

Netbench – large-scale network device testing with real-life traffic patterns

Monday, 9 July 2018 15:30 (15 minutes)

Network performance is key to the correct operation of any modern datacentre infrastructure or data acquisition (DAQ) system. Hence, it is crucial to ensure the devices employed in the network are carefully selected to meet the required needs.

The established benchmarking methodology [1,2] consists of various tests that create perfectly reproducible traffic patterns. This has the advantage of being able to consistently assess the performance differences between various devices, but comes at the disadvantage of always using known, pre-defined traffic patterns (frame sizes and traffic distribution) that do not stress the buffering capabilities of the devices to the same extent as real-life traffic would.

Netbench is a network-testing framework, relying on commodity servers and NICs, that aims at overcoming the previously mentioned shortcoming. While not providing identical conditions for every test, netbench enables assessing the devices' behaviour when handling multiple TCP flows, which closely resembles real-life usage. e.g.:

- a full-mesh traffic distribution [3] is most likely the best benchmarking pattern for a network device to be employed in a multi-purpose data-centre;
- a unidirectional partial-mesh traffic distribution [3] will closely mimic the event-building traffic pattern from a DAQ system.

Due to the prohibitive cost of specialized hardware equipment that implements RFC tests [1,2], few companies/organisations can afford a large-scale test setup. The compromise that is often employed is to use two hardware tester ports and feed the same traffic to the device multiple times through loop-back cables (the so called "snake-test"). This test fully exercises the per-port throughput capabilities, but barely stresses the switching engine of the device. The per-port price of a netbench test setup is significantly smaller than that of a testbed made using specialized hardware, especially if we take into account the fact that generic servers can be time-shared between netbench and day-to-day usage. Thus, a large-scale multi-port netbench setup is affordable, and enables organisations/companies to complement the snake test with benchmarks that stress test the switching fabric of network devices.

Netbench has a layered architecture and uses standard technologies. At its core, it relies on iperf3 [4] as an engine to drive TCP flows between servers, but it can easily be extended to support other traffic generators. The orchestration platform that sets up multiple iperf3 sessions is written in Python and relies on XML-RPC for fast provisioning of flows. Per-flow statistics are gathered into a PostgreSQL database, and the results visualisation is based on a Python REST API and a web page using JavaScript and the D3.js library for displaying graphs. Statistics are presented at different levels of detail allowing the human tester to quickly assess the overall state of a test from both per-node and per-pair (source-destination) statistics.

During its last call for tender for high-end routers, CERN has employed netbench for evaluating the behaviour of network devices when exposed to full and partial mesh TCP traffic. We will present sample results from the evaluation. Furthermore, during the evaluation it became apparent that, due to the temporary congestion caused by competing TCP flows, netbench provides a good estimation of the devices' buffering capabilities.

To summarize, we present netbench, a tool that allows provisioning TCP flows with various traffic distributions (pairs, partial and full-mesh). We consider netbench an essential complement to synthetic RFC tests [1][2], as it enables affordable, large-scale testing of network devices with traffic patterns that closely resemble real-life conditions.

[1] RFC 2544, Bradner, S. and McQuaid J., "Benchmarking Methodology for Network Interconnect Devices"

[2] RFC 2889, Mandeville, R. and Perser J., "Benchmarking Methodology for LAN Switching Devices"

[3] RFC 2285, Mandeville, R., "Benchmarking Terminology for LAN Switching Devices"

[4] iperf3 <http://software.es.net/iperf/>

Primary authors: STANCU, Stefan Nicolae (CERN); KRAJEWSKI, Adam Lukasz (CERN); Mr CAEDDU, Matia (Università degli Studi di Cagliari (IT)); ANTOSIK, Marta (CERN); PANZER-STEINDEL, Bernd (CERN)

Presenter: STANCU, Stefan Nicolae (CERN)

Session Classification: T8 - Networks and facilities

Track Classification: Track 8 –Networks and facilities