

Suricata IDPS and Linux kernel

É. Leblond, G. Longo

Stamus Networks

February 10, 2016

1

Suricata

- Introduction
- Streaming
- Performance

2

Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3

Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4

Conclusion

1

Suricata

- Introduction
- Streaming
- Performance

2

Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3

Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4

Conclusion

What is Suricata

- IDS and IPS engine
- Get it here:
<http://www.suricata-ids.org>
- Open Source (GPLv2)
- Initially publicly funded now funded by consortium members
- Run by Open Information Security Foundation (OISF)
- More information about OISF at
<http://www.openinfosecfoundation.org/>



Suricata Features

- High performance, scalable through multi threading
- Advanced Protocol handling
 - Protocol recognition
 - Protocol analysis: field extraction, filtering keywords
 - Transaction logging in extensible JSON format
- File identification, extraction, on the fly MD5 calculation
 - HTTP
 - SMTP
- TLS handshake analysis, detect/prevent things like Diginotar
- Lua scripting for detection
- Hardware acceleration support:
 - Endace
 - Napatech,
 - CUDA
 - PF_RING

Suricata capture modes

IDS

- pcap: multi OS capture
- af_packet: Linux high performance on vanilla kernel
- netmap: FreeBSD high performance
- NFLOG: Netfilter logging

IPS

- NFQUEUE: Using Netfilter on Linux
- ipfw: Use divert socket on FreeBSD
- af_packet: Level 2 software bridge

Offline analysis

- Pcap: Analyse pcap files
- Unix socket: Use Suricata for fast batch processing of pcap files

1

Suricata

- Introduction
- **Streaming**
- Performance

2

Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3

Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4

Conclusion

Fooling detection

- Get your activity unnoticed
- Complete your attack and stay in place

Principle

- Signature-based IDS rely on packet content
- Modification of traffic could be used to avoid detection
- Without changing the impact of the attack

Play on interpretation issue

OS-based evasion

- All OS do not react the same
 - RFC are incomplete. Improvisations have been made.
 - Variation of traffic for a same flow is possible
- Overlapping Fragments

Application-based evasion

- Different servers can treat the same request differently.
- No web server are treating a twice used argument the same way.

Personnality

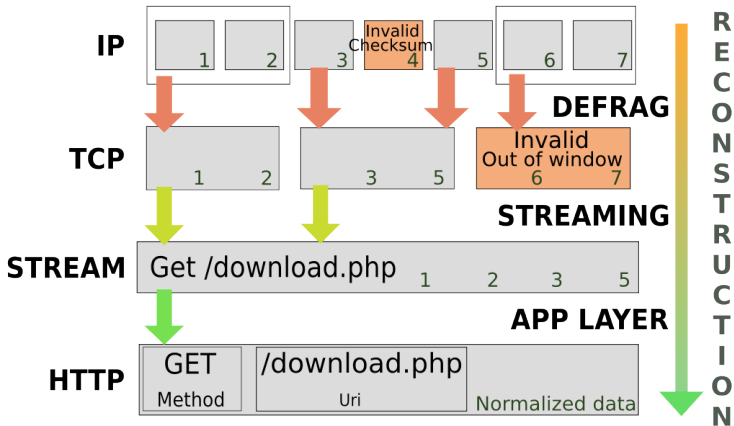
Personnality

- IDS implements personality
- It is possible to associate network and OS type
- For Suricata, HTTP servers can be personified too.

Suricata configuration

```
host-os-policy :  
  # Make the default policy windows.  
  windows: [0.0.0.0/0]  
  bsd: []  
  bsd-right: []  
  old-linux: []  
  linux: [10.0.0.0/8]
```

Suricata reconstruction and normalization



1

Suricata

- Introduction
- Streaming
- **Performance**

2

Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3

Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4

Conclusion

A typical signature example

Signature example: Chat facebook

```
alert http $HOME_NET any -> $EXTERNAL_NET any \
(
msg:"ET CHAT Facebook Chat (send message)"; \
flow:established,to_server; content:"POST"; http_method; \
content:"/ajax/chat/send.php"; http_uri; content:"facebook.com"; http_host; \
content:"netdev"; http_client_body;
reference:url,www.emergingthreats.net/cgi-bin/cvsweb.cgi/sigs/POLICY/POLICY_Facebook_Chat; \
sid:2010784; rev:4; \
)
```

