improved to support slower hardware and networks.
- An Elasticsearch ingest pipeline has been added for `suricata.ftp_data`.
- Elasticsearch’s `indices.query.bool.max_clause_count` value has been increased to accommodate a slightly larger number of fields (1024 -> 1500) when querying using a wildcard.
- On nodes being added to an existing grid, setup will compare the version currently being installed to the manager (`>=2.3.20`), pull the correct Security Onion version from the manager if there is a mismatch, and run that version.
- Setup will gather any errors found during a failed install into `/root/errors.log` for easy copy/paste and debugging.
- Selecting Suricata as the metadata engine no longer results in the install failing.
- `so-rule-update` now accepts arguments to `idstools`. For example, `so-rule-update -f` will force `idstools` to pull rules, ignoring the default 15-minute pull limit.

## 21.60 2.3.10 Changes

- UEFI installs with multiple disks should work as intended now.
- Telegraf scripts will now make sure they are not already running before execution.
- You are now prompted during setup if you want to change the docker IP range. If you change this it needs to be the same on all nodes in the grid.
- Soup will now download the new containers before stopping anything. If anything fails it will now exit and leave the grid at the current version.
- All containers are now hosted on quay.io to prevent pull limitations. We are now using GPG keys to determine if the image is from Security Onion.
- Osquery installers have been updated to osquery 4.5.1
- Fix for bug where Playbook was not removing the Elastalert rules for inactive Plays
- Exifdata reported by Strelka is now constrained to a single multi-valued field to prevent mapping explosion (`scan.exiftool`).

- Resolved issue with Navigator layer(s) not loading correctly.
- Wazuh authd is now started by default on port 1515/tcp.
- Wazuh API default credentials are now removed after setup. Scripts have been added for API user management.
- Upgraded Salt to 3002.2 due to CVEs.
- If salt-minion is unable to apply states after the defined threshold, we assume salt-minion is in a bad state and the salt-minion service will be restarted.
- Fixed bug that prevented mysql from installing for Fleet if Playbook wasn't also installed.
- so-status will now show `STARTING` or `WAIT_START`, instead of `ERROR` if so-status is run before a salt highstate has started or finished for the first time after system startup
- Stenographer can now be disabled on a sensor node by setting the pillar `steno:enabled:false` in its `minion.sls` file or globally if set in the `global.sls` file
- Added `so-ssh-harden` script that runs the commands listed in [SSH](#).
- NGINX now redirects the browser to the hostname/IP address/FQDN based on `global:url_base`
- MySQL state now waits for MySQL server to respond to a query before completing
- Added Analyst option to network installs
- Acknowledging (and Escalating) alerts did not consistently remove the alert from the visible list; this has been corrected.
- Escalating alerts that have a `rule.case_template` field defined will automatically assign that case template to the case generated in TheHive.
- Alerts and Hunt interface quick action bar has been converted into a vertical menu to improve quick action option clarity. Related changes also eliminated the issues that occurred when the quick action bar was appearing to the left of the visible browser area.
- Updated Go to newer version to fix a timezone, daylight savings time (DST) issue that resulted in Alerts and Hunt interfaces not consistently showing results.
- Improved Hunt and Alert table sorting.
- Alerts interface now allows absolute time searches.
- Alerts interface 'Hunt' quick action is now working as intended.
- Alerts interface 'Ack' icon tooltip has been changed from 'Dismiss' to 'Acknowledge' for consistency.
- Hunt interface bar charts will now show the quick action menu when clicked instead of assuming the click was intended to add an include filter.
- Hunt interface quick action will now cast a wider net on field searches.
- Now explicitly preventing the use of a dollar sign (\$) character in web user passwords during setup.
- Cortex container will now restart properly if the SO host was not gracefully shutdown.
- Added syslog plugin to the logstash container; this is not in-use by default but available for those users that choose to use it.
- Winlogbeat download package is now available from the SOC Downloads interface.
- Upgraded Kratos authentication system.
- Added new Reset Defaults button to the SOC Profile Settings interface which allows users to reset all local browser SOC customizations back to their defaults. This includes things like default sort column, sort order, items per page, etc.

