Internet Scale Reverse Traceroute

IETF 117 matrg

Kevin Vermeulen (LAAS-CNRS)
Ege Gurmericliler (Columbia University)
Italo Cunha (UFMG)
David Choffnes (Northeastern University)
Ethan Katz-Bassett (Columbia University)

Today, we cannot measure reverse paths at scale

- Traceroute only gives forward paths

- Reverse Traceroute (NSDI 2010) can measure the reverse path but has limitations

Hard to troubleshoot paths with half visibility

Traceroute reveals that the forward path is not circuitous

Reverse path remains invisible with traceroute

47% of paths are asymmetric at AS level, and much more at router level

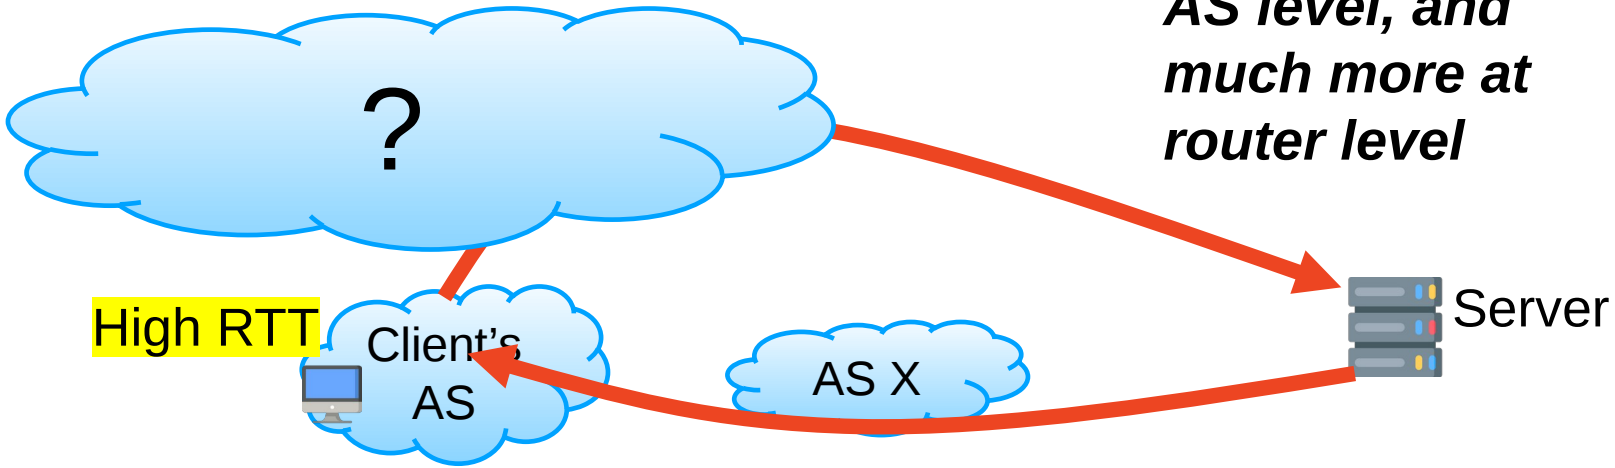

Hard to troubleshoot paths with half visibility

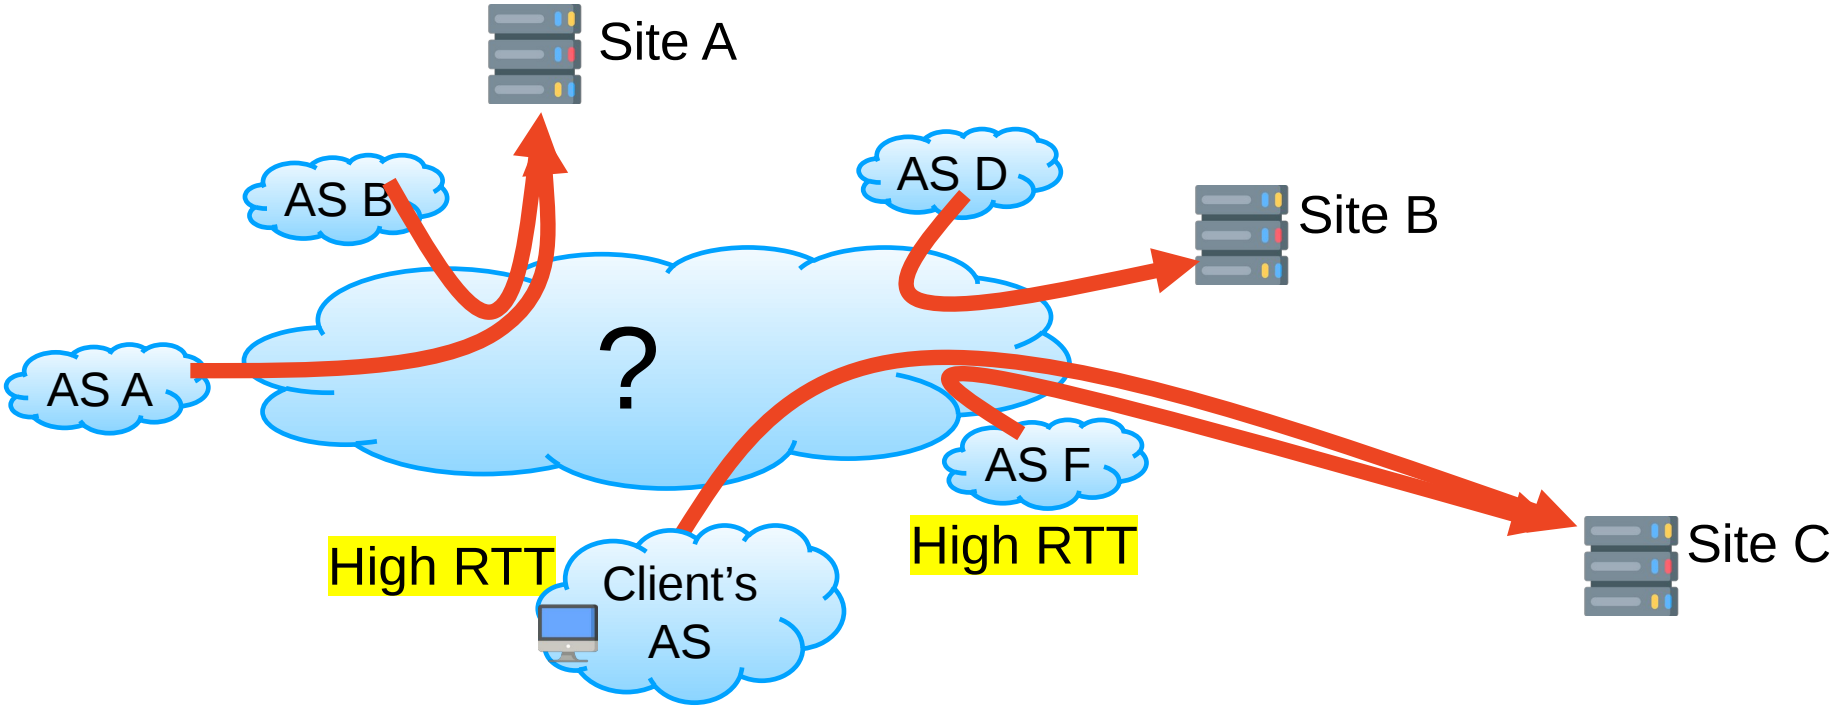

Reverse Traceroute (NSDI 2010 aka REVTR 1.0)

