



BUILDING 100G DDOS MITIGATION DEVICE WITH FPGA TECHNOLOGY

Martin Žádník
CESNET

2018
Brno



- DDoS attacks
- DDoS attacks as a service
- DDoS-for-hire industry
- Booters/Stresser service
- Mirai

Krebs on Security
In-depth security news and investigation



■ AKAMAI

- Several hundreds DDoS per year
- Largest more than 1 Tbps

■ CESNET

- Order of magnitude lower volume
- Similar amount
- Testing playground

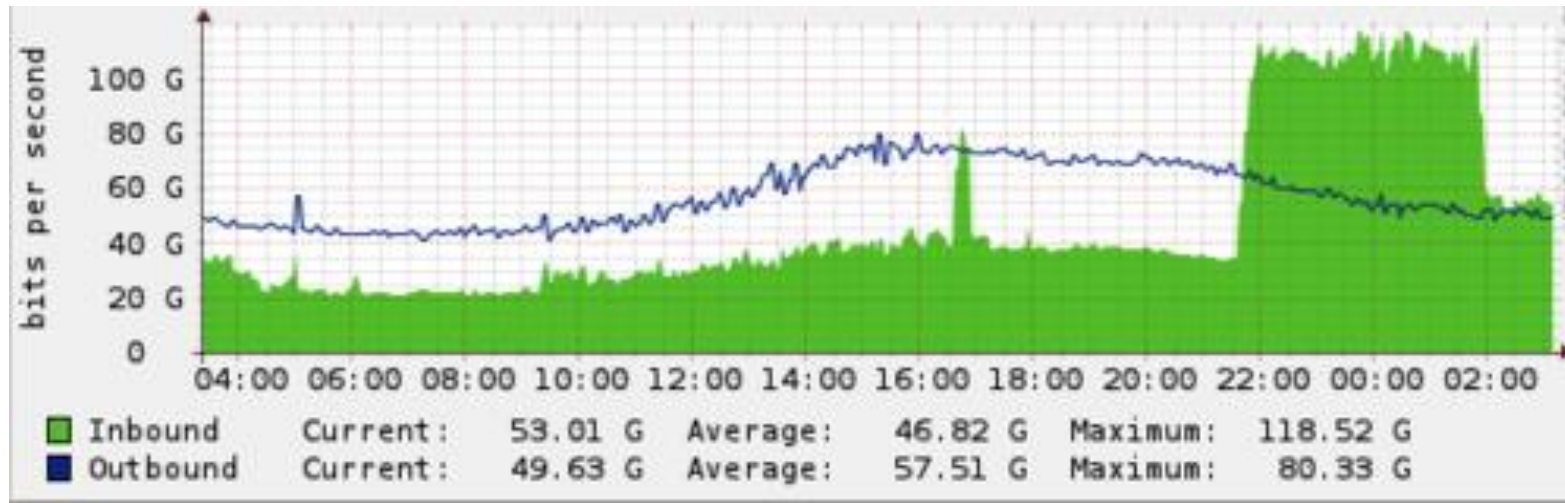
■ RTBH and Rate limiting at routers

- Too coarse grain
- Legitimate traffic is rate-limited together with attack

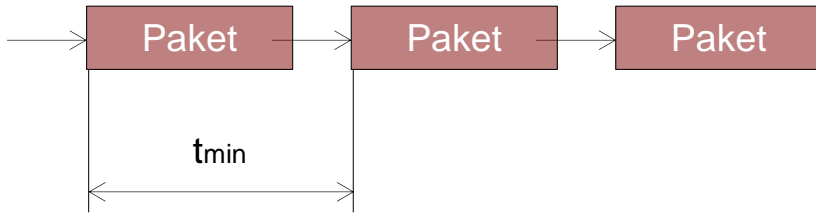
■ What's needed

- More fine grained
- Order of magnitude cheaper
- Customizable
- Own solution

- To protect infrastructure (connectivity)
- To reduce extensive amount of traffic targeting victim organization under the limit which can be actually processed by the organization

