First CONCORDIA Workshop on Collaborative DDoS Mitigation
15 September 2022

Anycast Agility: Network Playbooks to Fight DDoS

ASM Rizvi*
USC/ISI

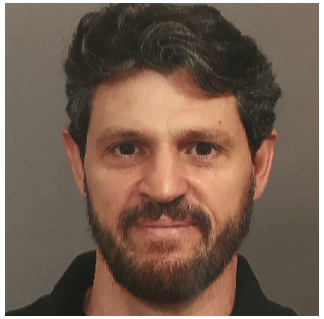

Leandro Bertholdo*
University of Twente

Joao M Ceron
SIDN Labs

John Heidemann
USC/ISI

* Indicates equal contribution

CONCORDIA Poster 2019

Planning for Anycast as Anti-DDoS

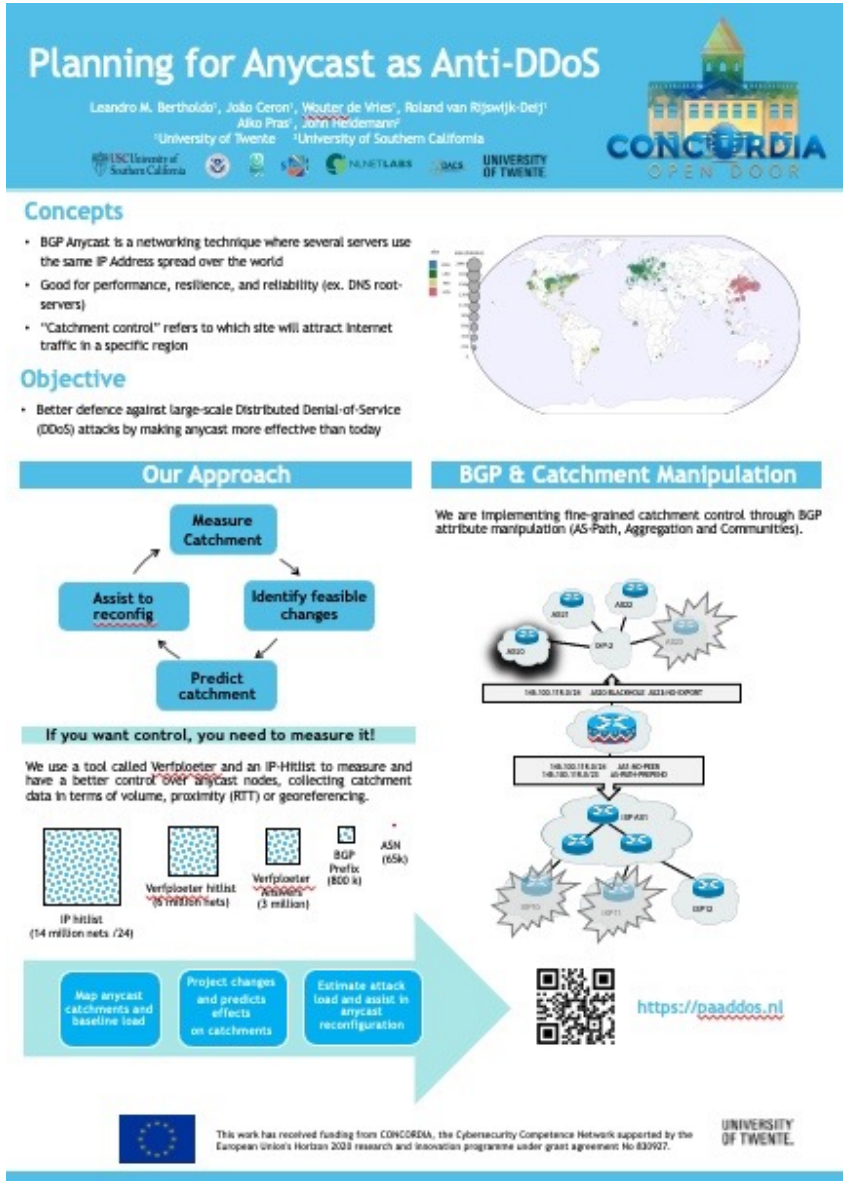

Distributed Denial-of-Service (DDoS) is Bad... and Getting Worse

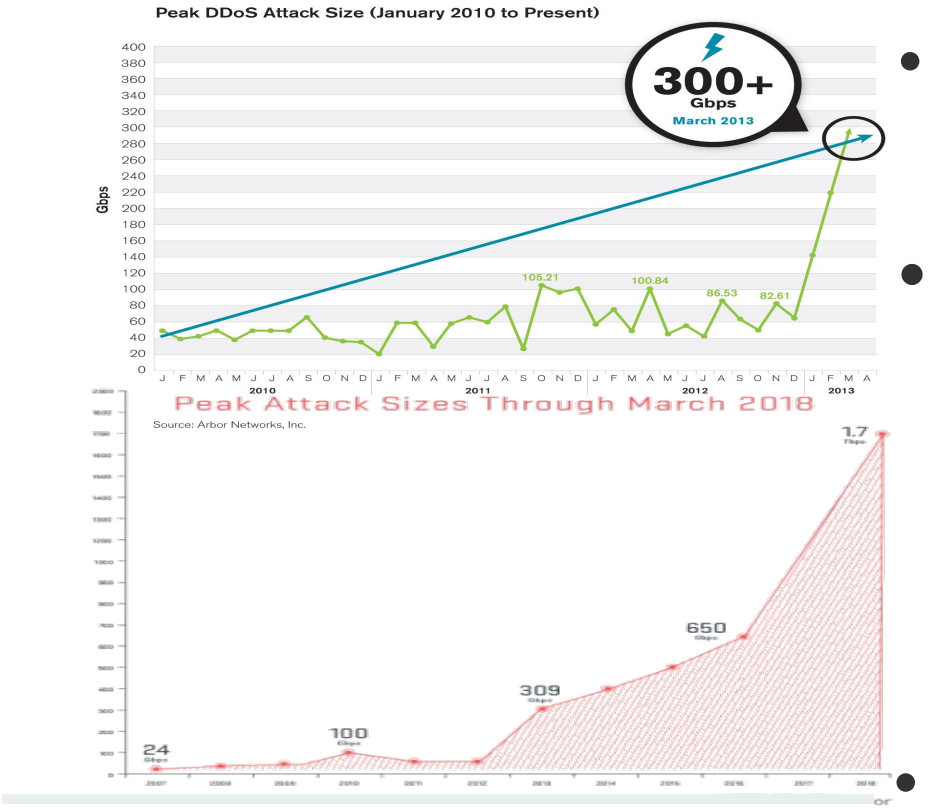

AWS said it mitigated a 2.3 Tbps DDoS attack, the largest ever

The previous record for the largest DDoS attack ever recorded was of 1.7 Tbps, recorded in March 2018.

- DDoS is big
  - Botnets
- DDoS is getting bigger
  - Github 1.35Tbps → Amazon gets 2.3 Tbps
  - **IoT & CPE devices**
  - **Reflection attacks from Cloudproviders**
- DDoS-as-a-service is cheap
  - starting at $1/attack [Santanna et al, 2015]

Why anycast?
Where do you use anycast in your daily life?

NETFLIX

CLOUDFLARE

Google™ DNS
8.8.8.8
8.8.4.4

aws

Microsoft Azure

NETSCOUT

DNS case study:
Where are universities hosting their DNS ?

