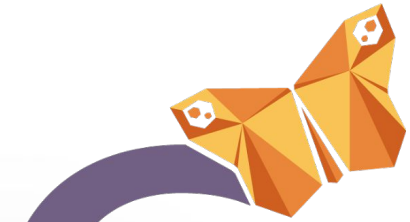

Gatekeeper

First Open Source DDoS Protection System

Michel Machado Cody Doucette Qiaobin Fu John W. Byers Andre Nathan

hostnet Raytheon BBN Google Boston University

GTER 49 | GTS 35 -- November 30th, 2020

Gatekeeper

Motivation -- Relevance of DDoS attacks

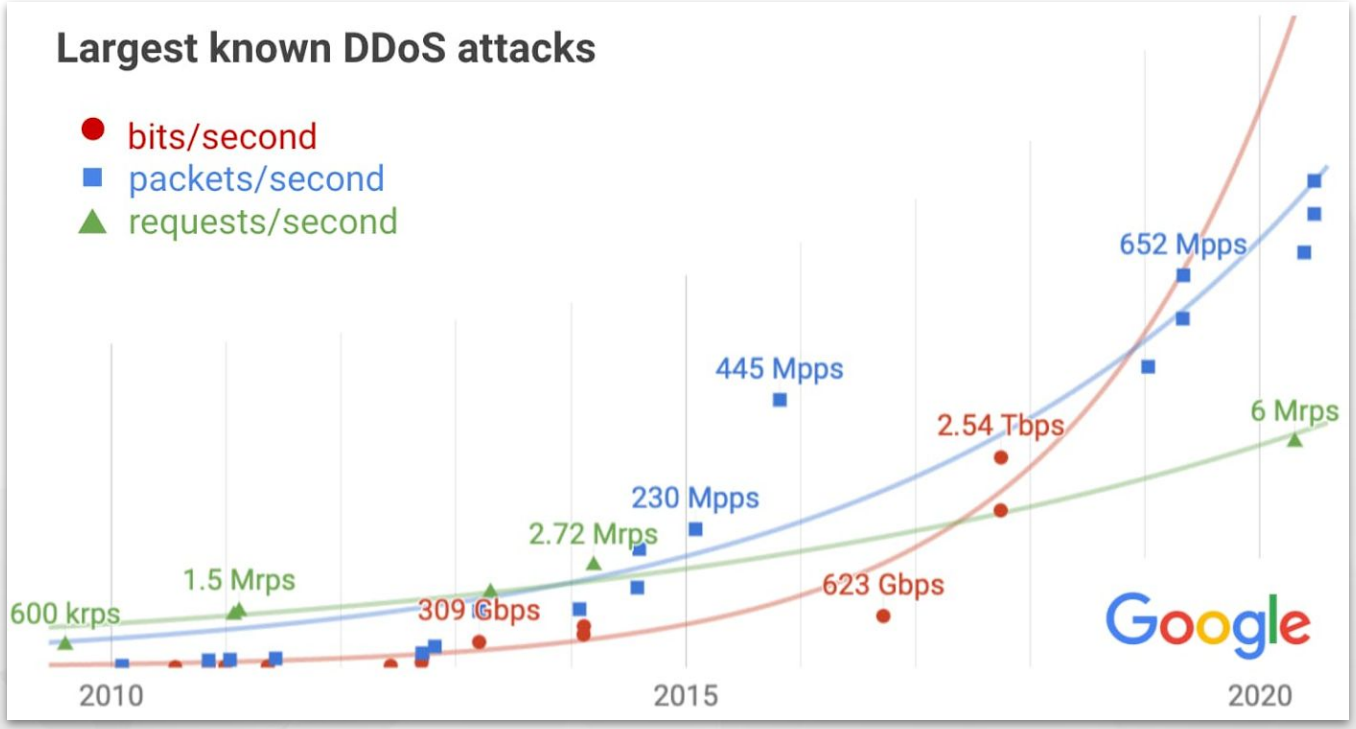

Gatekeeper

Motivation -- Largest DDoS attacks of 2020

Who	Peak	When
AWS	2.3 Tbps	February
Akamai	809 Mpps	June
Cloudflare	754 Mpps	June

809 Mpps is the newest packet-rate record

2.3 Tbps is close to the bandwidth record: 2.54 Tbps in Sep 2017

Gatekeeper

Motivation -- Why Gatekeeper?

Unparalleled multi-vector protection

⇒ All flows are monitored and all filters are active;
alternative solutions have limited filtering capacity;
See paper "The Catch-22 Attack" for details

Scalable

⇒ 1 Tbps deployment underway at Mail.ru

Mitigation in seconds

⇒ More than 80% of attacks last ≤ 4 min according to Kaspersky;
There is not much time for human intervention

Gatekeeper

Outline

- ✓ Motivation
- How Gatekeeper works
- How to write a destination policy
- Mitigating a SYN flood
- Conclusion

