



Realistic Traffic Generator

Hanoch Haim – Principal Engineer



FOSDEM



Agenda

- Overview
- Stateless
- Advance Stateful



TRex – Results

Open Source





Cisco Customers









TRex Usage Analytics monthly report (*)



(*) ~1200 distinct returning users,

TRex



TRex models of operation



 L7, DUT terminate TCP/SSL, flow based

 DUT inspect L7. does not change TCP. Flow based

 DUT L2/L3 Switch , packet based



What Problem is Being Solved?

- Networks include complex L4-7 features, such as
 - Load Balancer, DPI/AVC, Firewall, NAT



• Requires testing with stateful and realistic traffic mix



What Problem is Being Solved?

- Traffic generators for realistic traffic are
 - Expensive ~\$100-500K
 - Not scalable for high rates
 - Not flexible
- Implication
 - Limited and late testing
 - Different benchmarks and test methodologies
 - Real life bottlenecks and design issues

What is TRex?

- Linux user-space application uses DPDK library
- Stateless: Stream based uses Scapy
- Stateful: flow based

Generates, manipulates and amplifies based on templates of real, captured flows (W/O TCP stack)

- High performance: up to 200 Gb/sec
- Low cost: Standard server hardware
- Flexible and Open Software
- Virtualization
- Easy installation and deployment













Stateless



Stateless High level functionality

- High scale ~10M-35MPPS/core
- Profile can support multiple streams, scalable to 20K parallel streams
- Interactive support GUI/TUI
- Statistic per port/ stream (e.g. latency/ Jitter)
- Python automation support
- Multi-user support
- Capture to Wireshark
- Scalable services using plugins (e.g. DHCP, IPv6)

1.37% CPU	2.31 Hb/s Total Tx L2	3.03 Mb/s Total Tx L1		4.51 Kpkt/s Total RPS	Oh/s Drop Rate
0.10% Rx CPU	2.31 Mols Total Rx12	3 Total Stream		1 Active Ports	Opkts Queue Pul
Ports: All *	Ports Streams Latency Chart	3			
Port 0	Interval: 60 🔹				[
Port 1					
				Ť	
	40 45 40	45 40 45	-50 -35	41 .11	da da
			Time (c)		



Traffic Profile Example





Control plane High level





One stream with two directions



Python Automation example

```
c = STLClient(username = "itay", server = "10.0.0.10", verbose level = LoggerApi.VERBOSE HIGH)
```

```
try:
    # connect to server
   c.connect()
    # prepare our ports (my machine has 0 < --> 1 with static route)
    c.reset(ports = [0, 1])
    # add both streams to ports
    c.add streams(s1, ports = [0])
    # clear the stats before injecting
   c.clear stats()
   c.start(ports = [0, 1], mult = "5mpps", duration = 10)
    # block until done
    c.wait on traffic(ports = [0, 1])
    # check for any warnings
   if c.get warnings():
      # handle warnings here
      pass
```

finally:

```
c.disconnect()
```

Performance XL710 MPPS/Core



link

TRex



Advanced Stateful





User space TCP stack





TRex ASTF features

- High scale
- TCP is the core component
 - Can be tuned MSS/initwnd/delay-ack
- TCP is based on BSD with acceleration
- Interactive
- Accurate latency measurement usec
- Simulation of latency/jitter/drop in high rate
- OpenSSL integration
- L7 emulation layer
 - Emulate application using "micro-instructions
 - Field engine





L7 Emulation layer

Client

```
send(request)
wait_for_response(len<=1000)
delay(random(min=100,max=1000)) // in usec
send(request2)
wait_for_response(len<=2000)
close()</pre>
```

Server side

```
wait_for_request(len<=100)
send_response(response)
wait_for_request(len<=200)
send_response(response)
wait_for_peer_close()</pre>
```



