Next Gen Blackholing to Counter DDoS

NANOG75, San Francisco

Christoph Dietzel §*, Matthias Wichtlhuber*, Georgios Smaragdakis §, Anja Feldmann †

§TU Berlin, *DE-CIX, †MPI

DE-CIX

Where networks meet

www.de-cix.net

Volumetric DDoS Attacks

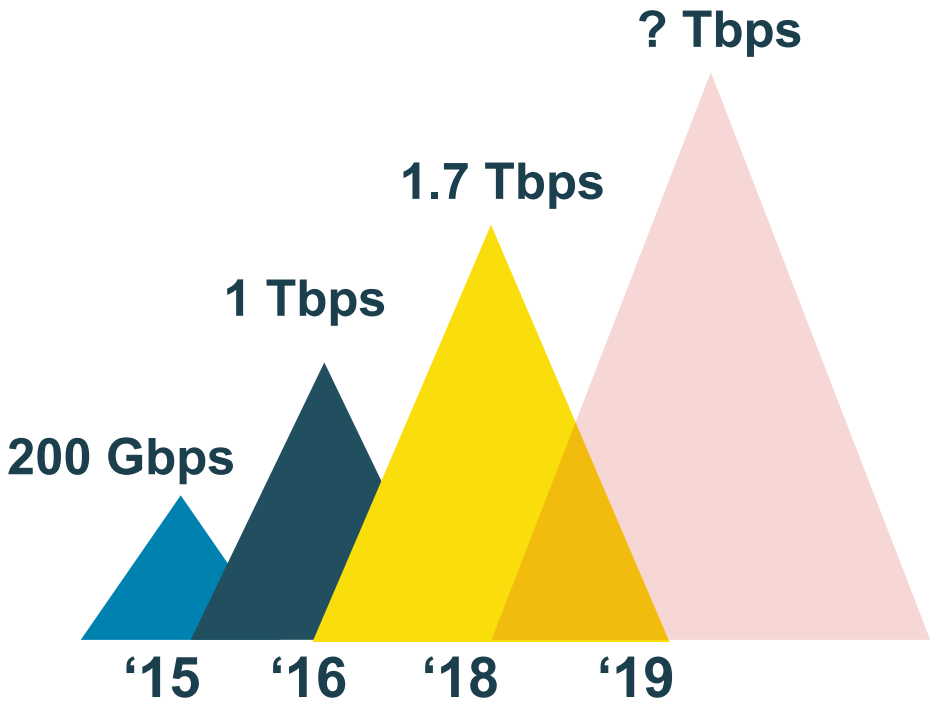

NETSCOUT

Attack Map | Archives | About

BLOG HOME | CORPORATE SITE | RSS

NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us

Carlos Morales on March 5, 2018.

A Frightening New Kind Of DDoS Attack Is Breaking Records

Lee Mathews Contributor ⓘ
Security
Observing, pondering, and writing about tech. Generally in that order.

Back in October of 2016, a denial-of-service attack against a service provider called Dyn crippled Americans' Internet access on the east coast. Its servers were bombarded with a jaw-dropping amount of traffic. Some estimates believed the data rate of the attack peaked at around 1.2Tbps, which was unheard of at the time.

ISP DDoS Defense Toolbox

- Filters at arbitrary granularity
- Vendor-specific
- Per device config

- Carefree service
- Redirects traffic to scrubbing centers
- On-demand vs. always on

- Configures rules at neighbor network
- Filters at arbitrary granularity
- Cooperation required

- Configures rules at neighbor network
- Filters at IP granularity
- Cooperation required