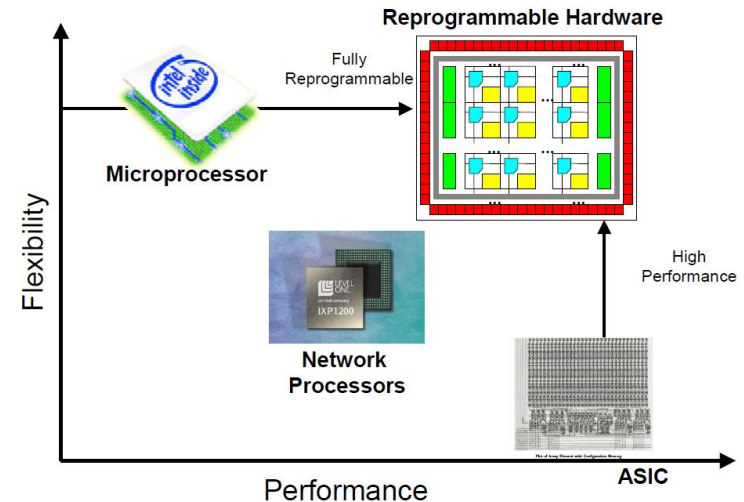


■ CESNET experience with network flow probes



40 Gb/s	12 ns	~	45 CPU clock cycles
100 Gb/s	5 ns	~	18 CPU clock cycles
400 Gb/s	1.25 ns	~	6 CPU clock cycles

3.6 GHz CPU

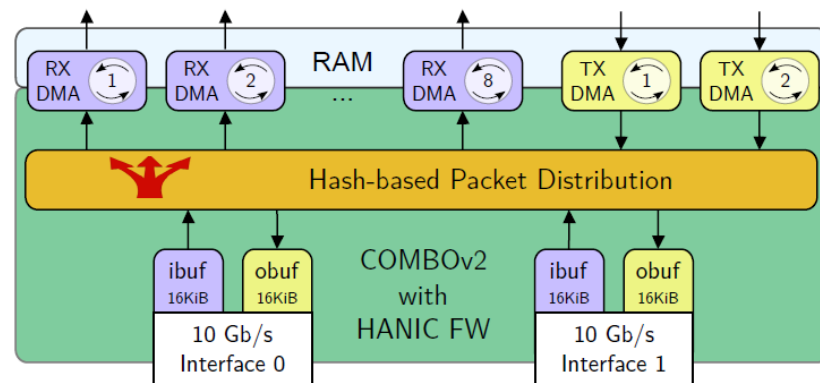


John Lockwood, Stanford University

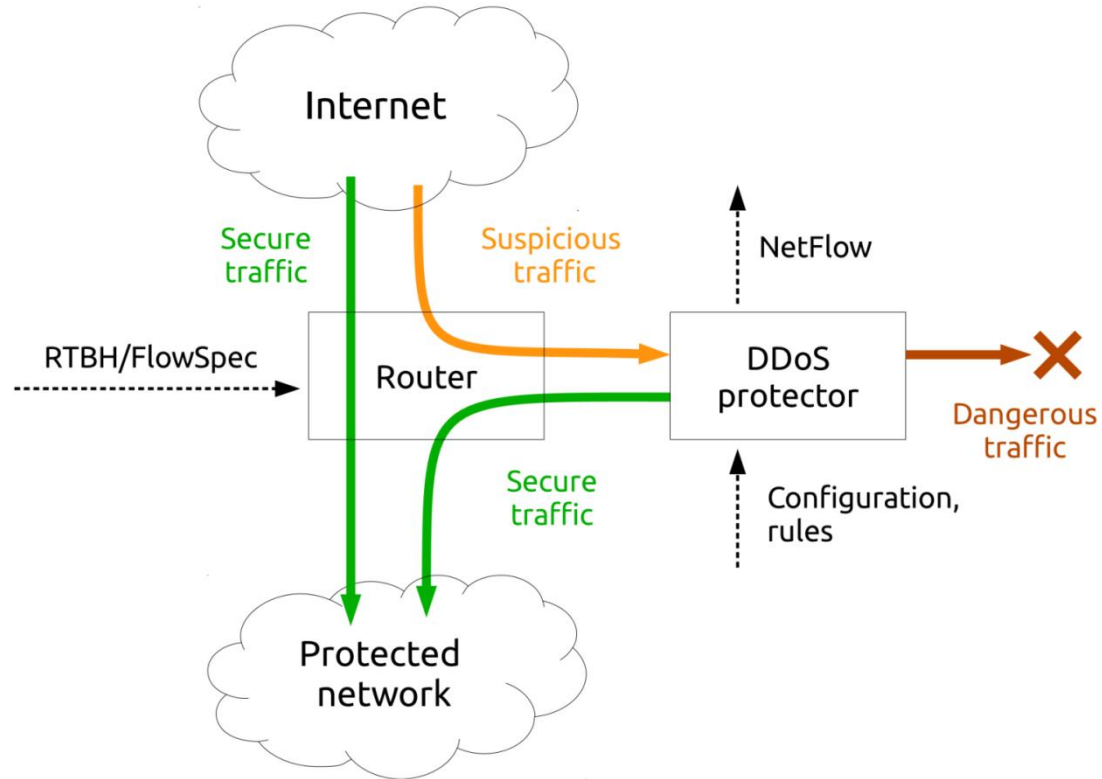
■ CESNET experience with network flow probes

