DDoS Defense Mechanisms for IXP Infrastructures

Tim Dijkhuizen
Lennart van Gijtenbeek

Supervisor: Stavros Konstantaras (AMS-IX)

SNE: Research Project II

03-07-2018

Introduction

- **D**istributed **D**enial **o**f **S**ervice
- DDoS attacks on banks in NL [1]
- DDoS launched via botnets/booters
- Increase in size and complexity [2]
- IXP is a central entity
- Challenges:
  - High traffic loads
  - IXP neutrality
  - Complex infrastructure

Research Question

What (automated) solution can be developed to identify and mitigate DDoS attacks in an IXP network?

Image source - thenounproject.com

Internet eXchange Points (IXPs)

- Peering LAN (BGP)
- Exchange of traffic
- Wide range of networks connected
  - Such as banks, content providers, etc.
- Layer 2 forwarding (no routing)
- Route servers

Amsterdam Internet Exchange (AMS-IX)

- ~820 peers
- 5 Tbit/s peaks each day
- Traffic forwarding: MPLS/VPLS
- Statistics collector: sFlow
- Route server: BIRD
- Current DDoS solution
  - Disable port(s), NaWas

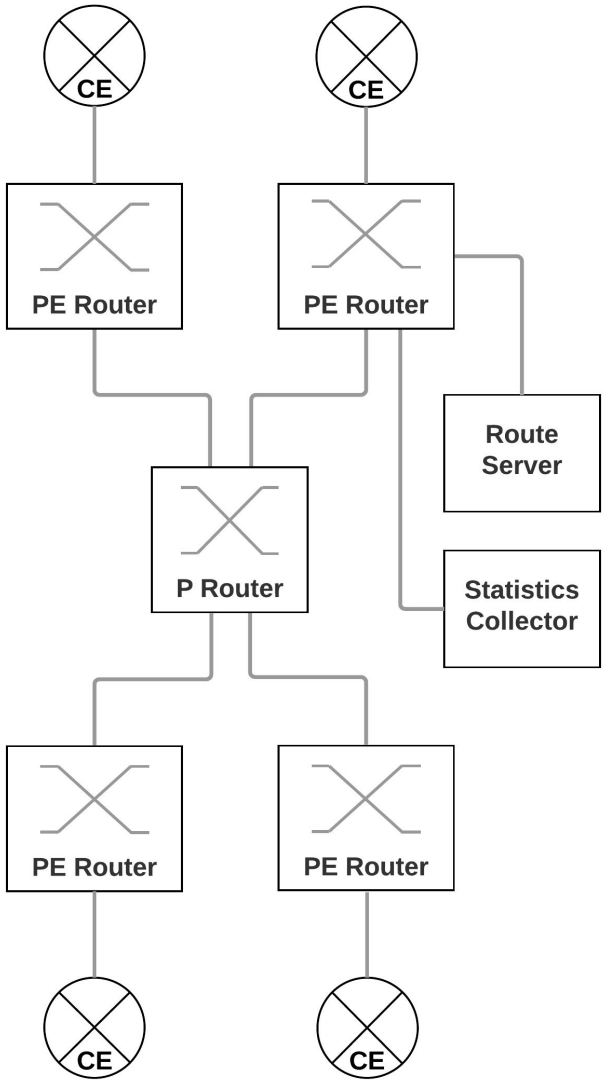

Types of DDoS Attacks

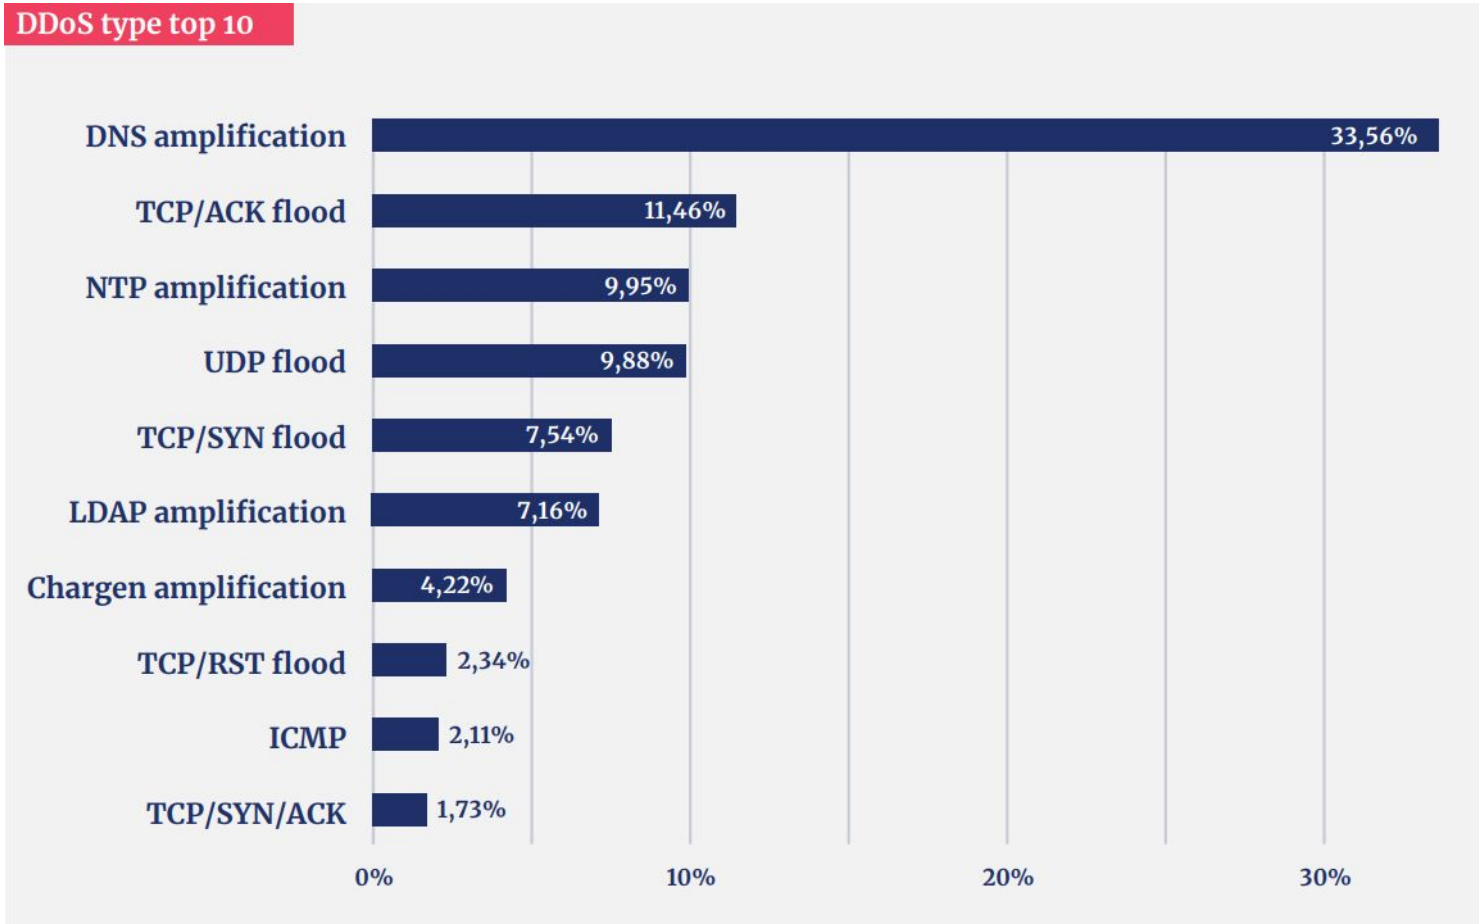

Image source - nbip.nl/nl/2018/05/16/nbip-ddos-data-report-2017-now-available/

Types of DDoS Attacks cont'd

- Volumetric attacks
  - Amplification attacks
    - E.g. DNS amplification
    - Small request, large response
- Protocol attacks
  - E.g. TCP SYN flood
  - State exhaustion
- Application attacks
  - Layer 7

