Improving the DDoS resistance of Internet services

Aapo Kalliola, Tuomas Aura
Aalto University, Finland

Flooding attacks

Distributed denial-of-service (DDoS) attacks by botnets and cyber terrorists threaten critical Internet-based services. Internet services such as web servers are flooded with millions of unnecessary requests or packets. DDoS mitigation currently requires heavy investment in network and server capacity, firewall hardware, or outsourced mitigation services. Manual traffic monitoring and filter configuration is a burden to system administrators, and manual defenses cannot keep up with fast-changing attacks.

DDoS filtering

The goal of our research is to create an *automatic filtering mechanism for DDoS attacks*. It should be inexpensive and easily deployable in the current services and networks.

Our approach is to use learn a model of the normal traffic and, during an attack, prioritize traffic that best fits the normal traffic model. In a way, we prioritize the "regular customers" of the service.

Cluster model

In the learning phase, we create a hierarchical resource usage tree of the normal traffic , mainly client IP address, using a variant of the Hierarchical Heavy Hitter algorithm.

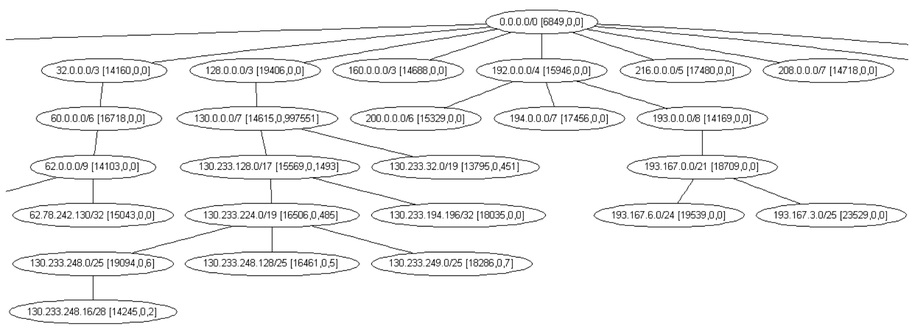

When the server load exceeds the capacity, traffic clusters with the smallest ratio of measured traffic to normal traffic are prioritized.

The figures below show the effectiveness of the defense mechanism when the attack traffic is *10 x* the normal traffic. Without the filtering, only ~10% of the honests requests are served during the attack (left graph), making the service unusable to all clients. With filtering, ~70% of the honest clients are served relatively well (right graph).

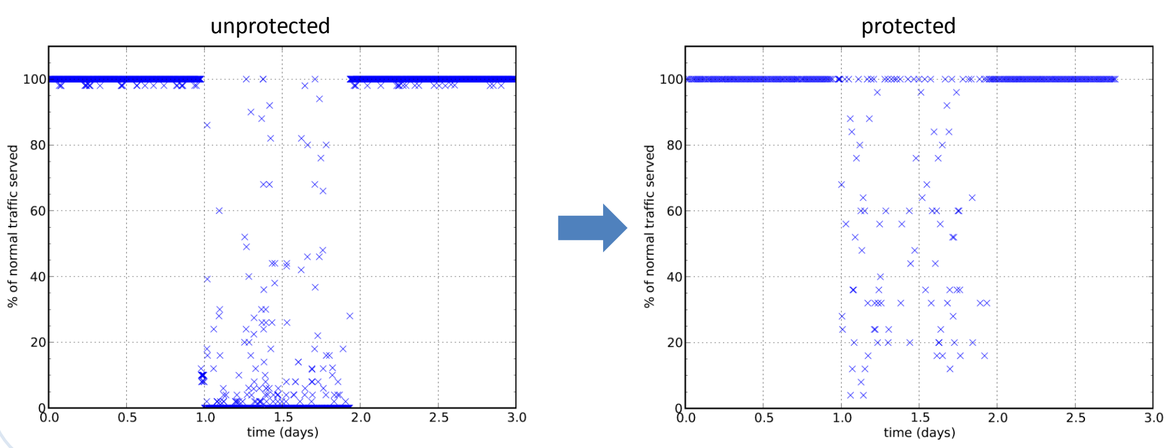

Simulation

Evaluation results from simulations are shown below. The percentage of honest clients served during the attack increases as the number of clusters grows. The optimal number of clusters depends on the variability of the normal traffic and the filter implementation.