DDoS Defense at IXPs

→ Combine good properties of existing solutions

→ Eradicate current shortcomings

\+ IXPs offer services to hundreds of Ases

\+ IXPs have multiple Tbps capacity

\+ Trusted part of the Internet community

Blackholing at IXPs

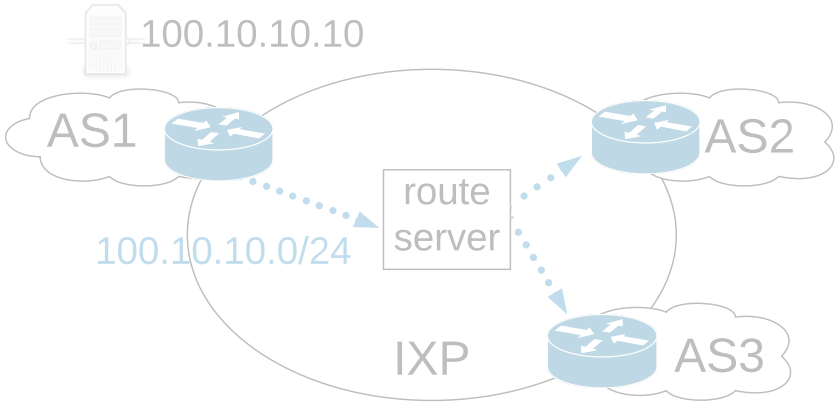

control plane

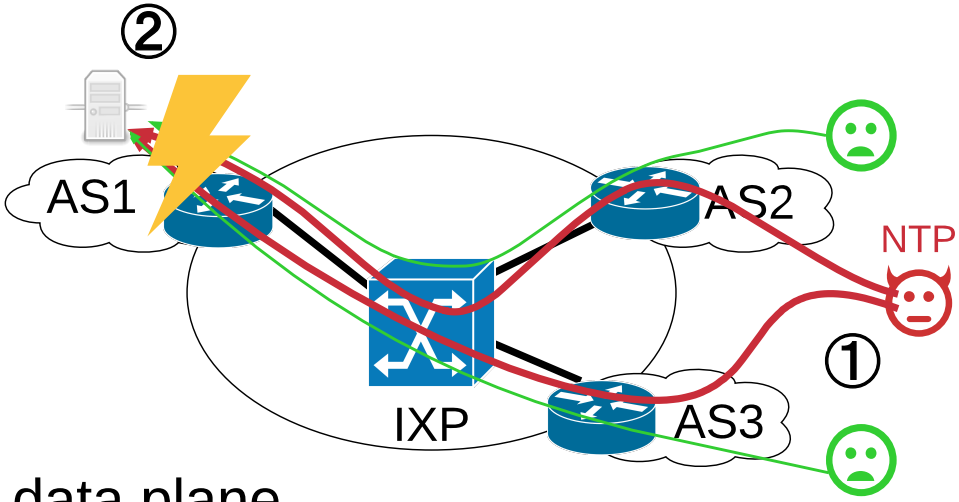

data plane

Blackholing at IXPs

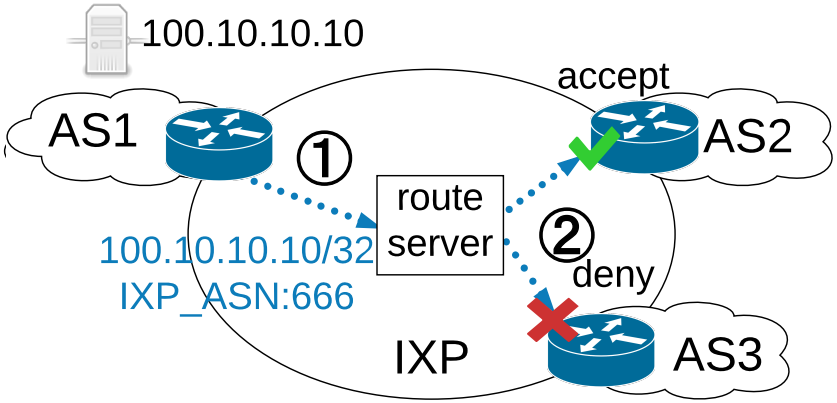

control plane

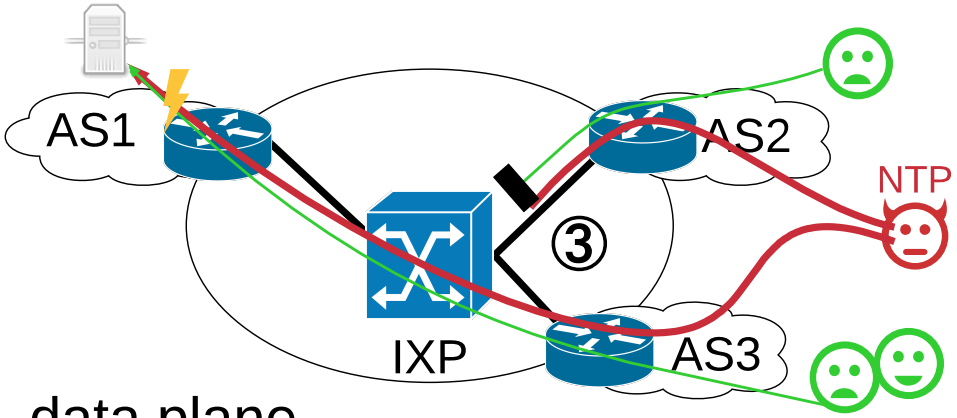

data plane

Blackholing – Limitations

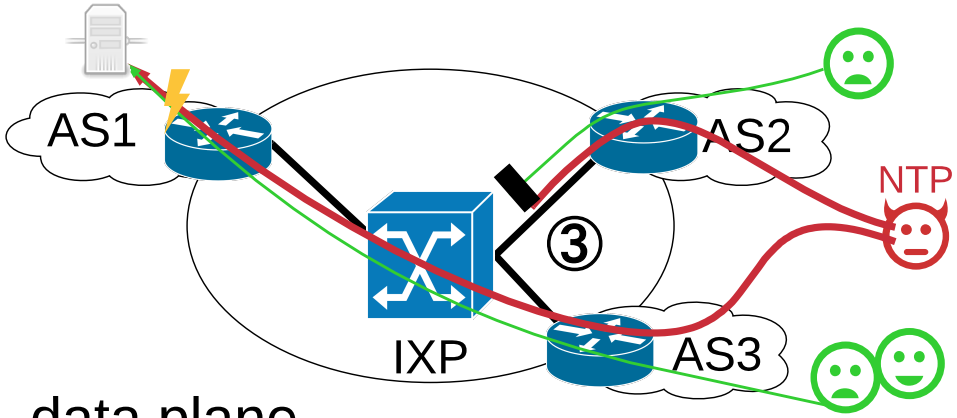