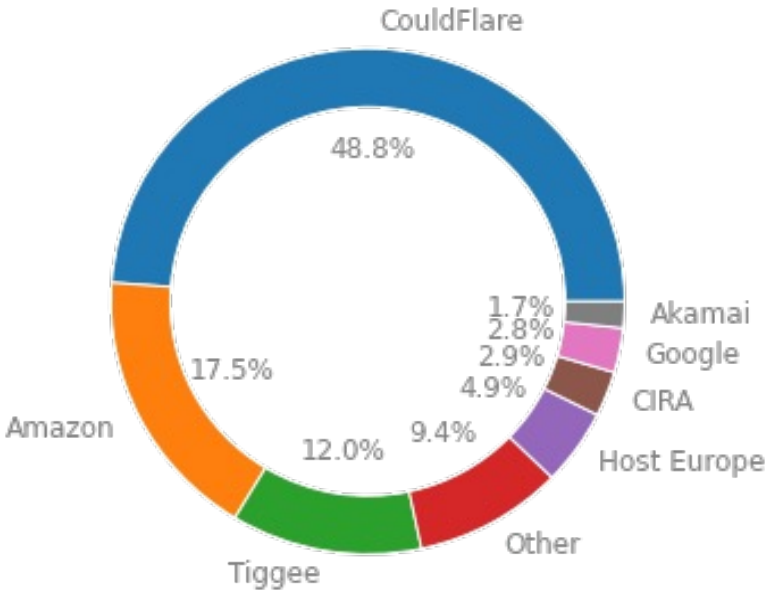

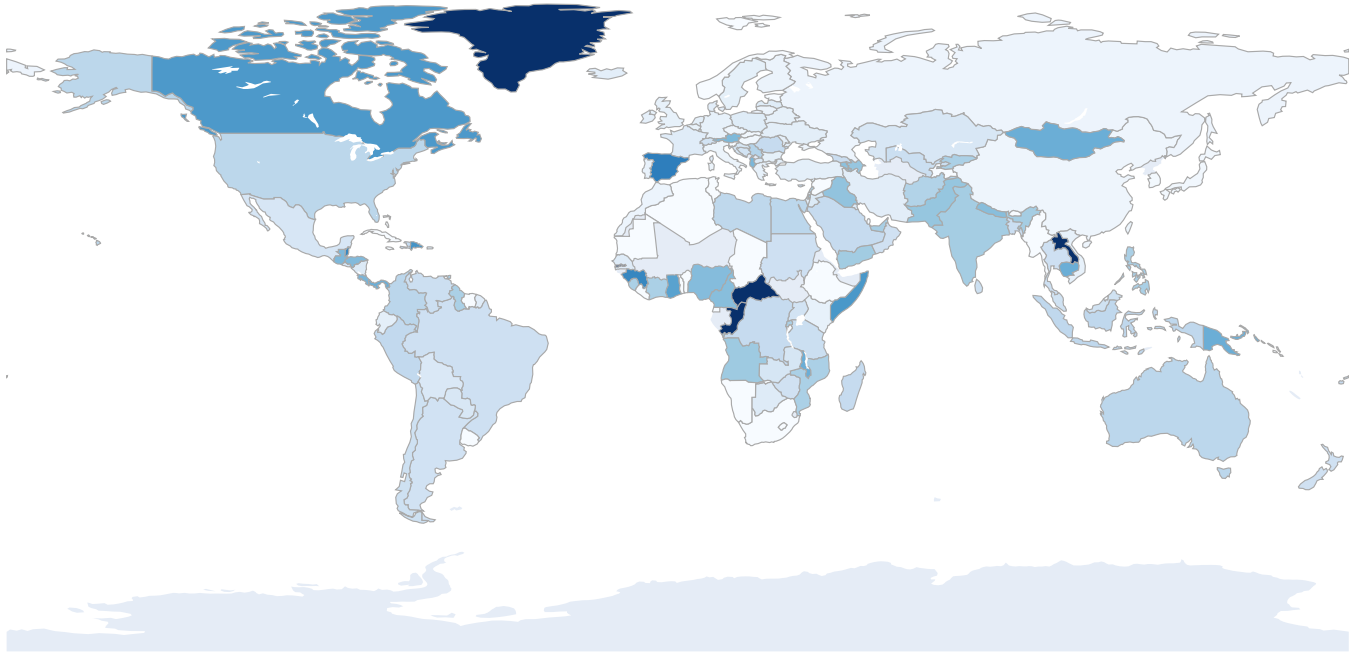

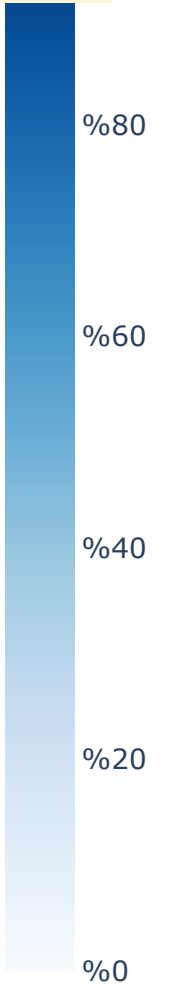

University Name Servers (NS) analyzed --> 15,218
University with anycasted name server (NS) --> 20 %

How Anycast works

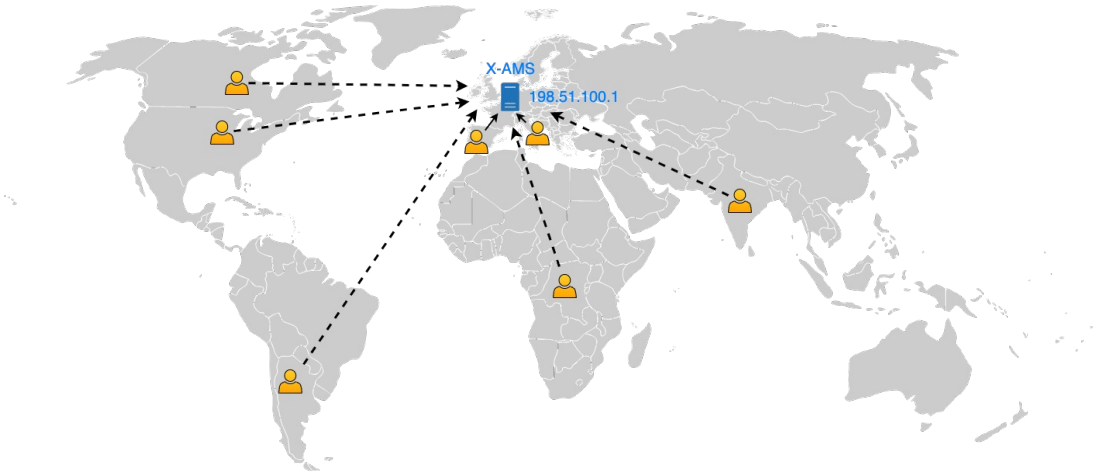

Unicast Service

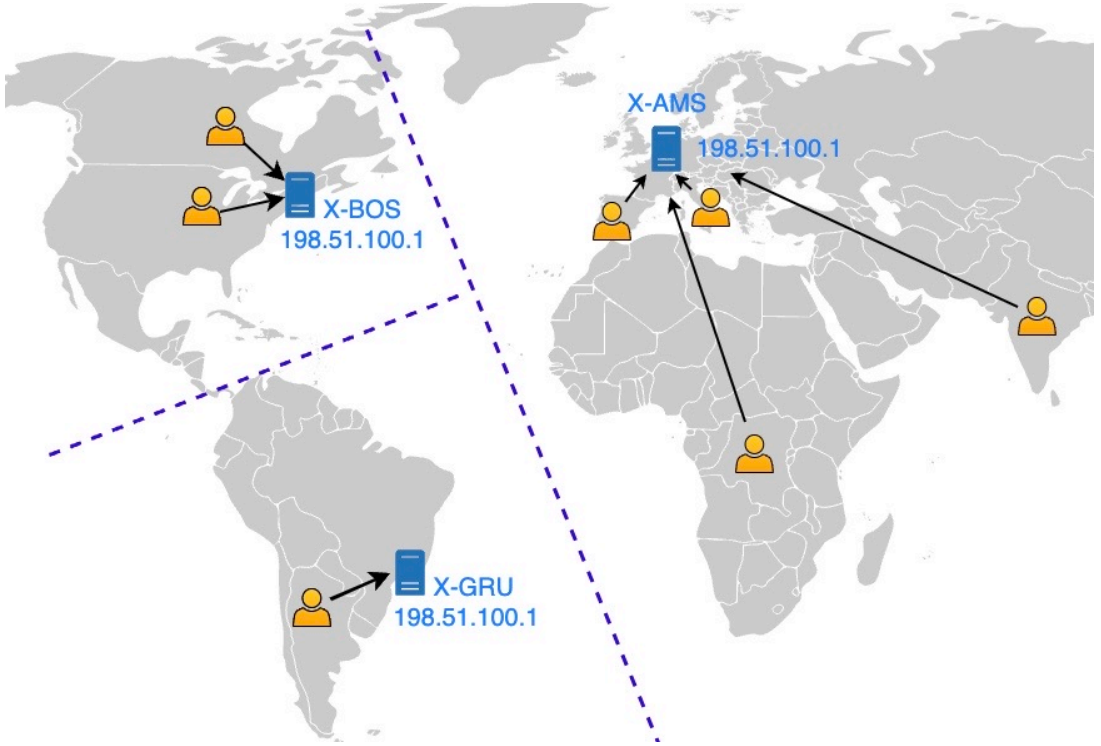

Anycast Service

-- Anycast as a defense mechanism -- more sites the better resilience !