- **Accuracy**: Sometimes returned incorrect paths without warning, so one could not trust any measurement
- **Throughput**: Could not measure many reverse paths per unit time because each required many probes
- **Coverage:** Could only measure reverse paths to PlanetLab and M-Lab sources, not open to outside sources

Contributions

- REVTR 2.0, an Internet Scale Reverse Traceroute system
  - Throughput: 15M reverse paths per day
  - Accuracy: < 1.5% incorrect paths at AS level
  - Coverage: ~40K (55%) ASes in the Internet and allows outside sources
- REVTR 2.0 enables new research
  - Example of troubleshooting poorly performing paths
  - Large-scale study of Internet path asymmetry

Outline

- Motivation
  - Hard to troubleshoot paths with half visibility
  - REVTR 1.0 limitations
- Contributions
- REVTR 2.0: Internet Scale Reverse Traceroute
- Evaluation
- Troubleshooting poorly performing paths with REVTR 2.0
- Conclusion

Building REVTR 2.0

- Uses basic techniques from REVTR 1.0
- Identify key design questions of a Reverse Traceroute system to overcome REVTR 1.0's limitations
- Insights from our earlier work
  - Destination-based routing allows to piece paths hop by hop (IMC 2012)
  - Record Route coverage allows to measure from many clients (IMC 2017)
  - Paths stability allows to cache measurements for a day (IMC 2020)
- New measurement techniques and studies

- Want path from **D** back to **S**, don’t control **D**, so can’t use traceroute

BASIC REVTR 1.0 TECHNIQUES

- Traceroute from all vantage points to **S**
- Gives atlas of paths to **S**;
  if we intersect one, we know rest of path
- Traceroute atlas gives baseline

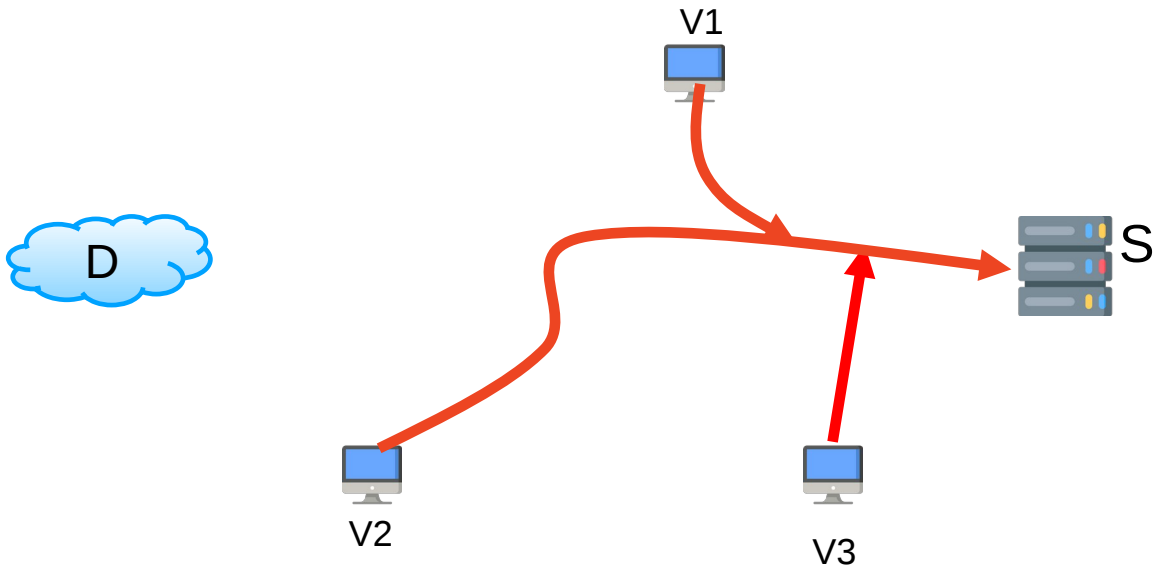

How do we find reverse hops?

- IP Options are reflected in reply
- Record Route IP Option (RR)
  - Record first 9 routers
  - If **D** within 8, reverse hops fill rest of slots
  - … but average path is > 8 hops
  - What if D is not within 8 RR hops from S?