data plane

→ Blocks unwanted and wanted traffic

→ Hard to predict behavior

→ No effect on a subset of peerings

Blackholing – Limitations

- → Relative traffic of 40GE IXP port
- → Mostly web traffic (80, 443, …)
- → Attack 70% memcached traffic
- → Still significant share of web traffic
- → **Collateral damage!**

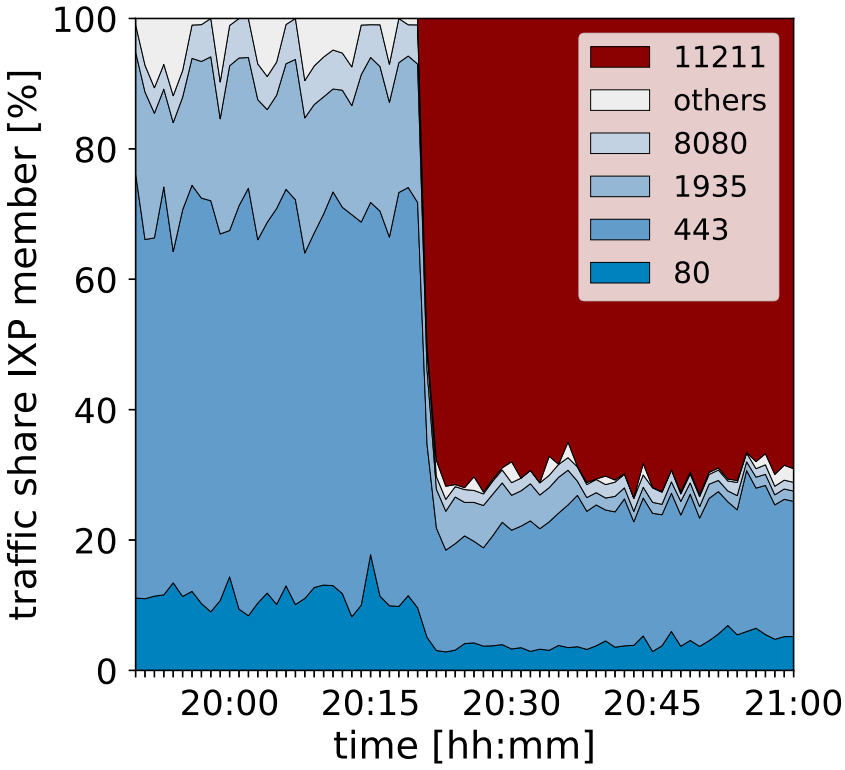

Blackholing – Limitations

- → All or nothing approach
  - → Prefix granularity
  - → Per peer selection at IXPs

- → Blackholing traffic:
  - → 99.94% UDP
  - → Expected L4 ports (NTP, LDAP, …)