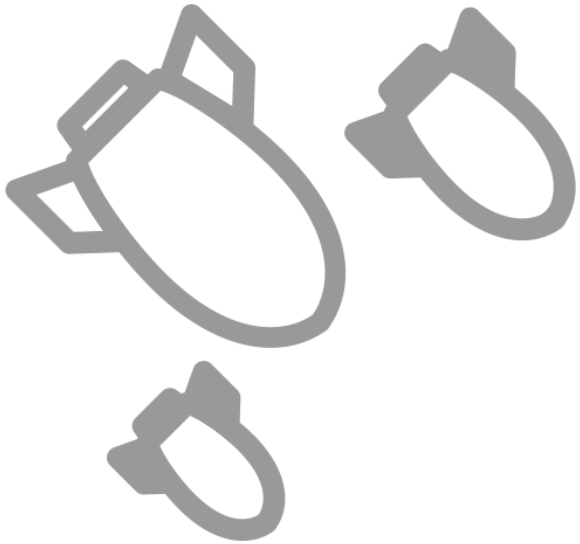

- No single detection method
- Distinct in: bandwidth and packets per second

Image source - thenounproject.com

Design Principles

1. Mitigate as close to the source as possible
2. No configuration required on the CEs
3. No congestion in the IXP core
4. Identification and mitigation on lower layers is preferred
5. Detect most common DDoS attacks
6. Intelligence resides in the IXP
7. Minimal impact on good traffic
8. IXP neutrality
9. Compatibility

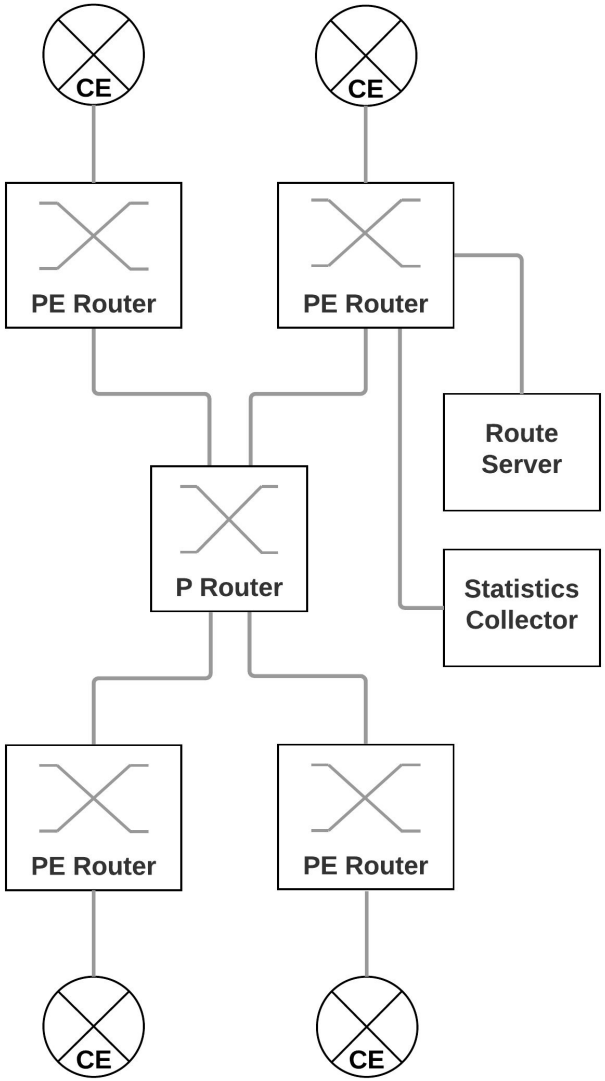

Detection Methods

- Traffic monitoring needed
  - PE switches
  - Sample data: sFlow/Netflow
- L2 detection
  - L2 headers are too limited
    - Frame size, CRC
  - Other parameters
    - Send rate, arrival interval
- L3/L4 detection

Detection Methods cont'd

- Threshold-based detection
  - Calculate thresholds based on destination IP(s)
    - Scalability: thresholds on prefixes
    - IXP environment: per source AS
  - Metrics:
    - L2/L3: BPS, PPS
    - L4: TCP flags, source ports, destination ports
- Fingerprint-based detection
  - DDoSDB [3]
  - False negatives

Mitigation Methods

- Scrubbing
  - On-site
    - Proprietary box
  - Off-site
    - NaWas
- Access Control Lists
- Software Defined Networking (SDN)
- BGP Blackholing

Blackholing Techniques with BGP

- Source-based blackholing
  - IXP neutrality
  - IP spoofing / false positives
- Destination-based blackholing on the CE
  1. Route withdrawal
  2. Static routing entry for prefix to Null0 and announce next-hop
- Destination-based blackholing on the PE
  - Set CE next-hop to ARP-dummy
  - L2 ACL