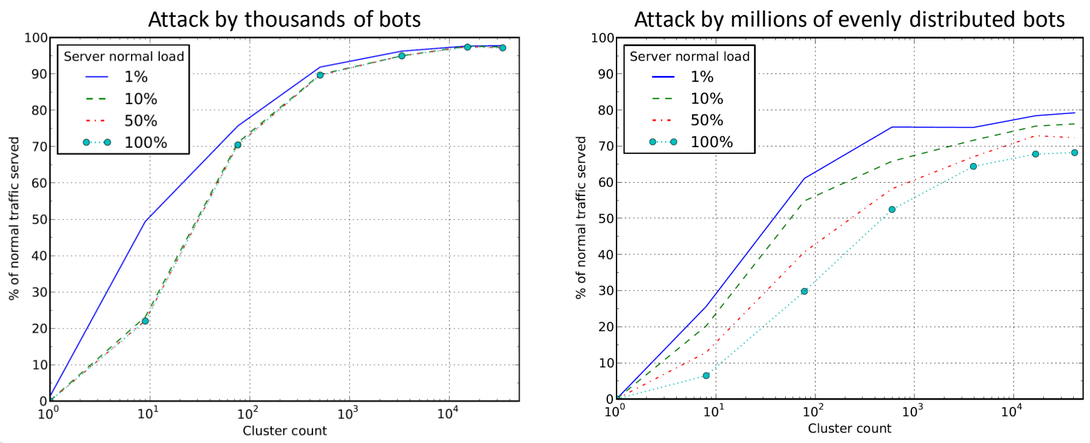

Prototype implementation

We implemented the filtering in a firewall. A separate computer samples the traffic and computes the normal traffic model. During an attack, it updates the IP prefix white list in the firewall filter in a tight control loop. Only a single prefix-tree lookup is needed on the performance-critical path in the firewall to accept or reject a single packet.

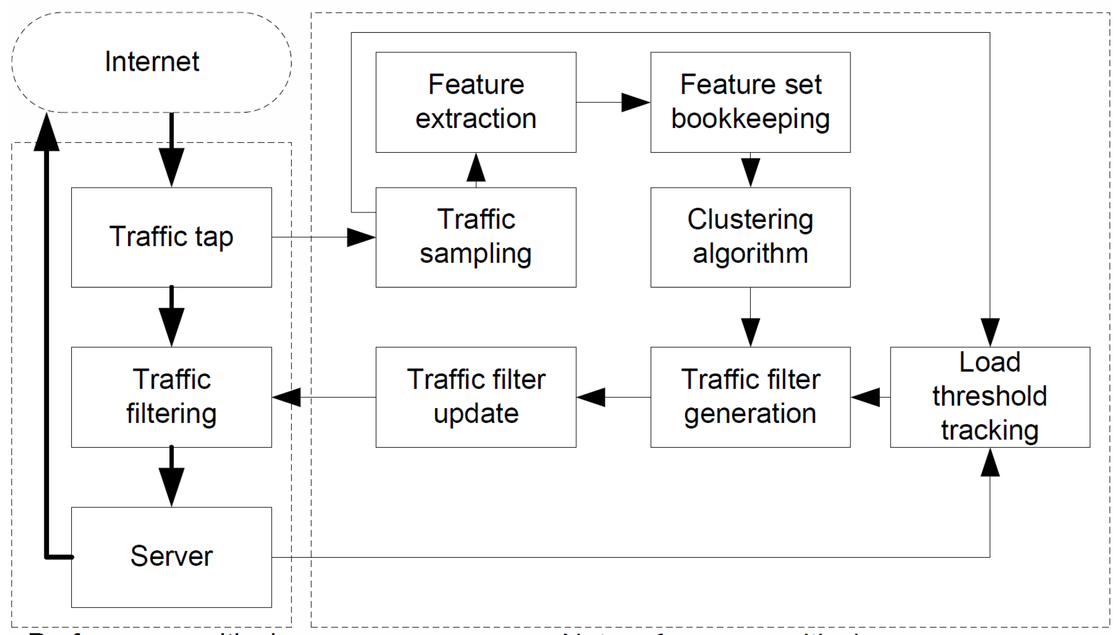

Summary

The simulations and experiments with the prototype implementation indicate that the learning and filtering mechanism is effective in automatic filtering of flooding DoS and DDoS attacks with quite different attack traffic distributions. It works for both IP packet flooding and HTTP request flooding against a web server.

The mechanism is easy to deploy in current networks and it works without human interaction. It is particularly useful for small service providers that cannot afford to pay for the outsourced mitigation services. The research continues on efficient filtering implementations for high-bandwidth networks.