- → **More granularity needed!**

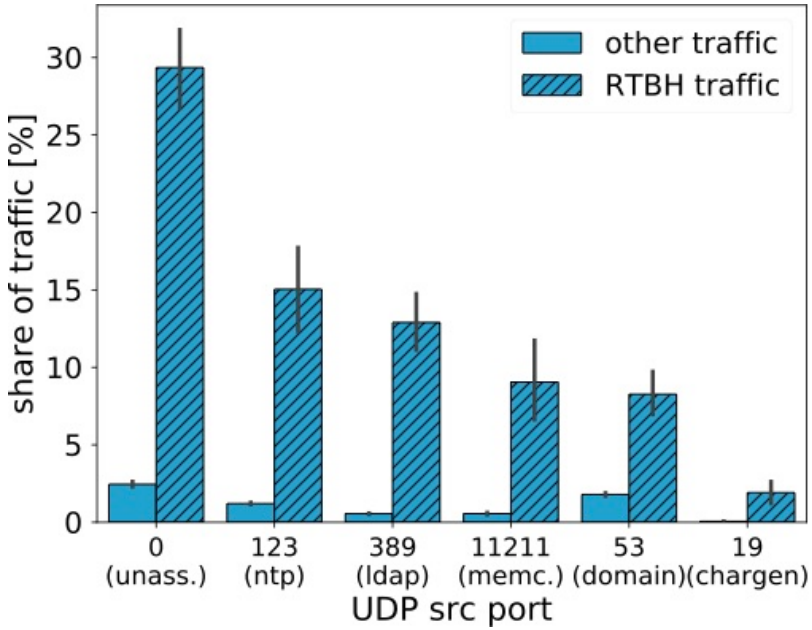

Blackholing – Limitations

- → How “ineffective“ can it be?
  - → NTP DDoS attack
  - → AS at IXP via ML peering
  - → Attacks for 10 min to /32

- → Drop all traffic to /32
- → Traffic: 800 to 600 Mbps
- → Peers: 38 to 26
- → **Signaling too complex!**

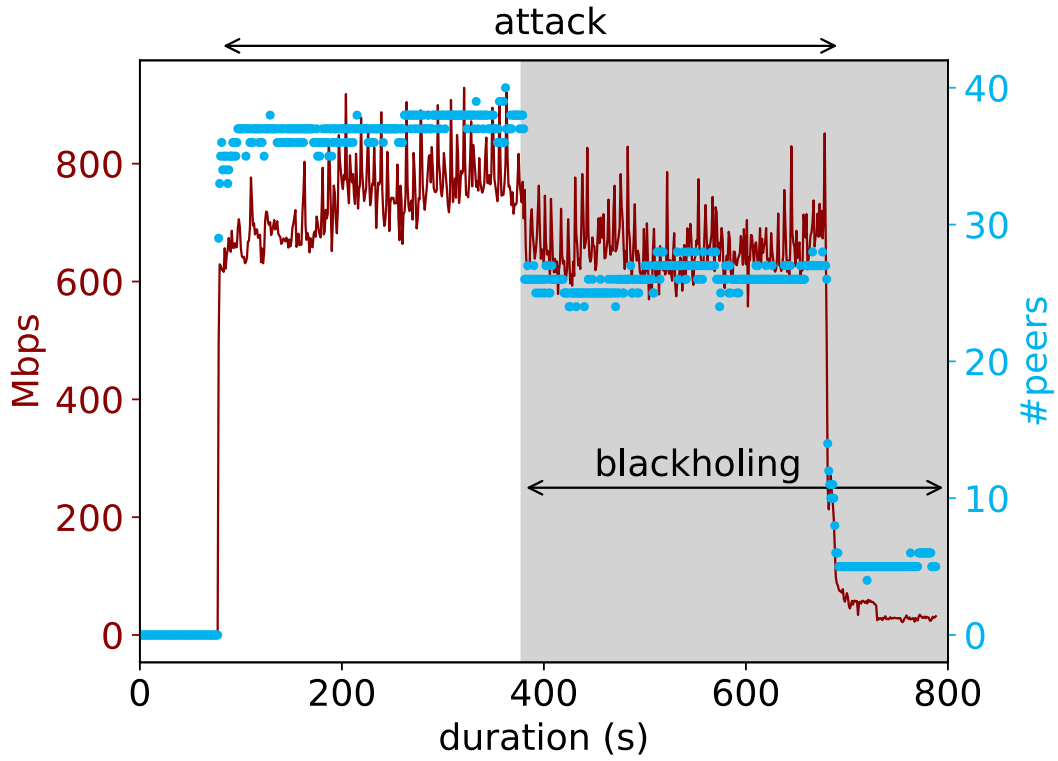

Advanced Blackholing Requirements

- → Granularity
  - → Fine-grained filtering (src/dst header fields)
- → Signaling complexity
  - → Easy to use, short setup time
- → Cooperation
  - → Lower levels of cooperation among the involved parties
- →Telemetry
  - → Feedback on the state of the attack at any time
- → Scalability
  - → Scale in terms of performance, filters, reaction time, config complexity
- → Cost
  - → Meeting all requirements with min. invest (CAPEX & OPEX)