Design Proposal

Image source - thenounproject.com

Added Components to IXP

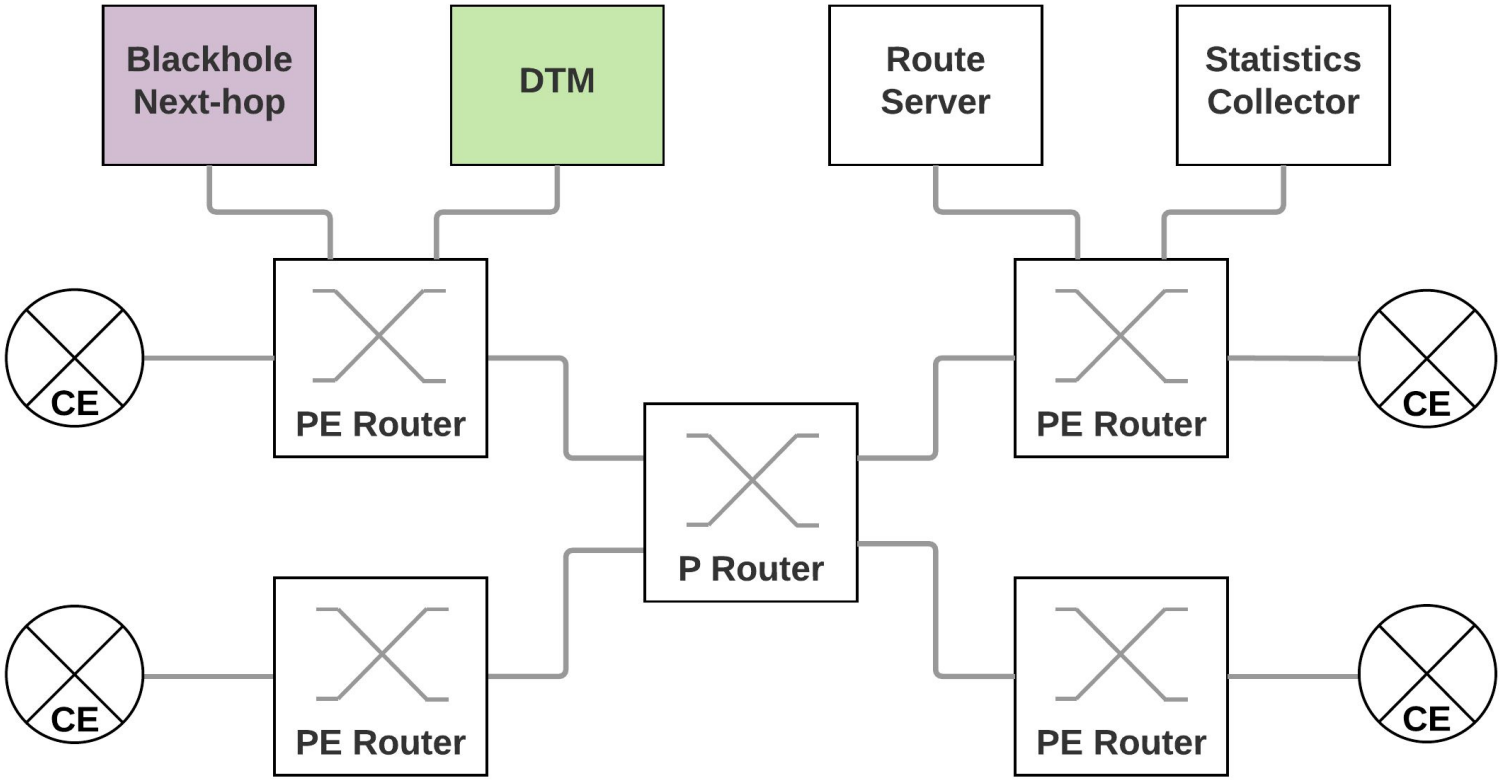

DTM = DDoS Threat Mitigator

Component Interaction

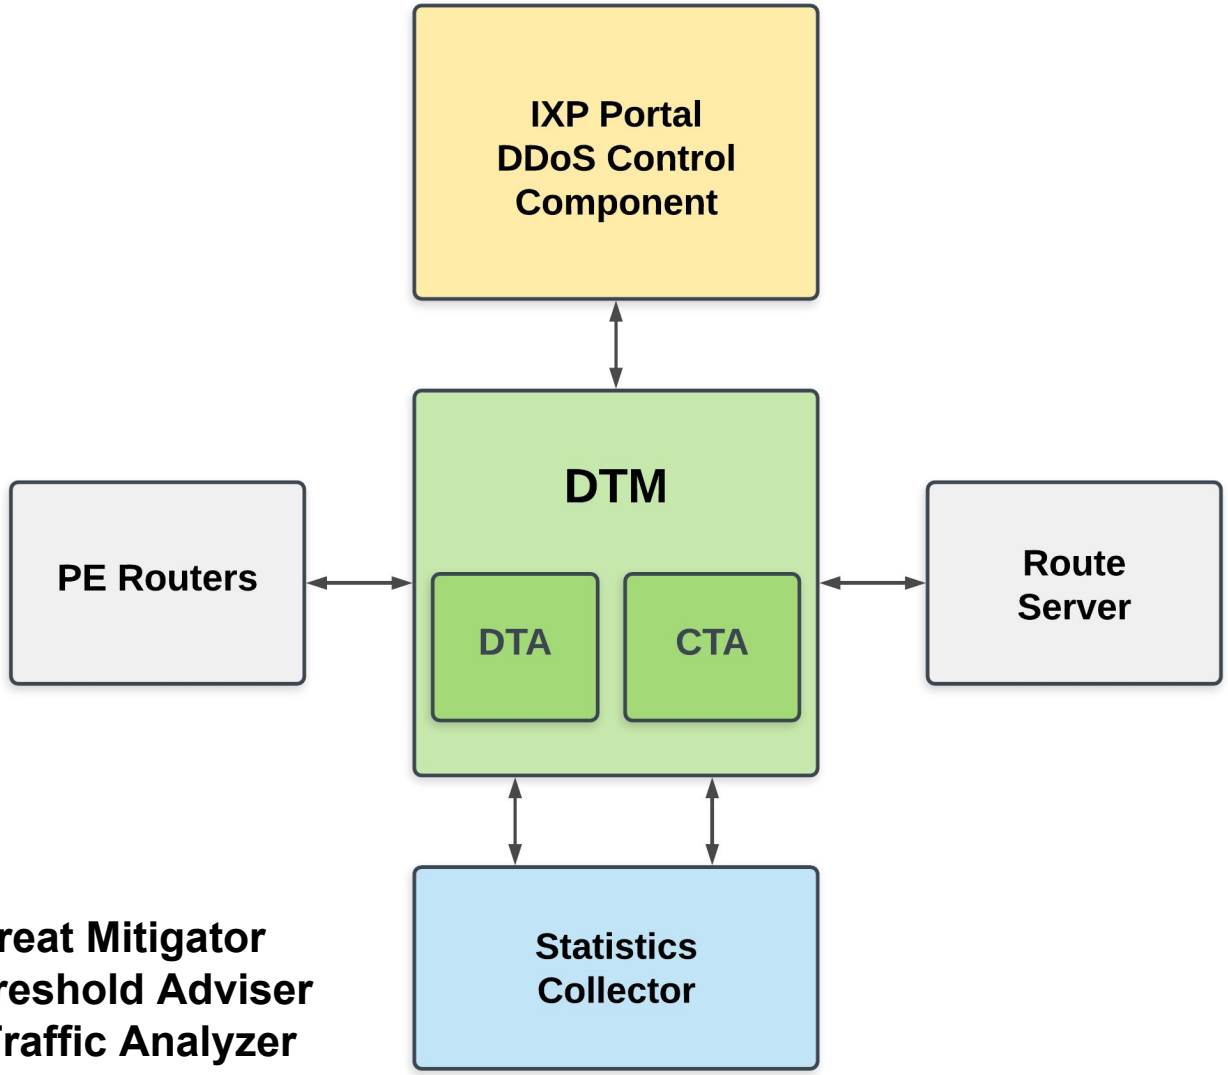

DTM = DDoS Threat Mitigator
DTA = DDoS Threshold Adviser
CTA = Current Traffic Analyzer

