Accurately Measuring Global Risk of Amplification Attacks using AmpMap


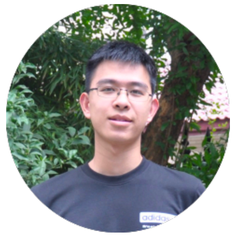



Soo-Jin Moon†
Yucheng Yin†, Rahul Anand Sharma†,
Yifei Yuan§, Jonathan M. Spring‡, Vyas Sekar†

†Carnegie Mellon University, §Alibaba Group,
‡CERT/CC®, SEI, Carnegie Mellon University

Carnegie Mellon University
CyLab
Security and Privacy Institute

Problem of DDoS Amplification Attacks

1.7 Tbps DDoS Attack — Memcached UDP Reflections Set New Record

Hackernews (Mar 2018)

Problem of DDoS Amplification Attacks

1.7 Tbps DDoS Attack — Memcached UDP Reflections Set New Record

Hackernews (Mar 2018)

SSDP amplification attacks rose 639%

Help Net Security (Jan 2019)

Problem of DDoS Amplification Attacks

1.7 Tbps DDoS Attack — Memcached UDP Reflections Set New Record

Hackernews (Mar 2018)

SSDP amplification attacks rose 639%

Help Net Security (Jan 2019)

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.

ZDNet (June 2020)

Primer on DDoS Amplification Attacks

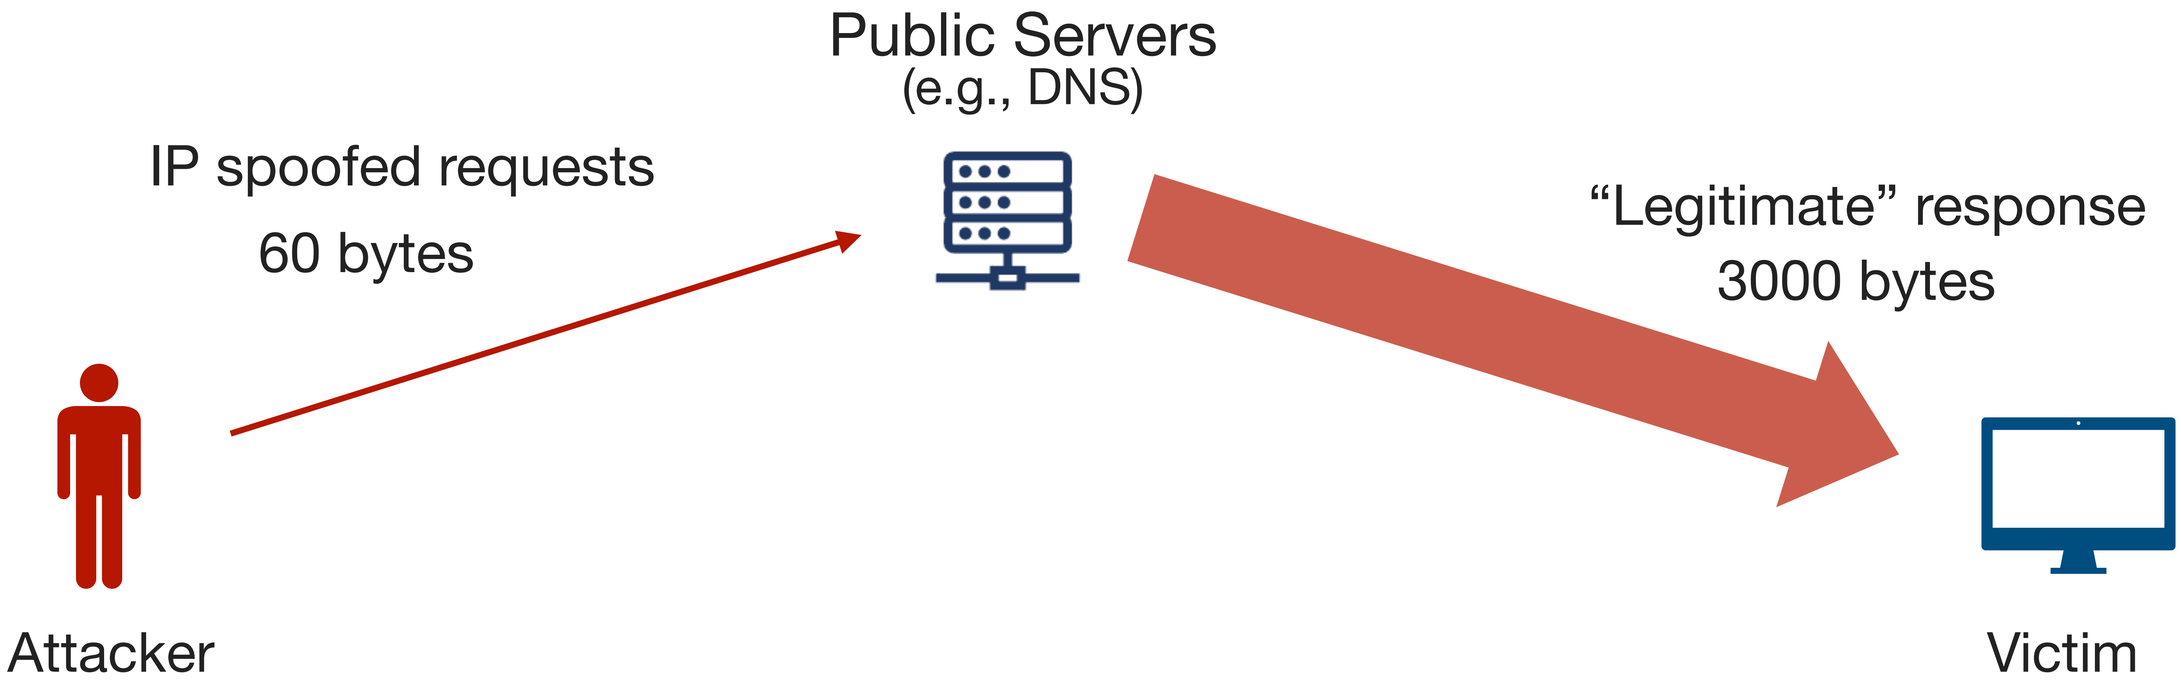

Primer on DDoS Amplification Attacks

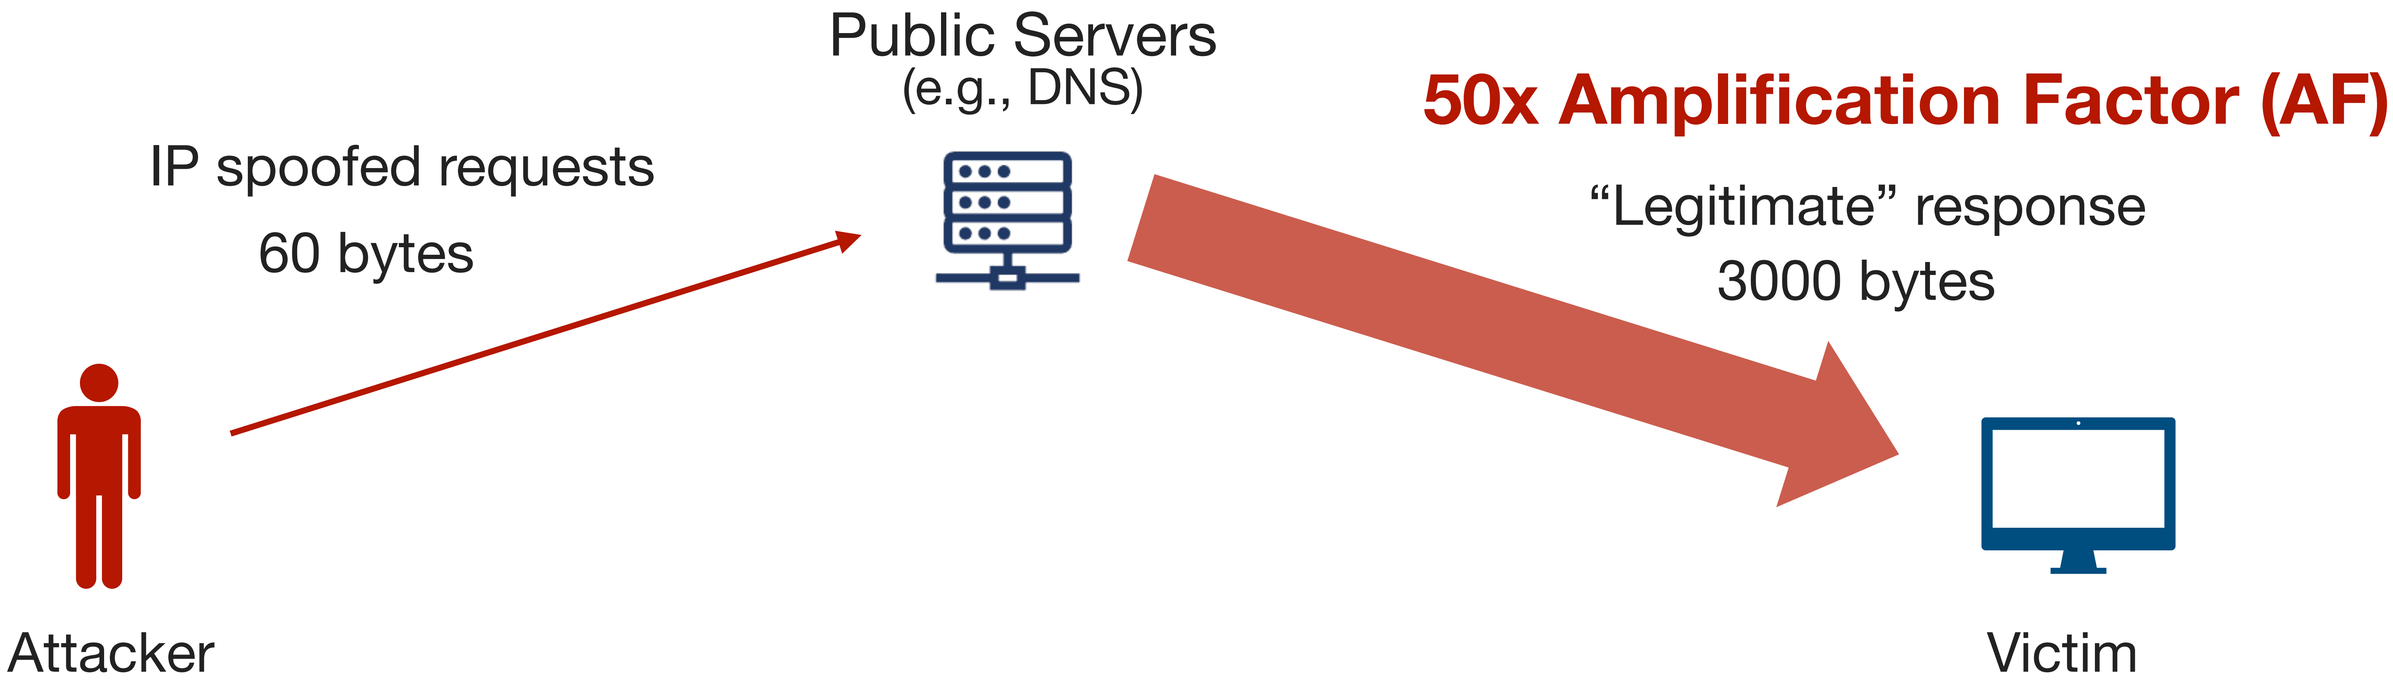

Primer on DDoS Amplification Attacks

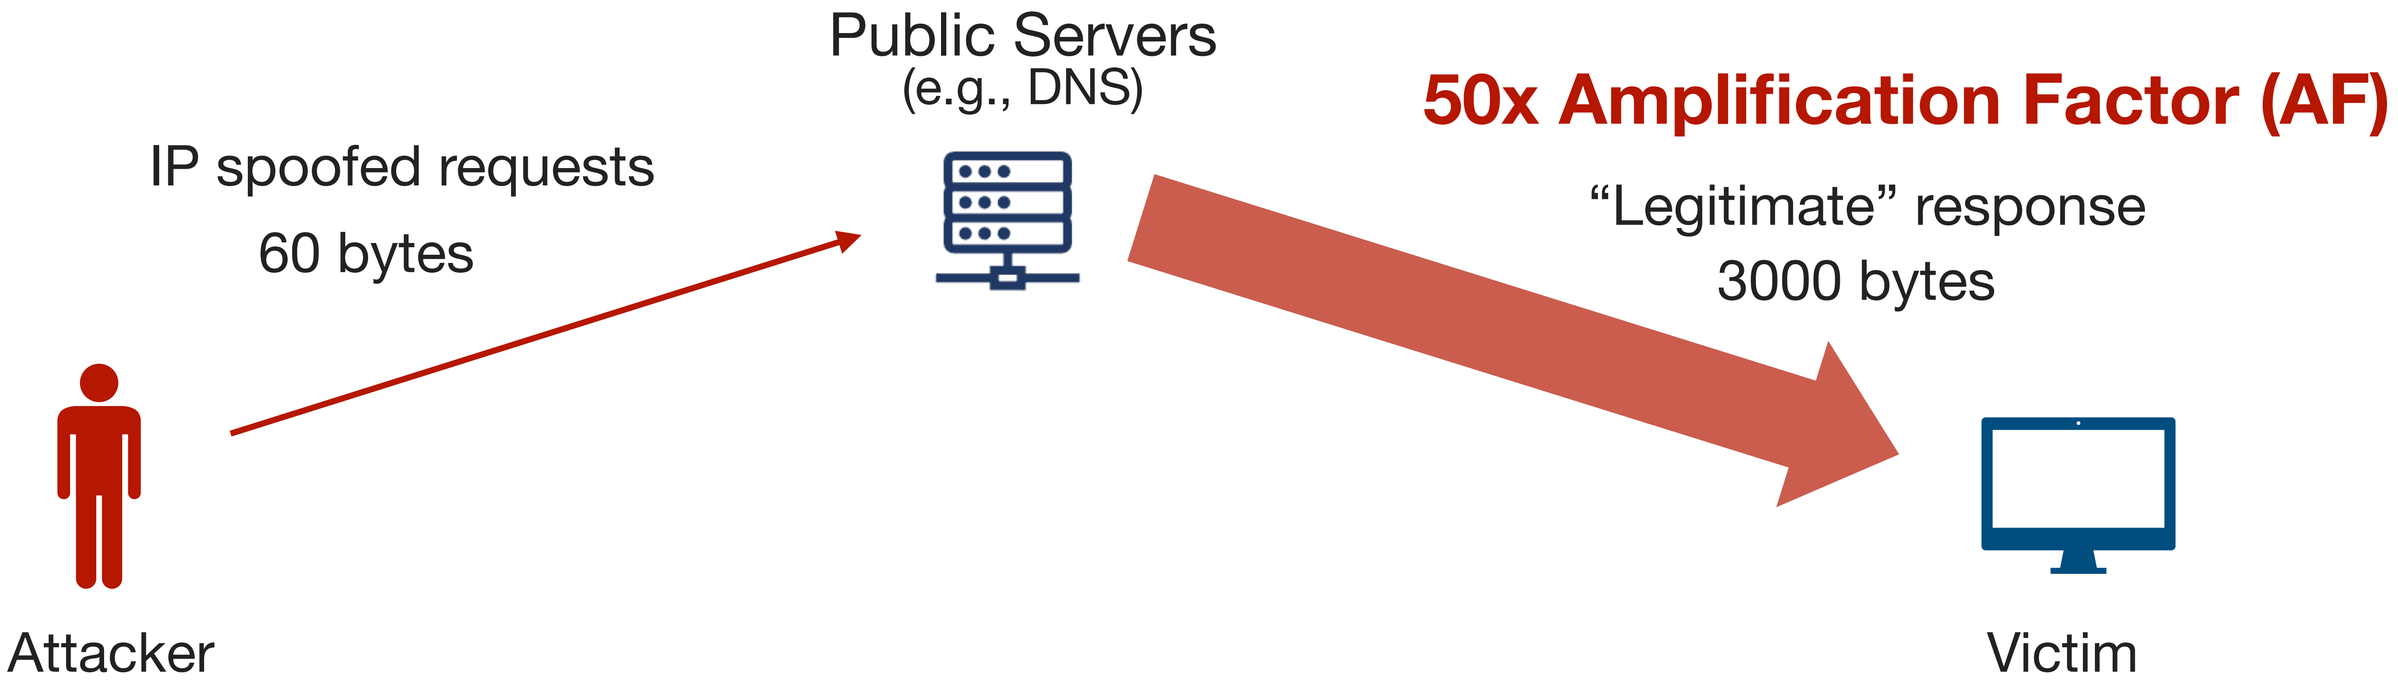

An example of an amplification mode for DNS:

- EDNS: 0
- Record type: ANY (255)
- EDNS maximum payload: > 4000

Primer on DDoS Amplification Attacks

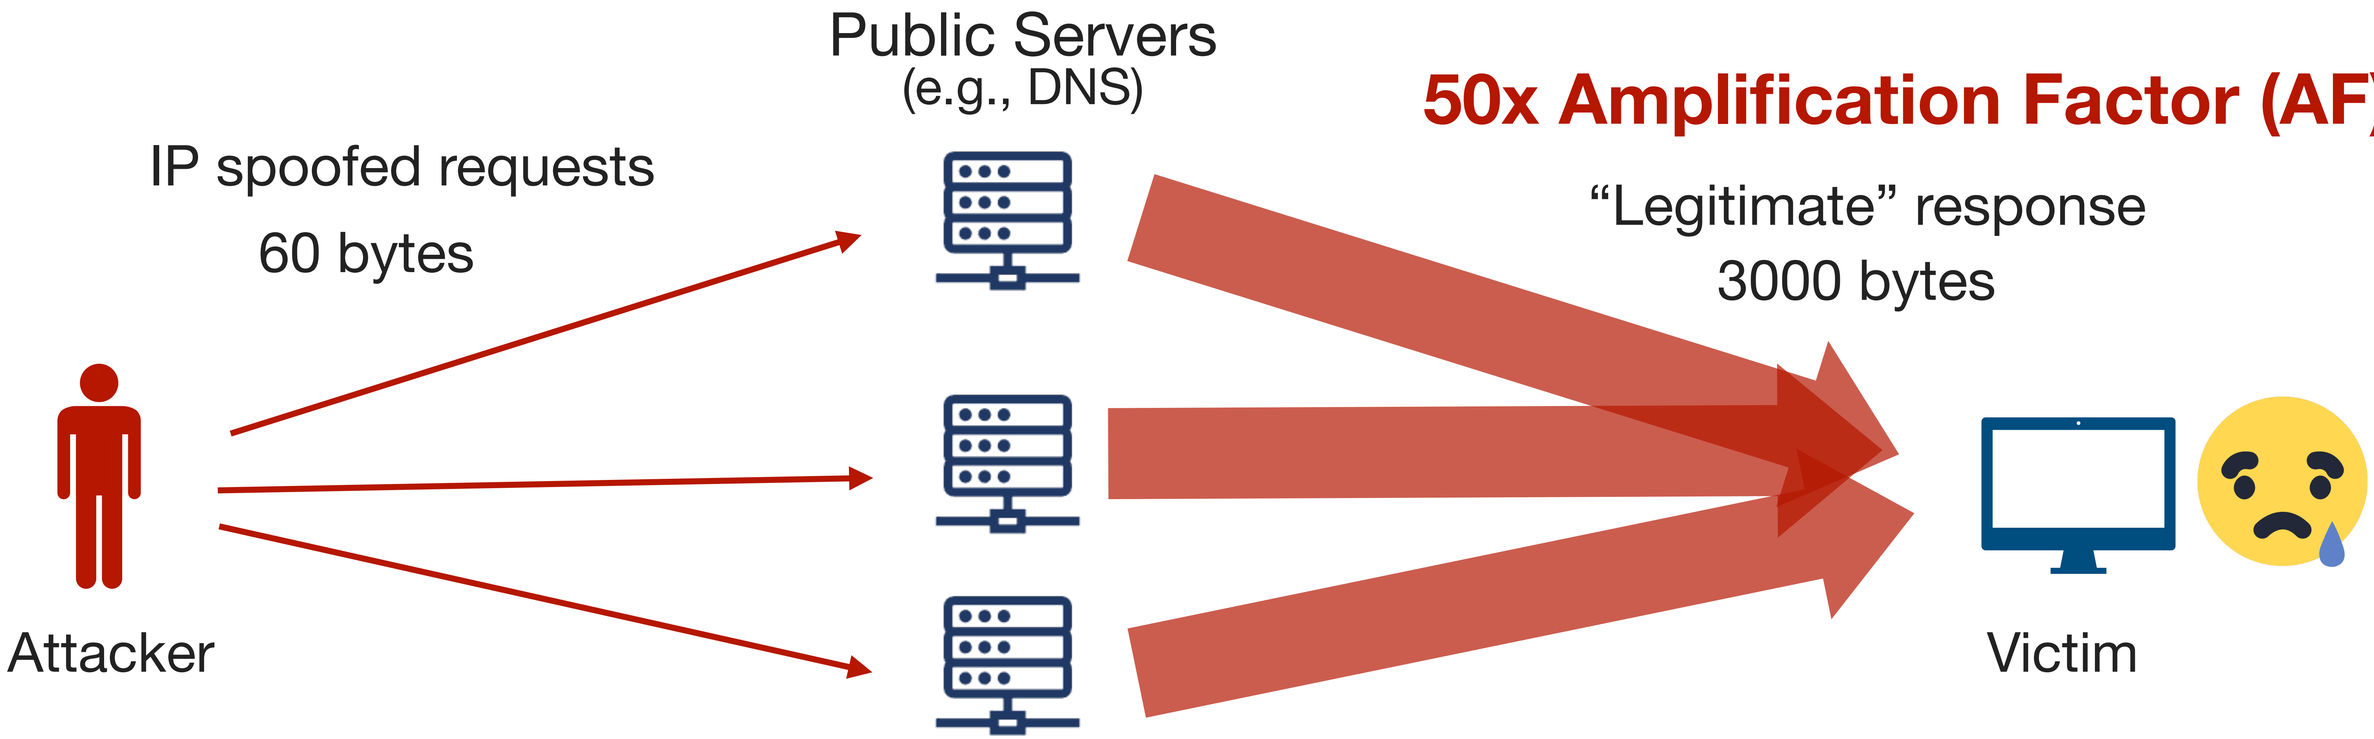

An example of an amplification mode for DNS:

- EDNS: 0
- Record type: ANY (255)
- EDNS maximum payload: > 4000

What We Need: Amplification Monitoring Service

Which modes (query patterns) induce high amplification?

How much amplification does each mode induce?

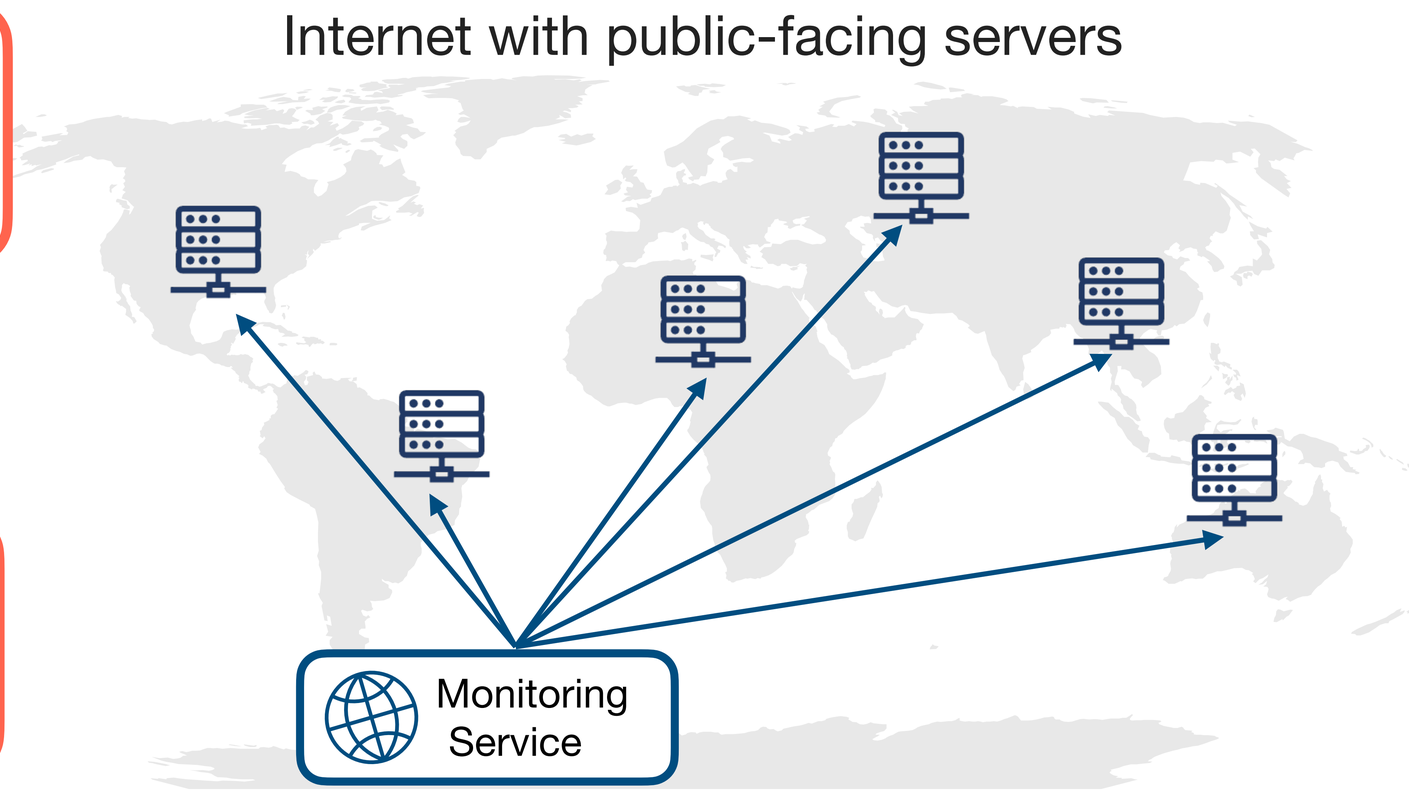

Strawman Solutions: Inaccurate or Incurs High Overhead!