BASIC REVTR 1.0 TECHNIQUES

- Traceroute atlas gives baseline
- Record Route IP option

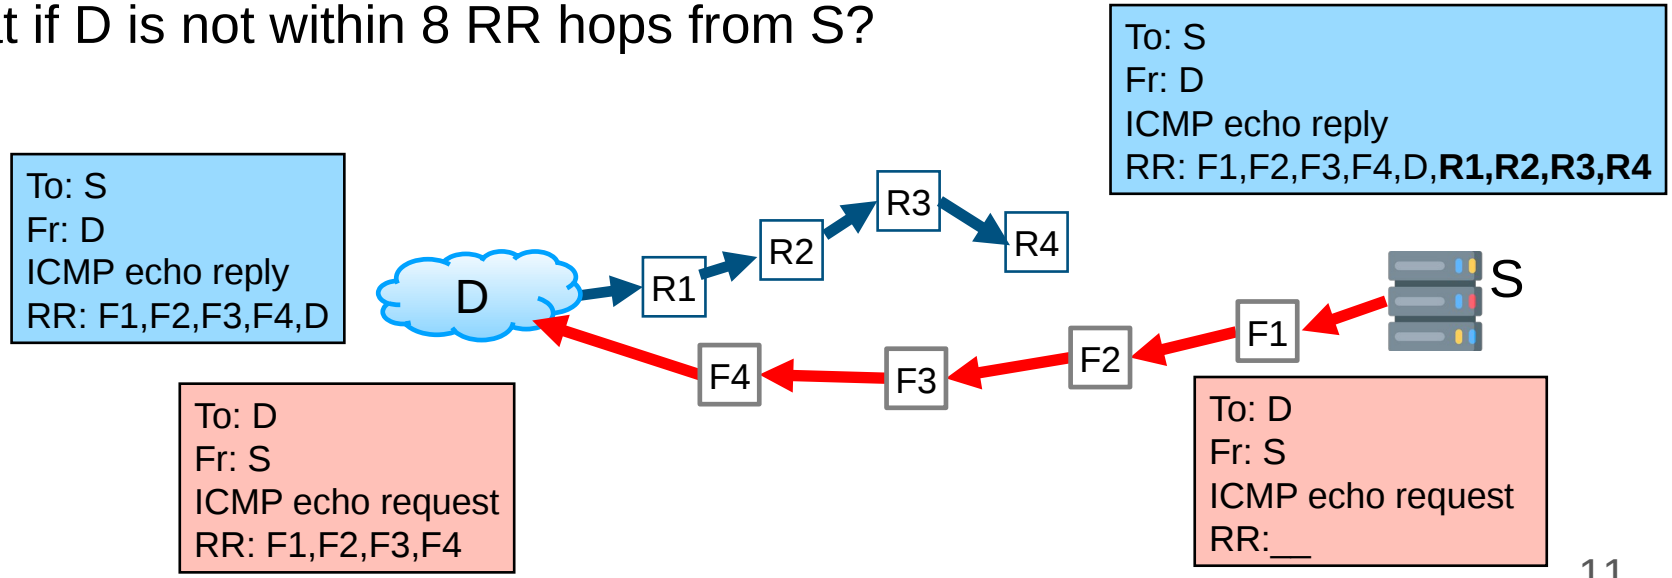

- From vantage point within 8 hops of **D**, ping **D** spoofing as **S** with Record Route
- **D**'s response records hop(s) on return path

BASIC REVTR 1.0 TECHNIQUES

- Traceroute atlas gives baseline
- Record Route IP option
- Spoofing as S

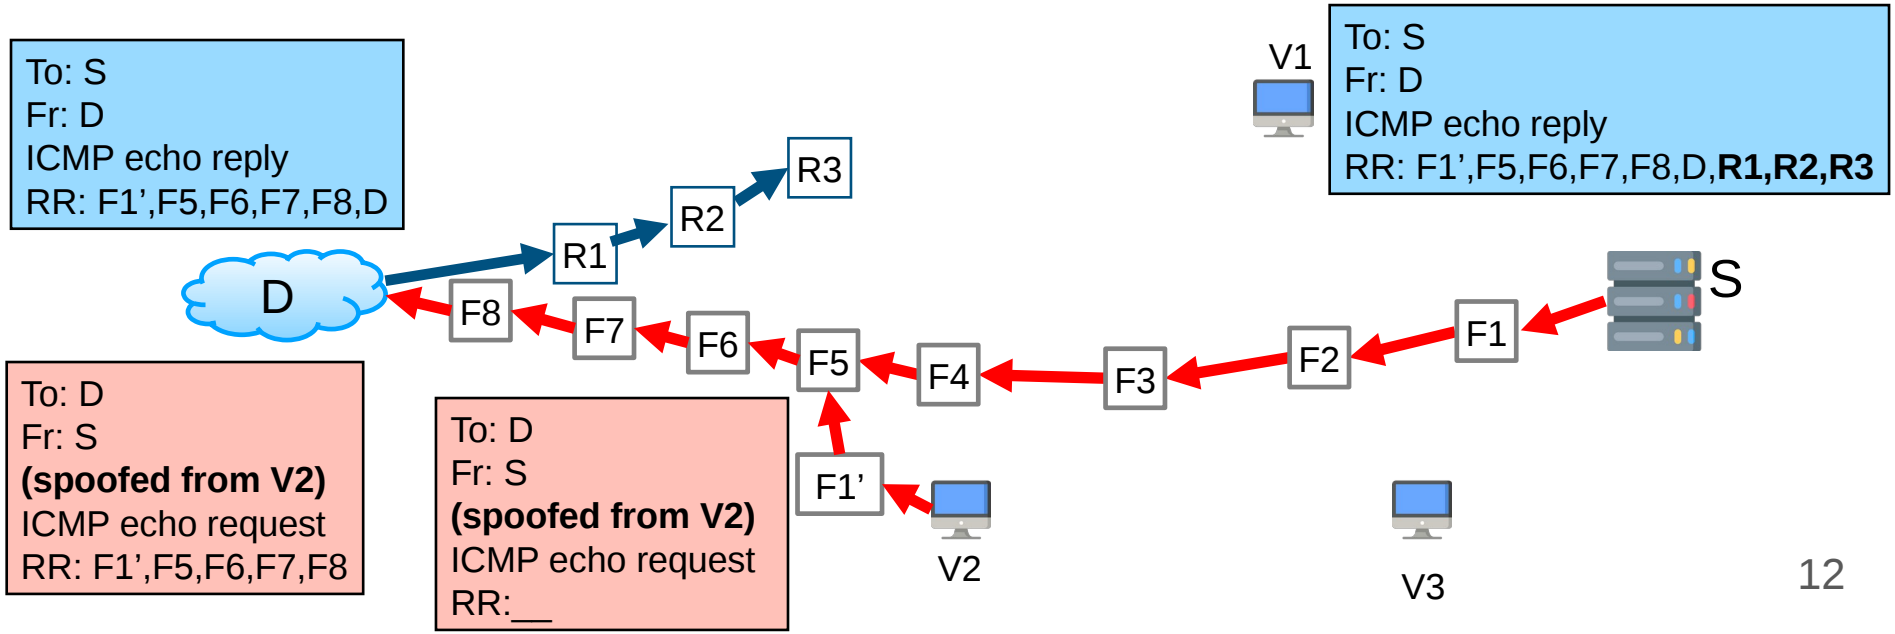

BASIC REVTR 1.0 TECHNIQUES

- Traceroute atlas gives baseline
- Record Route IP option
- Spoofing as S

- Continue to send Record Route until we intersect a traceroute
- Intersects a traceroute from V2 in R6
- The rest of the path follows the traceroute