Design Proposal

Threshold-based detection

Three-way mitigation

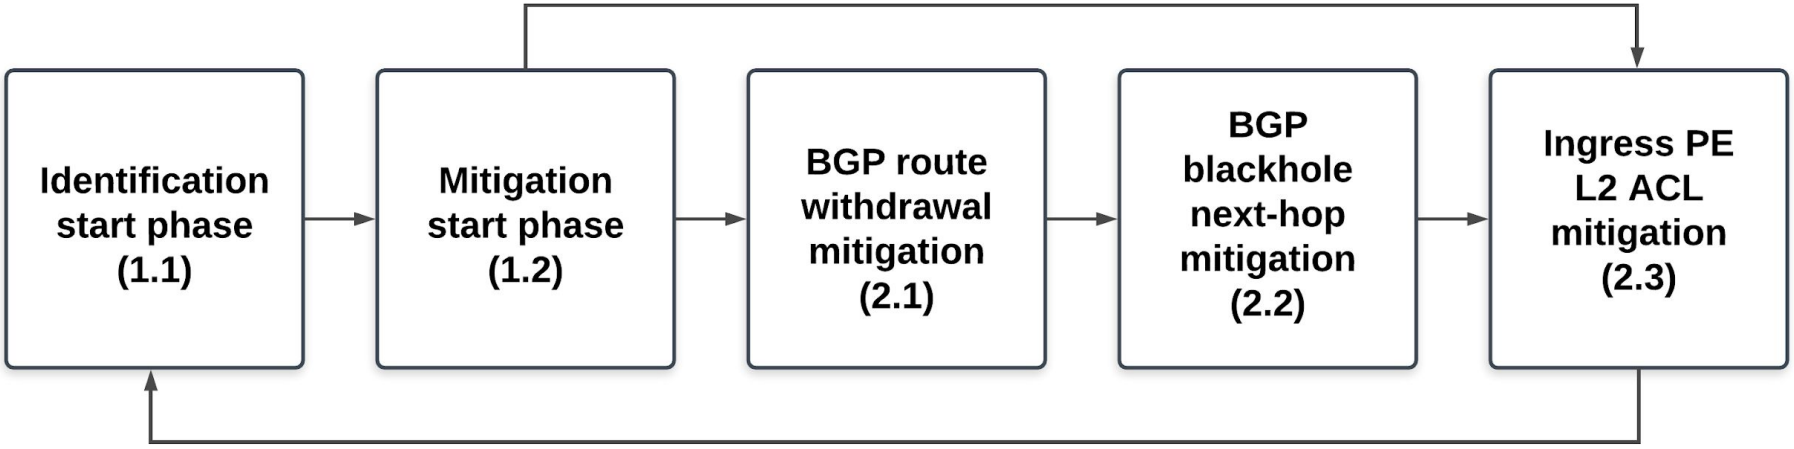

Design Workflow

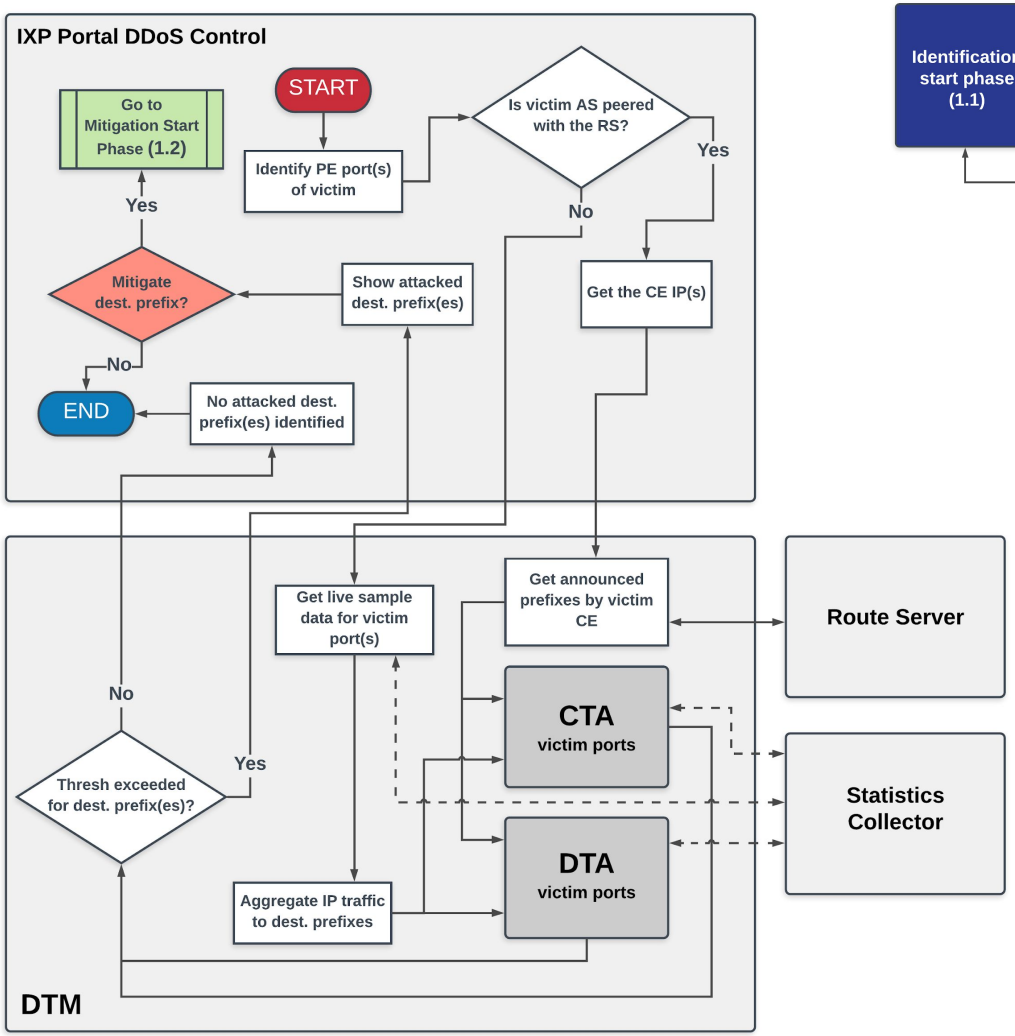

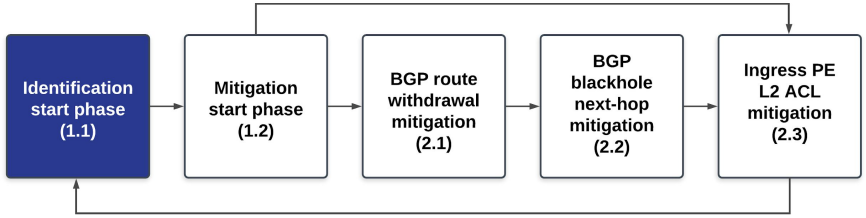

Identification Start Phase (1.1)

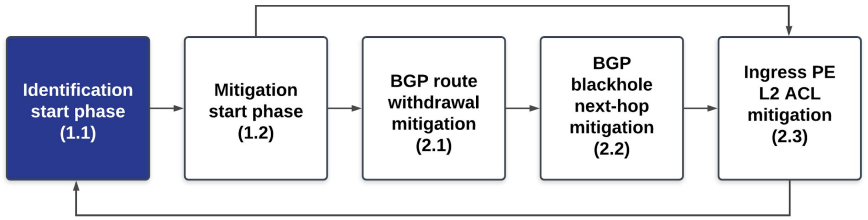

1. Peer starts the process
2. Identify PE port(s) of the victim
3. Get the CE IP, and announced prefixes (RS)
4. Start the DTA/CTA
    - Based on victim ports, and destination prefixes
5. Perform threshold comparisons
6. Present customer with exceeded prefixes
    - Customer decides which prefixes to mitigate