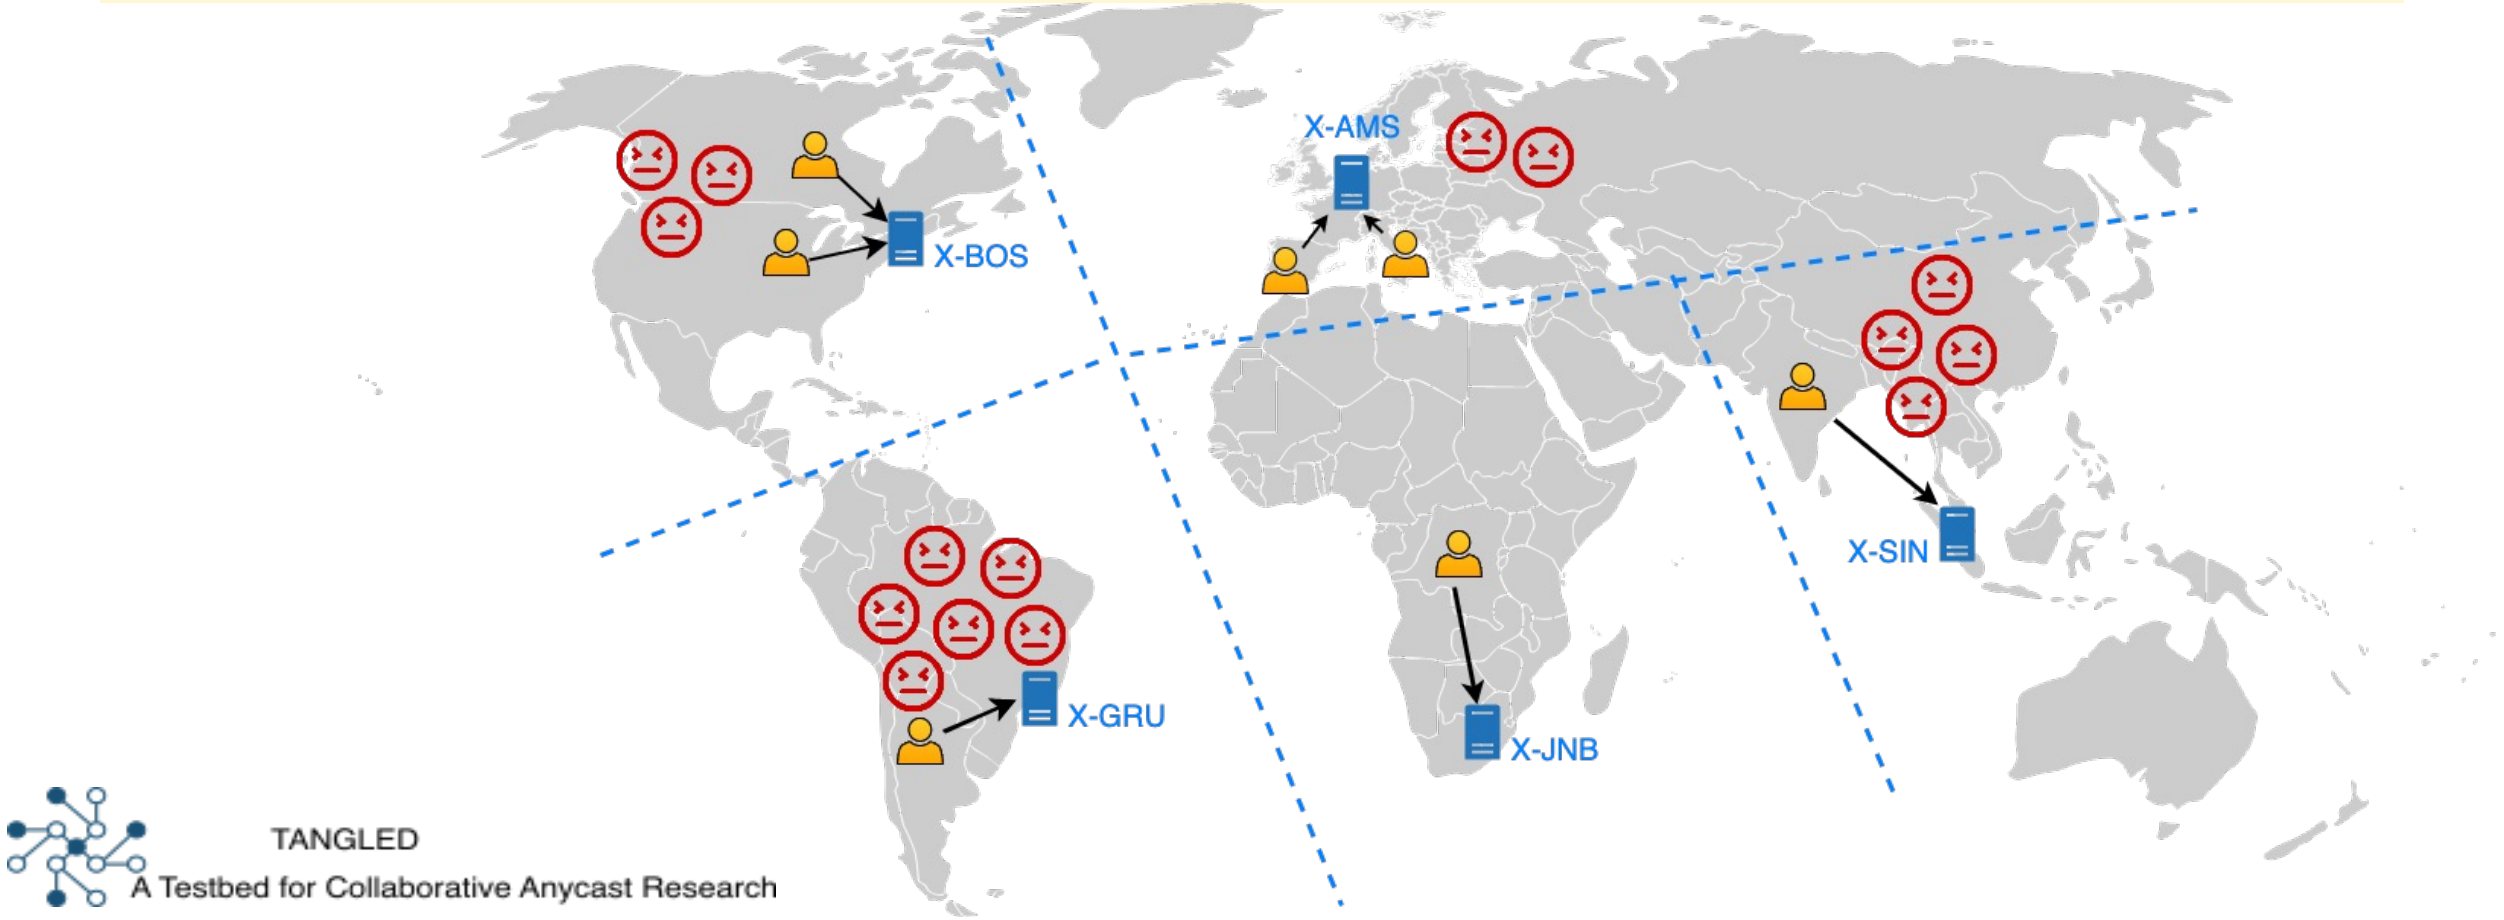

TANGLED
A Testbed for Collaborative Anycast Research

What we did…

https://youtu.be/ie5Gt7giMLw

Let's look the path to get there…

What happen in a DDoS Attack?

One site is **overwhelmed**

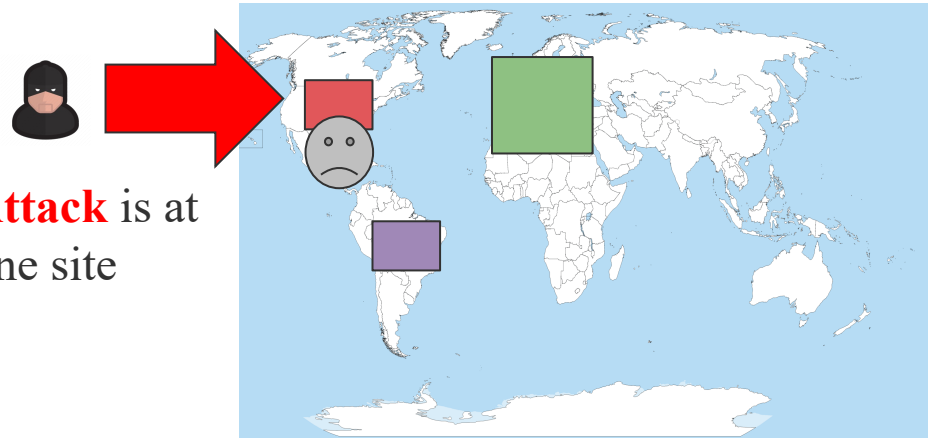

Attack is at one site

How to defend? 1- Absorb at One Site

One site is **overwhelmed**

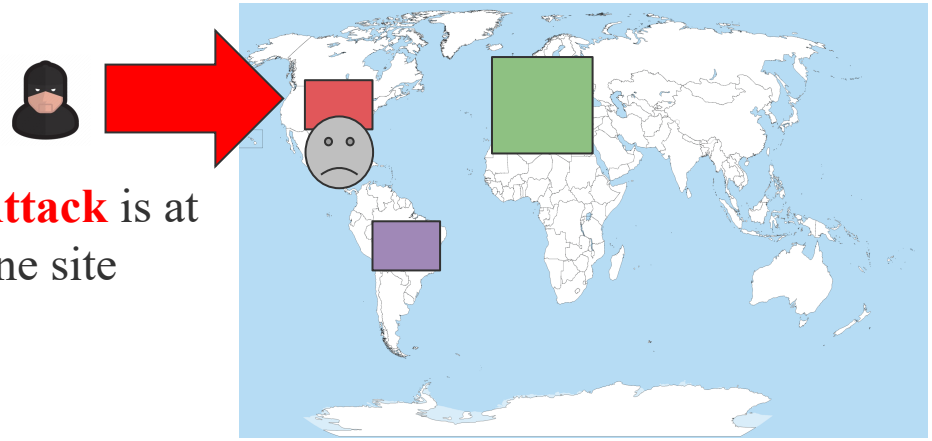

Attack is at one site

One site is hurt, but others are OK!

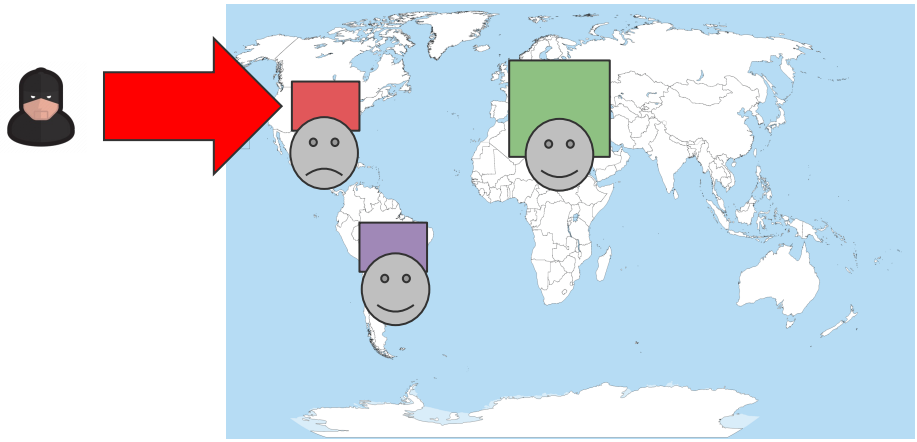

2- Spread Traffic

Rebalance the Network based on capacity

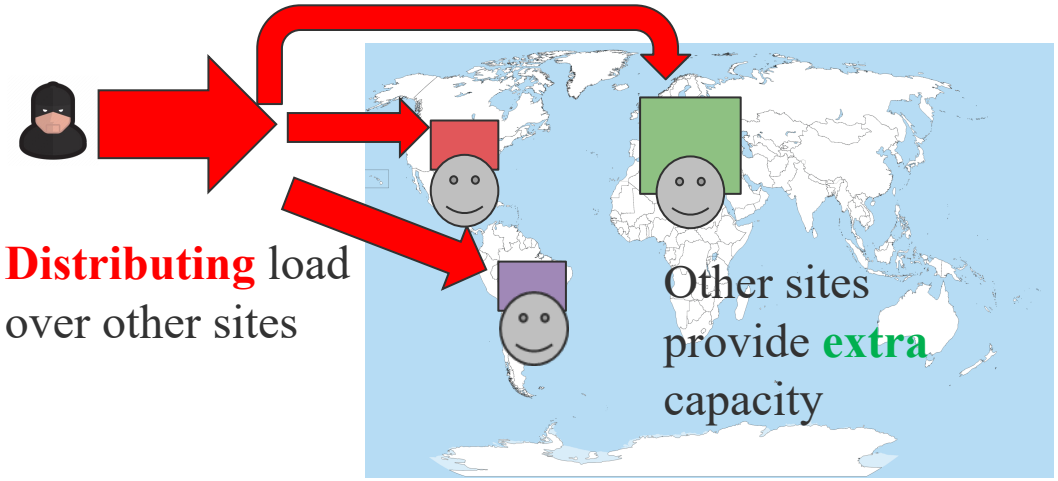

3- Shift Traffic

Shift to larger sites with spare/elastic capacity

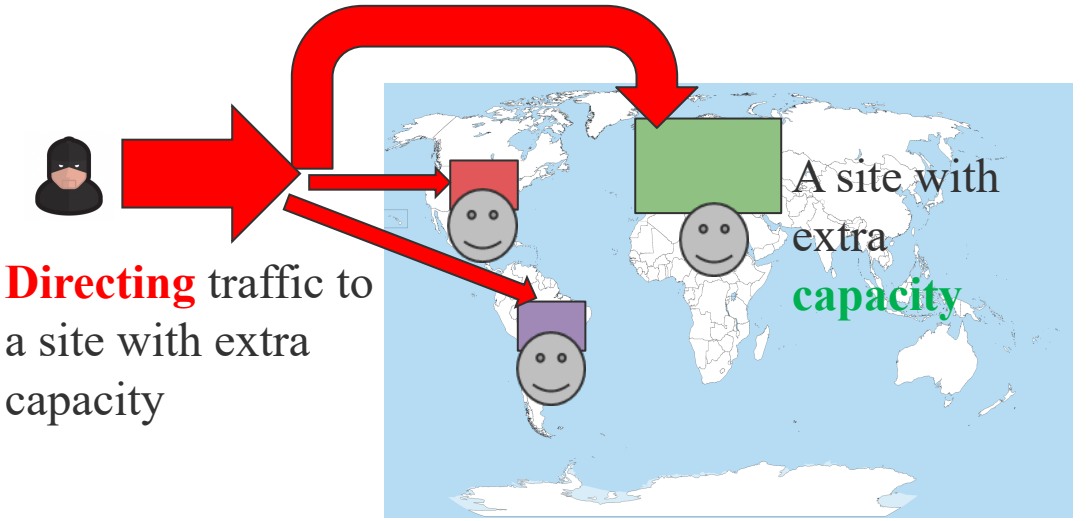

How does the redistribution? **BGP is unpredictable !**

What do you mean by “BGP unpredictable”?