## 21.61 2.3.10 Known Issues

- For Ubuntu, non master nodes, you may need to ssh to each node and run `salt-call state.highstate` in order initiate the update. To verify if this needs to be done on remote nodes, from the master, run `salt \* pkg.version salt-minion` after 30 minutes following the initial soup update. If the node does not return that is it running Salt 3002.2, then the node will need to manually be highstated locally from the node to complete the update.
- During soup, you may see the following during the first highstate run, it can be ignored: `Rendering SLS 'some_sls_here' failed: Jinja variable 'list object' has no attribute 'values'`. The second highstate will complete without that error.
- During install or soup, there is a false positive failure condition that can occur. It is caused by `[ERROR ] Failed to add job <job_name> to schedule..` This error indicates that Salt was unable to add a job to a schedule. If you see this in setup or soup log, it can be confirmed if this is false positive or not by running `salt-call schedule.list` on the node that saw the error. If the job isn't in the schedule list, run `salt-call state.highstate` and check if the job was added after it completes.

## 21.62 2.3.2 Changes

- Elastic components have been upgraded to 7.9.3.
- Fixed an issue where curator was unable to delete a closed index.
- Cheat sheet is now available for airgap installs.

## 21.63 2.3.1 Changes

- Fixed a SOC issue in airgap mode that was preventing people from logging in.
- Downloading Elastic features images will now download the correct images.
- Winlogbeat download no longer requires Internet access.
- Adjusted Alerts quick action bar to allow searching for a specific value while remaining in Alerts view.
- `/nsm` will properly display disk usage on the standalone Grafana dashboard.
- The manager node now has syslog listener enabled by default (you'll still need to allow syslog traffic through the firewall of course).
- Fixed an issue when creating host groups with `so-firewall`.

## 21.64 2.3.1 Known Issues

- It is still possible to update your grid from any release candidate to 2.3. However, if you have a true production deployment, then we recommend a fresh image and install for best results.
- In 2.3.0 we made some changes to data types in the elastic index templates. This will cause some errors in Kibana around field conflicts. You can address this in 2 ways:
  - Delete all the data on the ES nodes (preserving all of your other settings such as BPFs) by running `sudo so-elastic-clear` on all the search nodes.

- Re-index the data. This is not a quick process but you can find more information at <https://docs.securityonion.net/en/2.3/elasticsearch.html#re-indexing>
- Please be patient as we update our documentation. We have made a concerted effort to update as much as possible but some things still may be incorrect or omitted. If you have questions or feedback, please start a discussion at <https://securityonion.net/discuss>.
- Once you update your grid to 2.3, any new nodes that join the grid must be 2.3 so if you try to join an older node it will fail. For best results, use the latest 2.3 ISO (or 2.3 installer from github) when joining to a 2.3 grid.
- Shipping Windows Eventlogs with Osquery will fail intermittently with utf8 errors logged in the Application log. This is scheduled to be fixed in Osquery 4.5.
- When running soup to upgrade from older versions to 2.3, there is a Salt error that may occur during the final highstate. This error is related to the `patch_os_schedule` and can be ignored as it should not occur again in subsequent highstates.
- When Search Nodes are upgraded from older versions to 2.3, there is a chance of a race condition where certificates are missing. This will show errors in the manager log to the remote node. To fix this run the following on the search node that is having the issue:
  - Stop elasticsearch - `sudo so-elasticsearch-stop`
  - Run the SSL state - `sudo salt-call state.apply ssl`
  - Restart elasticsearch - `sudo so-elasticsearch-restart`
- If you are upgrading from RC1 you might see errors around registry:2 missing. This error does not break the actual upgrade. To fix, run the following on the manager:
  - Stop the Docker registry - `sudo docker stop so-dockerregistry`
  - Remove the container - `sudo docker rm so-dockerregistry`
  - Run the registry state - `sudo salt-call state.apply registry`