Mitigation Start Phase (1.2)

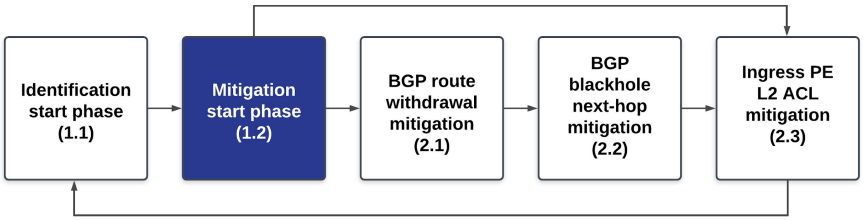

1. Determine the culprit AS(es)
   - Compare current to historical traffic
   - ASes to mitigation prefix
2. Determine mitigation workflow
   - Culprit AS is peered with RS:
     - Perform mitigation via BGP route withdrawal **(phase 2.1)**
   - Culprit AS is NOT peered with RS:
     - Perform mitigation via ACL on the ingress PE **(phase 2.3)**

CE Route Withdrawal Mitigation (2.1)

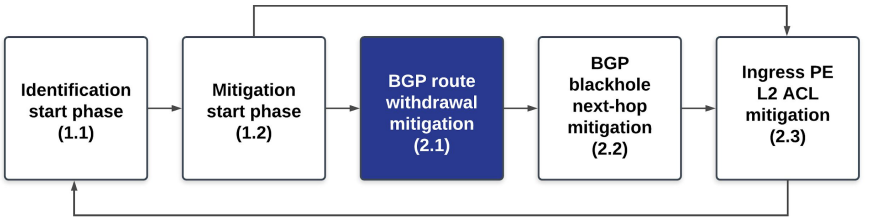

- Instruct the RS to withdraw the destination prefix to culprit
  - Wait for <*BGP_convergence_timeout*>
- Threshold is still exceeded:
  - Method *unsuccessful*, restore original BGP announcement
  - Perform mitigation via BGP blackhole nexthop (**phase 2.2)**
- Threshold is NOT exceeded:
  - Continue mitigation until DDoS no longer active
  - DDoS stopped or mitigation still working?

CE Blackhole Next-hop Mitigation (2.2)

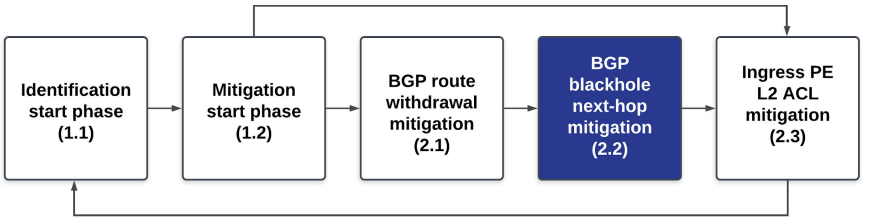

- Instruct the RS to announce blackhole next-hop to culprit
  - Wait for <*BGP_convergence_timeout*>
- Threshold is still exceeded:
  - Method *unsuccessful*, restore original BGP announcement
  - Perform mitigation via L2 ACL **(phase 2.3)**
- Threshold is NOT exceeded:
  - Continue mitigation until DDoS no longer active
  - Monitor on ingress PE

PE L2 ACL Mitigation (2.3)

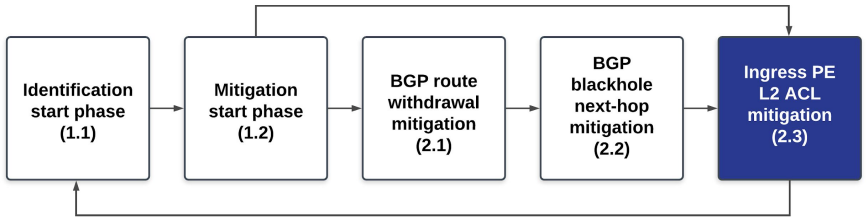

- Determine MAC addresses and DDoS ingress PE
- Instruct the PE to set up L2 ACL on the ingress PE
  - Based on source CE and destination CE
  - Wait for *<ACL_timeout>*
- Threshold is still exceeded:
  - Identification *unsuccessful*, remove ACL and go to **phase 1.1**
- Threshold is NOT exceeded:
  - Continue mitigation until DDoS no longer active
  - Monitor on ingress PE

Proof of Concept

- Focused on mitigation phases
  - Prefix identification, DTA, culprit AS identification
- Four different scenarios
  - Peered with RS:
    - 2.1 ✔
    - 2.1 ✘, 2.2 ✔
    - 2.1 ✘, 2.2 ✘, 2.3 ✔
  - Not peered with RS:
    - 2.3 ✔