Gatekeeper

Gatekeeper's components

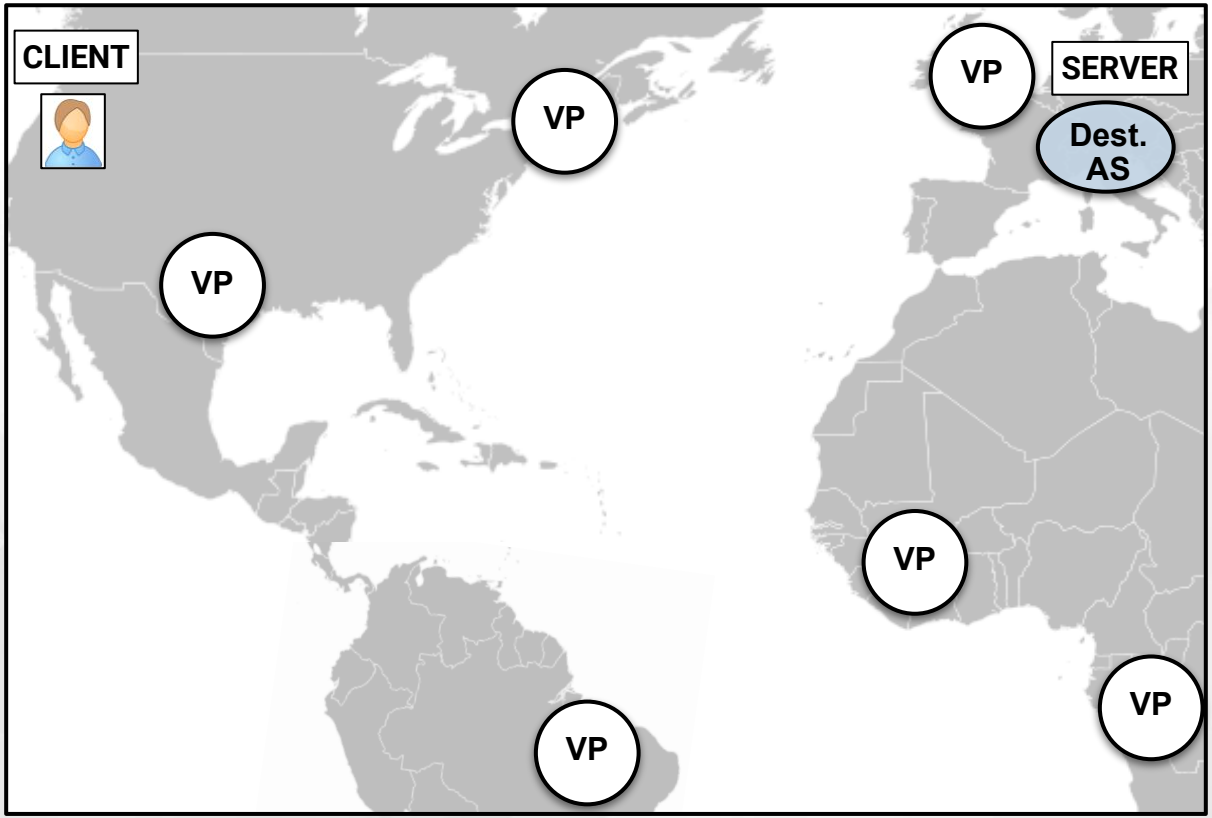

Vantage points: well-provisioned and geographically distributed locations

Gatekeeper

Gatekeeper's components

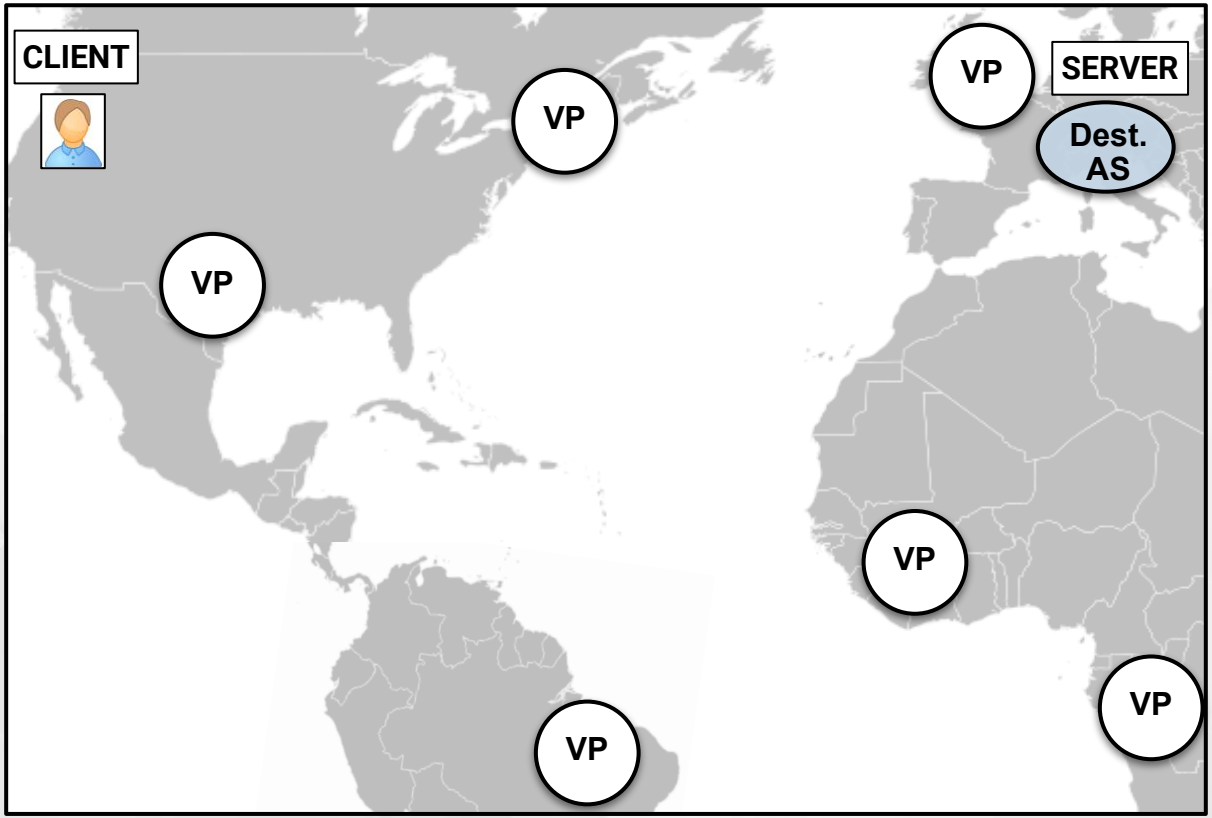

Vantage points: well-provisioned and geographically distributed locations

Requirements:

- computing capacity
- cheap ingress bandwidth
- BGP peering
- private links to the protected AS