- Count # servers & scale by a constant factor from prior work (e.g., Cybergreen[1])

[1] https://stats.cybergreen.net/

Strawman Solutions: Inaccurate or Incurs High Overhead!

- Count # servers & scale by a constant factor from prior work (e.g., Cybergreen[1])

  Do not account for server heterogeneity ✗

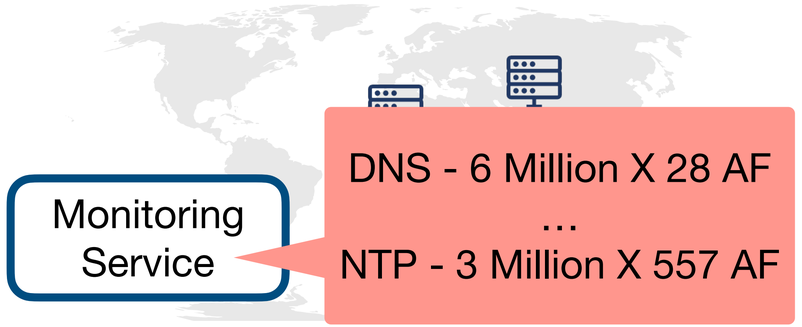

[1] https://stats.cybergreen.net/

Strawman Solutions: Inaccurate or Incurs High Overhead!

- Count # servers & scale by a constant factor from prior work (e.g., Cybergreen[1])

  Do not account for server heterogeneity ✗

- Identify one (or handful) amplification modes (e.g., [2])

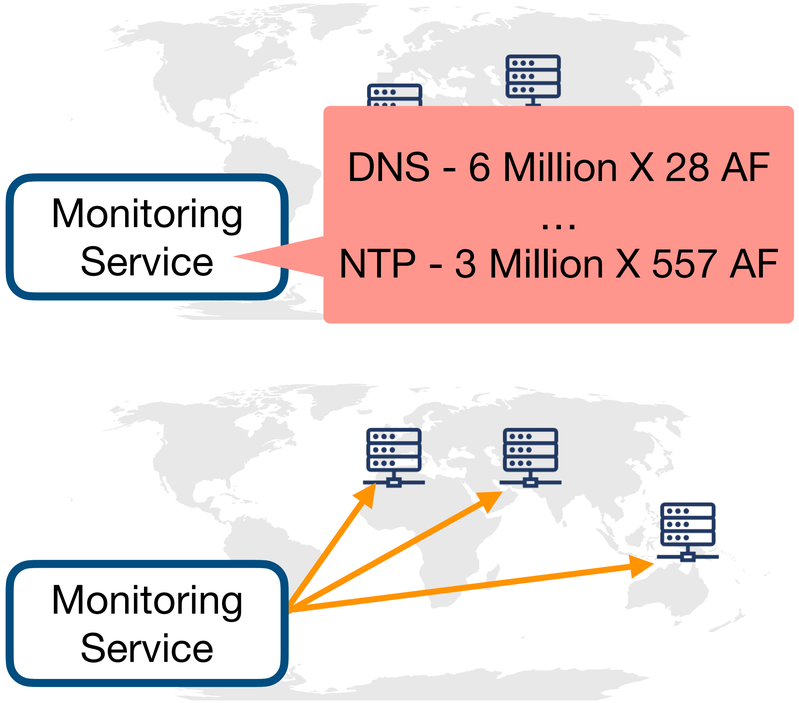

[1] https://stats.cybergreen.net/ [2] Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In Proc. NDSS 2014

Strawman Solutions: Inaccurate or Incurs High Overhead!

- Count # servers & scale by a constant factor from prior work (e.g., Cybergreen[1])

  Do not account for server heterogeneity ✗

- Identify one (or handful) amplification modes (e.g., [2])

  Lacks coverage across other modes ✗

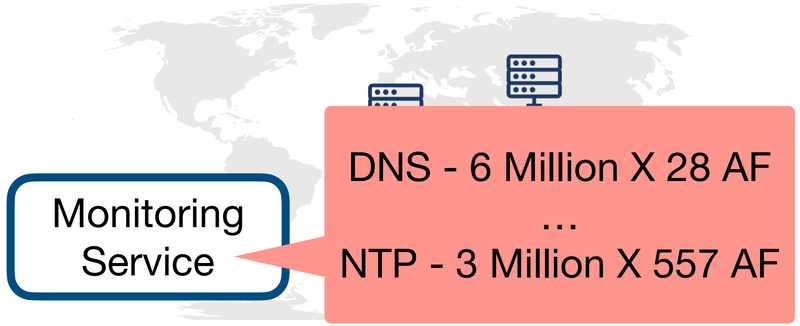

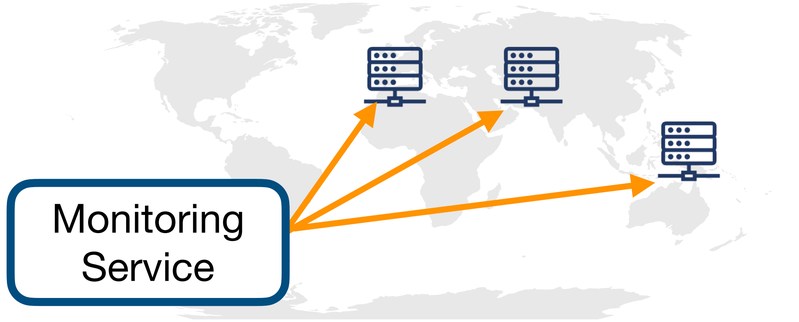

[1] https://stats.cybergreen.net/ [2] Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In Proc. NDSS 2014

Strawman Solutions: Inaccurate or Incurs High Overhead!

- Count # servers & scale by a constant factor from prior work (e.g., Cybergreen[1])

  Do not account for server heterogeneity ✗

- Identify one (or handful) amplification modes (e.g., [2])

  Lacks coverage across other modes ✗

- Brute-force query space for each server

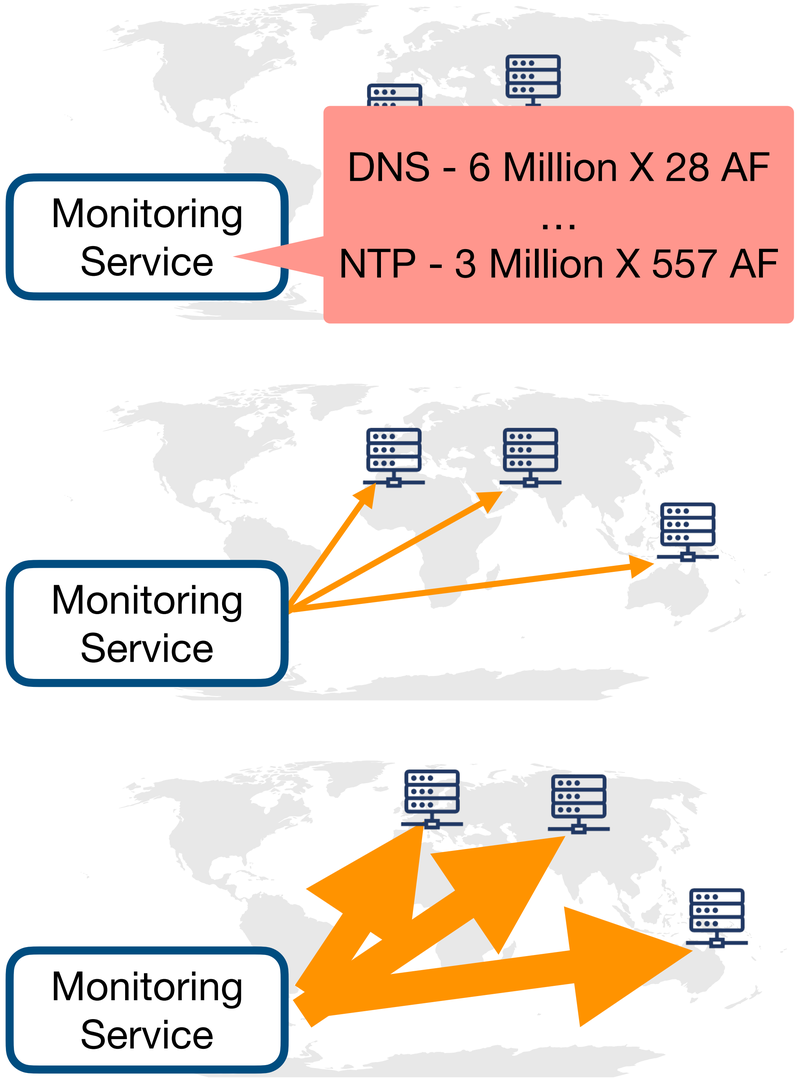

[1] https://stats.cybergreen.net/ [2] Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In Proc. NDSS 2014

Strawman Solutions: Inaccurate or Incurs High Overhead!

- Count # servers & scale by a constant factor from prior work (e.g., Cybergreen[1])

  Do not account for server heterogeneity ✗

- Identify one (or handful) amplification modes (e.g., [2])

  Lacks coverage across other modes ✗

- Brute-force query space for each server

  Infeasible ✗

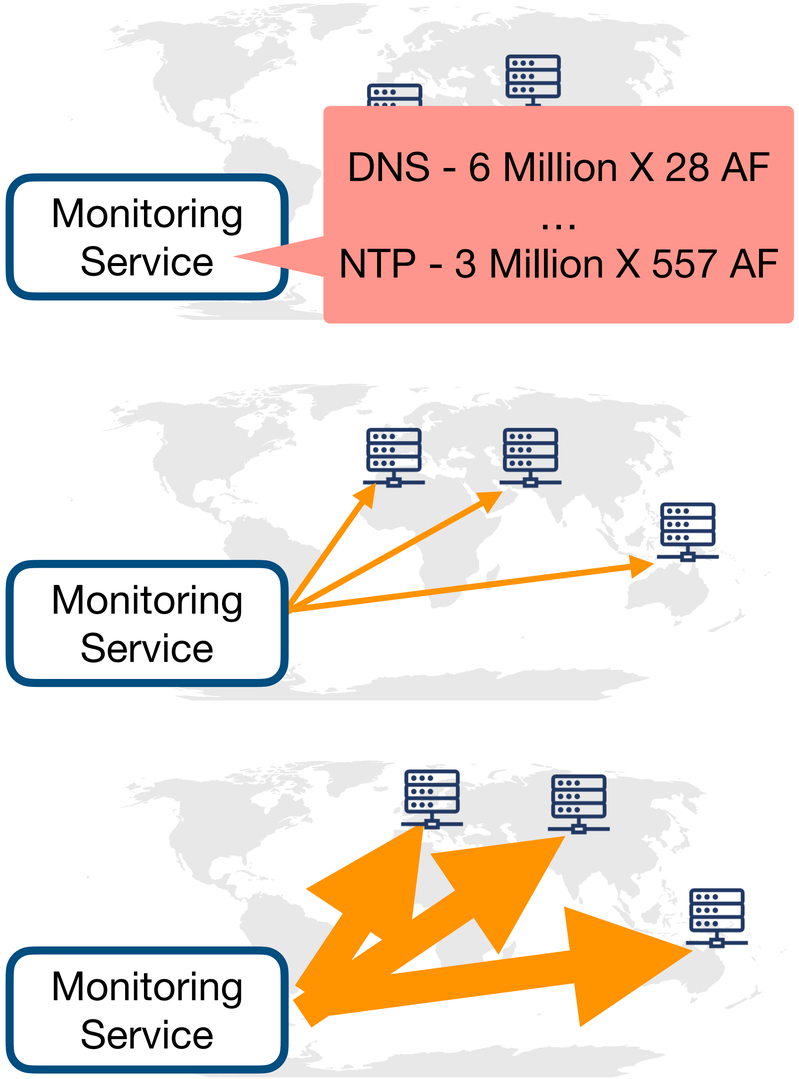

[1] https://stats.cybergreen.net/ [2] Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In Proc. NDSS 2014

Strawman Solutions: Inaccurate or Incurs High Overhead!

- Count # servers & scale by a constant factor from prior work (e.g., Cybergreen[1])

  Do not account for server heterogeneity ✗

- Identify one (or handful) amplification modes (e.g., [2])

  Lacks coverage across other modes ✗

- Brute-force query space for each server

  Infeasible ✗

↓ Accuracy & Coverage

↑ Overhead

[1] https://stats.cybergreen.net/ [2] Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In Proc. NDSS 2014

Motivating Question

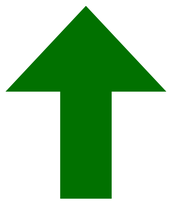

Accuracy &
Coverage

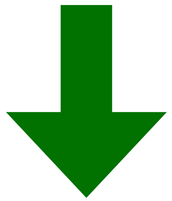

Overhead

Motivating Question

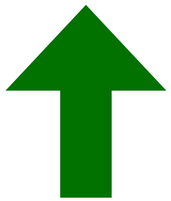

Accuracy & Coverage

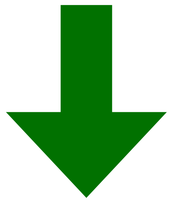

Overhead

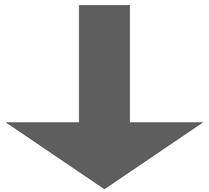

? Can we build an amplification monitoring service that achieves **high coverage** with **low network overhead**?