This signature tests:

- The HTTP method: *POST*
- The page: */ajax/chat/send.php*
- The domain: *facebook.com*
- The body content: *netdev*

No passthrough

All signatures are inspected

- Different from a firewall
- More than 15000 signatures in standard rulesets

Optimization on detection engine

- Tree pre filtering approach to limit the set of signatures to test
- Multi pattern matching on some buffers

```

 1 [|||||||||||||100.0%]  5 [|||||]          26.5%]    9 [|||||||||||||88.7%]  13 [|||||]          35.3%]
 2 [|||||||||||||100.0%]  6 [|||||]          18.8%]   10 [|||||||||||||62.9%]  14 [|||||]          22.4%]
 3 [|||||]           7 [|||||]          27.2%]   11 [|||||||||||||58.2%]  15 [|||||]          14.9%]
 4 [|||||]           8 [|||||]          33.8%]   12 [|||||]          38.6%]   16 [|||||]          23.8%]
Mem[|||||||||||||25891/64403MB]
Swp[|||||]          50/33377MB
Tasks: 43, 31 thr; 11 running
Load average: 7.40 7.24 7.32
Uptime: 82 days, 23:13:26

```

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
13679	root	20	0	34.6G	22.4G	5211M	S	710.	35.7	69h08:07	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13741	root	18	-2	34.6G	22.4G	5211M	R	102.	35.7	6h32:22	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13749	root	18	-2	34.6G	22.4G	5211M	R	90.9	35.7	4h02:18	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13758	root	18	-2	34.6G	22.4G	5211M	R	70.7	35.7	10h00:47	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13750	root	18	-2	34.6G	22.4G	5211M	R	63.3	35.7	3h40:08	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13751	root	18	-2	34.6G	22.4G	5211M	S	57.9	35.7	3h20:58	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13744	root	18	-2	34.6G	22.4G	5211M	R	35.7	35.7	3h22:16	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13753	root	18	-2	34.6G	22.4G	5211M	S	33.7	35.7	3h14:37	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13748	root	18	-2	34.6G	22.4G	5211M	S	33.0	35.7	3h11:43	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13742	root	18	-2	34.6G	22.4G	5211M	R	31.7	35.7	3h37:38	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13752	root	18	-2	34.6G	22.4G	5211M	S	29.6	35.7	3h44:17	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13756	root	18	-2	34.6G	22.4G	5211M	R	25.6	35.7	3h33:02	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13747	root	18	-2	34.6G	22.4G	5211M	S	25.6	35.7	3h10:13	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13745	root	18	-2	34.6G	22.4G	5211M	S	24.9	35.7	3h18:05	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13754	root	18	-2	34.6G	22.4G	5211M	R	22.2	35.7	3h15:22	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13746	root	18	-2	34.6G	22.4G	5211M	R	20.9	35.7	3h19:41	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13755	root	18	-2	34.6G	22.4G	5211M	S	18.9	35.7	3h21:57	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13743	root	18	-2	34.6G	22.4G	5211M	R	18.9	35.7	3h12:16	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13760	root	22	2	34.6G	22.4G	5211M	S	2.7	35.7	31:04.81	/usr/local/bin/suricata -c /etc/suricata/regit-ya
13761	root	22	2	34.6G	22.4G	5211M	S	2.7	35.7	27:38.31	/usr/local/bin/suricata -c /etc/suricata/regit-ya

```

Samples: 691K of event 'cycles', Event count (approx.): 256764876818
Overhead Shared Object          Symbol
 64.14%  suricata                      [.] SCACSearch
  3.20%  suricata                      [.] BoyerMoore
  1.16%  suricata                      [.] SigMatchSignatures
  0.90%  libc-2.19.so                 [.] memset
  0.87%  [kernel]                  [k] ixgbe_clean_rx_irq
  0.75%  suricata                      [.] IPOnlyMatchPacket
  0.68%  libpthread-2.19.so         [.] pthread_mutex_unlock
  0.64%  [kernel]                  [k] __netif_receive_skb_core
  0.62%  libpthread-2.19.so         [.] pthread_mutex_lock
  0.62%  suricata                      [.] AFPReadFromRing
  0.61%  [kernel]                  [k] irq_entries_start
  0.58%  [kernel]                  [k] tpacket_rcv
  0.55%  libc-2.19.so                 [.] __memcmp_sse4_1
  0.52%  [kernel]                  [k] memcpy
  0.42%  [kernel]                  [k] ixgbe_poll
  0.42%  [kernel]                  [k] menu_select
  0.40%  suricata                      [.] StreamTcpPacket
  0.36%  [kernel]                  [k] native_write_msr_safe
  0.35%  [kernel]                  [k] packet_lookup_frame.isra.56