Advanced Blackholing System

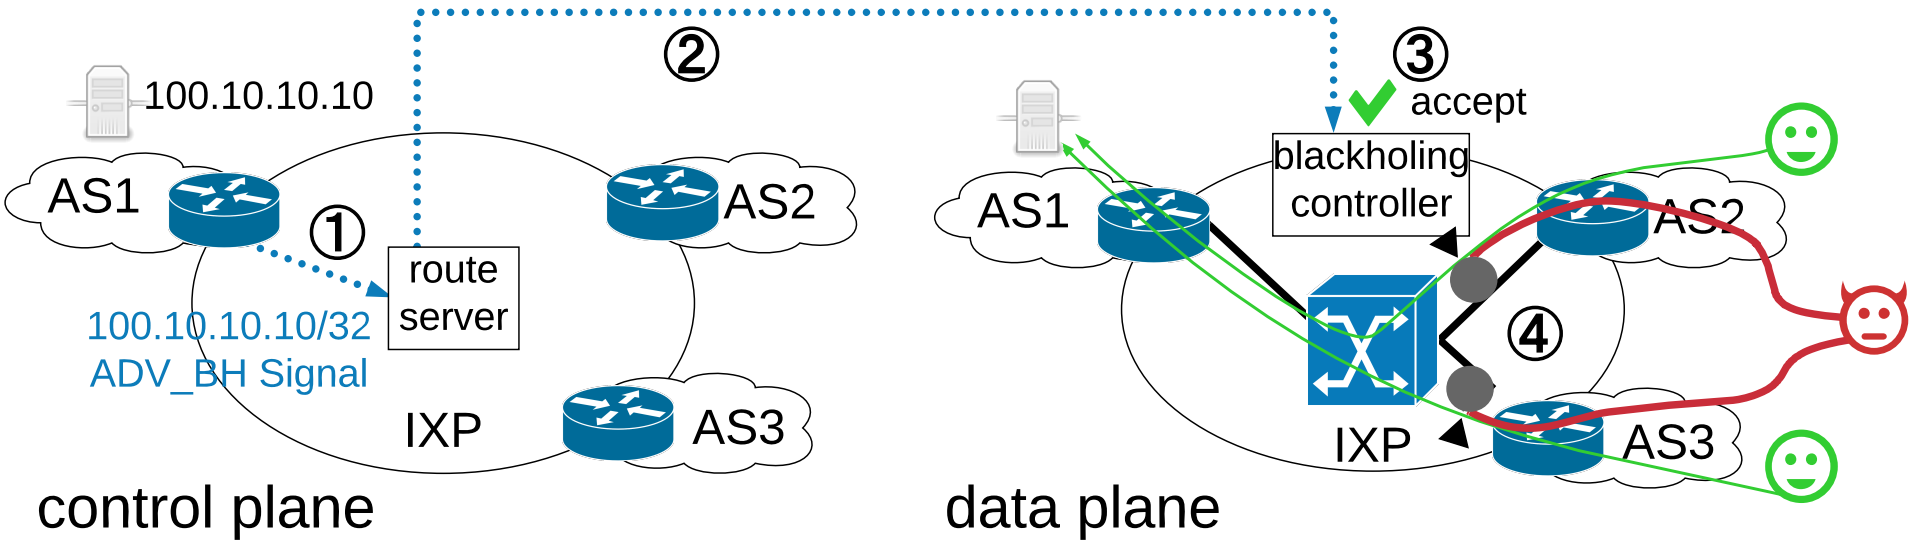

Advanced Blackholing System

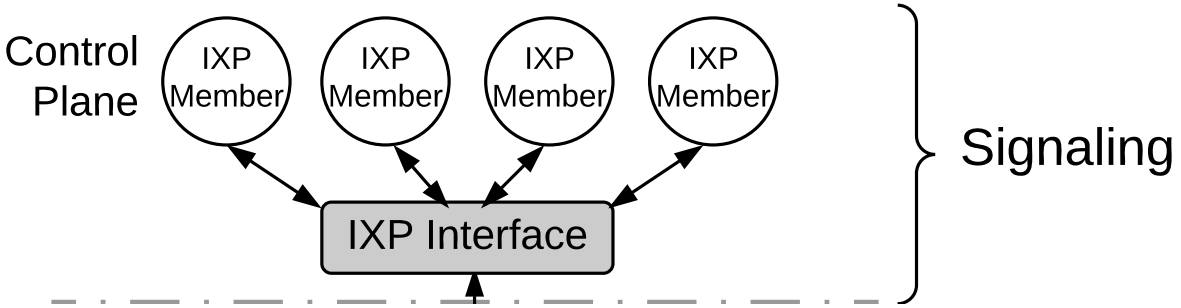

Advanced Blackholing System

