Lab-1001: Flows/IPFIX

---

*Evidence: /home/ndfir/labs/1001/upload-http-8000-**

Takeaways: Students will analyze network traffic produced by YAF (i.e., NetFlow)

Objective:

Analyze flow data produced by **yaf** and identify potential data exfiltration activity.

1. Navigate to the lab directory at /home/ndfir/labs/1001/ and find the various artifacts produced by yaf; there should be four files as shown in the screen capture below.

```
ndfir@ndfir-box:~/labs/1001$ ls -ll
total 20
-rwxrwxr-x 1 ndfir ndfir 1856 Sep 19 23:25 upload-http-8000-bidirection.yaf
-rwxrwxr-x 1 ndfir ndfir 2703 Sep 19 23:25 upload-http-8000-bidirection.yaf.txt
-rwxrwxr-x 1 ndfir ndfir 2132 Sep 19 23:25 upload-http-8000-uniflow.yaf
-rwxrwxr-x 1 ndfir ndfir 5124 Sep 19 23:25 upload-http-8000-uniflow.yaf.txt
```

YAF Output from 1.2GB PCAP

2. The following commands were used against a very large pcap file (1.2 GB) to create bidirectional uniflow Yaf output:

```
yaf --in upload-http-8000.pcap --out upload-http-8000-bidirection.yaf
yaf --in upload-http-8000.pcap --out upload-http-8000-uniflow.yaf --uniflow
```

   a. The yaf output can be imported by several tools for analysis, as detailed here: https://tools.netsa.cert.org/yaf/index.html
   b. Conversely, we can use the **yafscii** utility to convert the yaf output to human readable text.

3. The following commands were used to create human readable text logs from the yaf output:

```
yafscii --in upload-http-8000-bidirection.yaf --print-header --tabular
yafscii --in upload-http-8000-uniflow.yaf --print-header --tabular
```

   a. The above commands convert the yaf output to human readable text as well as adding a header row for field titles and in “tabular” format.

```
start-time            |end-time              |duration|rtt     |proto|sip
                |dp   |iflags  |uflags  |riflags |ruflags |isn     |risn    |tag|rtag|pkt    |oct     |rp
2021-09-19 19:02:50.945|2021-09-19 19:02:51.123|   0.178|   0.082|  6|                                 192.168.232.
    173.230.154.59| 8000|       S|     APF|      AS|     APF|b4096408|3827461d|000|000|       5|     710|
2021-09-19 19:02:57.190|2021-09-19 19:07:02.760| 245.570|   0.080|  6|                                 192.168.232.
    173.230.154.59| 8000|       A|    APSF|      AS|     APF|0000025c|5af6626b|000|000|  735553|1103278714|
2021-09-19 19:07:24.253|2021-09-19 19:07:24.621|   0.368|   0.079|  6|                                 192.168.232.
    173.230.154.59| 8000|       S|     APF|      AS|     APF|14ffa756|487c51a8|000|000|       5|     684|
2021-09-19 19:07:38.555|2021-09-19 19:08:05.836|  27.281|   0.079|  6|                                 192.168.232.
    173.230.154.59| 8000|       S|     APF|      AS|     APF|52f4b672|505d774e|000|000|   71983|107962321|
2021-09-19 19:08:22.606|2021-09-19 19:08:27.770|   5.164|   0.081|  6|                                 192.168.232.
    173.230.154.59| 8000|       S|      AF|      AS|     APF|cff31622|7fff32e3|000|000|       4|     180|
2021-09-19 19:08:27.910|2021-09-19 19:08:33.072|   5.162|   0.084|  6|                                 192.168.232.
```

