
Slips

Release 1.1.21

Stratosphere Laboratory

Jun 09, 2026

SLIPS

1	Installation	3
1.1	Table of Contents	3
1.2	Requirements	3
1.3	Slips in Docker	3
1.4	Installing Slips natively	7
1.5	Installing Slips on a Raspberry PI	9
2	Usage	11
2.1	Reading the input	11
2.2	Daemonized vs interactive mode	11
2.3	Running Several slips instances	12
2.4	Closing redis servers	13
2.5	Growing zeek directories	13
2.6	Reading the output	14
2.7	Saving the database	15
2.8	Whitelisting	15
2.9	Popup notifications	17
2.10	Slips permissions	18
2.11	Modifying the configuration file	18
2.12	Logging	20
2.13	Reading Input from stdin	21
2.14	Slips as an access point	22
2.15	Plug in a zeek script	22
2.16	Exporting strato letters	22
2.17	Slips parameters	22
2.18	Limiting Slips resource consumption	23
2.19	Running Slips from python	23
3	Architecture	25
3.1	Internal representation of data.	25
3.2	Alerts vs Evidence	26
3.3	Usage of Zeek.	27
3.4	Usage of Redis database.	27
3.5	Usage of SQLite database.	27
3.6	Threat Levels	27
3.7	How Slips Stops	27
4	Detection modules	29
4.1	Bruteforcing Module	29
4.2	HTTPS Anomaly Detection Module	29

4.3	RNN C&C Detection Module	30
4.4	Leak Detection Module	32
4.5	Blocking Module	32
4.6	Exporting Alerts Module	33
4.7	Flowalerts Module	33
4.8	Disabled alerts	33
4.9	Threat Intelligence Module	33
4.10	Update Manager Module	37
4.11	IP Info Module	37
4.12	Reverse DNS	38
4.13	ARP Module	38
4.14	CESNET sharing Module	39
4.15	HTTP Analyzer Module	39
4.16	Leak Detector Module	41
4.17	Network Service Discovery Module	41
5	Connections Made By Slips	43
6	Zeek Scripts	45
6.1	Detect DoH	45
6.2	Detect ICMP Scans	45
6.3	Detect the Gateway address	45
7	Brute force detector Module	47
7.1	Inputs	47
7.2	What It Detects	47
7.3	Detection Logic	47
7.4	Evidence Produced	48
7.5	Zeek Confirmation	49
7.6	Configuration	49
7.7	Relationship With Flow Alerts	49
8	HTTPS Anomaly Detection Module	51
8.1	Goal	51
8.2	Input data used	52
8.3	Traffic-time logic	52
8.4	Features	52
8.5	Baseline and training	52
8.6	Scoring	53
8.7	Adaptation states	53
8.8	New server vs JA3 behavior	54
8.9	Confidence and threat level	55
8.10	Evidence format	55
8.11	Configuration keys	55
8.12	Operational logs	56
9	flow_alerts Module	57
9.1	Long Connections	58
9.2	Connections without DNS resolution	58
9.3	Successful SSH connections	59
9.4	DNS resolutions without a connection	59
9.5	Connection to unknown ports	59
9.6	Data Upload	60
9.7	Tor Exit Nodes	60
9.8	Malicious JA3 and JA3s hashes	60

9.9	Connections to port 0	60
9.10	Multiple reconnection attempts	60
9.11	Zeek alerts	60
9.12	SMTP login bruteforce	61
9.13	SSH brute_force_detector	61
9.14	DGA	61
9.15	Connection to multiple ports	61
9.16	Malicious SSL certificates	61
9.17	Pastebin downloads	61
9.18	Young Domains	61
9.19	Bad SMTP logins	62
9.20	SMTP bruteforce	62
9.21	DNS ARPA Scans	62
9.22	SSH version changing	62
9.23	Incompatible CN	63
9.24	CN URL Mismatch	63
9.25	Weird HTTP methods	63
9.26	Non-SSL connections on port 443	63
9.27	Connection to private IPs	63
9.28	Connection to private IPs outside the current local network	63
9.29	High entropy DNS TXT answers	64
9.30	Devices changing IPs	64
9.31	GRE tunnels	64
9.32	GRE tunnel scans	64
9.33	Login log entries	64
9.34	Invalid DNS resolutions	65
10	Features	67
10.1	Flow Alerts Module	67
10.2	GRE tunnels	68
10.3	GRE tunnel scans	69
10.4	Non-HTTP connections on port 80.	70
10.5	Invalid DNS resolutions	74
10.6	Detection modules	74
10.7	Connections Made By Slips	82
10.8	Ensembling	82
10.9	Controlling Slips Sensitivity	83
10.10	Zeek Scripts	83
11	Training	85
12	Exporting Slips Alerts	87
12.1	Slack	87
12.2	STIX	88
12.3	JSON format	88
12.4	CESNET Sharing	88
12.5	Logstash	89
12.6	Text logs	89
12.7	TSV and json of labeled flows	89
13	P2P	91
13.1	Pigeon	91
13.2	Docker direct use	91
13.3	Installation:	91
13.4	Usage in Slips	92

13.5	Project sections	92
13.6	Dovecot experiments	93
13.7	How it works:	93
13.8	Logs	93
13.9	Limitations	94
13.10	TLDR;	94
14	Fides module. Global P2P Threat Ingelligence Sharing	95
14.1	Docker direct use	95
14.2	Conditions	95
14.3	Configuration	95
14.4	Usage in Slips	96
14.5	How it works:	96
14.6	Logs	96
14.7	Implementation notes and credit	96
14.8	Privacy	96
15	How to Create a New Slips Module	97
15.1	What is SLIPS and why are modules useful	97
15.2	Goal of this Blog	97
15.3	Creating a Module	98
15.4	Complete Code	103
15.5	Line by Line Explanation of the Module	106
15.6	Reading Input flows from an external module (Advanced)	106
16	Datasets	109
16.1	2017-3-8_win5.pcap	110
16.2	conn.log	110
16.3	CTU-Malware-Capture-Botnet-1	110
16.4	malicious-cc.conn.log	110
16.5	port-scans	110
16.6	test10-mixed-zeek-dir	110
16.7	test11-portscan.binetflow	110
16.8	test12-icmp-portscan.pcap	110
16.9	test13-malicious-dhcpscan-zeek-dir	110
16.10	test14-malicious-zeek-dir	110
16.11	test15-malicious-zeek-dir	110
16.12	test16-malicious-zeek-dir	110
16.13	test1-normal.nfdump	110
16.14	test2-malicious.binetflow	110
16.15	test3-mixed.binetflow	110
16.16	test4-malicious.binetflow	110
16.17	test5-mixed.binetflow	110
16.18	test6-malicious.suricata.json	110
16.19	test7-malicious.pcap	110
16.20	test8-malicious.pcap	110
16.21	test9-mixed-zeek-dir	110
16.22	test-cc	110
17	Slips Immune	113
17.1	Architecture	113
17.2	RPI performance	113
17.3	Updating Slips	113
17.4	Security & Network Configuration	113
17.5	Datasets & LLM Training	114

18 Slips In Action	115
18.1 Saefko RAT	115
18.2 Emotet	116
18.3 DroidJack v4.4 RAT	118
19 FAQ	121
19.1 Slips starting too slow in docker	121
19.2 Getting “Illegal instruction” error when running slips	121
19.3 Docker time is not in sync with that of the host	121
19.4 Redis WARNING Memory overcommit must be enabled! in docker	121
20 Contributing	123
20.1 How can you contribute?	123
20.2 Persistent Git Branches	123
20.3 Naming Git branches for Pull Requests	123
20.4 What branch should you base your contribution to Slips?	123
20.5 Creating a pull request	124
20.6 Beginner tips on how to create a PR in Slips	124
20.7 How to test the auto update functionality	124
20.8 Rejected PRs	125
20.9 FAQ	125
20.10 Global P2P - Fides contribution notes	128
21 Code documentation	131
22 Related Repositories	133



The tool is available on GitHub [here](#).

Slips is a Python-based intrusion prevention system that uses machine learning to detect malicious behaviors in the network traffic. Slips was designed to focus on targeted attacks, to detect of command and control channels, and to provide good visualisation for the analyst. Slips is able to analyze real live traffic from the device and the large network captures in the type of a pcap files, Suricata, Zeek/Bro and Argus flows. As a result, Slips highlights suspicious behaviour and connections that needs to be deeper analyzed.

This documentation gives an overview how Slips works, how to use it and how to help. To be specific, that table of contents goes as follows:

- **Installation.** Instructions to install Slips in a Docker and in a computer. See [Installation](#).
- **Usage.** Instructions and examples how to run Slips with different type of files and analyze the traffic using Slips and its web interface or the optional Kalipso submodule. See [Usage](#).
- **Detection modules.** Explanation of detection modules in Slips, types of input and output. See [Detection modules](#).
- **brute_force_detector.** Dedicated documentation for the SSH brute force detector module. See [brute_force_detector](#).
- **HTTPS anomaly detection.** Detailed design and behavior of the HTTPS anomaly detector. See [HTTPS anomaly detection](#).
- **Architecture.** Internal architecture of Slips (profiles, timewindows), the use of Zeek and connection to Redis. See [Architecture](#).
- **Training with your own data.** Explanation on how to re-train the machine learning system of Slips with your own traffic (normal or malicious). See [Training](#).
- **Detections per Flow.** Explanation on how Slips works to make detections on each flow with different techniques. See [flow_alerts](#).
- **Exporting.** The exporting module allows Slips to export to Slack and STIX servers. See [Exporting](#).
- **Slips in Action.** Example of using slips to analyze different PCAPs See [Slips in action](#).
- **Contributing.** Explanation how to contribute to Slips, and instructions how to implement new detection module in Slips. See [Contributing](#).

- **Create a new module.** Step by step guide on how to create a new Slips module See [Create a new module](#).
- **Code documentation.** Auto generated slips code documentation See [Code docs](#).
- **Datasets.** The folder *dataset* contains some testing datasets for you to try. See [Datasets](#).

INSTALLATION

There are two ways to install and run Slips: inside a Docker or in your own computer. We suggest to install and to run Slips inside a Docker since all dependencies are already installed in there. However, current version of docker with Slips does not allow to capture the traffic from the computer's interface. We will describe both ways of installation anyway.

1.1 Table of Contents

- Docker
 - Dockerhub (recommended)
 - * Linux and windows hosts
 - * MacOS hosts
 - * *Light Slips Image*
 - Docker-compose
 - Dockerfile
- Native
 - Using install.sh
 - Manually
- on RPI (Beta)

1.2 Requirements

- Python 3.10.12
- 5 GBs of disk space (for the docker image)
- at least 4 GBs of RAM

1.3 Slips in Docker

Slips can be run inside a Docker. Either using our docker image with from DockerHub (recommended) or building Slips image from the Dockerfile for more advanced users.

In both cases, you need to have the Docker platform installed in your computer. For instructions how to install Docker check <https://docs.docker.com/get-docker/>.

The recommended way of using slips would be to

- Run Slips from Dockerhub

For more advanced users, you can:

- *Run Slips using docker compose*
- *Build Slips using the dockerfile*

1.3.1 Running Slips from DockerHub

Linux And Windows Hosts

Analyse your own traffic

```
docker run --rm -it -p 55000:55000 --cpu-shares "700" --memory="8g" --memory-swap="8g" --
↪net=host --cap-add=NET_ADMIN -v $(pwd)/output:/StratosphereLinuxIPS/output -v $(pwd)/
↪dataset:/StratosphereLinuxIPS/dataset --name slips stratosphereips/slips:latest /
↪StratosphereLinuxIPS/slips.py -i eno1
```

Please change the name of the interface for your own. Check the alerts slips generated

```
tail -f output/eno1*/alerts.log
```

MacOS Hosts

Analyse your own traffic

```
docker run --platform linux/amd64 --rm -it -p 55000:55000 --cpu-shares "700" --memory="8g
↪" --memory-swap="8g" --net=host --cap-add=NET_ADMIN -v $(pwd)/output:/
↪StratosphereLinuxIPS/output -v $(pwd)/dataset:/StratosphereLinuxIPS/dataset --name
↪slips stratosphereips/slips:latest /StratosphereLinuxIPS/slips.py -i eno1
```

Please change the name of the interface to your own. Check the alerts slips generated

```
tail -f output/eno1*/alerts.log
```

To analyze your own files using slips, you can mount it to your docker using -v

```
mkdir ~/dataset
cp <some-place>/myfile.pcap ~/dataset
docker run -it --rm --net=host -v ~/dataset:/StratosphereLinuxIPS/dataset
↪stratosphereips/slips:latest
↪./slips.py -f dataset/myfile.pcap
```

Minimal Slips Docker Image

In addition to the full stratosphereips/slips:latest image, there is now a minimal Docker image available: using docker pull stratosphereips/slips_light:latest. This image excludes the following modules to reduce size and resource usage:

- rnn_cc_detection/
- timeline/
- p2p_trust/
- flow_ml_detection/

- cyst/
- cesnet/
- exporting_alerts/
- risk_iq/
- template/
- blocking/
- virustotal/

Additionally, several directories and files have been removed from this minimal image, including: dataset/, docs/, tests/ Slips' two GUIs, kalipso and the web interface, aren't available in this image. The only way to check the slips output there is by going through the generated logs.

The stratosphereips/slips_light:latest image is recommended for users who do not require these specific modules and want a more lightweight deployment.

1.3.2 Update slips image

```
docker pull stratosphereips/slips:latest
```

1.3.3 Run Slips sharing files between the host and the container

The following instructions will guide you on how to run a Slips docker container with file sharing between the host and the container.

```
# create a directory to load pcaps in your host computer
mkdir ~/dataset

# copy the pcap to analyze to the newly created folder
cp <some-place>/myfile.pcap ~/dataset

# create a new Slips container mapping the folder in the host to a folder in the
↪container
docker run -it --rm --net=host --name slips -v $(pwd)/dataset:/StratosphereLinuxIPS/
↪dataset stratosphereips/slips:latest

# run Slips on the pcap file mapped to the container
./slips.py -f dataset/myfile.pcap
```

1.3.4 Run Slips with access to block traffic on the host network

In Linux OS, the Slips can be used to analyze and **block** network traffic on the host network interface. To allow the container to see the host interface traffic and block malicious connections, it needs to run with the option `--cap-add=NET_ADMIN`. This option enables the container to interact with the network stack of the host computer. To block malicious behavior, run Slips with the parameter `-p`.

Change eno1 in the command below to your own interface

```
# run a new Slips container with the option to interact with the network stack of
↪the host
docker run -it --rm --net=host --cap-add=NET_ADMIN --name slips stratosphereips/
↪slips:latest
```

(continues on next page)

(continued from previous page)

```
# run Slips on the host interface `eno1` with active blocking `-p`  
./slips.py -i eno1 -p
```

1.3.5 Running Slips using docker compose

Change enp1s0 to your current interface in docker/docker-compose.yml and start slips using

```
docker compose -f docker/docker-compose.yml up
```

Now everything inside your host's config and dataset directories is mounted to /StratosphereLinuxIPS/config/ and /StratosphereLinuxIPS/dataset/ in Slips docker.

To run slips on a pcap instead of your interface you can do the following:

1. put the pcap in the dataset/ dir in your host
2. change the entrypoint in the docker compose file to ["python3","/StratosphereLinuxIPS/slips.py","-f","dataset/pcap"]
3. restart slips using `docker compose -f docker/docker-compose.yml up`

Limitations

The main limitation of running Slips in a Docker is that every time the container stops, all files inside the container are deleted, including the Redis database of cached data, and you lose all your Threat Intelligence (TI) data and previous detections. Next time you run Slips, it will start making detections without all the TI data until downloading the data again. The only solution is to keep the container up between scans.

1.3.6 Building Slips from the Dockerfile

Before building the docker locally from the Dockerfile, first you should clone Slips repo or download the code directly:

```
git clone https://github.com/stratosphereips/StratosphereLinuxIPS.git
```

NOTE: you have to be in the main Slips directory to build this.

If you cloned Slips in '~/StratosphereLinuxIPS', then you can build the Docker image with:

```
cd ~/StratosphereLinuxIPS  
docker build --target amd --no-cache -t slips -f docker/Dockerfile .  
docker run -it --rm --net=host slips  
cp config/slips.yaml config/my_slips.yaml  
./slips.py -c config/my_slips.yaml -f dataset/test3-mixed.binetflow
```

If you don't have Internet connection from inside your Docker image while building, you may have another set of networks defined in your Docker. For that try:

```
docker build --target amd --network=host --no-cache -t slips -f docker/Dockerfile .
```

You can also put your own files in the /dataset/ folder and analyze them with Slips:

```
cp some-pcap-file.pcap ~/StratosphereLinuxIPS/dataset
docker run -it --rm --net=host -v ~/StratosphereLinuxIPS/dataset:/StratosphereLinuxIPS/
↳dataset slips
./slips.py -f dataset/some-pcap-file.pcap
```

Note that some GPUs don't support tensorflow in docker which may cause "Illegal instruction" errors when running slips. You can read more about it [here](#)

1.4 Installing Slips natively

Slips depends on three major elements:

- Python 3.10.12
- Zeek 8.0.0
- Redis database v8

To install these elements, the script will use the APT package manager. After that, it will install python packages required for Slips to run and its modules to work.

Instructions to download everything for Slips are below.

1.4.1 Install Slips using shell script

You can install it using `install.sh`

```
sudo chmod +x install.sh
sudo ./install.sh
```

The script installs Slips core dependencies and builds p2p4slips. It does not install Kalipso.

1.4.2 Installing Slips manually

Installing Python, Redis, and required python libraries.

Update the repository of packages so you see the latest versions:

```
apt-get update
```

Install apt dependencies:

```
cat install/apt_dependencies.txt | xargs apt-get -y install
```

Even though we just installed pip3, the package installer for Python (3.10.12), we need to upgrade it to its latest version:

```
python3 -m pip install --upgrade pip
```

Now that pip3 is upgraded, we can proceed to install all required packages via pip3 python packet manager:

```
python3 -m pip3 install -r install/requirements.txt
```

Note: for those using a different base image, install the TensorFlow version pinned in `install/requirements.txt` via `pip3`.

Optional Kalipso submodule

Kalipso is maintained in a separate repository and checked out as the `modules/kalipso` submodule.

If you cloned Slips without submodules, initialize it with:

```
git submodule update --init --recursive modules/kalipso
```

If you are cloning Slips for the first time and want all optional components, use:

```
git clone --recurse-submodules --remote-submodules https://github.com/stratosphereips/  
↳StratosphereLinuxIPS -j4
```

Kalipso has its own installation instructions in `modules/kalipso/README.md`. It is optional and not required for Docker, CI, or core Slips execution.

Installing Zeek

The last requirement to run Slips is Zeek. Zeek is not directly available on Ubuntu or Debian. To install it, we will first add the repository source to our apt package manager source list. The following two commands are for Ubuntu 18.04:

check the repositories for the correct version if you are using a different OS:

<https://software.opensuse.org//download.html?project=security%3Azeek&package=zeek>

```
echo 'deb http://download.opensuse.org/repositories/security:/zeek/xUbuntu_18.04/ /' |  
↳tee /etc/apt/sources.list.d/security:zeek.list
```

We will download and store the gpg signature from the package for apt to read:

```
curl -fsSL http://download.opensuse.org/repositories/security:/zeek/xUbuntu_18.04/  
↳Release.key | gpg --dearmor | tee /etc/apt/trusted.gpg.d/security_zeek.gpg > /dev/null
```

Finally, we will update the package manager repositories and install zeek

```
apt-get update  
apt-get -y install zeek
```

To make sure that zeek can be found in the system we will add its link to a known path:

```
ln -s /opt/zeek/bin/zeek /usr/local/bin
```

PS: Slips supports the latest zeek release by default, make sure you install it and not the LTS.

Installing Redis v8

Install redis v8 from source as it's not available in their official repositories.

```
curl -L https://download.redis.io/redis-stable.tar.gz -o /tmp/redis-stable.tar.gz \  
  && mkdir -p /redis-stable \  
  && tar xzf redis-stable.tar.gz -C / \  
  && cd /redis-stable \  
  && make distclean \  
  && make MALLOC=libc
```

Please remember to add `/redis-stable/src` to your `PATH` for slips to be able to use it.

Running Slips for the First Time

Be aware that the first time you run Slips it will start updating all the databases and threat intelligence files in the background. However, it will give you as many detections as possible *while* updating. You may have more detections if you rerun Slips after the updates. Slips behaves like this, so you don't have to wait for the updates to finish to have some detections. however, you can change that in the config file by setting `wait_for_TI_to_finish` to yes.

Depending on the remote sites, downloading and updating the DB may take up to 4 minutes. Slips stores this information in a cache Redis database, which is kept in memory when Slips stops. Next time Slips runs, it will read from this database. The information in the DB is updated periodically according to the configuration file (usually one day).

You can check if the DB is running this by looking at your processes:

```
ps afx | grep redis
9078 ?          Ssl    1:25 redis-server *:6379
```

You can kill this redis database by running:

```
./slips.py -k
Choose which one to kill [0,1,2 etc..]
[0] Close all servers
[1] conn.log - port 6379
```

then choosing 1.

After these steps, if you need optional submodules, initialize them with `git submodule update --init --recursive` and then follow the instructions in their own READMEs.

1.5 Installing Slips on a Raspberry PI

The recommended way to install Slips on the RPI is using docker.

If you're using the 64-bit (arm64) version of the RPI, follow the official docker [installation instructions for Debian](#).

Slips now supports a native linux/arm64 docker image, you can pull it using

```
docker pull stratosphereips/slips:latest
```


To enable P2P, make sure of the following:

- You run Slips docker with `--net=host`
- You don't have redis running on the host and occupying Redis' default IP/Port 127.0.0.1:6379.

1.5.1 Protect your local network with Slips on the RPI

By installing Slips on your RPI and using it as an access point, you can extend its protection to your other connected devices.

Once Slips detects a malicious device, it will block all traffic to and from it using iptables. Meaning it will kick out the malicious device from the AP.



1. Connect your RPI to your router using an ethernet cable
2. Install `linux-wifi-hotspot`
3. Start the access point (in NAT mode)

```
sudo create_ap wlan0 eth0 rpi_wifi mysecurepassword -c 40
```

where `wlan0` is the wifi interface of your RPI, `eth0` is the ethernet interface and `-c 40` is the channel of the access point.

We chose channel 40 because it is a 5GHz channel, which is faster and less crowded than the 2.4GHz channels.

Note: Please make sure your RPI model supports 5GHz channels. If not, you can use `-c 1` for 2.4GHz.

If all goes well you should see `wlan0: AP-ENABLED` in the output of the command.

Check the [Debugging common AP errors](#) section if you have any issues.

4. Run Slips in the RPI using the command below to listen to the traffic from the access point.

```
./slips.py -i wlan0
```

5. (Optional) If you want to block malicious devices, run Slips with the `-p` parameter. Using this parameter will block all traffic to and from the malicious device when slips sets an alert.

```
./slips.py -i wlan0 -p
```

Now connect your devices to the `rpi_wifi` with “mysecurepassword” as the password, and enjoy the protection of Slips.

USAGE

Slips can read the packets directly from the **network interface** of the host machine, and packets and flows from different types of files, including

- Pcap files (internally using Zeek)
- Packets directly from an interface (internally using Zeek)
- Suricata flows (from JSON files created by Suricata, such as eve.json)
- Argus flows (CSV file separated by commas or TABs)
- Zeek/Bro flows from a Zeek folder with log files
- Nfdump flows from a binary nfdump file
- Text flows from stdin in zeek, argus or suricata form

It's recommended to use PCAPs.

All the input flows are converted to an internal format. So once read, Slips works the same with all of them.

After Slips runs on the given traffic, the output can be analyzed with the web interface. If the `modules/kalipso` submodule is checked out and installed separately, you can also use the Kalipso terminal interface.

In this section, we will explain how to execute each type of file in Slips.

Either you are [running Slips in docker](#) or [locally](#), you can run Slips using the same below commands and configurations.

2.1 Reading the input

The table below shows the different parameter Slips uses for different inputs.

(*) To find the interface in Linux, you can use the command `ifconfig`.

There is also a configuration file `config/slips.yaml` where the user can set up parameters for Slips execution and models separately. Configuration of the `config/slips.yaml` is described [here](#).

2.2 Daemonized vs interactive mode

Slips has 2 modes, interactive and daemonized.

Daemonized : Slips runs as a daemon in the background. which means all output, logs and alerts are written in files and nothing is displayed in the terminal.

In daemonized mode : Slips runs completely in the background, The output is written to `stdout`, `stderr` and `logfile` files specified in `config/slips.yaml`

by default, these are the paths used

`stdout = /var/log/slips/slips.log stderr = /var/log/slips/error.log logfile = /var/log/slips/slips.log`

NOTE: Since `/var/log/` is owned by root by default, If you want to store the logs in `/var/log/slips`, create `/var/log/slips` as root and slips will use it by default.

If slips can't write there, slips will store the logs in the `Slips/output/<some_dir>` dir by default.

NOTE: if `-o <output_dir>` is given when slips is in daemonized mode, the output log files will be stored in `<output_dir>` instead of the `output_dir` specified in `config/slips.yaml`

This is not the default mode, to use it, run slips with `-D`

```
./slips.py -i wlp3s0 -D
```

To stop the daemon run slips with `-S`, for example `./slips.py -S`

Only one instance of the daemon can run at a time.

Interactive : For viewing output, logs and alerts in your terminal.

This is the default mode, It doesn't require any flags.

Output files are stored in `output/<some_dir>` dir. the output directory used will be logged to your terminal.

By default you don't need root to run slips, but if you changed the default output directory to a dir that is owned by root, you will need to run Slips using `sudo` or give the current user enough permission so that slips can write to those files.

2.3 Running Several slips instances

By default, Slips will assume you are running only one instance and will use the redis port 6379 on each run.

You can run several instances of slips at the same time using the `-m` flag, and the output of each instance will be stored in `output/filename_timestamp/` directory.

If you want Slips to run on a certain port, you can use the `-P <portnumber>` parameter to specify the port you want Slips to use. but it will always use port 6379 db 1 for the cache db.

Each instance of Slips will connect to redis server on a randomly generated port in the range (32768 to 32850).

In macos, you will get a popup asking for permission to open and use that random port, press yes to allow it.

However, all instance share 1 cached redis database on `redis://localhost:6379 DB 1`, to store the IoCs taken from TI files.

Both redis servers, the main sever (DB 0) and the cache server (DB 1) are opened automatically by Slips.

When using the web interface, you can select among the active Redis-backed analyses:

```
To close all unused redis servers, run slips with --killall
You have 3 open redis servers, Choose which one to use [1,2,3 etc..]
[1] wlp3s0 - port 55879
[2] dataset/test7-malicious.pcap - port 59324
```

You can choose the corresponding file or interface from the web interface. You can also use the `Browse redis database` button in the web interface to load a saved `.rdb` file directly. The upload accepts only files with the `.rdb` extension.

Once you're done, you can run slips with `--killall` to close all the redis servers using the following command

```
./slips.py --killall
```

NOTICE: if you run more than one instance of Slips on the same file or the same interface, Slips will generate a new directory with the name of the file and the new timestamp inside the `output/` dir

2.4 Closing redis servers

Slips uses a random unused port in the range in the range (32768 to 32850).