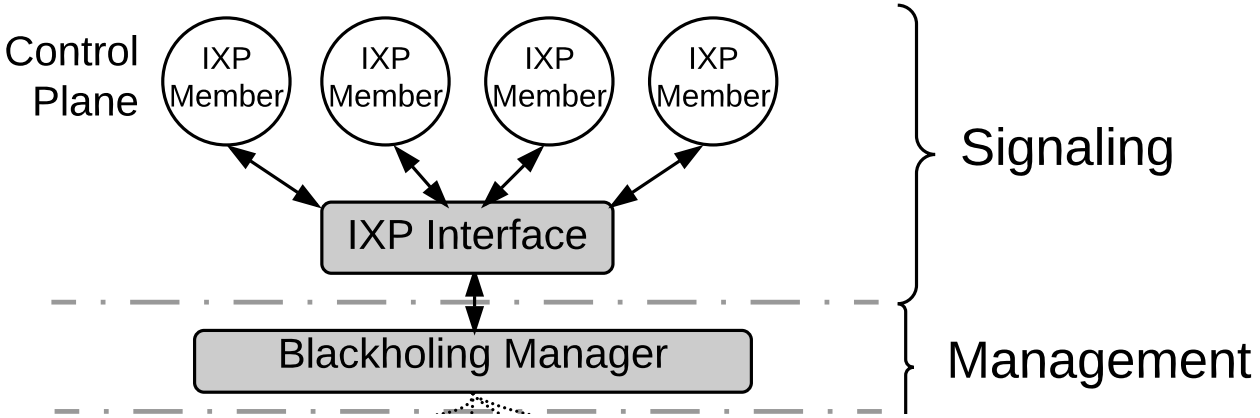

Advanced Blackholing System

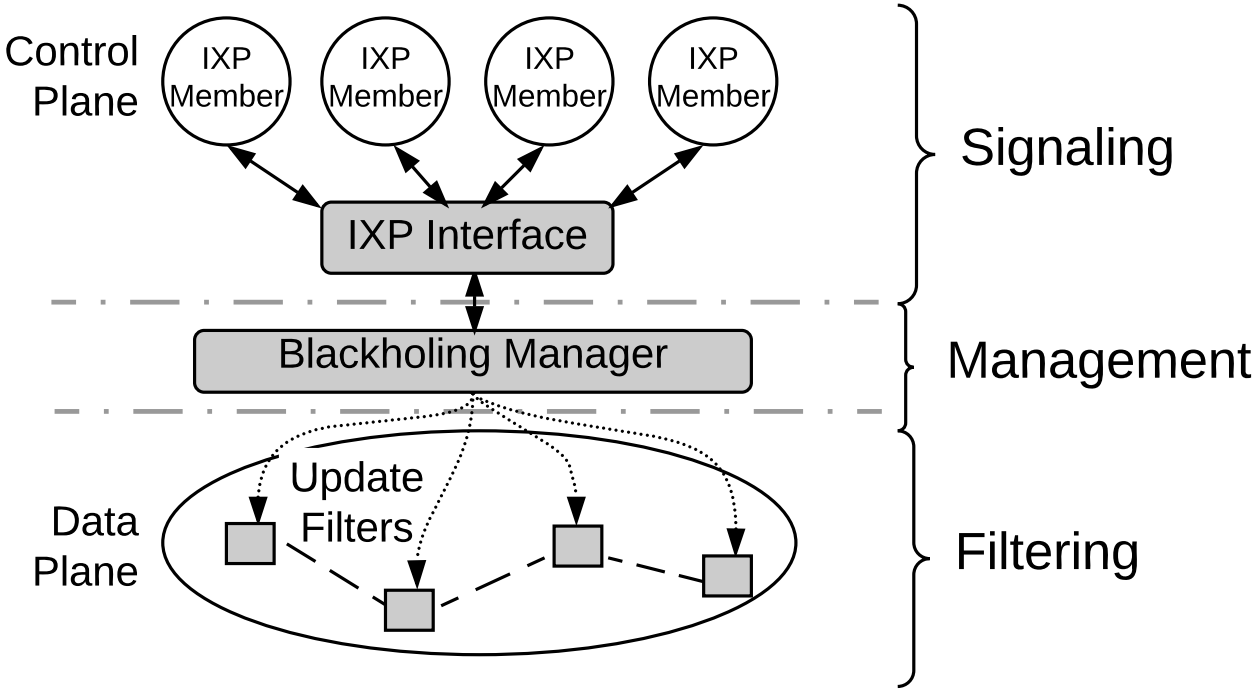

Advanced Blackholing Signaling (BGP part)

