Содержание

3 FAQ

- How to check the curent release (CCC)?

```
fastdpi -re
```

- How to check the curen version?

```
fastdpi -ve
```

- How to roll back to a previous verion?

```
rollback example, from 2.7 to 2.6 version:
yum downgrade fastdpi-2.6
```

- There is an error in the log: "error loading DSCP settings, res=-4"
  Error is displayed because of the lack of autonomous system dscp. It can be ignored.

- What are the lock counters stored in the statistics log?

- Commands aren't always processed and the following error is displayed, "ERROR : Can't connect to 127.0.0.1:29000, errcode=99 : Cannot assign requested address Autodetected fastdpi params : dev='lo', port=29000 connecting 127.0.0.1:29000 ..." We suspect that our way of subscribers loading to the VAS Experts DPI is not quite good enough. (we load each subscriber resulting in at least 50000 commands while initializing, that takes place in addition once a day just to be on the safe side)

- Migration scripts from SCE SM to VAS Experts DPI DB, the description is inside the archive

- How to check CPU core usage between available CPU cores

```
In order to view the CPU usage between available CPU cores, please launch
the "top" utility from terminal and press "1"
To get the CPU usage of a fastdpi's processes launch the following command:
ps -p `pidof fastdpi` H -o %cpu,lwp,pri,psr,comm
Example output:
%CPU   LWP PRI PSR COMMAND
0.0  23141  41   0 fastdpi_main
0.0  23146  41   0 fastdpi_dl
0.3  23147  41   0 fastdpi_ctrl
35.8 23148  41   0 fastdpi_ajb
32.7 23152  41   1 fastdpi_rx_1
34.1 23165  41   2 fastdpi_wrk0
34.1 23170  41   3 fastdpi_wrk1
The dpi tasks are allocated to PSR cores in order to avoid interference with
each other:
 wrk thread is responsible for the network packet data analysis
 rx thread is responsible for the data transfer over the network ports
 the rest threads handle application and auxiliary tasks
 (netflow generation, control commands reception, lists loading, pcap
recording etc.)
```

```
  and may cause peak CPU loads, so they are set to run at specific core.
```

- Got an error in fastdpi_alert.log, what should we do?:
  [CRITICAL][2017/10/06-16:36:44:616019][0x7fdb297ac700] metadata_storage : Can't allocate memory [repeat 1], cntr=188889, allocated=188889

```
 This error means, that fragmented packet cann't be reassembled. Usually it
indicates a DDoS attack.
 You can just ignore it. The "CRITICAL" level of the error is inaccurate,
it'll be changed to "WARNING" in future versions.
```

- Got an error in fastdpi_alert.log, what should we do?
  [CRITICAL][2017/10/06-16:36:44:616019][0x7fdb297ac700] metadata_storage : Can't allocate memory [repeat 1], cntr=188889, allocated=188889

```
DPI preallocates resourses before it gets started according to the given
subscribers number by default.
Эit is adjusted by the "mem_ip_metadata_recs" configuration option.
So to increase the number of subscribers to 500000 "mem_ip_metadata_recs"
opton should be edited in /etc/dpi/fastdpi.conf:
mem_ip_metadata_recs=500000
DPI restart is needed:
service fastdpi restart
```

- What files are suggested to be archived?

```
cp /etc/pf_ring/ /BACKUPDIR/pf_ring
cp /etc/dpi /BACKUPDIR/etc/
mdb_copy /var/db/dpi /BACKUPDIR/db/
```

- ipmi consume 100% CPU, so it disturb dpi operation

```
echo 100 > /sys/module/ipmi_si/parameters/kipmid_max_busy_us
To make your changes persistent (to avoid this setting reset after the next
server restart), you need to add the above command to the /etc/rc.local
file.
```