Practical Challenges & Dimensions to Consider

Building this service for a **single** server for a **single** protocol

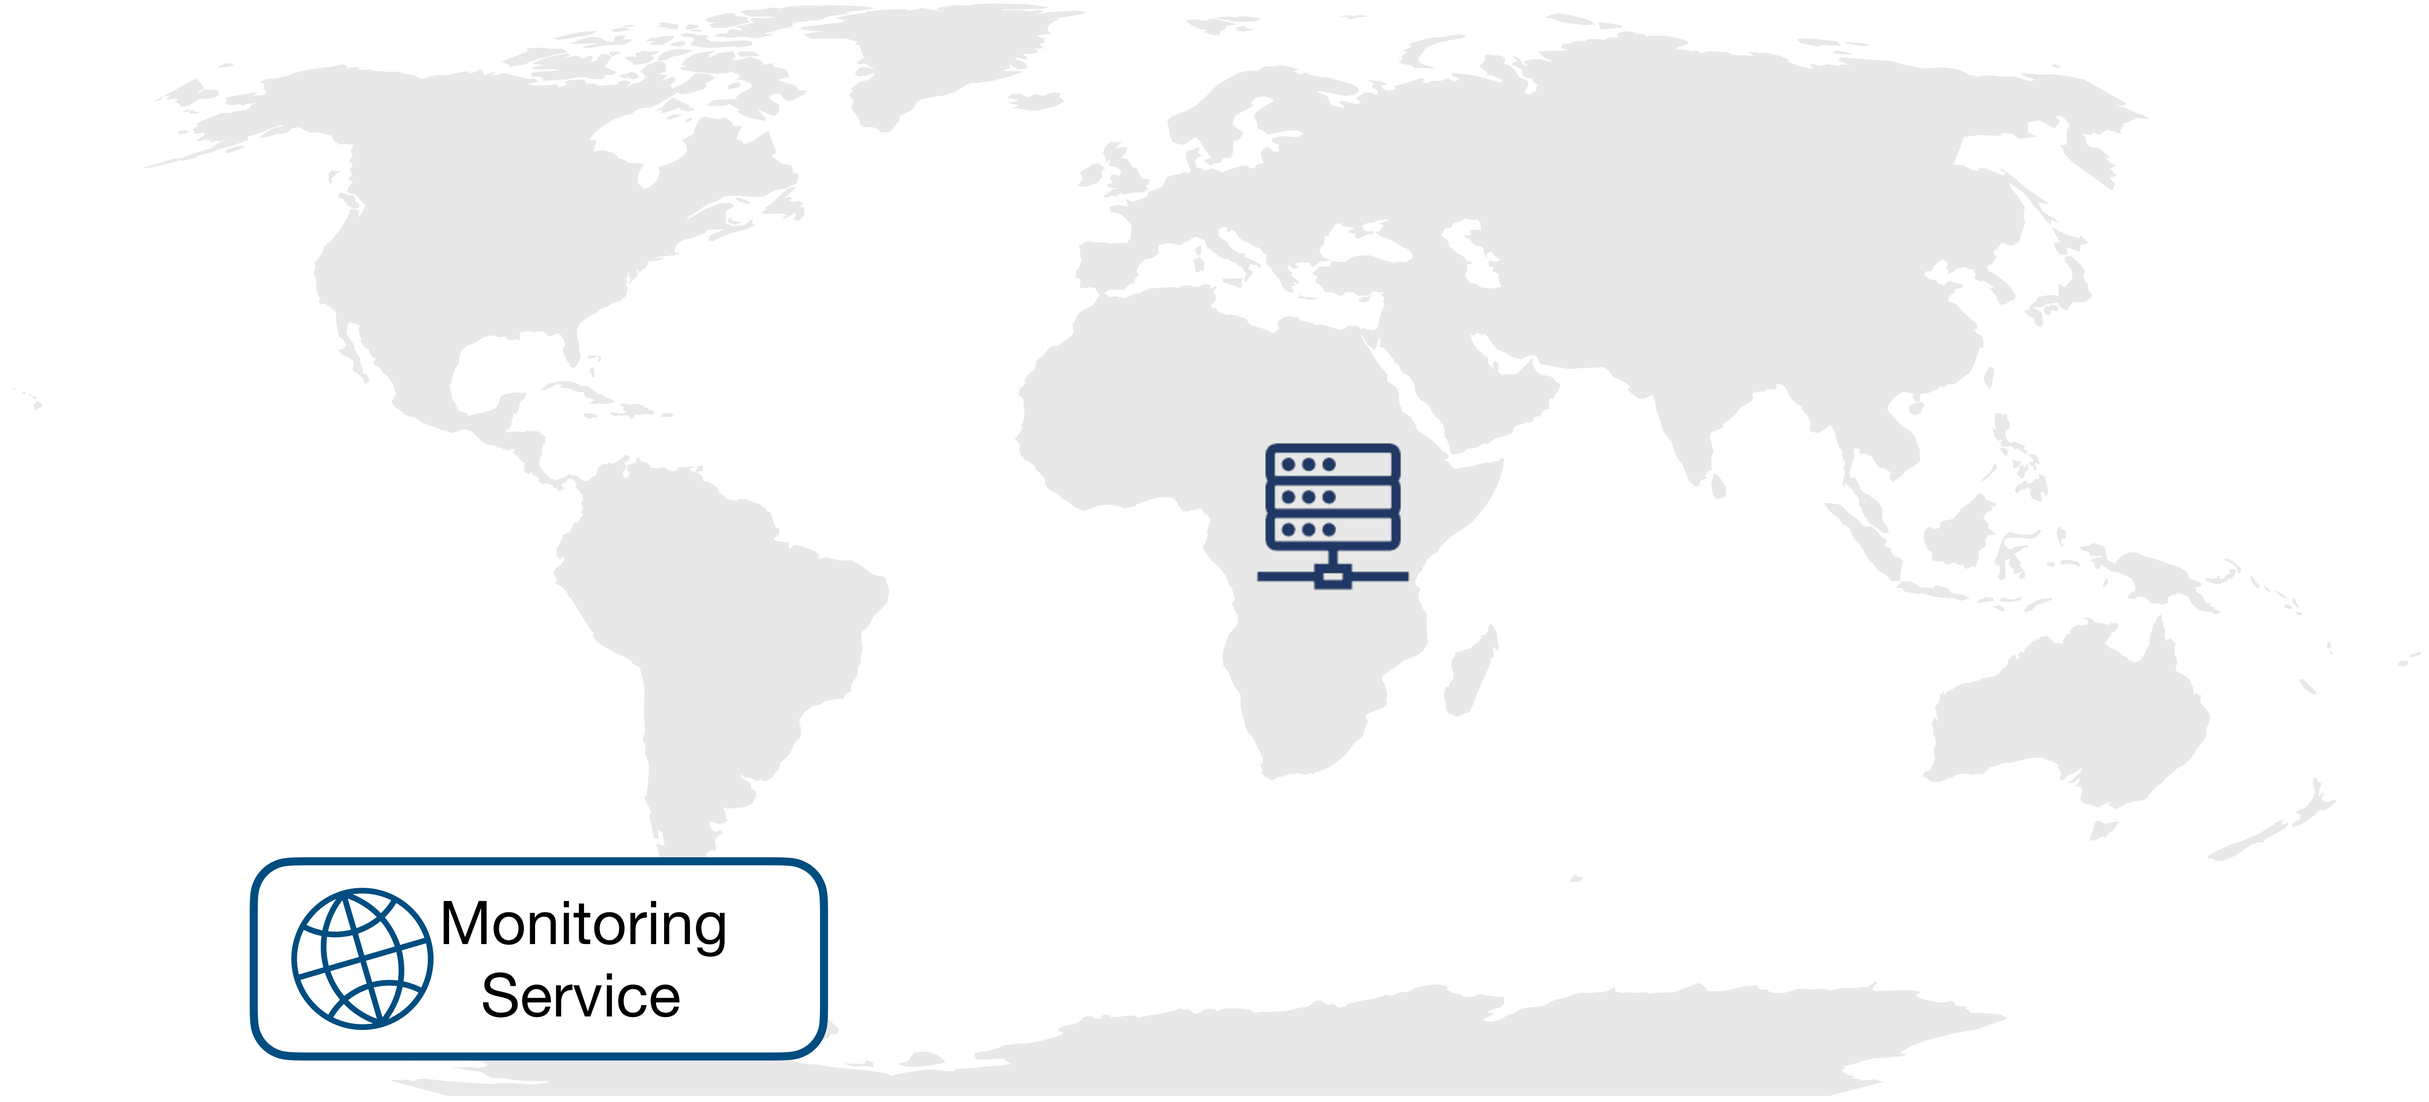

Monitoring Service

Practical Challenges & Dimensions to Consider

Building this service for a **single** server for a **single** protocol

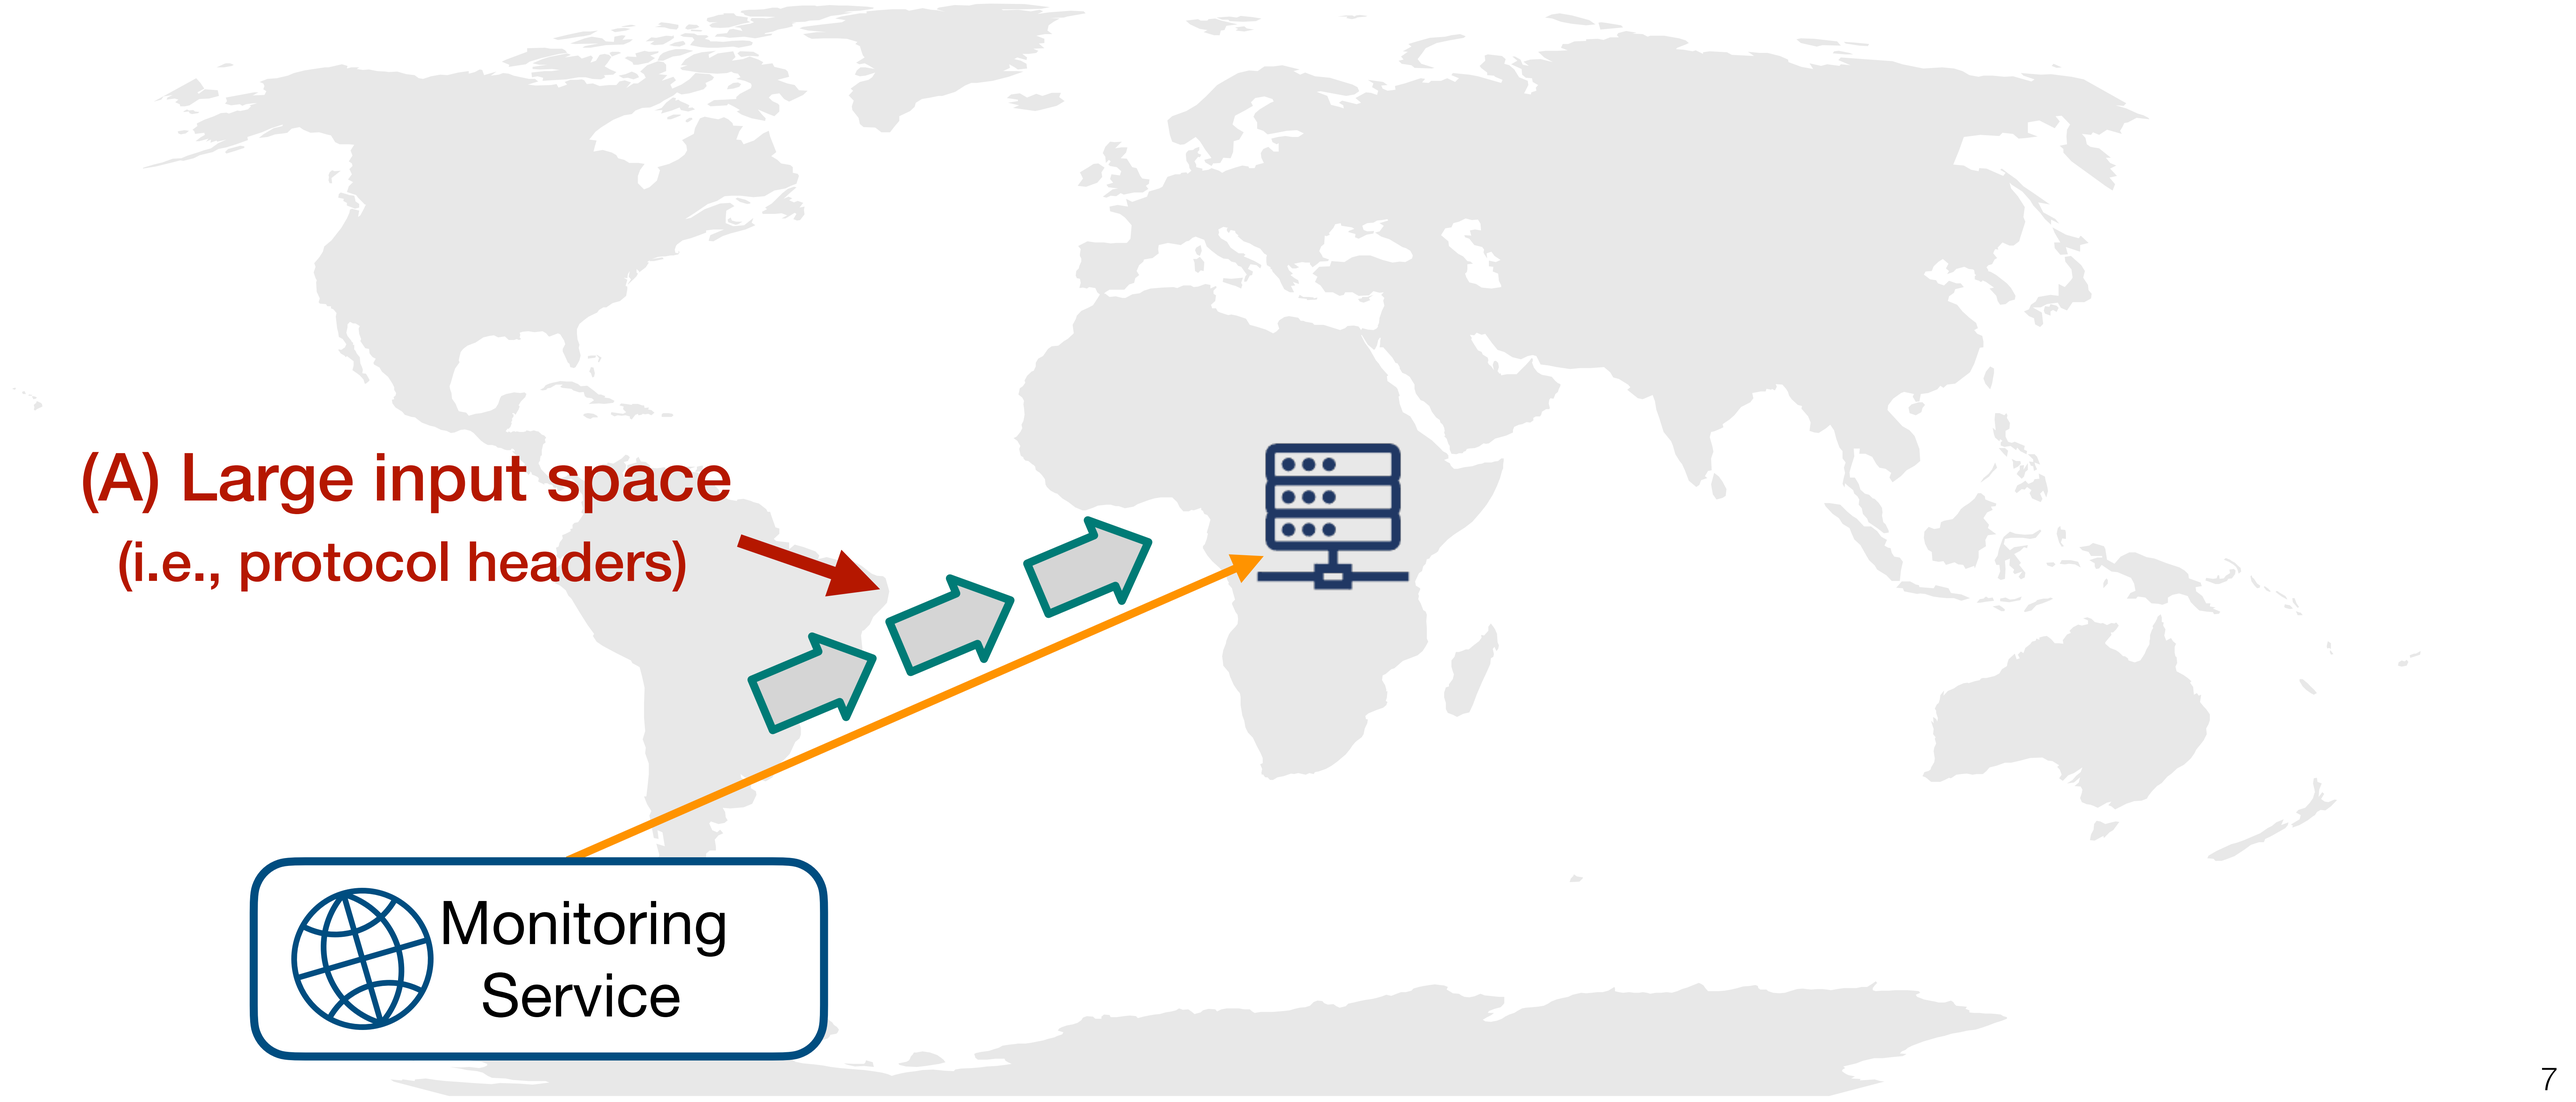

Practical Challenges & Dimensions to Consider

Building this service for a **single** server for a **single** protocol

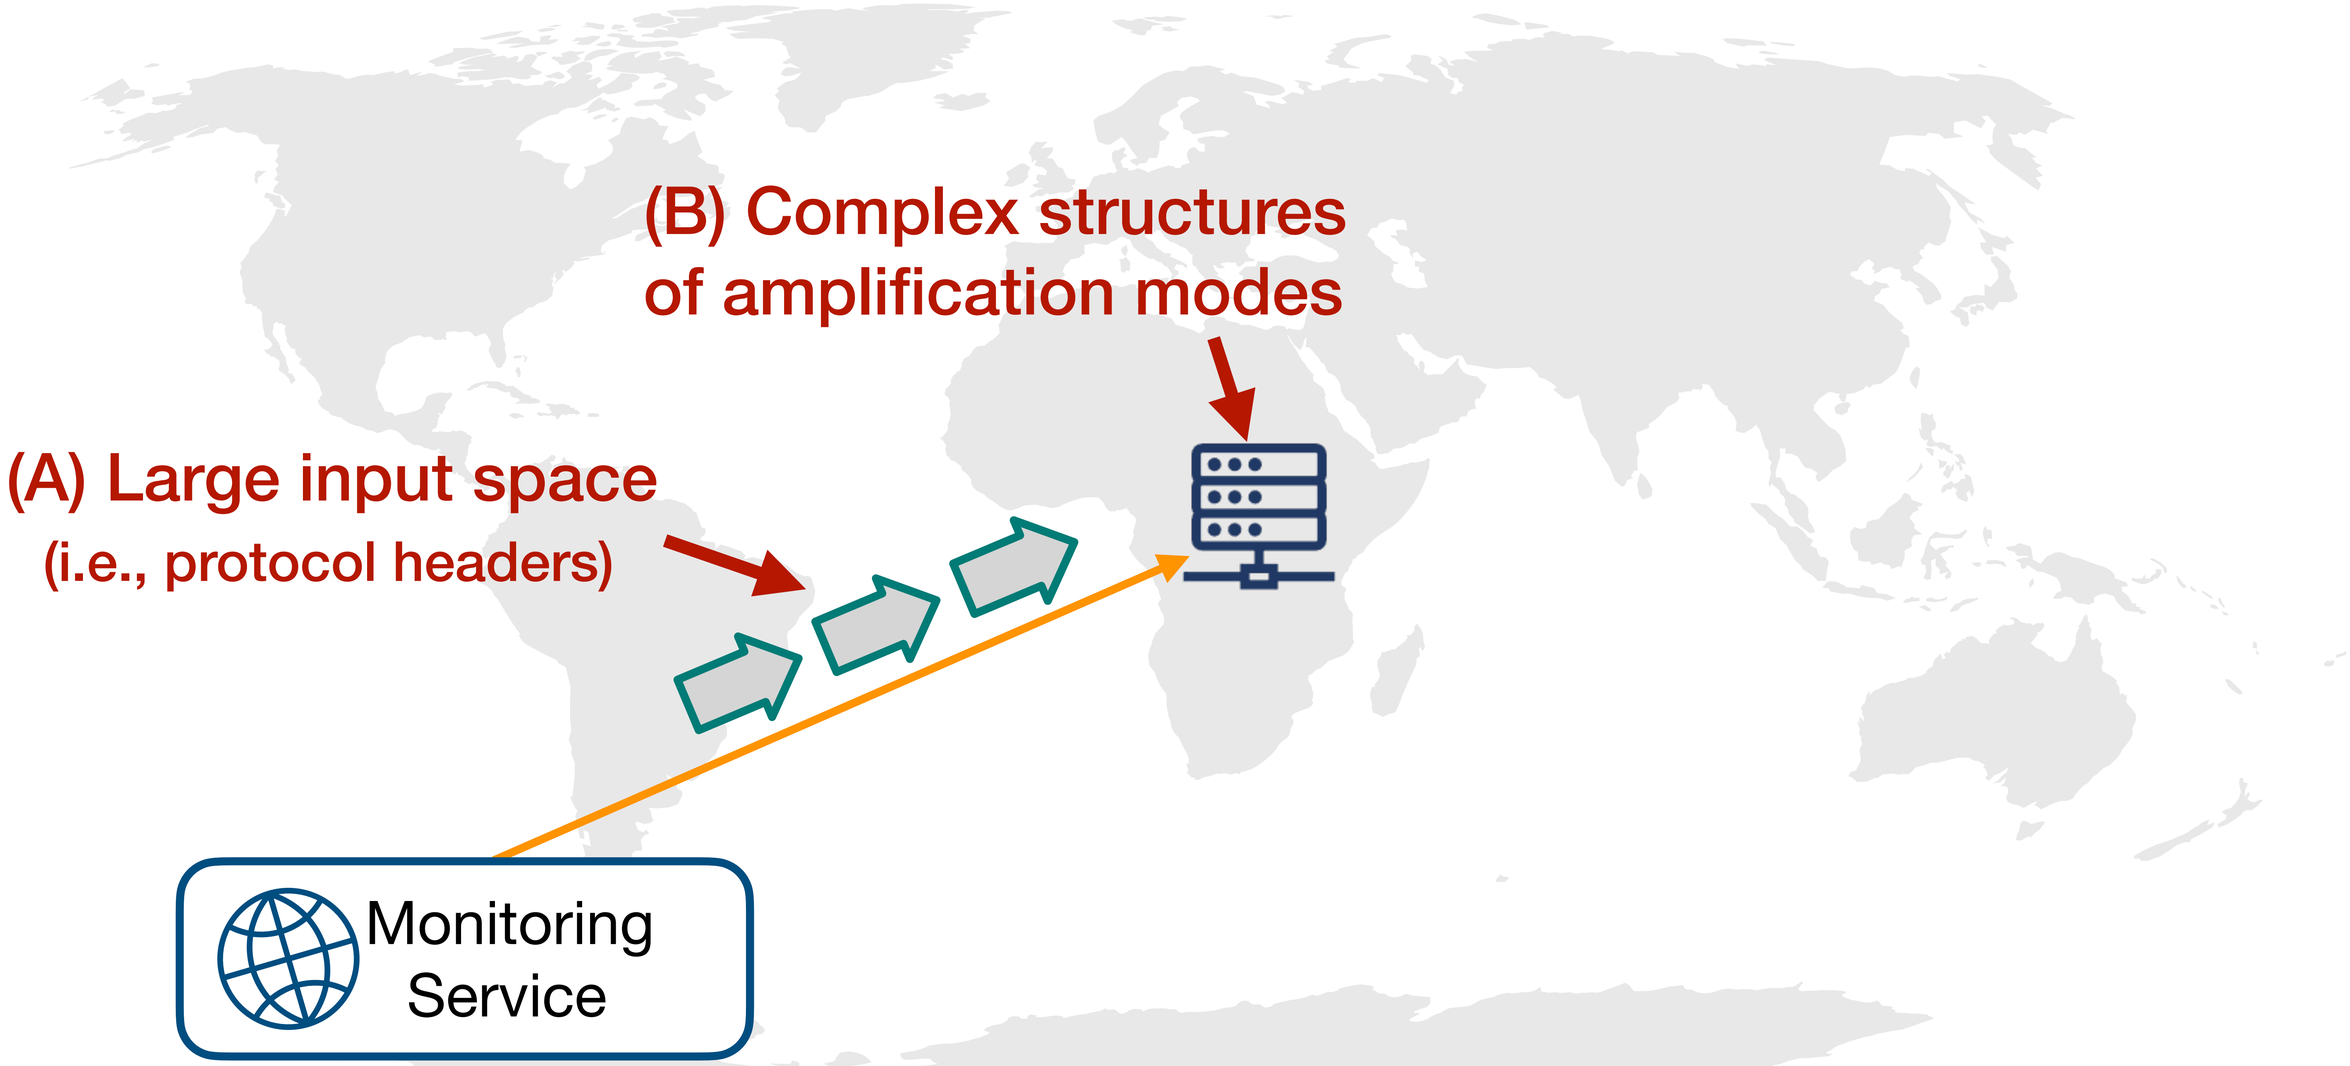

Practical Challenges & Dimensions to Consider

Building this service for ~~a **single** server~~ for a **single** protocol
multiple servers

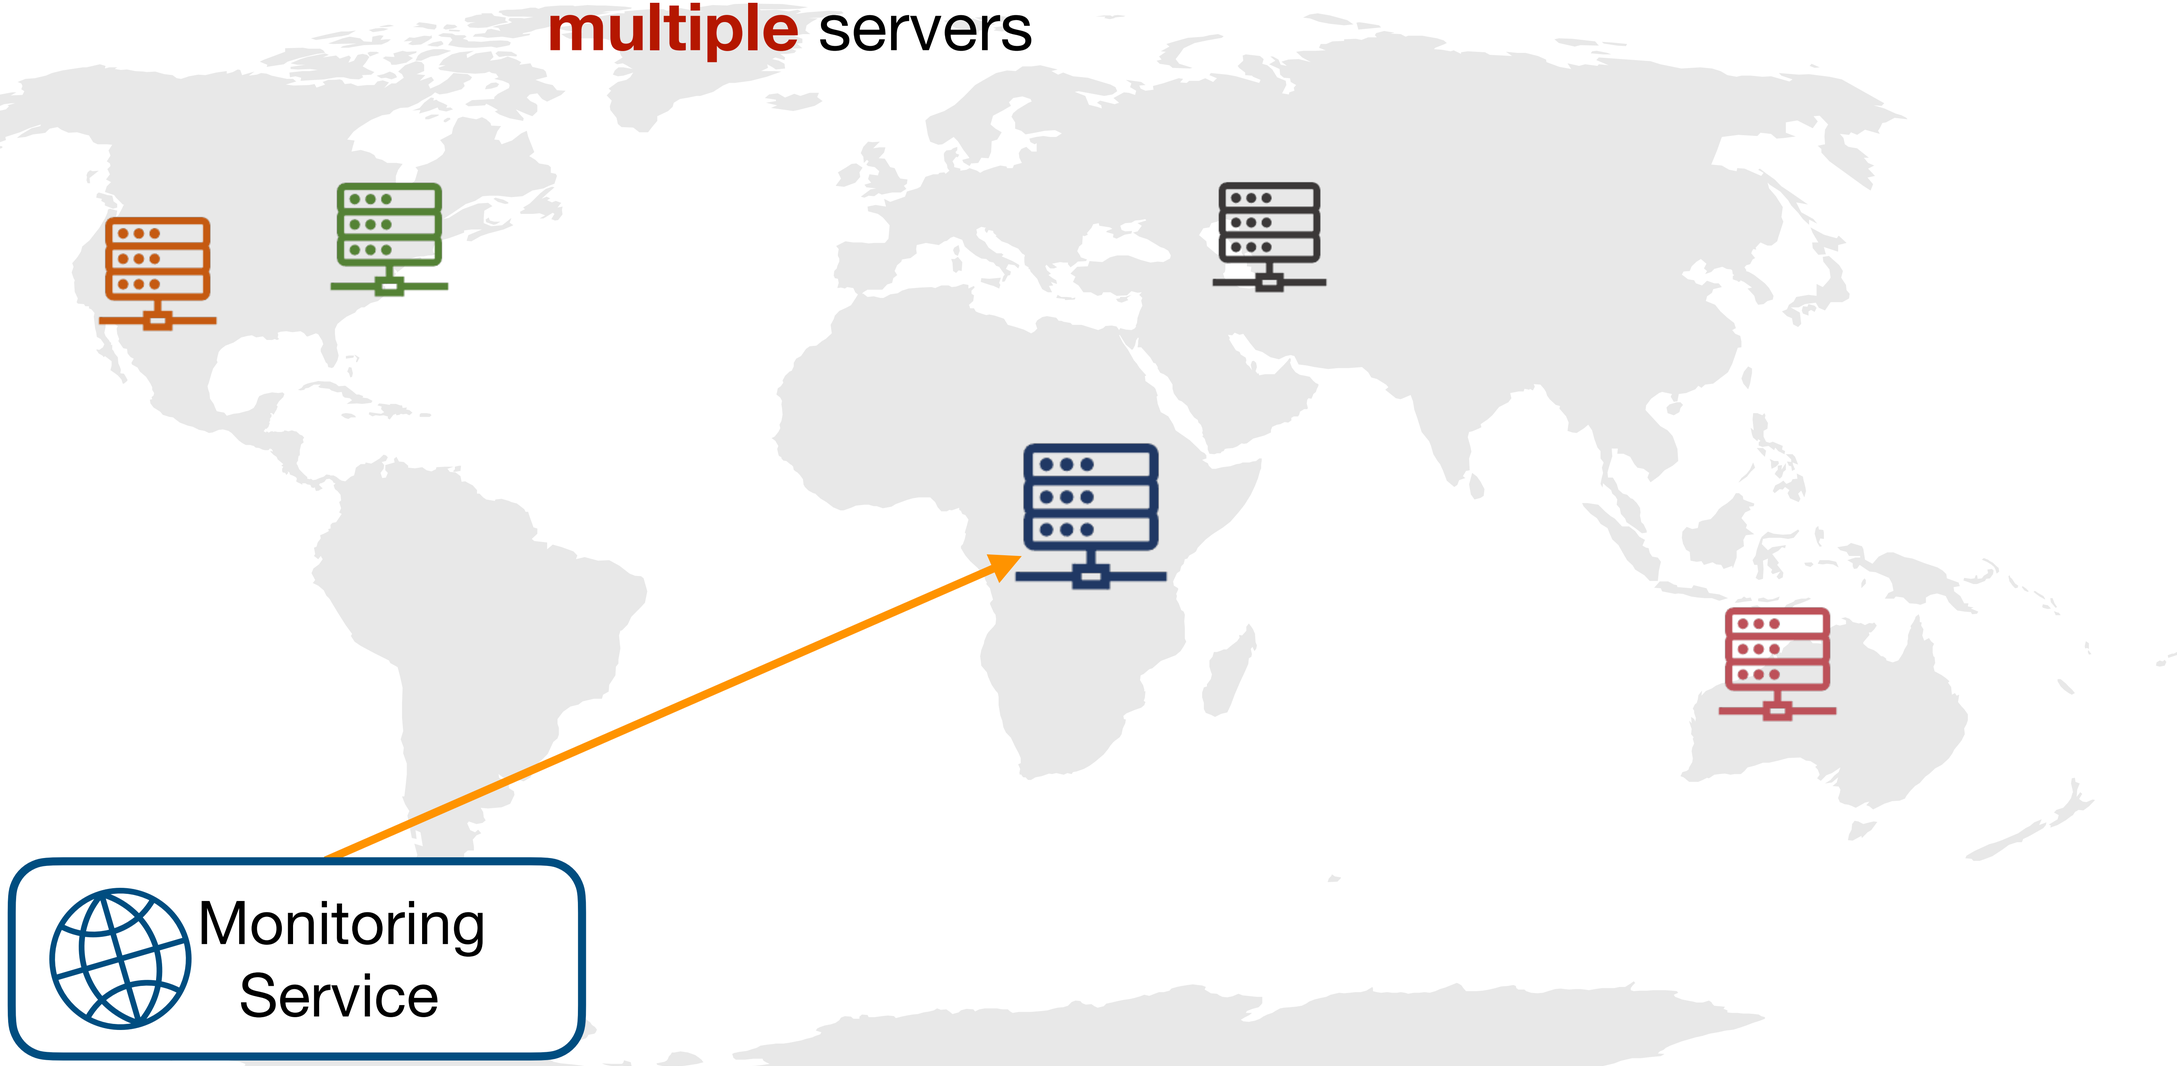

Practical Challenges & Dimensions to Consider

Building this service for ~~a **single** server~~ for a **single** protocol
multiple servers