Yafscii Output in Human Readable Text

b. The output is human readable, but formatting is still an issue. To address this, we use the 'tr' utility installed on our VM to remove all 'space characters' (0x20).

```
cat upload-http-8000-bidirection.yaf.txt |tr -d ' ' |less -S
```

```
start-time|end-time|duration|rtt|proto|sip|sp|dip|dp|iflags|uflags|riflags|ruflags|isn|risn|tag|rtag|pkt|oct|rpkt|r
2021-09-1919:02:50.945|2021-09-1919:02:51.123|0.178|0.082|6|192.168.232.130|49702|173.230.154.59|8000|S|APF|AS|APF|
2021-09-1919:02:57.190|2021-09-1919:07:02.760|245.570|0.080|6|192.168.232.130|49704|173.230.154.59|8000|A|APSF|AS|A
2021-09-1919:07:24.253|2021-09-1919:07:24.621|0.368|0.079|6|192.168.232.130|49724|173.230.154.59|8000|S|APF|AS|APF|
2021-09-1919:07:38.555|2021-09-1919:08:05.836|27.281|0.079|6|192.168.232.130|49726|173.230.154.59|8000|S|APF|AS|APF
2021-09-1919:08:22.606|2021-09-1919:08:27.770|5.164|0.081|6|192.168.232.130|49732|173.230.154.59|8000|S|AF|AS|APF|c
2021-09-1919:08:27.910|2021-09-1919:08:33.072|5.162|0.084|6|192.168.232.130|49736|173.230.154.59|8000|S|AF|AS|APF|f
2021-09-1919:08:27.909|2021-09-1919:08:33.075|5.166|0.084|6|192.168.232.130|49734|173.230.154.59|8000|S|APF|AS|APF|
```

Yafscii Output with 'spaces' removed

4. Although all fields are important at some point, for this exercise we'll want to focus on the following fields for Bidirectional flow traffic:

   start-time|end-time|duration|rtt|proto|sip|sp|dip|dp|pkt|oct|rpkt|roct

   start-time: Start time of the flow
   end-time: End time of the flow
   duration: Flow duration in fractional seconds. Only present if the flow has a non-zero duration
   rtt: Round-trip time estimate in milliseconds in decimal format
   proto: IP protocol identifier in decimal format
   sip: Source IPv4 address in dotted-quad format or IPv6 address in RFC 2373 format
   sp: Source transport port in decimal format
   dip: Destination IPv4 address in dotted-quad format or IPv6 address in RFC 2373 format
   dp: Destination transport port in decimal format
   pkt: Forward first-packet 802.1q VLAN tag in hexadecimal format
   oct: Forward octet count in decimal format (number of bytes)
   rpkt: Reverse first-packet 802.1q VLAN tag in hexadecimal format
   roct: Reverse octet count in decimal format (number of bytes)

Bidirectional:

```
cat upload-http-8000-bidirection.yaf.txt |tr -d ' ' |cut -f
   1,2,3,4,5,6,7,8,9,18,19,20,21 -d '|' |less -S
```

```
start-time|end-time|duration|rtt|proto|sip|sp|dip|dp|pkt|oct|rpkt|roct
2021-09-1919:02:50.945|2021-09-1919:02:51.123|0.178|0.082|6|192.168.232.130|49702|173.230.154.59|8000|5|710|5|664
2021-09-1919:02:57.190|2021-09-1919:07:02.760|245.570|0.080|6|192.168.232.130|49704|173.230.154.59|8000|735553|1103278714|772619|30905315
2021-09-1919:07:24.253|2021-09-1919:07:24.621|0.368|0.079|6|192.168.232.130|49724|173.230.154.59|8000|5|684|5|696
2021-09-1919:07:38.555|2021-09-1919:08:05.836|27.281|0.079|6|192.168.232.130|49726|173.230.154.59|8000|71983|107962321|78590|3144157
2021-09-1919:08:22.606|2021-09-1919:08:27.770|5.164|0.081|6|192.168.232.130|49732|173.230.154.59|8000|4|180|3|124
2021-09-1919:08:27.910|2021-09-1919:08:33.072|5.162|0.084|6|192.168.232.130|49736|173.230.154.59|8000|4|180|3|124
2021-09-1919:08:27.909|2021-09-1919:08:33.075|5.166|0.084|6|192.168.232.130|49734|173.230.154.59|8000|5|684|5|732
2021-09-1919:08:46.505|2021-09-1919:08:47.227|0.722|0.082|6|192.168.232.130|49740|173.230.154.59|8000|728|1080973|811|32995
2021-09-1919:08:51.312|2021-09-1919:08:51.752|0.440|0.081|6|192.168.232.130|49742|173.230.154.59|8000|6|724|5|764
```