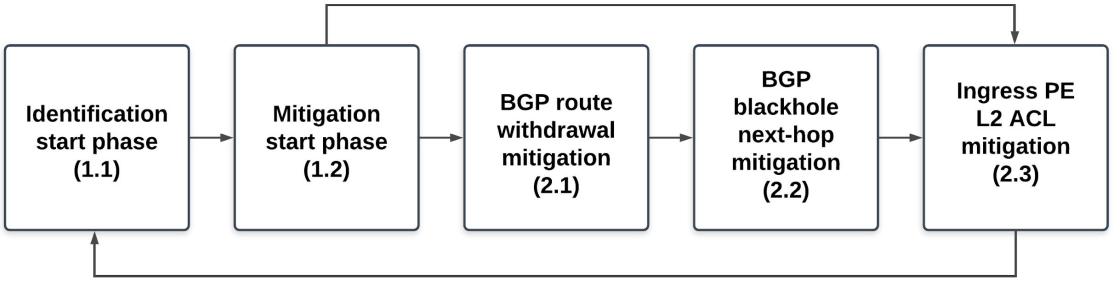

Proof of Concept cont'd

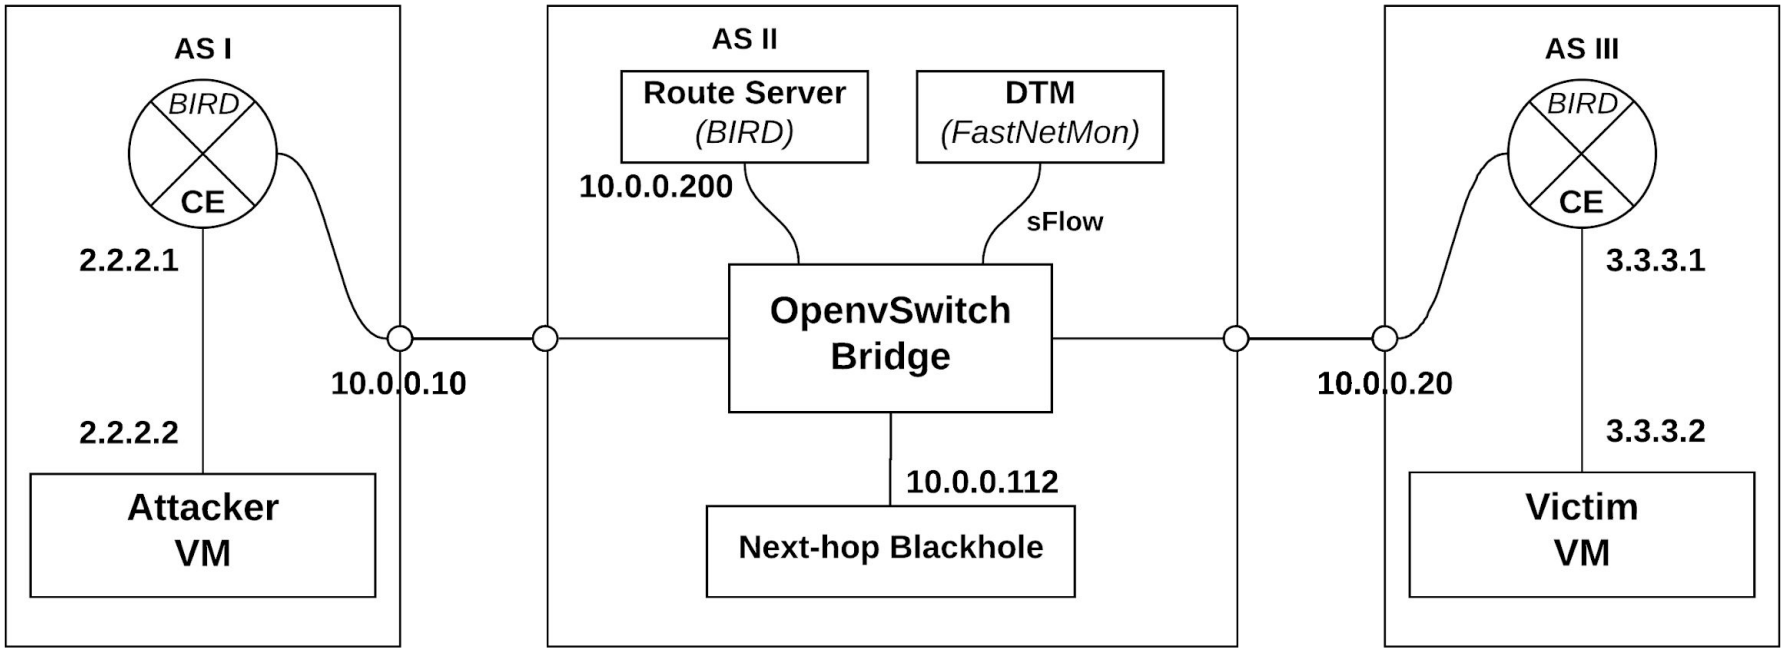

The **DTM** here also functions as the statistics collector
FastNetMon: DDoS detector that supports multiple packet capture engines
iPerf to generate traffic

Proof of Concept cont'd

- Culprit AS is peered with RS
- BGP route withdrawal mitigation (2.1)
- *Converge timeout: 10s, analysis: 4s*
- *50Mbit normal traffic, 150Mbit threshold*

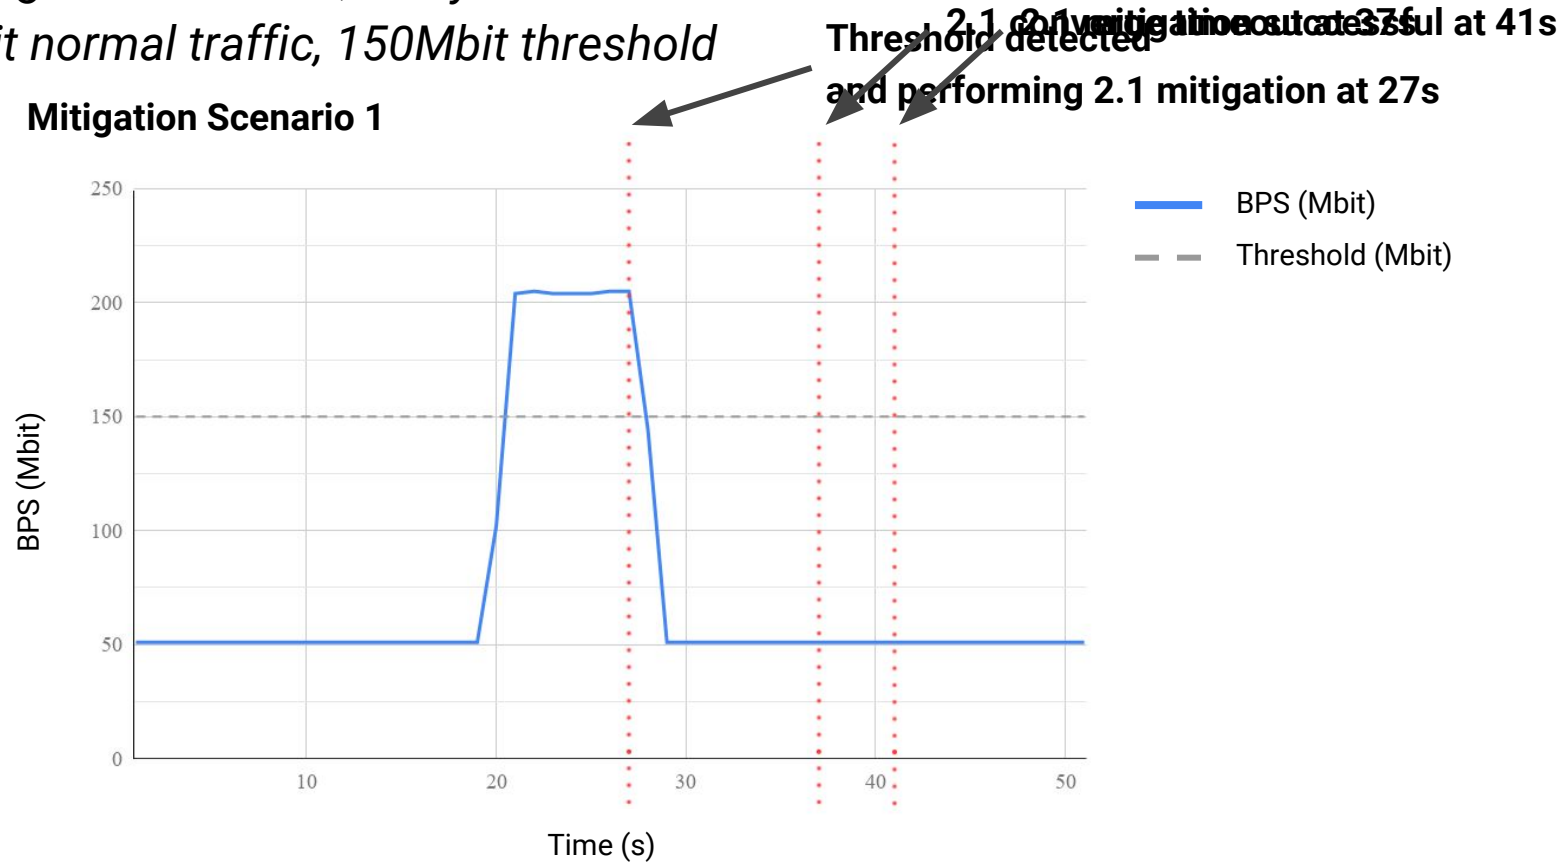

Proof of Concept cont'd

- Culprit AS is peered with RS
- BGP route withdrawal mitigation *unsuccessful* (2.1)
- BGP blackhole next-hop mitigation (2.2)