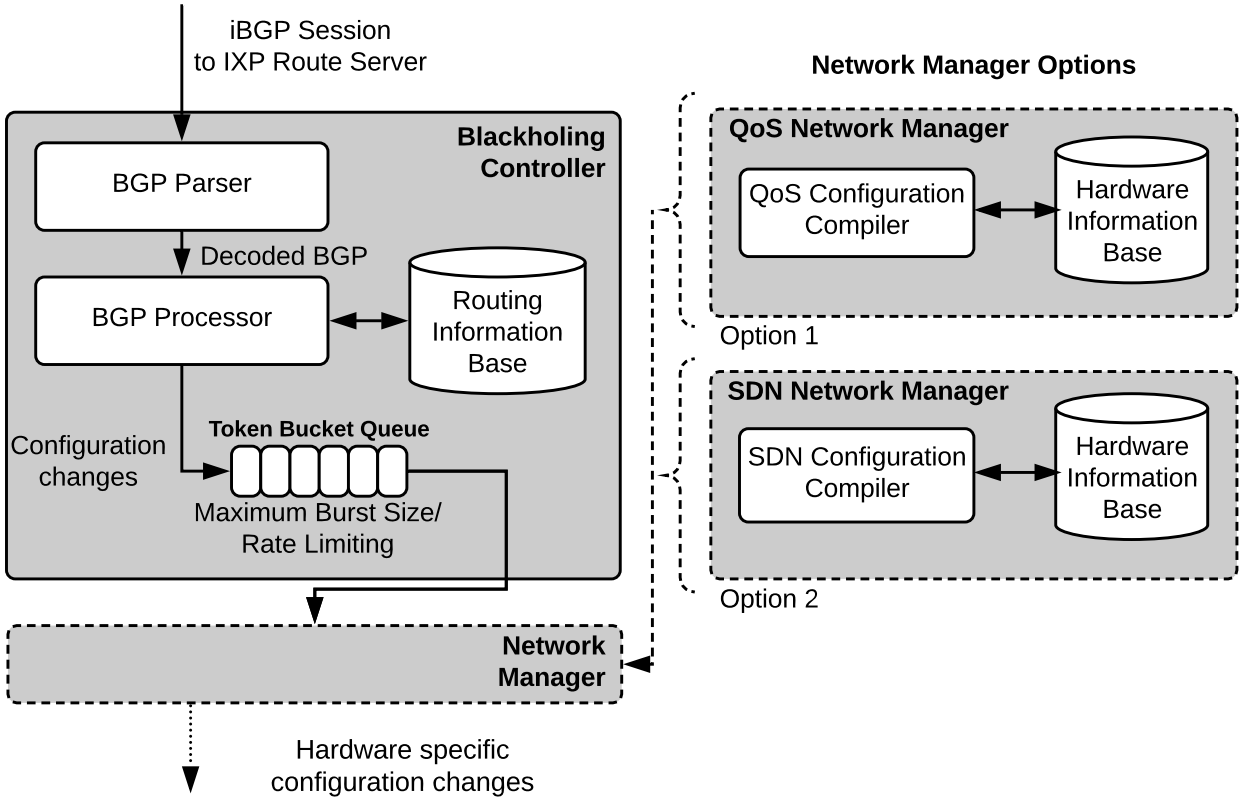

Building Blocks

- → Granularity
  - UDP, TCP, Ports, …
- → Signaling complexity
  - BGP communities or API
- → Cooperation
  - → Enforced by IXP
- → Telemetry
  - Monitoring with statistics
- → Scalability
  - Line-rate in hardware
- → Cost
  - Implemented in existing hardware

Implementation Challenges

- → BGP processing
- → Configuration proxy
- → Why not FlowSpec?

Does it Scale?

- → Scalability wrt. number of filters & IXP ports (of switches/routers)
  - → TCAM to match header fields
  - → Measuring system's limits & port's limits (max no. of filters)
  - → Results on next slide

- → Scalability wrt. configuration update frequency limits (of config proxy)
  - → Allows 4.33 filter updates per second
  - → 70% of BH updates below 1 second