When running slips, it will warn you if you have more than 1 redis serve open using the following msg

```
[Main] Warning: You have 2 redis servers running. Run Slips with --killall to stop them.
```

you can use the `-k` flag to kill 1 open redis server, or all of them using the following command

```
./slips.py -k
```

You will be prompted with the following options

```
Choose which one to kill [0,1,2 etc..]

[0] Close all servers
[1] dataset/sample_zeek_files - port 32768
[2] dataset/sample_zeek_files - port 32769
[3] dataset/sample_zeek_files - port 32770
```

you can select the number you want to kill or 0 to close all the servers.

Note that if all ports from (32768 to 32850) are unavailable, slips won't be able to start, and you will be asked to close all all of them using the following warning

```
All ports from 32768 to 32769 are used. Unable to start slips.

Press Enter to close all ports.
```

You can press enter to close all ports, then start slips again.

2.5 Growing zeek directories

Due to issues running Slips on an interface on MacOS, we added the option `-g` to run slips on a growing zeek directory so you can run Slips in docker, mount a into the container then start zeek inside it using the following command

```
bro -C -i <interface> tcp_inactivity_timeout=60mins tcp_attempt_delay=1min <path_to>/
slips/zeek-scripts
```

Then start Slips on your zeek dir using `-g` and specify the interface you're monitoring with `-i`.

```
./slips.py -e 1 -g <zeek_dir> -i <interface>
```

By using the `-g` parameter, slips will treat your given zeek directory as growing (the same way it treats zeek directories generated by using slips with `-i`) and will keep waiting for flows unless stopped by the user.

NOTE: When using `-g`, it is mandatory to give Slips the same interface zeek is running on with `-i`.

2.6 Reading the output

The output process collects output from the modules and handles the display of information on screen. Currently, Slips' analysis and detected malicious behaviour can be analyzed as following:

- **alerts.json and alerts.txt in the output folder** - collects all evidences and detections generated by Slips in a .txt and .json formats.
- **log files in a folder *current-date-time*** - separates the traffic into files according to a profile and timewindow and summarize the traffic according to each profile and timewindow.
- **Web interface** - browser based GUI for viewing slips detections, incoming and outgoing traffic, and an organized timeline of flows.

2.6.1 The Web Interface

You can use Slips' web interface by running slips with `-w` or running:

```
./webinteface.sh
```

Then navigate to `http://localhost:55000/` from your browser.

Use `Change DB` to switch between active Redis server. Use `Browse redis database` to upload and inspect a saved Redis `.rdb` file. If the file is not a Redis RDB file or Redis cannot load it, the web interface shows a warning and keeps the current database selected.

On the right column, you can see a list of all the IPs seen in your traffic.

The traffic of IP is splitted into time windows. each time window is 1h long of traffic.

IPs and timewindows that are marked in red are considered malicious in Slips.

You can view the traffic of each time window by clicking on it

- The timeline button shows the zeek logs formatted into a human readable format by Slips' timeline module
- The flows button shows the raw flows as seen in the input file, of by zeek in case of running Slips on a PCAP or on you rinterface
- The outgoing button shows flow sent from this IP/profile to other IPs only. it doesn't show traffic sent to the profile.
- The incoming button shows flow sent to this IP/profile to other IPs only. It doesn't show traffic sent from the profile.
- The Alerts button shows the alerts Slips saw for this IP, each alert is a bunch of evidence that the given profile is malicious. Slips decides to block the IP if an alert is generated for it (if running with `-p`). Clicking on each alert expands the evidence that resulted in the alert.
- The Evidence button shows all the evidence of the timewindow whether they were part of an alert or not.

If you're running slips in docker you will need to add one of the following parameters to docker to be able to use the web interface:

Use the host network by adding `--net=host`, for example:

```
docker run -it --rm --net=host --cap-add=NET_ADMIN --name slips stratosphereips/  
↔slips:latest
```

Use port mapping to access the container's port to the host's port by adding `-p 55000:55000`, for example:

```
docker run -it --rm -p 55000:55000 --name slips stratosphereips/slips:latest
```

2.7 Saving the database

Slips uses redis to store analysis information. you can save your analysis for later use by running slips with `-s`,

For example:

```
sudo ./slips.py -f dataset/test7-malicious.pcap -s
```

Your `.rdb` saved database will be stored in the default output dir.

Note: If you try to save the same file twice using `-s` the old backup will be overwritten.

You can load it again using `-d`, For example:

```
sudo ./slips.py -d redis_backups/hide-and-see-short.rdb -w
```

Then navigate to `http://localhost:55000/` and select the entry on port 32850 to view the loaded database.

Note: saving and loading the database requires **root privileges** and is only supported in linux.

This feature isn't supported in docker due to problems with redis in docker.

DISCLAIMER: When saving the database you will see the following warning

```
stop-writes-on-bgsave-error is set to no, information may be lost in the redis backup.
↪ file
```

This configuration is set by slips so that redis will continue working even if redis can't write to `dump.rdb`.

Your information will be lost only if you're out of space and redis can't write to `dump.rdb` or if you don't have permissions to write to `/var/lib/redis/dump.rdb`, otherwise you're fine and the saved database will contain all analyzed flows.

2.8 Whitelisting

Slips allows you to whitelist some pieces of data in order to avoid its processing. In particular, you can whitelist an IP address, a domain, a MAC address or a complete organization. You can choose to whitelist what is going **to** them, what is coming **from** them, or both directions. You can also choose to whitelist the flows, so they are not processed, or alerts, so you see the flows but don't receive alerts on them. The idea of whitelisting is to avoid processing any communication to or from these iocs, not to avoid any packet that contains that ioc. For example, if you whitelist the flow of the domain `slack.com`, then a DNS request to the DNS server 1.2.3.4 asking for `slack.com` will still be shown.

This whitelist can be enabled or disabled by changing the `enable_local_whitelist` key in `config/slips.yaml`.

Do not modify the default `config/whitelist.conf` in place. Create a copy, update your copy, and set `whitelists.local_whitelist_path` in the Slips config file you are using to point to that copy.

Example:

```
cp config/whitelist.conf config/my_whitelist.conf
cp config/slips.yaml config/my_slips.yaml
```

Then set `local_whitelist_path: config/my_whitelist.conf` in `config/my_slips.yaml` and run Slips with:

```
./slips.py -c config/my_slips.yaml -f dataset/test7-malicious.pcap
```

The attacker and victim of every evidence are checked against the whitelist. In addition to all the related IPs, DNS resolutions, SNI, and CNAMEs of the attacker and the victim. If any of them are whitelisted, the flow/evidence is discarded.

Whitelists now use bloom filters to speed up the process of checking if an IoC is whitelisted or not.

2.8.1 Flows Whitelist

If you whitelist an IP address, Slips will check all flows and see if you are whitelisting to them or from them.

If you whitelist a domain, Slips will check:

- Domains in HTTP Host header
- Domains in the SNI field of TLS flows
- All the Domains in the DNS resolution of IPs (there can be many) (be careful that the resolution takes time, which means that some flows may not be whitelisted because Slips doesn't know they belong to a domain).
- Domains in the CN of the certificates in TLS

If you whitelist an organization, then:

- Every IP is checked against all the known IP ranges of that organization
- Every domain (SNI/HTTP Host/IP Resolution/TLS CN certs) is checked against all the known domains of that organization
- ASNs of every IP are verified against the known ASN of that organization

If you whitelist a MAC address, then:

- The source and destination MAC addresses of all flows are checked against the whitelisted mac address.

2.8.2 Alerts Whitelist

If you whitelist some piece of data not to generate alerts, the process is the following:

- If you whitelisted an IP
 - We check if the source or destination IP of the flow that generated that alert is whitelisted.
 - We check if the content of the alert is related to the IP that is whitelisted.
- If you whitelisted a domain
 - We check if any domain in alerts related to DNS/HTTP Host/SNI is whitelisted.
 - We check also if any domain in the traffic is a subdomain of your whitelisted domain. So if you whitelist 'test.com', we also match 'one.test.com'
- If you whitelisted an organization
 - We check that the ASN of the IP in the alert belongs to that organization.
 - We check that the range of the IP in the alert belongs to that organization.
- If you whitelist a MAC address, then:
 - The source and destination MAC addresses of all flows are checked against the whitelisted mac address.

2.8.3 Tranco whitelist

In order to reduce the number of false positive alerts, Slips uses Tranco whitelist which contains a research-oriented top sites ranking hardened against manipulation here <https://tranco-list.eu/>

Slips download the top 10k domains from this list and by default whitelists all evidence and alerts from and to these domains. Slips still shows the flows to and from these IoC.

The tranco list is updated daily by default in Slips, but you can change how often to update it using the `online_whitelist_update_period` key in `config/slips.yaml`.

Tranco whitelist can be enabled or disabled by changing the `enable_online_whitelist` key in `config/slips.yaml`.

2.8.4 Whitelisting Example

Do not edit the default `config/whitelist.conf` directly. Copy it, set `local_whitelist_path` in your copied Slips config file to the copied whitelist file, and modify that copied whitelist instead.

For example, your copied whitelist file can contain:

```
"IoCType", "IoCValue", "Direction", "IgnoreType"
ip,1.2.3.4,both,alerts
domain,google.com,src,flows
domain,apple.com,both,both
ip,94.23.253.72,both,alerts
ip,91.121.83.118,both,alerts
mac,b1:b1:b1:c1:c2:c3,both,alerts
organization,microsoft,both,both
organization,facebook,both,both
organization,google,both,both
organization,apple,both,both
organization,twitter,both,both
```

The values for each column are the following:

```
Column IoCType
- Supported IoCTypes: ip, domain, organization, mac
Column IoCValue
- Supported organizations: google, microsoft, apple, facebook, twitter.
Column Direction
- Direction: src, dst or both
  - Src: Check if the IoCValue is the source
  - Dst: Check if the IoCValue is the destination
  - Both: Check if the IoCValue is the source or destination
Column IgnoreType
- IgnoreType: alerts or flows or both
  - Ignore alerts: slips reads all the flows, but it just ignores alerting if
↳ there is a match.
  - Ignore flows: the flow will be completely discarded.
```

2.9 Popup notifications

Slips Support displaying popup notifications whenever there's an alert.

This feature is disabled by default. You can enable it by changing `popup_alerts` to `yes` in `config/slips.yaml`

This feature is supported in Linux and it requires root privileges.

This feature is supported in MaOS without root privileges.

This feature is not supported in Docker

2.10 Slips permissions

Slips doesn't need root permissions unless you

1. use the blocking modules (Firewall and ARP poisoner) (with -p)
2. use slips notifications
3. are saving the database (with -d)

If you can't listen to an interface without sudo, you can run the following command to let any user use zeek to listen to an interface not just root.

```
sudo setcap cap_net_raw,cap_net_admin=eip /<path-to-zeek-bin/zeek
```

Even when Slips is run using sudo, it drops root privileges in modules that don't need them.

2.11 Modifying the configuration file

Slips has a `config/slips.yaml` the contains user configurations for different modules and general execution. Below are some of Slips features that can be modified with respect to the user preferences.

Do not modify the default `config/slips.yaml` in place. Keep it as the shipped baseline, create a copy for your local changes, and run Slips with that copy using `-c`.

Example:

```
cp config/slips.yaml config/my_slips.yaml
./slips.py -c config/my_slips.yaml -f dataset/test7-malicious.pcap
```

2.11.1 Generic configuration

Time window width.

Each IP address that appears in the network traffic of the input is represented as a profile in Slips. Each profile is divided into time windows. Each time window is 1 hour long by default, and it gathers the network traffic and its behaviour for the period of 1 hour. The duration of the timewindow can be changed in the `config/slips.yaml` using

`time_window_width`

Analysis Direction

`analysis_direction` can either be `out` or `all`

The `analysis_direction` parameter controls which traffic flows are processed and analyzed by SLIPS. It determines whether SLIPS should focus on outbound traffic (potential data exfiltration), inbound traffic (incoming attacks), or analyze traffic in both directions. In `all` mode, SLIPS creates profiles for both internal (A) and external (B) IP addresses, and analyzes traffic in both directions (inbound and outbound).

In `out` mode, SLIPS still creates profiles for internal (A) and external (B) IP addresses, but only analyzes the outgoing traffic from the internal (A) profiles to external (B) destinations.

This parameter allows you to tailor SLIPS's analysis focus based on your specific monitoring requirements, such as detecting potential data exfiltration attempts (out mode) or performing comprehensive network monitoring in both directions (all mode).

Persistent runtime data

Use `permanent_dir` to choose where Slips stores databases and runtime-generated files that must persist across different Slips runs and should not be overwritten.

This includes persistent artifacts such as `p2p_trust_runtime/` and shared module databases like the Fides cache.

Live Slips auto update

Use `update.auto_update_slips` to enable or disable automatic live updates of the installed Slips version.

```
update:
  auto_update_slips: false
```

This setting is separate from the runtime `feeds_update_manager` module, which only updates TI feeds and related files.

Automatic Slips updates may overwrite the default config files shipped with Slips. If you want to keep local config changes safe, do not modify the default config files. Create and use your own config files with different names instead.

Use `update.channel_to_update_slips_from` in `slips.yaml` to choose the update channel:

- `stable` -> `origin/master`
- `unstable` -> `origin/develop`
- `testing` -> uses the branch specified in `update.testing_branch_to_update_slips_from`

2.11.2 Disabling a module

You can disable modules easily by appending the module name to the `disable` list.

The module name to disable should be the same as the name of its directory name inside `modules/` directory

2.11.3 ML Detection

The `mode=train` should be used to tell the `MLdetection1` module that the flows received are all for training.

The `mode=test` should be used after training the models, to test unknown data.

You should have trained at least once with 'Normal' data and once with 'Malicious' data in order for the test to work.

2.11.4 Blocking

This module is enabled only using the `-p` parameter and needs an interface to run.

Usage example:

```
sudo ./slips.py -i wlp3s0 -p
```

Slips needs to be run as root so it can execute iptables commands.

In Docker, since there's no root, the environment variable `IS_IN_A_DOCKER_CONTAINER` should be set to `True` to use 'sudo' properly.

If you use the latest Dockerfile, it will be set by default. If not, you can set it manually by running this command in the docker container

```
export IS_IN_A_DOCKER_CONTAINER=True
```

2.11.5 VirusTotal

In order for virustotal module to work, you need to add your VirusTotal API key to the file `config/vt_api_key`.

You can specify the path to the file with VirusTotal API key in the `api_key_file` variable.

The file should contain the key at the start of the first line, and nothing more.

If no key is found, virustotal module will not start.

2.11.6 Exporting Alerts

Slips can export alerts to different systems.

Refer to the [exporting section of the docs](#) for detailed instructions on how to export.

2.12 Logging

To enable the creation of log files, there are two options:

1. Running Slips with `verbose` and `debug` flags
2. Using `errors.log` and `running.log`

2.12.1 Zeek log files

You can enable or disable deleting zeek log files after stopping slips by setting `delete_zeek_files` to `yes` or `no`.

DISCLAIMER: zeek generates log files that grow every second until they reach GBs, to save disk space, Slips deletes all zeek log files after 1 day when running on an interface. this is called zeek rotation and is enabled by default.

You can disable rotation by setting `rotation` to `no` in `config/slips.yaml`

Check [rotation section](#) for more info

But you can also enable storing a copy of zeek log files in the output directory after the analysis is done by setting `store_a_copy_of_zeek_files` to `yes`, or while zeek is still generating log files by setting `store_zeek_files_in_the_output_dir` to `yes`. this option stores a copy of the zeek files present in `zeek_files/` the moment slips stops. so this doesn't include deleted zeek logs.

Once slips is done, you will find a copy of your zeek files in `<output_dir>/zeek_files/`

DISCLAIMER: Once slips knows you do not want a copy of zeek log files after slips is done by enabling `delete_zeek_files` and disabling `store_a_copy_of_zeek_files` parameters, it deletes large log files periodically (like `arp.log`).

2.12.2 Rotation of zeek logs

Rotation is done in zeek files to avoid growing zeek log files and save disk space.

Rotating is only enabled when running on an interface.

By default, Slips rotates zeek files every 1 day. this can be changed in `config/slips.yaml` by changing the value of `rotation_period`

`rotation_period` value can be written as a numeric constant followed by a time unit where the time unit is one of `usec`, `msec`, `sec`, `min`, `hr`, or `day` which respectively represent microseconds, milliseconds, seconds, minutes, hours, and days. Whitespace between the numeric constant and time unit is optional. Appending the letter `s` to the time unit in order to pluralize it is also optional. Check [Zeek rotation interval](#) for more details

Slips has an option to not delete the rotated zeek files immediately by setting the `keep_rotated_files_for` parameter.

This is equivalent to telling slips, Delete all logs that happened before the last `keep_rotated_files_for` days. `keep_rotated_files_for` value supports days only.

2.12.3 Running Slips with verbose and debug flags

We use two variables for logging, `verbose` and `debug`, they both range from 0 to 3.

Default value for `debug` is 0. so no errors are printed by default.

Default value for `verbose` is 1. so basic operations and proof of work are printed by default.

To change them, We use `-v` for verbosity and `-e` for debugging

For example:

```
./slips.py -c config/slips.yaml -v 2 -e 1 -f zeek_dir
```

Verbosity is about less or more information on the normal work of slips.

For example: “Done writing logs to file x.”

Debug is only about errors.

For example: “Error reading threat intelligence file, line 4, column 2”

To more verbosity level, the more detailed info is printed.

The more debug level, the more errors are logged.

Below is a table showing each level of both.

2.12.4 Using errors.log and running.log

Slips also logs all errors to `output/errors.log` (interactive mode), and `/var/log/slips/error.log` (daemonized mode) whether you’re using `-e` or not. See [Daemonized vs interactive](#) for more information about the 2 modes

General slips logs are created in `/var/log/slips/running.log` in case of daemonized mode and `...#TODO` in case of interactive mode

2.13 Reading Input from stdin

Slips supports reading input flows in text format from stdin in interactive mode.

Supported flow from stdin are zeek conn.log files (json form), suricata and argus.

For example, you can run slips using:

```
./slips.py -f zeek
```

or

```
./slips.py -f suricata
```

and you once you see the following line:

```
[InputProcess] Receiving flows from stdin.
```

you can start giving slips flows in you desired format.

All zeek lines taken from stdin should be in json form and are treated as conn.log lines.

This feature is specifically designed to allow slips to interact with network simulators and scripts.

2.14 Slips as an access point

Slips supports monitoring two interface with Zeek when it's running as an access point.

To use this feature, simply pass the `--access-point` or `-ap` argument followed by a comma-separated list of interfaces to monitor when running slips.

The wifi interface should be listed first, followed by the ethernet interface.

```
./slips.py --access-point wlan0,eth0
```

or

```
./slips.py -ap wlan0,eth0
```

Slips will produce zeek logs in two separate directories inside your output directory: `zeek_files/eth0/` and `zeek_files/wlan0/`. and will be able to detect the host IP, gateway IP and zeek logs of each interface separately.

2.15 Plug in a zeek script

Slips supports automatically running a custom zeek script by adding it to `zeek-scripts` dir and adding the file name in `zeek-scripts/__load__.zeek`.

For example, if you want to add a zeek script called `arp.zeek` you should add it to `__load__.zeek` like this:

```
@load ./arp.zeek
```

Zeek output is suppressed by default, so if your script has errors, Slips will fail silently.

2.16 Exporting strato letters

Exporting the strato letters can be done by enabling the `export_strato_letters` option in `config/slips.yaml`. once enabled, Slips will export the strato letters to `strato_letters.tsv` in the output directory. this file can be used for training Slips RNN module.

2.17 Slips parameters

- `-c` or `--config` Used for changing then path to the Slips config file. default is `config/slips.yaml`. It is recommended to copy `config/slips.yaml` and pass your copy with `-c` instead of editing the default file.
- `-v` or `--verbose` Verbosity level. This logs more info about Slips.
- `-e` or `--debug` Debugging level. This shows more detailed errors.
- `-f` or `--filepath` Read and automatically recognize a Zeek dir, a Zeek conn.log file, a Suricata JSON file, Argus, PCAP.
- `-i` or `--interface` Read packets from an interface.
- `-l` or `--createlogfiles` Create log files with all the traffic info and detections.
- `-F` or `--pcapfilter` Packet filter for Zeek. BPF style.
- `-cc` or `--clearcache` Clear the cache database.
- `-p` or `--blocking` Allow Slips to block malicious IPs. Requires root access. Supported only on Linux.
- `-cb` or `--clearblocking` Flush and delete slipsBlocking iptables chain
- `-o` or `--output` Store alerts.json and alerts.txt in the given folder.

- `-s` or `--save` Save the analysed file db to disk.
- `-d` or `--db` Read an analysed file (rdb) from disk.
- `-D` or `--daemon` Run slips in daemon mode
- `-S` or `--stopdaemon` Stop slips daemon
- `-k` or `--killall` Kill all unused redis servers
- `-m` or `--multiinstance` Run multiple instances of slips, don't overwrite the old one
- `-P` or `--port` The redis-server port to use
- `-g` or `--growing` Treat the given zeek directory as growing. eg. zeek dirs generated when running onan interface
- `-w` or `--webinterface` Start Slips web interface automatically
- `-V` or `--version` Used for checking your running Slips version flags.
- `-im` or `--input-module` Used for reading flows from a module other than input process.
- `-ap` or `--access-point` Used for reading packets from two interfaces when Slips is running as an access point.

2.18 Limiting Slips resource consumption

When given a very a large pcap, slips may use more memory/CPU than it should. to fix that you can reduce the niceness of Slips by running:

```
renice -n 6 -p <Slips-PID>
```

2.19 Running Slips from python

You can run Slips from python using the following script

```
import subprocess
command = './slips.py -f dataset/test3-mixed.binetflow -o /data/test'
args = command.split()
process = subprocess.run(args, stdout=subprocess.PIPE)
```


ARCHITECTURE

The architecture of Slips is basically: - To receive some data as input - To process it to a common format - To enrich it (gather all possible info about the IPs/MAC/User-Agents etc.) - To apply detection modules - To output results

Slips is heavily based on the Zeek monitoring tool as input tool for packets from the interface and pcap file, due to its excellent recognition of protocols and easiness to identify the content of the traffic.

Figure 1 shows how the data is analyzed by Slips. As we can see, Slips internally uses Zeek, an open source network security monitoring tool. Slips divides flows into profiles and each profile into a timewindows, timewindows are numbered from 1 to infinity. Slips runs detection modules on each flow and stores all evidence, alerts and features in an appropriate profile structure. All profile info, performed detections, profiles and timewindows' data, is stored inside a Redis database. All flows are read, interpreted by Slips, labeled, and stored in the SQLite database in the output/ dir of each run. The output of Slips is a folder with logs (output/ directory) that has alert.json, alerts.log, errors.log. Kalipso, a terminal graphical user interface. or the Web interface.

Below is more explanation on internal representation of data, usage of Zeek and usage of Redis inside Slips.

3.1 Internal representation of data.

Slips works at a flow level, instead of a packet level, gaining a high level view of behaviors. Slips creates traffic profiles for each IP that appears in the traffic. A profile contains the complete behavior of an IP address. Each profile is divided into time windows. Each time window is 1 hour long by default and contains dozens of features computed for all connections that start in that time window. Detections are done in each time window, allowing the profile to be marked as uninfected in the next time window.

This is what slips stores for each IP/Profile it creates:

- Ipv4 - ipv4 of this profile
- IPv6 - list of ipv6 used by this profile
- Threat_level - the threat level of this profile, updated every TW.
- Confidence - how confident slips is that the threat level is correct
- Past threat levels - history of past threat levels
- Used software - list of software used by this profile, for example SSH, Browser, etc.
- MAC and MAC Vendor - Ether MAC of the IP and the name of the vendor
- Host-name - the name of the IP
- first User-agent - First UA seen use dby this profile.
- OS Type - Type of OS used by this profile as extracted from the user agent
- OS Name - Name of OS used by this profile as extracted from the user agent

- Browser - Name of the browser used by this profile as extracted from the user agent
- User-agents history - history of the all user agents used by this profile
- DHCP - if the IP is a dhcp or not
- Starttime - epoch formatted timestamp of when the profile first appeared
- Duration - the standard duration of every TW in this profile
- Modules labels - the labels assigned to this profile by each module
- Gateway - if the IP is the gateway (router) of the network
- Timewindow count - Amount of timewindows in this profile
- ASN - autonomous service number of the IP
- Asnorg - name of the org that own the ASN of this IP
- ASN Number
- SNI - Server name indicator
- Reverse DNS - name of the IP in reverse dns
- Threat Intelligence - If the IP appeared in any of Slips blacklist
- Description - Description of this IP as taken from the blacklist
- Blacklist Threat level - threat level of the blacklisted that has this IP
- Passive DNS - All the domains that resolved into this IP
- Certificates - All the certificates that were used by this IP
- Geocountry - Country of this IP
- VirusTotal - contains virustotal scores of this IP
 - Down_file: files in virustotal downloaded from this IP
 - Ref_file: files in VT that referenced this IP
 - Com_file : files in VT communicating with this IP
 - Url ratio: The higher the score the more malicious this IP is

3.2 Alerts vs Evidence

When running Slips, the alerts you see in red in the CLI or the web interface are a bunch of evidence. Evidence in slips are detections caused by a specific IP in a specific timeframe. Slips doesn't alert on every evidence/detection. it accumulates evidence and only generates and alert when the amount of gathered evidence crosses a threshold. After this threshold Slips generates an alert, marks the timewindow as malicious in the web interface and blocks the IP causing the alert if iptables is enabled.

Each alert has a threat level and confidence; the Threat level of each alert is Critical by default, and the confidence is the accumulated threat level of all the evidence of the alert normalized to a value ranging from 0 to 1. The more evidence the higher the confidence of the alert.

NOTE: When slips is run with -ap, evidence are tied to one interface, meaning each evidence belongs to a flow that belongs to one interface. Meanwhile an alert contains evidence from different interfaces.

3.3 Usage of Zeek.

Slips uses Zeek to generate files for most input types, and this data is used to create the profiles. For example, Slips uses this data to create a visual timeline of activities for each time window. This timeline consists of Zeek generated flows and additional interpretation from other logs like dns log and http log.

3.4 Usage of Redis database.

All the data inside Slips is stored in Redis, an in-memory data structure. Redis allows all the modules in Slips to access the data in parallel. Apart from read and write operations, Slips takes advantage of the Redis messaging system called Redis PUB/SUB. Processes may publish data into the channels, while others subscribe to these channels and process the new data when it is published.

When Slips uses a Redis port other than the default Redis port 6379, it closes the opened redis server when analysis is done, and keeps a copy of the Redis database for later analysis in `output_dir/databases/dump.rdb`. To inspect this database, start a Redis server with that RDB file and connect to the same port:

```
redis-server --port 6380 --dir output_dir/databases --dbfilename dump.rdb
redis-cli -p 6380
```

3.5 Usage of SQLite database.

Slips uses SQLite database to store all flows in Slips interpreted format. The SQLite database is stored in the `output/ dir` and each flow is labeled to either 'malicious' or 'benign' based on slips detections. all the labeled flows in the SQLite database can be exported to tsv or json format.