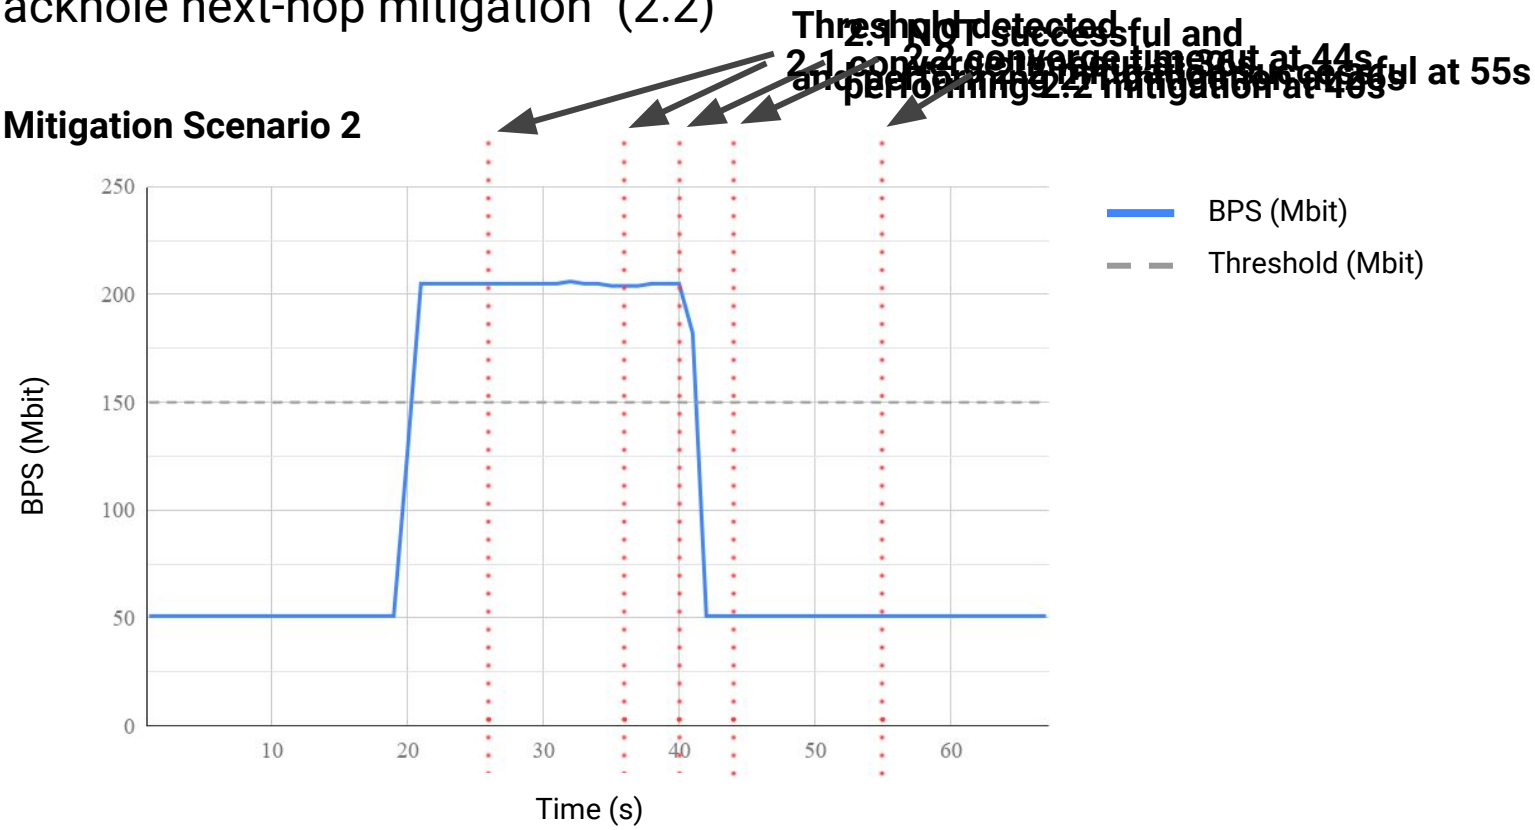

Proof of Concept cont'd

- Culprit AS is peered with RS
- BGP route withdrawal mitigation *unsuccessful* (2.1)
- BGP blackhole next-hop mitigation *unsuccessful* (2.2)
- Ingress PE L2 ACL mitigation (2.3)