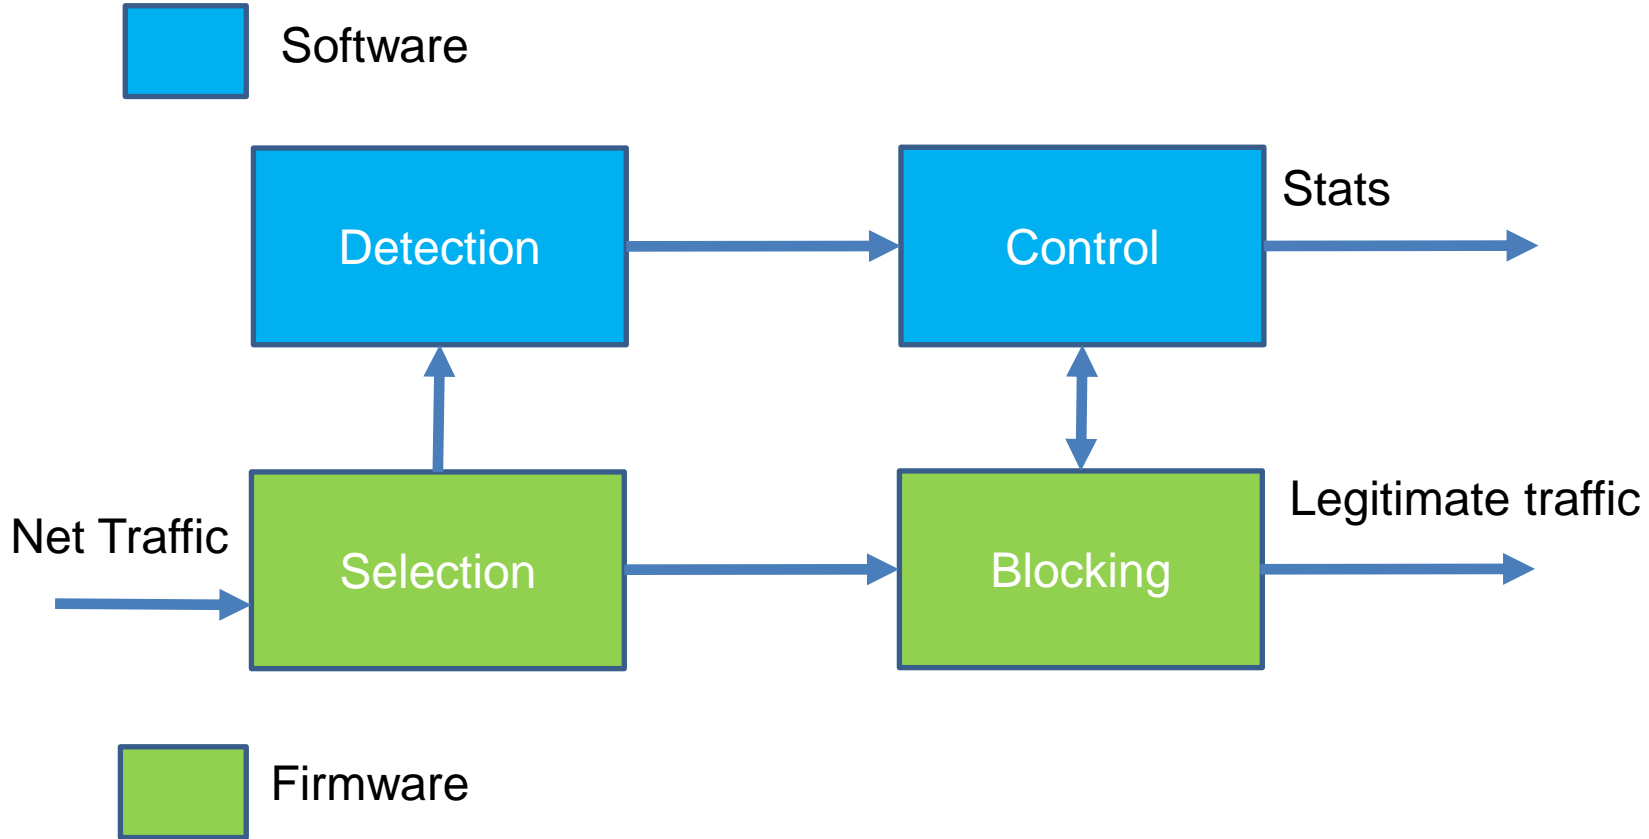
■ Platform

- Network card with programmable FPGA
- Own firmware into FPGA
- Decent server with threaded software



- 10x 10Gbps
- 1x 100 Gbps






■ Deal with how to deploy

- Support of VLAN translation
- Support of routing
- Support of ARP, ND
- Dead-man's vigilance device

■ Utilize what is already available

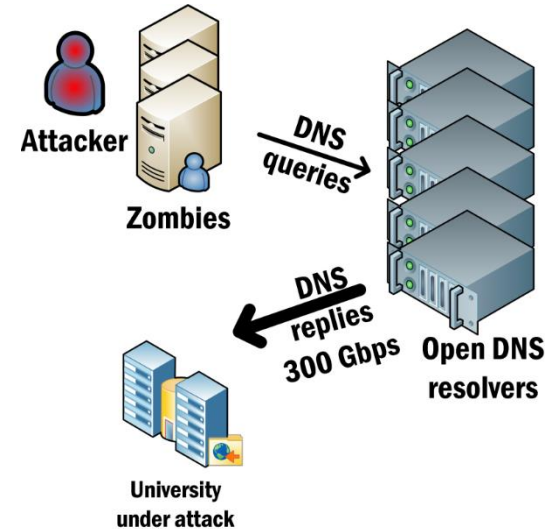
- BIRD, Suricata (to be utilized)

■ Practical and straight-forward approach usually works well

- Single-direction only
 - Heuristics to deal with various types of attacks
- 
- A decorative horizontal bar at the bottom of the slide consisting of a series of small blue squares of varying heights and widths, creating a pixelated or digital effect.