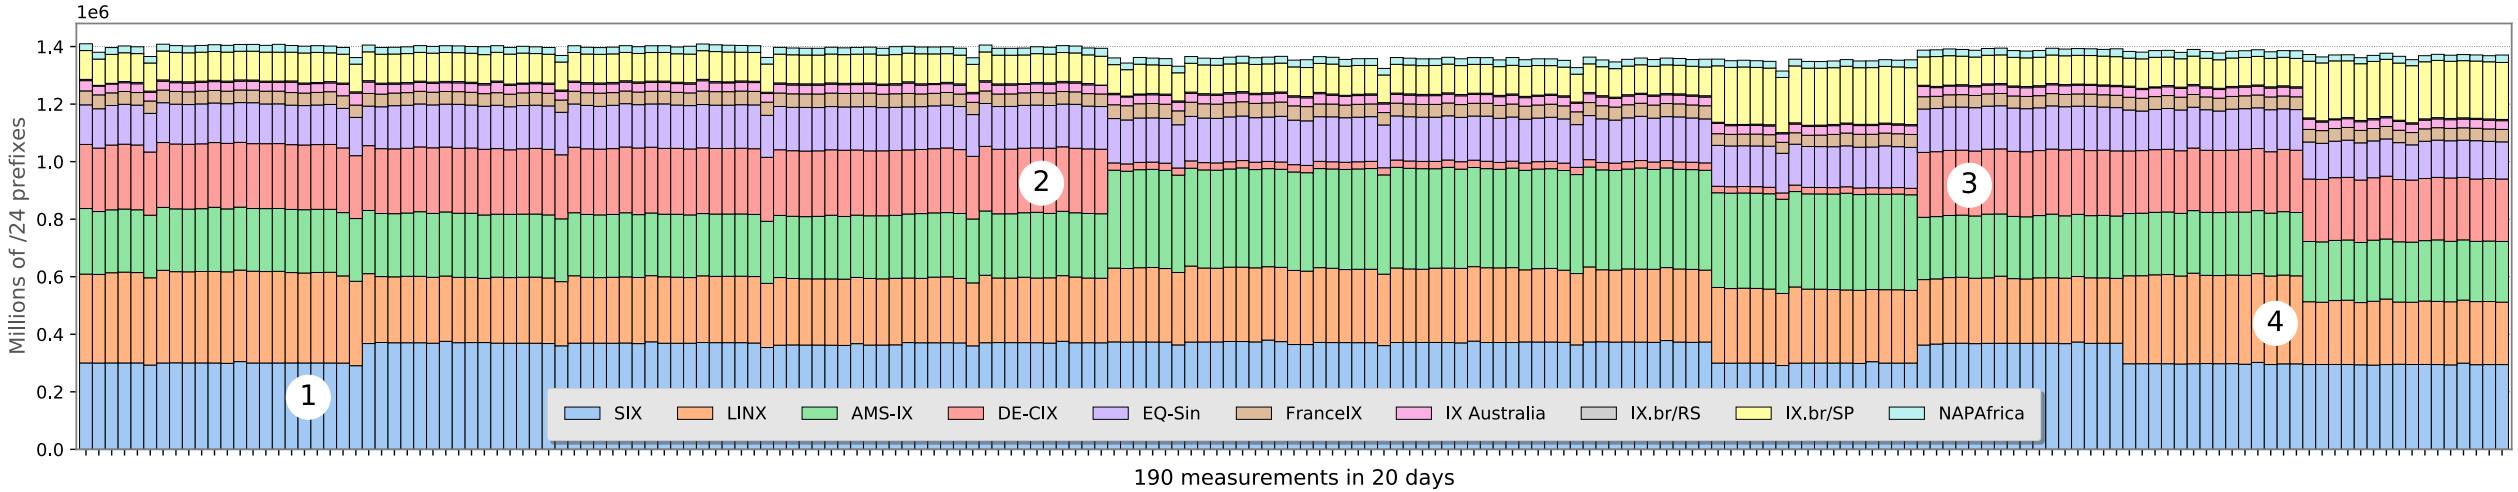

TANGLED

A Testbed for Collaborative Anycast Research

What do you mean by “BGP unpredictable”?