Stress Test on the IXP‘s Hardware

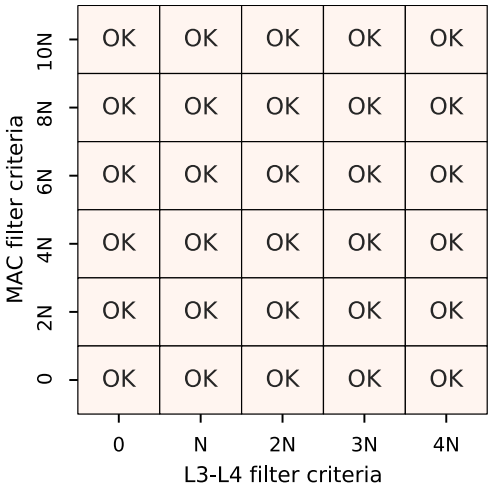

20% of IXP member ASes

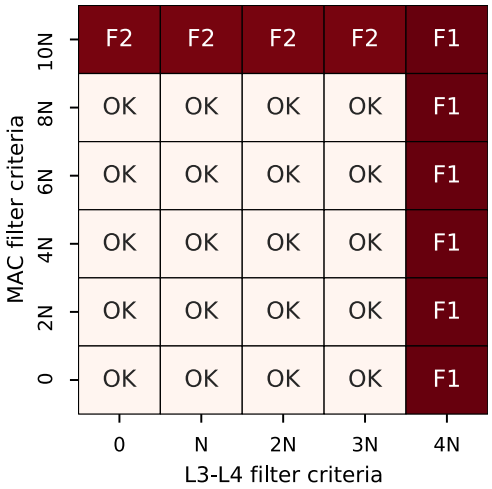

60% of IXP member ASes

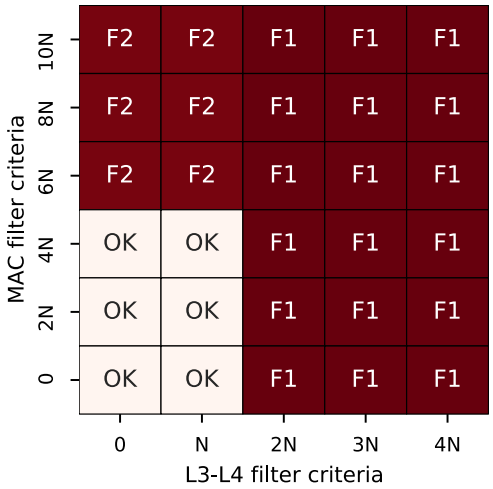

100% of IXP member ASes

Measurement Experiment

- How “effective“ is it
  - NTP DDoS attack
  - AS at IXP via ML peering
  - Attacks for 10 min to /32

- Drop / shape UDP NTP
- Traffic: 1000 to 200 to 0 Mbps
- Peers: 60 to (almost) 0

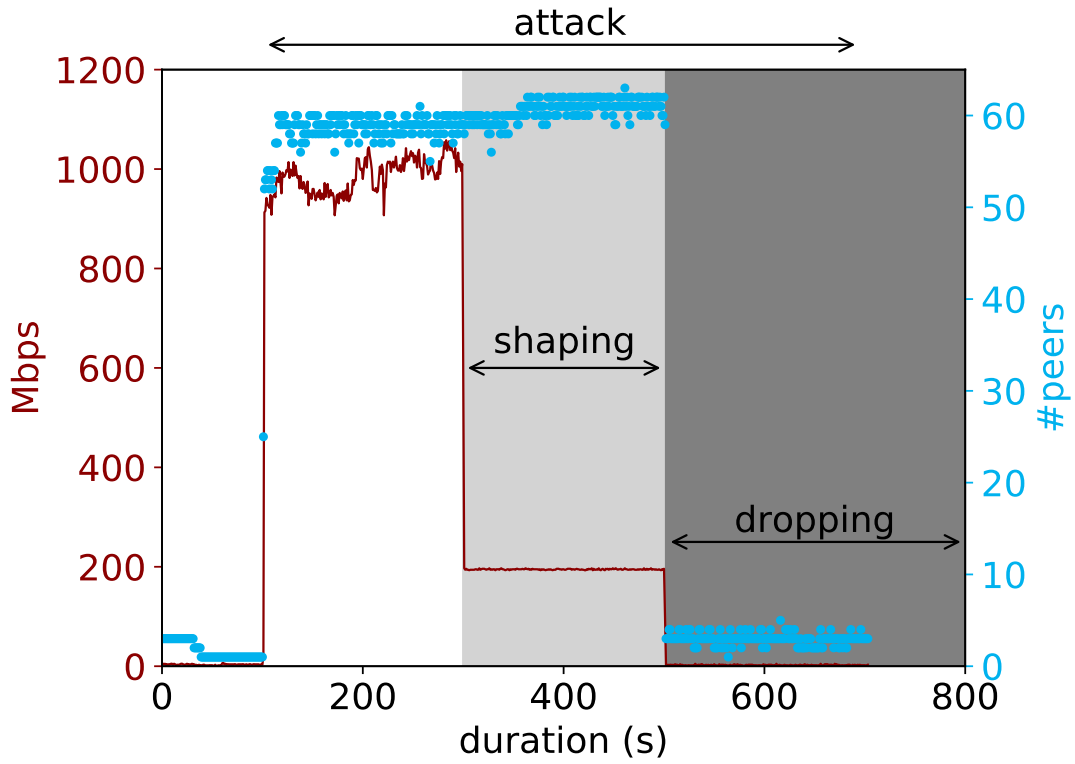

Summary

→ A number of DDoS mitigation solutions exist, but …

→ We identify and measure Blackholing limitations

→ We propose Advanced Blackholing, combining the benefits and overcome problems of today's DDoS defense

→ We implement a new system with a BGP and API interface

→ We evaluated and proved good scales scaling