HTTP simple profile

```
from trex astf lib.api import *
class Prof1():
    def init (self):
        pass
    def get profile(self):
        # ip generator
        ip gen_c = ASTFIPGenDist(ip_range=["16.0.0.0", "16.0.0.255"],
                                 distribution="seq")
        ip gen s = ASTFIPGenDist(ip range=["48.0.0.0", "48.0.255.255"],
                                  distribution="seq")
        ip gen = ASTFIPGen(glob=ASTFIPGenGlobal(ip offset="1.0.0.0"),
                                                                           0
                           dist client=ip gen c,
                           dist server=ip gen s)
        return ASTFProfile (default ip gen=ip gen,
                            cap_list=[ASTFCapInfo(
                                      file="../avl/delay 10 http browsing 0.pcap"
                                      cps=1)
                                                                            0
                                     1)
def register():
    return Prof1()
```



Client side pseudo code

Client side pseudo code

<pre>template = choose_template()</pre>	0
<pre>src_ip,dest_ip,src_port = gene dst_port = temp</pre>	erate from pool of client plate.get_dest_port()
s = socket.socket(socket.AF_IN	IET, SOCKESTREAM)
<pre>s.connect(dest_ip,dst_port)</pre>	0
# program	Ø
s.write(template.request)	#write the following taken from the pcap file
	# GET /3384 HTTP/1.1
	# Host: 22.0.0.3
	# Connection: Keep-Alive
	# User-Agent: Mozilla/4.0
	# Accept: */*
	# Accept-Language: en-us
	# Accept-Encoding: gzip, deflate, compress
s.read(template.request_size)	# wait for 32K bytes and compare some of it
	#HTTP/1.1 200 OK
	#Server: Microsoft-IIS/6.0
	#Content-Type: text/html
	#Content-Length: 32000
	# body
s.close();	



Server side pseudo code

# if this is SYN for flow that alrea	ndy exist, let TCP handle it
<pre>if (flow_table.lookup(pkt) == False # first SYN in the right direct compare (pkt.src_ip/dst_ip to th</pre>	:) : .on with no flow we generator ranges) # check that it is in the range or valia
server IP (src_ip,dest_ip)	
template= lookup_template(pkt.de	<pre>st_port) #get template for the dest_port</pre>
<pre># create a socket for TCP server s = socket.socket(socket.AF_INET</pre>	:, socket.SOCK_STREAM)
<pre># bind to the port s.bind(pkt.dst_ip, pkt.dst_port)</pre>	
s.listen(1)	
#program of the template	0
<pre>s.read(template.request_size)</pre>	# just wait for x bytes, don't check them
	# GET /3384 HTTP/1.1
	# Host: 22.0.0.3
	# Connection: Keep-Alive
	# User-Agent: Mozilla/4.0
	# Accept: */*
	<pre># Accept-Language: en-us</pre>
	# Accept-Encoding: gzip, deflate, compress
s.write(template.response)	# just wait for x bytes,
	<pre># don't check them (TCP check the seq and checksum)</pre>
	#HTTP/1.1 200 OK
	#Server: Microsoft-IIS/6.0
	#Content-Type: text/html
	#Content-Length: 32000
	# body
s.close()	



Profile with two template

```
class Prof1():
       def init (self):
        pass
       def get profile(self):
       # ip generator
       ip_gen_c = ASTFIPGenDist(ip_range=["16.0.0.0", "16.0.0.255"],
                                 distribution="seq")
       ip gen s = ASTFIPGenDist(ip range=["48.0.0.0", "48.0.255.255"],
                                 distribution="seq")
        ip gen = ASTFIPGen(glob=ASTFIPGenGlobal(ip_offset="1.0.0.0"),
                           dist client=ip gen c,
                           dist server=ip gen s)
       return ASTFProfile(default ip gen=ip gen,
                            cap list=[
                             ASTFCapInfo(file="../avl/delay 10 http browsing 0.pcap",
                                         cps=1),
                                                    0
                             ASTFCapInfo(file="avl/delay 10 https 0.pcap",
                                                    0
                                         cps=2)
                                      (1)
def register():
   return Prof1()
```