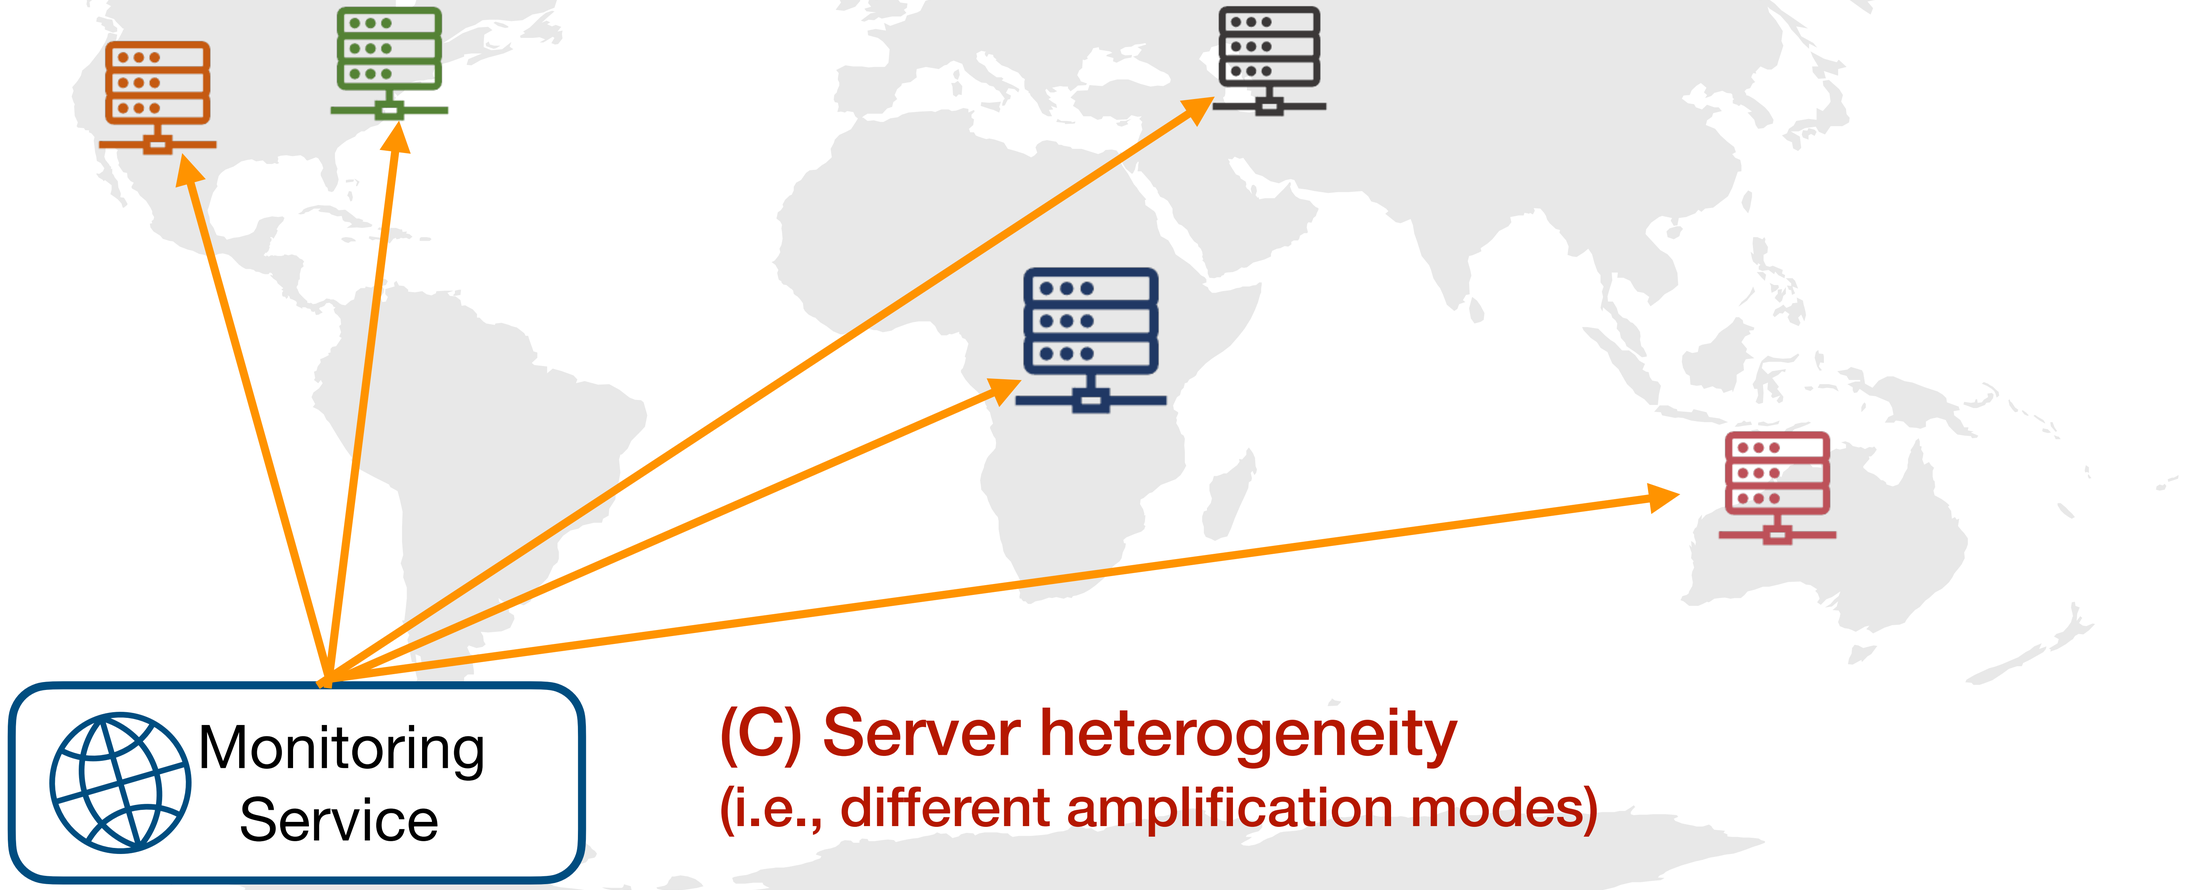

(C) Server heterogeneity
(i.e., different amplification modes)

Practical Challenges & Dimensions to Consider

Building this service for ~~a **single** server~~ for a ~~**single** protocol~~

multiple servers **multiple** protocols

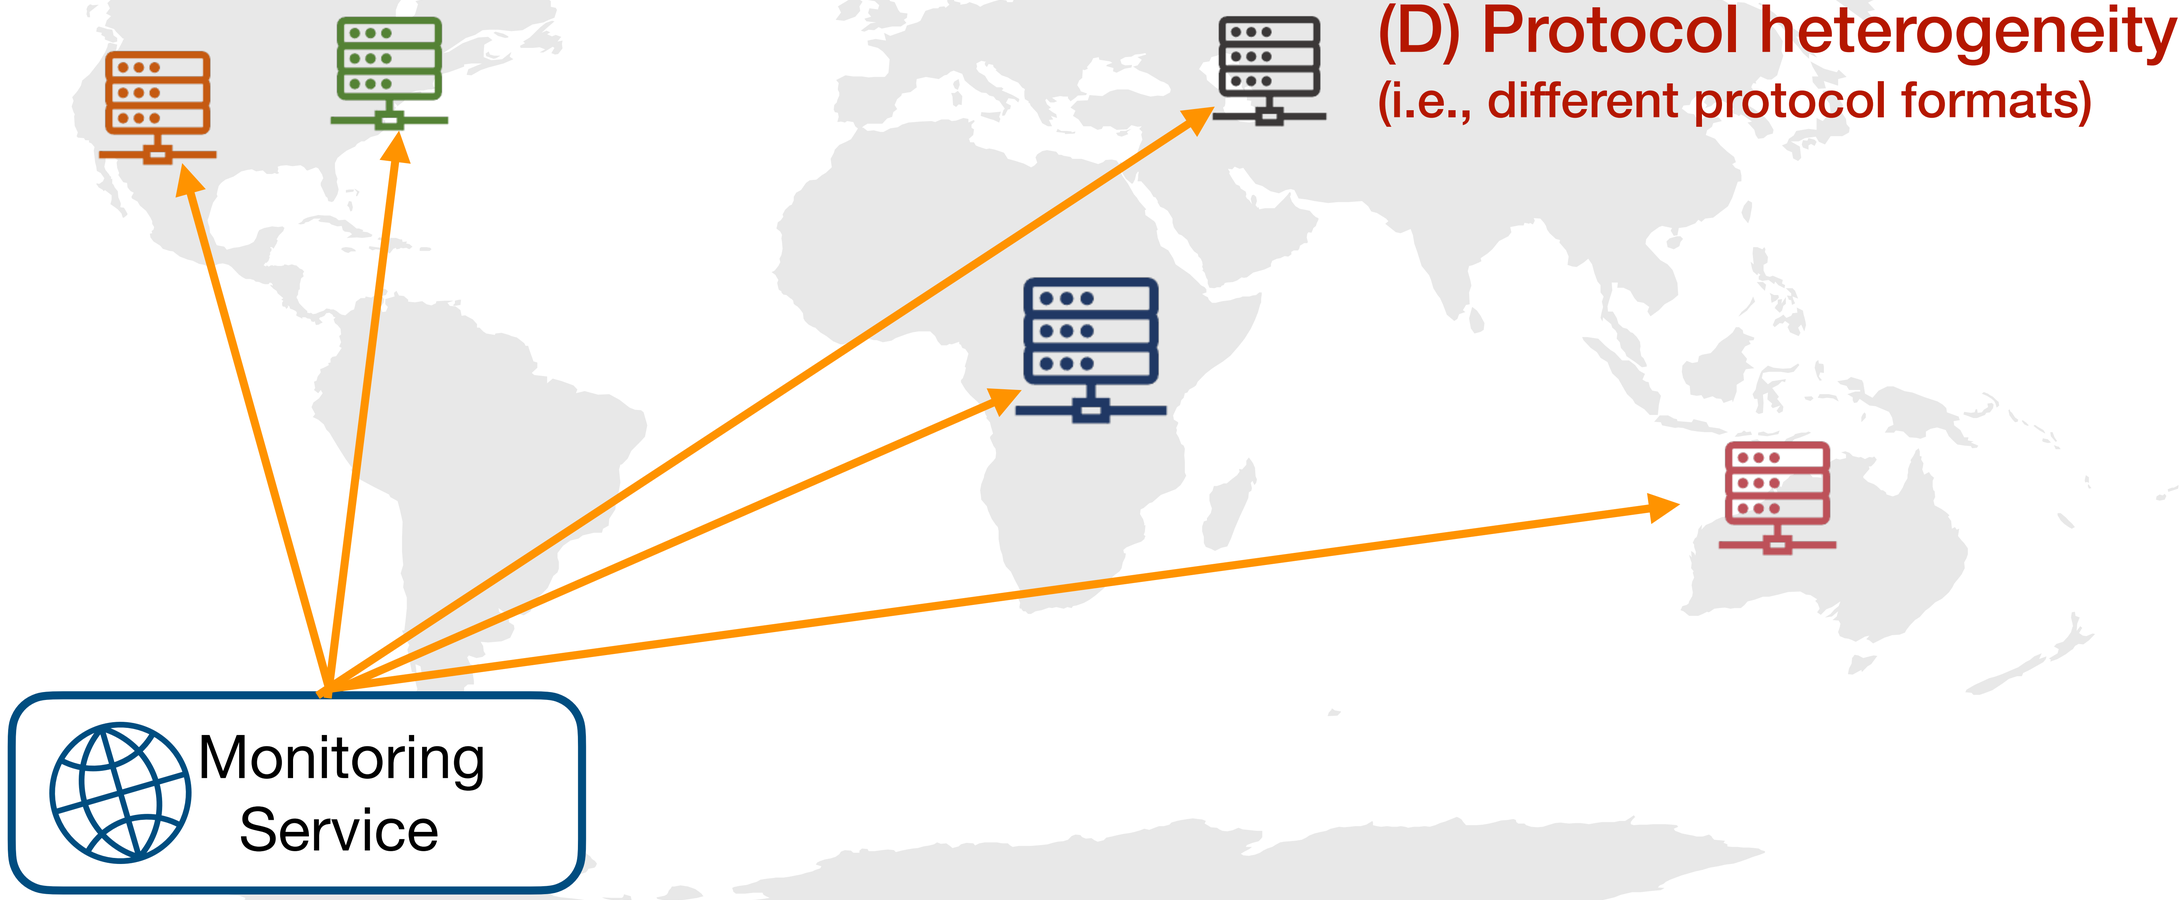

Our Work: AmpMap

A **low-footprint amplification monitoring service** to quantify risk

Our Work: AmpMap

A **low-footprint amplification monitoring service** to quantify risk

Leverages **structural properties** across packet header & server space to improve **coverage** with **low overhead**

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

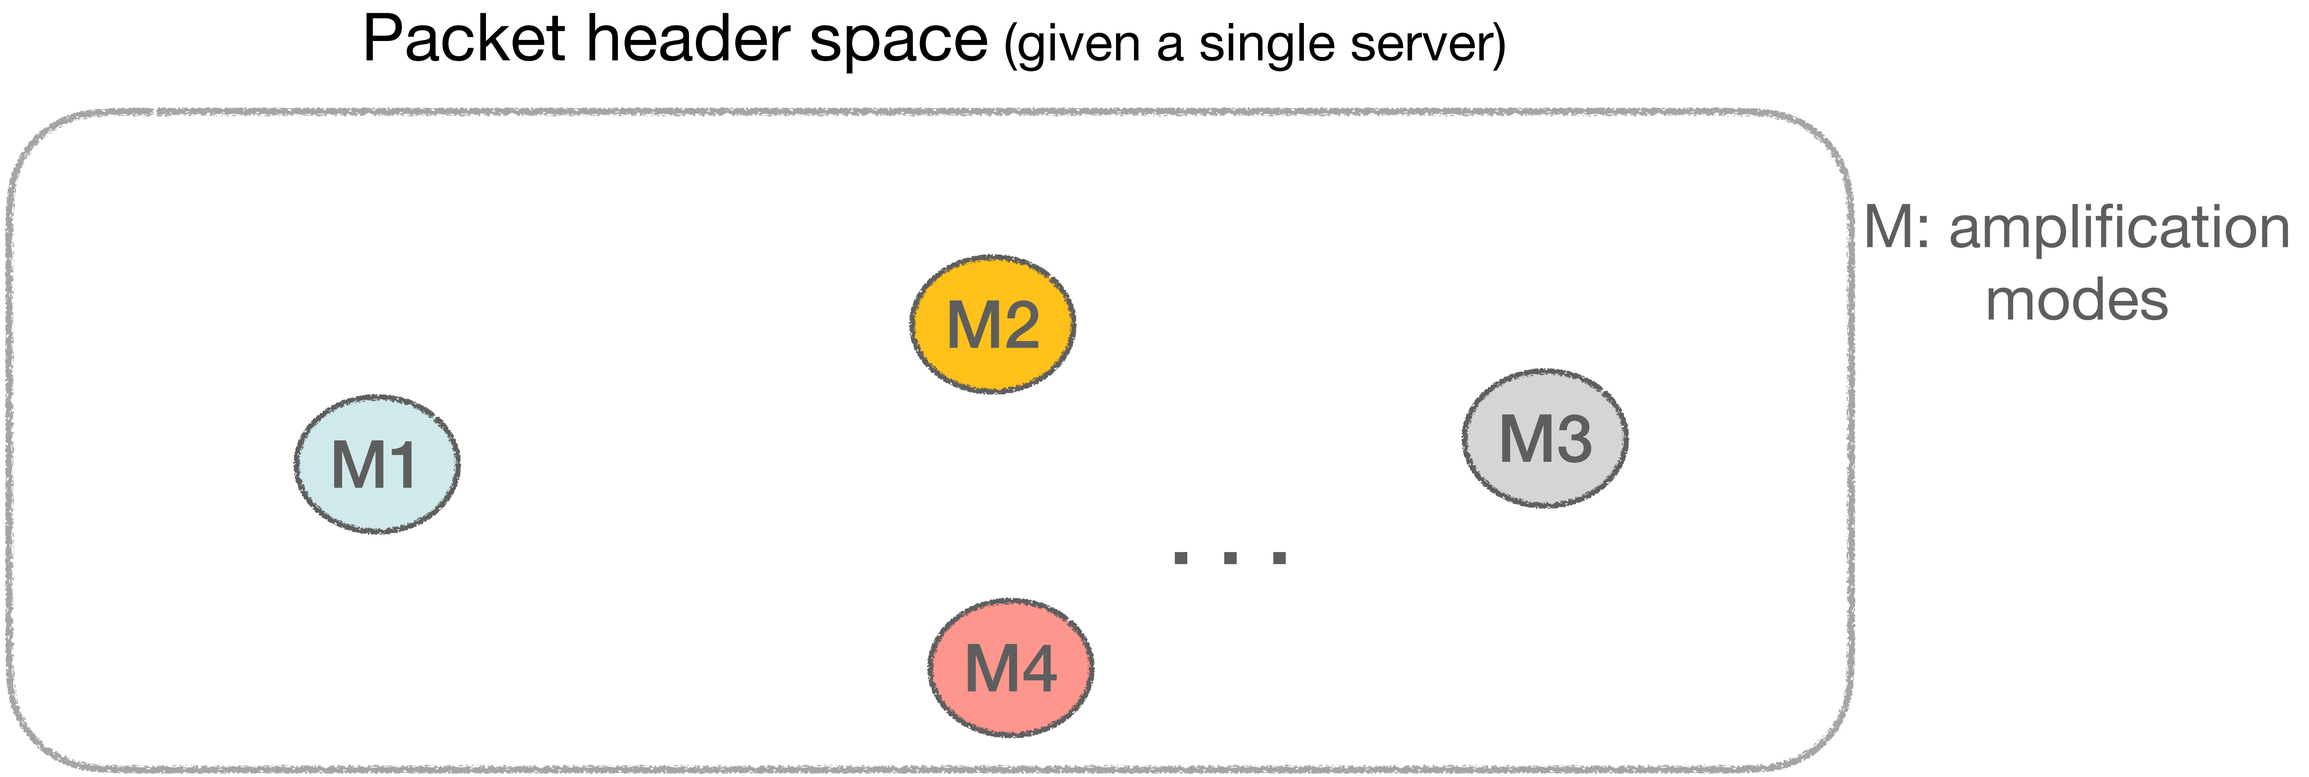

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

Packet header space (given a single server)

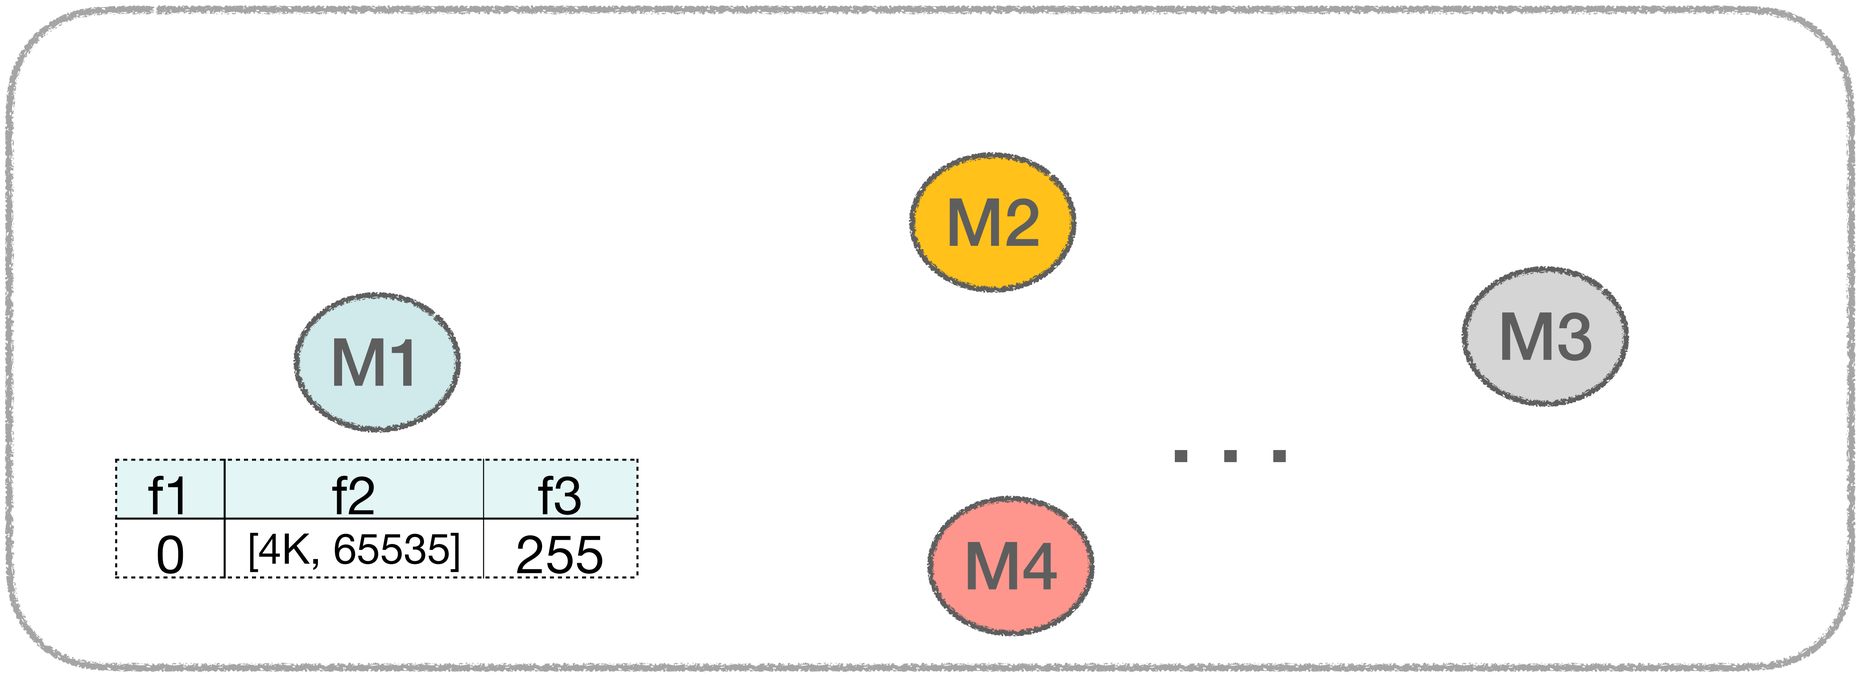

f1	f2	f3
0	[4K, 65535]	255

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

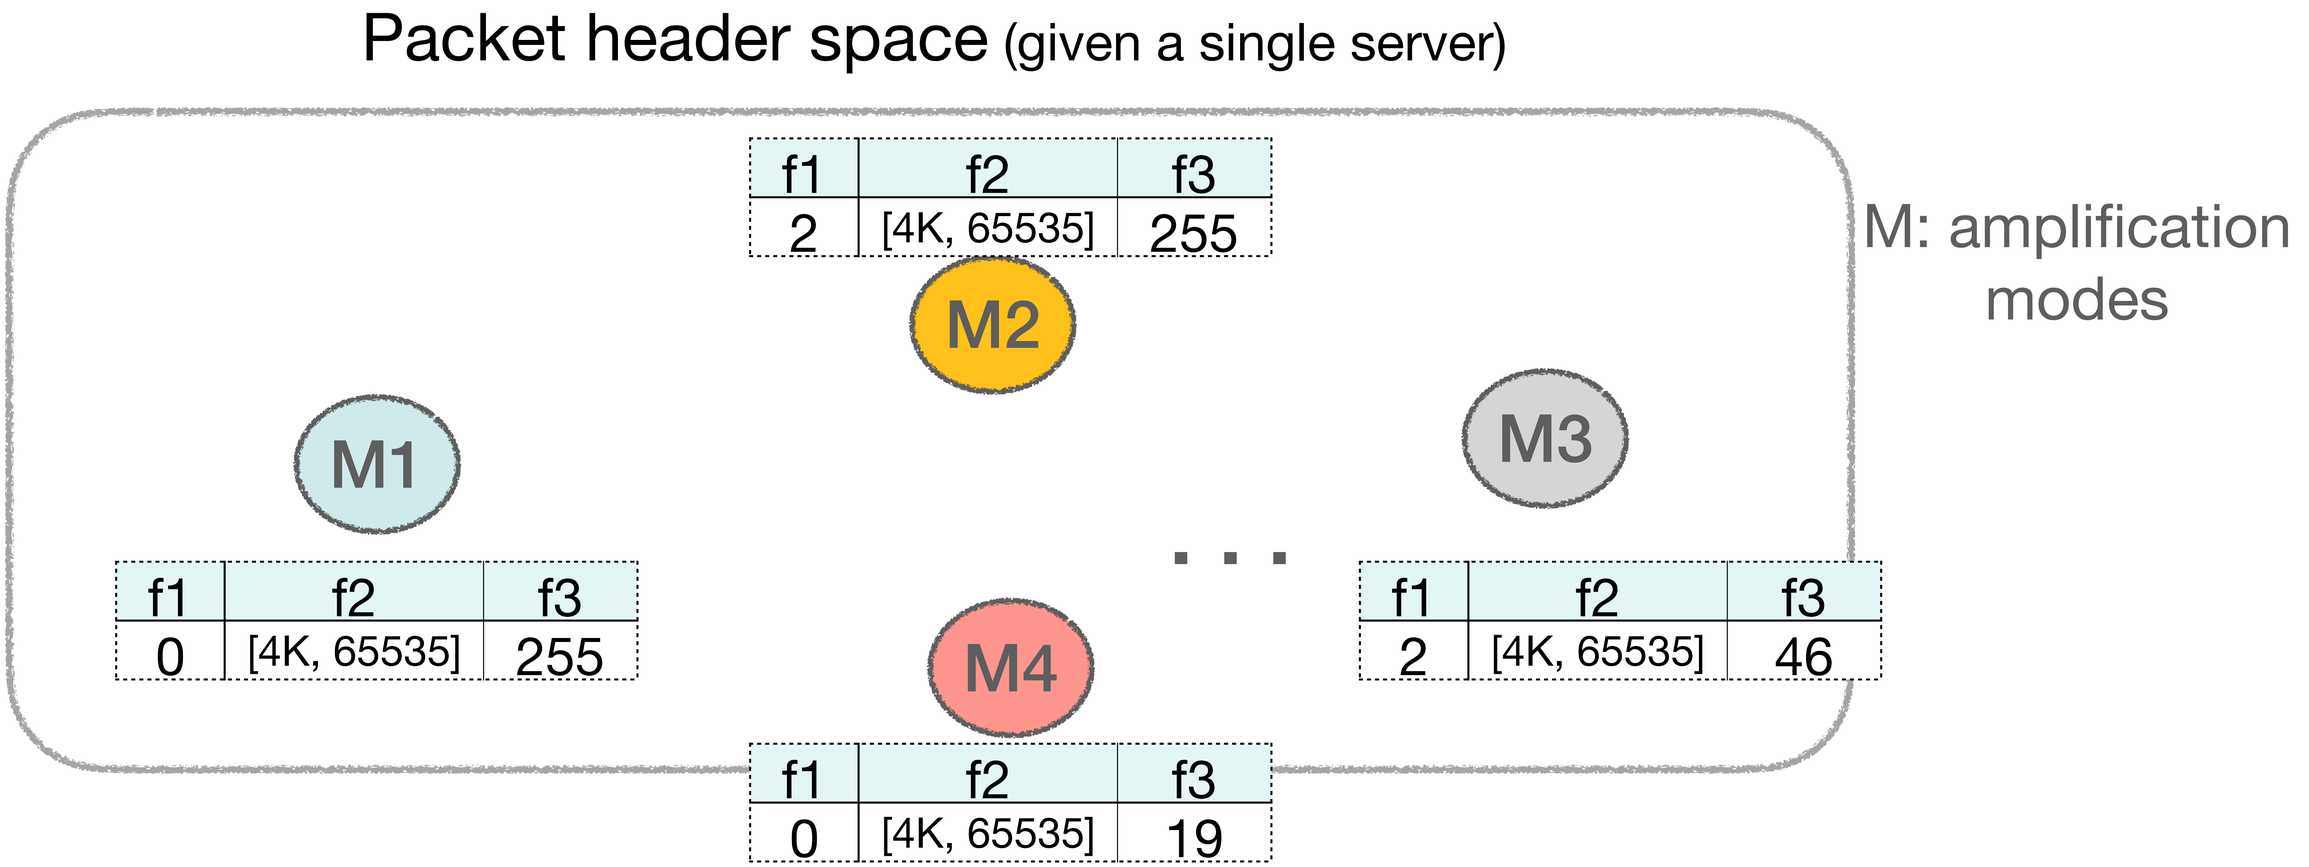

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

Amplification modes overlap in their field values.