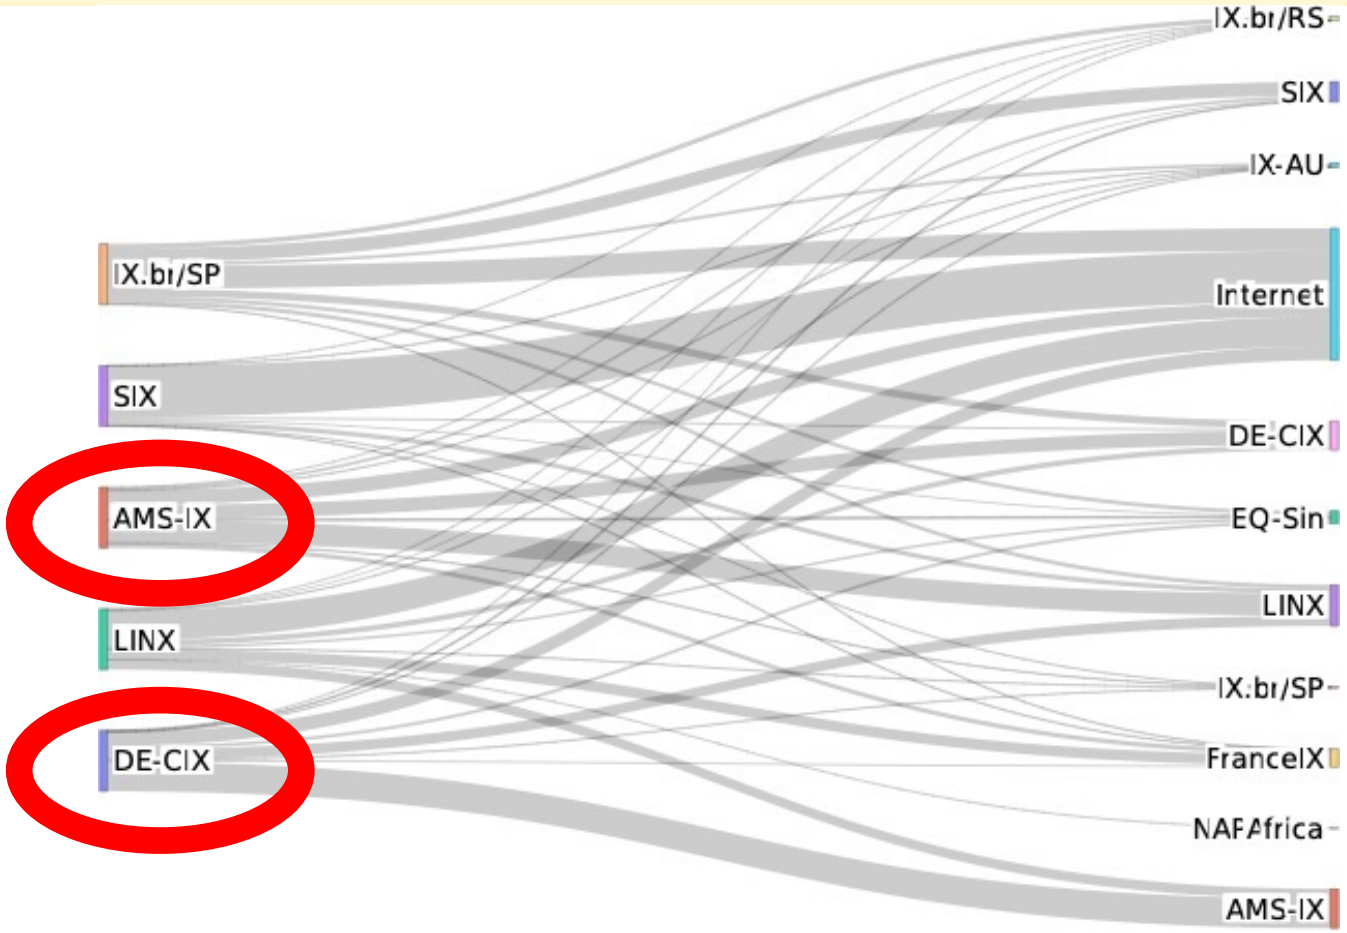

The Challenges

Challenge 1: Unknown Load

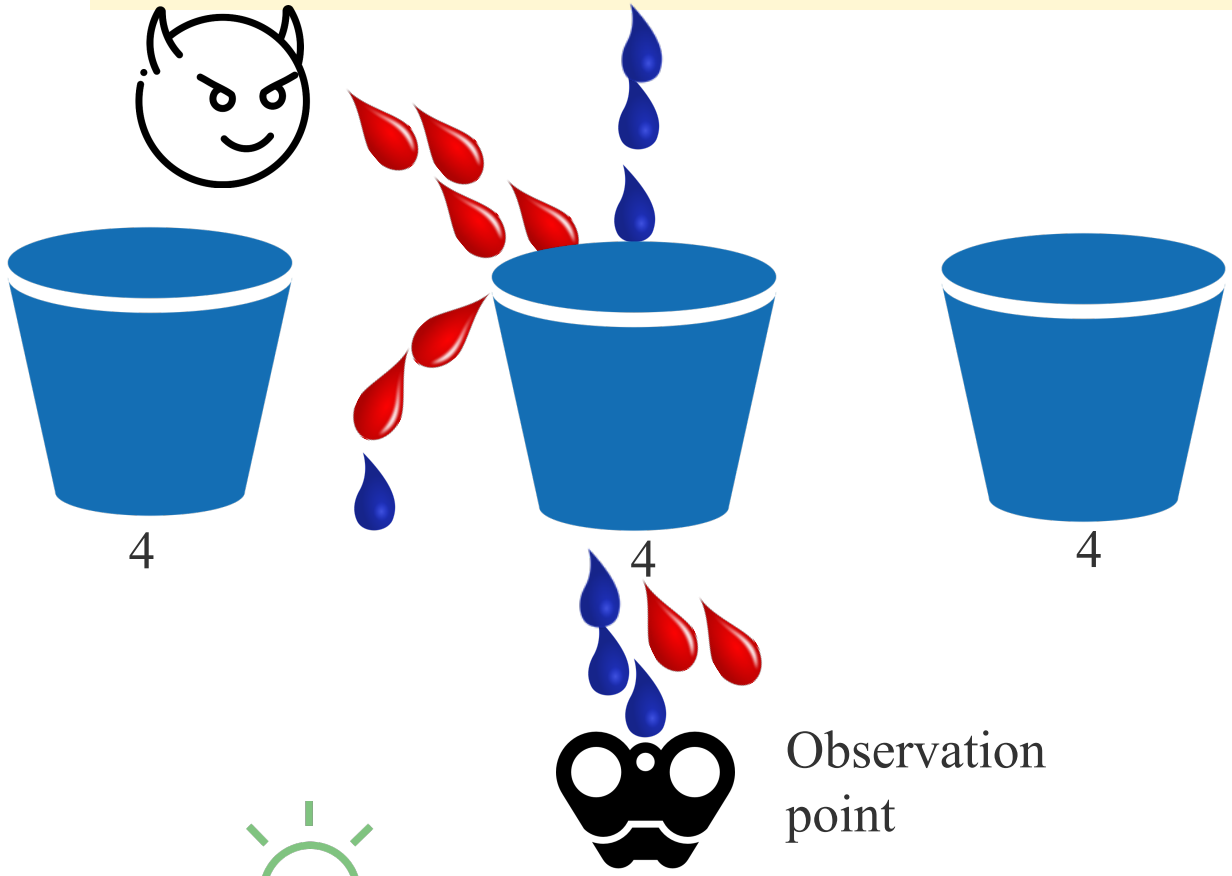

- What you see is
  - At full capacity: 50% attack traffic
- The truth is
  - At 175% capacity
  - 100% attack traffic
  - 75% legitimate
  - Lost 25% of legitimate traffic

Site observation under-estimates attack

Our **contribution**: proposing a way to estimate the **attack x offered load**

Challenge 2: Controlled Traffic Engineering

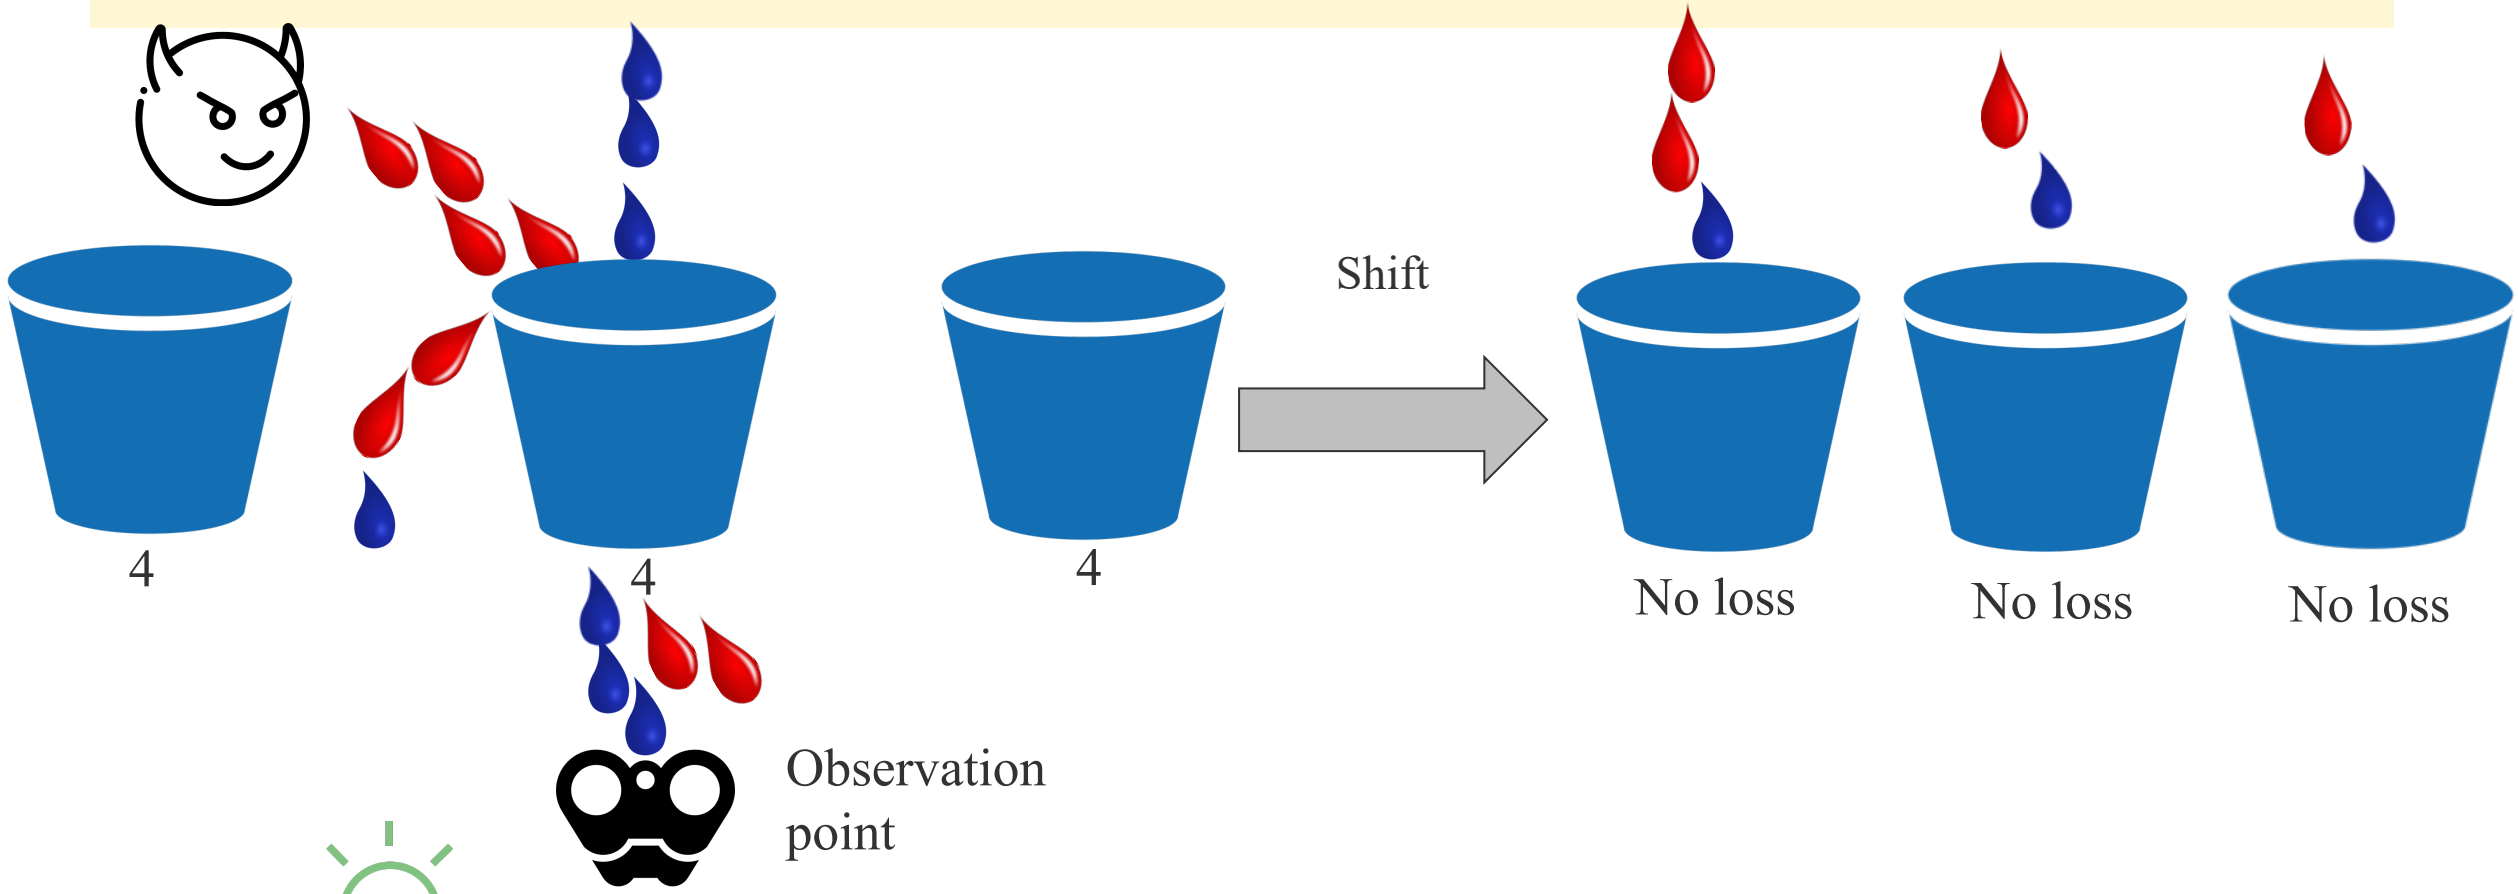

Our **contribution**: we help the operator to get the **right** shift

Challenge 3: How to redistribute?

?? **BGP assignment of traffic to anycast can be unpredictable** ??

