High-speed Networks, Cybersecurity, and Software-defined Networking Workshop

Jorge Crichigno
University of South Carolina

2020 Western Academy Support and Training Conference
Summer Conference
June 15 – June 19

Lab 3: Emulating WAN with NETEM I: Latency, Jitter

Content

- Network Emulators
- WANs and Delays
- NETEM

Section 1: Network Emulators

Introduction to network emulators

- Network emulators play an important role for the research and development of network protocols and applications
- Network emulators provide the ability to perform tests of realistic scenarios in a controlled manner, which is very difficult on production networks
- This is particularly complex for researchers who develop and test tools for Wide Area Networks (WANs) and for multi-domain environments

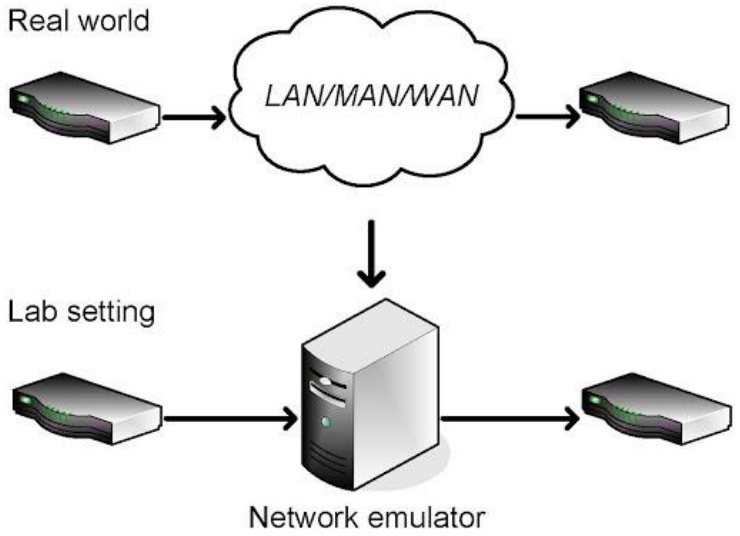

NETEM

- One of the most popular network emulators is NETEM, Linux network emulator for testing the performance of real applications over a virtual network
- The virtual network may reproduce long-distance WANs in the lab environment
- These scenarios facilitate the test and evaluation of protocols and devices from the application layer to the data-link layer under a variety of conditions
- NETEM allows the user to modify parameters such as:
  - Delay
  - Jitter
  - Packet loss
  - Duplication
  - Re-ordering of packet
- NETEM is implemented in Linux and consists of two portions, a small kernel module for a queuing discipline and a command line utility to configure it

Queuing Disciplines

- Queuing disciplines exist between the IP protocol output and the network device
- A queuing discipline is a simple object with two interfaces
- One interface queues packets to be sent and the other interface releases packets to the network device
- The queuing discipline makes the policy decision of which packets to send, which packets to delay, and which packets to drop

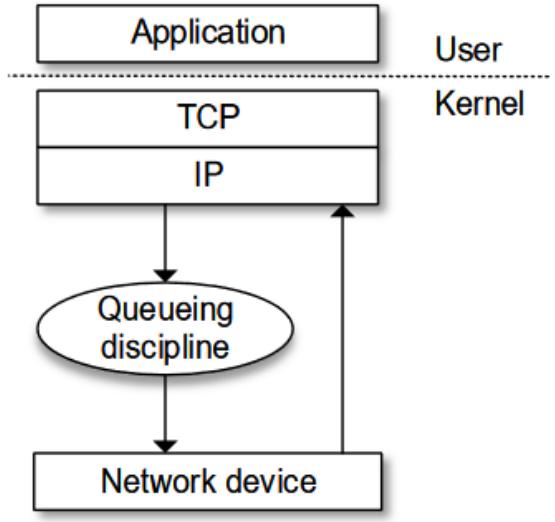

Section 2: WANs and Delay

WANs and Delay

- In networks, there are several processes and devices that contribute to the end-to-end delay between a sender node and a destination node
- Many times, the end-to-end delay is dominated by the WAN's propagation delay
- Consider two adjacent switches A and B connected by a WAN
- Once a bit is pushed on to the WAN by switch A, it needs to propagate to switch B
- The time required to propagate from the beginning of the WAN to switch B is the propagation delay
- The bit propagates at the propagation speed of the WAN's link
- The propagation speed depends on the physical medium
- The propagation delay is the distance between two switches divided by the propagation speed

