

Author: Carlo Augusto Grazia

Title: FAQ on Flent

This document is a short guide of FAQ created to share my experience in the use of the Flent tool, in order to avoid other users to text 200 times the Flent creator, Toke Høiland-Jørgensen, many times like me :)

```
#####  
##### ORGANIZE THE TESTS #####  
#####
```

How can I run multiple runs?

You can create your own script (as I've done in the beginning) and configure each run as you want, but it is DEFINITELY better to get familiar with the batch scripting supported by flent because it simplifies a lot the entire process from the run to the plots/results.

```
$ flent -B flent-file.batch -b ALL --batch-title 'logical name'
```

The example.batch available on the flent website is complete enough to start your own tests.

N.B. to specify multiple hosts in rtt_var* tests in the batch file, a comma separation between hosts is needed. For multiple test_parameters in batch, instead, the separation between each parameter must be a semicolon.

How can I get my own testbed?

You can play with Flent by pointing it at some remote machines like bufferbloat servers, but if you want to play with the bottleneck and extract some secondary data with Flent you need access to the "middle" nodes, and you will surely need your own testbed with full access on all the machines. Linux Distro suggested: Arch and Debian

How to collect secondary data like CPU consumption and Qdisc data?

Those info are all collected by auxiliary datasets. You'll need to add the 'qdisc_stats_hosts', 'qdisc_stats_interfaces' and 'cpu_stats_hosts' test parameters to the Flent invocation. The hosts are anything SSH understands (so IP address, hostname, or a config in ~/.ssh/config), and the test client needs to have passwordless login access to the host in question. When monitoring multiple interfaces on the same host, repeat the hostname (you'll need the same number of hostnames and interfaces names).

```
$ flent <your params> --test-parameter qdisc_stats_hosts=myrouter,myrouter \  
--test-parameter qdisc_stats_interfaces=eth0,eth1 --test-parameter \  
cpu_stats_hosts=myrouter
```

or in a batch file

```
test_parameters = qdisc_stats_hosts=myrouter,myrouter; \  
qdisc_stats_interfaces=eth0,eth1; cpu_stats_hosts=myrouter
```

Flent will look for an executable called tc_iterate on the target host to get the tc stats. You can compile that from tc_iterate.c in the 'misc' directory with the Flent source code. Flent will fall back to a shell script-based monitoring loop, but the C program gives a more accurate sampling.

Why my delay plot does not show anything?

It'll only get delay if you're running Codel, FQ-CoDel or PIE as those are the only qdiscs that report queueing delay

Can I point to different servers and have more test simultaneously?

Sure. You can use the rtt_var* test suit, you don't need to have different RTT path, rtt_var* tests are good to point to different hosts

Change QDisc between runs

Fleint does not have a built-in command to change the local (or remote) qdisc as well as it does with the TCP Congestion Control, it must be changed (if needed) before each run, as done in the example batch-file

Am I sure that the Congestion Control I've selected is running?

Just check always in the metadata the congestion control of the TCP flows in the test field "SERIES_META"

I've a lot of spikes in the plots

Turn off the offloads on the the Qdisc with

```
$ ethtool -K INTERFACENAME tso off gso off gro off
```

```
#####  
##### DEAL WITH THE DATA #####  
#####
```

How can I have the Jain's Fairness Index plot?

The JFi plot is automatically extracted by the run if you are using a rtt_var* test

How can I see the variance on box/bar plots?

Fleint aggregates by itself if you name the test properly like NAME-1, NAME-2 etc

The dash "-" separation is very important with Fleint and the final "-NUMBER" gives the iteration.

The individual runs (signified by the numbers at the end) should then be summed to get you variances in the bar plots as well.

Is important to have exactly the same name for each run, this means that also date and time must be the same. It is better to get the different run from a batch configuration in order to have the same date and time in the title.

How can I change the name of a line plot or a histogram bar?

Changing names is done by liberal application of `--filter-regexp` and `--replace-legend`.

How can I group plots in the graph?

The combination logic relies on specific features of the filename. Basically, it splits the filename into parts separated by dashes ('-') and tries to group them on repeated parts, with a special case for numerical suffixes. But since you're using underscores between the parts of your file names, that doesn't work and everything becomes its own group.

The group separation must be enabled in the Flent invocation adding `--override-group-by both`

How can I change grouping logic in the graph?

You just need to switch the grouping mode, which is what the 'override group_by attribute' is for. Put 'series' into that, and it will switch the plot around. If you want to filter out the average, you put 'TCP download avg' into the 'filter series' box.

How can I have a different color for each data series?

just add to the Flent invocation `--override-colour-mode` which you can set to 'series' to get what you want

How can I add borders to the histograms?

You can customize the plot appearance by creating a matplotlib rc file (`~/.config/matplotlib/matplotlibrc`). See <http://matplotlib.org/users/customizing.html>

The value you'll want to set to get borders around bars is:

```
patch.force_edgecolor: true
```

And if you want to change the colour:

```
patch.edgecolor: 0000ff
```

Flent will override some matplotlib settings (unless you tell it not to via `--no-matplotlibrc`), but not those two.

general tip

always grab the last version of Flent on git, it is linked to the Flent main page