```


- Bandwidth per core is limited
 - From 150Mb/s
 - To 500Mb/s
- Scaling
 - Using RSS
 - Splitting load on workers

- 1 Suricata
 - Introduction
 - Streaming
 - Performance
- 2 Suricata and Linux kernel
 - AF_PACKET
 - NFQUEUE

- 3 Suricata and offloading
 - Interest of offloading
 - Implementation of framework
 - Use it with NFQ
 - Other Methods

- 4 Conclusion

1 Suricata

- Introduction
- Streaming
- Performance

2 Suricata and Linux kernel

- **AF_PACKET**
- NFQUEUE

3 Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4 Conclusion

Linux raw socket

- Raw packet capture method
- Socket based or mmap based


Linux raw socket

- Raw packet capture method
- Socket based or mmap based

Fanout mode

- Load balancing over multiple sockets
- Multiple load balancing functions
 - Flow based
 - CPU based
 - RSS based

Suricata workers mode



The rollover option

Concept

- Ring buffer can fill in burst or single flow
- Capture would gain of splitting single intensive flow
- Rollover mode switch to next socket when ring is full


The rollover option

Concept

- Ring buffer can fill in burst or single flow
- Capture would gain of splitting single intensive flow
- Rollover mode switch to next socket when ring is full

Problem with Suricata

- Suricata reconstruct the stream
- Rollover mode causes reordering of stream
- Massive accuracy loss



1 Suricata

- Introduction
- Streaming
- Performance

2 Suricata and Linux kernel

- AF_PACKET
- **NFQUEUE**

3 Suricata and offloading


- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4 Conclusion

- It is used in Suricata to work in IPS mode, performing action like DROP or ACCEPT on the packets, permitting us to delegate the verdict on the packets.
- With NFQUEUE we are able to delegate the verdict on the packet to a userspace software.
- The following rules will ask a userspace software connected to queue 0 for a decision.
 - **nft add filter forward queue num 0**
 - **iptables -A FORWARD -j NFQUEUE --queue-num 0**

The following steps explains how NFQUEUE works with Suricata in IPS mode:

- Incoming packet matched by a rule is sent to Suricata through nfnetlink
- Suricata receives the packet and issues a verdict depending on our ruleset
- The packet is either transmitted or rejected by kernel



- NFQUEUE number of packets per second on a single queue is limited due to the nature of nfnetlink communication.
- Batching verdict can help but without an efficient improvement.
- Starting Suricata with multiple queue could improve it:

```
suricata -c /etc/suricata/suricata.yaml -q 0 -q 1
```

- 1 Suricata
 - Introduction
 - Streaming
 - Performance
- 2 Suricata and Linux kernel
 - AF_PACKET
 - NFQUEUE
- 3 Suricata and offloading
 - Interest of offloading
 - Implementation of framework
 - Use it with NFQ
 - Other Methods
- 4 Conclusion

1 Suricata

- Introduction
- Streaming
- Performance

2 Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3 Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4 Conclusion

Stream depth

Attacks characteristic

- In most cases attack is done at start of TCP session
- Generation of requests prior to attack is not common
- Multiple requests are often not even possible on same TCP session