■ Large reflection attacks

- DNS
- NTP
- LDAP
- SSDP
- SNMP
- CharGEN



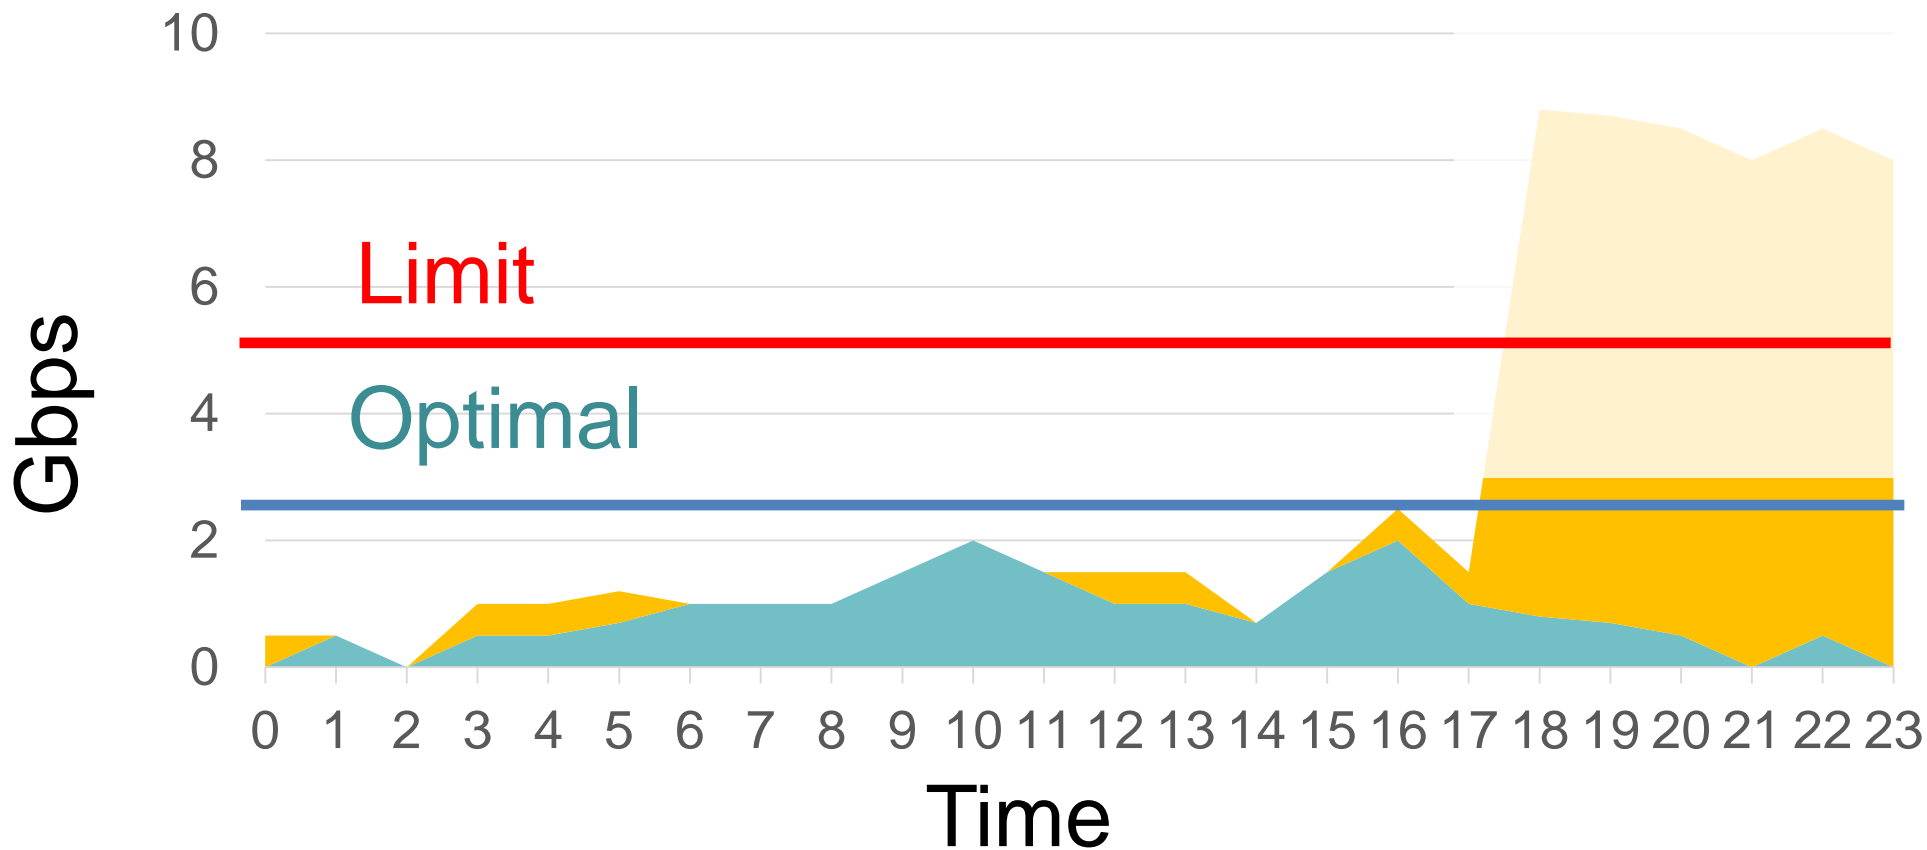
■ TCP SYN flood

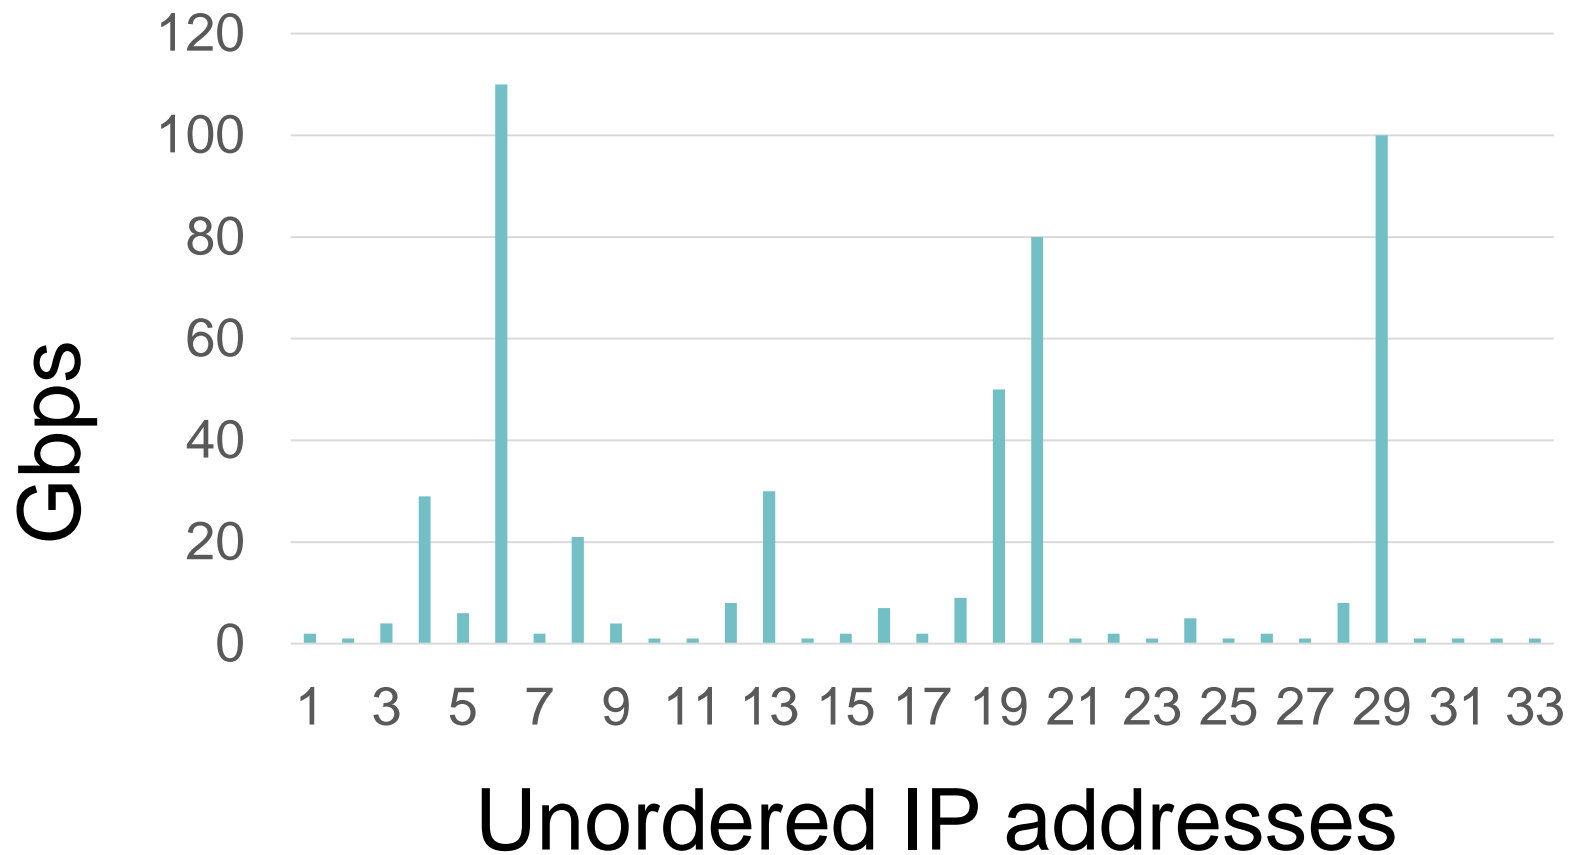
- Protector looks for exceeding traffic thresholds per IP prefixes
- Time window is configurable (default 1 s)
- Simple static rules set by administrator