WANs and Delay

- Network tools usually estimate delay for troubleshooting and performance measurements
- For example, an estimate of end-to-end delay is the Round-Trip Time (RTT), which is the time it takes for a small packet to travel from sender to receiver and then back to the sender
- The RTT includes packet-propagation delays, packet-queuing delays in intermediate routers and switches, and packet-processing
- As mentioned above, if the propagation delay dominates other delay components (as in the case of many WANs), then RTT is also an estimate of the propagation delay

Section 3: NETEM

NETEM (Delay Emulation)

- Network emulators emulate delays by introducing them to an interface
- For example, delay introduced to a switch A's interface that is connected to a switch B's interface may represent the propagation delay of a WAN connecting both switches
- The command below adds a delay of 100 milliseconds (ms) to the output interface

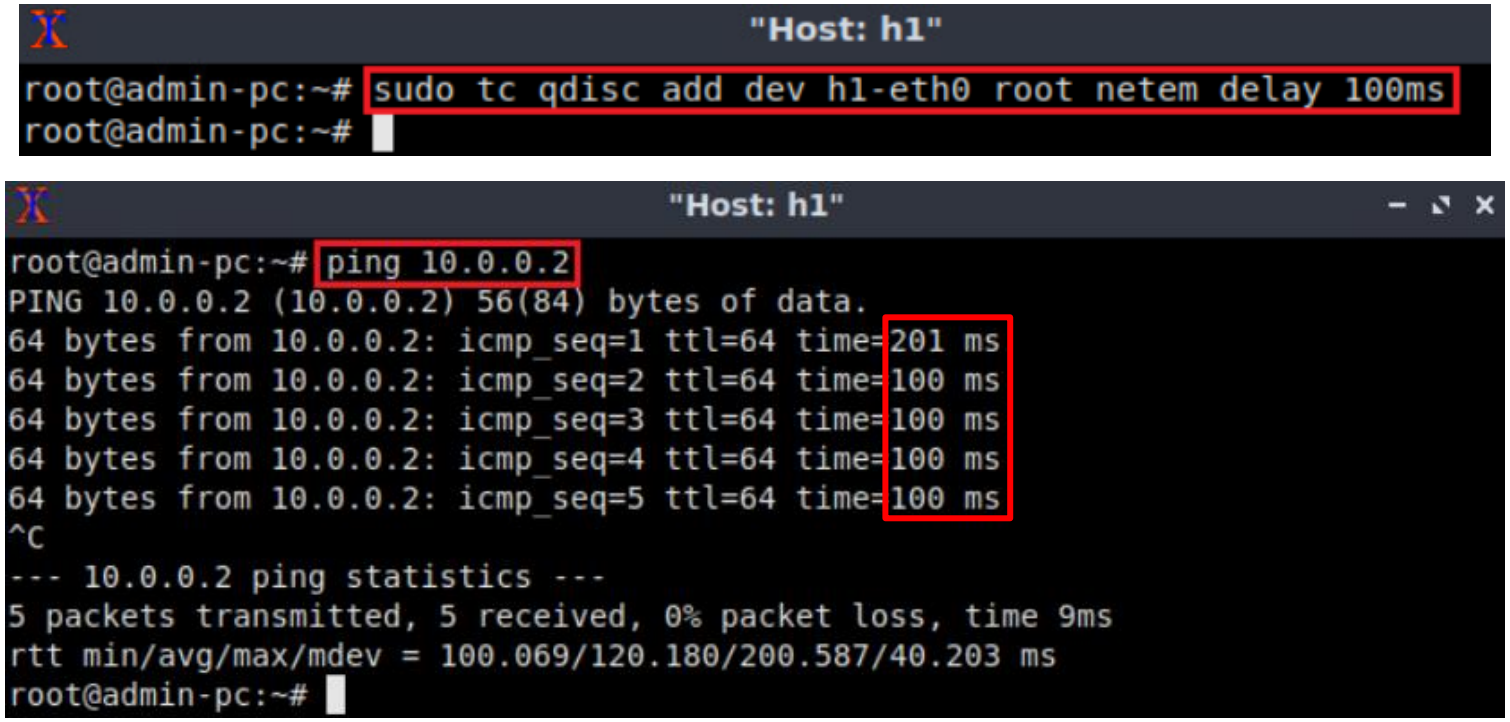

NETEM (Jitter Emulation)

- Networks do not exhibit constant delay
- the delay may vary based on other traffic flows contending for the same path
- Jitter is the variation of delay time
- The delay parameters are described by the average value, standard deviation, and correlation
- By default, NETEM uses a uniform distribution
- The new value added here represents jitter which defines the delay variation
