Mitigation of DoS Attacks Using Machine Learning

Author: *Ing. Patrik Goldschmidt*
Supervisor: *Ing. Jan Kučera*

BRNO UNIVERSITY OF TECHNOLOGY — FACULTY OF INFORMATION TECHNOLOGY
cesnet

Thesis Aims

- Propose, implement, and evaluate a method able to detect and mitigate (D)DoS attacks in real-time
- Legitimate clients must not be significantly affected

DoS Attacks

Denial of Service (DoS) is a cybersecurity attack aiming to disrupt the availability of the computer system. This is typically achieved by sending specific messages that trigger an undesired state on the victim's machine or overload all available resources.

Its distributed variant (DDoS) employs a large quantity of attacking computers, making its impact much bigger and defense considerably more challenging. DoS/DDoS attacks have been around for more than 20 years now, but their utilization is still growing, often with devastating results.

Recent advances in machine learning have slightly improved the attack detection process, but performing it efficiently in real-time while keeping the ratio of false positives low is still a subject of extensive scientific research nowadays.

Our Approach

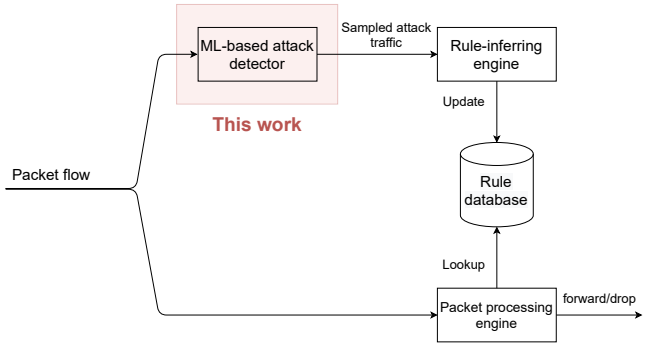

Method

1. ***Detect the attack near real-time***
   (few seconds delay to gather traffic statistics)
2. Sample the attacker's data
3. Infer packet decision rules
4. Use them to forward/drop packets with *ns* delays

Attack Detection Mechanism

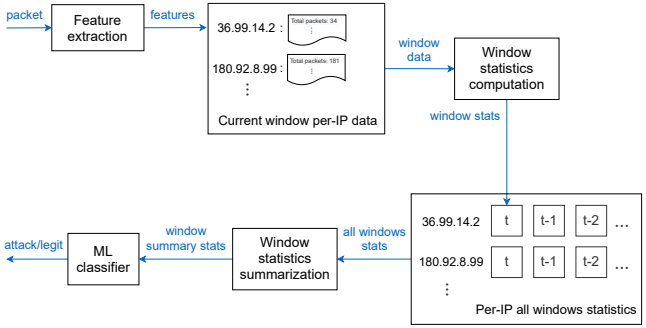

- Custom per-packet feature extraction
  - Arrival timestamp, packet size, IPs, L4 protocol, L4 ports
- Data aggregation by IP using time windows
- Statistics computation
  - Within each window
  - Between several windows
- 32 statistical features fed into the ML classifier
  - E.g.: byte rate, packet arrivals time variance, packet sizes variance, headers to payload size ratio, src port entropy,...
- Used techniques:
  - Data mining in streams (aggregation, approximation algs., sampling, sketching, windowing)
  - Supervised machine learning

Datasets

- Similar research is mostly evaluated on 1 dataset
- Increase credibility and usability of the proposed system in practice by combining 7 different datasets
  - *Real data (2):* CAIDA 2007 DDoS Attack, CESNET's internal
  - *Synthetic data (5):* Canadian Institute for Cybersecurity (CIC)
- Additionally, any raw packet capture file (PCAP) can be used to generate new datasets on demand

Evaluation

- Method evaluated with 12 different ML models
- Despite combining several datasets, the results were biased towards some features (ex. ICMP ratio)
- Even after dropping these low-generalizing features, the results were still excellent (>99%)
- Nevertheless, such high accuracy is nothing uncommon in the ML DDoS detection field

System performance with respect to the number of features

6 to 8 features are seemingly enough to distinguish between attack and legitimate traffic very reliably.

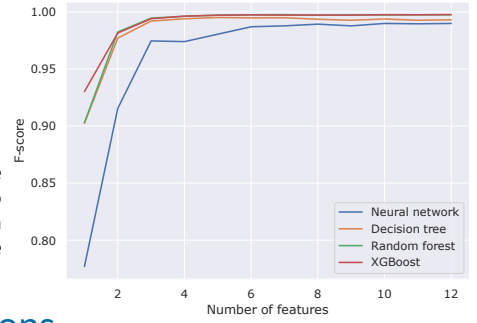

Contributions

- Unique approach for attack detection
- Reacts well on both volumetric and Slow DoS-es
- Detection within 4 seconds of the attack start
- Further development and usage in CESNET's research project "DDoS Protector" granted by MVČR