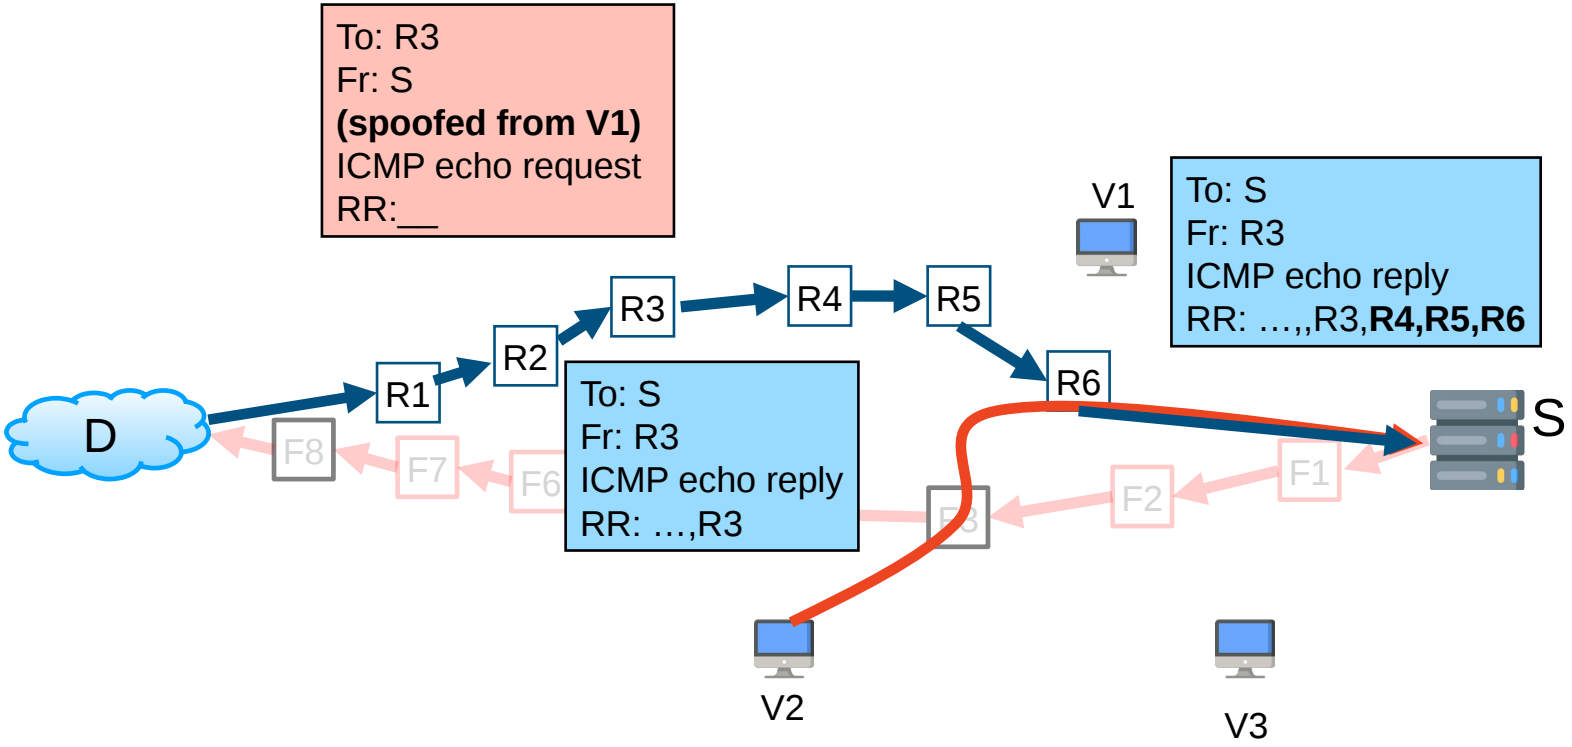

REVTR 2.0 DESIGN QUESTIONS

- Which traceroutes to issue for the traceroute atlas?
- How to identify intersections between Record Route and the traceroute atlas?
- Which vantage points to choose when spoofing?

BASIC REVTR 1.0 TECHNIQUES

- Traceroute atlas gives baseline
- Record Route IP option
- Spoofing as S

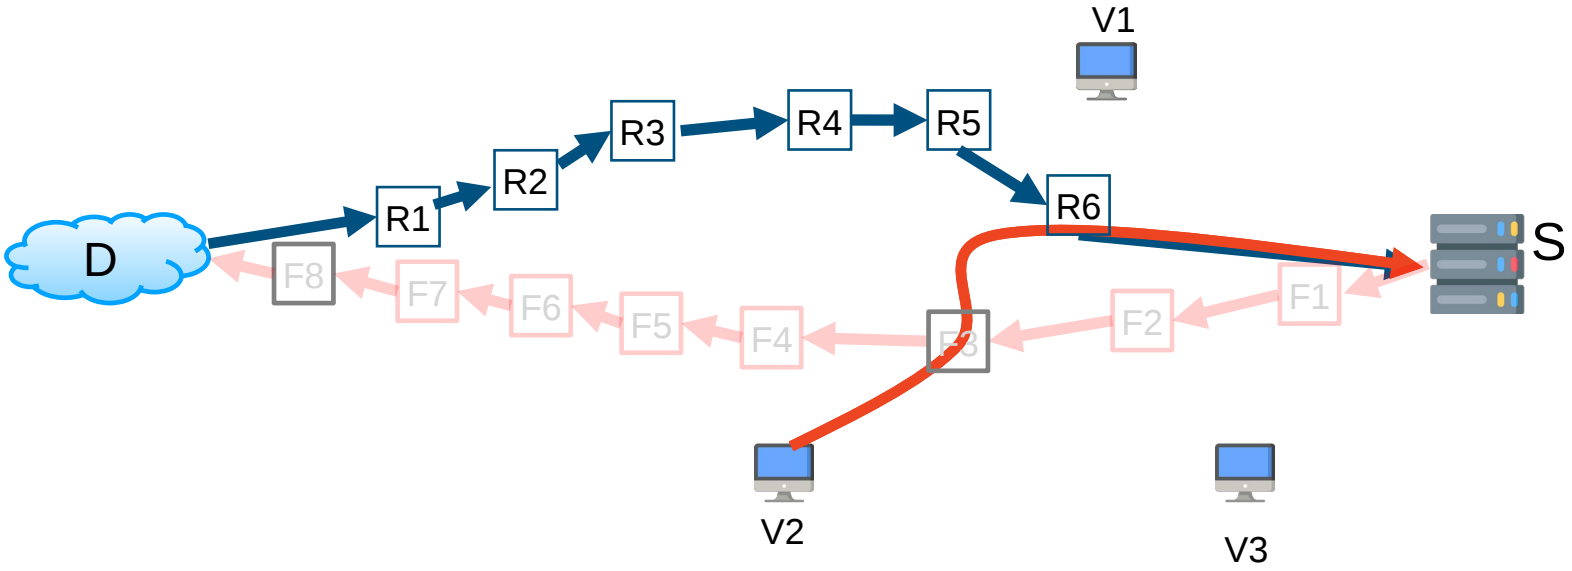

Outline

- Motivation
- Contributions
- REVTR 2.0: Internet Scale Reverse Traceroute
  - Studies and techniques to answer design questions
- Evaluation
- Troubleshooting poorly performing paths with REVTR 2.0
- Conclusion

REVTR 2.0 DESIGN QUESTIONS

- **Which traceroutes to issue for the traceroute atlas?**
- How to identify intersections between Record Route and the traceroute atlas?
- Which vantage points to choose when spoofing?