3.6 Threat Levels

Slips has 5 threat levels.

3.7 How Slips Stops

- When slips is running on an interface or a growing zeek directory, slips keeps running forever until the user presses ctrl+c
- When Slips is analyzing a PCAP or a zeek directory or any other supported file, It keeps running until no more flows are received.
- After the modules receive that signal that says “no more new flows are coming”, all modules keep processing the existing flows normally until they run out of msgs and stop.
- Modules stop only if no more msgs are received in their Redis channels, and if they receive the signal that slips is no longer receiving new flows.
- Slips knows that no more flows are arriving when it reaches the end of the given zeek/suricata/nfdump logs.
- If some processes are hanging in memory, slips wait by default 1 week before killing them. This can be modified in the `config.yaml`.

For more technical details about this check <https://stratospherelinuxips.readthedocs.io/en/develop/contributing.html#faq>

DETECTION MODULES

Slips is a behavioral-based IPS that uses machine learning to detect malicious behaviors in the network traffic. It is a modular software that can be extended. When Slips is run, it spawns several child processes to manage the I/O, to profile attackers and to run the detection modules.

Here we describe what detection modules are run on the traffic to detect malicious behaviour.

Modules are Python-based files that allow any developer to extend the functionality of Slips. They process and analyze data, perform additional detections and store data in Redis for other modules to consume. Currently, Slips has the following modules:

4.1 Bruteforcing Module

The `Bruteforcing` module is responsible for SSH bruteforcing detection.

It consumes:

- `ssh.log`
- `software.log`
- `notice.log`

It correlates repeated SSH sessions by source IP, destination IP, destination port, and time window. It starts alerting at 9 attempts by default, reports sparsely as the count grows, uses the SSH client banner to adjust confidence, and uses Zeek `SSH::Password_Guessing` notices as confirmation.

For the full design and configuration details, see:

- *Bruteforcing Module*

4.2 HTTPS Anomaly Detection Module

For the full technical description of the HTTPS anomaly detector (features, training, adaptation, z-score logic, evidence format, and configuration), see:

- *HTTPS Anomaly Detection*

The module ships with a local HTML report generator that visualizes detections, confidence levels, and top anomaly reasons from `anomaly_detection_https.log`.

4.2.1 Virustotal Module

This module is used to lookup IPs, domains, and URLs on virustotal.

To use it you need to add your virustotal API key in `config/vt_api_key`

4.2.2 RiskIQ Module

This module is used to get different information (passive DNS, IoCs, etc.) from RiskIQ To use this module your RiskIQ email and API key should be stored in `config/RiskIQ_credentials`

the format of this file should be the following:

```
example@domain.com
e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855
```

The hash should be your 64 character API Key.

The path of the file can be modified by changing the `RiskIQ_credentials_path` parameter in `config/slips.yaml`

4.3 RNN C&C Detection Module

This module is used to detect command and control channels in a network by analyzing the features of the network flows and representing them as Stratosphere Behavioral Letters(Stratoletters). This is achieved through the use of a recurrent neural network, which is trained on these letters to identify and classify potential C&C traffic.

4.3.1 Strato letters

Stratoletters is a method used to represent network flows in a concise and standardized manner. Stratoletters encodes information about the periodicity, duration, and size of network flows into a string of letter(s) and character(s).

The **letter** is the key part in a Stratoletter string. It is derived from a dictionary and defined based on the features of the flow such as periodicity, size and duration.

periodicity : A number that denotes how frequent the flow is, calculated based on time of past flows

-1 = No previous data

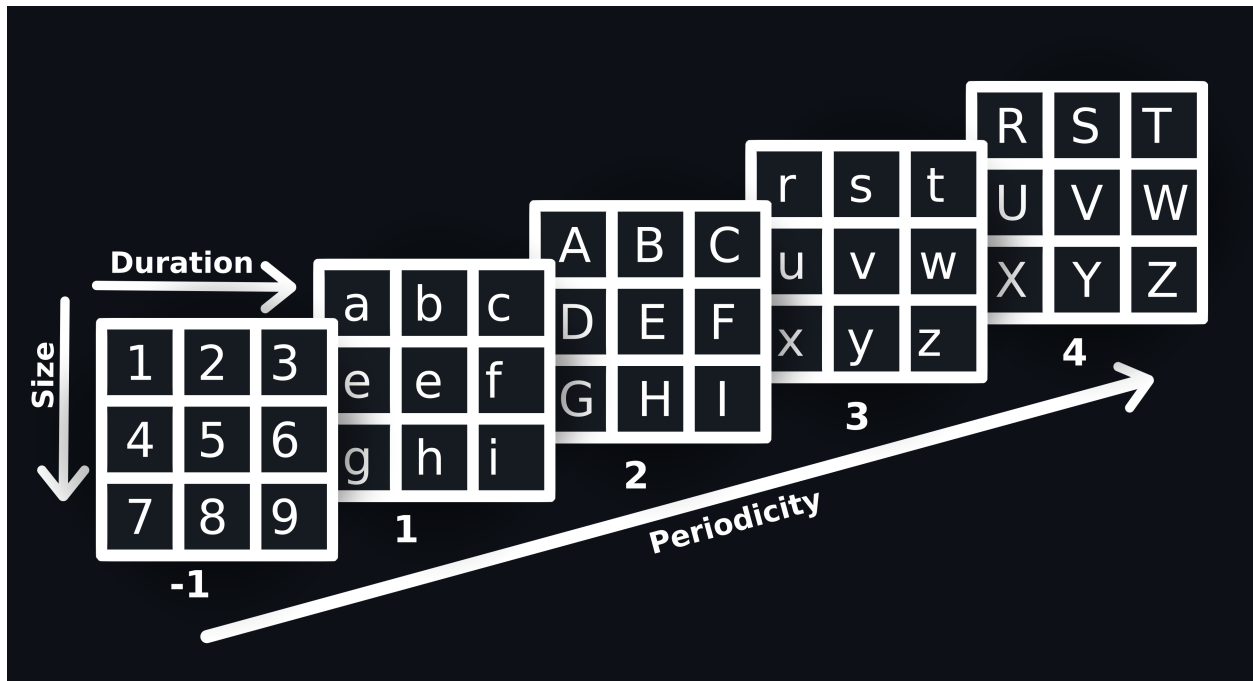
1-4 = 1 strongly periodic to 4 strongly not periodic

size : Number denotes the size of flow, range from 1 to 3

duration : Number that denotes the duration of flow, range from 1 to 3

A visual representation of the dictionary from which letter is derived

In this image, each block represents the possible values of the letter, We choose the block based on the periodicity and choose the letter from the block based on the duration(number of row) and size (number of column).



Example:

```
# Slips computed value of the flow
periodicity = 1      # Strongly periodic
duration = 1
size = 3
letter = g          # lowercase letter for periodicity 1(Strongly periodic) and 3(Weakly_
↳not periodic)
```

```
periodicity = -1   # no previous flow data
duration = 2
size = 3
letter = 8         # the letter will be an integer if there is no previous data
```

```
periodicity = 4    # Weakly not periodic
duration = 3
size = 3
letter = Z         # uppercase letter for periodicity 2(Weakly periodicity) and_
↳4(Strongly not periodic)
```

Stratoletters represent details of current flow and past available flow with latest on left. The current flow symbol is concise of three parts.

each symbol consists of hrs passed since last flow + a letter that represents the periodicity, size and dur of the flow + a char showing the time passed since last flow

```
symbol = zeros + letter + timechar
```

zero : hours passed since last flow, each hour is represented by 1 zero. for example, 2 hours = 00

letter : chosen based on the periodicity, size, and dur of the flow eg: 1,w,H

timechar : character to denote the time elapsed since last flow, can be: ., ,, +, * or null

Ultimately this is how a Stratoletter is formed

```
No of hours passed since last flow = 2
periodicity = 2      # Weakly not periodicity
duration = 1
size = 1
timechar =
stratoletter of last flow = 9*z*
letter = A
symbol = 00A9*z*

No of hours passed since last flow = 0
periodicity = 3      # Weakly not periodic
duration = 1
size = 2
timechar = *
stratoletter of last flow = e.
letter = u
symbol = u*e.
```

Then the model will predict how secure each flow is based on the Stratoletter

```
symbol = 99*z*i.i*
model_score = 0.9573354
symbol = 99.
model_score = 0.9063127
symbol = 77*g.g*g*g.g.g.g*x*x*x*g.g.
model_score = 0.96772265
```

In first example **9** is Stratoletter of current flow. **9*** is previous one, **z*** is before that and so on.

4.4 Leak Detection Module

This module on runs on pcaps, it uses YARA rules to detect leaks.

You can add your own YARA rule in `modules/leak_detector/yara_rules/rules` and it will be automatically compiled and stored in `modules/leak_detector/yara_rules/compiled` and matched against every pcap.

4.5 Blocking Module

Blocking in Slips is done for any IP that results in an alert. If an IP is detected as malicious and is blocked, it stays blocked forever, unless it is unblocked manually.

The feature of unblocking IPs after a while is not supported yet.

The blocking is done using iptables, and the blocked IPs are stored in the database for future reference.

Blocking is disabled by default. To enable blocking in slips, start slips with the `-p` flag.

This feature is only supported in linux using iptables when running on an interface.

4.6 Exporting Alerts Module

Slips supports exporting alerts to other systems using different modules (ExportingAlerts, CESNET sharing etc.)

For now the supported systems are:

- Slack
- TAXII Servers (STIX format)
- Warden servers
- suricata-like JSON format
- Logstash

Refer to the [exporting section of the docs](#) for detailed instructions on how to export.

4.7 Flowalerts Module

This module is responsible for detecting malicious behaviours in your traffic.

Refer to the [Flowalerts section of the docs](#) for detailed explanation of what Slips detects and how it detects.

4.8 Disabled alerts

All Slips detections are turned on by default, You can configure which alerts you want to enable/disable in `config/slips.yaml`

Slips support disabling unwanted alerts, simply add the detection you want to disable in the `disabled_detections` list and slips will not generate any alerts of this type.

for example:

```
disabled_detections = [MaliciousJA3, DataExfiltration, SelfSignedCertificate]
```

Supported detections are:

ARPScan, ARP-outside-localnet, UnsolicitedARP, MITM-ARP-attack, SSHSuccessful, LongConnection, MultipleReconnectionAttempts, ConnectionToMultiplePorts, InvalidCertificate, UnknownPort, Port0Connection, ConnectionWithoutDNS, DNSWithoutConnection, MaliciousJA3, DataExfiltration, SelfSignedCertificate, VerticalPortscan, HorizontalPortscan, Password_Guessing, MaliciousFlow, SuspiciousUserAgent, multiple_google_connections, NETWORK_gps_location_leaked, Command-and-Control-channels-detection, ThreatIntelligenceBlacklistDomain, ThreatIntelligenceBlacklistIP, MaliciousDownloadedFile, DGA, MaliciousSSLCert, YoungDomain, MultipleSSHVersions, DNS-ARPA-Scan, SMTPLoginBruteforce, BadSMTPLogin, IncompatibleUserAgent, ICMP-Timestamp-Scan, ICMP-AddressScan, ICMP-AddressMaskScan

4.9 Threat Intelligence Module

Slips has a complex system to deal with Threat Intelligence feeds. The Threat Intelligence Module in Slips is designed to enhance detection capabilities, utilizing both external and internal thread intelligence sources.

Slips supports various indicators of compromise (IoCs) from Threat Intelligence (TI) feeds, including IPs, IP ranges, domains, JA3 hashes, and SSL hashes. This provides comprehensive coverage against potential threats

While file hashes and URLs aren't supported in TI feeds directly, Slips compensates by integrating with specialized external services for these types of IoCs.

4.9.1 External Threat Intelligence Services

Besides searching 40+ TI files for every IP/domain Slips encounters, Slips integrates with the following external threat intelligence services to enrich its detection capabilities:

URLhaus: This service is utilized for checking URLs observed in `http.log` and files observed in `files.log` against known malicious URLs and files. URLhaus provides a comprehensive database of malicious URLs, which Slips queries to determine if observed URLs or files are associated with known malware or phishing campaigns.

Spamhaus: Spamhaus is used for inbound traffic IP lookups to assess the reputation of IP addresses encountered during the analysis. By querying Spamhaus, Slips can identify IP addresses associated with spamming activities, botnets, and other malicious behaviors, enhancing its ability to detect and alert on suspicious network traffic.

Circl.lu: Circl.lu's service is leveraged for hash lookups, particularly for downloaded files. Each file hash extracted from `files.log` is checked against Circl.lu's extensive database of known malicious file hashes. This integration allows Slips to identify and react to the transfer or presence of known malicious files within the monitored network environment.

Circl.lu returns scores (`hashlookup:trust`) for each md5. this score ranges from 0 to 100, with 0 being the most malicious. This score is converted to a value that slips can deal with using the following equation

```
malicious_percentage = 100 - circll_score
# scale the benign percentage from 0 to 1
threat_level = float(malicious_percentage) / 100
```

And then it's converted to a string threat level using the table in <https://stratospherelinuxips.readthedocs.io/en/develop/architecture.html#threat-levels>

URLhaus Access:

- **Purpose:** Identify malicious URLs and files.
- **Method:** Slips queries the URLhaus API with URLs and file hashes observed in network traffic logs.
- **Response Handling:** If a URL or file is found in the URLhaus database, Slips generates an alert indicating the presence of a known threat.

Spamhaus Access:

- **Purpose:** Assess the reputation of IP addresses.
- **Method:** IP addresses are queried against Spamhaus's DNSBL (DNS-based Block List).
- **Response Handling:** Slips interprets the DNSBL response to determine if an IP address is associated with known malicious activities, triggering alerts accordingly.

Circl.lu Access:

- **Purpose:** Perform hash lookups for downloaded files.
- **Method:** File hashes are checked against Circl.lu's database via their API.
- **Response Handling:** Matches with known malicious hashes result in the generation of alerts to inform about potential threats.

By integrating these external services, Slips significantly enhances its detection capabilities, allowing for real-time alerting on threats identified through global intelligence feeds. This integration not only broadens the scope of detectable threats but also contributes to the overall security posture by enabling proactive responses to emerging threats.

4.9.2 Matching of IPs

Slips gets every IP it can find in the network (DNS answers, HTTP destination IPs, SSH destination IPs, etc.) and tries to see if it is in any blacklist.

If a match is found, it generates an evidence, if no exact match is found, it searches the Blacklisted ranges taken from different TI feeds.

4.9.3 Matching of Domains

Slips gets every domain that can find in the network and tries to see if it is in any blacklist. The domains are currently taken from:

- DNS requests
- DNS responses
- HTTP host names
- TLS SNI

Once a domain is found, it is verified against the downloaded list of domains from the blacklists defined in `ti_files` path in the configuration file `config/slips.yaml`. which is `config/TI_feeds.csv` by default. If an exact match is found, then an evidence is generated.

If an exact match is not found, then Slips verifies if the verified domain is a subdomain of any domain in the blacklist.

For example, if the domain in the traffic is `here.testing.com`, Slips first checks if the exact domain `here.testing.com` is in any blacklist, and if there is no match, it checks if the domain `testing.com` is in any blacklists too.

Slips also sets evidence about DNS answers of blacklisted domains. If `test.com` is blacklisted, and `test.com` resolves to 1.2.3.4, slips sets an evidence when it encounters the DNS answer with 1.2.3.4.

4.9.4 Matching of JA3 Hashes

Every time Slips encounters an TLS flow, it compares each JA3 and JA3s with the feeds of malicious JA3 and alerts when there's a match. Slips is shipped with the `Abuse.ch` JA3 feed by default You can add your own SSL feed by appending to the `ja3_feeds` key in `config/slips.yaml`

4.9.5 Matching of SSL SHA1 Hashes

Every time Slips encounters an SSL flow, it tries to get the certificate hash from zeek ssl.log, then it compares the hash with our list of blacklisted SSL certificates

Slips is shipped with the `Abuse.ch` SSL feed by default,

You can add your own SSL feed by appending to the `ssl_feeds` key in `config/slips.yaml`

4.9.6 Matching of ASNs

Every time Slips sees a new IP, it stores info about it in the db, for example its organization, RDNs, and ASN. If the ASN of an IP matches a blacklisted ASN, slips alerts.

Blacklisted ASNs are read from our local TI file `config/local_data_files/own_malicious_iocs.csv`, so you can update them or add your own.

4.9.7 Local Threat Intelligence files

Slips has a local file for adding IoCs of your own, it's located in `config/local_data_files/own_malicious_iocs.csv` by default, this path can be changed by changing `download_path_for_local_threat_intelligence` in `config/slips.yaml`

The format of the file is "IP address/IP Range/domain/ASN", "Threat level", "Description"

Threat level available options: info, low, medium, high, critical

Refer to the [architecture section of the docs](#) for detailed explanation of Slips threat levels.

Example:

```
"23.253.126.58", "high", "Simda CC"
"bncv00.no-ip.info", "critical", "Variant.Zusy"
```

4.9.8 Local JA3 hashes

Slips has a local file for adding JA3 hashes of your own, it's located in `config/local_data_files/own_malicious_JA3.csv` by default.

The format of the file is "JA3 hash", "Threat level", "Description"

Threat level available options: info, low, medium, high, critical

Refer to the [architecture section of the docs](#) for detailed explanation of Slips threat levels.

Example:

```
"e7d705a3286e19ea42f587b344ee6865", "medium", "Standard tor client"
"6734f37431670b3ab4292b8f60f29984", "high", "Trickbot Malwar"
```

4.9.9 Whitelisting known FP hashes

To avoid false positive "Malicious downloaded file" detections, before looking up MD5 hashes of each downloaded file online, Slips checks if the given hash is part of a known FP.

The list of known FP MD5 hashes is at `config/local_ti_files/known_fp_md5_hashes.csv`. This list is taken from <https://github.com/Neo23x0/ti-falsepositives/tree/master>

If the hash is a part of that list, Slips doesn't look it up.

4.9.10 Adding your own remote feed

We update the remote ones regularly. The list of remote threat intelligence files is set in the path of `ti_files` variable in `config/slips.yaml`. The path of all the TI feeds is in `config/TI_feeds.csv` by default.

You can add your own remote threat intelligence feeds in this variable. Supported extensions are: `.txt`, `.csv`, `.netset`, `ipsum` feeds, or `.intel`.

Each URL should be added with a `threat_level` and a tag, the format is (url,threat_level,tag)

tag is which category is this feed e.g. phishing, adtrackers, etc..

Threat level available options: info, low, medium, high, critical

Refer to the [architecture section of the docs](#) for detailed explanation of Slips threat levels.

TI files commented using # may be processed as they're still in our database.

Use ; for commenting TI files in `config/slips.yaml` instead of #.

Commented TI files (lines starting with ;) will be completely removed from our database.

The remote files are downloaded to the path set in the `download_path_for_local_threat_intelligence`. By default, the files are stored in the Slips directory `modules/ThreatIntelligence1/remote_data_files/` are deleted after slips is done reading them.

Domains found in remote feeds are considered invalid, and therefore discarded by Slips, if they have suffix that doesn't exist in https://publicsuffix.org/list/public_suffix_list.dat

4.9.11 Commenting a remote TI feed

If you have a remote file link that you wish to comment and remove from the database you can do so by adding ';' to the line that contains the feed link in `config/TI_feeds.csv`, don't use the '#' for example to comment the `bruteforcelogin` feed you should do the following:

```
;https://lists.blocklist.de/lists/bruteforcelogin.txt, medium, ['honeypot']
```

instead of:

```
#https://lists.blocklist.de/lists/bruteforcelogin.txt,medium,['honeypot']
```

4.10 Update Manager Module

To make sure Slips is up to date with the most recent IoCs in all feeds, all feeds are loaded, parsed and updated periodically and automatically by Slips every 24 hours, which requires no user interaction.

The 24 hours interval can be changed by changing the `TI_files_update_period` key in `config/slips.yaml`

Update manager is responsible for updating all remote TI files (including SSL and JA3 etc.)

By default, local slips files (`organization_info`, `ports_info`, etc.) are cached to avoid loading and parsing

then everytime we start slips. However, they are updated automatically by the update manager if they were changed on disk.

Only one slips instance is allowed to be using the update manager at a time to avoid race conditions.

By default, slips starts without the TI files, and runs the Update Manager in the background if the `wait_for_TI_to_finish` option in `slips.yaml` is set to yes, slips will not start until the update manager is done

and all TI files are loaded successfully, this is useful if you want to ensure that slips doesn't miss the detection of any blacklisted IPs, but it adds some time to the startup of slips since it will be downloading, parsing, and caching 45+ different TI feeds.

4.11 IP Info Module

The IP info module has several ways of getting information about an IP address, it includes:

- ASN
- Country by Geolocation
- Given a MAC, its Vendor
- Reverse DNS

4.11.1 ASN

Slips is shipped with an offline database (GeoLite2) in `databases/GeoLite2-ASN.mmdb` to search for ASNs, if the ASN of a given IP is not in the GeoLite2 database, we try to get the ASN online using the online database using the `ipwhois` library. However, to reduce the amount of requests, we retrieve the range of the IP and we cache the whole range. To search and cache the whole range of an IP, the module uses the `ipwhois` library. The `ipwhois` library gets the range of this IP by making a connection to the server `cymru.com` using a TXT DNS query. The DNS server is the one set up in the operating system. For example to get the ASN of the IP 13.32.98.150, you will see a DNS connection asking for the TXT record of the domain `150.98.32.13.origin.asn.cymru.com`.

4.11.2 Country by Geolocation

Slips is shipped with an offline database (GeoLite2) in `databases/GeoLite2-Country.mmdb` to search for Geolocation.

4.11.3 Mac Vendors

Slips is shipped with an offline database `databases/macaddress-db.json` for MAC address vendor mapping.

Slips updates this database by default every 2 weeks using the following online db

<https://maclookup.app/downloads/json-database/get-db?t=22-08-19&h=d1d39c52de447a7e7194331f379e1e99f94f35f1>

You can change how often this db is updated by changing the value of `mac_db_update` in `config/slips.yaml`.

Slips gets the MAC address of each IP from `dhcp.log` and `arp.log` and then searches the offline database using the OUI.

If the vendor isn't found in the offline MAC database, Slips tries to get the MAC using the online database <https://www.macvendorlookup.com>

The offline database is updated manually and shipped with slips, you can find it in the `databases/` dir.

Slips makes sure it doesn't perform duplicate searches of the same MAC Address either online, or offline.

4.12 Reverse DNS

This is obtained by doing a standard `in-addr.arpa` DNS request.

4.13 ARP Module

This module is used to check for ARP attacks in your network traffic.

By default, zeek doesn't generate and log ARP flows, but Slips is shipped with it's own zeek scripts that enable the logging of ARP flows in `arp.log`

The detection techniques are:

- ARP scans
- ARP to a destination IP outside of local network
- Unsolicited ARP
- MITM ARP attack

4.13.1 ARP Scans

Slips considers an IP performing an ARP scan if it sends 5 or more non-gratuitous ARP to different destination addresses in 30 seconds or less.

4.13.2 ARP to a destination IP outside of local network

Slips alerts when an ARP flow is being sent to an IP outside of local network as it's a weird behaviour that shouldn't be happening.

4.13.3 Unsolicited ARP

Unsolicited ARP is used to update the neighbours' ARP caches but can also be used in ARP spoofing, we detect it with threat level 'info', so we don't consider it malicious, we simply notify you about it.

4.13.4 MITM ARP attack

Slips detects when a MAC with IP A, is trying to tell others that now that MAC is also for IP B (ARP cache attack)

4.14 CESNET sharing Module

This module is responsible for importing and exporting alerts from and to warden server

Refer to the [exporting section of the docs](#) for detailed instructions on CESNET exporting and the format of the configuration files.

To enable the importing alerts from warden servers, set `receive_alerts` to `yes` in `config/slips.yaml`

Slips imports 100 alerts from warden servers each day, and automatically stores the IoCs in our database

Time to wait before receiving alerts from warden server is 1 day by default, you can change this by changing the `receive_delay` in `config/slips.yaml`

These are the categories Slips imports: ['Availability', 'Abusive.Spam','Attempt.Login', 'Attempt', 'Information', 'Fraud.Scam', 'Information', 'Fraud.Scam']

4.15 HTTP Analyzer Module

This module handles the detections of HTTP flows

Available detection are:

- Multiple empty connections
- Suspicious user agents
- Incompatible user agents
- Multiple user agents
- Pastebin downloads
- Unencrypted HTTP traffic
- Non-HTTP connections on port 80.

4.15.1 Multiple empty connections

Due to the usage of empty connections to popular site by malware to check for internet connectivity, We consider this type of behaviour suspicious activity that shouldn't happen

We detect empty connection to 'bing.com', 'google.com', 'yandex.com', 'yahoo.com', 'duckduckgo.com' etc.

If Google is whitelisted in `whitelist.conf`, this detection will be suppressed.

4.15.2 Suspicious user agents

Slips has a list of suspicious user agents, whenever one of them is found in the traffic, slips generates and evidence.

Our current list of user agents has: ['httpsend', 'chm_msdn', 'pb', 'jndi', 'tesseract']

4.15.3 Incompatible user agents

Slips uses and offline MAC address database to detect the type of device based on the MAC OUI.

First, Slips store the MAC address and vendor of every IP it sees (if available)

Second, When slips encounters a user agent in HTTP traffic it performs an online query to <http://useragentstring.com> to get more info about this user agent, like the os type, name and browser.

Third, When slips has both information available (MAC vendor and user agent), it compares them to detect incompatibility using a list of keywords for each operating system.

Available keywords for Apple: ('macos', 'ios', 'apple', 'os x', 'mac', 'macintosh', 'darwin')

Available keywords for Microsoft: ('microsoft', 'windows', 'nt')

Available keywords for Android: ('android', 'google')

4.15.4 Multiple user agents

Slips stores the MAC address and vendor of every IP it sees (if available) in the redis database. Then, when an IP is seen using a different user agent than the one stored in the database, it tries to extract os info from the user agent string, either by performing an online query to <http://useragentstring.com> or by using `zeek`.

If an IP is detected using different user agents that refer to different operating systems, an alert of type 'Multiple user agents' is made

for example, if an IP is detected using a macOS user agent then an android user agent, slips detects this with 'low' threat level

4.15.5 Pastebin downloads

Some malware use pastebin as the host of their malicious payloads.

Slips detects downloads of files from pastebin through HTTP with size \geq 700 bytes.

This value can be customized in `slips.yaml` by changing `pastebin_download_threshold`

When found, slips alerts pastebin download with threat level low because not all downloads from pastebin are malicious.

4.15.6 Unencrypted HTTP traffic

When slip sees an HTTP unencrypted traffic in `zeek`'s `http.log` it generates an evidence with `threat_level` low

4.15.7 Non-HTTP connections on port 80

Slips detects established connections on port 80 that are not using HTTP using zeek's conn.log flows

if slips finds a flow using destination port 80 and the 'service' field in conn.log isn't set to 'http', it means zeek didn't recognize that flow as http. Slips makes sure no matching flows were detected as HTTP by zeek within 5 mins before or after the given flow. if not, slips sets an evidence saying "non http established conn on port 80"

4.16 Leak Detector Module

This module work only when slips is given a PCAP

The leak detector module uses YARA rules to detect leaks in PCAPs

4.16.1 Module requirements

In order for this module to run you need:

using `sudo apt install yara`

You can install tshark by running

`sudo apt install wireshark`

4.16.2 How it works

This module works by

1. Compiling the YARA rules in the `modules/leak_detector/yara_rules/rules/` directory
2. Saving the compiled rules in `modules/leak_detector/yara_rules/compiled/`
3. Running the compiled rules on the given PCAP
4. Once we find a match, we get the packet containing this match and set evidence.