## 21.65 2.3.0 Changes

- We have a new *Alerts* interface for reviewing alerts and acknowledging or escalating them. Escalating creates a new case in TheHive. Please note that TheHive no longer receives alerts directly.
- Kibana no longer presents the option to create alerts from events, but instead allows creation of cases from events.
- Our Security Onion ISO now works for UEFI as well as Secure Boot.
- *Airgap* deployments can now be updated using the latest ISO. Please read this documentation carefully.
- *Suricata* has been updated to version 5.0.4.
- *Zeek* has been updated to version 3.0.11.
- *Stenographer* has been updated to the latest version.
- *soup* will now attempt to clean up old docker images to free up space.
- *Hunt* actions can be customized via `hunt.actions.json`.
- *Hunt* queries can be customized via `hunt.queries.json`.
- *Hunt* event fields can be customized via `hunt.eventfields.json`.
- *Alerts* actions can be customized via `alerts.actions.json`.
- *Alerts* queries can be customized via `alerts.queries.json`.

- *Alerts* event fields can be customized via `alerts.eventfields.json`.
- This help documentation is now viewable offline for airgap installations.
- The script *so-user-add* will now validate the password is acceptable before attempting to create the user.
- *Playbook* and *Grafana* no longer use static passwords for their admin accounts.
- *Analyst VM* now comes with NetworkMiner 2.6 installed.
- *Strelka* YARA matches now generate alerts that can be viewed through the Alerts interface .

## 21.66 2.2.0 Changes

- Setup now includes an option for airgap installations
- Playbook now works properly when installed in airgap mode
- Added *so-analyst* script to create an analyst workstation with GNOME desktop, Chromium browser, Wireshark, and NetworkMiner
- Upgraded Zeek to version 3.0.10 to address a recent security issue
- Upgraded Docker to latest version
- Re-worked IDSTools to make it easier to modify
- Added *so-\** tools to the default path so you can now tab complete
- *so-status* can now be run from a manager node to get the status of a remote node. Run `salt <target> so.status`
- Salt now prevents states from running on a node that it shouldn't so you can't, for example, accidentally apply the elasticsearch state on a forward node
- Added logic to check for Salt mine corruption and recover automatically
- Collapsed Hunt filter icons and action links into a new quick action bar that will appear when a field value is clicked; actions include:
  - Filtering the hunt query
  - Pivot to PCAP
  - Create an alert in TheHive
  - Google search for the value
  - Analyze the value on VirusTotal.com
- Fixed minor bugs in Hunt user interface relating to most-recently used queries, tooltips, and more
- *so-user-add* now automatically adds users to Fleet and TheHive (in addition to SOC)
- Introduced *so-user-disable* and *so-user-enable* commands which allows administrators to lock out users that are no longer permitted to use Security Onion
- Added icon to SOC Users list representing their active or locked out status
- Removed User delete action from SOC interface in favor of disabling users for audit purposes
- Prune old PCAP job data from sensors once the results are streamed back to the manager node
- Hunt filtering to a specific value will search across all fields instead of only the field that was originally clicked
- Limiting PCAP jobs to extract at most 2GB from a sensor to avoid users accidentally requesting unreasonably large PCAP via the web interface

- `so-test` is back - run it to easily replay PCAPs and verify that all the components are working as expected
- New Elasticsearch subfield (`.security`) based on the new community-driven analyzer from @neu5ron - [https://github.com/neu5ron/es\\_stk](https://github.com/neu5ron/es_stk)
- Playbook now uses the new `.security` subfield for case-insensitive wildcard searches

## 21.67 2.1.0 Changes

- Fixed an issue where the console was timing out and making it appear that the installer was hung
- Introduced Import node type ideal for running `so-import-pcap` to import pcap files and view the resulting logs in Hunt or Kibana
- Moved `static.sls` to `global.sls` to align the name with the functionality
- Traffic between nodes in a distributed deployment is now fully encrypted
- Playbook
  - Elastalert now runs active Plays every 3 minutes
  - Changed default rule-update config to only import Windows rules from the Sigma Community repo
  - Lots of bug fixes & stability improvements
- Ingest Node parsing updates for Osquery and Winlogbeat - implemented single pipeline for Windows eventlogs & sysmon logs
- Upgraded Osquery to 4.4 and re-enabled auto-updates
- Upgraded to Salt 3001.1
- Upgraded Wazuh to 3.13.1
- Hunt interface now shows the timezone being used for the selected date range
- Fixed Cortex initialization so that TheHive integration and initial user set is correctly configured
- Improved management of TheHive/Cortex credentials
- SOC now allows for arbitrary, time-bounded PCAP job creation, with optional filtering by host and port