Packet header space (given a single server)

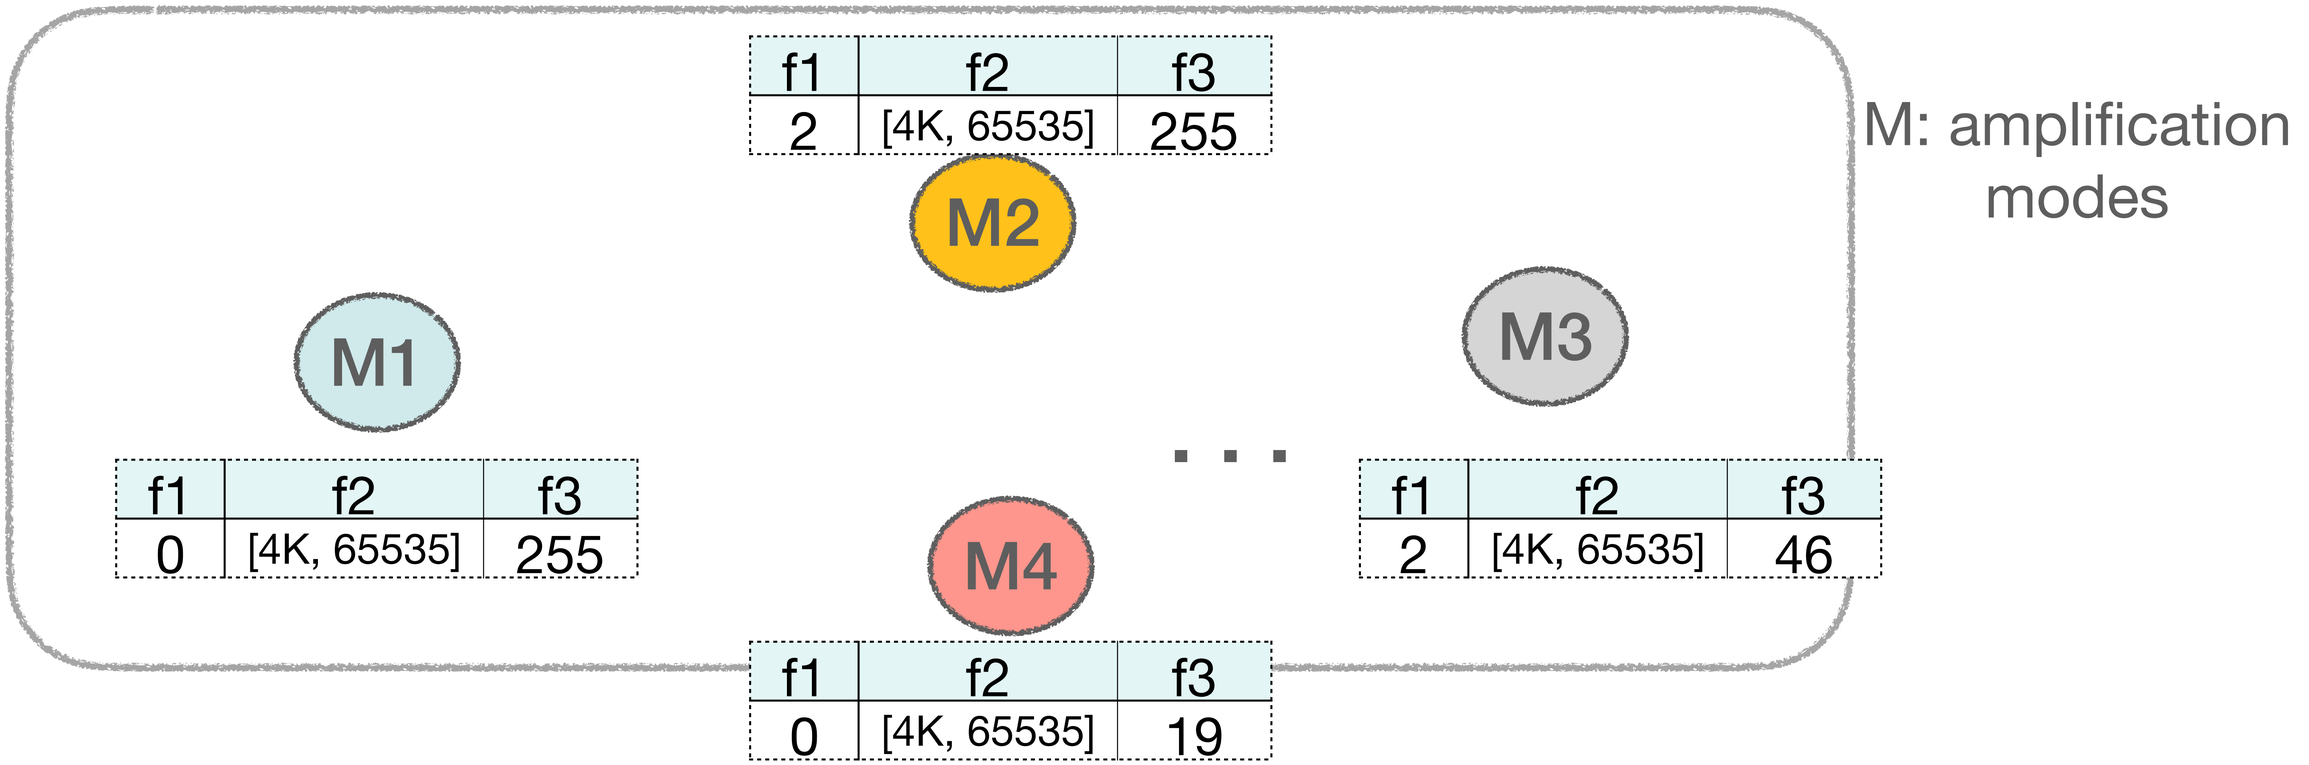

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

Amplification modes overlap in their field values.

Packet header space (given a single server)

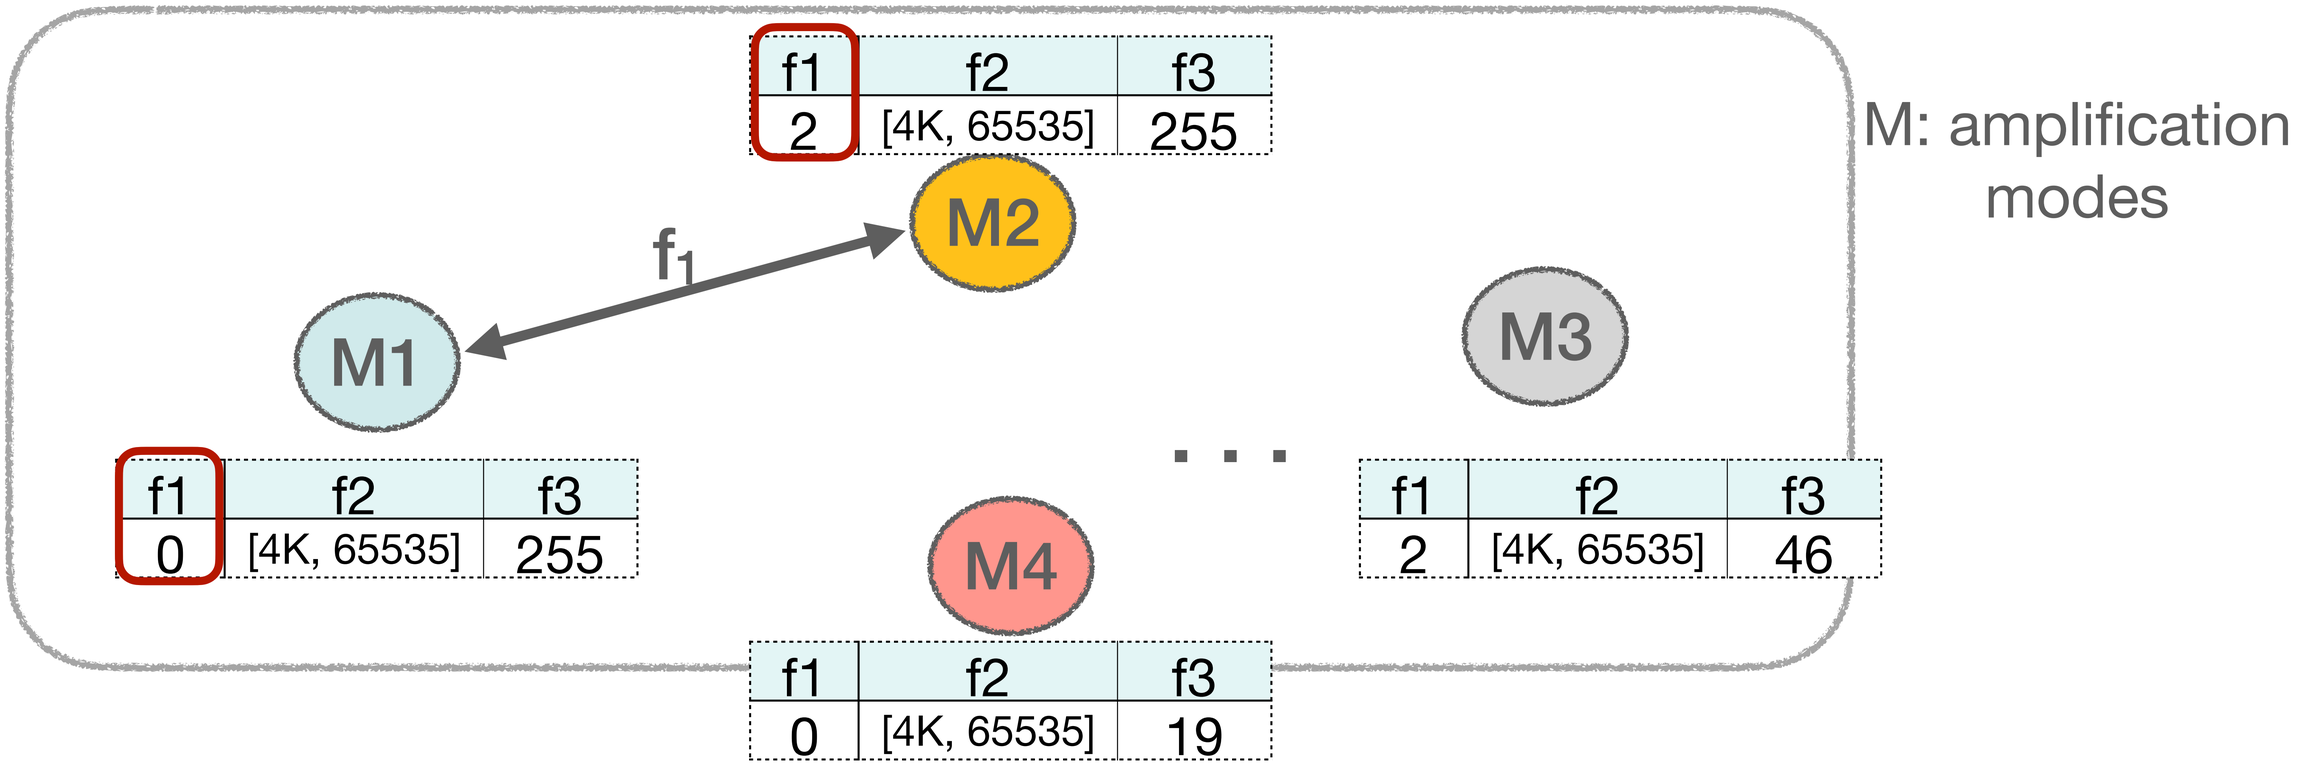

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

Amplification modes overlap in their field values.

Packet header space (given a single server)

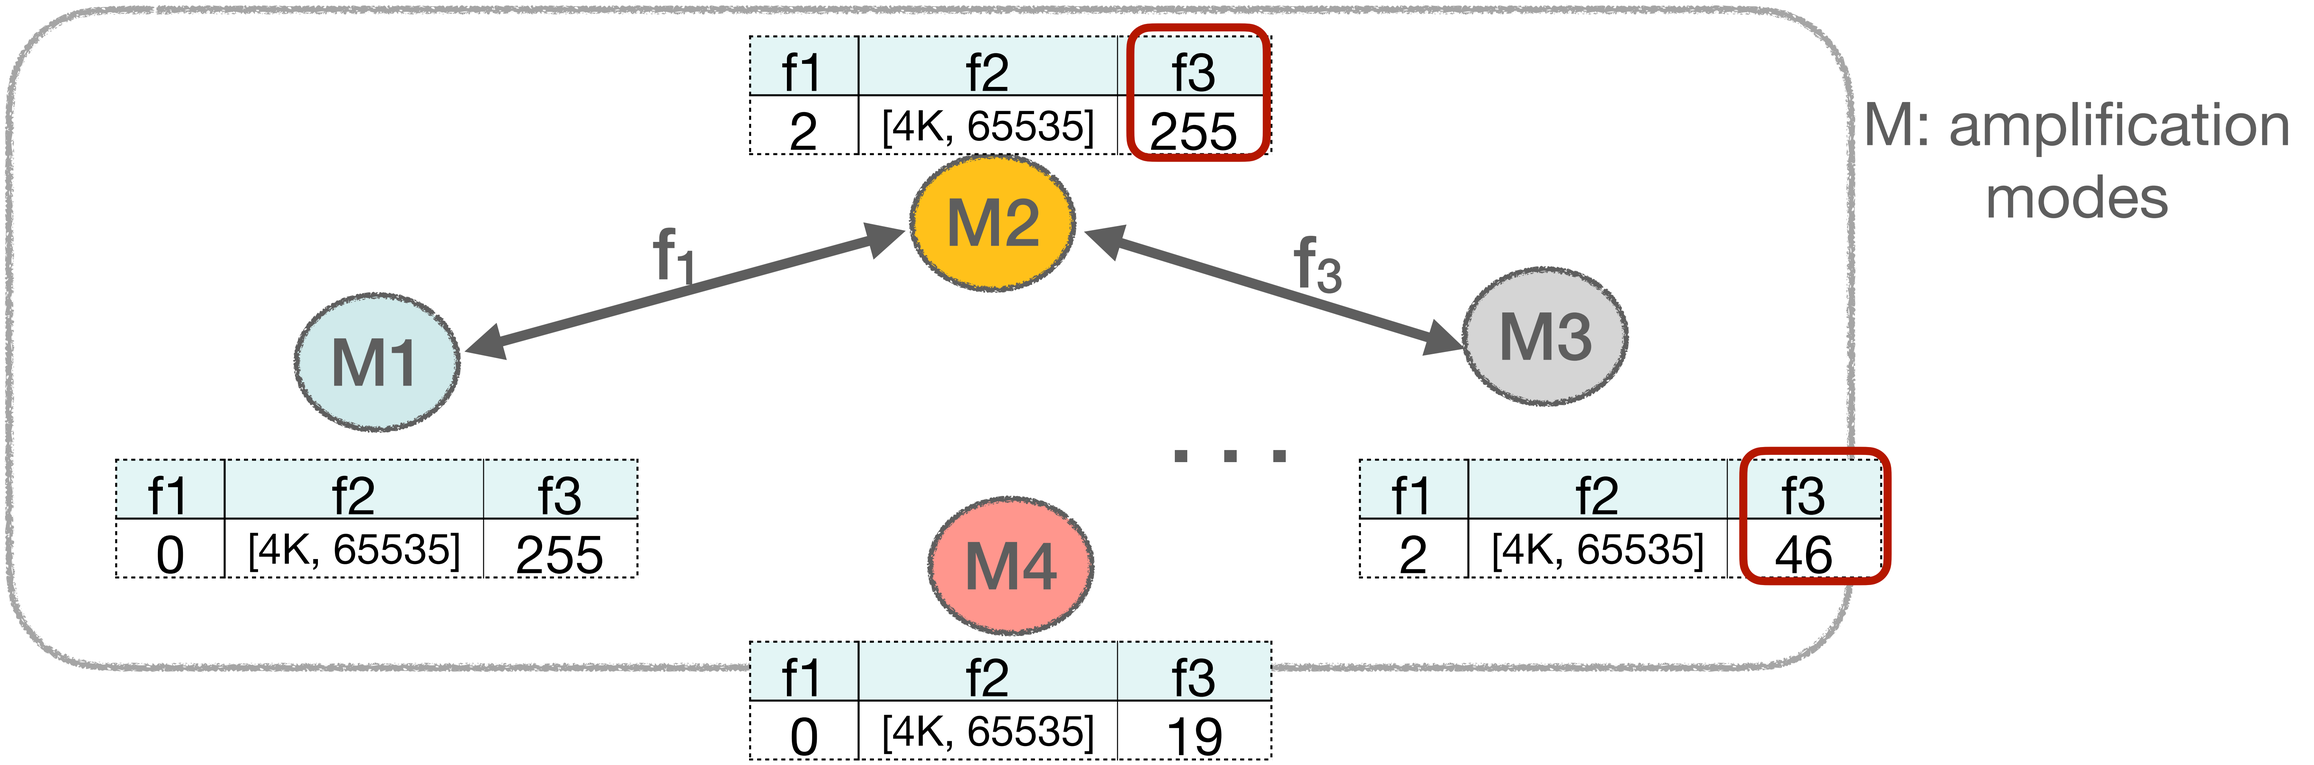

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

Amplification modes overlap in their field values.

Packet header space (given a single server)

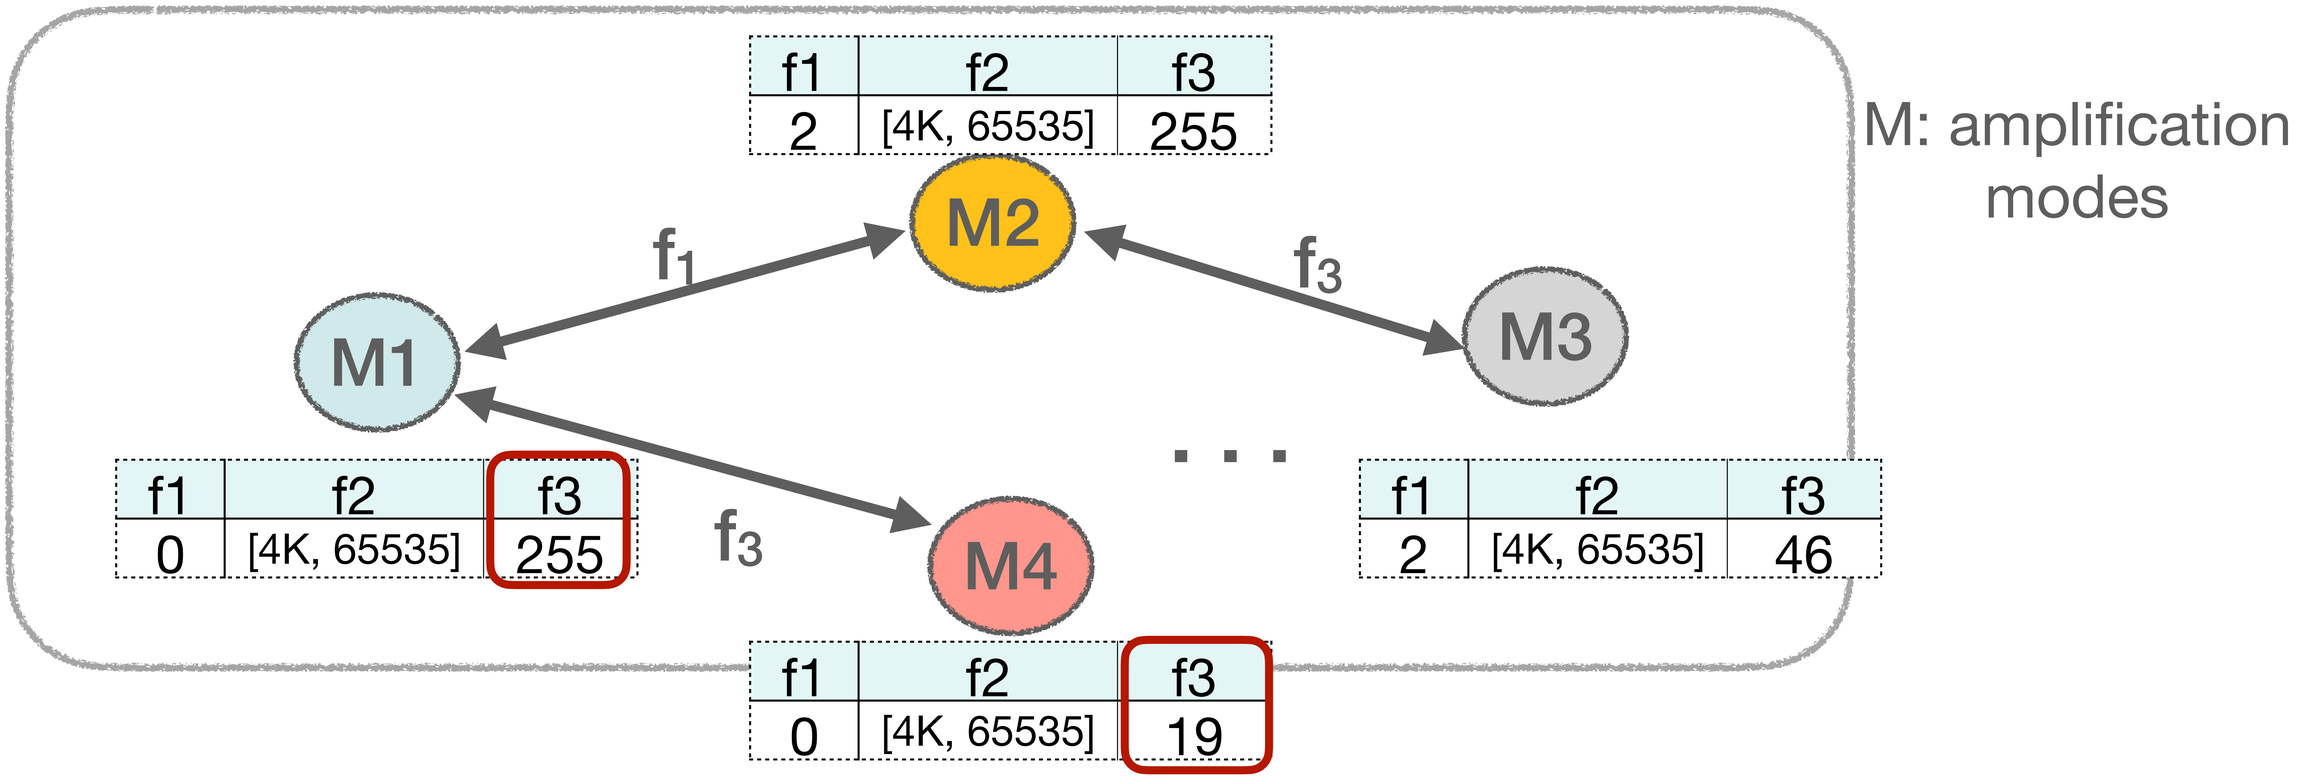

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

Amplification modes overlap in their field values.