4.16.3 Extending

You can extend the module be adding more YARA rules in `modules/leak_detector/yara_rules/rules/`.

The rules will be automatically detected, compiled and run on the given PCAP.

If you want to contribute, improve existing Slips detection modules or implement your own detection modules, see section `:doc:Contributing <contributing>`.

4.17 Network Service Discovery Module

This module is responsible for detecting scans such as:

- Vertical port scans
- Horizontal port scans
- PING sweeps
- DHCP Scans

4.17.1 Vertical port scans

Slips considers an IP performing a vertical port scan if it contacts 5 or more different destination ports to the same destination IP in at least one time window (usually 1hs). The flows can be both, Non-Established TCP or UDP flows. On each arriving flow this check is performed.

The first portscan is detected as soon as it happens so the analysts knows with minimum 5 ports.

Slips reports evidence of a vertical portscan based on a log scale. meaning every 10, 100, 1000 ports scanned will generate an evidence. This avoids generating one port scan alert per flow in a long scan.

The total number of *packets* in all flows in the scan give us the confidence of the scan.

To minimize false positives, Slips ignores the broadcast IP 255.255.255.255 and the multicast IP if it's the source of vertical port scan.

4.17.2 Horizontal port scans

Slips detects TCP and UDP horizontal port scans. It considers an IP performing a horizontal port scan if it contacted 6 or more destination IPs on the same port with not established connections.

Slips reports evidence of a horizontal portscan based on a log scale. meaning every 10, 100, 1000 ports scanned will generate an evidence. This avoids generating one port scan alert per flow in a long scan.

To minimize false positives, Slips ignores the broadcast IP 255.255.255.255 if it's the source or the destination of horizontal port scans, and ignores all resolved IPs if they're the destination of port scans.

4.17.3 PING Sweeps

ICMP messages can be used to find out which hosts are alive in a network. Slips relies on Zeek detections for this, but it is done with our own Zeek scripts located in `zeek-scripts/icmps-scans.zeek`. The scripts detect three types of ICMP scans: 'ICMP-Timestamp', 'ICMP-Address', 'ICMP-AddressMask'.

We detect a scan every threshold. So we generate an evidence when there is 5,10,15, .. etc. ICMP established connections to different IPs.

Slips does this detection using Slips' own zeek script located in `zeek-scripts/icmps-scans.zeek` for zeek and pcap files and using the portscan module for binetflow files.

4.17.4 DHCP Scans

DHCP requests can be used to find out which IPs are taken in a network. Slips detects when an IP is requesting 4, 8, 12, etc. different IPs from the DHCP server within the same twimewindow (1 hour by default)

CONNECTIONS MADE BY SLIPS

Slips uses online databases to query information about many different things, for example (user agents, mac vendors etc.)

The list below contains all connections made by Slips

[useragentstring.com](#) -> For getting user agent info if no info was found in Zeek [macvendorlookup.com](#) -> For getting MAC vendor info if no info was found in the local maxmind db [maclookup.app](#) -> For getting MAC vendor info if no info was found in the local maxmind db [ip-api.com](#) -> For getting ASN info about IPs if no info was found in our Redis DB [ipinfo.io](#) -> For getting your public IP [virustotal.com](#) -> For getting scores about domains, IPs and URLs [urlhaus-api.abuse.ch](#) -> For getting info about URLs and downloaded files [check.torproject.org](#) -> For getting info about tor exist nodes. [cert.pl](#) -> Used in our list of TI files. [abuse.ch](#) -> Used by urlhaus for getting info about contacted domains and downloaded files.

If you want to contribute: improve existing Slips detection modules or implement your own detection modules, see section `:doc:Contributing` `<contributing>`.

ZEEK SCRIPTS

Slips is shipped with its own custom zeek scripts to be able to extend zeek functionality and customize the detections

6.1 Detect DoH

In the `detect_DoH.zeek` script, slips has its own list of ips that belong to dns/doh servers,

When slips encounters a connection to any IP of that list on port 443/tcp, it assumes it's a DoH connection, and times out the connection after 1h so that the connection won't take too long to appear in slips.

6.2 Detect ICMP Scans

In the `zeek-scripts/icmpps-scans.zeek` script, we check the type of ICMP in every ICMP packet seen in the network,

and we detect 3 types of ICMP scans: `ICMP-Timestamp-Scan`, `ICMP-AddressScan`, and `ICMP-AddressMaskScan` based on the `icmp` type

We detect a scan every threshold. So we generate an evidence when there is 5,10,15, .. etc. ICMP established connections to different IPs.

6.3 Detect the Gateway address

The `zeek-scripts/log_gw.zeek` script is responsible for recognizing the gateway address using zeek, and logging it to `notice.log`

BRUTE FORCE DETECTOR MODULE

The `brute_force_detector` module detects SSH brute forcing by combining repeated SSH sessions, Zeek SSH metadata, client software banners, and Zeek notice confirmations.

This module is loaded automatically by Slips like the rest of the modules in `modules/`, unless it is explicitly disabled in `config/slips.yaml`.

7.1 Inputs

The module subscribes to the following Slips channels:

- `new_ssh`
- `new_software`
- `new_notice`
- `tw_closed`

These channels are populated from Zeek logs:

- `ssh.log`
- `software.log`
- `notice.log`

7.2 What It Detects

The module tracks repeated SSH activity from the same source IP to the same destination IP and destination port inside the same time window.

It uses the following inputs:

- `ssh.log` to count repeated SSH sessions and authentication attempts
- `software.log` to extract the `SSH::CLIENT` banner and identify likely automation libraries such as `libssh`, `libssh2`, `paramiko`, `hydra`, `medusa`, or `ncrack`
- `notice.log` to consume Zeek `SSH::Password_Guessing` confirmations

7.3 Detection Logic

7.3.1 Counting Attempts

For each SSH flow, the module first checks the Zeek SSH authentication outcome:

- If `auth_success` is `true` or `T`, the flow is ignored for `brute_force_detector`.
- If `auth_attempts` is greater than `0`, that value is added to the bruteforce campaign counter.
- If `auth_attempts` is `0` or missing, but the SSH session is not marked successful, the module counts the session as one suspected password attempt.

The last rule is important for datasets where Zeek records repeated SSH handshakes without recording explicit authentication attempts, such as the `malicious-ssh-bruteforce.pcap` sample.

7.3.2 Threshold and Reporting

The default SSH brute force detector threshold is 9 attempts.

After the threshold is reached, the module does not alert on every new attempt. Instead, it uses sparse bucketed reporting so alerts become less frequent over time but never completely stop. With the default threshold, the alert points are:

- 9
- 10
- 12
- 16
- 24
- 40
- ...

7.3.3 Confidence

The evidence threat level is `medium`.

Confidence grows with the number of attempted passwords:

- first brute force detector evidence starts at the configured threshold
- full confidence is reached at `30` attempts
- suspicious SSH client banners add a small confidence bonus
- a Zeek `SSH: :Password_Guessing` notice acts as confirmation and promotes confidence using Zeek's confirmed connection count

7.4 Evidence Produced

The module emits `PASSWORD_GUESSING` evidence with:

- source attacker IP
- destination victim IP when available
- TCP destination port
- time window
- accumulated UIDs
- threat level `medium`
- confidence based on the number of attempts and confirmation data

Example description:

```
SSH brute force detector from 147.32.80.40 to 147.32.80.37 on SSH 902/tcp. Attempts_
↳observed: 24. Client banner: libssh libssh2_1.11.0 from software.log. Confidence: 0.89.
↳ by Slips
```

7.5 Zeek Confirmation

If Zeek raises `SSH::Password_Guessing` in `notice.log`, the module:

- emits an evidence immediately based on the notice
- stores the notice as confirmation for later `brute_force_detector` evidence
- uses the confirmed connection count from the Zeek notice to increase confidence

If Zeek does not generate `notice.log` for SSH password guessing, the module still detects `brute_force_detector` events from `ssh.log` and `software.log`.

7.6 Configuration

The module currently exposes:

```
brute_force_detector:
  ssh_attempt_threshold: 9
```

This value is read from `config/slips.yaml`.

7.7 Relationship With Flow Alerts

SSH brute force detector is now handled by the `brute_force_detector` module.

The `flow_alerts` module still handles:


- successful SSH detections
- Zeek port-scan notices
- certificate alerts
- DNS and connection heuristics
- SMTP bruteforce and the rest of the single-flow detections

It no longer owns SSH password guessing detection.


HTTPS ANOMALY DETECTION MODULE

This document describes how the `anomaly_detection_https` module detects anomalies from TLS/HTTPS traffic in Slips.

Event Volume Over Time (traffic time)



Detections Over Time (by confidence)



8.1 Goal

Detect unusual HTTPS behavior per host, using:

- Hourly behavior changes (volume and novelty patterns).
- Flow-level deviations (for known servers).
- Adaptive baselines that update over time, with poisoning resistance.

8.2 Input data used

The module subscribes to SSL/TLS events and reads related connection metadata from DB for the same UID.

Main fields used:

- SSL: uid, server_name (SNI), ja3, ja3s, dport, sport
- Conn (correlated): destination IP, total bytes, timing info

8.3 Traffic-time logic

All detection windows are based on **traffic timestamps** (packet/log time), not wall clock time.

This keeps behavior consistent for:

- live interface capture,
- live Zeek folder input,
- offline PCAP,
- offline Zeek logs.

8.4 Features

The module computes per-host hourly features:

- `ssl_flows`: number of SSL flows in the hour.
- `unique_servers`: number of distinct destination servers.
- `new_servers`: number of servers not seen before for that host.
- `ja3_changes`: number of new JA3 variants seen per server in the hour.
- `known_server_avg_bytes`: mean bytes for flows to already-known servers.

Flow-level feature:

- `bytes_to_known_server`: per-server bytes deviation on each flow.

8.5 Baseline and training

Each host has independent models.

8.5.1 Training phase (`training_hours > 0`)

For the first configured benign hours, the module does **fit-only** (Welford online moments):

- no detection decisions are emitted from hourly z-score rules before training ends,
- baseline mean/variance are learned strongly from this period.

Training fit strength is configurable with `training_alpha`:

Training fit technique is selected by `training_fit_method`:

- `training_fit_method = welford` -> Welford benign fit.
- `training_fit_method = ewma` -> EWMA-style training adaptation.

When `training_fit_method = ewma`, `training_alpha` controls strength:

- higher `training_alpha` = faster adaptation,
- lower `training_alpha` = slower adaptation.

8.5.2 No explicit training (`training_hours = 0`)

Detection starts immediately using online adaptation.

Special fallback only for `ja3_changes`:

- if hourly `ja3_changes < ja3_min_variants_per_server`, that hourly signal is ignored until enough activity exists.

8.6 Scoring

Each modeled feature uses robust scoring in three explicit steps:

1. Transform heavy-tail signals: $y = \log(1 + x)$ (`log1p`) for non-negative count/bytes features.
2. Estimate robust center/scale on recent transformed values:
 - `m = median(y)`
 - `MAD = median(|y - m|)`
 - `sigma_robust = max(1.4826 * MAD, min_std_floor)`
3. Score deviation:
 - `z_robust = |y_t - m| / sigma_robust`

Why this is used:

- HTTPS counts and byte volumes are typically right-skewed and heavy-tailed,
- mean/std-only scoring overreacts to bursts and underreacts after outliers,
- `log1p + median/MAD` is more stable under non-Gaussian traffic.

Thresholds:

- empirical thresholds calibrated from benign training when `training_hours > 0`,
- otherwise defaults (`hourly_zscore_threshold`, `flow_zscore_threshold`).

Calibration rule:

- per signal, collect robust z-scores on confirmed benign training data,
- set threshold to high benign quantile (`empirical_threshold_quantile`, default 0.995),
- fallback to defaults if training data is insufficient.

8.7 Adaptation states

After each hour closes, the module chooses model update mode.

Update event semantics:

- `training_fit`: initial benign baseline fit while `trained_hours < training_hours`; uses training fit method (Welford-style), not EWMA alpha.
- `baseline_update`: normal post-training adaptation; uses EWMA with `baseline_alpha`. In ADWIN mode, this is used when ADWIN does not signal drift.

- **drift_update**: post-training drift adaptation; uses EWMA with `drift_alpha`. In ADWIN mode, this is used only after ADWIN drift signal and small/drift-like classification.
- **suspicious_update**: post-training conservative adaptation; uses EWMA with `suspicious_alpha`. In ADWIN mode, this is used only after ADWIN drift signal and suspicious classification.

When `use_adwin_drift=false`:

1. **training_fit**
During benign training: Welford fit (no EWMA alpha).
2. **drift_update**
If anomaly score is small (`hourly_score <= adaptation_score_threshold`) and flow anomaly count is small (`<= max_small_flow_anomalies`), update with `drift_alpha`.
3. **suspicious_update**
Otherwise update with `suspicious_alpha` (much smaller), to limit poisoning.

For normal non-anomalous periods outside training, per-feature EWMA uses `baseline_alpha`.

8.7.1 ADWIN drift trigger (`use_adwin_drift=true`)

If `use_adwin_drift=true` and `river` is installed, ADWIN is the only drift trigger in both paths:

- **Hourly path**: ADWIN receives each raw hourly feature stream.
- **Flow path**: ADWIN receives each raw per-flow signal stream.
- ADWIN drift detected -> classify as `drift_update` or `suspicious_update` using existing thresholds.
- No ADWIN drift -> use `baseline_update` (`baseline_alpha`), even if anomalies exist.
- During benign training, ADWIN is still warmed with benign scores to reduce cold-start noise after training.

Why raw signals:

- drift is a distribution change in the observed variables, so ADWIN tracks the raw feature streams directly,
- z-scores are still used for anomaly magnitude and evidence reasons, but not as the primary drift input.

Performance note:

- hourly ADWIN cost scales with hourly feature count,
- flow ADWIN cost scales with per-flow signal count,
- both are constant-time scalar updates and usually lightweight.

Current tuned defaults for faster ADWIN reaction:

- `adwin_delta`: 0.01
- `adwin_clock`: 1
- `adwin_grace_period`: 5
- `adwin_min_window_length`: 5

8.8 New server vs JA3 behavior

- `new_servers` is modeled as an hourly statistical feature and adapted over time.
- `new_server` can also appear as a direct flow-level novelty reason.
- `ja3_changes` is handled statistically at hourly level (with fallback gate only when training is zero).

- `new_ja3s` can appear as direct flow-level novelty reason.

8.9 Confidence and threat level

Each detection computes confidence score $[0, 1]$ from multiple factors:

- anomaly severity,
- persistence in recent history,
- baseline quality,
- multi-signal agreement.

Mapped levels:

- low / medium / high confidence

Threat level used in evidence:

- low for low or medium confidence
- medium for high confidence

8.10 Evidence format

Evidence description is human-readable and concise:

HTTPS anomaly: `type=<type>; confidence=<level> (<score>); reason=<reason>; value=<value>; why=<explanation>.`

Examples of reasons:

- New Server
- New JA3S
- Bytes to Known Server
- Hourly feature deviations (e.g., New Servers Count, JA3 Changes)

8.11 Configuration keys

Section: `anomaly_detection_https` in `config/slips.yaml`.

Main keys:

- `training_hours`
- `training_fit_method`
- `training_alpha`
- `hourly_zscore_threshold`
- `flow_zscore_threshold`
- `adaptation_score_threshold`
- `baseline_alpha`
- `drift_alpha`
- `suspicious_alpha`

- `min_baseline_points`
- `max_small_flow_anomalies`
- `ja3_min_variants_per_server`
- `use_adwin_drift`
- `adwin_delta`
- `adwin_clock`
- `adwin_grace_period`
- `adwin_min_window_length`
- `empirical_threshold_quantile`
- `log_verbosity`

Defaults (from parser/config):

- `training_alpha: 1.0`
- `training_fit_method: welford`
- `use_adwin_drift: true`
- `adwin_delta: 0.01`
- `adwin_clock: 1`
- `adwin_grace_period: 5`
- `adwin_min_window_length: 5`

Reference:

- River ADWIN: <https://riverml.xyz/latest/api/drift/ADWIN/>
- Data transformations for skew/heavy tails: <https://otexts.com/fpp3/transformations.html>
- Robust scale (MAD): https://en.wikipedia.org/wiki/Median_absolute_deviation

8.12 Operational logs

The module logs key events such as:

- flow arrivals,
- hour close and computed features,
- training fit updates,
- drift updates,
- suspicious updates,
- detections and emitted evidence.

FLOW_ALERTS MODULE

The module of flow alerts has several behavioral techniques to detect attacks by analyzing the content of each flow alone.

The detection techniques are:

- Long connections
- Successful SSH connections
- Connections without DNS resolution
- DNS resolutions without a connection
- Connections to unknown ports
- Data exfiltration
- Malicious JA3 hashes
- Connections to port 0
- Multiple reconnection attempts
- Alerts from Zeek: Self-signed certs, invalid certs, port-scans and address scans
- DGA
- Connection to multiple ports
- Malicious SSL certificates
- Pastebin downloads
- Young domains
- Bad SMTP logins
- SMTP login bruteforce
- DNS ARPA Scans
- SSH version changing
- Incompatible CN
- CN URL Mismatch
- Weird HTTP methods
- Non-SSL connections on port 443
- Connection to private IPs
- Connection to private IPs outside the current local network

- High entropy DNS TXT answers
- Devices changing IPs
- GRE tunnels
- GRE tunnel scan
- Invalid DNS answers
- Tor exit nodes The details of each detection follows.

9.1 Long Connections

Detect connections that are long, except the multicast connections.

By default, A connection is considered long if it exceeds 1500 seconds (25 Minutes).

This threshold can be changed in `slips.yaml` by changing the value of `long_connection_threshold`

9.2 Connections without DNS resolution

This will detect connections done without a previous DNS resolution. The idea is that a connection without a DNS resolution is slightly suspicious.

If Slips runs by capturing packets directly from a network device (as opposed to, for example, a PCAP file), this detection will ignore all connections that happen in the first 3 minutes of operation of Slips. This is because most times Slips is started when the computer is already running, and many DNS connections were already done. So waiting 3 minutes decreases the amount of False Positives.

This detection will ignore certain IP addresses for which a connection without DNS is ok. The exceptions are:

- Private IPs
- Localhost IPs (127.0.0.0/8)
- Reserved IPs (including the super broadcast 255.255.255.255)
- IPv6 local-link IPs
- Multicast IPs
- Broadcast IPs only if they are private
- Well known organizations

DNS resolutions of well known orgs might be done using DoH, in this case, slips doesn't know about the DNS resolution because the resolved domain won't be in `dns.log` so we simply ignore alerts of this type when connected to well known organizations. In particular Facebook, Apple, Google, Twitter, and Microsoft.

Slips uses its own lists of organizations and information about them (IPs, IP ranges, domains, and ASNs). They are stored in `slips_files/organizations_info` and they are used to check whether the IP/domain of each flow belongs to a known org or not.

Slips also doesn't detect connection without DNS to any domain in the tranco whitelist.

Slips doesn't detect 'connection without DNS' when running on an interface except for when it's done by this instance's own IP and only after 30 minutes has passed to avoid false positives (assuming the DNS resolution of these connections did happen before slips started).

check the [DoH section](#) of the docs for info on how slips detects DoH.

9.3 Successful SSH connections

Slips detects successful SSH connections using 2 ways

1. Using Zeek. Zeek logs successful SSH connection to `ssh.log` by default
2. If all bytes sent in a SSH connection is more than 4290 bytes

9.4 DNS resolutions without a connection

This will detect DNS resolutions for which no further connection was done. A resolution without a usage is slightly suspicious.

The domains that are excepted are:

- All reverse DNS resolutions using the `in-addr.arpa` domain.
- All `.local` domains
- The wild card domain *
- Subdomains of `cymru.com`, since it is used by the `ipwhois` library to get the ASN of an IP and its range.
- WPAD domain from Windows
- domains without a TLD such as the Chrome test domains.
- DNS resolutions of any type other than AAAA and A

Slips doesn't detect 'DNS resolutions without a connection' when running on an interface except for when it's done by this instance's own IP and only after 30 minutes has passed to avoid false positives (assuming the connection did happen and yet to be logged).

When running on interface and files. For each DNS flow found, slips waits 30 mins zeek time for the connection to be found before setting an evidence.

This is done by comparing each ts of every new dns flow to the pending detection, once 30 mins difference between the 2 flows is detected, slips sets the evidence.

To avoid accumulating so many pending DNS flows for 30 mins, slips checks if the connection of the pending DNS flows arrived every 10 and 20 mins too, if not found, slips waits extra 10 mins (so that would be 30 mins total) and sets the evidence.

9.5 Connection to unknown ports

Slips has a list of known ports located in `slips_files/ports_info/services.csv`

and a list of ports that belong to a specific organization in `slips_files/ports_info/ports_used_by_specific_orgs.csv`

These are the cases where Slips marks the port as known and doesn't trigger an alert

1. If the port is in the list of well known ports in `services.csv`.
2. If Slips has that port's info in `ports_used_by_specific_orgs.csv` and the source and destination addresses belong to that organization.

Slips considers an IP belongs to an org if:

1. Both `saddr` and `daddr` have the organization's name in their MAC vendor (e.g. Apple.)
2. Both `saddr` and `daddr` belong to the range specified in `theports_used_by_specific_orgs.csv` for that organization.

3. If the SNI, hostname, rDNS, ASN of this IP belong to this organization.
4. If the IP is hardcoded in any of the organizations IPs in `slips_files/organizations_info/`.

Otherwise, Slips triggers and “unknown port” evidence.

For example, even though 5223/TCP isn't a well known port, Apple uses it in Apple Push Notification Service (APNS).

The threat level of this evidence depends on the state of the flow. established connections have higher threat levels.

9.6 Data Upload

Slips generates ‘possible data upload’ alerts when the number of uploaded bytes to any IP exceeds 100 MBs over the timewindow period which is, by default, 1h.

See detailed explanation of timewindows [here](#).

The number of MBs can be modified by changing the value of `data_exfiltration_threshold` in `slips.yaml`

Slips also detects data upload when an IP uploads ≥ 100 MBs to any IP in 1 connections.

9.7 Tor Exit Nodes

Slips generates informational evidence when it detects a tor exit node from the list <https://check.torproject.org/torbulkexitlist>.

9.8 Malicious JA3 and JA3s hashes

Slips uses JA3 hashes to detect C&C servers (JA3s) and infected clients (JA3)

Slips is shipped with its own zeek scripts that add JA3 and JA3s fingerprints to the SSL log files generated by zeek.

Slips supports JA3 feeds in addition to having more than 40 different threat intelligence feeds. The JA3 feeds contain JA3 fingerprints that are identified as malicious. The JA3 threat intelligence feed used by Slips now is [Abuse.ch](#) JA3 feed. And you can add other JA3 TI feeds in `ja3_feeds` in `slips.yaml`.

9.9 Connections to port 0

There has been a significant rise in the number of attacks listed as Port 0. Last year, these equated to 10% of all attacks, but now it's up to almost 25%.

Slips detects any connection to port 0 using any protocol other than ‘IGMP’ and ‘ICMP’ as malicious.

9.10 Multiple reconnection attempts

Multiple reconnection attempts in Slips are 5 or more not established flows (reconnections) to the same destination IP on the same destination port.

9.11 Zeek alerts

By default, Slips depends on Zeek for detecting different behaviours, for example Self-signed certs, invalid certs, port-scans, and address scans.

Some scans are also detected by Slips independently of Zeek, like ICMP sweeps and vertical/horizontal portscans. Check [PING Sweeps](#) section for more info

9.12 SMTP login bruteforce

Slips alerts when 3+ invalid SMTP login attempts occurs within 10s

9.13 SSH brute_force_detector

SSH brute force detector is documented in the dedicated `brute_force_detector` module page:

- *brute_force_detector Module*

The `flow_alerts` module still detects successful SSH sessions, but SSH password guessing is no longer owned by `flow_alerts`.

9.14 DGA

When the DNS server fails to resolve a domain, it responds back with NXDOMAIN code.

To detect DGA, Slips will count the amount of NXDOMAINs met in the DNS traffic of each source IP.

Then we alert when there is 10 or more NXDOMAINs.

Every 10,15,20 ..etc slips generates an evidence.

9.15 Connection to multiple ports

When Slips encounters a connection to or from a specific IP and a specific port, it scans previous connections looking for connection to/from that same IP using a different port.

It alerts when finding two or more connections to the same IP.

9.16 Malicious SSL certificates

Slips uses SSL certificates sha1 hashes to detect C&C servers.

Slips supports SSL feeds and is shipped with [Abuse.ch](#) feed of malicious SSL hashes by default. And you can add other SSL feeds in `ssl_feeds` in `slips.yaml`.

9.17 Pastebin downloads

Slips detects downloads from pastebin using SSL and HTTP

It alerts when a downloaded file from pastebin exceeds 700 bytes

This value can be customized in `slips.yaml` by changing `pastebin_download_threshold`

Slips detects the pastebin download once the SSL connection is over , which may take hours.

9.18 Young Domains

Slips uses whois python library to get the creation date of every domain met in the dns flows.

If a domain's age is less than 60 days, slips sets an alert.

Not all domains are supported, here's the list of supported TLDs.

```
[ '.ac_uk', '.am', '.amsterdam', '.ar', '.at', '.au',
'.bank', '.be', '.biz', '.br', '.by', '.ca', '.cc',
'.cl', '.club', '.cn', '.co', '.co_il', '.co_jp', '.com',
'.com_au', '.com_tr', '.cr', '.cz', '.de', '.download', '.edu',
'.education', '.eu', '.fi', '.fm', '.fr', '.frl', '.game', '.global_',
'.hk', '.id_', '.ie', '.im', '.in_', '.info', '.ink', '.io',
'.ir', '.is_', '.it', '.jp', '.kr', '.kz', '.link', '.lt', '.lv',
'.me', '.mobi', '.mu', '.mx', '.name', '.net', '.ninja',
'.nl', '.nu', '.nyc', '.nz', '.online', '.org', '.pe',
'.pharmacy', '.pl', '.press', '.pro', '.pt', '.pub', '.pw',
'.rest', '.ru', '.ru_rf', '.rw', '.sale', '.se', '.security',
'.sh', '.site', '.space', '.store', '.tech', '.tel', '.theatre',
'.tickets', '.trade', '.tv', '.ua', '.uk', '.us', '.uz', '.video',
'.website', '.wiki', '.work', '.xyz', '.za' ]
```

9.19 Bad SMTP logins

Slips uses zeek to detect SMTP connections. When zeek detects a bad smtp login, it logs it to smtp.log, then slips reads this file and sets an evidence.

9.20 SMTP bruteforce

Slips detects a SMTP bruteforce when 3 or more bad SMTP logins happen within 10 seconds.

With every generated evidence, Slips gathers as much info about the malicious IP and prints it with the alert.