## 21.68 2.0.3 Changes

- Resolved an issue with large drives and the ISO install
- Modified ISO installation to use Logical Volume Management (LVM) for disk partitioning
- Updated Elastic Stack components to version 7.8.1
- Updated Zeek to version 3.0.8



## 21.69 2.0.2 Changes

- Sensoroni fails on 2.0.1 ISO EVAL installation #1089  
<https://github.com/Security-Onion-Solutions/securityonion/issues/1089>

## 21.70 2.0.1 Changes

- Security Fix: variables.txt from ISO install stays on disk for 10 days  
<https://github.com/Security-Onion-Solutions/securityonion/issues/1067>
- Security Fix: Remove user values from static.sls  
<https://github.com/Security-Onion-Solutions/securityonion/issues/1068>
- Fix distributed deployment sensor interval issue allowing PCAP  
<https://github.com/Security-Onion-Solutions/securityonion/issues/1059>
- Support for passwords that start with special characters  
<https://github.com/Security-Onion-Solutions/securityonion/issues/1058>
- Minor soup updates

## 21.71 2.0.0 Changes

- This version requires a fresh install, but there is good news - we have brought back *soup*! From now on, you should be able to run *soup* on the manager to upgrade your environment to RC2 and beyond!
- Re-branded 2.0 to give it a fresh look
- All documentation has moved to our docs site
- soup is alive! Note: This tool only updates Security Onion components. Please use the built-in OS update process to keep the OS and other components up to date
- so-import-pcap is back! See the docs here
- Fixed issue with so-features-enable
- Users can now pivot to PCAP from Suricata alerts
- ISO install now prompts users to create an admin/sudo user instead of using a default account name
- The web email & password set during setup is now used to create the initial accounts for TheHive, Cortex, and Fleet
- Fixed issue with disk cleanup
- Changed the default permissions for /opt/so to keep non-privileged users from accessing salt and related files
- Locked down access to certain SSL keys
- Suricata logs now compress after they roll over
- Users can now easily customize shard counts per index
- Improved Elastic ingest parsers including Windows event logs and Sysmon logs shipped with WinLogbeat and Osquery (ECS)
- Elastic nodes are now “hot” by default, making it easier to add a warm node later

- so-allow now runs at the end of an install so users can enable access right away
- Alert severities across Wazuh, Suricata and Playbook (Sigma) have been standardized and copied to event.severity:
  - 1-Low / 2-Medium / 3-High / 4-Critical
- Initial implementation of alerting queues:
  - Low & Medium alerts are accessible through Kibana & Hunt
  - High & Critical alerts are accessible through Kibana, Hunt and sent to TheHive for immediate analysis
- ATT&CK Navigator is now a statically-hosted site in the nginx container
- Playbook
  - All Sigma rules in the community repo (500+) are now imported and kept up to date
  - Initial implementation of automated testing when a Play's detection logic has been edited (i.e., Unit Testing)
  - Updated UI Theme
  - Once authenticated through SOC, users can now access Playbook with analyst permissions without login
- Kolide Launcher has been updated to include the ability to pass arbitrary flags - new functionality sponsored by SOS
- Fixed issue with Wazuh authd registration service port not being correctly exposed
- Added option for exposure of Elasticsearch REST API (port 9200) to so-allow for easier external querying/integration with other tools
- Added option to so-allow for external Strelka file uploads (e.g., via strelka-fileshot)
- Added default YARA rules for Strelka – default rules are maintained by Florian Roth and pulled from <https://github.com/Neo23x0/signature-base>
- Added the ability to use custom Zeek scripts
- Renamed “master server” to “manager node”
- Improved unification of Zeek and Strelka file data

## **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/>