Packet header space (given a single server)

Random sampling
(e.g., few hundreds packets)

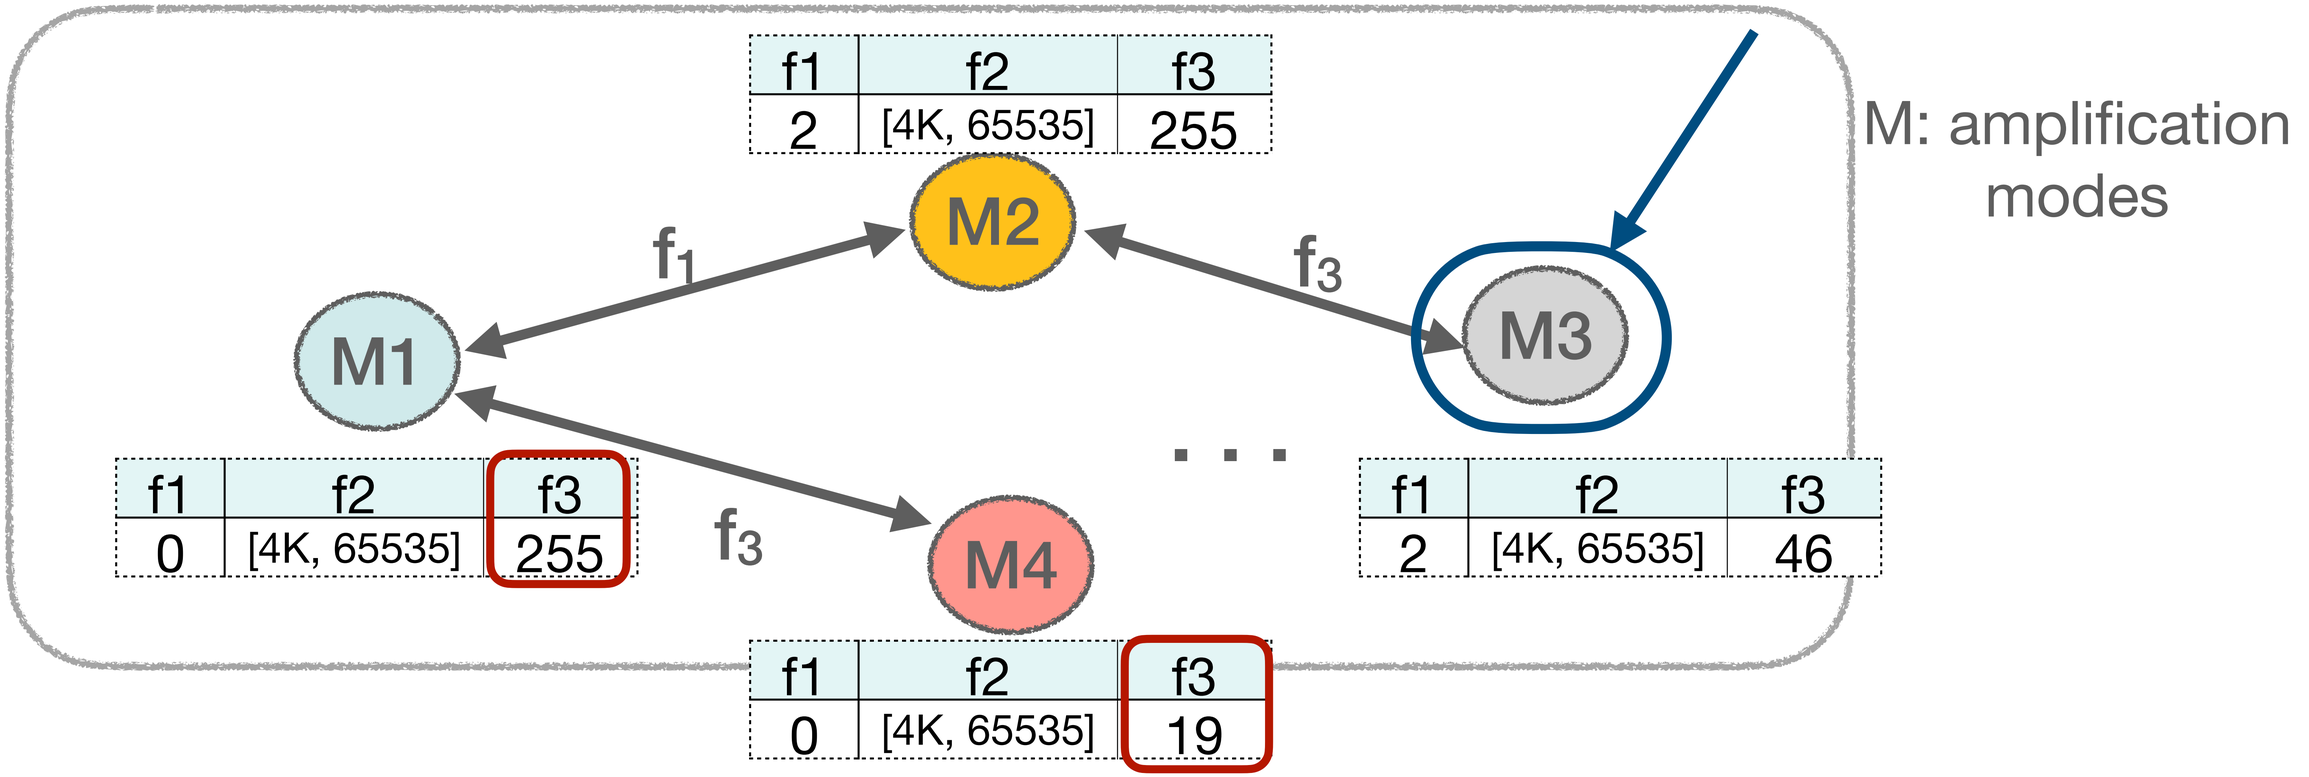

Insight: Leverage Locality across Amplification Modes to Achieve Coverage

Amplification modes overlap in their field values.

Packet header space (given a single server)

Random sampling
(e.g., few hundreds packets)

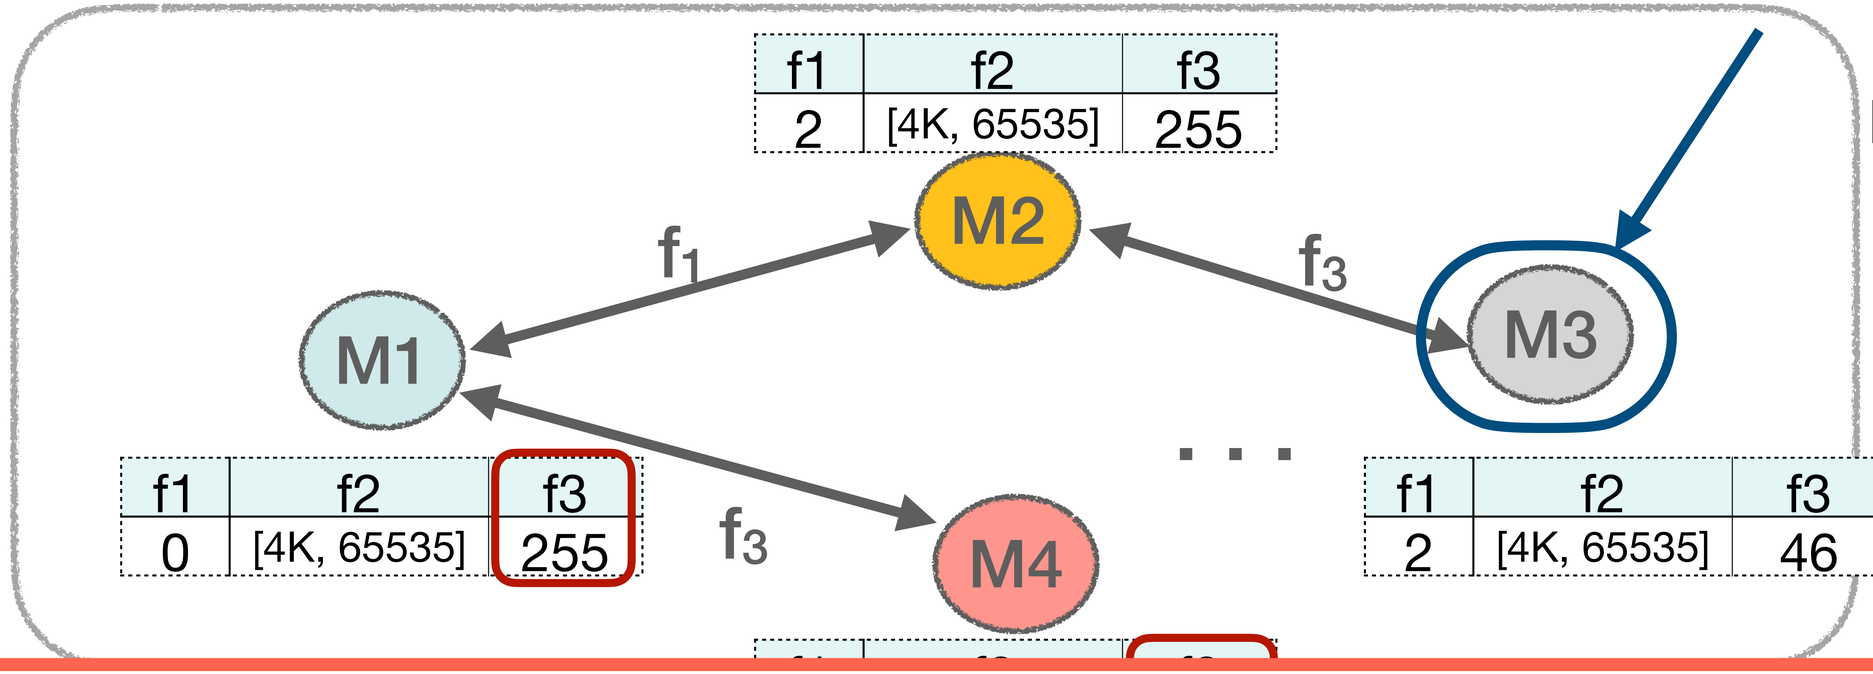

Exploit locality to achieve coverage by using a **per-field search**

Insight: Share Mode Insights across Servers!

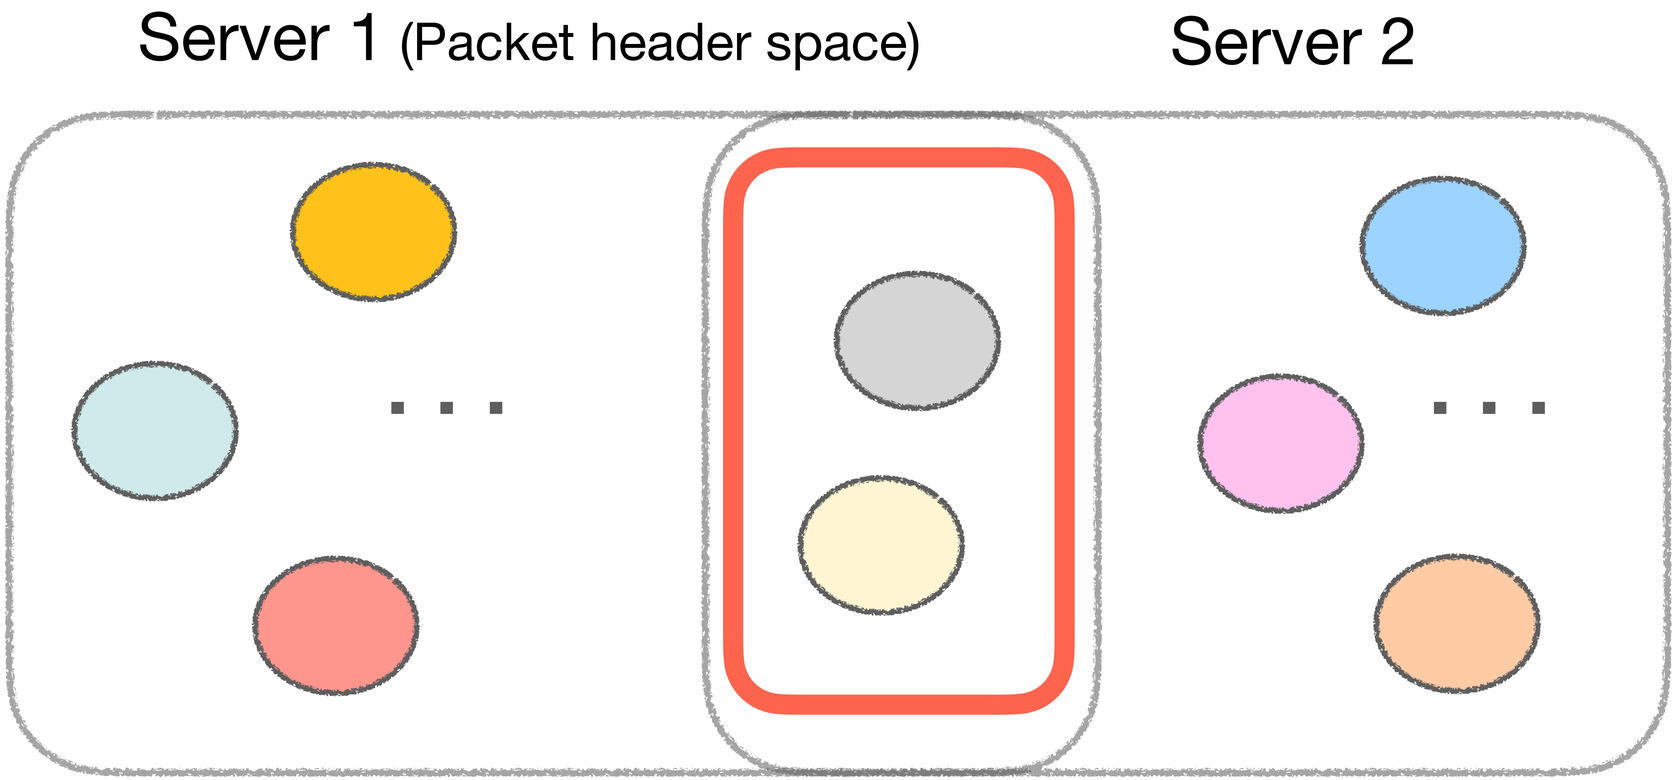

Insight: Share Mode Insights across Servers!

While servers are heterogeneous, some share a subset of amplification modes.

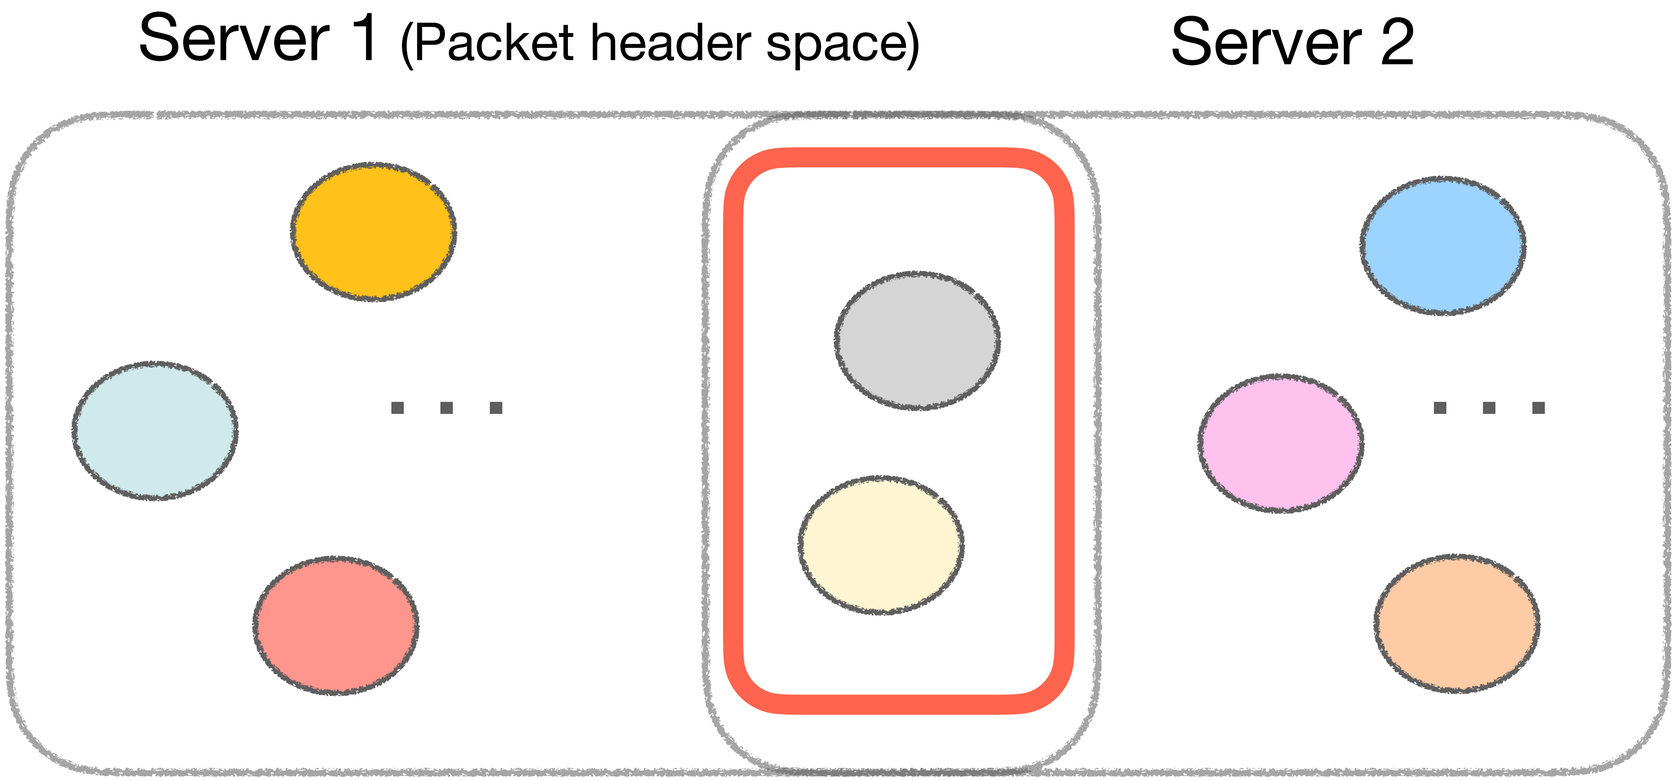

Insight: Share Mode Insights across Servers!

While servers are heterogeneous, some share a subset of amplification modes.

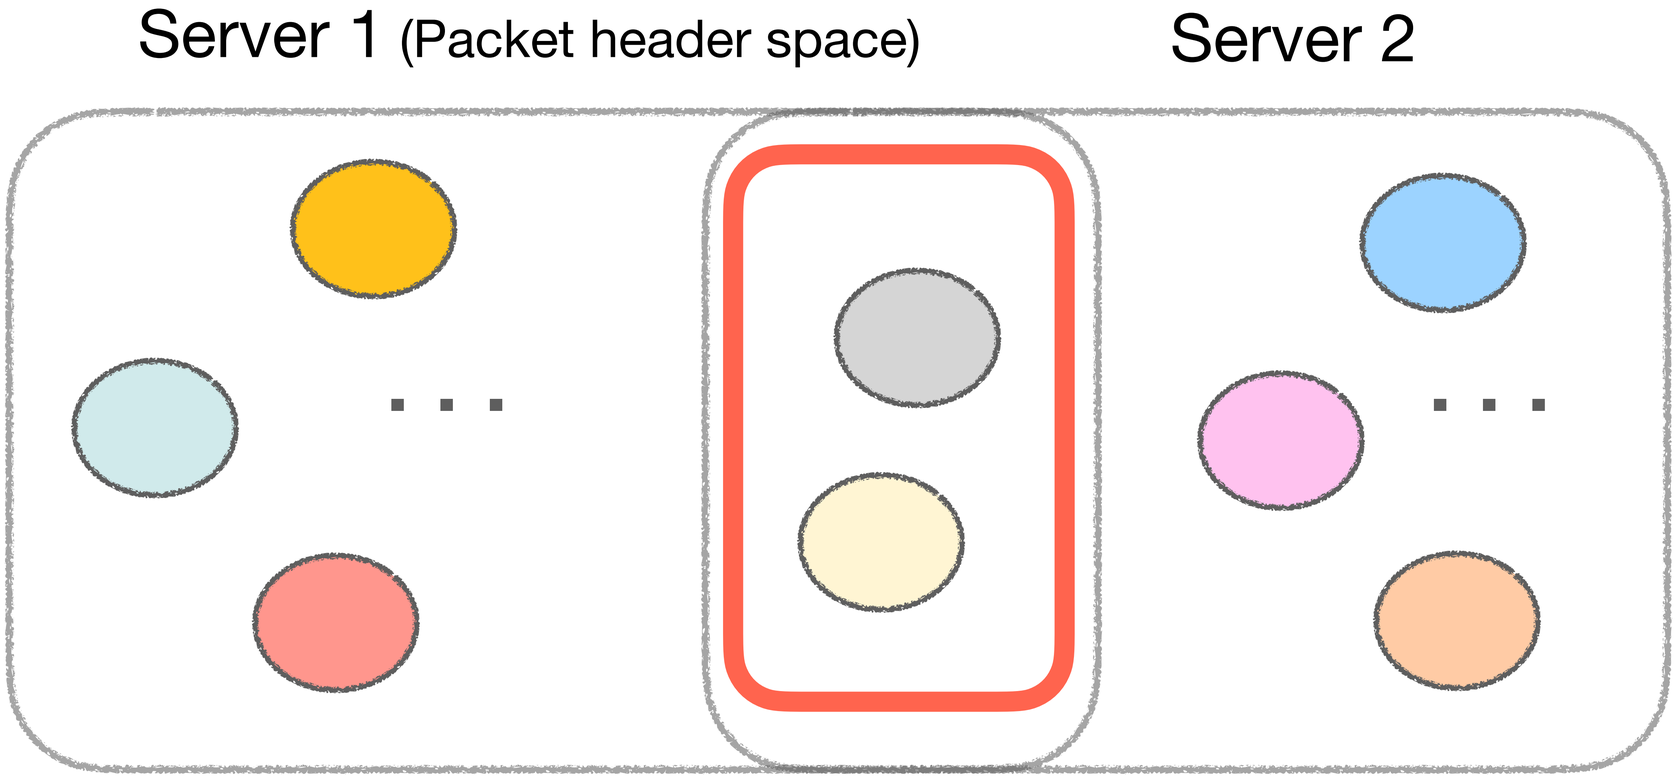

Reduce query overhead by sharing insights across servers!