So instead of having an alerts saying:

```
Detected SSL certificate validation failed with (certificate has expired) Destination_
↪IP: 216.58.201.70.
```

Slips gathers AS, hostname, SNI, rDNS and any available data about this IP and you get an alert saying:

```
Detected SSL certificate validation failed with (certificate has expired) Destination IP:
216.58.201.70. AS: GOOGLE, US, SNI: 2542116.fl.doubleclick.net, rDNS: prg03s01-in-f70.
↪1e100.net
```

9.21 DNS ARPA Scans

Whenever slips sees a new domain in dns.log, if the domain ends with '.in-addr.arpa' slips keeps track of this domain and the source IP that made the DNS request.

Then, if the source IP is seen doing 10 or more ARPA queries within 2 seconds, slips generates an ARPA scan detection.

9.22 SSH version changing

Zeek logs the used software and software versions in software.log, so slips knows from this file the software used by different IPs, like whether it's an SSH::CLIENT, an HTTP::BROWSER, or an HTTP::SERVER

When slips detects an SSH client or an SSH server, it stores it with the IP and the SSH versions used in the database

Then whenever slips sees the same IP using another SSH version, it compares the stored SSH versions with the current SSH versions

If they are different, slips generates an alert

9.23 Incompatible CN

Zeek logs each Certificate CN in ssl.log

When slips encounters a cn that claims to belong to any of Slips supported orgs (Google, Microsoft, Apple or Twitter) Slips checks if the destination address or the destination server name belongs to these org.

If not, slips generates an evidence.

9.24 CN URL Mismatch

Zeek logs each Certificate CN in ssl.log For each CN Slips encounters, it checks if the server name is the same as the CN Or if it belongs to the same org as the CN. if not, slips triggers an evidence

9.25 Weird HTTP methods

Slips uses zeek's weird.log where zeek logs weird HTTP methods seen in http.log

When there's a weird HTTP method, slips detects it as well.

9.26 Non-SSL connections on port 443

Slips detects established connections on port 443 that are not using SSL using zeek's conn.log flows

if slips finds a flow using destination port 443 and the 'service' field in conn.log isn't set to 'ssl', it alerts.

Sometimes zeek detects a connection from a source to a destination IP on port 443 as SSL, and another connection within 5 minutes later as non-SSL. Slips detects that and does not set an evidence for any of them.

Here's how it works

9.27 Connection to private IPs

Slips detects when a private IP is connected to another private IP with threat level info.

But it skips this alert when it's a DNS or a DHCP connection on port 53, 67 or 68 UDP to the gateway IP.

Slips also skips this alert for DNS traffic on port 53 and for DHCPv6 traffic on ports 546 and 547. When Slips sees a private DNS flow in the analyzed traffic, it records that destination address as a DNS server and prints it to the display and to slips.log.

9.28 Connection to private IPs outside the current local network

Slips detects the currently used local network and alerts if it find a connection to/from a private IP that doesn't belong to it.

For example if the currently used local network is: 192.168.1.0/24

and slips sees a forged packet going from 192.168.1.2 to 10.0.0.1, it will alert

Slips detects the current local network by using the local network of the private ips specified in `client_ips` parameter in `slips.yaml`

If no IPs are specified, slips uses the local network of the first private source ip found in the traffic.

This threat level of this detection is low if the source ip is the one outside of local network because it's unlikely. and high if the destination ip is the one outside of local network.

Slips ignores evidence of this type when the destination IP is a private IP outside of local network and is communicating on port 53/UDP. Slips marks that destination address as the DNS server as soon as it sees a private DNS flow in `dns.log` and prints `Detected DNS server by traffic heuristic: <ip>` to the display and `slips.log`. The same heuristic applies to private IPv4 and IPv6 DNS servers. this is likely a DNS misconfiguration hence a FP.

Slips also ignores this evidence when the checked IP is one of the DNS servers detected from the analyzed DNS traffic and the matching DNS server port is 53. This applies to both IPv4 and IPv6 resolver addresses and covers requests sent to the resolver and replies coming back from it.

For `CONNECTION_TO_PRIVATE_IP`, Slips suppresses DHCPv6 traffic on ports 546 and 547. This avoids false positives when the same local network device provides both DNS and DHCPv6 services but answers with a different IPv6 address, such as a link-local address.

9.29 High entropy DNS TXT answers

Slips check every DNS answer with TXT record for high entropy strings. Encoded or encrypted strings with entropy higher than or equal 5 will then be detected using shannon entropy and alerted by slips.

the entropy threshold can be changed in `slips.yaml` by changing the value of `entropy_threshold`

9.30 Devices changing IPs

Slips stores the MAC of each new IP it sees in `conn.log`.

Then for every source address in `conn.log`, slips checks if the MAC of it was used by another IP.

If so, it alerts "Device changing IPs".

9.31 GRE tunnels

Slips uses `zeek tunnel.log` to alert on GRE tunnels when found. Whenever one any action other than "Tunnel::DISCOVER" is found, slips sets an evidence with threat level low

9.32 GRE tunnel scans

Slips uses `zeek tunnel.log` to alert on GRE tunnels scan. Slips considers any log with "Tunnel::DISCOVER" action a GRE scan.

The threat level of this evidence is low.

9.33 Login log entries

Slips reads Zeek `login.log` entries generated by `zeek-scripts/login.zeek`. It can detect telnet, rlogin, or rsh logins. For each entry in the zeek logs, the profiler publishes a `new_login` message and `flow_alerts` sets informational evidence for the server receiving the login.

9.34 Invalid DNS resolutions

Some DNS resolvers answer the DNS query to adservers with 0.0.0.0 or 127.0.0.1 as the ip of the domain to block the domain. Slips detects this and sets an informational evidence.

This detection doesn't apply to queries ending with ".arpa" or ".local"

FEATURES

This Section will contain a list of all features and detections listed in the [flowalerts section](#) and the [detection modules section](#) and a brief description of how slips works.

10.1 Flow Alerts Module

The module of flow alerts has several behavioral techniques to detect attacks by analyzing the content of each flow alone.

The detection techniques are:

- Long connections
- Successful SSH connections
- Connections without DNS resolution
- DNS resolutions to IPs that were never used
- Connections to unknown ports
- Data exfiltration
- Malicious JA3 hashes
- Connections to port 0
- Multiple reconnection attempts
- Alerts from Zeek: Self-signed certs, invalid certs, port-scans and address scans
- DGA
- Connection to multiple ports
- Malicious SSL certificates
- Young domains
- Bad SMTP logins
- SMTP login bruteforce
- DNS ARPA Scans
- Multiple SSH versions
- Incompatible CN
- Weird HTTP methods
- Non-SSL connections on port 443

- Connection to private IPs
- Connection to private IPs outside the current local network
- High entropy DNS TXT answers
- Devices changing IPs
- GRE tunnels
- GRE tunnel scan
- SSH version changing
- Invalid DNS resolutions

The details of each detection follows.

10.1.1 Long Connections

Detect connections that are long, except the multicast connections.

By default, A connection is considered long if it exceeds 1500 seconds (25 Minutes).

This threshold can be changed in `config/slips.yaml` by changing the value of `long_connection_threshold`

10.1.2 Incompatible CN

Zeek logs each Certificate CN in `ssl.log`

When slips encounters a cn that claims to belong to any of Slips supported orgs (Google, Microsoft, Apple or Twitter) Slips checks if the destination address or the destination server name belongs to these org.

If not, slips generates an alert.

10.1.3 CN URL Mismatch

Zeek logs each Certificate CN in `ssl.log` For each CN Slips encounters, it checks if the server name is the same as the CN Or if it belongs to the same org as the CN. if not, slips triggers an evidence

10.1.4 High entropy DNS TXT answers

Slips check every DNS answer with TXT record for high entropy strings. Encoded or encrypted strings with entropy higher than or equal 5 will then be detected using shannon entropy and alerted by slips.

the entropy threshold can be changed in `slips.yaml` by changing the value of `entropy_threshold`

10.1.5 Devices changing IPs

Slips stores the MAC of each new IP it sees in `conn.log`.

Then for every source address in `conn.log`, slips checks if the MAC of it was used by another IP.

If so, it alerts “Device changing IPs”.

10.2 GRE tunnels

Slips uses `zeek tunnel.log` to alert on GRE tunnels when found. Whenever one any action other than “Tunnel::DISCOVER” is found, slips sets an evidence with threat level low

10.3 GRE tunnel scans

Slips uses zeek tunnel.log to alert on GRE tunnels scan. Slips considers any log with “Tunnel::DISCOVER” action a GRE scan.

The threat level of this evidence is low.

10.3.1 SMTP login bruteforce

Slips alerts when 3+ invalid SMTP login attempts occurs within 10s

10.3.2 Connection to private IPs

Slips detects when a private IP is connected to another private IP with threat level info.

But it skips this alert when it’s a DNS or a DHCP connection on port 53, 67 or 68 UDP to the gateway IP.

Slips also skips this alert for DNS traffic on port 53 and for DHCPv6 traffic on ports 546 and 547. When Slips sees a private DNS flow in the analyzed traffic, it records that destination address as a DNS server and prints it to the display and to slips.log.

10.3.3 Connection to private IPs outside the current local network

Slips detects the currently used local network and alerts if it find a connection to/from a private IP that doesn’t belong to it.

For example if the currently used local network is: 192.168.1.0/24

and slips sees a forged packet going from 192.168.1.2 to 10.0.0.1, it will alert

Slips detects the current local network by using the local network of the private ips specified in `client_ips` parameter in `slips.yaml`

If no IPs are specified, slips uses the local network of the first private source ip found in the traffic.

This threat level of this detection is low if the source ip is the one outside of local network because it’s unlikely. and high if the destination ip is the one outside of local network.

Slips ignores evidence of this type when the destination IP is a private IP outside of local network and is communicating on port 53/UDP. Slips marks that destination address as the DNS server as soon as it sees a private DNS flow in `dns.log` and prints `Detected DNS server by traffic heuristic: <ip>` to the display and `slips.log`. The same heuristic applies to private IPv4 and IPv6 DNS servers. this is likely a DNS misconfiguration hence a FP.

Slips also ignores this evidence when the checked IP is one of the DNS servers detected from the analyzed DNS traffic and the matching DNS server port is 53. This applies to both IPv4 and IPv6 resolver addresses and covers requests sent to the resolver and replies coming back from it.

For `CONNECTION_TO_PRIVATE_IP`, Slips suppresses DHCPv6 traffic on ports 546 and 547. This avoids false positives when the same local network device provides both DNS and DHCPv6 services but answers with a different IPv6 address, such as a link-local address.

10.3.4 Weird HTTP methods

Slips uses zeek’s weird.log where zeek logs weird HTTP methods seen in http.log

When there’s a weird HTTP method, slips detects it as well.

10.3.5 Non-SSL connections on port 443

Slips detects established connections on port 443 that are not using SSL using zeek's conn.log flows

if slips finds a flow using destination port 443 and the 'service' field in conn.log isn't set to 'ssl', it alerts.

Sometimes zeek detects a connection from a source to a destination IP on port 443 as SSL, and another connection within 5 minutes later as non-SSL. Slips detects that and does not set an evidence for any of them.

Here's how it works

10.4 Non-HTTP connections on port 80.

Slips detects established connections on port 80 that are not using SSL using zeek's conn.log flows.

if slips finds a flow using destination port 80 and the 'service' field in conn.log isn't set to 'http', it sets an evidence.

If a flow without http as a service is found, slips first checks past and future flows from the same src ip + dst ip + port to see if there's a flow with http as a service, if there is, slips ignores the alert. This is done to avoid FPs coming from zeek.

10.4.1 Connections without DNS resolution

This will detect connections done without a previous DNS resolution. The idea is that a connection without a DNS resolution is slightly suspicious.

If Slips runs by capturing packets directly from a network device (as opposed to, for example, a PCAP file), this detection will ignore all connections that happen in the first 3 minute of operation of Slips. This is because most times Slips is started when the computer is already running, and many DNS connections were already done. So waiting 3 minutes decreases the amount of False Positives.

This detection will ignore certain IP addresses for which a connection without DNS is ok. The exceptions are:

- Private IPs
- Localhost IPs (127.0.0.0/8)
- Reserved IPs (including the super broadcast 255.255.255.255)
- IPv6 local-link IPs
- Multicast IPs
- Broadcast IPs only if they are private
- Well known organizations

DNS resolutions of well known orgs might be done using DoH, in this case, slips doesn't know about the DNS resolution because the resolved domain won't be in dns.log so we simply ignore alerts of this time about well known org such as (facebook, apple, google, twitter, and microsoft)

Slips uses its own lists of organizations info (IPs, IP ranges, domains, and ASNs) stored in `slips_files/organizations_info` to check whether the IP/domain of each flow belong to a known org or not.

Slips doesn't detect 'connection without DNS' when running on an interface except for when it's done by this instance's own IP.

check [DoH section](#) of the docs for info on how slips detects DoH.

10.4.2 Successful SSH connections

Slips detects successful SSH connections using 2 ways

1. Using Zeek. Zeek logs successful SSH connection to `ssh.log` by default
2. if all bytes sent in a SSH connection is more than 4290 bytes

10.4.3 DNS resolutions without a connection

This will detect DNS resolutions for which no further connection was done. A resolution without a usage is slightly suspicious.

The domains that are excepted are:

- All reverse DNS resolutions using the `in-addr.arpa` domain.
- All `.local` domains
- The wild card domain *
- Subdomains of `cymru.com`, since it is used by the `ipwhois` library to get the ASN of an IP and its range.
- WPAD domain from Windows
- domains without a TLD such as the Chrome test domains.
- DNS resolutions of any type other than AAAA and A

Slips doesn't detect 'DNS resolutions without a connection' when running on an interface except for when it's done by this instance's own IP and only after 30 minutes has passed to avoid false positives (assuming the connection did happen and yet to be logged).

When running on interface and files. For each DNS flow found, slips waits 30 mins zeek time for the connection to be found before setting an evidence.

This is done by comparing each ts of every new dns flow to the pending detection, once 30 mins difference between the 2 flows is detected, slips sets the evidence.

To avoid accumulating so many pending DNS flows for 30 mins, slips checks if the connection of the pending DNS flows arrived every 10 and 20 mins too, if not found, slips waits extra 10 mins (so that would be 30 mins total) and sets the evidence.

10.4.4 Connection to unknown ports

Slips has a list of known ports located in `slips_files/ports_info/services.csv`

and a list of ports that belong to a specific organization in `slips_files/ports_info/ports_used_by_specific_orgs.csv`

These are the cases where Slips marks the port as known and doesn't trigger an alert

1. If the port is in the list of well known ports in `services.csv`.
2. If Slips has that port's info in `ports_used_by_specific_orgs.csv` and the source and destination addresses belong to that organization.

Slips considers an IP belongs to an org if:

1. Both `saddr` and `daddr` have the organization's name in their MAC vendor (e.g. Apple.)
2. Both `saddr` and `daddr` belong to the range specified in `theports_used_by_specific_orgs.csv` for that organization.
3. If the SNI, hostname, rDNS, ASN of this IP belong to this organization.

4. If the IP is hardcoded in any of the organizations IPs in `slips_files/organizations_info/`.

Otherwise, Slips triggers and “unknown port” evidence.

For example, even though 5223/TCP isn’t a well known port, Apple uses it in Apple Push Notification Service (APNS).

The threat level of this evidence depends on the state of the flow. established connections have higher threat levels.

10.4.5 Data exfiltration

Slips generate a ‘possible data exfiltration alerts when the number of uploaded files to any IP exceeds 700 MBs.

The number of MBs can be modified by editing the value of `data_exfiltration_threshold` in `config/slips.yaml`

10.4.6 Malicious JA3 and JA3s hashes

Slips uses JA3 hashes to detect C&C servers (JA3s) and infected clients (JA3)

Slips is shipped with it’s own zeek scripts that add JA3 and JA3s fingerprints to the SSL log files generated by zeek.

Slips supports JA3 feeds in addition to having more than 40 different threat intelligence feeds. The JA3 feeds contain JA3 fingerprints that are identified as malicious. The JA3 threat intelligence feed used by Slips now is [Abuse.ch](#) JA3 feed. And you can add other JA3 TI feeds in `ja3_feeds` in `config/slips.yaml`.

10.4.7 Connections to port 0

There has been a significant rise in the number of attacks listed as Port 0. Last year, these equated to 10% of all attacks, but now it’s up to almost 25%.

Slips detects any connection to port 0 using any protocol other than ‘IGMP’ and ‘ICMP’ as malicious.

10.4.8 Multiple reconnection attempts

Multiple reconnection attempts in Slips are 5 or more not established flows (reconnections) to the same destination IP.

10.4.9 Zeek alerts

By default, Slips depends on Zeek for detecting different behaviours, for example Self-signed certs, invalid certs, port-scans and address scans, and password guessing.

Some scans are also detected by Slips independently of Zeek, like ICMP sweeps and vertical/horizontal portscans. Check `section` for more info #todo

10.4.10 DGA

When the DNS server fails to resolve a domain, it responds back with NXDOMAIN code.

To detect DGA, Slips will count the amount of NXDOMAINs met in the DNS traffic of each source IP.

Then we alert when there is 10 or more NXDOMAINs.

Every 10,15,20 ..etc slips generates an evidence.

10.4.11 Connection to multiple ports

When Slips encounters a connection to or from a specific IP and a specific port, it scans previous connections looking for connection to/from that same IP using a different port.

It alerts when finding two or more connections to the same IP.

10.4.12 Malicious SSL certificates

Slips uses SSL certificates sha1 hashes to detect C&C servers.

Slips supports SSL feeds and is shipped with [Abuse.ch](#) feed of malicious SSL hashes by default. And you can add other SSL feeds in `ssl_feeds` in `config/slips.yaml`.

10.4.13 Young Domains

Slips uses whois python library to get the creation date of every domain met in the dns flows.

If a domain's age is less than 60 days, slips sets an alert.

Not all domains are supported, here's the list of supported TLDs.

```
[ '.ac_uk', '.am', '.amsterdam', '.ar', '.at', '.au',
  '.bank', '.be', '.biz', '.br', '.by', '.ca', '.cc',
  '.cl', '.club', '.cn', '.co', '.co_il', '.co_jp', '.com',
  '.com_au', '.com_tr', '.cr', '.cz', '.de', '.download', '.edu',
  '.education', '.eu', '.fi', '.fm', '.fr', '.frl', '.game', '.global_',
  '.hk', '.id_', '.ie', '.im', '.in_', '.info', '.ink', '.io',
  '.ir', '.is_', '.it', '.jp', '.kr', '.kz', '.link', '.lt', '.lv',
  '.me', '.mobi', '.mu', '.mx', '.name', '.net', '.ninja',
  '.nl', '.nu', '.nyc', '.nz', '.online', '.org', '.pe',
  '.pharmacy', '.pl', '.press', '.pro', '.pt', '.pub', '.pw',
  '.rest', '.ru', '.ru_rf', '.rw', '.sale', '.se', '.security',
  '.sh', '.site', '.space', '.store', '.tech', '.tel', '.theatre',
  '.tickets', '.trade', '.tv', '.ua', '.uk', '.us', '.uz', '.video',
  '.website', '.wiki', '.work', '.xyz', '.za' ]
```

10.4.14 Bad SMTP logins

Slips uses zeek to detect SMTP connections, When zeek detects a bad smtp login, it logs it to `smtp.log`, then slips reads this file and sets an evidence.

10.4.15 SMTP bruteforce

Slips detects a SMTP bruteforce when 3 or more bad SMTP logins happen within 10 seconds.

With every generated evidence, Slips gathers as much info about the malicious IP and prints it with the alert.

So instead of having an alerts saying:

```
Detected SSL certificate validation failed with (certificate has expired) Destination_
↪ IP: 216.58.201.70.
```

Slips gathers AS, hostname, SNI, rDNS and any available data about this IP and you get an alert saying:

```
Detected SSL certificate validation failed with (certificate has expired) Destination IP:
216.58.201.70. AS: GOOGLE, US, SNI: 2542116.fl.doubleclick.net, rDNS: prg03s01-in-f70.
↪ 1e100.net
```

10.4.16 DNS ARPA Scans

Whenever slips sees a new domain in `dns.log`, if the domain ends with `.in-addr.arpa` slips keeps track of this domain and the source IP that made the DNS request.

Then, if the source IP is seen doing 10 or more ARPA queries within 2 seconds, slips generates an ARPA scan detection.

10.4.17 SSH version changing

Zeek logs the used software and software versions in `software.log`, so slips knows from this file the software used by different IPs, like whether it's an `SSH::CLIENT`, an `HTTP::BROWSER`, or an `HTTP::SERVER`

When slips detects an SSH client or an SSH server, it stores it with the IP and the SSH versions used in the database

Then whenever slips sees the same IP using another SSH version, it compares the stored SSH versions with the current SSH versions

If they are different, slips generates an alert

10.5 Invalid DNS resolutions

Some DNS resolvers answer the DNS query to ad servers with `0.0.0.0` or `127.0.0.1` as the ip of the domain to block the domain. Slips detects this and sets an informational evidence.

This detection doesn't apply to queries ending with `.arpa` or `.local`

10.6 Detection modules

Slips is a behavioral-based IPS that uses machine learning to detect malicious behaviors in the network traffic. It is a modular software that can be extended. When Slips is run, it spawns several child processes to manage the I/O, to profile attackers and to run the detection modules.

Here we describe what detection modules are run on the traffic to detect malicious behaviour.

Modules are Python-based files that allow any developer to extend the functionality of Slips. They process and analyze data, perform additional detections and store data in Redis for other modules to consume. Currently, Slips has the following modules:

10.6.1 Virustotal Module

This module is used to lookup IPs, domains, and URLs on virustotal

To use it you need to add your virustotal API key in `config/vt_api_key`

10.6.2 RiskIQ Module

This module is used to get different information (passive DNS, IoCs, etc.) from RiskIQ To use this module your RiskIQ email and API key should be stored in `config/RiskIQ_credentials`

the format of this file should be the following:

```
example@domain.com
e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855
```

The hash should be your 64 character API Key.

The path of the file can be modified by changing the `RiskIQ_credentials_path` parameter in `config/slips.yaml`

10.6.3 Leak Detection Module

This module runs on pcaps, it uses YARA rules to detect leaks.

You can add your own YARA rule in `modules/leak_detector/yara_rules/rules` and it will be automatically compiled and stored in `modules/leak_detector/yara_rules/compiled` and matched against every pcap.

10.6.4 Blocking Module

To enable blocking in slips, start slips with the `-p` flag.

This feature is only supported in linux using iptables natively and using docker.

10.6.5 Exporting Alerts Module

Slips supports exporting alerts to other systems using different modules (ExportingAlerts, CESNET sharing etc.)

For now the supported systems are:

- Slack
- TAXII Servers (STIX format)
- Warden servers
- IDEA JSON format
- Logstash

Refer to the [exporting section of the docs](#) for detailed instructions on how to export.

10.6.6 Flowalerts Module

This module is responsible for detecting malicious behaviours in your traffic.

Refer to the [Flowalerts section of the docs](#) for detailed explanation of what Slips detects and how it detects.

10.6.7 Disabled alerts

All Slips detections are turned on by default, You can configure which alerts you want to enable/disable in `config/slips.yaml`

Slips support disabling unwanted alerts, simply add the detection you want to disable in the `disabled_detections` list and slips will not generate any alerts of this type.

for example:

```
disabled_detections = [MaliciousJA3, DataExfiltration, SelfSignedCertificate]
```

Supported detections are:

ARPScan, ARP-outside-localnet, UnsolicitedARP, MITM-ARP-attack, SSHSuccessful, LongConnection, MultipleReconnectionAttempts, ConnectionToMultiplePorts, InvalidCertificate, UnknownPort, Port0Connection, ConnectionWithoutDNS, DNSWithoutConnection, MaliciousJA3, DataExfiltration, SelfSignedCertificate, PortScanType1, PortScanType2, Password_Guessing, MaliciousFlow, SuspiciousUserAgent, multiple_google_connections, NETWORK_gps_location_leaked, Command-and-Control-channels-detection, ThreatIntelligenceBlacklistDomain, ThreatIntelligenceBlacklistIP, MaliciousDownloadedFile, DGA, MaliciousSSLCert, YoungDomain, MultipleSSHVersions, DNS-ARPA-Scan, SMTPLoginBruteforce, BadSMTPLogin, IncompatibleUserAgent, ICMP-Timestamp-Scan, ICMP-AddressScan, ICMP-AddressMaskScan

10.6.8 Threat Intelligence Module

Slips has a complex system to deal with Threat Intelligence feeds.

Slips supports different kinds of IoCs from TI feeds (IPs, IP ranges, domains, JA3 hashes, SSL hashes)

File hashes and URLs aren't supported.

Matching of IPs

Slips gets every IP it can find in the network and tries to see if it is in any blacklist.

If a match is found, it generates an evidence, if no exact match is found, it searches the Blacklisted ranges taken from different TI feeds

Matching of Domains

Slips gets every domain that can find in the network and tries to see if it is in any blacklist. The domains are currently taken from:

- DNS requests
- DNS responses
- HTTP host names
- TLS SNI

Once a domain is found, it is verified against the downloaded list of domains from the blacklists defined in `ti_files` in the configuration file `config/slips.yaml`. If an exact match is found, then an evidence is generated.

If an exact match is not found, then Slips verifies if the verified domain is a subdomain of any domain in the blacklist.

For example, if the domain in the traffic is `here.testing.com`, Slips first checks if the exact domain `here.testing.com` is in any blacklist, and if there is no match, it checks if the domain `testing.com` is in any blacklists too.

Matching of JA3 Hashes

Every time Slips encounters an TLS flow, it compares each JA3 and JA3s with the feeds of malicious JA3 and alerts when there's a match. Slips is shipped with the [Abuse.ch](#) JA3 feed by default You can add your own SSL feed by appending to the file defined by the `ja3_feeds` key in `config/slips.yaml`, which is by default `config/JA3_feeds.csv`

Matching of SSL SHA1 Hashes

Every time Slips encounters an SSL flow, it tries to get the certificate hash from zeek `ssl.log`, then it compares the hash with our list of blacklisted SSL certificates

Slips is shipped with the [Abuse.ch](#) SSL feed by default,

You can add your own SSL feed by appending to the file defined by the `ssl_feeds` key in `config/slips.yaml`, which is by default `config/SSL_feeds.csv`

Local Threat Intelligence files

Slips has a local file for adding IoCs of your own, it's located in `config/local_data_files/own_malicious_iocs.csv` by default, this path can be changed by changing `download_path_for_local_threat_intelligence` in `config/slips.yaml`.

The format of the file is "IP address","Threat level", "Description"

Threat level available options: info, low, medium, high, critical

Refer to the [architecture section of the docs](#) for detailed explanation of Slips threat levels.

Example:

```
"23.253.126.58", "high", "Simda CC"
"bncv00.no-ip.info", "critical", "Variant.Zusy"
```

Local JA3 hashes

Slips has a local file for adding JA3 hashes of your own, it's located in `config/local_data_files/own_malicious_JA3.csv` by default.

The format of the file is “JA3 hash”, “Threat level”, “Description”

Threat level available options: info, low, medium, high, critical

Refer to the [architecture section of the docs](#) for detailed explanation of Slips threat levels.