Stream reassembly depth

- Suricata reassemble TCP sessions till `stream.reassembly.depth` bytes.
- Stream is not analyzed once limit is reached

Introducing offloading

Principle

- No need to get packet from kernel after stream depth is reached
- If there is
 - no file store
 - or other operation

Usage

Set `stream.offloading` option to `yes` in suricata config file to offload

1 Suricata

- Introduction
- Streaming
- Performance

2 Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3 Suricata and offloading

- Interest of offloading
- **Implementation of framework**
- Use it with NFQ
- Other Methods

4 Conclusion

Implementation

Suricata update

- Add callback function
- Capture method register itself and provide a callback
- Suricata calls callback when it wants to offload

Implementation

Suricata update

- Add callback function
- Capture method register itself and provide a callback
- Suricata calls callback when it wants to offload

Coded for NFQ

- Update capture register function
- Written callback function
 - Set a mark with respect to a mask on packet
 - Mark is set on packet when issuing the verdict

- 1 Suricata
 - Introduction
 - Streaming
 - Performance
- 2 Suricata and Linux kernel
 - AF_PACKET
 - NFQUEUE
- 3 Suricata and offloading
 - Interest of offloading
 - Implementation of framework
 - **Use it with NFQ**
 - Other Methods
- 4 Conclusion

nftables ruleset

```
table ip filter {
    chain forward {
        type filter hook forward priority 0;
        # usual ruleset
    }

    chain ips {
        type filter hook forward priority 10;
        meta mark set ct mark
        mark 0x00000001 accept
        queue num 0
    }

    chain connmark_save {
        type filter hook forward priority 20;
        ct mark set mark
    }
}
```

Results of iperf3 tests

Local testing

```
eric@ice-age2:~/git/oisf (dev-offloading-v7)$ iperf3 -c 127.0.0.1 -p 6666
Connecting to host 127.0.0.1, port 6666
[ 4] local 127.0.0.1 port 38532 connected to 127.0.0.1 port 6666
[ ID] Interval      Transfer    Bandwidth  Retr  Cwnd
[ 4]  0.00-1.00    sec   86.3 MBytes  724 Mbits/sec    0   1.03 MBytes
[ 4]  1.00-2.00    sec   84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4]  2.00-3.00    sec   84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4]  3.00-4.00    sec   84.8 MBytes  712 Mbits/sec    0   1.03 MBytes
[ 4]  4.00-5.00    sec   84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4]  5.00-6.00    sec   84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4]  6.00-7.00    sec   84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4]  7.00-8.00    sec   84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4]  8.00-9.00    sec   84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4]  9.00-10.00   sec   84.2 MBytes  706 Mbits/sec    0   1.03 MBytes
-----
[ ID] Interval      Transfer    Bandwidth  Retr
[ 4]  0.00-10.00   sec   844 MBytes  708 Mbits/sec    0           sender
[ 4]  0.00-10.00   sec   842 MBytes  707 Mbits/sec                   receiver
```

Results of iperf3 tests

Local testing

```
eric@ice-age2:~/git/oisf (dev-offloading-v7)$ iperf3 -c 127.0.0.1 -p 6666
Connecting to host 127.0.0.1, port 6666
[ 4] local 127.0.0.1 port 38532 connected to 127.0.0.1 port 6666
[ ID] Interval           Transfer     Bandwidth   Retr  Cwnd
[ 4] 0.00-1.00 sec      86.3 MBytes  724 Mbits/sec    0   1.03 MBytes
[ 4] 1.00-2.00 sec      84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4] 2.00-3.00 sec      84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4] 3.00-4.00 sec      84.8 MBytes  712 Mbits/sec    0   1.03 MBytes
[ 4] 4.00-5.00 sec      84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4] 5.00-6.00 sec      84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4] 6.00-7.00 sec      84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4] 7.00-8.00 sec      84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4] 8.00-9.00 sec      84.0 MBytes  705 Mbits/sec    0   1.03 MBytes
[ 4] 9.00-10.00 sec     84.2 MBytes  706 Mbits/sec    0   1.03 MBytes
-----
[ ID] Interval           Transfer     Bandwidth   Retr
[ 4] 0.00-10.00 sec     844 MBytes  708 Mbits/sec    0
[ 4] 0.00-10.00 sec     842 MBytes  707 Mbits/sec    0
                               sender
                               receiver
```