AmpMap: A Low-Footprint Amplification Monitoring Service

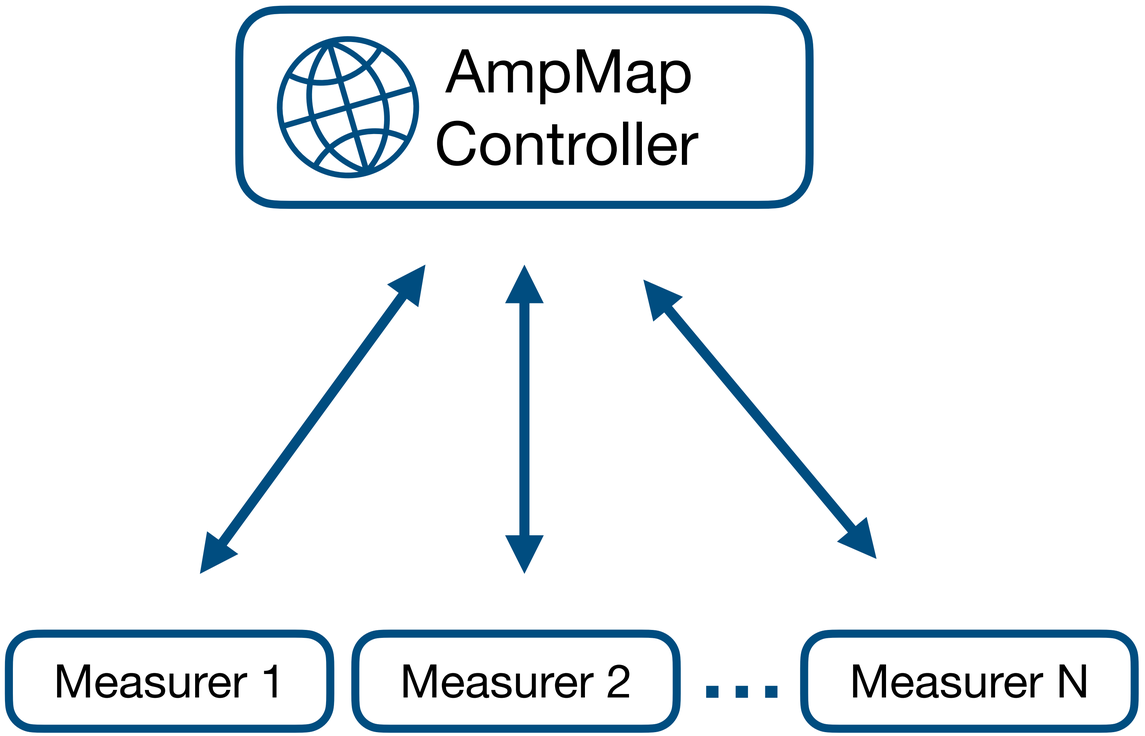

AmpMap: A Low-Footprint Amplification Monitoring Service

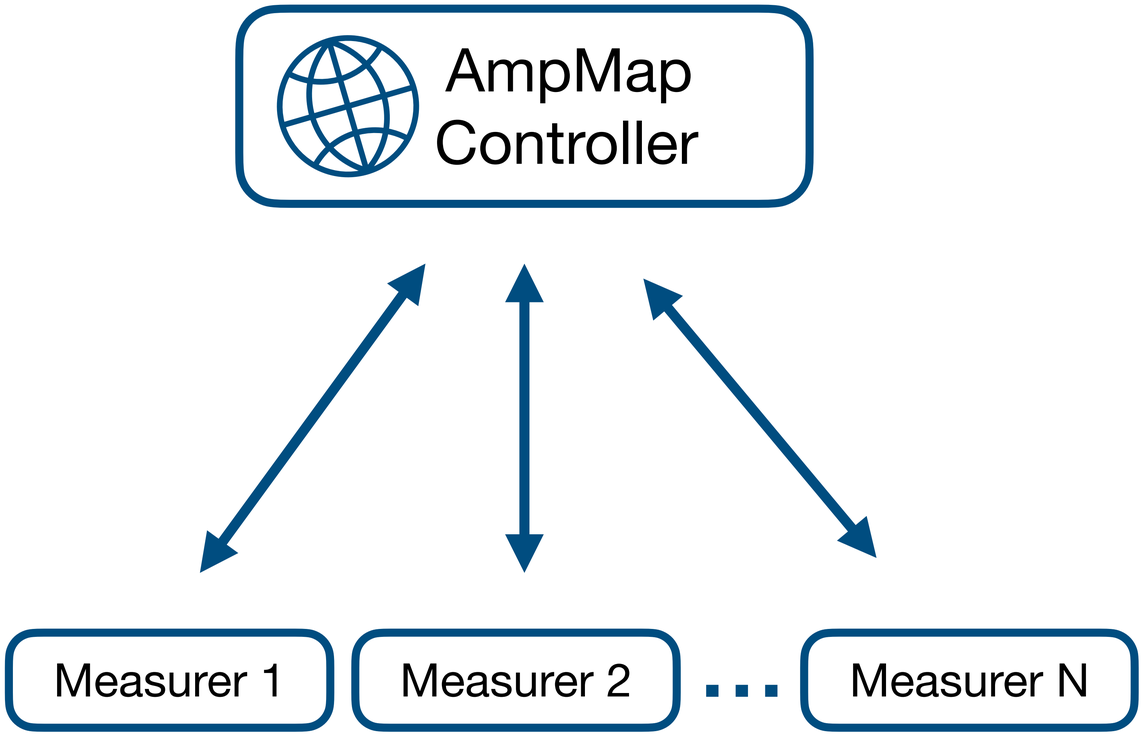

- Scanned 10K servers each for 6 popular UDP protocols (10K servers obtained from Censys[3] and Shodan[4]).

[3] Z. Durumeric et al. A search engine backed by internet-wide scanning. In Proc. CCS, 2015. [4] https://www.shodan.io/

AmpMap: A Low-Footprint Amplification Monitoring Service

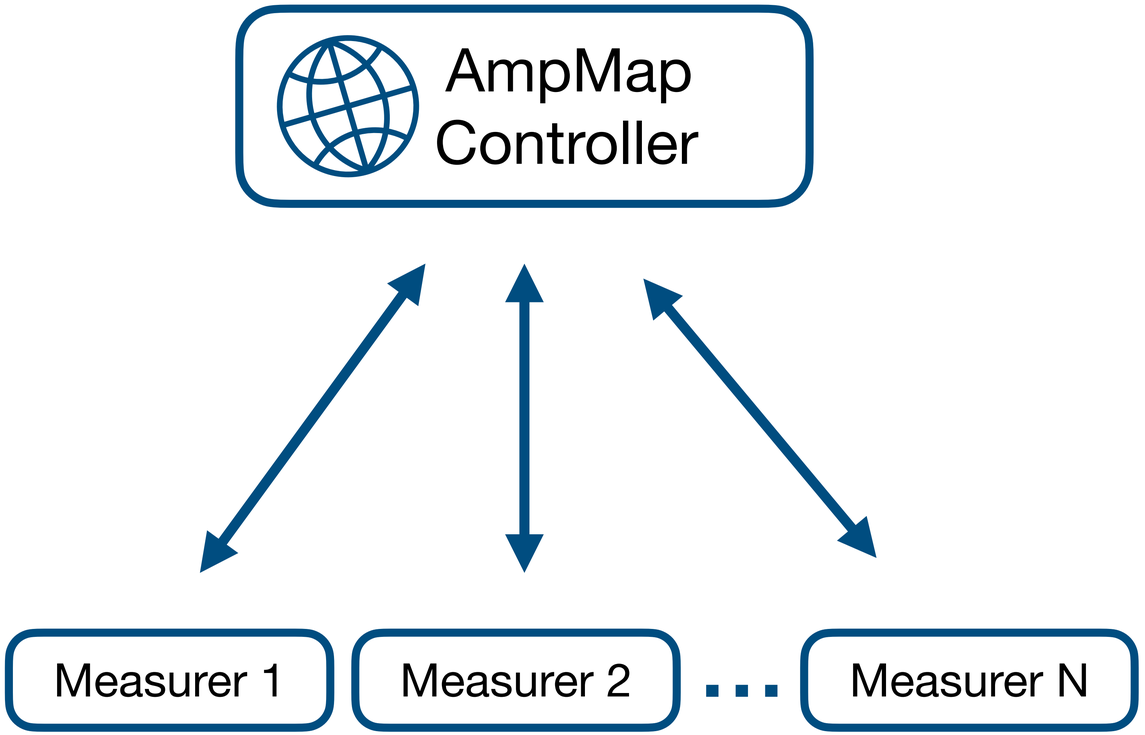

- Scanned 10K servers each for 6 popular UDP protocols (10K servers obtained from Censys[3] and Shodan[4]).
- **Low footprint:** 48 kbps across 30 measurers for 3 days to scan ~10K servers

[3] Z. Durumeric et al. A search engine backed by internet-wide scanning. In Proc. CCS, 2015. [4] https://www.shodan.io/

AmpMap: A Low-Footprint Amplification Monitoring Service

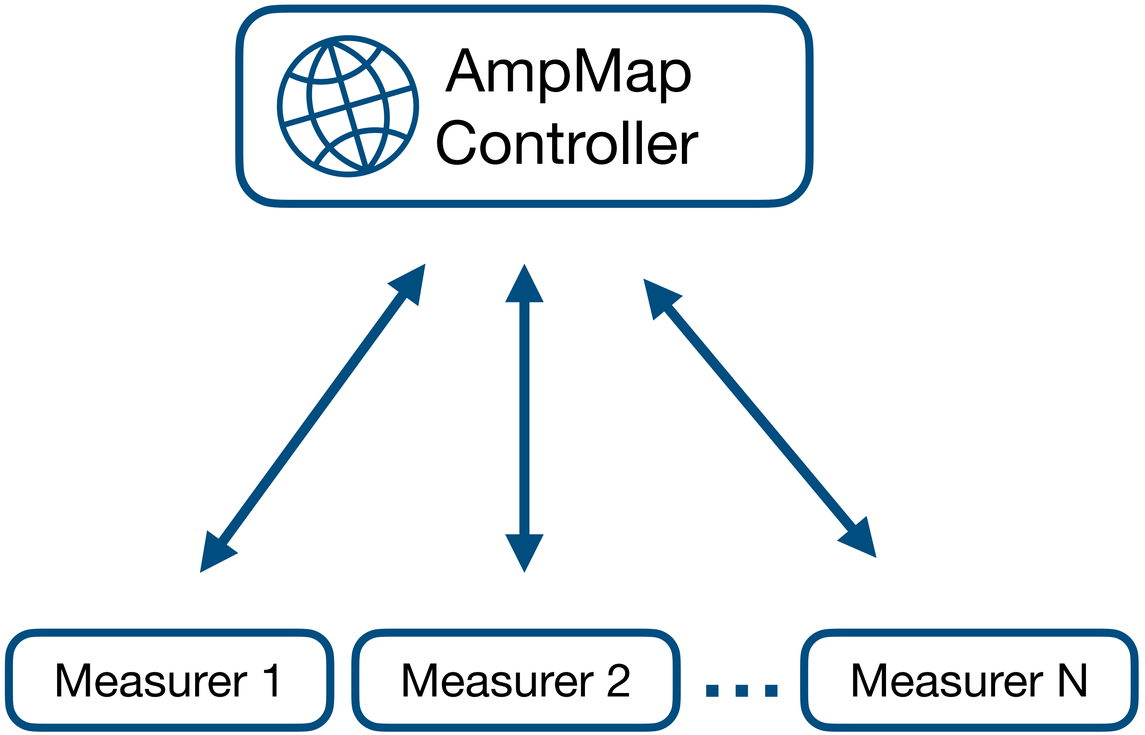

- Scanned 10K servers each for 6 popular UDP protocols (10K servers obtained from Censys[3] and Shodan[4]).
- **Low footprint:** 48 kbps across 30 measurers for 3 days to scan ~10K servers
- **General & Extensible** across future protocols

[3] Z. Durumeric et al. A search engine backed by internet-wide scanning. In Proc. CCS, 2015. [4] https://www.shodan.io/

AmpMap: A Low-Footprint Amplification Monitoring Service

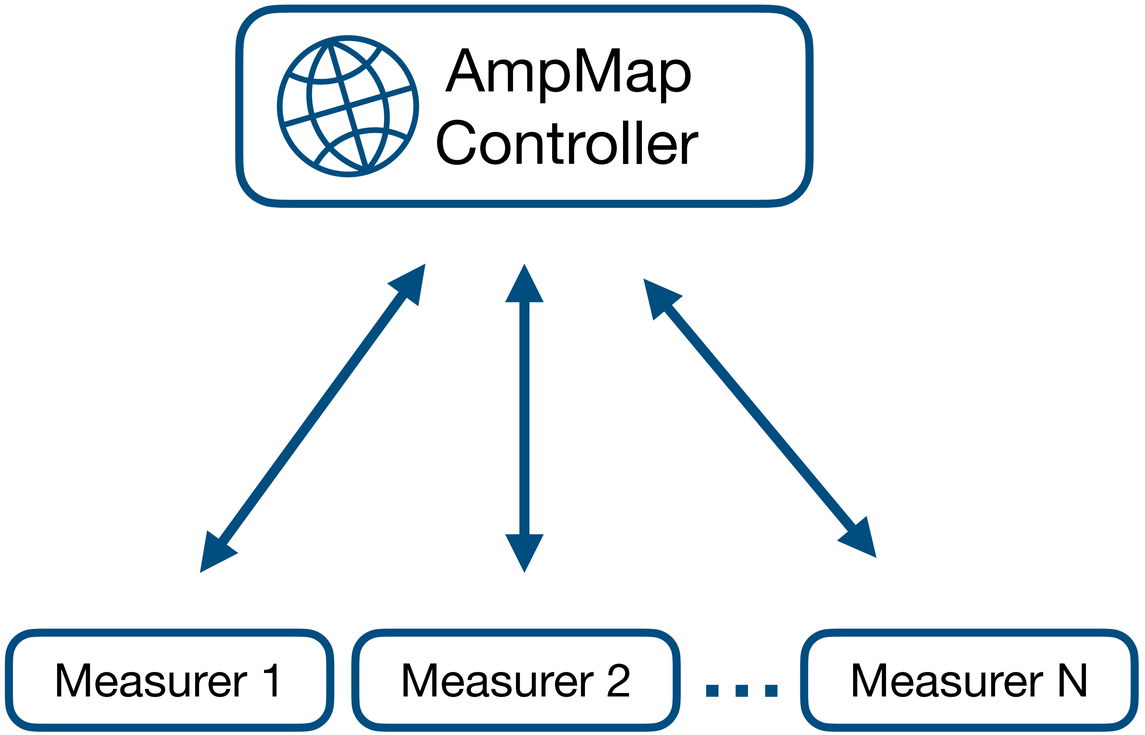

- Scanned 10K servers each for 6 popular UDP protocols (10K servers obtained from Censys[3] and Shodan[4]).
- **Low footprint:** 48 kbps across 30 measurers for 3 days to scan ~10K servers
- **General & Extensible** across future protocols
- **Configurable & horizontally scalable** by configuring the packet sending rate and # measurers

[3] Z. Durumeric et al. A search engine backed by internet-wide scanning. In Proc. CCS, 2015. [4] https://www.shodan.io/

DNS: New Amplification Modes

- Previously known amplification modes for **DNS** are: ANY[3] or TXT[4] record types.

[3] Alert (TA13-088A) UDP-Based Amplification Attacks. https://us-cert.cisa.gov/ncas/alerts/TA13-088A [4] Security Bulletin: Crafted DNS Text Attack. https://tinyurl.com/y9zpevuy

DNS: New Amplification Modes

- Previously known amplification modes for **DNS** are: ANY[3] or TXT[4] record types.

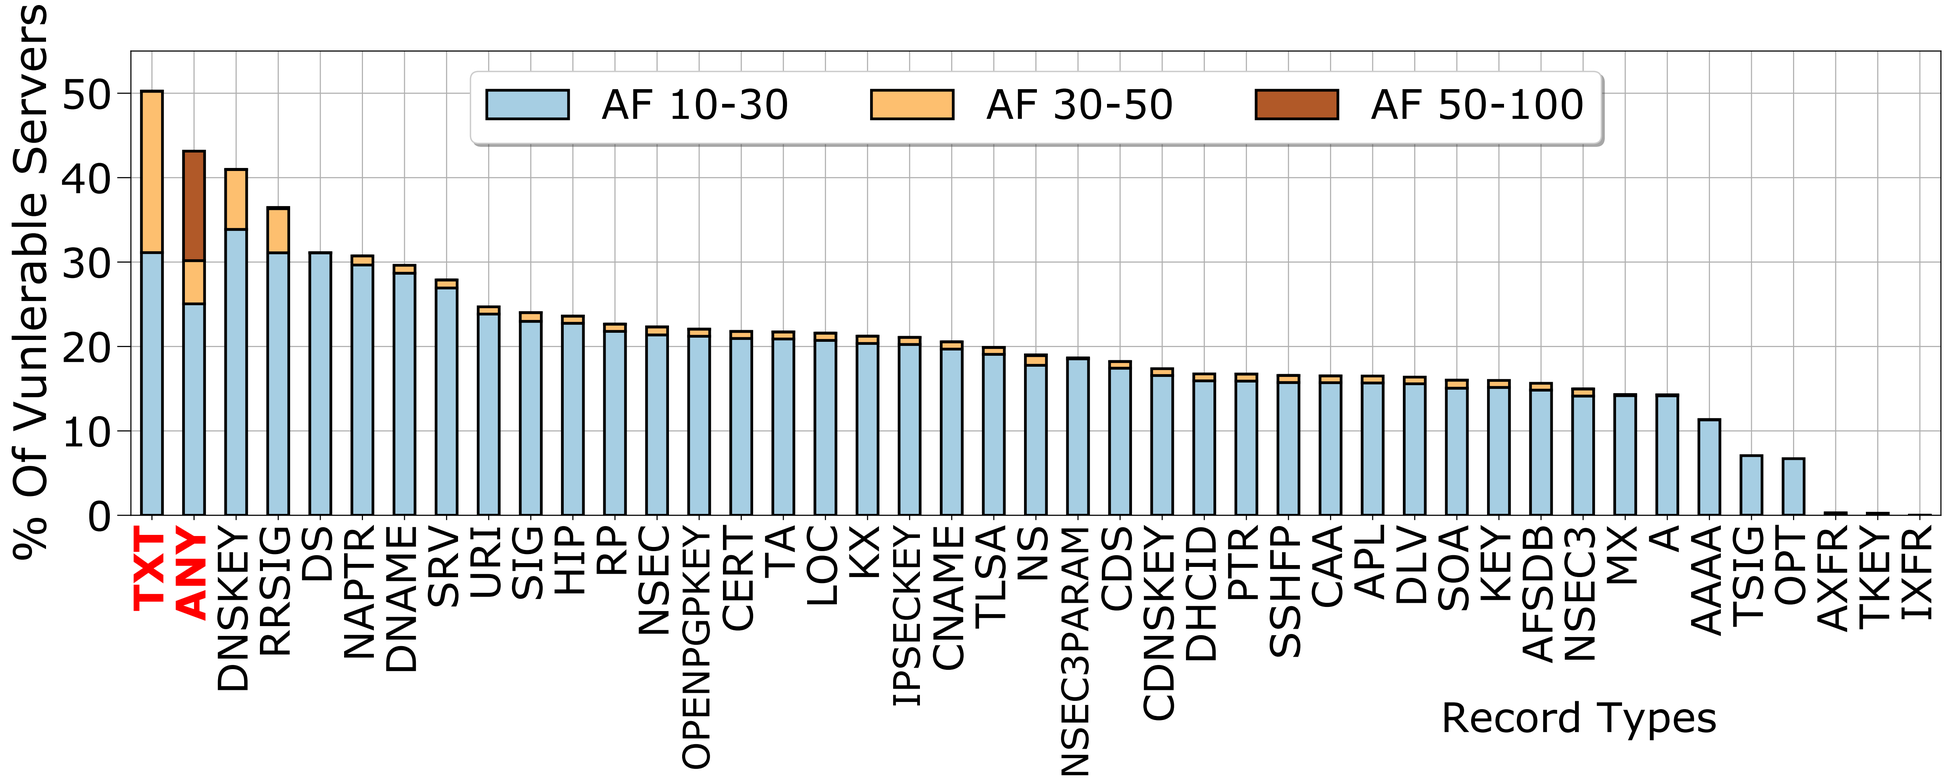

$$\text{AF (amp factor)} = \frac{|response|}{|request|}$$

[3] Alert (TA13-088A) UDP-Based Amplification Attacks. https://us-cert.cisa.gov/ncas/alerts/TA13-088A [4] Security Bulletin: Crafted DNS Text Attack. https://tinyurl.com/y9zpevuy

DNS: New Amplification Modes

- Previously known amplification modes for **DNS** are: ANY[3] or TXT[4] record types.

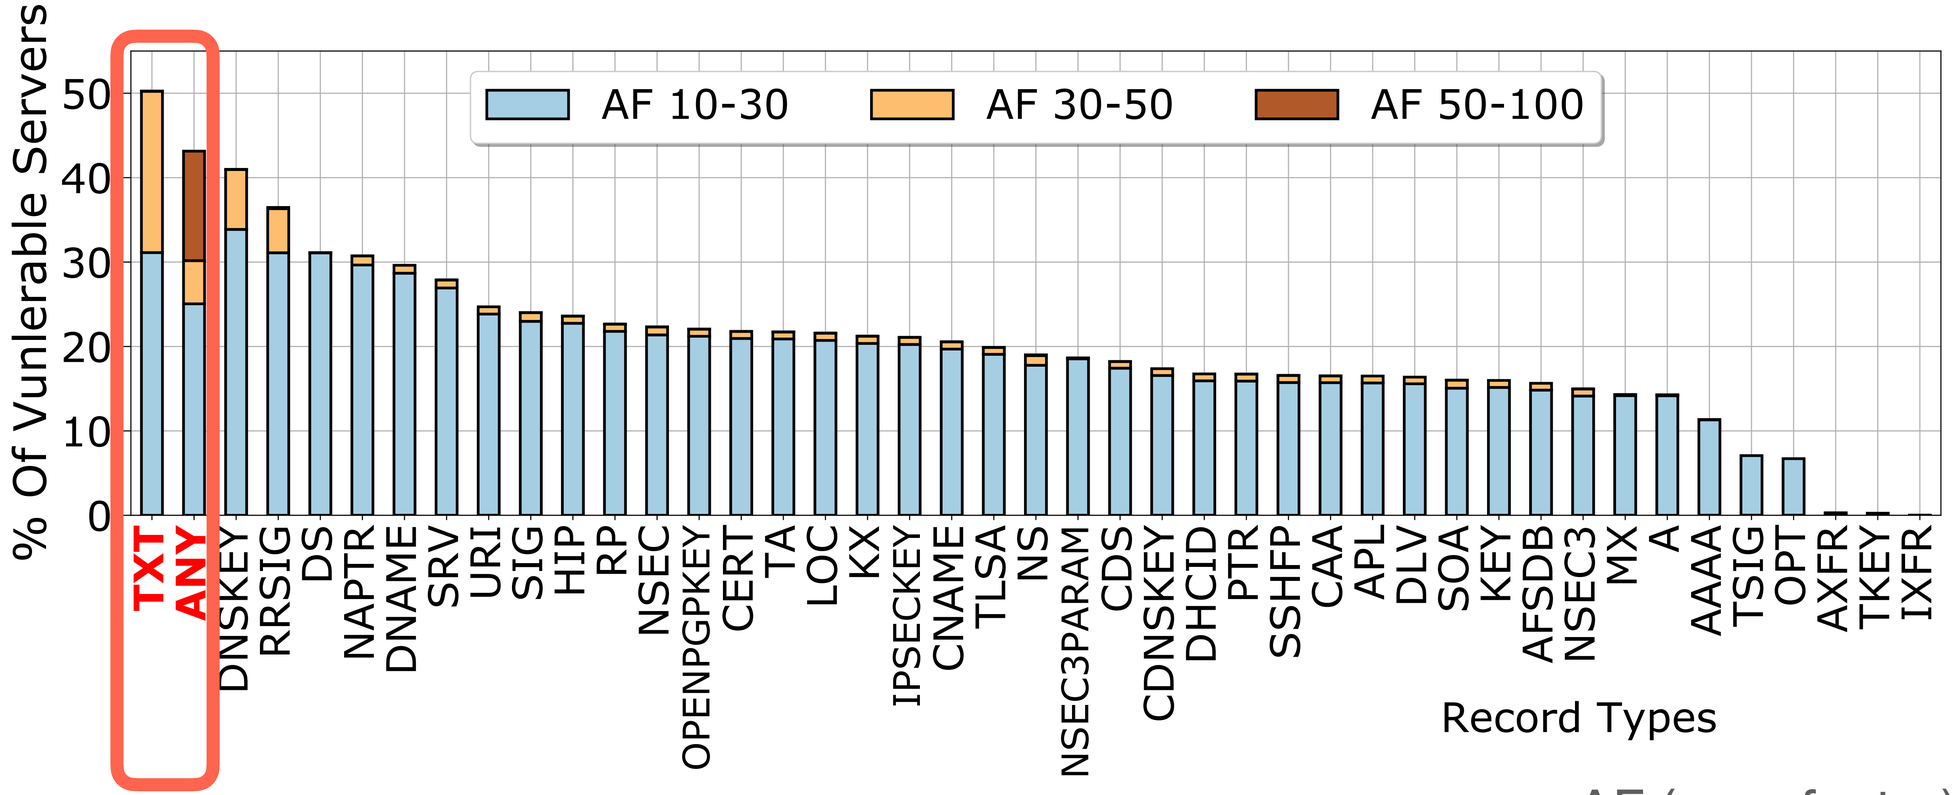

$$\text{AF (amp factor)} = \frac{|response|}{|request|}$$

[3] Alert (TA13-088A) UDP-Based Amplification Attacks. https://us-cert.cisa.gov/ncas/alerts/TA13-088A [4] Security Bulletin: Crafted DNS Text Attack. https://tinyurl.com/y9zpevuy

DNS: New Amplification Modes

- Previously known amplification modes for **DNS** are: ANY[3] or TXT[4] record types.

