

Table of Contents

Collection and Analysis of Statistics by Protocols and Directions	3
--	----------

Collection and Analysis of Statistics by Protocols and Directions



Product Description

1. After setting up flow transmission, not all information is sent. What could be the problem?

The Netflow v5 protocol does not guarantee delivery because it operates over UDP; hence, there is no retransmission of packets in case of network losses or losses at the collector.

Ensure the following:

1. There are no network losses between SSG and the collector. For example, check if the traffic from the control channel to the collector passes through shaping or if there are interface limitations below the Netflow output speed of SSG.
2. Ensure that the collector can accept data at the SSG output speed. Use the `netflow_rate_limit` parameter to limit the speed; for diagnostic purposes, you can set the Netflow SSG output speed to minimum values. If there are no issues at minimal values, the losses are on the collector level.

Losses on the collector can be checked with the command

```
grep "Sequence Errors" /var/log/messages | grep -v "Sequence Errors: 0"
```

Non-zero values indicate the presence of losses.

To eliminate losses:

1. Set the `netflow_rate_limit` parameter according to the information flow and collector capabilities. If set too low, losses will occur for another reason — not all information will be sent in time.
2. [Tuning the network stack](#)
3. Install nfsen on a more powerful computer, avoid virtualization.
4. Switch to the TCP version of the IPFIX (Netflow) protocol.

In the `statistics` log `var/log/dpi/fastdpi_stat.log`, information on Netflow data transmission is output, which can help diagnose issues.

```
[STAT      ][2019/02/01-17:21:28:938274] Statistics on NFLW_Full :  
{0/0/1668468}  
NFLW_Full_IPv4{3948181/939339852}{3111140/3415836963}{7760/13036/6640}  
The first 3 digits - {0/0/1668468} : { connect errors/flow released/nothing  
to send - packet counters unchanged }  
NFLW_Full_IPv4{3948181/939339852}{3111140/3415836963}{7760/13036/6640} :  
{3948181/939339852} : packets/bytes for direction = 0 ( ip_src < ip_dst )  
{3111140/3415836963} : packets/bytes for direction = 1  
{7760/13036/6640} : not sent by full netflow/ipfix - number of  
flows/packets direction==0/packets direction==1
```

For IPv6, it is similar but called `NFLW_Full_IPv6`.

[2. No losses with `netflow_timeout=1`](#)

This indicates that losses occur on the collector; setting the parameter to 1 smooths out Netflow output peaks. Losses without smoothing are likely due to the collector's receive buffer overflow.

Details: What the `netflow_timeout` parameter does.

Start transmission at moment t_1 , determine the time for the next transmission t_2 . If necessary, send statistical changes:

1. by ports
2. by AS
3. by billing
4. by sessions. Session changes are sent considering active and passive timeouts.

Then check: if the current time t_n is greater than t_2 , start a new transmission cycle immediately. Otherwise, sleep for t_2-t_n .

Then presumably the following occurs:

Losses can only be determined on the collector through the sequence value in the header. If there are no losses with `netflow_timeout==1`, the volume of sent data has decreased.

Fewer sessions change in 1 second than in 10, so the collector cannot handle it.

Suppose all packets from SSG reach the collector, which can only process, for example, 10 MB. As a result, the socket's receive buffer will fill up, and incoming packets will simply be discarded.

Attention: If setting this parameter to this value, check for the absence of errors in the alert log during peak hours.

Alternatively, set `netflow_timeout` to 10 and the transmission rate to `netflow_rate_limit=10` for testing.

[3. How to export data in a format suitable for loading into Excel?](#)

The simplest option is to cut the required data by column width:

```
nfdump -R /usr/local/nfsen/profiles-data/live/petrosvaz/2015/07/20 -s
dpipr/bytes -n 50 |grep "(" |awk -v FIELDWIDTHS='40 40 28 16' -v OFS=';';
'{print $2,$4 }'|tr -d '[:blank:]'
```

The result is loaded into Excel.

Similarly for autonomous systems:

```
nfdump -R /usr/local/nfsen/profiles-data/live/petrosvaz_as/2015/07/20 -s
asn/bytes -n 50 |grep "(" |awk -v FIELDWIDTHS='38 65 28 16' -v OFS=';';
'{print $2,$4 }'|tr -d '[:blank:]'
```

TOP 50 protocols:

```
nfdump -R /usr/local/nfsen/profiles-data/live/protocols/2015/07/20 -s
dpipr/bytes -n 50 |grep "(" |awk -v FIELDWIDTHS='40 40 28 16' -v OFS=';'
```

```
'{print $2,$4 }'|tr -d '[:blank:]' > top_proto.csv
```

TOP 50 autonomous systems:

```
ndfdump -R /usr/local/nfsen/profiles-data/live/directions/2015/07/20 -s  
asn/bytes -n 50 |grep "(" |awk -v FIELDWIDTHS='38 65 28 16' -v OFS='';'  
'{print $2,$4 }'|tr -d '[:blank:]' > top_asn.csv
```

Attention:

When using the summation option to get TOP results -s dpipr/bytes, the -o format does not work: -o fmt:"%ts %td %pr %sap → %dap %flg %tos %pkt %byt %fl"