Gatekeeper

Gatekeeper's components

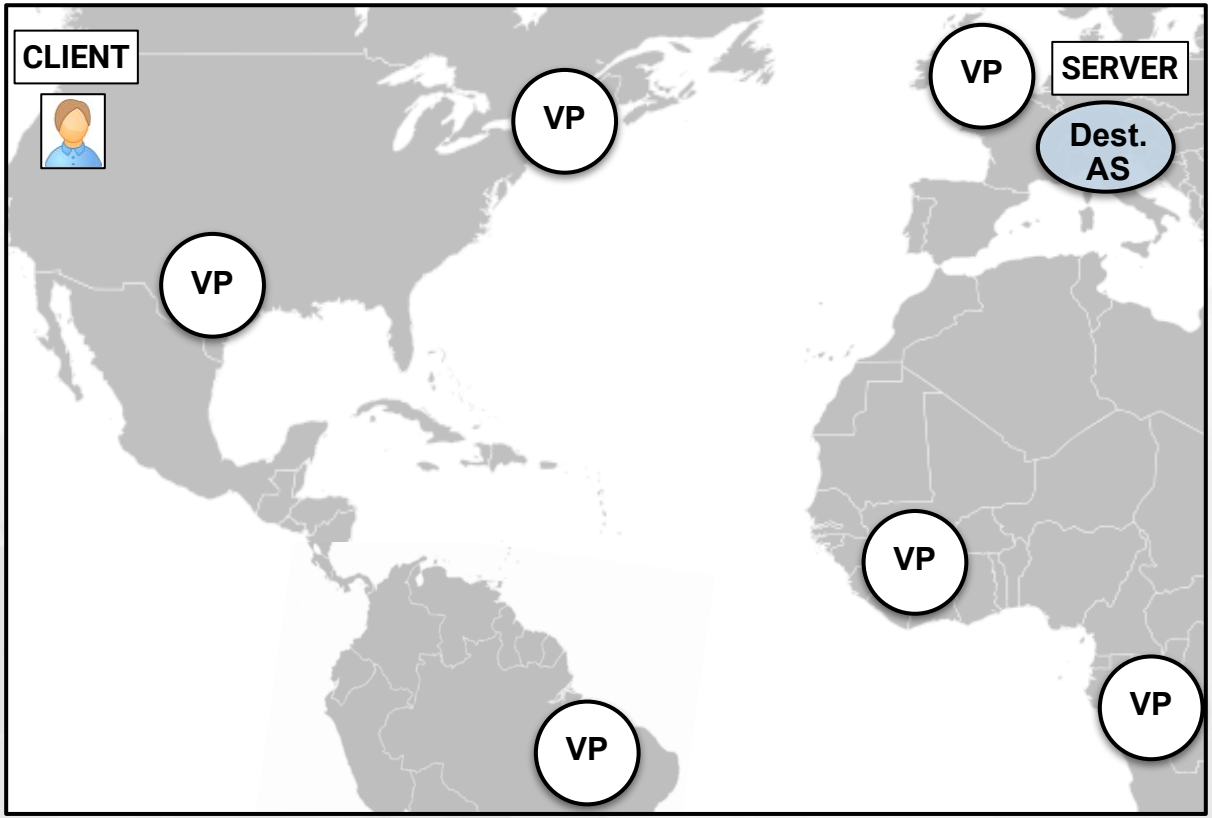

Vantage points:
well-provisioned and geographically distributed locations

Requirements:

- computing capacity
- cheap ingress bandwidth
- BGP peering
- private links to the protected AS

Examples:

- Internet exchanges
- Peering link
- Some cloud providers

Gatekeeper's components

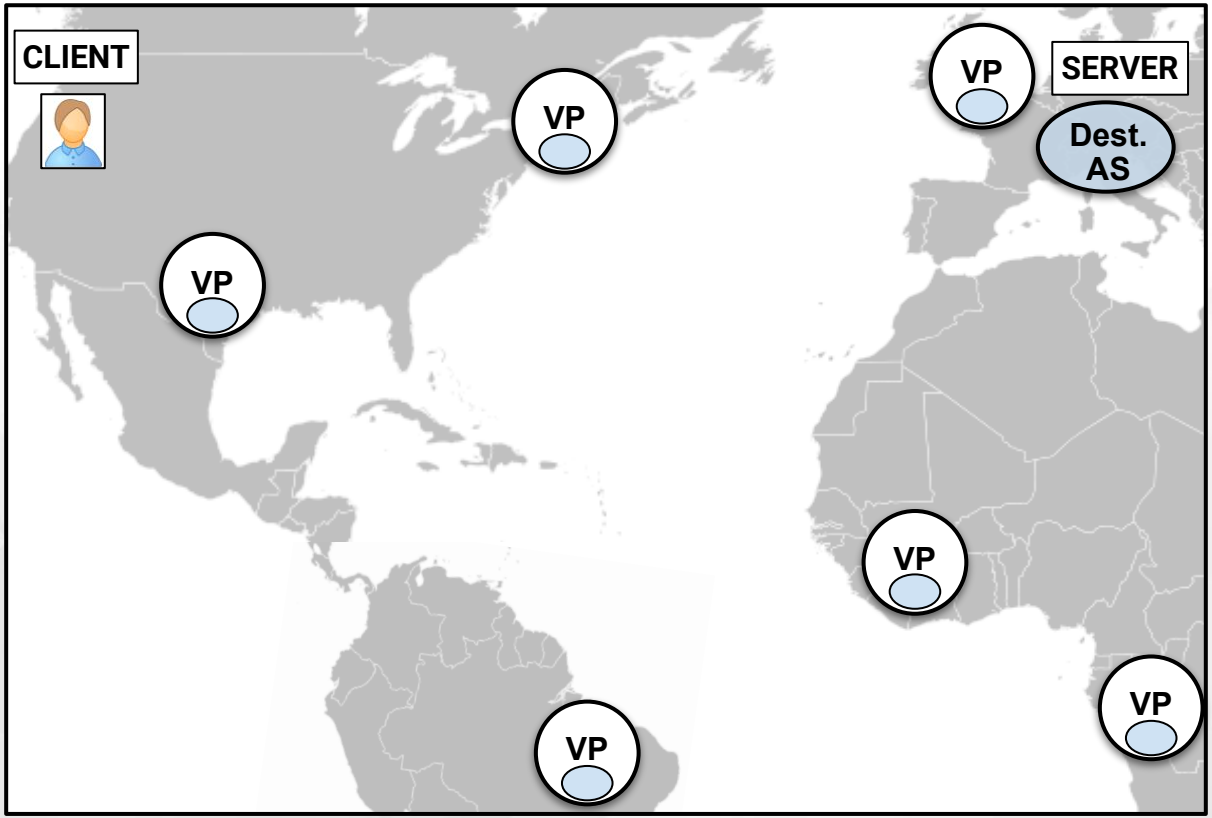

Vantage points: well-provisioned and geographically distributed locations

Requirements:

- computing capacity
- cheap ingress bandwidth
- BGP peering
- private links to the protected AS

Examples:

- Internet exchanges
- Peering link
- Some cloud providers

Gatekeeper's components

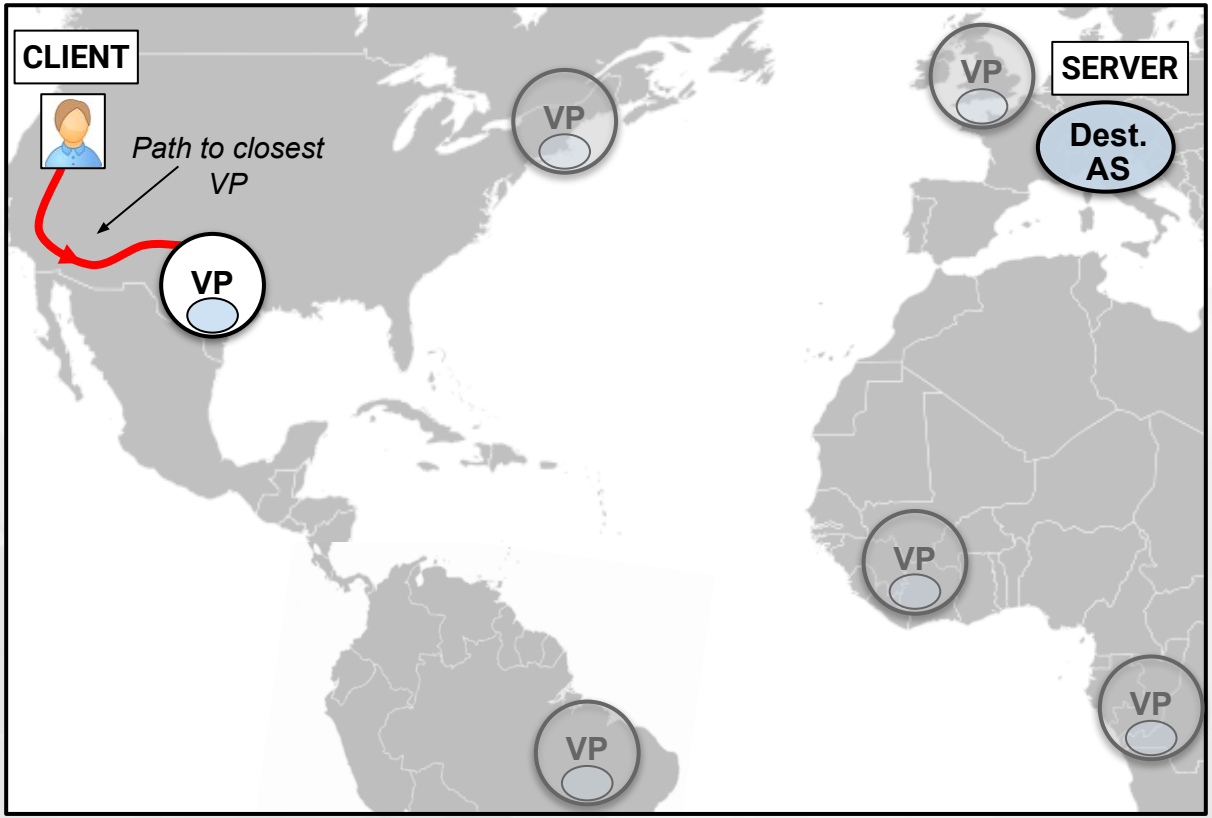

Vantage points: well-provisioned and geographically distributed locations

Requirements:

- computing capacity
- cheap ingress bandwidth
- BGP peering
- private links to the protected AS

Examples:

- Internet exchanges
- Peering link
- Some cloud providers

Gatekeeper

Gatekeeper's components

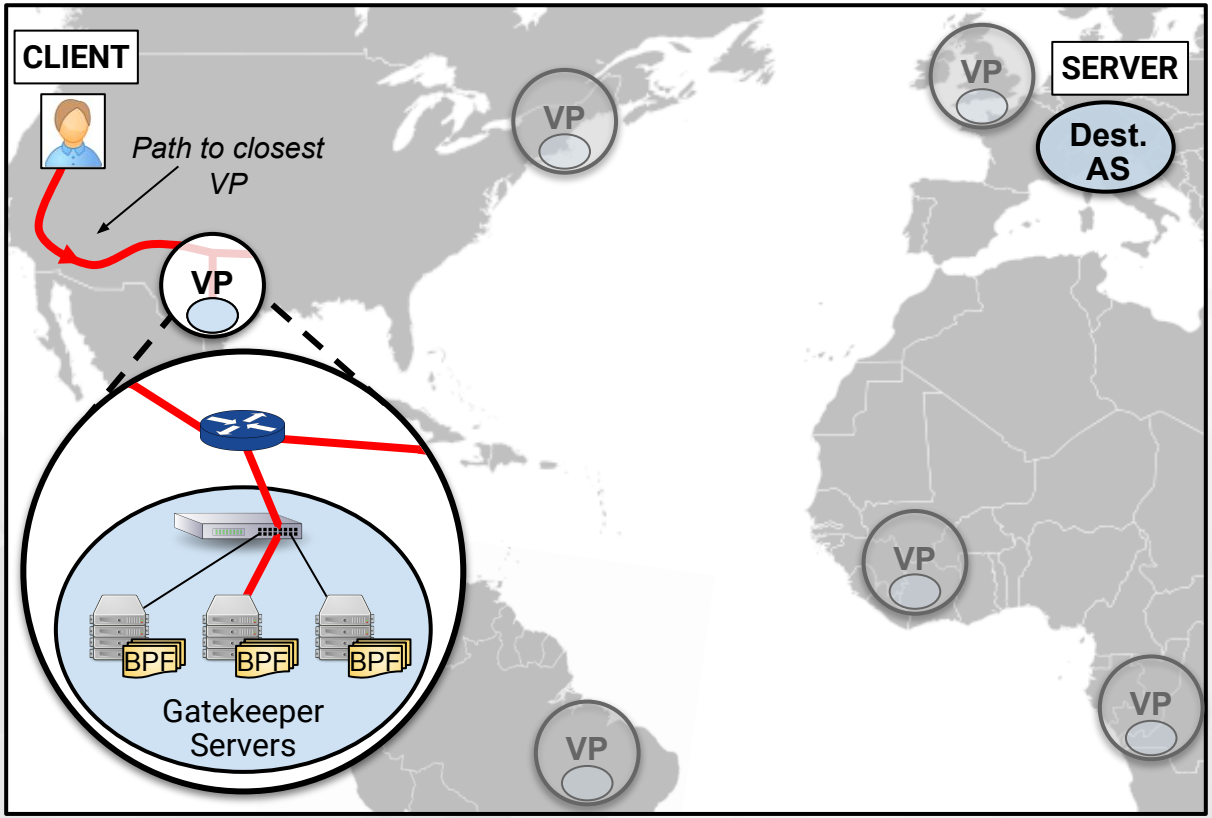

Gatekeeper servers: upstream policy enforcement

Responsibilities:

- Forwarding requests (new flows)
- Dropping or rate-limiting according to per-flow policy enforcement program
- Encapsulating

Gatekeeper's components

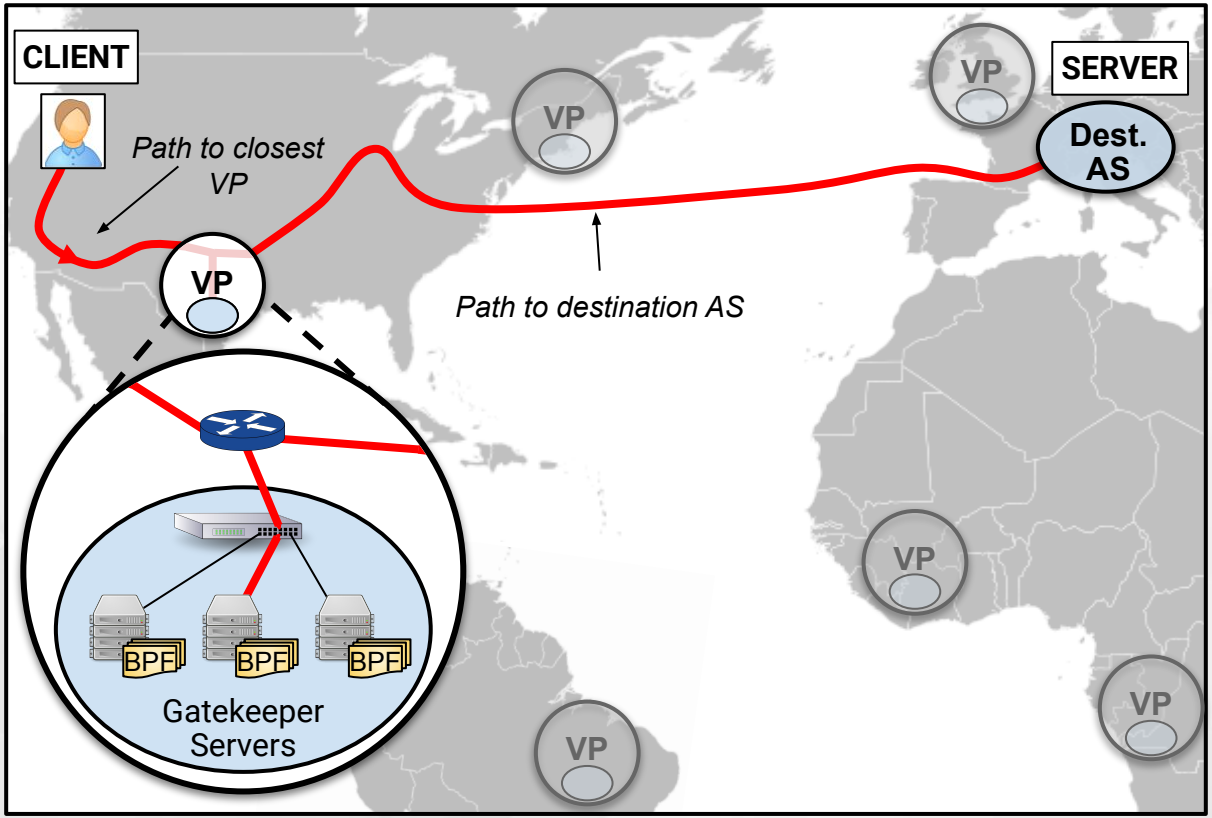

Gatekeeper servers: upstream policy enforcement

Responsibilities:

- Forwarding requests (new flows)
- Dropping or rate-limiting according to per-flow policy enforcement program
- Encapsulating

Gatekeeper's components

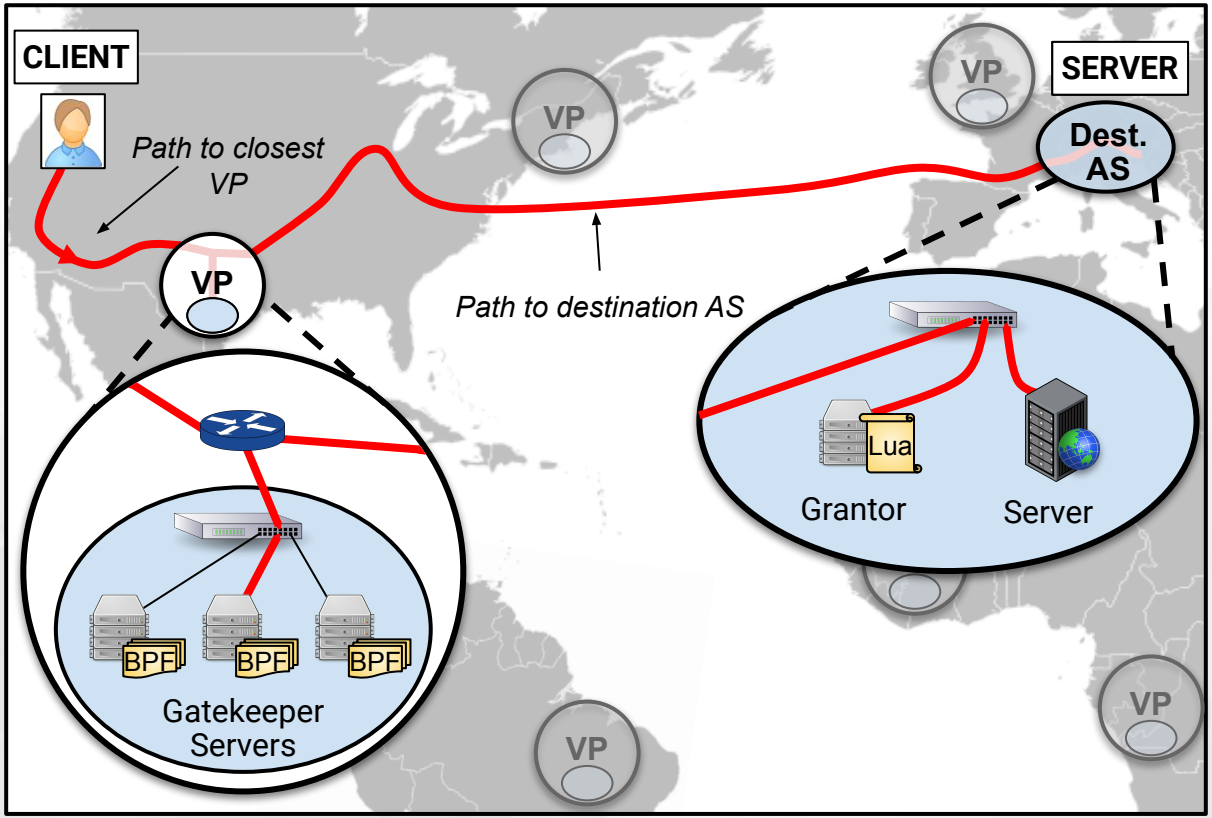

Grantor servers: centralized policy decision making

Responsibilities:

- Making policy decisions about requests and installing those decisions at Gatekeeper
- Decapsulating and sending to destination server

Gatekeeper

Quick summary

1. Packets from clients are forwarded to the closest VPs
2. Gatekeeper servers forward packets of new flows to Grantor servers, or run BPF programs to decide what to do
3. Grantor servers run a policy to map flows to BPF programs, and forward granted packets to destinations
4. Grantor servers notify Gatekeeper servers of all policy decisions
5. Gatekeeper servers enforce the police decisions

Gatekeeper

Quick summary

1. Packets from clients are forwarded to the closest VPs
2. Gatekeeper servers forward packets of new flows to Grantor servers, or run BPF programs to decide what to do
3. Grantor servers run a policy to map flows to BPF programs, and forward granted packets to destinations
4. Grantor servers notify Gatekeeper servers of all policy decisions
5. Gatekeeper servers enforce the police decisions

Gatekeeper

Outline

- ✓ Motivation
- ✓ How Gatekeeper works
- How to write a destination policy
- Mitigating a SYN flood
- Conclusion

Gatekeeper

Step 1: identify ALL your network profiles

A profile may apply:

to a single server, a group of servers, or

to blocks of IP addresses

Example of a profile: outgoing email servers

- No listening sockets
- Very small ingress traffic footprints

Sources: config files, production servers, docs

Gatekeeper

Step 2: write an BPF program for each profile

Classify packets into one of these bins:

- **Primary:** main purpose of the service
- **Secondary:** needed packets (e.g. TCP SYN, ICMP)
- **Unwanted:** please guess :-)

Enforce primary bandwidth limit before classification

Enforce secondary bandwidth limit after classification on secondary packets

Step 1: Network profiles ➔ Step 2: BPF programs ➔ Step 3: Lua Policy

Gatekeeper

Step 3: map flows to your BPF programs

Just classify flows using the destination IP address

Example: 10.99.99.128/25 are outgoing email servers

This information is a byproduct of Step 1

Grantor servers run this part of the policy (Lua policy)

Gatekeeper

Step 3: map flows to your BPF programs (bonus)

Classify source IP addresses too!

- Reject bogons, abusers, malware
- Tune bandwidth to partners, countries, end users
- Return different profiles to CDNs, crawlers, offices

Manage all your IP ranges with Drib:
https://github.com/andrenth/drib

Gatekeeper

Outline

- ✓ Motivation
- ✓ How Gatekeeper works
- ✓ How to write a destination policy
- Mitigating a SYN flood
- Conclusion

Gatekeeper

Mitigation example: beginning

A1. **Attack vector:** each source bursts SYN packets (secondary traffic) above the primary limit

A2. BPF programs are associated to these flows as described in section "How Gatekeeper works"