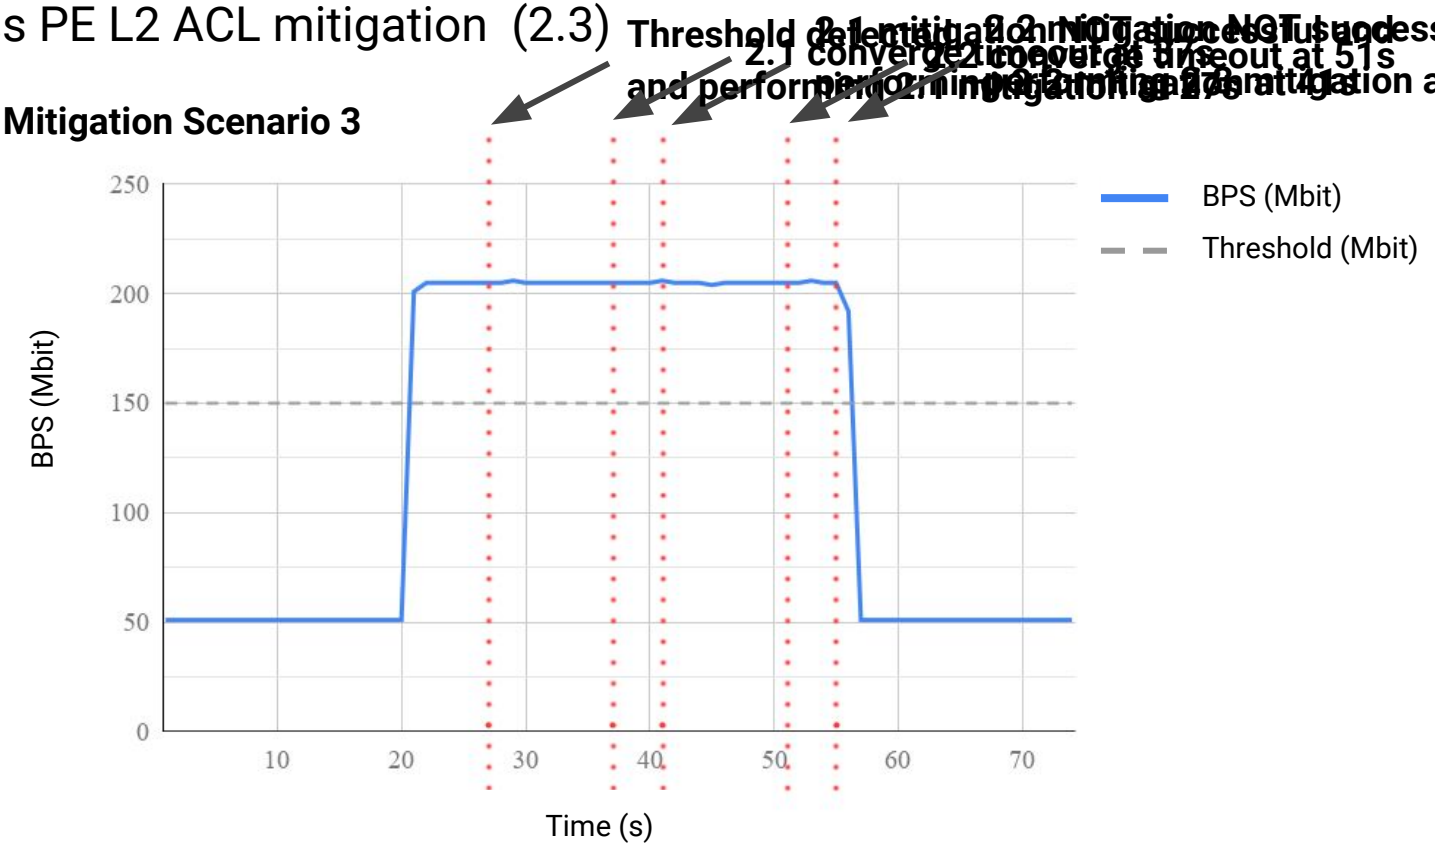

Proof of Concept cont'd

- Culprit AS is NOT peered with RS
- Ingress PE L2 ACL mitigation (2.3)

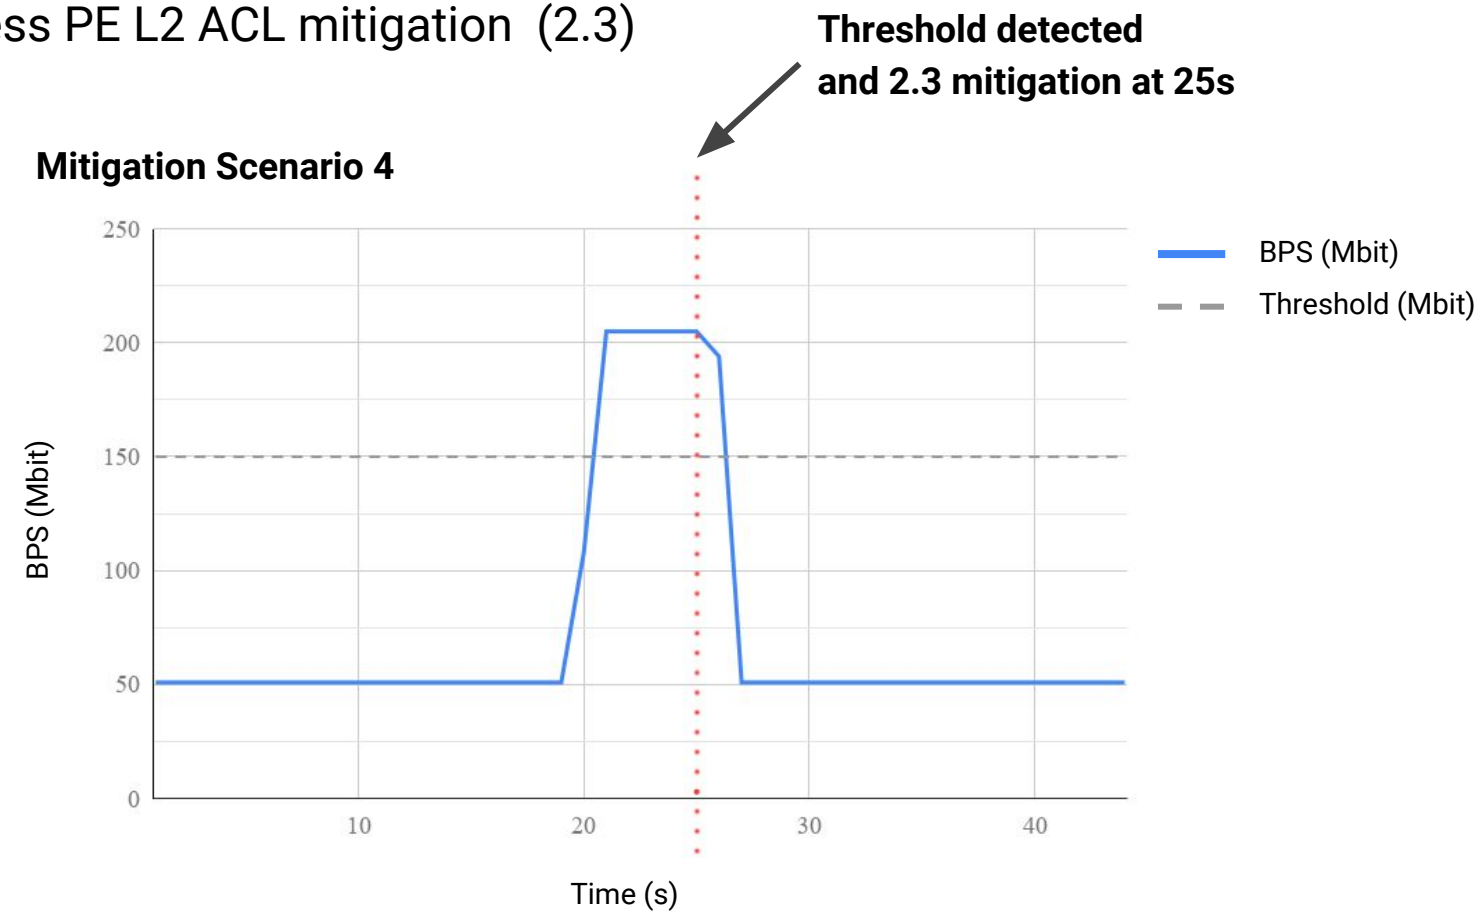

Discussion

- Usage of route server and statistics collector
- BGP convergence time (too long?)
- Layer 3 ACL
  - IXP environment: focus on layer 2 mitigation
- Fine-grained thresholds (time of day)
- Present more details to customer

Conclusion

- Thresholds and Three-way mitigation
- Identification requires layer 3 analysis (prefixes)
- Mitigation achieved on layer 2
  - BGP TE
  - IXP perspective

Future Work

- Different mitigations per type of attack
  - More advanced threshold metrics
- Testing with different sample rates
- Test scalability of the design
- Expand proof of concept
  - Identification phase
- Other methods of identification
  - Unsupervised/supervised learning

Questions

Image source - thenounproject.com

References

[1] ABN AMRO Group. Service temporarily disrupted by DDoS attacks (Jan 2018). Available at https://www.abnamro.com/en/newsroom/newsarticles/2018/service-temporarily-disrupted-by-ddos-attacks.html (Accessed on 01/06/2018)

[2] Cyberscoop. Arbor: DDoS attacks growing faster in size, complexity (Jan 2018). Available at https://www.cyberscoop.com/ddos-attacks-growing-arbor-networks/ (Accessed on 01/06/2018)

[3] DDoSDB. Collecting and Sharing the most important information of DDoS attacks. https://ddosdb.org/ (Accessed on 14/06/2018)