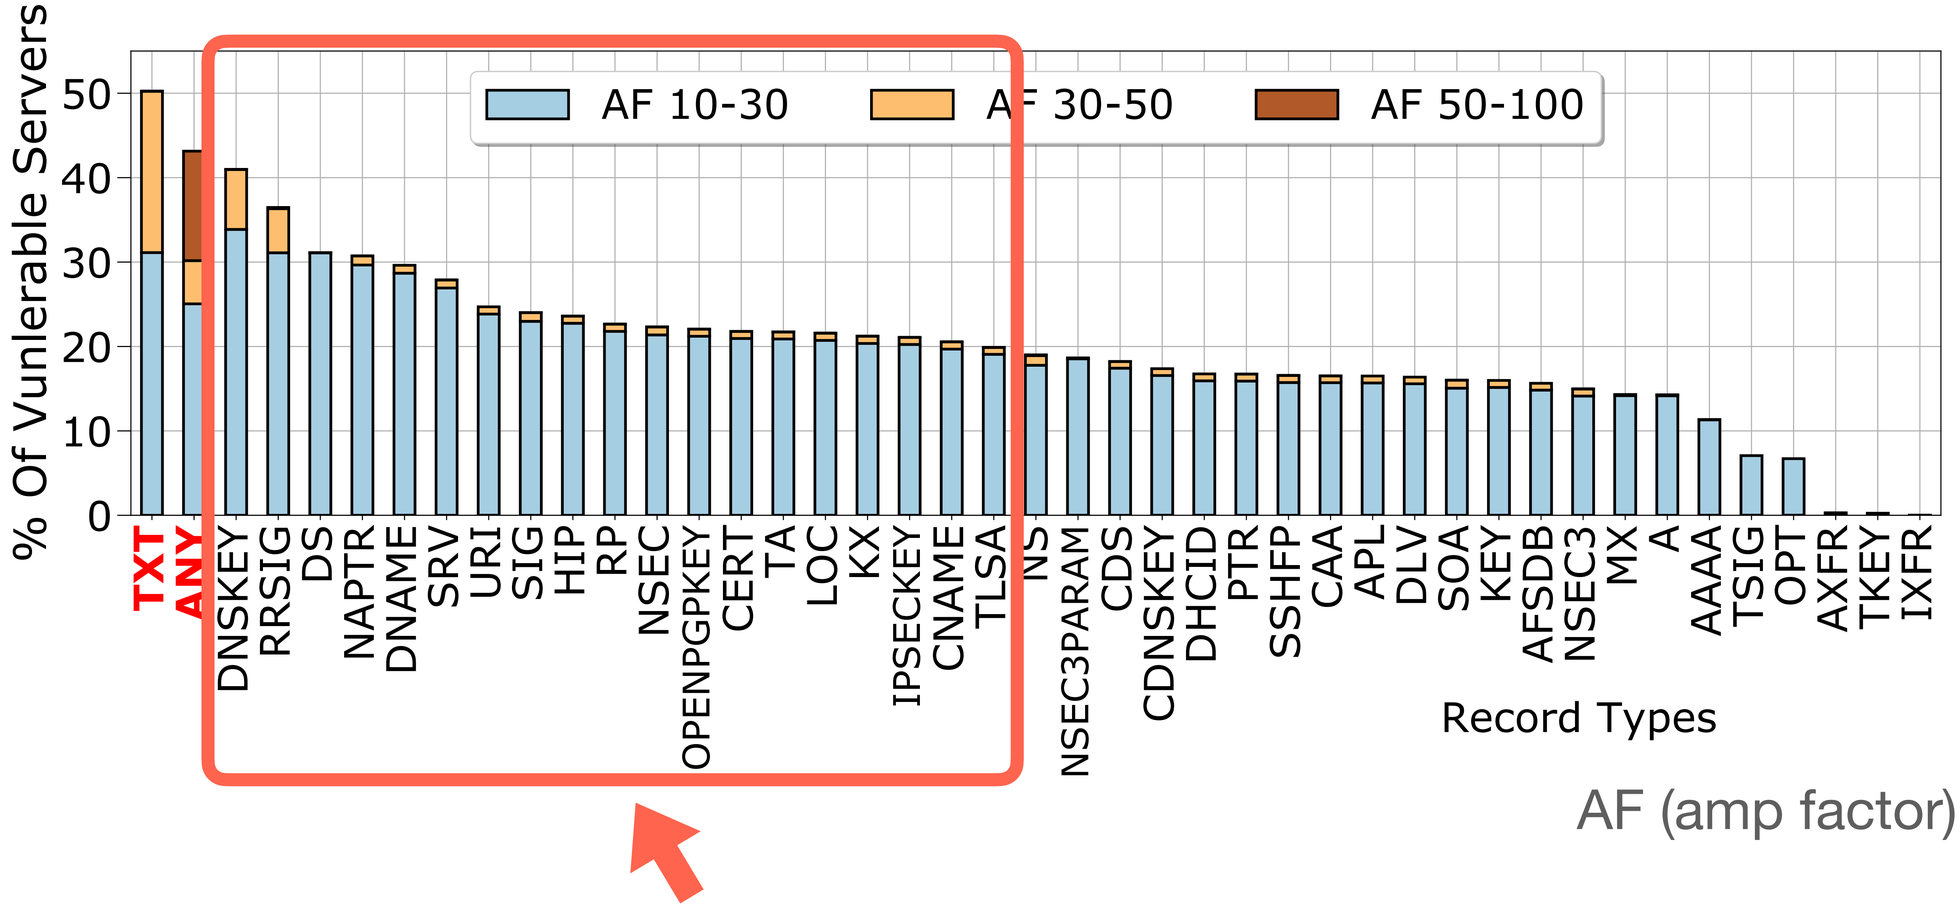

$$\text{AF (amp factor)} = \frac{|response|}{|request|}$$

More than 20% of (scanned) servers can incur AF ≥10 with 19 other record types!

[3] Alert (TA13-088A) UDP-Based Amplification Attacks. https://us-cert.cisa.gov/ncas/alerts/TA13-088A [4] Security Bulletin: Crafted DNS Text Attack. https://tinyurl.com/y9zpevuy

Other Protocols: New Amplification Modes

DNS

Known → EDNS:0, RecordType: ANY | TXT

New → EDNS:1 or Other RecordTypes

Other Protocols: New Amplification Modes

DNS

Known → EDNS:0, RecordType: ANY | TXT

New → EDNS:1 or Other RecordTypes

NTP

Known → Monlist

New → GetRestrict, If Stats, etc.

AF ≥ 500 for certain servers!

Other Protocols: New Amplification Modes

DNS

Known → EDNS:0, RecordType: ANY | TXT

New → EDNS:1 or Other RecordTypes

NTP

Known → Monlist

New → GetRestrict, If Stats, etc.

AF ≥ 500 for certain servers!

SNMP

Known → GetBulk

New → GetNext

AF ≥ 200 for certain servers!

Other Protocols: New Amplification Modes

DNS

Known → EDNS:0, RecordType: ANY | TXT

New → EDNS:1 or Other RecordTypes

NTP

Known → Monlist

New → GetRestrict, If Stats, etc.

AF ≥ 500 for certain servers!

SNMP

Known → GetBulk

New → GetNext

AF ≥ 200 for certain servers!

Blocking these known modes still leave many other vectors for attackers

Significant Diversity across Servers & Protocols!

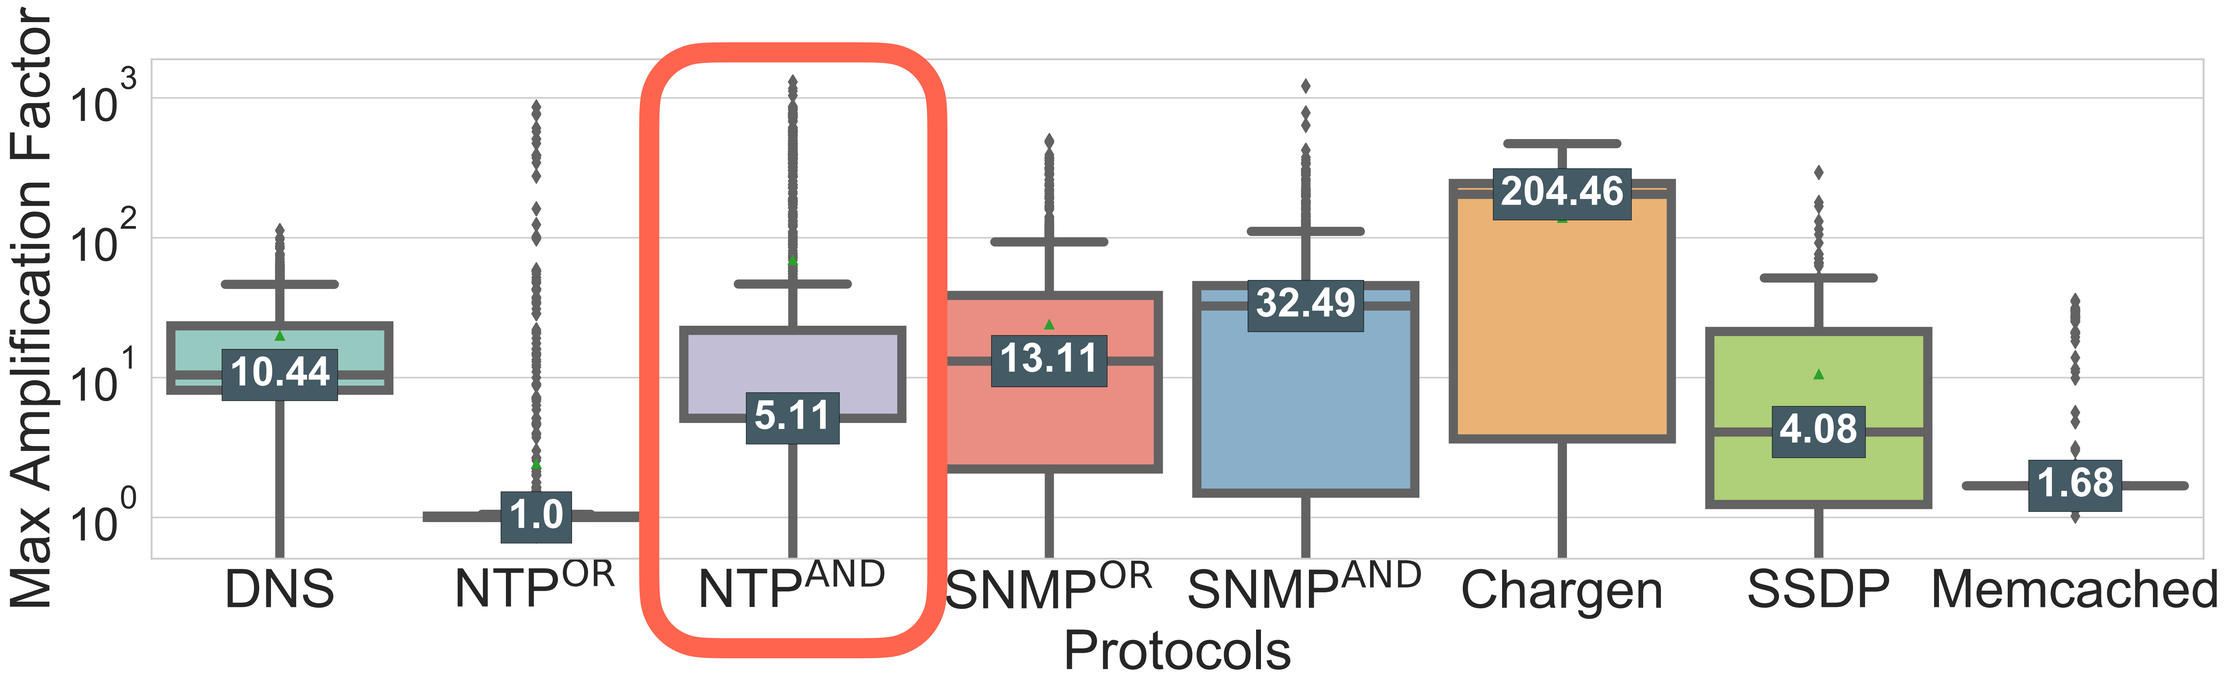

- **Across servers:**
  - NTP's median AF is only 5.11 but 1,300 AF for the max across measured servers.

Significant Diversity across Servers & Protocols!

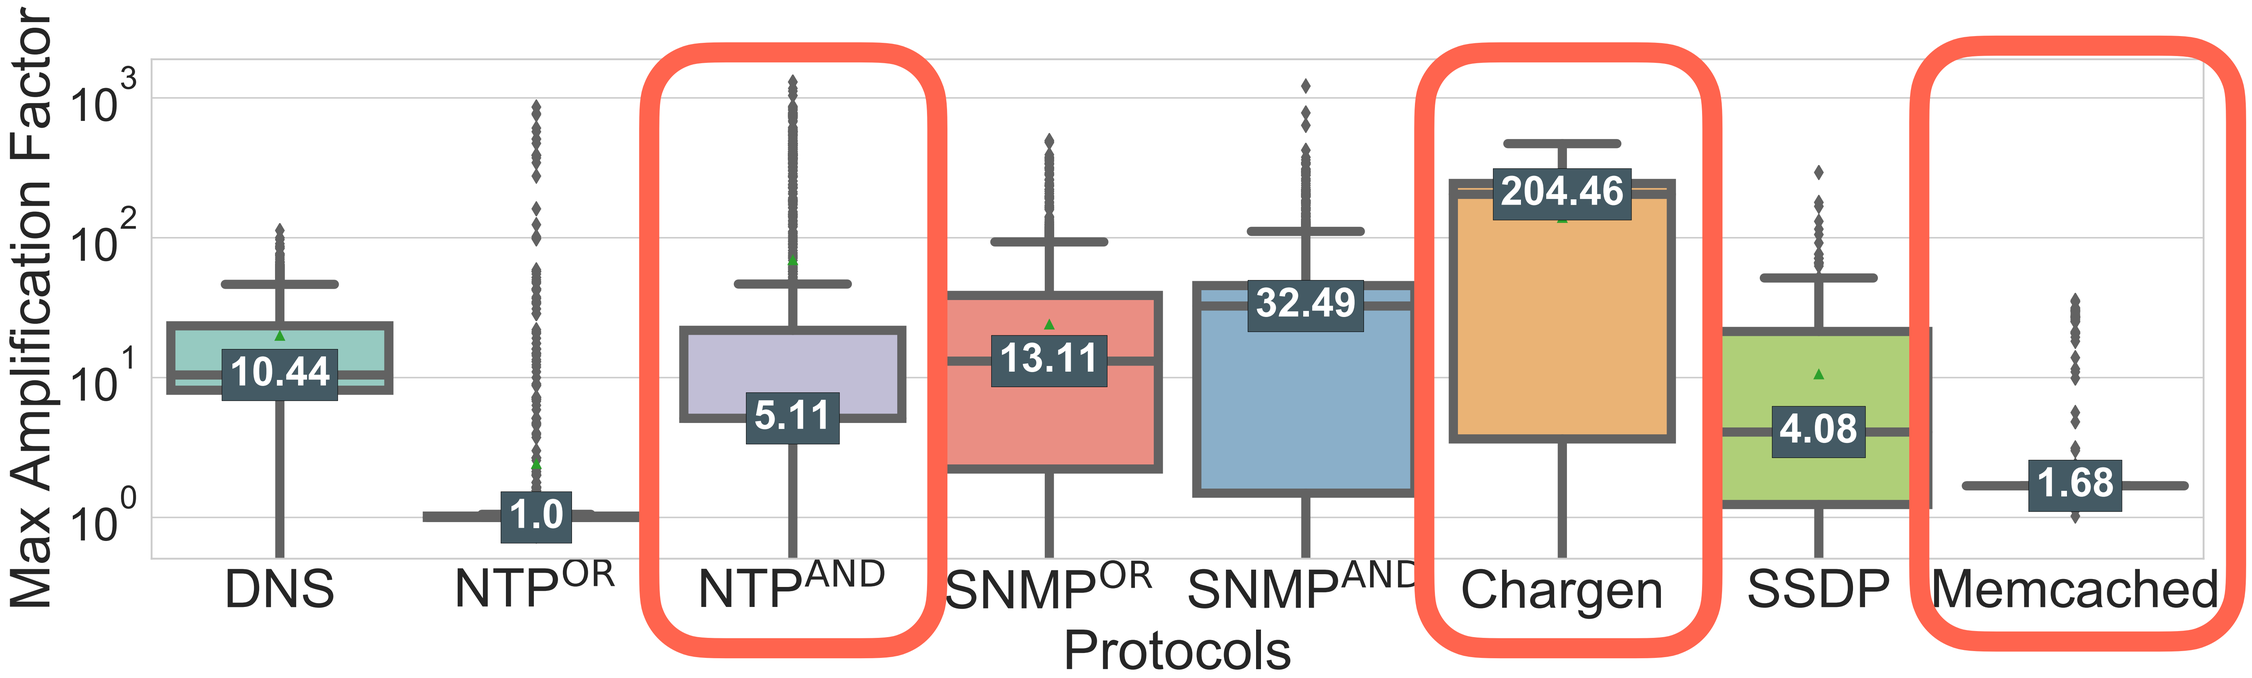

- **Across servers:**
  - NTP's median AF is only 5.11 but 1,300 AF for the max across measured servers.
- **Across protocols:**
  - Median AF for 5.11 for NTP vs. Chargen is 204.46 vs. 1.68 for Memcached.

Significant Diversity across Servers & Protocols!

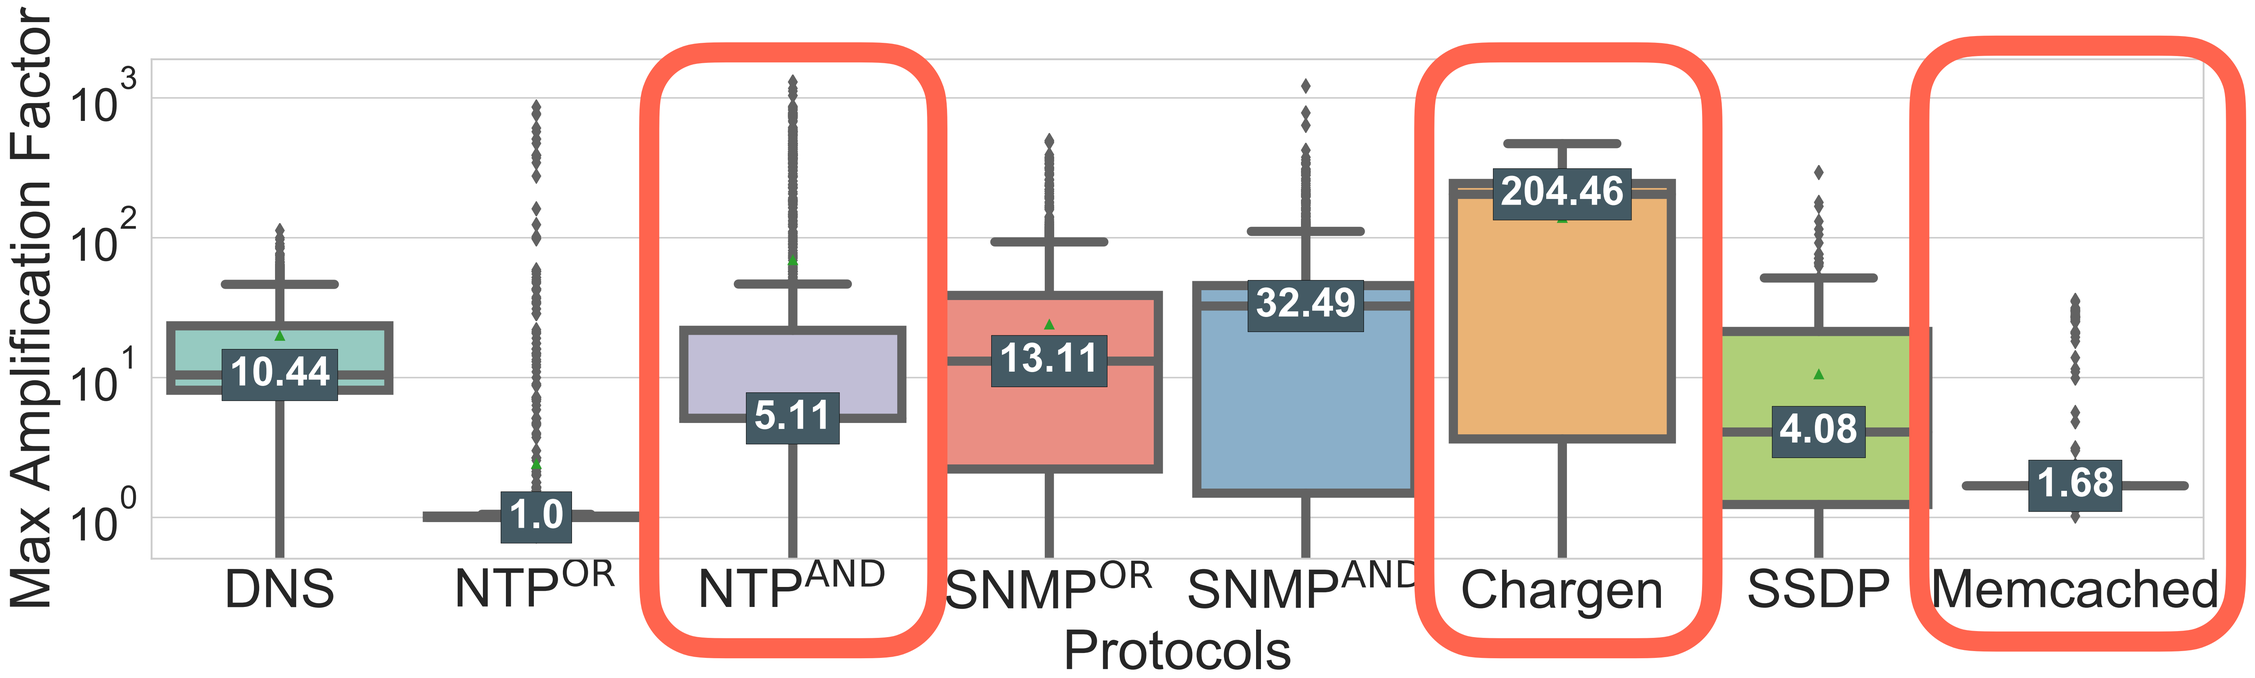

- **Across servers:**
  - NTP's median AF is only 5.11 but 1,300 AF for the max across measured servers.
- **Across protocols:**

Cannot assign the identical risk across servers and protocols

Prior Analysis Misestimation

DNS

Known → EDNS:0, RecordType: ANY | TXT (1.9x **over**-approx.)

New → EDNS:1 or Other RecordTypes

NTP

Known → Monlist (427x **over**-approx.)

New → GetRestrict, If Stats, etc.

SNMP

Known → GetBulk (3.5 X **under**-approx.)

New → GetNext

Prior Analysis Misestimation

DNS

Known → EDNS:0, RecordType: ANY | TXT (1.9x **over**-approx.)

New → EDNS:1 or Other RecordTypes

(**21.9x more risk** than the known modes)

NTP

Known → Monlist (427x **over**-approx.)

New → GetRestrict, If Stats, etc.

(**3.3x more risk** than the known mode)

SNMP

Known → GetBulk (3.5 X **under**-approx.)

New → GetNext

(**0.27x** risk of the known mode)

Implications of Our Findings

- Our **findings** imply:
  - Blocking or rate-limiting one mode still leave significant residual risk
  - Need to consider new defenses (e.g., new signature generation)

Implications of Our Findings

- Our **findings** imply:
  - Blocking or rate-limiting one mode still leave significant residual risk
  - Need to consider new defenses (e.g., new signature generation)

- To **accurately quantify amplification risk:**
  - Need to handle server heterogeneity (given a single mode)
  - Need to achieve coverages across multiple (unforeseen) modes

Conclusions & Takeaways

- DDoS amplification attacks continue to cripple our Internet
- Today: lack a systematic mechanism to precisely quantify the amplification risk
- **AmpMap:** A low-footprint measurement system to quantify amplification risk
  - Use **structural insights** to tackle the combinatorial explosion of input & server space
- Our measurements reveal:
  - Uncovered new amplification modes across protocols.
  - Uncovered significant diversity in amplification risk across servers and protocols.
  - Demonstrated that using prior analysis significantly mis-estimates the risk.
- Our findings imply the need for new defenses (e.g., new signature generation)

www.ampmap.net

https://github.com/ampmap-cmu/ampmap