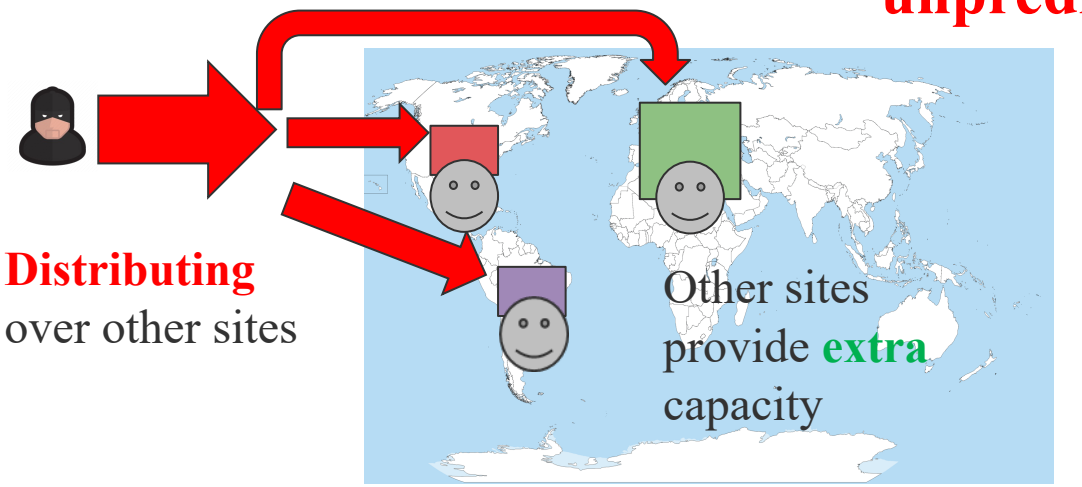

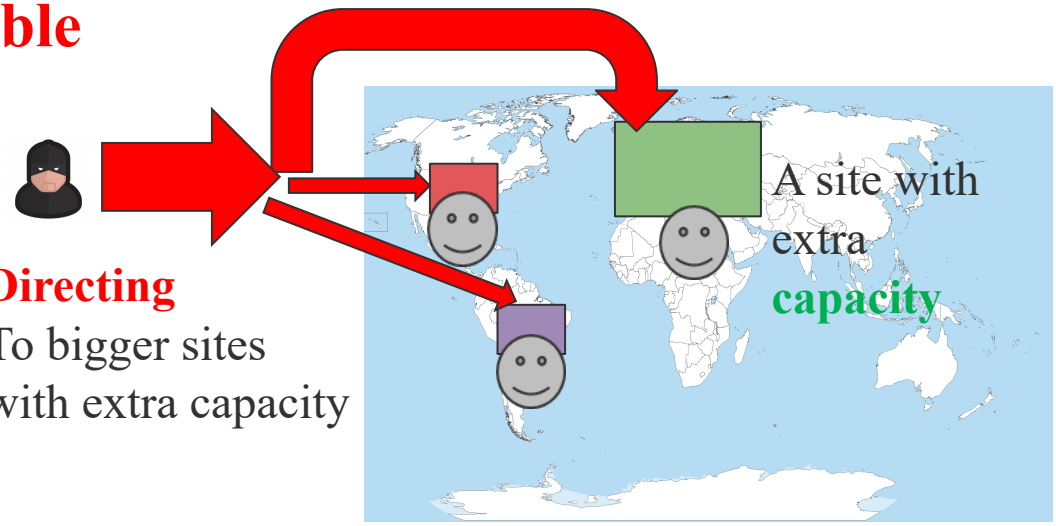

Our **contribution**: how to build a **BGP playbook** to predict anycast ahead of time

Our Contribution

- New approach to **estimate the load** (challenge 1)
  - **Allows us to plan a defense**
- Define a method to build **BGP playbook** (challenge 2)
  - Allows us to **execute the correct defense**
- Show a **BGP playbook works in a real DDoS event** (challenge 3)
  - **Effectiveness of our approach in real attacks.**

How it works?

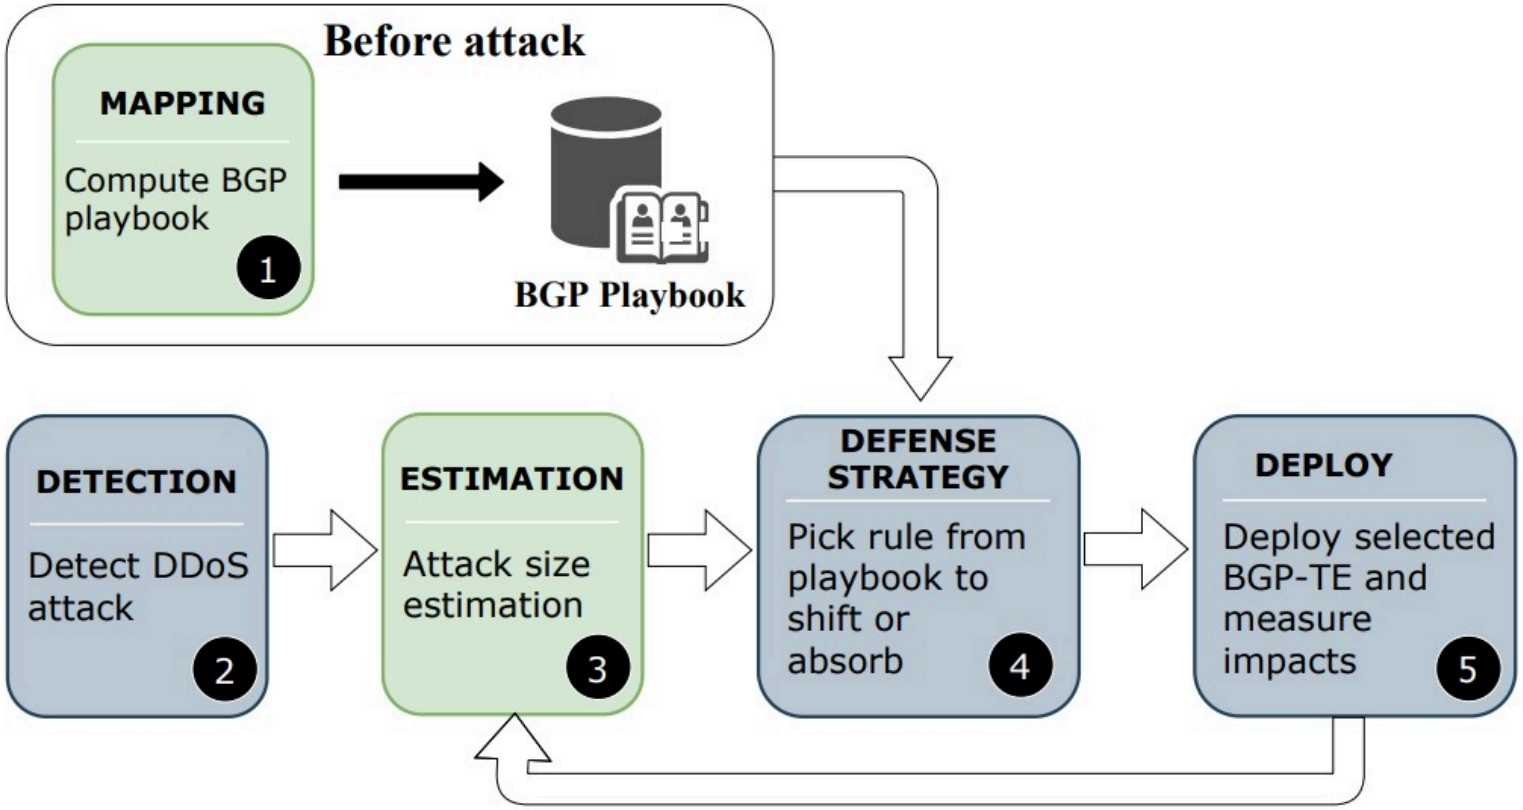
Before attack
MAPPING
Compute BGP playbook
1
BGP Playbook
DETECTION
Detect DDoS attack
2
ESTIMATION
Attack size estimation
3
DEFENSE STRATEGY
Pick rule from playbook to shift or absorb
4
DEPLOY
Deploy selected BGP-TE and measure impacts
5

Methodology: Estimating Load

- **Problem:**
  - **upstream loss is invisible**
- Insight:
  - **Heavy hitters**
  - Sites have predictable known good traffic
  - Infer attack size by change in this traffic

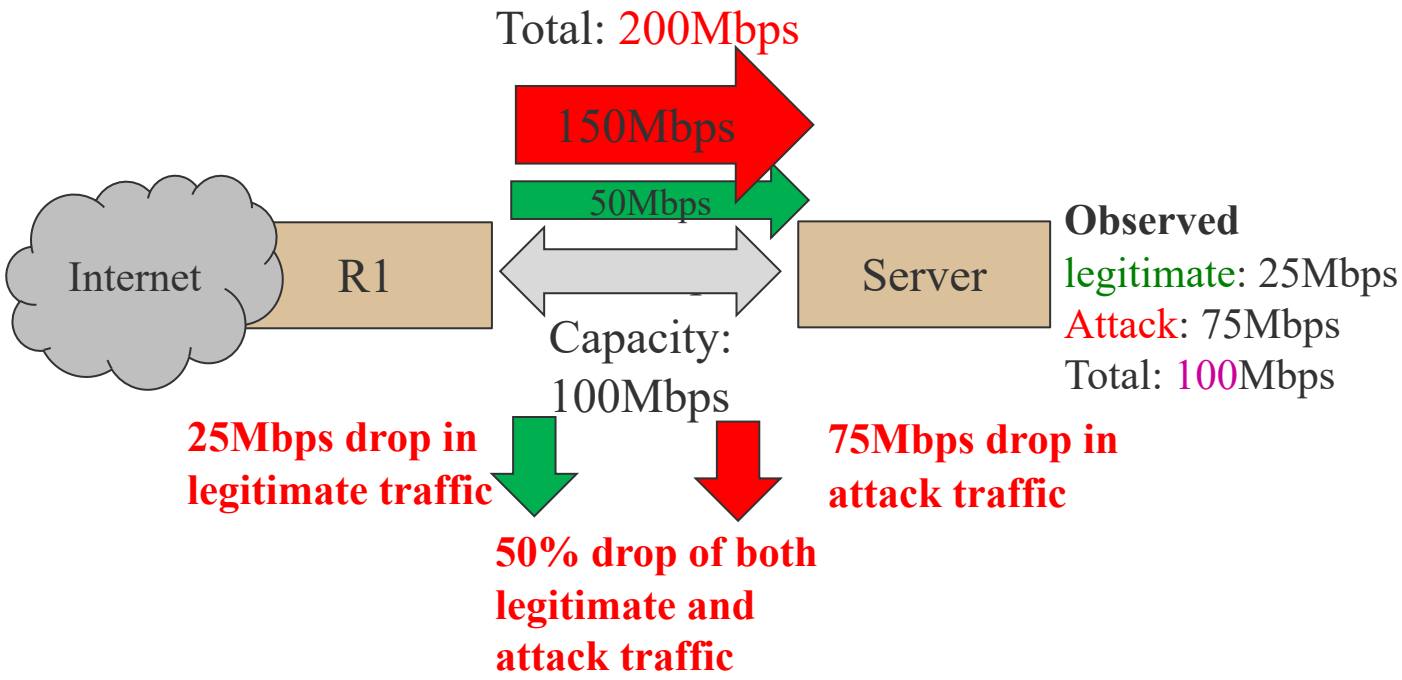

Methodology: Understanding Traffic Engineering (TE)