A3. Packets are limited to the secondary limit

Gatekeeper

Mitigation example: attack sent

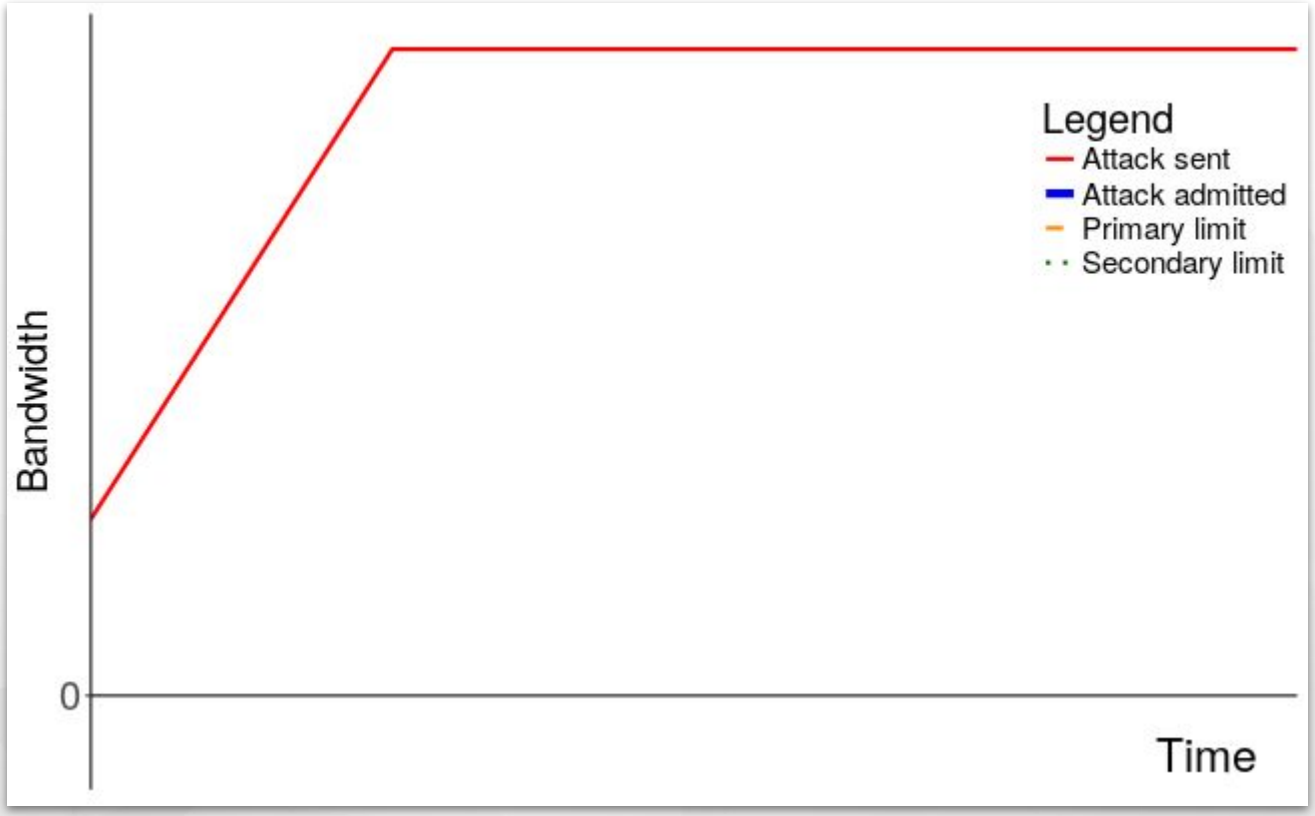

Gatekeeper

Mitigation example: secondary limit

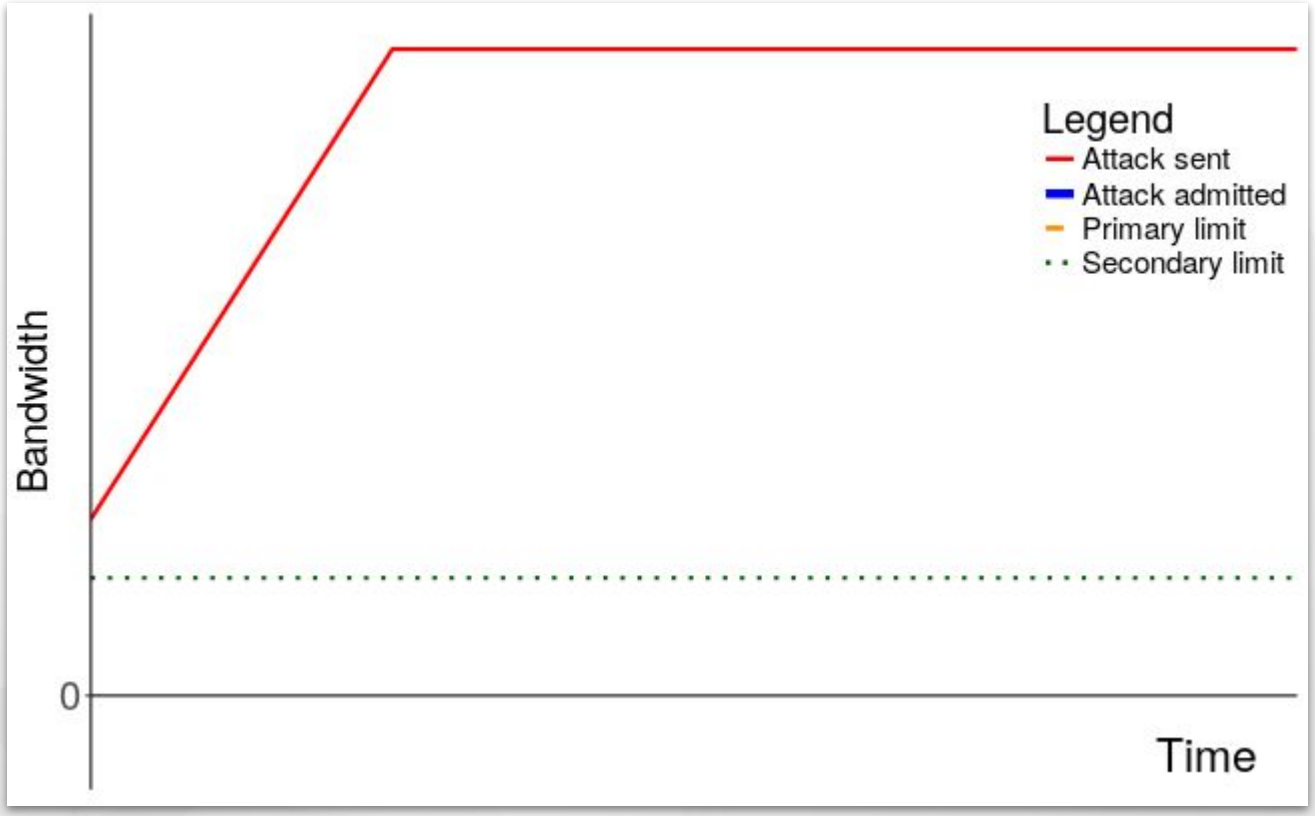

Gatekeeper

Mitigation example: primary limit

B1. The primary limit decreases when
the attack traffic > the maximum of the primary limit