BASIC REVTR 1.0 TECHNIQUES

- Traceroute atlas gives baseline
- Record Route
- Spoofing

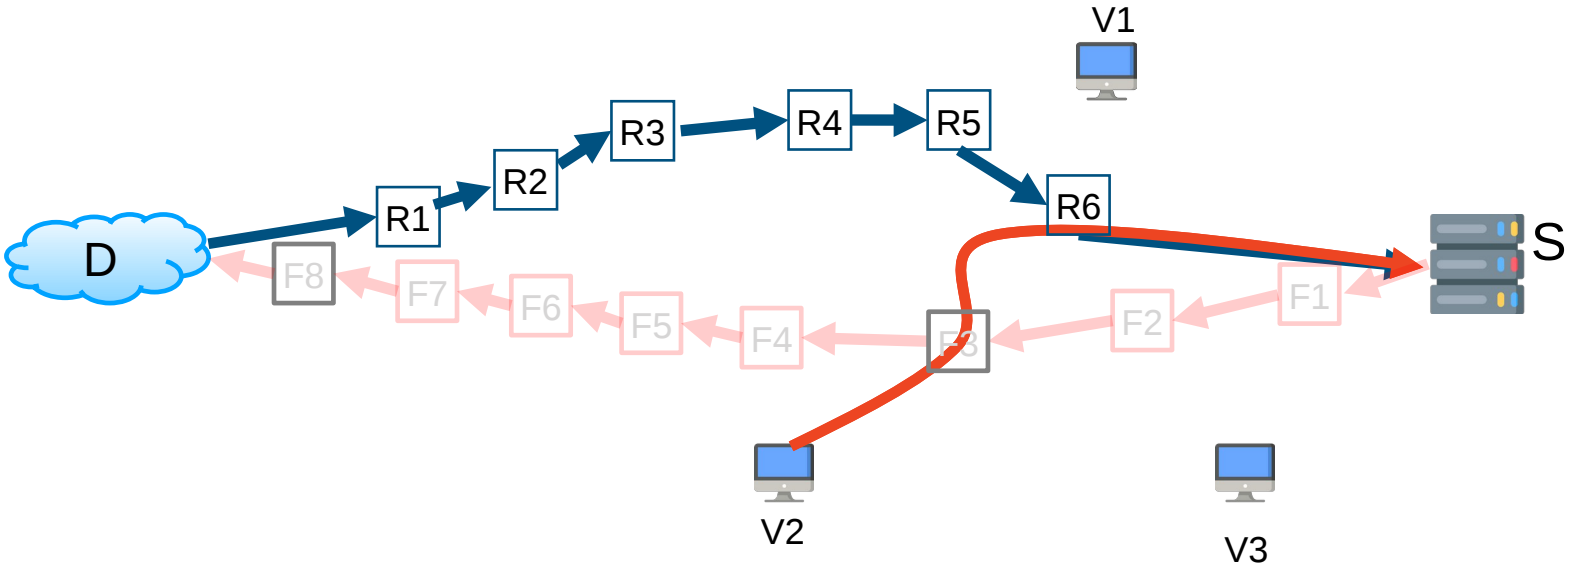

Which traceroutes to issue for the atlas?

- Challenge
  - RIPE Atlas is rate limited
- Goals:
  - Maximize the hops coming from traceroutes rather than from Record Route probes to improve throughput
  - Minimize the number of stale traceroutes to keep high accuracy

A random selection of 1,000 traceroutes per source is almost optimal to capture the most important hops

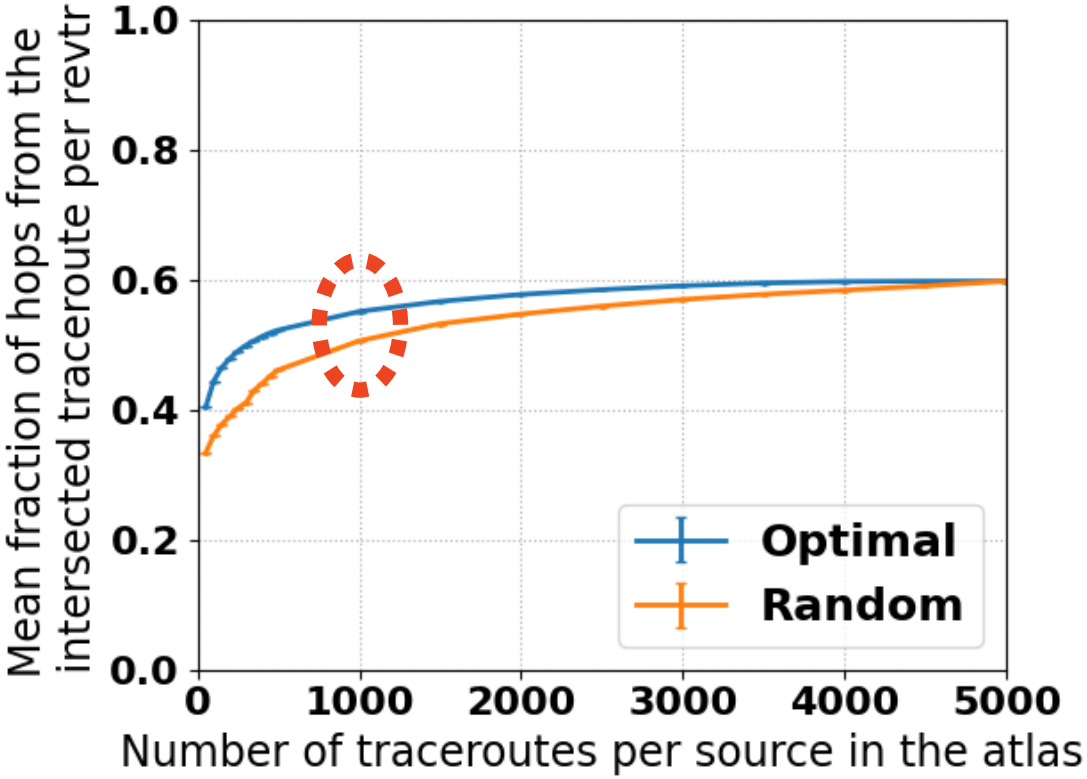

- Only 0.7% of reverse traceroutes intersect a stale traceroute if we keep traceroutes in the atlas for a day
- The atlas converges to optimal in 5 days

Which traceroutes to issue for the atlas?