<marketing>Local testing with offload is 90 times faster</marketing>

```
eric@ice-age2:~/git/oisf (dev-offloading-v7)$ iperf3 -c 127.0.0.1 -p 6666
Connecting to host 127.0.0.1, port 6666
[ 4] local 127.0.0.1 port 38624 connected to 127.0.0.1 port 6666
[ ID] Interval           Transfer     Bandwidth   Retr  Cwnd
[ 4] 0.00-1.00 sec      7.42 GBytes  63.7 Gbits/sec    0   1.09 MBytes
[ 4] 1.00-2.00 sec      7.04 GBytes  60.5 Gbits/sec    0   1.09 MBytes
[ 4] 2.00-3.00 sec      7.06 GBytes  60.7 Gbits/sec    0   1.09 MBytes
[ 4] 3.00-4.00 sec      7.30 GBytes  62.7 Gbits/sec    0   1.09 MBytes
[ 4] 4.00-5.00 sec      8.35 GBytes  71.7 Gbits/sec    0   1.26 MBytes
[ 4] 5.00-6.00 sec      8.45 GBytes  72.6 Gbits/sec    0   1.26 MBytes
[ 4] 6.00-7.00 sec      7.22 GBytes  62.0 Gbits/sec    0   1.26 MBytes
[ 4] 7.00-8.00 sec      6.99 GBytes  60.0 Gbits/sec    0   1.26 MBytes
[ 4] 8.00-9.00 sec      6.98 GBytes  60.0 Gbits/sec    0   1.26 MBytes
[ 4] 9.00-10.00 sec     7.05 GBytes  60.5 Gbits/sec    0   1.26 MBytes
-----
[ ID] Interval           Transfer     Bandwidth   Retr
[ 4] 0.00-10.00 sec     73.9 GBytes  63.4 Gbits/sec    0
[ 4] 0.00-10.00 sec     73.9 GBytes  63.4 Gbits/sec    0
                               sender
                               receiver
```

Selective offloading

Ignore some traffic

- Ignore intensive traffic like Netflix
- Can be done independently of stream depth
- Can be done using generic or custom signatures

Selective offloading

Ignore some traffic

- Ignore intensive traffic like Netflix
- Can be done independently of stream depth
- Can be done using generic or custom signatures

The offload keyword

- A new `offload` signature keyword
- Trigger offloading when signature match
- Example of signature

```
alert http any any -> any any (content:"netdevconf.org"; \\
    http_host; offload; sid:6666; rev:1;)
```

1 Suricata

- Introduction
- Streaming
- Performance

2 Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3 Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- **Other Methods**

4 Conclusion

Implementation for other captures

Possibilities

- AF_PACKET
- Signaling Openvswitch
- Custom HW
- ...

Constraint

- Method needs to be fast
- It needs to handle
 - Huge amount of flow/items
 - Rapid change rate

1 Suricata

- Introduction
- Streaming
- Performance

2 Suricata and Linux kernel

- AF_PACKET
- NFQUEUE

3 Suricata and offloading

- Interest of offloading
- Implementation of framework
- Use it with NFQ
- Other Methods

4 Conclusion

Suricata and Linux

- A deep imbrication
- IDS constraint causes some generic features to fail
- Offloading looks promising

More information

- **Suricata:** <http://www.suricata-ids.org/>
- **Netfilter:** <http://www.netfilter.org/>
- **Stamus Networks:** <https://www.stamus-networks.com/>



Contact us

- **Éric Leblond:**
eleblond@stamus-networks.com
- **Giuseppe Longo:**
glongo@stamus-networks.com
- **Twitter:** @regiteric and @theglongo
- `https://www.stamus-networks.com/`