POCS OP4 : Proactive bandwidth flooding resistant internet

November 6, 2019

In a bandwidth flooding attack, a large amount of compromised end systems (*adversaries*) send packets to a victim end system (*client*).

A proactive solution against such attacks targeted at a client tries to mitigate the flooding before the client tail's circuit is overflowed. We discuss here such a solution based on mutual trust of routers and throughput control at the network layer.

The main idea behind this solution is to split the traffic incoming to the client gateway by selecting several other routers as *gatekeepers* for the traffic.

To select them, the client's gateway first broadcasts its need for gatekeepers using a network layer protocol such as BGP. The willing gatekeepers use BGP to lead traffic intended for the client's IP to themselves and forward it to the gateway.

These gatekeepers each receive a fraction of the traffic sent to the client, this way, the traffic sent to the client is guaranteed to go through one of the gatekeepers first.

To make sure that the traffic is split as evenly as possible, gatekeepers need to be at a sufficiently large hop distance from each other and from the gateway, they can even be in a different AS. Gatekeepers and the gateway need to check each other's liveliness regularly and the gateway can ask for more gatekeepers if needed. If it is not enough, gatekeepers with more client traffic can redirect a fraction of it to another gatekeeper using the gateway to retrieve gatekeepers' information.

The client provides a throughput threshold for its tail circuit to the gateway. The gateway can tune each gatekeeper's allowed throughput and make them comply to it willingly or not. If the incoming client traffic exceeds its tail circuit threshold, the gateway will ask gatekeepers with the most throughput to reduce it to a desired amount. If a gatekeeper doesn't cooperate, the gateway will arbitrarily drop packets coming from it to reach this amount.

As the goal of routers on the internet is to have good interconnections to other routers they will have no incentive to let the gateway drop their packets. Thus to compensate the decrease in throughput the gateway asked, the gatekeeper will in turn ask its neighbouring routers to reduce their throughput and recursively propagate until reaching a router not willing to lower its throughput : an adversary.

The adversary traffic will be arbitrary decreased by the last router in the path from the client and every router on this path will have lowered its throughput to the client. After some amount of time, provided that the threshold is not reached, the gateway can set back a higher limit for the gatekeeper.

This throughput limiting mechanism is akin to a trust system, in which the allowed amount of traffic coming from each gatekeeper roughly translates to their trust factor. An adversary sending too much traffic will be muted and its trust lowered.