- Therefore, all packets leaving host h1 via interface h1-eth0 will experience a delay of 100ms, with a random variation of ± 10ms

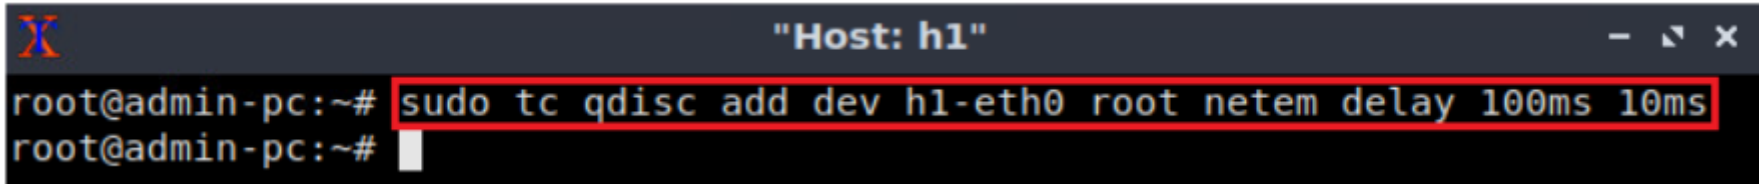

NETEM (Correlation value for Jitter and Delay)

- The correlation parameter controls the relationship between successive pseudo-random values
- The user will add a delay of 100 milliseconds with a variation of ± 10 milliseconds while adding a correlation value
- The new value added here represents the correlation value for jitter and delay
- Therefore, all packets leaving the device host h1 on the interface h1-eth0 will experience a 100ms delay time, with a random variation of ± 10 millisecond with the next random packet depending 25% on the previous one

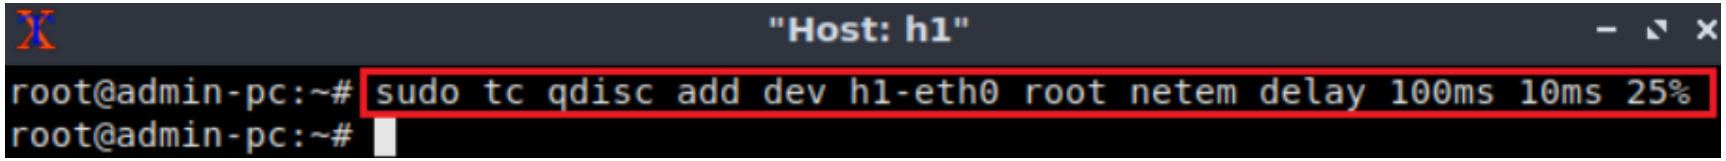

NETEM (Delay Distribution)

- NETEM permits user to specify a distribution that describes how delays vary in the network
- Usually delays are not uniform, so it may be convenient to use a non-uniform distribution such as normal, pareto, or pareto-normal
- The new option added here (distribution) represents the delay distribution type
- We define the delay to have a normal distribution, which provides a more realistic emulation of WAN networks
- As a result, all packets leaving the host h1 on the interface h1-eth0 will experience delay time which is normally distributed between the range of 100ms ± 20ms