Example:

```
"e7d705a3286e19ea42f587b344ee6865", "medium", "Standard tor client"
"6734f37431670b3ab4292b8f60f29984", "high", "Trickbot Malwar"
```

Adding your own remote feed

We update the remote ones regularly. The list of remote threat intelligence files is set in the variables `ti_files` variable in `config/slips.yaml`. You can add your own remote threat intelligence feeds in this variable. Supported extensions are: `.txt`, `.csv`, `.netset`, `ipsum` feeds, or `.intel`.

Each URL should be added with a `threat_level` and a tag, the format is `(url,threat_level,tag)`

tag is which category is this feed e.g. `phishing`, `adtrackers`, etc..

Threat level available options: info, low, medium, high, critical

Refer to the [architecture section of the docs](#) for detailed explanation of Slips threat levels.

TI files commented using `#` may be processed as they're still in our database.

Use `;` for commenting TI files in `config/slips.yaml` instead of `#`.

Commented TI files (lines starting with `;`) will be completely removed from our database.

The remote files are downloaded to the path set in the `download_path_for_local_threat_intelligence`. By default, the files are stored in the Slips directory `modules/ThreatIntelligence1/remote_data_files/`

Commenting a remote TI feed

If you have a remote file link that you wish to comment and remove from the database you can do so by adding `;` to the line that contains the feed link in `config/slips.yaml`, don't use the `#` for example to comment the `bruteforcelogin` feed you should do the following:

```
;https://lists.blocklist.de/lists/bruteforcelogin.txt,medium,['honeypot']
```

instead of:

```
#https://lists.blocklist.de/lists/bruteforcelogin.txt,medium,['honeypot']
```

10.6.9 Update Manager Module

To make sure Slips is up to date with the most recent IoCs in all feeds, all feeds are loaded, parsed and updated periodically and automatically by Slips every 24 hours, which requires no user interaction.

The 24 hours interval can be changed by changing the `TI_files_update_period` key in `config/slips.yaml`

Update manager is responsible for updating all remote TI files (including SSL and JA3 etc.)

By default, local slips files (`organization_info`, `ports_info`, etc.) are cached to avoid loading and parsing them everytime we start slips. However, they are updated automatically by the update manager if they were changed on disk.

10.6.10 IP Info Module

The IP info module has several ways of getting information about IP and MAC address, it includes:

- ASN
- Country by Geolocation
- MAC Vendors
- Reverse DNS

ASN

Slips is shipped with an offline database (GeoLite2) in `databases/GeoLite2-ASN.mmdb` to search for ASNs, if the ASN of a given IP is not in the GeoLite2 database, we try to get the ASN online using the online database using the `ipwhois` library. However, to reduce the amount of requests, we retrieve the range of the IP and we cache the whole range. To search and cache the whole range of an IP, the module uses the `ipwhois` library. The `ipwhois` library gets the range of this IP by making a connection to the server `cymru.com` using a TXT DNS query. The DNS server is the one set up in the operating system. For example to get the ASN of the IP 13.32.98.150, you will see a DNS connection asking for the TXT record of the domain `150.98.32.13.origin.asn.cymru.com`.

Country by Geolocation

Slips is shipped with an offline database (GeoLite2) in `databases/GeoLite2-Country.mmdb` to search for Geolocation.

Mac Vendors

Slips is shipped with an offline database `databases/macaddress-db.json` for MAC address vendor mapping.

This database is a combination of 2 different online databases, but the format of them is changed to a format slips understands and to reduce the size of the db.

Slips gets the MAC address of each IP from `dhcp.log` and `arp.log` and then searches the offline database using the OUI.

If the vendor isn't found in the offline MAC database, Slips tries to get the MAC using the online database <https://www.macvendorlookup.com>

The offline database is updated manually and shipped with slips, you can find it in the `databases/` dir.

Slips makes sure it doesn't perform duplicate searches of the same MAC Address either online, or offline.

Reverse DNS

This is obtained by doing a standard `in-addr.arpa` DNS request.

10.6.11 ARP Module

This module is used to check for ARP attacks in your network traffic.

By default, zeek doesn't generate and log ARP flows, but Slips is shipped with its own zeek scripts that enable the logging of ARP flows in `arp.log`

The detection techniques are:

- ARP scans
- ARP to a destination IP outside of local network
- Unsolicited ARP
- MITM ARP attack

ARP Scans

Slips considers an IP performing an ARP scan if it sends 5 or more non-gratuitous ARP to different destination addresses in 30 seconds or less.

ARP to a destination IP outside of local network

Slips alerts when an ARP flow is being sent to an IP outside of local network as it's a weird behaviour that shouldn't be happening.

Unsolicited ARP

Unsolicited ARP is used to update the neighbours' ARP caches but can also be used in ARP spoofing, we detect it with threat level 'info', so we don't consider it malicious, we simply notify you about it.

MITM ARP attack

Slips detects when a MAC with IP A, is trying to tell others that now that MAC is also for IP B (ARP cache attack)

10.6.12 CESNET sharing Module

This module is responsible for importing and exporting alerts from and to warden server

Refer to the [exporting section of the docs](#) for detailed instructions on CESNET exporting and the format of the configuration files.

To enable the importing alerts from warden servers, set `receive_alerts` to `yes` in `config/slips.yaml`

Slips imports 100 alerts from warden servers each day, and automatically stores the alerts in our database

Time to wait before receiving alerts from warden server is 1 day by default, you can change this by changing the `receive_delay` in `config/slips.yaml`

These are the categories we import: ['Availability', 'Abusive.Spam', 'Attempt.Login', 'Attempt', 'Information', 'Fraud.Scam', 'Information', 'Fraud.Scam']

10.6.13 HTTP Analyzer Module

This module handles the detections of HTTP flows

Available detection are:

- Multiple empty connections
- Suspicious user agents

- Incompatible user agents
- Multiple user agents
- Downloads from pastebin
- Executable downloads

Multiple empty connections

Due to the usage of empty connections to popular site by malware to check for internet connectivity, We consider this type of behaviour suspicious activity that shouldn't happen

We detect empty connection to 'bing.com', 'google.com', 'yandex.com', 'yahoo.com', 'duckduckgo.com' etc.

If Google is whitelisted in `whitelist.conf`, this detection will be suppressed.

Suspicious user agents

Slips has a list of suspicious user agents, whenever one of them is found in the traffic, slips generates and evidence.

Our current list of user agents has: ['httpsend', 'chm_msdn', 'pb', 'jndi', 'tesseract']

Incompatible user agents

Slips uses and offline MAC address database to detect the type of device based on the MAC OUI.

First, Slips store the MAC address and vendor of every IP it sees (if available)

Second, When slips encounters a user agent in HTTP traffic it performs an online query to <http://useragentstring.com> to get more info about this user agent, like the os type, name and browser.

Third, When slips has both information available (MAC vendor and user agent), it compares them to detect incompatibility using a list of keywords for each operating system.

Available keywords for Apple: ('macos', 'ios', 'apple', 'os x', 'mac', 'macintosh', 'darwin')

Available keywords for Microsoft: ('microsoft', 'windows', 'nt')

Available keywords for Android: ('android', 'google')

Multiple user agents

Slips stores the MAC address and vendor of every IP it sees (if available) in the redis database. Then, when an IP is seen using a different user agent than the one stored in the database, it tries to extract os info from the user agent string, either by performing an online query to <http://useragentstring.com> or by using zeek.

If an IP is detected using different user agents that refer to different operating systems, an alert of type 'Multiple user agents' is made

for example, if an IP is detected using a macOS user agent then an android user agent, slips detects this with 'low' threat level

Pastebin downloads

Some malware use pastebin as the host of their malicious payloads.

Slips detects downloads of files from pastebin with size ≥ 700 bytes.

This value can be customized in `slips.yaml` by changing `pastebin_download_threshold`

When found, slips alerts pastebin download with threat level low because not all downloads from pastebin are malicious.

Executable downloads

Slips generates an evidence everytime there's an executable download from an HTTP website.

10.6.14 Leak Detector Module

This module work only when slips is given a PCAP

The leak detector module uses YARA rules to detect leaks in PCAPs

Module requirements

In order for this module to run you need:

You can install YARA by running

```
sudo apt install yara
```

You can install tshark by running

```
sudo apt install wireshark
```

How it works

This module works by

1. Compiling the YARA rules in the `modules/leak_detector/yara_rules/rules/` directory
2. Saving the compiled rules in `modules/leak_detector/yara_rules/compiled/`
3. Running the compiled rules on the given PCAP
4. Once we find a match, we get the packet containing this match and set evidence.

Extending

You can extend the module be adding more YARA rules in `modules/leak_detector/yara_rules/rules/`.

The rules will be automatically detected, compiled and run on the given PCAP.

If you want to contribute, improve existing Slips detection modules or implement your own detection modules, see section [:doc:Contributing <contributing>](#).

10.6.15 Network service discovery Module

This module is responsibe for detecting scans such as:

- Vertical port scans
- Horizontal port scans
- PING sweeps
- DHCP scans

Vertical port scans

Slips checks both TCP and UDP connections for port scans.

Slips considers an IP performing a vertical port scan if it scans 6 or more different destination ports

Slips reports evidence of a vertical portscan based on a log scale. meaning every 10, 100, 1000 ports scanned will generate an evidence. This avoids generating one port scan alert per flow in a long scan.

Horizontal port scans

Slips checks both TCP and UDP connections for horizontal port scans.

Slips considers an IP performing a horizontal port scan if it contacted more than 3 destination IPs on a specific port with not established connections.

Slips reports evidence of a horizontal portscan based on a log scale. meaning every 10, 100, 1000 ports scanned will generate an evidence. This avoids generating one port scan alert per flow in a long scan.

PING Sweeps

PING sweeps or ICMP sweeps is used to find out which hosts are alive in a network or large number of IP addresses using PING/ICMP.

We detect a scan every threshold. So we generate an evidence when there is 5,10,15, .. etc. ICMP established connections to different IPs.

We detect 3 types of ICMP scans: ICMP-Timestamp-Scan, ICMP-AddressScan, and ICMP-AddressMaskScan

Slips does this detection using Slips' own zeek script located in `zeek-scripts/icmps-scans.zeek` for zeek and pcap files and using the portscan module for binetflow files.

10.7 Connections Made By Slips

Slips uses online databases to query information about many different things, for example (user agents, mac vendors etc.)

The list below contains all connections made by Slips

[useragentstring.com](#) -> For getting user agent info if no info was found in Zeek [macvendorlookup.com](#) -> For getting MAC vendor info if no info was found in the local maxmind db [ip-api.com/json/](#) -> For getting ASN info about IPs if no info was found in our Redis DB [ipinfo.io/json](#) -> For getting your public IP [virustotal.com](#) -> For getting scores about downloaded files, domains, IPs and URLs [cymru.com](#) -> For getting the range of a specific IP to cache the ASN of this range. TXT DNS queries are made to this domain.

By default, slips whitelists alerts from or to any of the above domains, witch means that if an alert was detected to one of the above alerts, slips does not detect it assuming it's a false positive and the connection was made internally by slips.

You can change this behaviour by copying `config/whitelist.conf`, updating the copy, and pointing `whitelists.local_whitelist_path` in your copied `slips.yaml` file to that copied whitelist.

10.8 Ensembling

Ensembling in Slips is done by the Evidence Process.

Every time the evidence process gets a new evidence from the detection modules, it retrieves all the past evidence by this the source IP, in the current timewindow from the database.

Then, Slips uses the following equation to get the score of each evidence

$$\text{threat_level} = \text{threat_level} * \text{confidence}$$

Slips accumulates the threat level of all evidenc, then, it checks if the accumulated threat level reached a certain threshold or not.

If the accumulated threat level reached the threshold specified in `evidence_detection_threshold`, Slips generates and alert. If not, slips waits for the next evidence, accumulates threat levels, and checks again until the threshold is reached.

10.9 Controlling Slips Sensitivity

The threshold that controls Slips sensitivity is determined by the `evidence_detection_threshold` key in `config/slips.yaml`, by default it is set to `0.25`.

This threshold is used in slips according to the following equation

$$\text{threshold per width} = \text{detection_threshold} * \text{width} / 60$$

For example, if you're using the default slips width 3600, the threshold used in slips will be

$$0.25 * 3600 / 60 = 15$$

This equation's goal is to make it more sensitive on smaller tws, and less sensitive on longer tws

When the accumulated threat levels of all evidence detected in a timewindow exceeds 15, slips will generate an alert.

In simple terms, it means slips will alert when users get the equivalent of 1 alert per minute.

The default threshold of 0.25 gives you balanced detections with the optimal false positive rate and accuracy.

Here are more options

- 0.08: Use this threshold If you want Slips to be super sensitive with higher FPR, using this means you are less likely to miss a detection but more likely to get false positives
- 0.25: Optimal threshold, has the most optimal FPR and TPR.
- 0.43: Use this threshold If you want Slips to be insensitive. Using this means Slips will need so many evidence to trigger an alert. May lead to false negatives

10.10 Zeek Scripts

Slips is shipped with it's own custom zeek scripts to be able to extend zeek functionality and customize the detections

10.10.1 Detect DoH

In the `detect_DoH.zeek` script, slips has it's own list of ips that belong to dns/doh servers,

When slips encounters a connection to any IP of that list on port 443/tcp, it assumes it's a DoH connetion,

and times out the connection after 1h so that the connection won't take too long to appear in slips.

10.10.2 Detect ICMP Scans

In the `zeek-scripts/icmps-scans.zeek` script, we check the type of ICMP in every ICMP packet seen in the network,

and we detect 3 types of ICMP scans: `ICMP-Timestamp-Scan`, `ICMP-AddressScan`, and `ICMP-AddressMaskScan` based on the `icmp` type

We detect a scan every threshold. So we generate an evidence when there is 5,10,15, .. etc. ICMP established connections to different IPs.

10.10.3 CPU Profiling

Slips is shipped with its own tool for CPU Profiling, it can be found it `slips_files/common/cpu_profiler.py`

CPU Profiling supports 2 modes: live and development mode

Live mode:

The main purpose of this mode is to show live CPU stats in the web interface. “live” mode publishes updates during the runtime of the program to the redis channel ‘cpu_profile’ so that the web interface can use them

Development mode:

Setting the mode to “dev” outputs a JSON file of the CPU usage at the end of the program run. It is recommended to only use dev mode for static file inputs (pcaps, suricata files, binetflows, etc.) instead of interface and growing zeek dirs, because longer runs result in profiling data loss and not everything will get recorded. The JSON file created in this mode is placed in the output dir of the current run and can be viewed by running the following command

```
vizviewer results.json
```

then going to <http://127.0.0.1:9001/> in your browser for seeing the visualizations of the CPU usage

Options to enable cpu profiling can be found under the [Profiling] section of the `slips.yaml` file. `cpu_profiler_enable` set to “yes” enables cpu profiling, or “no” to disable it. `cpu_profiler_mode` can be set to “live” or “dev”. Setting to `cpu_profiler_multiprocess` can be set to “yes” or “no” and only affects the dev mode profiling. If set to “yes” then all processes will be profiled. If set to “no” then only the main process (`slips.py`) will be profiled. `cpu_profiler_output_limit` is set to an integer value and only affects the live mode profiling. This option sets the limit on the number of processes output for live mode profiling updates. `cpu_profiler_sampling_interval` is set to an integer value and only affects the live mode profiling. This option sets the duration in seconds of live mode sampling intervals. It is recommended to set this option greater than 10 seconds otherwise there won’t be much useful information captured during sampling.

10.10.4 Memory Profiling

Memory profiling can be found in `slips_files/common/memory_profiler.py`

Just like CPU profiling, it also has supports live and development mode. Set `memory_profiler_enable` to `yes` to enable this feature. Set `memory_profiler_mode` to `live` to use live mode or `dev` to use development mode profiling.

Live Mode

This mode shows memory usage stats during the runtime of the program. `memory_profiler_multiprocess` controls whether live mode tracks all processes or only the main process. If set to `no`, the program will wait for you to connect from a different terminal using the command `memray live <port_number>`, where `port_number` is 5000 by default. After connection, the program will continue with its run and the terminal that is connected will receive a feed of the memory statistics. If set to `yes`, the redis channel “memory_profile” can be used to set pid of the process to be tracked. Only a single process can be tracked at a time. The interface is cumbersome to use from the command line so multiprocessing live profiling is intended to be used primarily from the web interface.

Development Mode

When enabled, the profiler will output the profile data into the output directory. The data will be in the `memoryprofile` directory of the output directory of the run. Each process during the run of the program will have an associated binary file. Each of the generated binaries will automatically be converted to viewable html files, with each process converted to a flamegraph and table format. All generated files will be denoted by their PID.

If you want to contribute: improve existing Slips detection modules or implement your own detection modules, see section `:doc:Contributing <contributing>`.

TRAINING

Slips has one machine learning module that can be retrained by users. This is done by putting Slips in training mode so you can re-train the machine learning models with your own traffic. By default Slips includes an already trained model with our data, but it is sometimes necessary to adapt it to your own circumstances.

Until Slips 0.7.3, there is only one module for now that can do this, the one called 'flow_ml_detection'. This module analyzes flows one by one, as formatted similarly as in a conn.log Zeek file. This module is enabled by default in testing mode. This module uses by default the SGDClassifier with a linear support vector machine (SVM). The decision to use SVM was done because it is one of the few algorithms that can be used for online learning and that can extend a current model with new data.

To re-train this machine learning algorithm, you need to do the following:

1- Edit the config/slips.yaml file to put Slips in train mode. Search the word **train** in the section **[flow_ml_detection]** and uncomment the **mode = train** and comment **mode = test**. It should look like

```
[flow_ml_detection]
# The mode 'train' should be used to tell the flow_ml_detection module that the flows
↳received are all for training.
# A label should be provided in the [Parameters] section
mode = train

# The mode 'test' should be used after training the models, to test in unknown data.
# You should have trained at least once with 'Normal' data and once with 'Malicious'
↳data in order for the test to work.
#mode = test
```

2- Establish the general label for all the traffic that you want to re-train with. For now we only support 1 label per file. Search in the [parameters] section and choose the type of traffic you will send to Slips.

```
# Set the label for all the flows that are being read. For now only normal and malware
↳directly. No option for setting labels with a filter
label = normal
#label = malicious
#label = unknown
```

After these edits, just run Slips as usual with any type of input, for example with a Zeek folder.

```
./slips.py -c config/slips.yaml -f ~/my-computer-normal/
```

Or with a pcap file.

```
./slips.py -c config/slips.yaml -f ~/my-computer-normal2.pcap
```

3- If you have also malicious traffic, first change the label to malicious in config/slips.yaml

```
# Set the label for all the flows that are being read. For now only normal and malware.
↳directly. No option for setting labels with a filter
#label = normal
label = malicious
#label = unknown

./slips.py -c config/slips.yaml -f ~/my-computer-normal2.pcap
```

After this edits, just run Slips as usual with any type of input, for example another pcap

```
./slips.py -c config/slips.yaml -f ~/malware1.pcap
```

You can also run slips in an interface and train it directly with your data

```
./slips.py -c config/slips.yaml -i eth0
```

4- Finally to use the model, put back the **test** mode in the configuration config/slips.yaml

```
[flow_ml_detection]
# The mode 'train' should be used to tell the flow_ml_detection module that the flows.
↳received are all for training.
# A label should be provided in the [Parameters] section
#mode = train

# The mode 'test' should be used after training the models, to test in unknown data.
# You should have trained at least once with 'Normal' data and once with 'Malicious'.
↳data in order for the test to work.
mode = test
```

5- Use slips normally in files or interfaces

```
./slips.py -c config/slips.yaml -i eth0
```


10. Edit the `config/slips.yaml` file, put `slack` in the `export_to` variable, and add the channel's name to which you want to send.

```
[exporting_alerts] export_to = [slack] slack_channel_name = SlipsAlertsChannel
```

12.2 STIX

If you want to export alerts to your TAXII 2 server using STIX 2.1 format, set `export_to` to `stix` and Slips will automatically generate a `STIX_data.json` bundle containing the indicators it detects and push it to your collection.

```
[ExportingAlerts]
export_to = [stix]
```

Configure the TAXII client by editing the following variables:

`TAXII_server`: host name or IP address of the TAXII server.

`port`: TCP port (optional, defaults to 80/443).

`use_https`: set to true to connect over HTTPS (be careful that the default TAXII server in SlipsWeb, Medallion, do not support HTTPS yet)

`discovery_path`: TAXII discovery endpoint path or full URL (for example `/taxii2/`).

`taxii_version`: set to 2 for TAXII 2.1 (Medallion) or 1 for TAXII 1.x (OpenTAXII). TAXII 1 export uses the inbox service and is sent via direct export.

`collection_name`: ID or title of the TAXII collection that should receive your indicators. Be default Alerts.

`push_delay`: time between automatic pushes (in seconds) when Slips is running continuously.

`taxii_username` / `taxii_password`: credentials used for HTTP Basic authentication.

Change the default config password of the TAXII servers you are going to export to in `config/medallion_config.yaml`

Slips stores the generated bundle for each run in the output directory of that execution (for example `output/<run_id>/STIX_data.json`), so you can inspect the exact STIX objects that were pushed.

If running on a file, Slips will export once before shutdown. If running on an interface, Slips will export to the server every `push_delay` seconds (default 1 hour).

12.3 JSON format

By default Slips logs all alerts to `output/alerts.json` in CESNET's [IDEA0 format](#) which is also a JSON format.

12.4 CESNET Sharing

Slips supports exporting alerts to warden servers, as well as importing alerts.

To enable the exporting, set `receive_alerts` to `yes` in `config/slips.yaml`

The default configuration file path is specified in the `configuration_file` variable in `config/slips.yaml`

The default path is `config/warden.conf`

The format of `warden.conf` should be the following:

```
{ "url": "https://example.com/warden3",
  "certfile": "cert.pem",
  "keyfile": "key.pem",
  "cafile": "/etc/ssl/certs/DigiCert_Assured_ID_Root_CA.pem",
  "timeout": 600,
  "errlog": {"file": "output/warden_logs/warden.err", "level": "debug"},
  "filelog": {"file": "output/warden_logs/warden.log", "level": "warning"},
  "name": "com.example.warden.test" }
```

To get your key and the certificate, you need to run `warden_apply.sh` with you registered `client_name` and password. [Full instructions here](#)

The name key is your registered warden node name.

All evidence causing an alert are exported to warden server once an alert is generated. See the [difference between alerts and evidence](#)) in Slips architecture section.

You can change how often you get alerts (import) from warden server

By default Slips imports alerts every 1 day, you can change this by changing the `receive_delay` value in `config/slips.yaml`

Slips logs all alerts to `output/alerts.json` in CESNET's IDEA0 format by default.

Make sure that the `DigiCert_Assured_ID_Root_CA` is somewhere accessible by slips. or run slips with root if you want to leave it in `/etc/ssl/certs/`

Refer to the [Detection modules section of the docs](#) for detailed instructions on how CESNET importing.

12.5 Logstash

Slips has `logstash.conf` file that exports our alerts.json to a given output file, you can change the output to your preference (for example: elastic search, stdout, etc.)

12.6 Text logs

By default, the output of Slips is stored in the `output/` directory in two files:

1. `alert.json` in IDEA0 format
2. `alerts.log` human readable text format

12.7 TSV and json of labeled flows

Slips supports exporting all the labeled flows and altflows stored in the sqlite database the sqlite database can be exported to json or tsv format.

Each labeled flow has an [AID fingerprint](#), which is used to identify the flow based on the ts, source and destination address, source and destination port and protocol.

this can be done by setting the `export_labeled_flows` parameter to `yes` in `slips.yaml` and changing the `export_format` parameter to your desired format. for now, the `export_format` parameter supports tsv or json formats only.

the exported flows are stored in a file called `labeled_flows.json` or `labeled_flows.tsv` in the output directory.

P2P

The P2P module makes Slips be a peer in a peer to peer network of computers in the local network. The peers are only in the local network and they communicate using multicast packets. The P2P module is a highly complex system of data sharing, reports on malicious computers, asking about IoC to the peers and a complex trust model that is designed to resist adversarial peers in the network. Adversarial peers are malicious peers that lie about the data being shared (like saying that a computer is malicious when it is not, or that an attacker is benign).

This module was designed and partially implemented in a [Master Thesis](#) on CTU FEL by [Dita Hollmannova](#). The goal was to enable Slips instances in a local network to share detections and collectively improve blocking decisions. While the thesis succeeded in creating a framework and a trust model, the project is far from stable. The final implementation in Slips was finished by [Alya Goma](#).

This readme provides a shallow overview of the code structure, to briefly document the code for future developers. The whole architecture was thoroughly documented in the thesis itself, which can be downloaded from the link above.

The basic structure of the P2P system is (i) an Slips P2P module in python (called Dovecot), and (ii) a P2P communication system done in Golang (called Pigeon).

13.1 Pigeon

Pigeon is written in golang and is developed in an independent repository from Slips, but is included as a submodule of Slips repository. <https://github.com/stratosphereips/p2p4slips>

Pigeon handles the P2P communication in the network using the libp2p library, and provides a simple interface to the Slips module. A compiled Pigeon binary is included in the module for convenience.

Pigeon uses the JSON format to communicate with the module or with other Pigeons. For details on the communication format, see the thesis.

13.2 Docker direct use

You can use Slips with P2P directly in a special docker image by doing:

```
docker pull stratosphereips/slips
docker run -it --rm --net=host --cap-add=NET_ADMIN stratosphereips/slips
```

For the p2p to be able to listen on the network interfaces and receive packets you should use `--cap-add=NET_ADMIN`

13.3 Installation:

1. download and install go:

```
apt install golang
```

or by hand

```
curl https://dl.google.com/go/go1.18.linux-amd64.tar.gz --output go.tar.gz
rm -rf /usr/local/go && tar -C /usr/local -xzf go.tar.gz
export PATH=$PATH:/usr/local/go/bin
```

2. build the pigeon:

- if you installed slips with the submodules using

```
git clone --recurse-submodules --remote-submodules https://github.com/stratosphereips/
↳StratosphereLinuxIPS -j4
```

then you should only build the pigeon using: `cd p2p4slips && go build -buildvcs=false`

- If you installed Slips without the submodules then you should download and build the pigeon using:

```
git submodule init && git submodule update && cd p2p4slips && go build -buildvcs=false
```

The p2p binary should now be in `p2p4slips/` dir and slips will be able to find it.

NOTE

If you installed the p2p4slips submodule anywhere other than slips main directory, remember to add it to PATH by using the following commands:

```
echo "export PATH=$PATH:/path/to/StratosphereLinuxIPS/p2p4slips/" >> ~/.bashrc
source ~/.bashrc
```

13.4 Usage in Slips

The P2P module is disabled by default in Slips.