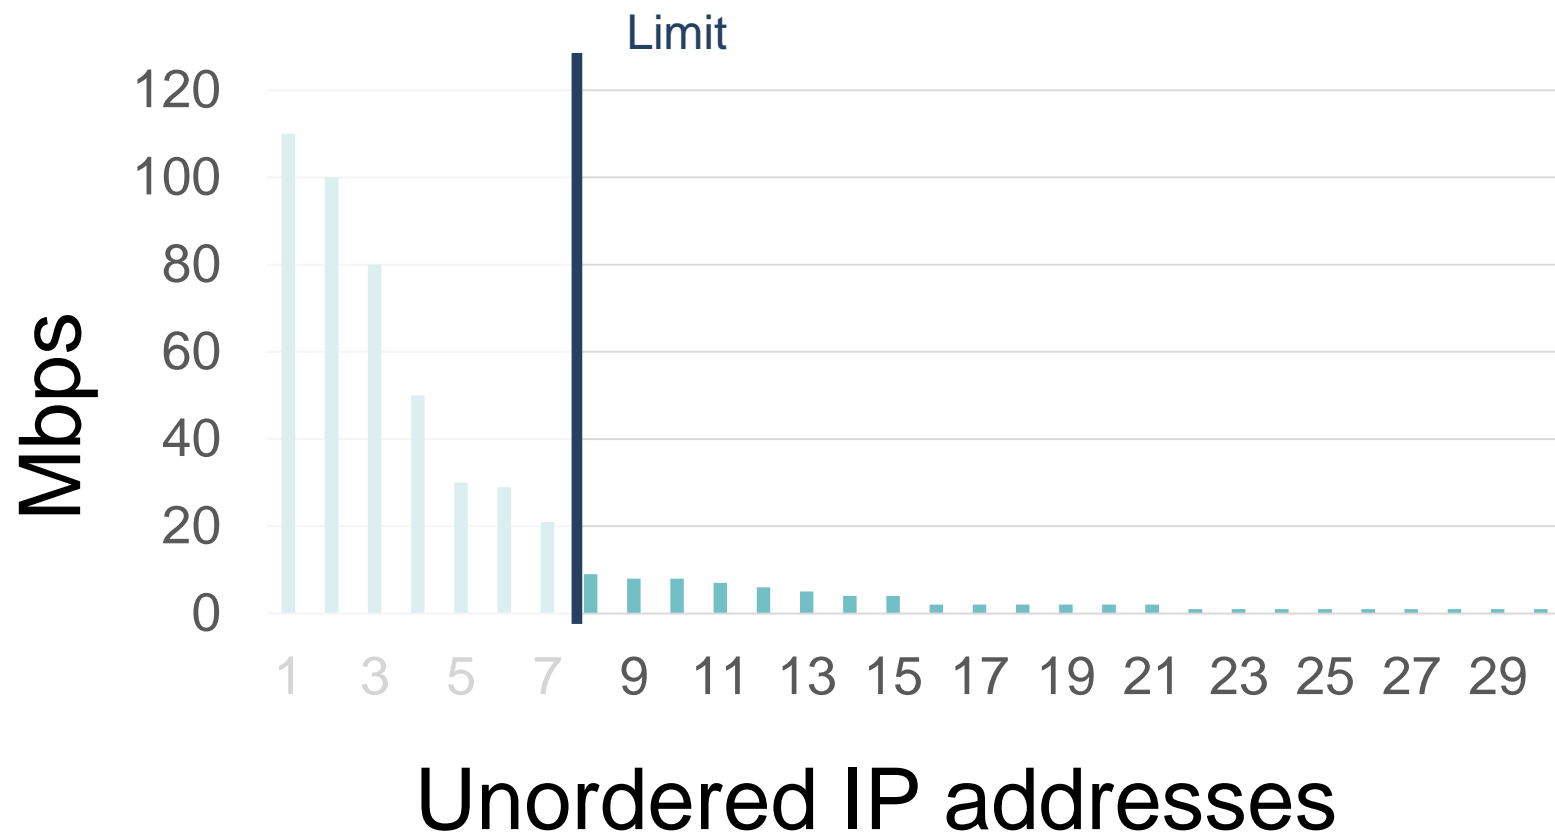
„VUT UDP“ dst net 147.229.0.0/16 protocol 17 src port 53
threshold 1 Gbps limit 100 Mbps

If matching traffic
exceeds 1+ Gbps
then I reduce it to
100 Mbps

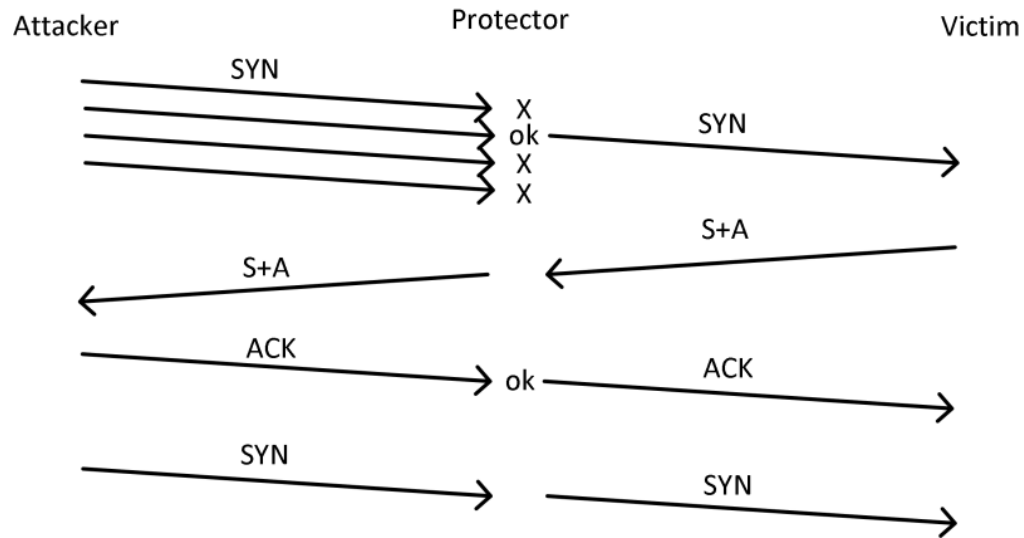
- Drop matching traffic from IP addresses that contributed the most to exceeding the threshold
- To this end
 - Keep contribution of each IP address
 - If threshold is exceeded choose such a number of IP address to reduce the traffic below limit