- Answer:
  - REVTR 2.0 measures from 1,000 random RIPE Atlas vantage points to each source daily, replacing redundant traceroutes with new vantage points each day.
- Impact:
  - Throughput
    - 56% of the hops of a REVTR 2.0 measurement comes from the traceroute atlas on average, saving Record Route probes

Outline

- Motivation
- Contributions
- REVTR 2.0: Internet Scale Reverse Traceroute
- Evaluation
- Troubleshooting poorly performing paths with REVTR 2.0
- Conclusion

REVTR 2.0 vs REVTR 1.0

	Throughput (paths per day)	Accuracy (errors at AS level)	Coverage (# of ASes)	Outside sources
REVTR 2.0	15M	< 1.5%	39,544	Yes
REVTR 1.0	0.35M	8.3%	Unknown	No

REVTR 2.0 vs other approaches

- More complete than RIPE Atlas
- More accurate than assuming symmetry on forward traceroutes

Outline

- Motivation
- Contributions
- REVTR 2.0: Internet Scale Reverse Traceroute
- Evaluation
- Troubleshooting poorly performing paths with REVTR 2.0
- Conclusion

Troubleshoot poorly performing paths with REVTR 2.0

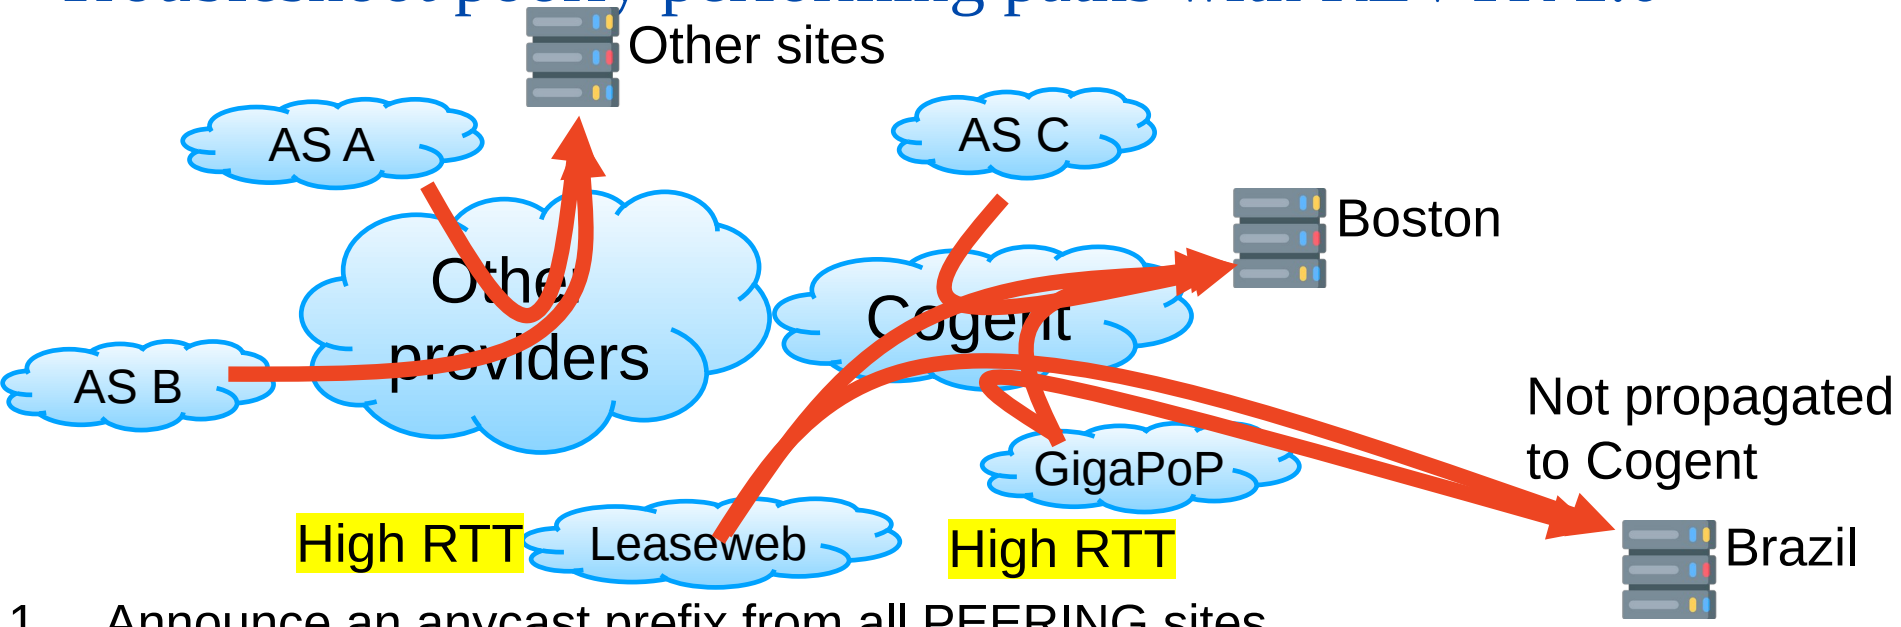

1. Announce an anycast prefix from all PEERING sites
2. Measure the performance of the clients: Leaseweb and GigaPoP (SE US) are routed to Brazil and have bad performance
3. Run REVTR 2.0 from clients: Leaseweb and GigaPoP are some Cogent clients routed to Brazil
4. Restrict the Brazil announcement to not propagate to Cogent

Troubleshoot poorly performing paths with REVTR 2.0

- With REVTR 2.0, you can find which routing decisions make some clients end up at a suboptimal site
- Accuracy allows you to trust the results
- Coverage allows you to measure reverse paths from many clients
- Throughput allows you to include reverse path measurements in reactive traffic engineering systems

Outline

- Motivation
- Contributions
- REVTR 2.0: Internet Scale Reverse Traceroute
- Evaluation
- Troubleshooting poorly performing paths with REVTR 2.0
- Conclusion

Takeaways

- Traceroute cannot give reverse path
- REVTR 2.0 overcomes REVTR 1.0’s limitations by improving accuracy and throughput and allowing outside sources
- Has been used in other contexts:
  - Traffic Engineering during failover (IMC 2022)
  - Used by a security company to uncover the providers of malicious ASes

REVTR 2.0 is an open system!