Statistic

|--|

	client	T	server	
m_active_flows	39965	T	39966	active flows
m_est_flows	39950	1	39952	active est flows
m_tx_bw_17_r	31.14 Mbps	1	4.09 Gbps	tx bw
m_rx_bw_17_r	4.09 Gbps	1	31.14 Mops	rx bw
m_tx_pps_r	140.36 Kpps	1	124.82 Kpps	tx pps
m_rx_pps_r	156.05 Kpps	1	155.87 Kpps	rx pps
m_avg_size	1.74 KB	1	1.84 KB	average pkt size
-		1		
TCP	1	1		
-	1	1		
tcps_connattempt	73936	1	0	connections initiated
tcps_accepts	1 0	1	73924	connections accepted
tcps_connects	73921	1	73910	connections established
tcps_closed	33971	1	33958	conn. closed (includes drops)
tcps_segstimed	213451	1	558085	segs where we tried to get rtt
tcps_rttupdated	213416	1	549736	times we succeeded
tcps_delack	344742	1	0	delayed acks sent
tcps_sndtotal	623780	1	558085	total packets sent
tcps_sndpack	73921	1	418569	data packets sent
tcps_sndbyte	18406329	1	2270136936	data bytes sent
tcps_sndctr1	73936	1	0	control (SYN,FIN,RST) packets sent
tcps_sndacks	475923	1	139516	ack-only packets sent
tcps_rcvpack	550465	1	139502	packets received in sequence
tcps_rcvbyte	2269941776	1	18403590	bytes received in sequence
tcps_rcvackpack	139495	1	549736	rcvd ack packets
tcps_rcvackbyte	18468679	1	2222057965	tx bytes acked by rovd acks
tcps_preddat	410970	1	0	times hdr predict ok for data pkts
tcps_rcvoopack	1 0	1	0	*out-of-order packets received #
-	1	1		
Flow Table	1	1		
-	1	1		
redirect rx ok	1 0	1	1	redirect to rx OK



Under the hood





TCP stack – Flow Scale –TX





TCP stack - Flow Scale issue - RX



TCP stack – Delay/Jitter/Drop simulation



TRe



TRex vs NGINX



3



Performance setup #2





Performance numbers

	TRex one DP core									
	CPU (1DP)							Memory		
m	%%	cps	rps	rx (mb/sec)	pps(tx+rx)	active flows	drop	тср/мв		
1000	0.6	1000	1000	265	34444	600	0	0.3		
5000	2.7	5000	5000	1320	172222	3000	0	1.4		
10000	5.7	10000	10000	2210	344444	6000	0	2.9		
20000	13.5	20000	20000	5410	688889	12000	0	5.7		
50000	40.8	50000	50000	13480	1722222	29870	0	14.2		
87500	79.0	87500	87500	23670	3013889	55329	0	26.4		
90000	86.1	90000	90000	24271	3100000	56217	0.10%	26.8		

x80 faster x2000 less memory

Figure 3. TRex with 1 DP core

					Linux 16 cores				
								Total	
							Kernel	Memory	
			rx				memory	used (free-	
m	cps	rps	(mb/sec)	active flows	16xCPU	drop	SLAB (MB)	h) MB	
1000	1000	1000	265	600		no		21000	
5000	5000	5000	1320	3000		no		22000	
10000	10000	10000	2210	6000		no		25000	
					20%/25% 8x				
20000	20000	20000	5410	12000	cores IRQ break	yes	800	31000	
50000	50000	50000	13480	29870					
87500	87500	87500	23670	55329					
90000	90000	90000	24271	56217					

Figure 4. NGINX 16 cores

https://trex-tgn.cisco.com/trex/doc/trex_astf_vs_nginx.html