Yafscii Output of Bidirectional Flows (spaces removed)

5. For uniflow output, there is no “reverse” data as each directional flow is recorded on a separate line; thus, the fields rtt, rpkt, and roct are not applicable. The following fields can be used on uniflow logs.

   start-time|end-time|duration|proto|sip|sp|dip|dp|pkt|oct

Uniflow:

```
cat upload-http-8000-uniflow.yaf.txt |tr -d ' ' |cut -f 1,2,3,5,6,7,8,9,18,19 -d '|'
   |less -S
```

```
start-time|end-time|duration|proto|sip|sp|dip|dp|pkt|oct
2021-09-1919:02:50.945|2021-09-1919:02:51.123|0.178|6|192.168.232.130|49702|173.230.154.59|8000|5|710
2021-09-1919:02:51.027|2021-09-1919:02:51.123|0.096|6|173.230.154.59|8000|192.168.232.130|49702|5|664
2021-09-1919:02:57.190|2021-09-1919:07:02.760|245.570|6|192.168.232.130|49704|173.230.154.59|8000|735553|1103278714
2021-09-1919:02:57.270|2021-09-1919:07:02.760|245.490|6|173.230.154.59|8000|192.168.232.130|49704|772619|30905315
2021-09-1919:07:24.253|2021-09-1919:07:24.621|0.368|6|192.168.232.130|49724|173.230.154.59|8000|5|684
2021-09-1919:07:24.332|2021-09-1919:07:24.621|0.289|6|173.230.154.59|8000|192.168.232.130|49724|5|696
2021-09-1919:07:38.555|2021-09-1919:08:05.836|27.281|6|192.168.232.130|49726|173.230.154.59|8000|71983|107962321
2021-09-1919:07:38.634|2021-09-1919:08:05.836|27.202|6|173.230.154.59|8000|192.168.232.130|49726|78590|3144157
2021-09-1919:08:22.606|2021-09-1919:08:27.770|5.164|6|192.168.232.130|49732|173.230.154.59|8000|4|180
2021-09-1919:08:22.687|2021-09-1919:08:27.770|5.083|6|173.230.154.59|8000|192.168.232.130|49732|3|124
2021-09-1919:08:27.910|2021-09-1919:08:33.072|5.162|6|192.168.232.130|49736|173.230.154.59|8000|4|180
2021-09-1919:08:27.994|2021-09-1919:08:33.072|5.078|6|173.230.154.59|8000|192.168.232.130|49736|3|124
2021-09-1919:08:27.909|2021-09-1919:08:33.075|5.166|6|192.168.232.130|49734|173.230.154.59|8000|5|684
2021-09-1919:08:27.993|2021-09-1919:08:33.075|5.082|6|173.230.154.59|8000|192.168.232.130|49734|5|732
2021-09-1919:08:46.505|2021-09-1919:08:47.227|0.722|6|192.168.232.130|49740|173.230.154.59|8000|728|1080973
2021-09-1919:08:46.587|2021-09-1919:08:47.227|0.640|6|173.230.154.59|8000|192.168.232.130|49740|811|32995
2021-09-1919:08:51.312|2021-09-1919:08:51.752|0.440|6|192.168.232.130|49742|173.230.154.59|8000|6|724
2021-09-1919:08:51.393|2021-09-1919:08:51.752|0.359|6|173.230.154.59|8000|192.168.232.130|49742|5|764
```