To enable it, change `use_p2p=no` to `use_p2p=yes` in `config/slips.yaml`

P2P is only available when running slips in you local network using an interface. (with `-i`)

You don't have to do anything in particular for the P2P module to work, just enable it and Slips will: 1- Automatically find other peers in the network (and remember them even if they go offline for days)

2- Ask the group of peers (the network) about what they think of some IoC

3- Group the answers and give Slips an aggregated, balanced, normalized view of the network opinion on each IoC

4- Send blame reports to the whole network about attackers

5- Receive blame reports on attackers from the network, balanced by the trust model

6- Keep a trust score on each peer, which varies in time based on the interactions and quality of data shared

13.5 Project sections

The project is built into Slips as a module and uses Redis for communication. Integration with Slips is seamless, and it should be easy to adjust the module for use with other IPSs. The following code is related to Dovecot:

- Slips, the Intrusion Prevention System

- Dovecot module, the module for Slips
- Pigeon, a P2P wrapper written in golang
- Dovecot experiments, a framework for evaluating trust models (optional)

13.6 Dovecot experiments

Experiments are not essential to the module, and the whole project runs just fine without them. They are useful for development of new trust models and modelling behavior of the P2P network.

To use the experiments, clone the <https://github.com/stratosphereips/p2p4slips-experiments> repository into `modules/p2p_trust/testing/experiments`.

The experiments run independently (outside of Slips) and start all processes that are needed, including relevant parts of Slips. The code needs to be placed inside the module, so that necessary dependencies are accessible. This is not the best design choice, but it was the simplest quick solution.

13.7 How it works:

Slips interacts with other slips peers for the following purposes:

13.7.1 Blaming IPs

If slips finds that an IP is malicious given enough evidence, it blocks it and tells other peers that this IP is malicious and they need to block it. this is called sending a blame report.

13.7.2 Receiving Blames

When slips receives a blame report from the network, which means some other slips instance in th network set an evidence about an IP and is letting other peers know about it.

Slips then generates an evidence about the reported IP and takes the report into consideration when deciding to block the attacker's IP.

13.7.3 Asking the network about an IP

Whenever slips sees a new IP, it asks other peers about it, and waits 3 seconds for them to reply.

The network then replies with a score and confidence for the IP. The higher the score the more malicious this IP is.

Once we get the score of the IP, we store it in the database, and we alert if the score of this IP is more than 0 (threat level=info).

The persistent local P2P runtime directory is stored under the directory configured by `parameters.permanent_dir` in `config/slips.yaml`. By default, this is `permanent/p2p_trust_runtime/`.

13.7.4 Answering the network's request about an IP

When asked about an ip, slips shares the score of it and the confidence with the requesting peer. the scores are generated by slips and saved in the database.

13.8 Logs

Slips contains a minimal log file for reports received by other peers and peer updates in `output/p2p_reports.log`

For a more detailed p2p logs, for example (peer ping pongs, peer lists, errors, etc.) you can enable p2p.log in slips.yaml by setting `create_p2p_logfile` to `yes` and a `p2p.log` will be available in the output dir

Slips rotates the p2p.log every 1 day by default, and keeps the logs of 1 past day only.

13.9 Limitations

For now, slips only supports requests and blames about IPs.

Domains, URLs, or hashes are not supported, but can easily be added in the future.

13.10 TLDR;

Slips only shares scores and confidence (numbers) generated by slips about IPs to the network, no private information is shared.

FIDES MODULE. GLOBAL P2P THREAT INTELLIGENCE SHARING

Slips implements an internet global P2P system for Threat Intelligence sharing and alerting.

The Fides module implements the Global P2P system in Slips.

Traditional network defense systems depend on centralized threat intelligence, which has limitations like single points of failure, inflexibility, and reliance on trust in centralized authorities. Peer-to-peer networks offer an alternative for sharing threat intelligence but face challenges in verifying the trustworthiness of participants, including potential malicious actors.

The Fides Module is based on the [Master Thesis](#) of Lukáš Forst and implemented in Slips in the Master Thesis of David Otta. The goal of the Fides module is to address the challenge of trust of peers in P2P networks by providing several trust evaluation models. It evaluates peer behavior, considers membership in trusted organizations, and assesses incoming threat data to determine reliability. Fides aggregates and weights data to enhance intrusion prevention systems, even in adversarial scenarios. Experiments show that Fides can maintain accurate threat intelligence even when 75% of the network is controlled by malicious actors, assuming the remaining 25% are trusted.

The whole architecture is thoroughly documented in the thesis itself, which can be downloaded from the link above.

14.1 Docker direct use

You can use Slips with the Fides Module by allowing it in the Slips config file or by using the following commands.

```
docker pull stratosphereips/slips
docker run -it --rm --net=host --use_fides=True --cap-add=NET_ADMIN stratosphereips/slips
```

To be able to use the fides module, you should use `--cap-add=NET_ADMIN`

14.2 Conditions

If you plan on using the Fides Module, please be aware that it is used only if Slips is running on an interface OR on a growing Zeek directory. The `--use_fides=True` is ignored when Slips is run on a file.

14.3 Configuration

The evaluation model used, the evaluation thresholds, and other configurations are located in `modules/fides/config/fides.conf.yml`.

If you need a Slips run to use a different Fides configuration file, set `global_p2p.fides_conf` in Slips config to the relative path of that alternate YAML file.

Possible threat intelligence evaluation models

Model Name	Description
average	Average Confidence Trust Intelligence Aggregation
weightedAverage	Weighted Average Confidence Trust Intelligence Aggregation
stdevFromScore	Standard Deviation From Score Trust Intelligence Aggregation

14.4 Usage in Slips

Fides is inactive by default in Slips.

To enable it, change `use_fides=False` to `use_fides=True` in `config/slips.yaml`.

And start Slips on your interface.

The Fides shared SQLite cache is stored under the directory configured by `parameters.permanent_dir`. By default, it is created in `permanent/databases/`, so it persists across different Slips runs.

14.5 How it works:

Slips interacts with other slips peers for the following purposes:

14.5.1 Sharing an opinion with peers

If peer A is asked for its opinion on peer B by peer C, peer A sends its opinion on peer B to peer C, if there is any.

14.5.2 Asking for an opinion

Peers can ask other peers what they think about an IP address or domain.

14.5.3 Dispatching alerts

If a peer generates an alert based on evidence of an attack, it can alert other peers by sending an **Alert message** in the P2P network.

14.6 Logs

Slips contains a minimal log file for reports received by other peers and peer updates in the `output` directory if not manually specified using the appropriate `slips` parameter upon start. The custom logger `modules/fides/utills/logger.py` code is used by the Fides Module for internal logging.

14.7 Implementation notes and credit

The mathematical models for the trust evaluation were written by Lukáš Forst as part of his [Master Thesis](#).

14.8 Privacy

Slips only shares the trust level and confidence values generated by Slips about IPs to the network, no more information.

HOW TO CREATE A NEW SLIPS MODULE

15.1 What is SLIPS and why are modules useful

Slips is a machine learning-based intrusion prevention system for Linux and MacOS, developed at the Stratosphere Laboratories from the Czech Technical University in Prague. Slips reads network traffic flows from several sources, applies multiple detections (including machine learning detections) and detects infected computers and attackers in the network. It is easy to extend the functionality of Slips by writing a new module. This blog shows how to create a new module for Slips from scratch.

15.2 Goal of this Blog

This blog creates an example module to detect when any private IP address communicates with another private IP address. What we want is to know if, for example, the IP 192.168.4.2, is communicating with the IP 192.168.4.87. This simple idea, but still useful, is going to be the purpose of our module. Also, it will generate an alert for Slips to consider this situation. Our module will be called `local_connection_detector`.

15.2.1 High-level View of how a Module Works

All modules implement the `IModule` interface located at `slips_files/common/abstracts/module.py`. Abstract methods in this interface must be implemented in every new slips module, the rest are optional.

Below is a detailed description of each abstract method.

The Module consists of the `init()` function for initializations, like subscribing to channels, reading API files, etc.

The main function of each module is the `main()`, this function is run in a loop that keeps looping as long as Slips is running so that the module doesn't terminate.

In case of errors in the module, the `main()` function should return 1 which will cause the module to immediately terminate.

any initializations that should be run only once should be placed in the `init()` function OR the `pre_main()`. the `pre_main()` is a function that acts as a hook for the main function. it runs only once and then the main starts running in a loop. the `pre_main()` is the place for initialization logic that cannot be done in the `init`, for example dropping the root privileges from a module. we'll discuss this in detail later.

Printing in all modules is handled by a common `print()` method, the one implemented in the `IModule` interface. All this common `print()` does is acts as a proxy between the module responsible for printing, `output.py`, and all slips modules.



Each Module has its own `shutdown_gracefully()` function that handles cleaning up after the module is done processing. It handles for example:

- Saving a model before Slips stops
- Saving alerts in a .txt file if the module's job is to export alerts
- Telling the main module (`slips.py`) that the module is done processing so `slips.py` can kill it etc.

15.3 Creating a Module

When Slips runs, it automatically loads all the modules inside the `modules/` directory. Therefore, our new module should be placed there.

Slips has a template module directory that we are going to copy and then modify for our purposes.

```
cp -a modules/template modules/local_connection_detector
```

15.3.1 Changing the Name of the Module

Each module in Slips should have a name, author and description. Use a snake_case module package and main file name in the `x_y_doer` style already used in the repository, for example `http_analyzer` or `local_connection_detector`.

We should change the name inside the python file by finding the lines with the name and description in the class 'Module' and changing them:

```
name = 'local_connection_detector'
description = (
    'detects connections to other devices in your local network'
)
authors = ['Your name']
```

Also change the name of the class to something like this

```
class LocalConnectionDetector(IModule):
    ...
```

At the end you should have a structure like this:

```
modules/
├─ local_connection_detector/
│   └─ __init__.py
│   └─ local_connection_detector.py
```

The `init.py` is to make sure the module is treated as a python package, don't delete it.

Remember to delete the `pycache` dir if it's copied to the new module using:

```
rm -r modules/local_connection_detector/__pycache__
```

15.3.2 Redis Pub/Sub

First, some initialization in the `init()`:

1. we need to subscribe to the channel `new_flow`
2. to be able to convert the flows received in the above channel from dict format to objects, we'll need a classifier

```
self.c1 = self.db.subscribe('new_flow')

# add this channel to the module's list of channels
# this list will be used to get msgs from the channel later
self.channels = {
    'new_flow': self.c1,
}

# to be able to convert flows from dict format to objects
self.classifier = FlowClassifier()
```

So now everytime slips sees a new flow, you can access it from your module using the following line

```
msg = self.get_msg('new_flow')
```

the implementation of the `get_msg()` is placed in the abstract module in `slips_files/common/abstracts/module.py` and is inherited by all modules.

The above line checks if a message was received from the `new_flow` channel that you subscribed to.

Now, you can access the content of the flow using

```
flow = msg['data']
```

Thus far, we have the following code to prep the module for receiving new flows

```
def init(self):
    self.c1 = self.db.subscribe('new_flow')
    self.channels = {
        'new_flow': self.c1,
    }
    # to be able to convert flows from dict format to objects
    self.classifier = FlowClassifier()
```

```
def pre_main(
    self
):
    """
    Initializations that run only once before the main() function runs in a loop
    """
    utils.drop_root_privs_permanently()
```

```
def main(self):
    """Main loop function"""
    if msg:= self.get_msg('new_flow'):
        #TODO
        pass
```

15.3.3 Detecting connections to local devices

Now that we have the flow, we need to:

- Extract the source IP

- Extract the destination IP
- Check if both of them are private
- Generate an evidence

Extracting IPs is done by the following:

```
msg = json.loads(msg['data'])
# convert the given dict flow to a flow object
flow = self.classifier.convert_to_flow_obj(msg["flow"])
saddr = flow.saddr
daddr = flow.daddr
timestamp = flow.starttime
```

The above snippet should be in the main() function, since we wanna repeat it in a loop everytime we get a new flow

Now we need to check if both of them are private.

```
import ipaddress
srcip_obj = ipaddress.ip_address(saddr)
dstip_obj = ipaddress.ip_address(daddr)
if srcip_obj.is_private and dstip_obj.is_private:
    #TODO
    pass
```

Now that we're sure both IPs are private, we need to generate an alert.

Slips requires certain info about the evidence to be able to deal with them.

First, since we are creating a new type of evidence that is not defined in the EvidenceType Enum in slips_files/core/structure/evidence.py, we need to add a new type there.

so the EvidenceType Enum in slips_files/core/structures/evidence.py would look something like this

```
class EvidenceType(Enum):
    """
    These are the types of evidence slips can detect
    """
    ...
    CONNECTION_TO_LOCAL_DEVICE = auto()
    ...
```

Now we have our evidence type supported. it's time to set the evidence!

Now we need to use the Evidence structure of slips, to do that, First import the necessary dataclasses

```
from slips_files.core.evidence_structure.evidence import \
(
    Evidence,
    ProfileID,
    TimeWindow,
    Victim,
    Attacker,
    ThreatLevel,
    EvidenceType,
    IoCType,
```

(continues on next page)

(continued from previous page)

```

    Direction,
)

```

now use them,

```

# on a scale of 0 to 1, how confident you are of this evidence
confidence = 0.8
# how dangerous is this evidence? info, low, medium, high, critical?
threat_level = ThreatLevel.HIGH

# the name of your evidence, you can put any descriptive string here
# this is the type we just created
evidence_type = EvidenceType.CONNECTION_TO_LOCAL_DEVICE
# which ip is the attacker here?
attacker = Attacker(
    direction=Direction.SRC, # who's the attacker the src or the dst?
    attacker_type=IoCType.IP, # is it an IP? is it a domain? etc.
    value=saddr # the actual ip/domain/url of the attacker, in our case, this is the_
↪IP
)
victim = Victim(
    direction=Direction.SRC,
    victim_type=IoCType.IP,
    value=daddr,
)
# describe the evidence
description = f'A connection to a local device {daddr}'
# the current profile is the source ip,
# this comes in the msg received in the channel
# the profile this evidence should be in, should be the profile of the attacker
# because this is evidence that this profile is attacker others right?
profile = ProfileID(ip=saddr)
# Profiles are split into timewindows, each timewindow is 1h,
# this if of the timewindow comes in the msg received in the channel
twid_number = int(
    msg['twid'].replace("timewindow", '')
)
timewindow = TimeWindow(number=twid_number)
# how many flows formed this evidence?
# in the case of scans, it can be way more than 1
conn_count = 1
# list of uids of the flows that are part of this evidence
uid_list = [flow.uid]
# no use the above info to create the evidence obj
evidence = Evidence(
    evidence_type=evidence_type,
    attacker=attacker,
    threat_level=threat_level,
    description=description,
    victim=victim,
    profile=profile,
    timewindow=timewindow,

```

(continues on next page)

(continued from previous page)

```
        uid=uid_list,  
        # when did this evidence happen? use the  
        # flow's ts detected by zeek  
        # this comes in the msg received in the channel  
        timestamp=timestamp,  
        conn_count=conn_count,  
        confidence=confidence  
    )  
self.db.set_evidence(evidence)
```

15.3.4 Testing the Module

The module is now ready to be used. You can copy/paste the complete code that is [here](#)

First we start Slips by using the following command:

```
./slips.py -i wlp3s0 -o local_conn_detector
```

-o is to store the output in the local_conn_detector/ dir.

Then we make a connection to a local ip (change it to a host you know is up in your network)

```
ping 192.168.1.18
```

And you should see your alerts in ./local_conn_detector/alerts.log by using

```
cat local_conn_detector/alerts.log
```

```
Using develop - 9f5f9412a3c941b3146d92c8cb2f1f12aab3699e - 2022-06-02 16:51:43.989778  
2022/06/02-16:51:57: Src IP 192.168.1.18 . Detected a connection to a local_  
↪device 192.168.1.12  
2022/06/02-16:51:57: Src IP 192.168.1.12 . Detected a connection to a local_  
↪device 192.168.1.18
```

15.3.5 Conclusion

Due to the high modularity of slips, adding a new slips module is as easy as modifying a few lines in our template module, and slips handles running your module and integrating it for you.

This is the [list of the modules](#) Slips currently have. You can enhance them, add detections, suggest new ideas using our [Discord](#) or by opening a PR.

For more info about the threat levels, [check the docs](#)

Detailed explanation of [IDEA categories](#) here

Detailed explanation of [Slips profiles and timewindows](#) here

[Contributing guidelines](#)

15.4 Complete Code

Here is the whole `local_connection_detector.py` code for copy/paste.

```
import ipaddress
import json

from slips_files.common.flow_classifier import FlowClassifier
from slips_files.core.structures.evidence import

(
    Evidence,
    ProfileID,
    TimeWindow,
    Victim,
    Attacker,
    ThreatLevel,
    EvidenceType,
    IoCType,
    Direction,
)
from slips_files.common.parsers.config_parser import ConfigParser
from slips_files.common.slips_utils import utils
from slips_files.common.abstracts.imodule import IModule

class LocalConnectionDetector(
    IModule
):
    # Name: short name of the module. Do not use spaces
    name = 'local_connection_detector'
    description = 'detects connections to other devices in your local network'
    authors = ['Template Author']

    def init(
        self
    ):
        # To which channels do you want to subscribe? When a message
        # arrives on the channel the module will receive a msg

        # You can find the full list of channels at
        # slips_files/core/database/redis_db/database.py
        self.c1 = self.db.subscribe(
            'new_flow'
        )
        self.channels = {
            'new_flow': self.c1,
        }
        # to be able to convert flows from dict format to objects
        self.classifier = FlowClassifier()
```

(continues on next page)

(continued from previous page)

```

def pre_main(
    self
):
    """
    Initializations that run only once before the main() function runs in a loop
    """
    utils.drop_root_privs_permanently()

def main(
    self
):
    """Main loop function"""
    if msg := self.get_msg(
        'new_flow'
    ):
        msg = json.loads(
            msg['data']
        )
        # convert the given dict flow to a flow object
        flow = self.classifier.convert_to_flow_obj(
            msg["flow"]
        )
        saddr = flow.saddr
        daddr = flow.daddr
        timestamp = flow.starttime
        srcip_obj = ipaddress.ip_address(
            saddr
        )
        dstip_obj = ipaddress.ip_address(
            daddr
        )
        if srcip_obj.is_private and dstip_obj.is_private:
            # on a scale of 0 to 1, how confident you are of this.
            ↪evidence
            confidence = 0.8
            # how dangerous is this evidence? info, low, medium,
            ↪high, critical?
            threat_level = ThreatLevel.HIGH

            # the name of your evidence, you can put any descriptive
            ↪string here

            # this is the type we just created
            evidence_type = EvidenceType.CONNECTION_TO_LOCAL_DEVICE
            # which ip is the attacker here?
            attacker = Attacker(
                direction=Direction.SRC,
                # who's the attacker the src or the dst?
                attacker_type=Ioctype.IP,
                # is it an IP? is it a domain? etc.
                value=saddr
                # the actual ip/domain/url of the attacker, in

```

(continues on next page)

(continued from previous page)

```

→our case, this is the IP
    )
    victim = Victim(
        direction=Direction.SRC,
        ioc_type=IoCType.IP,
        value=daddr,
    )
    # describe the evidence
    description = f'A connection to a local device {daddr}'
    # the current profile is the source ip,
    # this comes in the msg received in the channel
    # the profile this evidence should be in, should be the
→profile of the attacker
    # because this is evidence that this profile is attacker
→others right?
    profile = ProfileID(
        ip=saddr
    )
    # Profiles are split into timewindows, each timewindow
→is 1h,
    # this if of the timewindow comes in the msg received in
→the channel
    twid_number = int(
        msg['twid'].replace(
            "timewindow",
            ''
        )
    )
    timewindow = TimeWindow(
        number=twid_number
    )
    # list of uids of the flows that are part of this
→evidence
    uid_list = [flow.uid]
    # no use the above info to create the evidence obj
    evidence = Evidence(
        evidence_type=evidence_type,
        attacker=attacker,
        threat_level=threat_level,
        description=description,
        victim=victim,
        profile=profile,
        timewindow=timewindow,
        uid=uid_list,
        # when did this evidence happen? use the
        # flow's ts detected by zeek
        # this comes in the msg received in the channel
        timestamp=timestamp,
        confidence=confidence
    )
    self.db.set_evidence(
        evidence

```

(continues on next page)

(continued from previous page)

```
        )
    self.print(
        "Done setting evidence!!!"
    )
```

All good, you can find your evidence now in `alerts.json` and `alerts.log` of the output directory.

15.5 Line by Line Explanation of the Module

This section is for more detailed explanation of what each line of the module does.

In order to print in your module, you simply use the following line

```
self.print("some text", 1, 0)
```

and the text will be sent to the output process for logging and printing to the terminal.

Now here's the `pre_main()` function, all initializations like dropping root privs, checking for API keys, etc should be done here

```
utils.drop_root_privs_permanently()
```

the above line is responsible for dropping root privileges, so if slips starts with `sudo` and the module doesn't need the root permissions, we drop them.

Now here's the `main()` function, this is the main function of each module, it's the one that gets executed in a loop when the module starts.

All the code in this function is run in a loop as long as the module is up.

in case of an error, the module's main should return non-zero and the module will finish execution and terminate. if there's no errors, the module will keep looping until it runs out of msgs in the redis channels and will call `shutdown_gracefully()` and terminate.

```
if msg := self.get_msg('new_flow'):
```

The above line listens on the channel called `new_flow` that we subscribed to earlier.

The messages received in the channel are flows the slips read by the input process.

15.6 Reading Input flows from an external module (Advanced)

Slips relies on input process for reading flows, either from an interface, a pcap, or zeek files, etc.

If you want to add your own module that reads flows from somewhere else, for example from a simulation framework like the `CYST` module, you can easily do that using the `--input-module <module_name>` parameter

Reading flows should be handled by that module, then sent to the `inputprocess` for processing using the `new_module_flow` channel.

For now, this feature only supports reading flows in zeek json format, but feel free to extend it.

15.6.1 How to shutdown_gracefully()

So, for example if you're training a ML model in your module, and you want to save it before the module stops, You should place your `save_model()` function in the `shutdown_gracefully()` function.

15.6.2 Troubleshooting

- If the module does not start at all, make sure it is not disabled in the `config/slips.yaml` file.
- Check that the `__init__.py` file is present in module directory
- Read the output files (`errors.log` and `slips.log`) - if there were any errors (eg. import errors), they would prevent the module from starting.
- If the module started, but did not receive any messages from the channel, make sure that:
 - The channel is properly subscribed to in the module
 - Messages are being sent in this channel
 - Other modules subscribed to the channel get the message
 - The channel name is present in the `supported_channels` list in `slips_files/core/database/redis_db/database.py`

15.6.3 Testing

Slips has 2 kinds of tests, unit tests and integration tests.

integration tests are in `tests/integration/`, In there we test all files in our `dataset/` dir.

Before pushing, run the unit tests and integration tests by:

- 1- Make sure you're in slips main dir
- 2- Run all tests `./tests/run_all_tests.sh`

Slips supports the `-P` flag to run redis on your port of choice. this flag is used so that slips can keep track of the ports it opened while testing and close them later.

15.6.4 Adding your own unit tests

Slips uses `pytest` as the main testing framework, You can add your own unit tests by:

- 1- create a file called `test_module_name.py` in the `tests/` dir
- 2- create a method for initializing your module in `tests/module_factory.py`
- 3- every function should start with `test_`
- 4- go to the main slips dir and run `./tests/run_all_tests.sh` and every test file in the `tests/` dir will run

15.6.5 Getting in touch

Feel free to join our [Discord server](#) and ask questions, suggest new features or give us feedback.

PRs and Issues are welcomed in our repo.

15.6.6 Conclusion

Adding a new feature to SLIPS is an easy task. The template is ready for everyone to use and there is not much to learn about Slips to be able to write a module.

If you wish to add a new module to the Slips repository, issue a pull request and wait for a review.

DATASETS

Slips comes with some datasets for you to try on the folder `dataset`. They are a mix of real malware, real normal, both malicious and benign, in Argus format, Zeek, pcap, etc.

- 16.1 2017-3-8_win5.pcap**
- 16.2 conn.log**
- 16.3 CTU-Malware-Capture-Botnet-1**
- 16.4 malicious-cc.conn.log**
- 16.5 port-scans**
- 16.6 test10-mixed-zeek-dir**
- 16.7 test11-portscan.binetflow**
- 16.8 test12-icmp-portscan.pcap**
- 16.9 test13-malicious-dhcpscan-zeek-dir**
- 16.10 test14-malicious-zeek-dir**
- 16.11 test15-malicious-zeek-dir**
- 16.12 test16-malicious-zeek-dir**
- 16.13 test1-normal.nfdump**
- 16.14 test2-malicious.binetflow**
- 16.15 test3-mixed.binetflow**
- 16.16 test4-malicious.binetflow**
- 16.17 test5-mixed.binetflow**
- 16.18 test6-malicious.suricata.json**
- 16.19 test7-malicious.pcap**
- 16.20 test8-malicious.pcap**
- 16.21 test9-mixed-zeek-dir**
- 16.22 test-cc**

This is a test for detecting command and control channels. It is a synthetic dataset created by capturing very periodic and semi-periodic connections.

16.22.1 test-cc-capture-1.pcap

Very periodic every 2 seconds.

Capture

- `sudo tcpdump -n -s0 -i eno1 host 147.32.80.37 and host testing.com -v -w test-cc-capture-1.pcap`

Connection

- `while [1]; do curl https://testing.com; sleep 2; done`

16.22.2 test-cc-capture-2.pcap

Semi-periodic, from 2 to 3 second[s]

Capture

- `sudo tcpdump -n -s0 -i eno1 port 53 or (host 147.32.80.37 and host testing.com) -v -w test-cc-capture-2.pcap`

Connection

- `while [1]; do curl https://testing.com; sleep $(echo "scale=2; 2+$RANDOM / 20000" | bc); done`

16.22.3 test18-malicious-ctu-sme-11-win

This capture is a short part of the Dataset [CTU-SME-11](#), capture Experiment-VM-Microsoft-Windows7full-3, day 2023-02-22. It consist of only the first 5000 packets.

Labels

The labels were assigned by an expert by hand. The configuration file is `labels.config` and it was labeled using the tool [netflowlabeler](#).