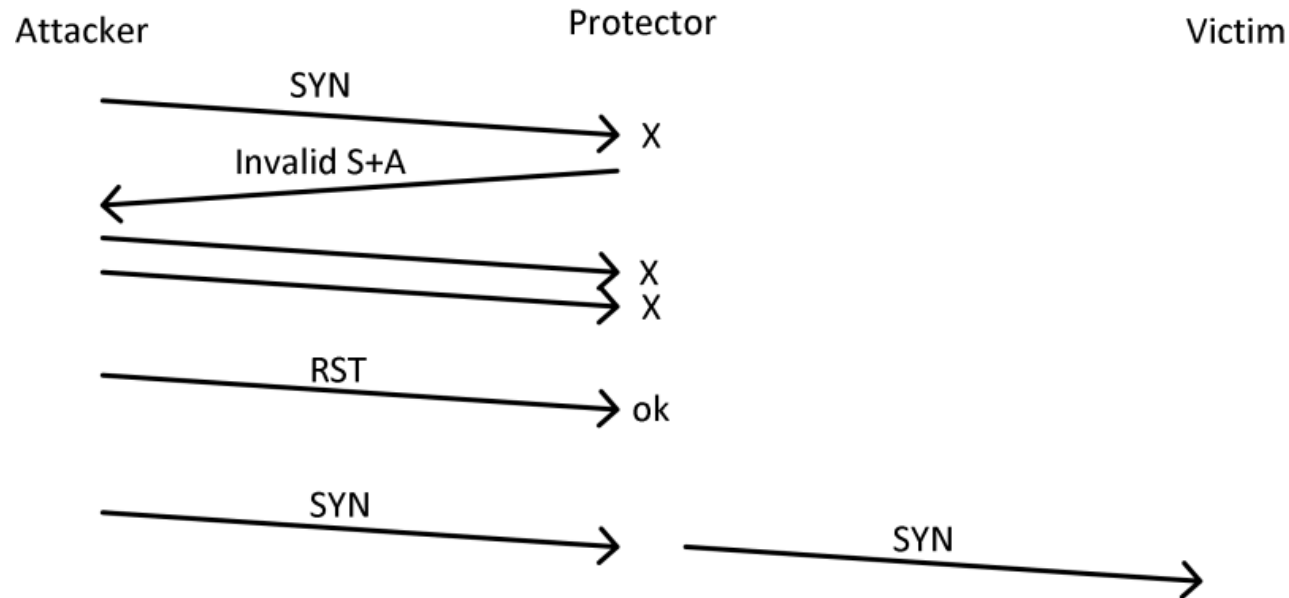




■ SYN drop heuristic



- RST cookies – Alternative to SYN drop
- Protector generates non-valid SYN-ACK packet
- If a client sends RST then whitelisted



- Wire speed throughput 100Gbps
- Extremely low latency (microseconds)
- Support IPv6
- TCP flags
- Fragments
- Configuration: Linux CLI + database rules
- Stats: SNMP, logs

- Extended blocking capacity
- Support various heuristics
- Build less proprietary interface
 - BGP FlowSpec
 - Cisco-like CLI
- Release
 - Polish it till anyone can use it
 - Offer to others

- Straightforward extensible and customizable solution
- Deployed in productional CESNET backbone
- Interest of other entities



THANK YOU FOR YOUR ATTENTION

- Forward suspicious traffic to Protector
- Return cleansed traffic to target destination