B2. When a primary limit < the secondary limit
⇒ more packets are dropped

B3. When a primary limit < 0
⇒ all packets of the flow are dropped

Gatekeeper

Mitigation example: primary limit

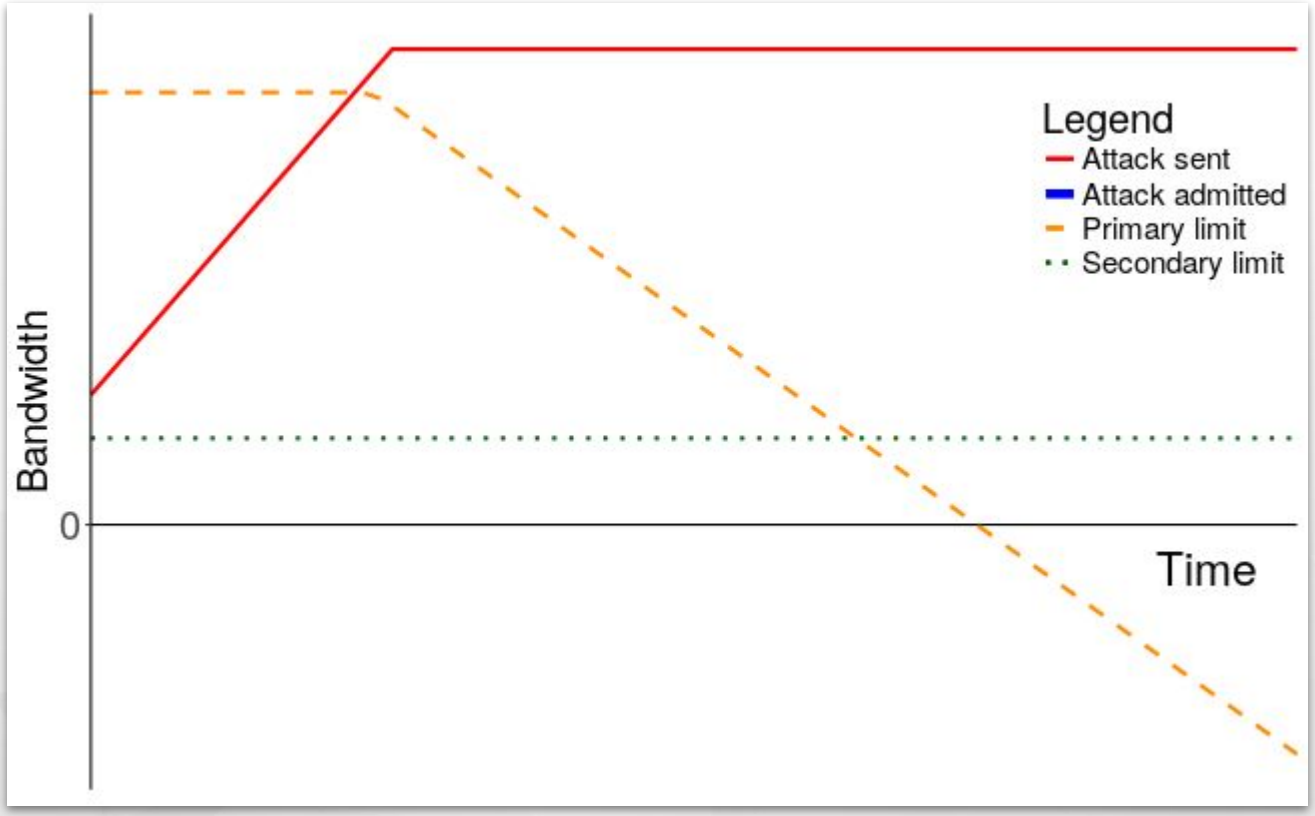

Gatekeeper

Mitigation example: attack admitted

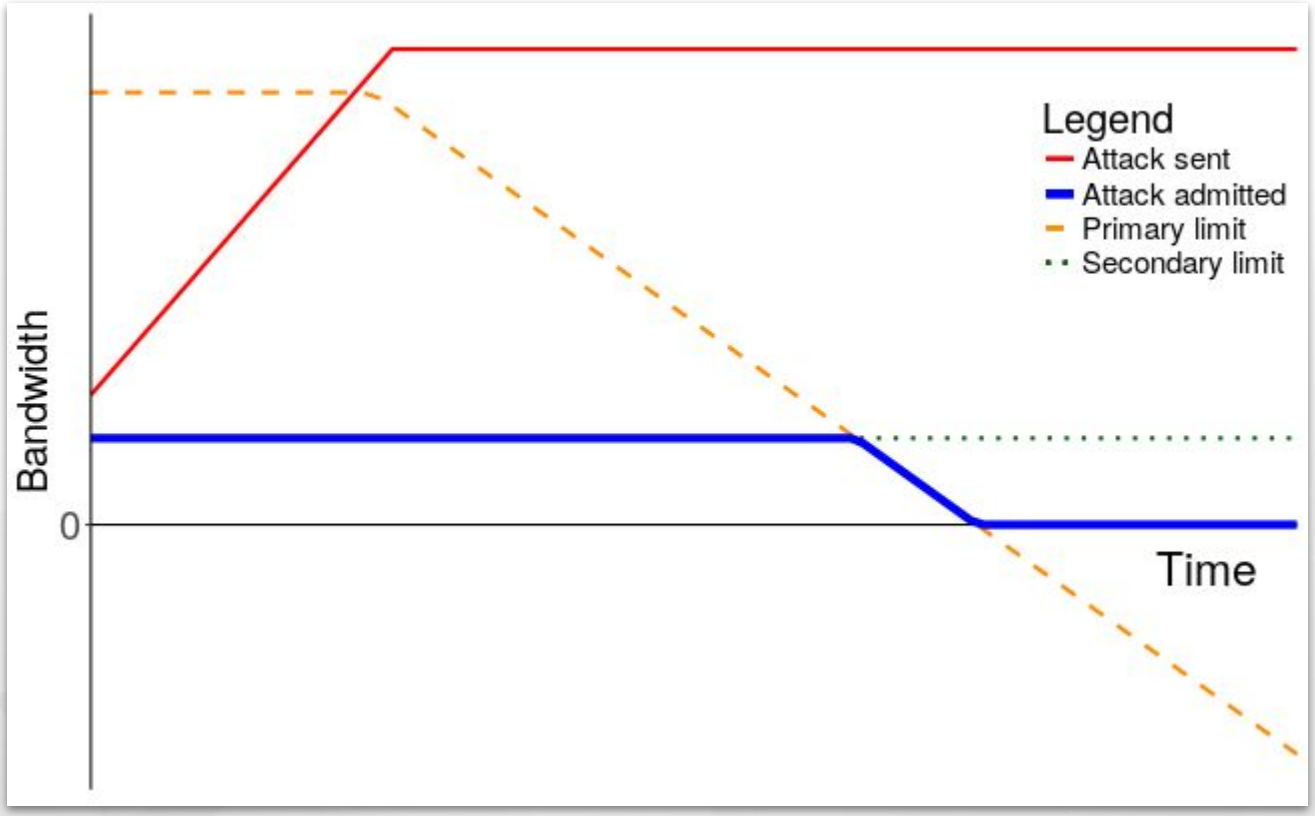

Gatekeeper

Mitigation example: punishment

C1. Once a flow stops the attack, its packets will be dropped until the primary limit goes positive
⇒ Punishment is proportional to the offense

C2. **OR** the flow entry expires
⇒ Whichever is shorter

Gatekeeper

Outline

- ✓ Motivation
- ✓ How Gatekeeper works
- ✓ How to write a destination policy
- ✓ Mitigating a SYN flood
- Conclusion

Gatekeeper

Limitations

The more sophisticated attacks become,
the less effective a policy can be

Policy complements:

1. Distributed database (e.g. anti-spoofing)
2. The protected applications (e.g. Cerber Security)
3. Intrusion detection systems (e.g. Suricata)
4. … (where does it stop?)

Gatekeeper

Limitations -- the endgame for attackers

The endgame is when

the cost to identify attack traffic is ≥

the cost to serve it

Typical example nowadays: DNS queries over UDP

The best action is to serve as many users as possible

Flow orchestration is the last resort

Gatekeeper

Some future work

Supporting 100 Gbps NICs at line speed

⇒ Cheaper deployments

Supporting load balancing in policies

⇒ Better return on investment

Flow orchestration

⇒ Insurance for endgame

Gatekeeper

Wrap-up

Unparalleled multi-vector protection

Mitigation in seconds

Scalable, open source, and ready for deployment

Impactful features in store for the future

Gatekeeper

https://github.com/AltraMayor/gatekeeper