- We used three TE techniques
- Each TE method has tradeoffs (details in section 6)
  - Path prepending
    - Available in all sites
    - no granular control
  - Community strings
    - Not available in all sites
    - provide granular control
  - Path poisoning
    - Filtered when poisoning Tier-1 Ases
    - provide limited control

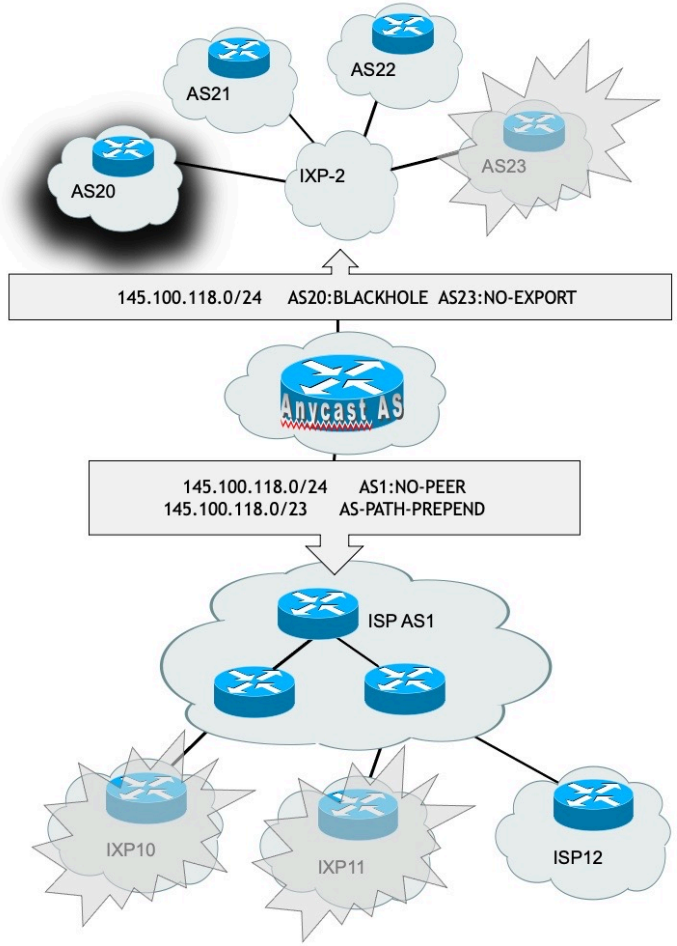

How we evaluate TE impact ?

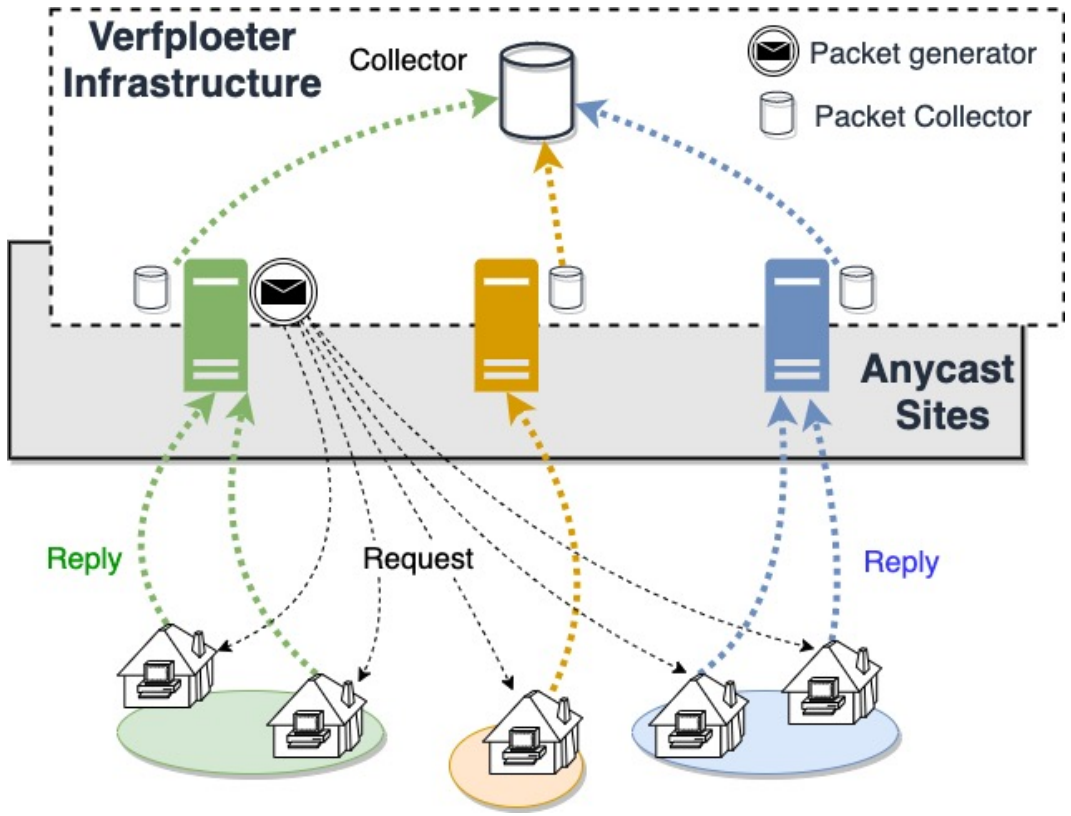

How a playbook looks like?

Routing Policy	Traffic to Site (%) AMS	BOS	CNF
(a) Route-server	15	35	55
(b) All-IXP-Peers/Poison transits	15	35	45
(c) 2xPrepend AMS	25	35	45
(d) 1xPrepend AMS	35	25	35
(e) -1xPrepend BOS	45	45	15
(f) -1xPrepend CNF	45	5	45
(g) Transit-1	45	25	35
(h) Transit-2	55	15	25
(i) Poison Tier-1/Transit-2	35	25	35
(j) Poison Transit-1	55	25	25
(k) Baseline	65	15	15
(l) 1,2xPrepend BOS	65	5	25
(m) 1,2,3xPrepend CNF	75	15	5
(n) -1,-2,-3xPrepend AMS	85	5	5

A sample playbook

Announcing only to **Transit-2:**
AMS: 55% traffic
BOS: 15% traffic
CNF: 25% traffic

Validation and Results

Offered Load Estimates are Accurate

- Question: **does estimation work?**
- Experiment:
  - **Replayed packet trace**
  - Measured observed traffic rate and access fraction to estimate
  - Compared the estimation with the reported rate
- **Answer: yes**

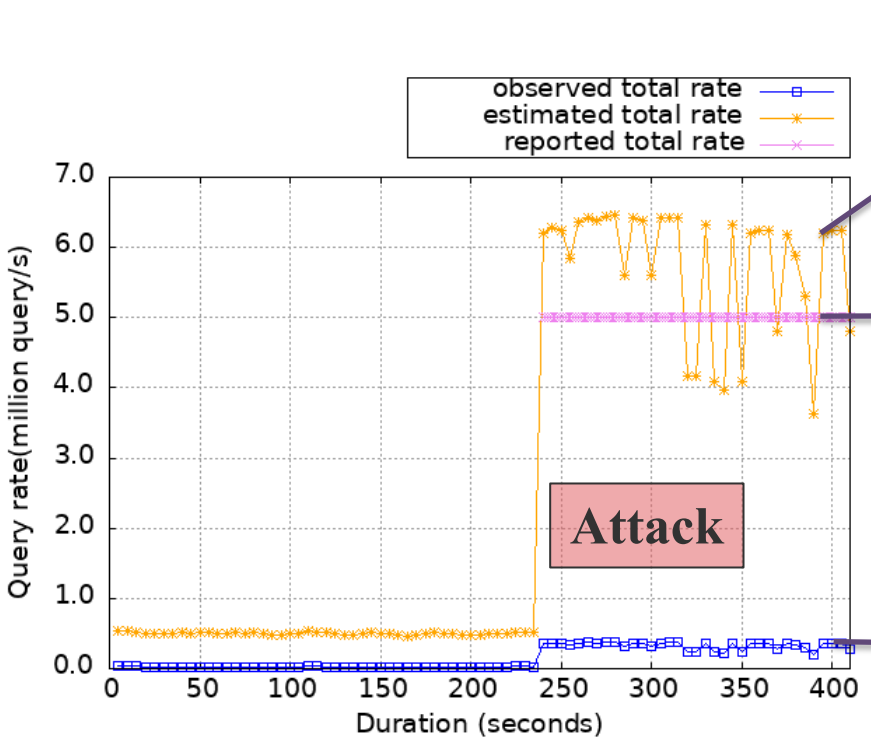

Our estimate: ~5 M query/s

True Offered load 5 M query/s

Observation: 0.35 M query/s (very low) because loss in upstream

Attack is root DNS attack from 2015-11-30 with data from B-root

Using a Playbook to Defend

- **Question: how to use a playbook during an attack?**
- Experiment:
  - Simulate a DNS attack
    - B-root event from 2017-03-06
    - More events in section 8 of the paper
  - Against a 3-site anycast system
    - Each site has ~60k queries/s capacity