SLIPS IMMUNE

This is the main guide to the documentation related to the changes done to Slips as part of incorporating the immunology ideas

17.1 Architecture

- Main Architecture of Slips Immune
- Immunology-inspired Model for Slips

17.2 RPI performance

- Research RPI Limitations
- Slips Compatibility In The RPI
- Installing Slips On the RPI
- Performance Evaluation
- Testing
- LLM Research and Selection
- LLM RPI Performance
- Stress Testing

17.3 Updating Slips

- Updating Slips

17.4 Security & Network Configuration

- ARP Poisoning
- ARP Poisoning Risks
- Blocking with Slips as an Access Point
- IDS-in-the-middle Traffic routing
- RPI Failover Mechanisms
- Security Audit

17.5 Datasets & LLM Training

Report Documents:

- [Summarization Dataset Report](#) - Event summarization and behavior analysis
- [Risk Analysis Dataset Report](#) - Root cause and risk assessment

Workflow Guides:

- [Summarization Workflow Implementation](#) - Step-by-step guide for generating summarization datasets
- [Risk Analysis Workflow Implementation](#) - Step-by-step guide for generating risk datasets
- [Alert DAG Parser Documentation](#) - DAG structural analysis reference

Datasets Evaluation (LLM-as-a-judge):

- [LLM Evaluation Guide](#) - How to evaluate and compare LLM models
- [LLM-as-Judge Rubric](#) - Scoring criteria for cause analysis and risk assessment evaluation (two-stage rubric, max 60 pts)
- [Summarization Dataset Evaluation Results](#) - Performance metrics for summarization models.
- [Risk Analysis Dataset Evaluation Results](#) - Performance metrics for risk assessment models

LLM finetuning

- [LLM RPI Finetuning Frameworks](#) - Framework selection rationale (Unsloth vs alternatives)
- [Fine-Tuning Approach](#) - General pipeline: LoRA, GGUF export, hardware setup
- [Fine-Tuning Evaluation Methodology](#) - LLM-as-judge pipeline, metrics, and breakdown dimensions
- [Quantization and Deployment](#) - GGUF conversion, Ollama publication, and quantization performance analysis

Incident Summarization task:

- [Summarization Training Procedure](#) - Dataset filtering, ground truth selection, system prompt
- [Summarization Fine-Tuned Model: Evaluation Results](#) - Benchmark results for the finetuned Qwen2.5-1.5B

Risk Assessment & Cause Analysis task:

- [Risk Assessment Training Procedure](#) - Dataset filtering, best-of-N selection, combined adapter training
- [Risk Fine-Tuned Model: Evaluation Results](#) - Benchmark results for the finetuned Qwen2.5-1.5B risk model

Unified model (Summary + Cause + Risk):

- [Unified Training Procedure](#) - Single adapter for all three tasks: dataset merging, lora_r=128 + RSLoRA, version history
- [Unified Fine-Tuned Model: Evaluation Results](#) - Benchmark results vs standalone models and quantization impact

SLIPS IN ACTION

To demonstrate the capabilities of Slips, we will give it real life malware traffic and checking how Slips analyses it.

18.1 Saefko RAT

We provide the analysis of the network traffic of the RAT06-Saefko [download here](#) using Slips.

The capture contains different actions done by the RAT controller (e.g. upload a file, get GPS location, monitor files, etc.). For detailed analysis, check Kamila Babayeva's blog [Dissecting a RAT. Analysis of the Saefko RAT](#).

Disclaimer: The used Slips version in this demo is 1.0.2, alerts and evidence generated in this demo may be different than the alerts you may see using the latest version of Slips.

From the analysis we know that:

- The controller IP address: 192.168.131.1 and 2001:718:2:903:f410:3340:d02b:b918
- The victim's IP address: 192.168.131.2 and 2001:718:2:903:b877:48ae:9531:fbfc

First we run slips using the following command:

```
./slips.py -e 1 -f RAT06_Saefko.pcap
```

First, Slips will start by updating all the remote TI feeds added in slips.yaml

To make sure Slips is up to date with the most recent IoCs in all feeds, all feeds are loaded, parsed and updated periodically and automatically by Slips every 24 hours by our [update_manager](#), which requires no user interaction.

After updating, slips modules start and print the PID of every successfully started module.

Then, we see the alert

Alerts are printed by the evidence module, Slips detected IP 2001:718:2:903:b877:48ae:9531:fbfc as infected due to the above evidence See the difference between alerts and evidence [here](#)

Slips splits does detections in timewindows, each time window is 1 hour long by default and contains dozens of features computed for all connections that start in that time window. So if an IP behaves maliciously at 4 PM, it will be marked as infected only during that hour, the next hour if no malicious behaviour occurs, slips will treat the traffic as normal. This explains the start and stop timestamps in the alert `start 2021-04-10T16:44:43.285478+02:00, stop 2021-04-10T17:44:43.285478+02:00`. This is the period (timewindow) in which this IP was behaving maliciously.

The difference between infected and normal timewindows is shown better in the web interface.

Start Slips with `-w` to inspect the results in the web interface.

We can see that IP 2001:718:2:903:b877:48ae:9531:fbfc is infected only in timewindow1 as it's marked in red and is behaving normally in timewindow0 as it's colored in green.

We can see all the flows done by this IP in the infected timewindow in the web interface by opening timewindow1.

In the interface evidence view we can inspect what slips detected, this is the same evidence printed in Figure 3.

We can see the following detections in the evidence:

```
Detected Malicious JA3: 807fca46d9d0cf63adf4e5e80e414bbe from source address
2001:718:2:903:b877:48ae:9531:fbfc AS: CESNET z.s.p.o. description: Tofsee ['malicious']
```

JA3 fingerprint the client part of the SSL certificate. This indicates that the source IP 2001:718:2:903:b877:48ae:9531:fbfc was infected with one of the Tofsee malware family

Slips also detected the connection to the database:

```
SSL certificate validation failed with (certificate has expired) Destination IP:
↪ 2a02:4780:dead:d8f::1. SNI: experimentsas.000webhostapp.com
```

From the RAT analysis, we know that 000webhostapp.com is the web hosting service used by the C&C server.

Slips also detected

```
Connection to unknown destination port 6669/TCP destination IP 2001:67c:2564:a191::fff:1.
↪ (['open.ircnet.net'])
Connection to unknown destination port 8000/TCP destination IP 192.168.131.1.
```

From the APK list of IRC servers shown in the RAT analysis, we know that the phone connects on port 6669/TCP and 8000/TCP to different IRC servers to receive the malicious commands. The rDNS of the server is also printed in the alert open.ircnet.net

Our machine learning module rnn-cc-detection detected the C&C server using recurrent neural network

```
Detected C&C channel, destination IP: 192.168.131.1 port: 8000/tcp score: 0.9871
```

Slips also detected

```
Possible DGA or domain scanning. 192.168.131.2 failed to resolve 15 domains
```

The above detections are evidence that when accumulated, resulted in an alert.

To view all evidence that slips detected including those that weren't enough to generate an alert, you can

```
cat output/alerts.log
```

Slips also has another log file in JSON format so they can be easily parsed and exported. See [the exporting section](#) of the documentation.

The generated alerts in this file follow [CESNET's IDEA0 format](#).

```
cat output/alerts.json
```

18.2 Emotet

We will be analysing several Emotet PCAPs starting from infection, until Trickbot and Qakbot malwares are dropped.

The captures contain different actions done by the Emotet and trickbot controller. For detailed analysis, check Paloalto's blog [Examining Emotet Infection Traffic](#).

18.2.1 Emotet infection

We will be analysing this Emotet PCAP [download here](#). password: infected

When running Slips on the PCAP

```
./slips.py -f Example-1-2021-01-06-Emotet-infection.pcap
```

We get the following alerts

The reconnection attempts shown in the analysis

are detected by Slips in the following evidence

```
Detected a connection without DNS resolution to IP: 46.101.230.194
Detected Multiple reconnection attempts to Destination IP: 46.101.230.194 from
↪IP: 10.1.6.206
```

18.2.2 Trickbot

Analyzing the next PCAP [download here](#) that contains the Trickbot traffic. password: infected

Running slips on the pcap

```
./slips.py -f Example-4-2021-01-05-Emotet-infection-with-Trickbot.pcap
```

Slips detects a self signed SSL certificate to 102.164.208.44 which is the trickbot IP associated with data exfiltration

```
Detected SSL certificate validation failed with (self signed certificate)
↪Destination IP: 102.164.208.44
```

Slips also detected

```
Detected a connection without DNS resolution to IP: 102.164.208.44.
```

and

```
Detected Connection to unknown destination port 449/TCP destination IP 102.164.208.44.
```

18.2.3 Qakbot

Analyzing the next PCAP [download here](#) that contains the Qakbot traffic. password: infected

Running slips on the pcap

```
./slips.py -f Example-5-2020-08-18-Emotet-infection-with-Qakbot.pcap
```

Slips detected that the victim 192.168.100.101 is infected with Qakbo using JA3

```
Detected Malicious JA3: 7dd50e112cd23734a310b90f6f44a7cd from source address 192.168.100.
↪101 description: Quakbot ['malicious']
Detected Malicious JA3: 57f3642b4e37e28f5cbe3020c9331b4c from source address 192.168.100.
↪101 description: Gozi ['malicious']
```

We can also see a Domain generation algorithm detection by the same victim 192.168.100.101

```
Detected possible DGA or domain scanning. 192.168.100.101 failed to resolve 40 domains
```

And an expired certificate to [samaritantec.com](#). This domain was reported as hosting an Emotet binary on the same date.

```
Detected SSL certificate validation failed with (certificate has expired) Destination
↪IP: 43.255.154.32. SNI: samaritantec.com
```

Slips then detected

```
Detected C&C channel, destination IP: 71.80.66.107 port: 443/tcp score: 0.9601
ected a connection without DNS resolution to IP: 71.80.66.107. AS: CHARTER-
↪20115, rDNS: 071-080-066-107.res.spectrum.com
```

a quick search in [virustotal](#) shows that this IP 71.80.66.107 is associated with [qakbot](#)

and a port scan

```
Detected horizontal port scan to port 443/TCP. From 192.168.100.101 to 6 unique dst IPs.
↪Tot pkts: 21. Threat Level: medium
```

18.3 DroidJack v4.4 RAT

Running Slips on DroidJack v4.4 RAT [download here](#). password: `infected`.

The capture contains different actions done by the RAT controller(e.g. upload a file, get GPS location, monitor files, etc.). For detailed analysis, check Kamila Babayeva's blog [Analysis of DroidJack v4.4 RAT network traffic](#).

From the analysis we know that:

- The controller IP address: 147.32.83.253
- The victim's IP address: 10.8.0.57

When running slips on the PCAP

```
./slips.py -f RAT02.pcap
```

We get the following alerts

Slips detected the connection to the C&C server using an unknown port

```
Detected Connection to unknown destination port 1334/TCP destination IP 147.32.83.253.
```

Slips also detected the reconnection attempts made from the victim to the C&C server

```
Multiple reconnection attempts to Destination IP: 147.32.83.253 from IP: 10.8.0.57
```

Slips also detects connections without resolutions due to their wide usages among malware to either check internet connectivity or get commands from the C&C servers.

```
Detected a connection without DNS resolution to IP: 147.32.83.253. AS: CESNET z.
↪s.p.o., rDNS: dhcp-83-253.felk.cvut.cz
```

The identification (AS, SNI, rDNS) of each IP, if available, is printed in every evidence generated by Slips.

Our Threat intelligence feed [Abuse.ch](#) detected a malicious JA3 indicating that the victim 10.8.0.57 is infected

```
Detected Malicious JA3: 7a29c223fb122ec64d10f0a159e07996 from source address 10.8.0.57  
↔description: ['malicious']
```

And our machine learning models detected the C&C server

```
Detected C&C channel, destination IP: 147.32.83.253 port: 1334/tcp score: 0.9755
```

Slips creates a profile per each IP that appeared in the traffic. Each profile contains flows sent from this IP. Each flow is described with a specific letter which description can be found [here](#).

Considering that, Slips detects the C&C channel over 1334/TCP. Slips' machine learning module called LSTM detecting C&C channel is shown below

Slips did not detect periodic connection over 1337/UDP because the LSTM module focuses on TCP. But from the behavioral model of the connections over 1337/UDP shown below, we can conclude that the model is periodic and most of connections are of a small size.

19.1 Slips starting too slow in docker

Make sure you're not running many containers at the same time because they share kernel resources even though they're isolated.

19.2 Getting “Illegal instruction” error when running slips

If the tensorflow version you're using isn't compatible with your architecture, you will get the “Illegal instruction” error and slips will terminate.

To fix this you can disable the modules that use tensorflow by adding `rnn-cc-detection`, `flow_ml_detection` to the `disable` key in `config/slips.yaml`

19.3 Docker time is not in sync with that of the host

You can add your local `/etc/localtime` as volume in Slips Docker container by using:

```
docker run -it --rm --net=host --cap-add=NET_ADMIN -v /etc/localtime:/etc/localtime:ro --  
->name slips stratosphereips/slips:latest
```

19.4 Redis WARNING Memory overcommit must be enabled! in docker

This is a redis known issue, you can find the fix here [https://redis.io/docs/latest/operate/oss_and_stack/management/admin/#:~:text=Redis setup tips-,Linux,-Deploy Redis using](https://redis.io/docs/latest/operate/oss_and_stack/management/admin/#:~:text=Redis%20setup%20tips-,Linux,-Deploy%20Redis%20using)

CONTRIBUTING

All contributions are welcomed, thank you for taking the time to contribute to this project! These are a set of guidelines for contributing to the development of Slips [1].

20.1 How can you contribute?

- Run Slips and report bugs and needed features, and suggest ideas
- Pull requests with a solved GitHub issue and new feature
- Pull request with a new detection module.

20.2 Persistent Git Branches

The following git branches in the Slips repository are permanent:

- **master**: contains the stable version of Slips, with new versions at least once a month.
- **develop**: contains the latest unstable version of Slips and also its latest features. All new features should be based on this branch.

20.3 Naming Git branches for Pull Requests

To keep the Git history clean and facilitate the revision of contributions we ask all branches to follow concise namings. These are the branch-naming patterns to follow when contributing to Slips:

- **author-bugfix-**: pull request branch, contains one bugfix,
- **author-docs-**: pull request branch, contains documentation work,
- **author-enhance-**: pull request branch, contains one enhancement (not a new feature, but improvement nonetheless)
- **author-feature-**: pull request branch, contains a new feature,
- **author-refactor-**: pull request branch, contains code refactoring,

20.4 What branch should you base your contribution to Slips?

As a general rule, base your contributions to the `develop` branch.

20.5 Creating a pull request

Commits:

- Commits should follow the KISS principle: do one thing, and do it well (keep it simple, stupid).
- Commit messages should be easily readable, imperative style (“Fix memory leak in...”, not “FixES mem...”)

Pull Requests:

- If you have developed multiple features and/or bugfixes, create separate branches for each one of them, and request merges for each branch;
- Each PR to develop will trigger the develop Github checks, these checks will run Slips unit tests and integration tests locally in a ubuntu VM and in docker to make sure the branch is ready to merge.
- PRs won't be merged unless the checks pass.
- The cleaner you code/change/changeset is, the faster it will be merged.

20.6 Beginner tips on how to create a PR in Slips

Here's a very simple beginner-level steps on how to create your PR in Slips

1. Fork the Slips repo
2. Clone the forked repo
3. In your clone, checkout origin/develop: `git checkout origin develop`
4. Install slips pre-commit hooks `pre-commit install`
5. Generate a baseline for detecting secrets before they're committed `detect-secrets scan --exclude-files ". *dataset/* | (?x)(^config/local_ti_files/own_malicious_JA3.csv$|.*test.*|slips.yaml|.*\.md$)" > .secrets.baseline`
6. Create your own branch off develop using your name and the feature name: `git checkout -b <yourname>_<feature_name> develop`
7. Change the code, add the feature or fix the bug, etc. then commit with a descriptive msg `git commit -m "descriptive msg here"`
8. Test your code: this is a very important step. you shouldn't open a PR with code that is not working: `./tests/run_all_tests.sh`
9. If some tests didn't pass, it means you need to fix something in your branch.
10. Push to your own repo: `git push -u origin <yourname>_<feature_name>`
11. Open a PR in Slips, remember to set the base branch as develop.
12. List your changes in the PR description

20.7 How to test the auto update functionality

To test Slips auto update using the testing channel, follow these steps:

1. Create a new temporary branch from the branch you want to test, for example `origin/yourname-autoupdate-test`.
2. Edit [update.json] in your new branch and update the testing channel entry:
 - increase the version value so this branch appears newer than the version currently available

- keep "backwards_compatible": true
 - keep "has_new_dependencies": false
 - set the testing branch entry to your new branch name e.g origin/yourname-autoupdate-test
3. Commit the update.json change in that temporary branch and push the branch to your remote.
 4. Check out your original old branch again. Before continuing, make sure there are no merge conflicts and no modified Slips files in this branch.
 5. Create a copy of `config/slips.yaml`.
 6. In your copied config file, enable auto update and configure the testing channel:

```
update:
  auto_update_slips: true
  channel_to_update_slips_from: testing
  testing_branch_to_update_slips_from: origin/<yourname-autoupdate-test>
```

7. Run Slips on an interface and point it to your copied config file using `-c`, for example:

```
./slips.py -i <interface_name> -c config/<your_test_config>.yaml
```

8. While Slips is running, it should detect the higher version in the testing channel and auto update to the branch you pushed in step 3.

20.8 Rejected PRs

We will not review PRs that have the following:

- Code that's not tested. a screenshot of the passed tests is required for each PR.
- PRs without steps to reproduce your proposed changes.
- Asking for a step by step guide on how to solve the problem. It is ok to ask us clarifications after putting some effort into reading the code and the docs. but asking how exactly should i do X shows that you didn't look at the code

Some IDEs like [PyCharm](#) and [vscode](#) have the option to open a PR from within the IDE.

That's it, now you have a ready-to-merge PR!

[1] These contributions guidelines are inspired by the project [Snoopy](#)

20.9 FAQ

20.9.1 How does slips work?

- `slips.py` is the entry point, it's responsible for starting all modules, and keeping slips up until the analysis is finished.
- `slips.py` starts the input process, which is the one responsible for reading the flows from the files given to slips using `-f` it detects the type of file, reads it and passes the flows to the profiler process. if slips was given a PCAP or is running on an interface, the input process starts a zeek thread that analyzes the pcap/interface using slips' own zeek configuration and sends the generated zeek flows to the profiler process.

- `slips.py` also starts the update manager, it updates slips local TI files, like the ones stored in `slips_files/organizations_info` and `slips_files/ports_info`. later, when slips is starting all the modules, slips also starts the update manager but to update remote TI files in the background in this case.
- Once the profiler process receives the flows read by the input process, it starts to convert them to a structure that slips can deal with. it creates profiles and time windows for each IP it encounters.
- Profiler process gives each flow to the appropriate module to deal with it. for example flows from `http.log` will be sent to `http_analyzer.py` to analyze them.
- Profiler process stores the flows, profiles, etc. in slips databases for later processing. the info stored in the dbs will be used by all modules later. Slips has 2 databases, Redis and SQLite. it uses the sqlite db to store all the flows read and labeled. and uses redis for all other operations. the sqlite db is created in the output directory, meanwhile the redis database is in-memory. 7-8. using the flows stored in the db in step 6 and with the help of the timeline module, slips puts the given flows in a human-readable form which is then used by the web UI.
- when a module finds a detection, it sends the detection to the evidence process to deal with it (step 10) but first, this evidence is checked by the whitelist to see if it's whitelisted in our `config/whitelist.conf` or not. if the evidence is whitelisted, it will be discarded and won't go through the next steps
- now that we're sure that the evidence isn't whitelisted, the evidence process logs it to slips log files and gives the evidence to all modules responsible for exporting evidence. so, if CEST, Exporting modules, or CYST is enabled, the evidence process notifies them through redis channels that it found an evidence and it's time to share the evidence.
- if the blocking module is enabled using `-p`, the evidence process shares all detected alerts to the blocking module. and the blocking module handles the blocking of the attacker IP through the linux firewall (supported in linux only)
- if `p2p` is enabled in `config/slips.yaml`, the `p2p` module shares the IP of the attacker, its' score and blocking requests sent by the evidence process with other peers in the network so they can block the attackers before they reach them.
- The output process is slips custom logging framework. all alerts, warnings and info printed are sent here first for proper formatting and printing.

20.9.2 What is the recommended development environment?

For minimum hassle when developing it's recommended to use ubuntu, install slips natively, and use your favorite IDE

20.9.3 While developing, my module is not working and there's no errors shown to the CLI or printed to `errors.log`

Always make sure to run slips with `-e 1`, for example `./slips.py -e 1 -f <some_pcap> -o some_output_dir`

The goal of suppressing errors by default is the most errors should be handled by the developers and modules should recover and continue working normally afterwards (if possible), so no need to show minor errors to users by default.

20.9.4 What are all these Databases? Redis cache db, redis main database, SQLite, and Database manager?

- We use SQLite for storing all the flows and alflows, so if you want to store or retrieve something you will most probably find the function you need already implemented there (in `slips_files/core/database/sqlite_db/database.py`)
- Any other info goes in Redis.
- The DB manager is a Facade which acts as a proxy to both the sqlite and the redis databases. The goal of this is to add an abstraction layer between the developers and the dbs. To avoid the confusion of "i need to do X, is it in redis or sqlite?"

- The point above means that for each function you add to Redis or SQLite, you need to add a wrapper for it in the `database_manager.py` to be accessible to all modules.

20.9.5 How does Redis communication work?

- If you run `slips` without any special arguments, Slips starts redis cache db (redis server port 6379 db 1) and Redis main db (redis port 6379 db 1)
- You can start Slips with `-m`, which starts redis on a random available redis port in the range (32768 to 10000), or `-P` if you want to start redis on a specific port.
- Slips starts the redis server if it's not started by default.
- Slips uses its own `redis.conf`, it doesn't use the default one. you can find it in `config/redis.conf.template`.
- The cache db is shared among all running `slips` instances, and is persistent, meaning it is not deleted on each run unlike the main redis db (redis port 6379 db 1), which is overwritten every run.
- If you're gonna add a new redis channel to `slips`, remember to add it to the list of `supported_channels` in `slips_files/core/database/redis_db/database.py`

20.9.6 How are the modules loaded?

- All modules in the `modules/` directory that implement the `IModule` interface are automatically imported by `slips`, for more technical details check the `load_modules()` function in `managers/process_manager.py`
- Module names must use snake_case and follow the `x_y_doer` style used by names such as `http_analyzer`.
- Set the module class name to the same snake_case identifier as the module directory and main file.
- Keep the module directory name, the main module file, the config section name, and any module references in docs aligned with that same snake_case name.

20.9.7 There's some missing code in all modules, what's happening?

- All modules implement the `IModule` interface in `slips_files/common/abstracts/module.py`, it ensures that all modules behave the same, for example they all shutdown the same, they all keep track of the redis channels they're using, they all have a common `init()`, they all forward msgs to the printer in the same way, etc.
- Any logic that will be duplicated across all modules should be in this interface

20.9.8 How does slips stop?

It all begins when `input.py` realizes there's no more flows arriving from the `zeek files/suricata/nfdump` file it's reading.

It's a good idea to read the code before checking this graph

Evidence Handler is the only process that stops but keeps waiting in memory for new msgs to arrive until all other modules are done. because if any of the modules added an evidence, `EvidenceHandler` should be up to report and handle it or else it will be discarded. Once all modules are done processing, `EvidenceHandler` is killed by the `Process manager`.

20.9.9 How does the tests work?

- Running the tests locally should be done using `./tests/run_all_tests.sh`
- It runs the unit tests first, then the integration tests.
- Please get familiar with `pytest` first <https://docs.pytest.org/en/stable/how-to/output.html>

20.9.10 What does generate_performance_plots do?

- Debug.generate_performance_plots in *config/slips.yaml* is a developer-only debugging switch for performance investigations.
- When it is true, Slips writes extra CSVs under *output/performance_plots/csv/*, including alert latency (*latency.csv*), profiler worker latency (*profiler_worker_*_latency.csv*), and throughput (*flows_per_minute.csv*).
- On shutdown, the process manager turns those CSVs into plots and metrics under *output/performance_plots/* and *output/metrics.txt*.
- Leave it false for normal development and production-style runs. Enabling it adds Redis bookkeeping, file writes, and plot-generation work that are only useful when diagnosing throughput or latency behavior.
- The plots shown in *docs/immune/stress_testing.md* were generated with this parameter enabled.

20.9.11 Where and how do we get the GW info?

Using one of these 3 ways

20.10 Global P2P - Fides contribution notes

Variables used in the trust evaluation and its accompanied processes, such as database-backup in persistent SQLite storage and memory persistent Redis database of Slips, are strings, integers and floats grouped into custom dataclasses. Aforementioned data classes can be found in *modules/fides/model*. The reader may find that all of the floating variables are in the interval $<-1; 1>$ and some of them are between $<0; 1>$, please refer to the *modules/fides/model* directory.

The Fides Module is designed to cooperate with a global-peer-to-peer module. The communication is done using Slips' Redis channel, for more information please refer to communication and messages sections above.

An example of a message answering Fides-Module's opinion request follows.

```
import redis

# connect to redis database 0
redis_client = redis.StrictRedis(host='localhost', port=6379, db=0)

message = '''
{
  "type": "nl2tl_intelligence_response",
  "version": 1,
  "data": [
    {
      "sender": {
        "id": "peer1",
        "organisations": ["org_123", "org_456"],
        "ip": "192.168.1.1"
      },
      "payload": {
        "intelligence": {
          "target": {"type": "server", "value": "192.168.1.10"},
          "confidentiality": {"level": 0.8},
          "score": 0.5,
          "confidence": 0.95
        }
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

        "target": "stratosphere.org"
    }
},
{
    "sender": {
        "id": "peer2",
        "organisations": ["org_789"],
        "ip": "192.168.1.2"
    },
    "payload": {
        "intelligence": {
            "target": {"type": "workstation", "value": "192.168.1.20"},
            "confidentiality": {"level": 0.7},
            "score": -0.85,
            "confidence": 0.92
        },
        "target": "stratosphere.org"
    }
}
]
}
'''

# publish the message to the "network2fides" channel
channel = "network2fides"
redis_client.publish(channel, message)

print(f"Message published to channel '{channel}'.")

```

For more information about message handling, please also refer to `modules/fides/messaging/message_handler.py` and to `modules/fides/messaging/dacite/core.py` for message parsing.

20.10.1 Communication

The module uses Slips' Redis to receive and send messages related to trust intelligence, evaluation of trust in peers and alert message dispatch.

Used Channels `modules/fides/messaging/message_handler.py`

Slips Channel Name	Purpose
<code>slips2fides</code>	Provides communication channel from Slips to Fides
<code>fides2slips</code>	Enables the Fides Module to answer requests from <code>slips2fides</code>
<code>network2fides</code>	Facilitates communication from network (P2P) module to the Fides Module
<code>fides2network</code>	Lets the Fides Module request network opinions form network modules

For more details, the code [here](#) may be read.

20.10.2 Messages

Message type (data['type'])	Channel	Call/Handle	Description
alert	slips2	FidesModule as self.__alerts.dispatch_alert(target=data['target'], confidence=data['confidence'], score=data['score'])	Triggers sending an alert to the network, about given target, which SLips believes to be compromised.
intelligence	slips2	FidesModule as self.__intelligence.request_data(target=data['target'])	Triggers request of trust intelligence on given target.
t12n1_alert	fides2:	call dispatch_alert() of AlertProtocol class instance	Broadcasts alert through the network about the target.
t12n1_intell:	fides2:	NetworkBridge.send_intelligence_response(...)	Shares Intelligence with peer that requested it.
t12n1_intell:	fides2:	NetworkBridge.send_intelligence_request(...)	Requests network intelligence from the network regarding this target.
t12n1_recomm:	fides2:	NetworkBridge.send_recommendation_response(...)	Responds to given request_id to recipient with recommendation on target.
t12n1_recomm:	fides2:	NetworkBridge.send_recommendation_request(...)	Request recommendation from recipients on given peer.
t12n1_peers_reliability:	fides2:	NetworkBridge.send_peers_reliability(...)	Sends peer reliability, this message is only for network layer and is not dispatched to the network.

Implementations of Fides_Module-network-communication can be found in `modules/fides/messaging/network_bridge.py`.

CODE DOCUMENTATION

Slips auto-generated code documentation

RELATED REPOSITORIES

- [Slips-tools](#): repo is to store all the tools and scripts needed to test and evaluate Slips