Yafscii Output of Uniflow Flows (spaces removed)

6. Reviewing each output, there are three data transfers that are exponentially larger than all other flows.

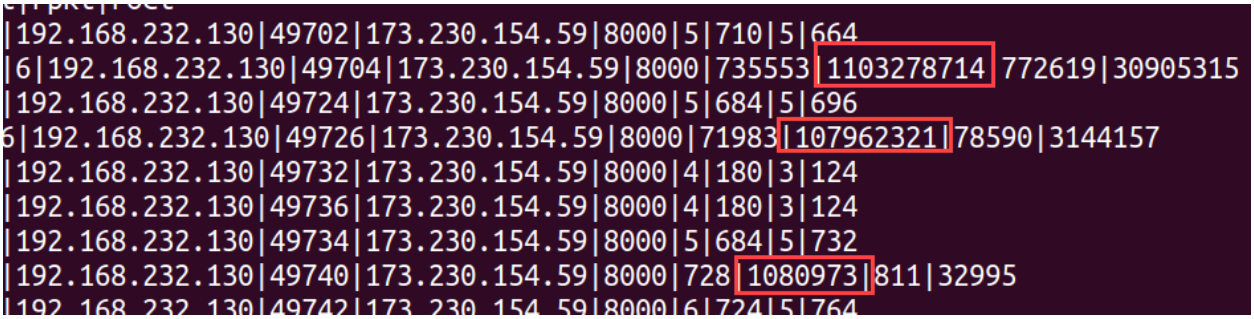

Yafscii Output of Bidirectional Flows – Large Data Transfers

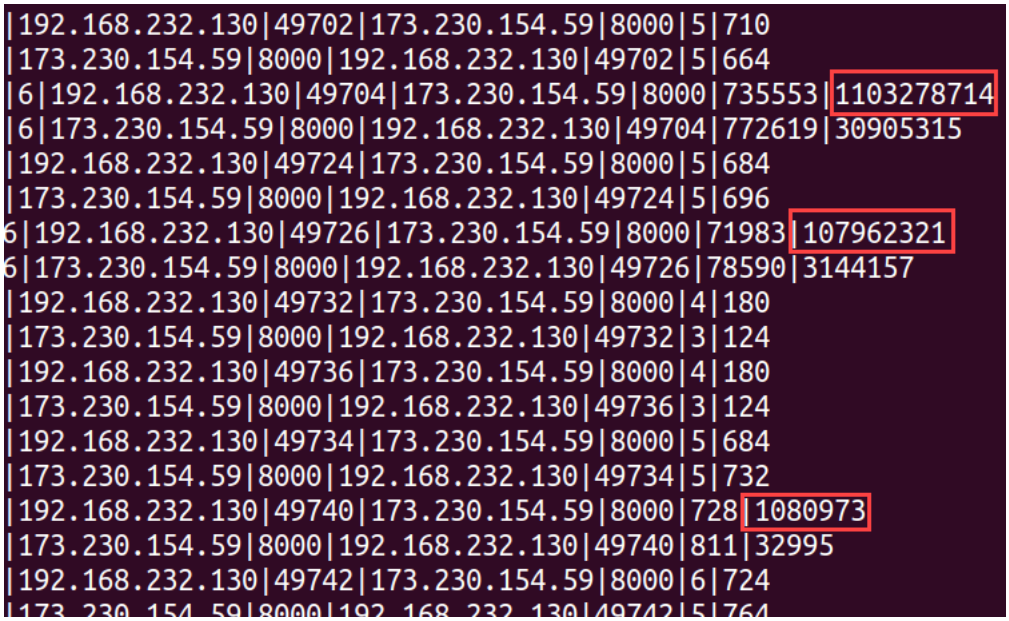

Yafscii Output of Uniflow Flows – Large Data Transfers