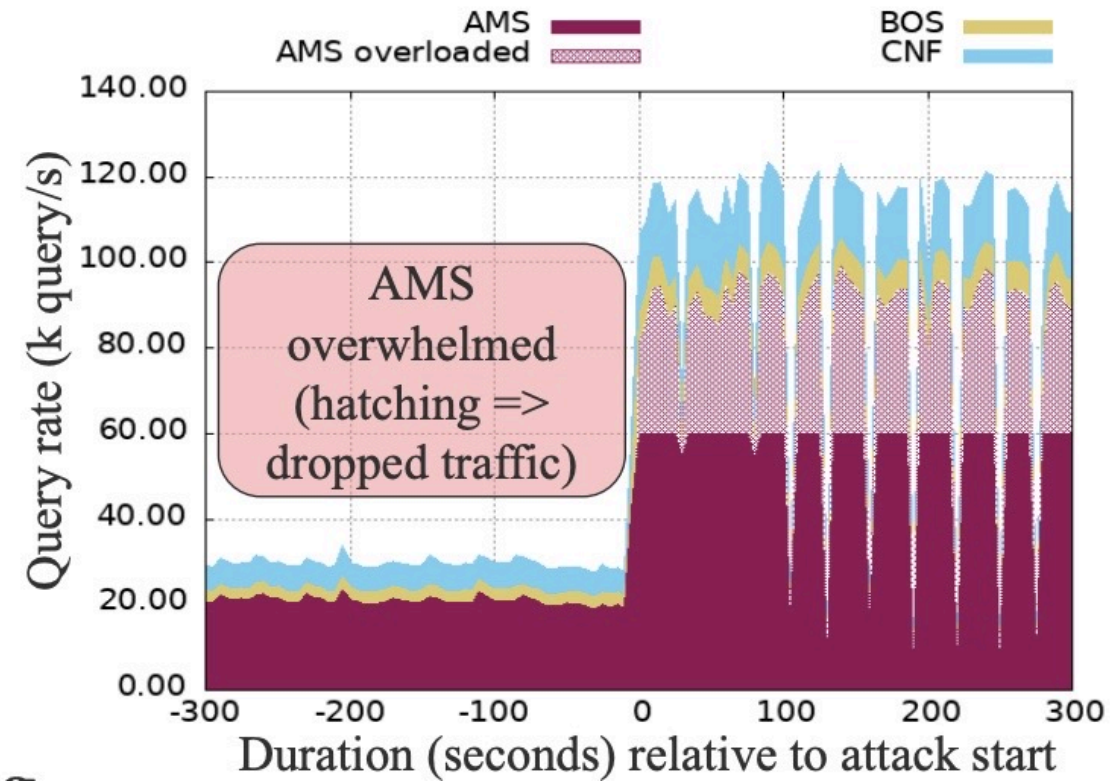

Let's look at the BGP playbook.

Solution: Playbook to Get Routing Options

Routing Policy	Traffic to Site (%) AMS	BOS	CNF
~~(a) Route-server~~	15	35	55
~~(b) All-IXP-Peers/Poison transits~~	15	35	45
(c) 2xPrepend AMS	25	35	45
(d) 1xPrepend AMS	35	25	35
(e) -1xPrepend BOS	45	45	15
(f) -1xPrepend CNF	45	5	45
(g) Transit-1	45	25	35
~~(h) Transit-2~~	55	15	25
(i) Poison Tier-1/Transit-2	35	25	35
~~(j) Poison Transit-1~~	55	25	25
~~(k) Baseline~~	65	15	15
~~(l) 1,2xPrepend BOS~~	65	5	25
~~(m) 1,2,3xPrepend CNF~~	75	15	5
~~(n) 1, 2, 3xPrepend AMS~~	85	5	5

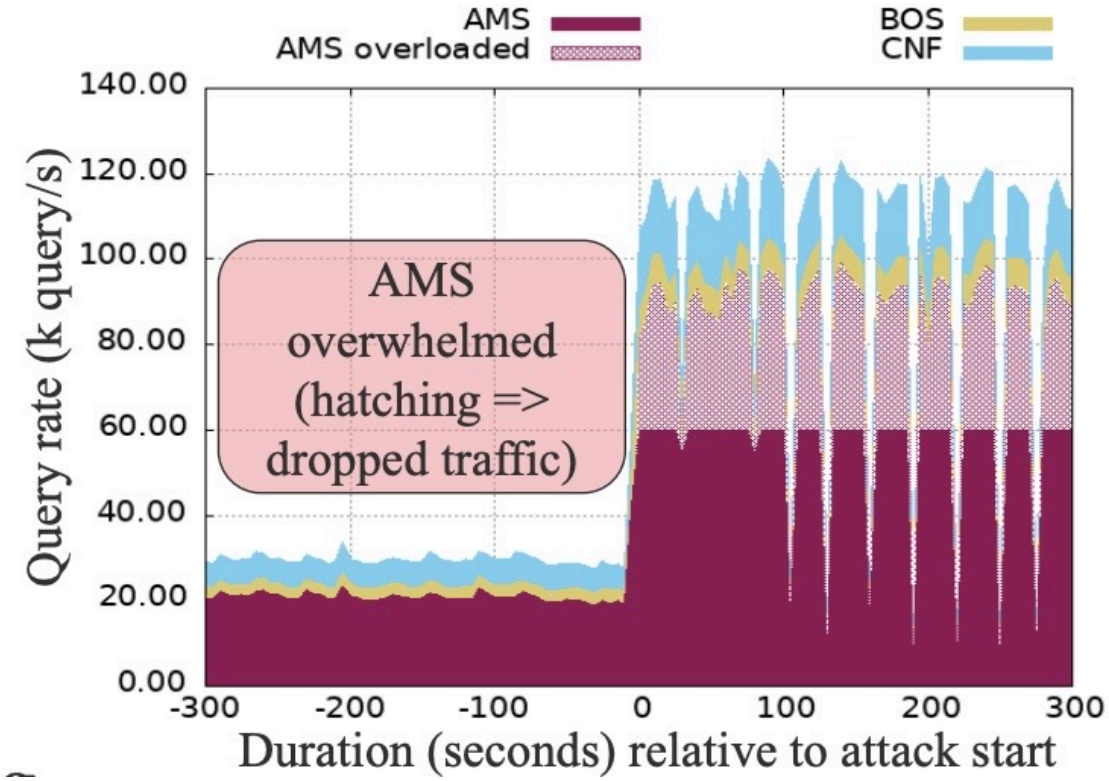

- Goal: lower traffic at AMS
- Several options work: c, d, e, f. g
- We pick **d** to avoid overloading other sites

Outcome after Applying a New BGP Policy

BGP changes at 300s; new traffic balance => no more drops (no hatching)

Routing Policy	Traffic to Site (%) AMS	BOS	CNF
~~(a) Route-server~~	~~15~~	~~35~~	~~55~~
~~(b) All-IXP-Peers/Poison transits~~	~~15~~	~~35~~	~~45~~
(c) 2xPrepend AMS	25	35	45
(d) 1xPrepend AMS	35	25	35
~~(e) -1xPrepend BOS~~	~~45~~	~~45~~	~~15~~
(f) -1xPrepend CNF	45	5	45
(g) Transit-1	45	25	35
~~(h) Transit-2~~	~~55~~	~~15~~	~~25~~
(i) Poison Tier-1/Transit-2	35	25	35
~~(j) Poison Transit-1~~	~~55~~	~~25~~	~~25~~
~~**(k) Baseline**~~	~~65~~	~~15~~	~~15~~
~~(l) 1,2xPrepend BOS~~	~~65~~	~~5~~	~~25~~
~~(m) 1,2,3xPrepend CNF~~	~~75~~	~~15~~	~~5~~
~~(n) 1, 2, 3xPrepend AMS~~	~~85~~	~~5~~	~~5~~

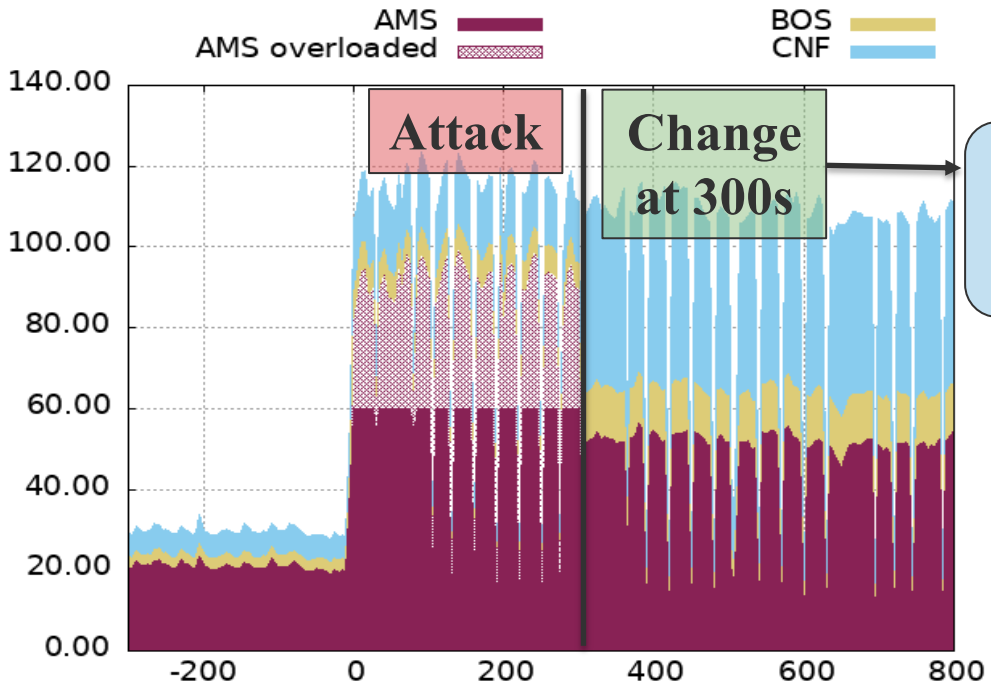

No loss
Some AMS traffic => CNF

Conclusion

- New method to **estimate attack size** from known good traffic
- Propose **BGP playbook** to plan reactions to DDoS
- **Evaluations against real attacks**
- More information about software
  - **Paper** https://www.usenix.org/system/files/sec22-rizvi.pdf
  - **Artifacts:** https://zenodo.org/record/6473023