- Deployed on M-Lab (special thanks to the M-Lab team)
- Want to run reverse traceroutes?
  - Ask for API access to revtr@ccs.neu.edu
- Want to run reverse traceroutes to your own source?
  - Send an email to revtr@ccs.neu.edu with name, company/institution, and public IP address

Hard to troubleshoot paths with half visibility

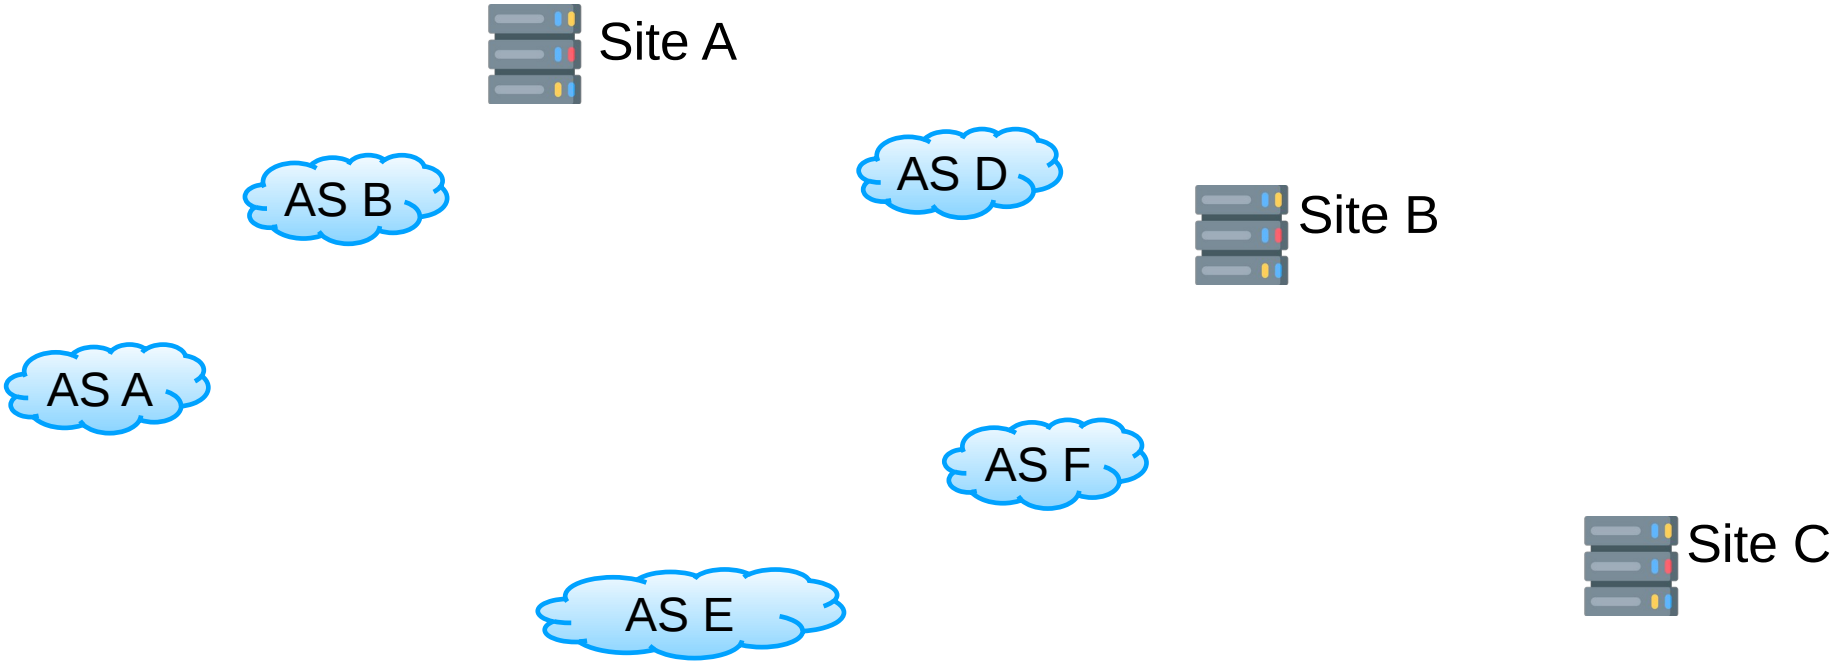

Hard to troubleshoot paths with half visibility

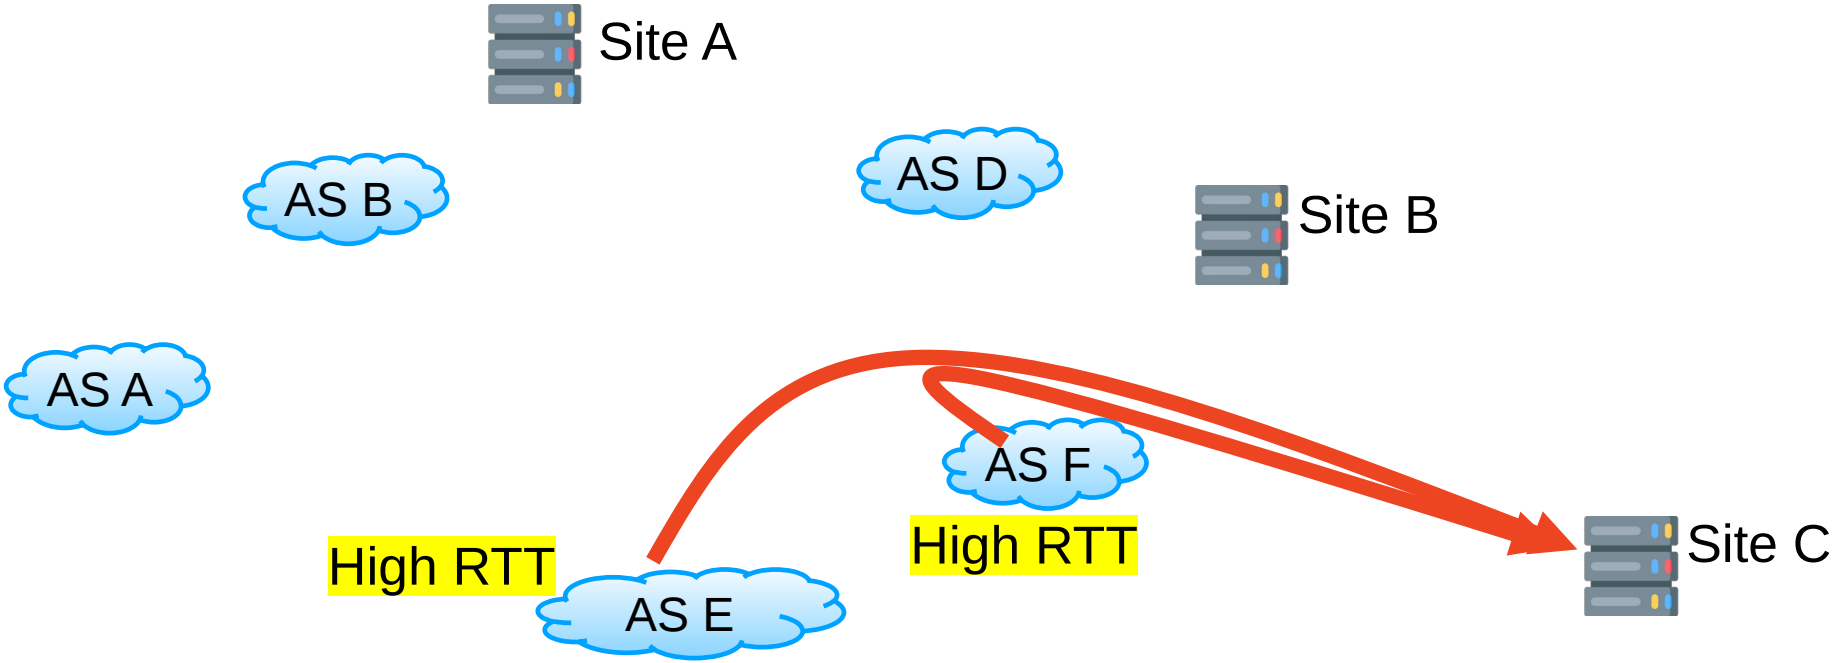

Hard to troubleshoot paths with half visibility

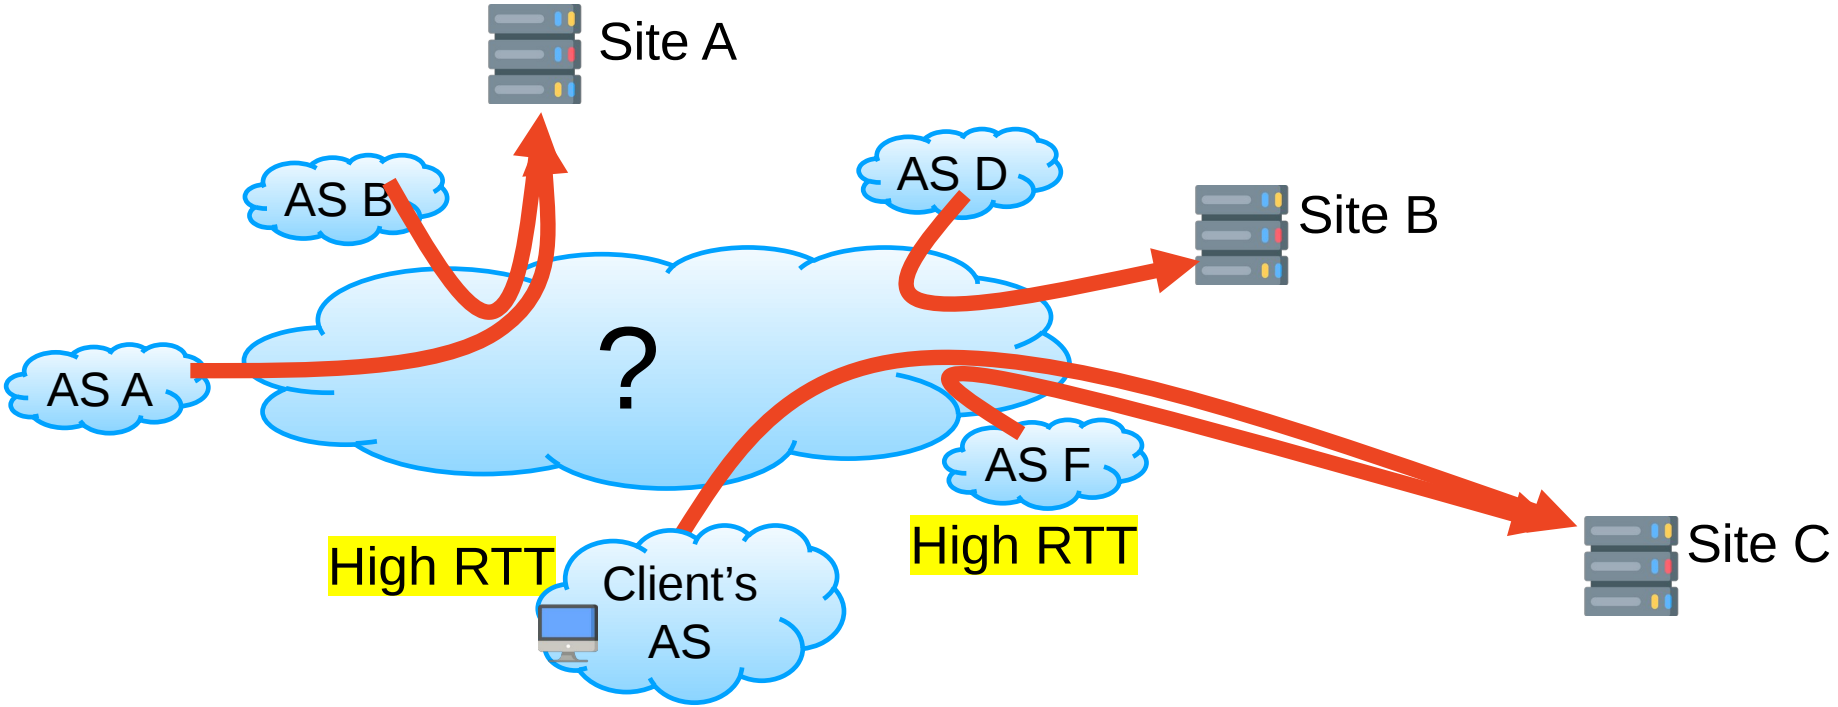

Hard to troubleshoot paths with half visibility