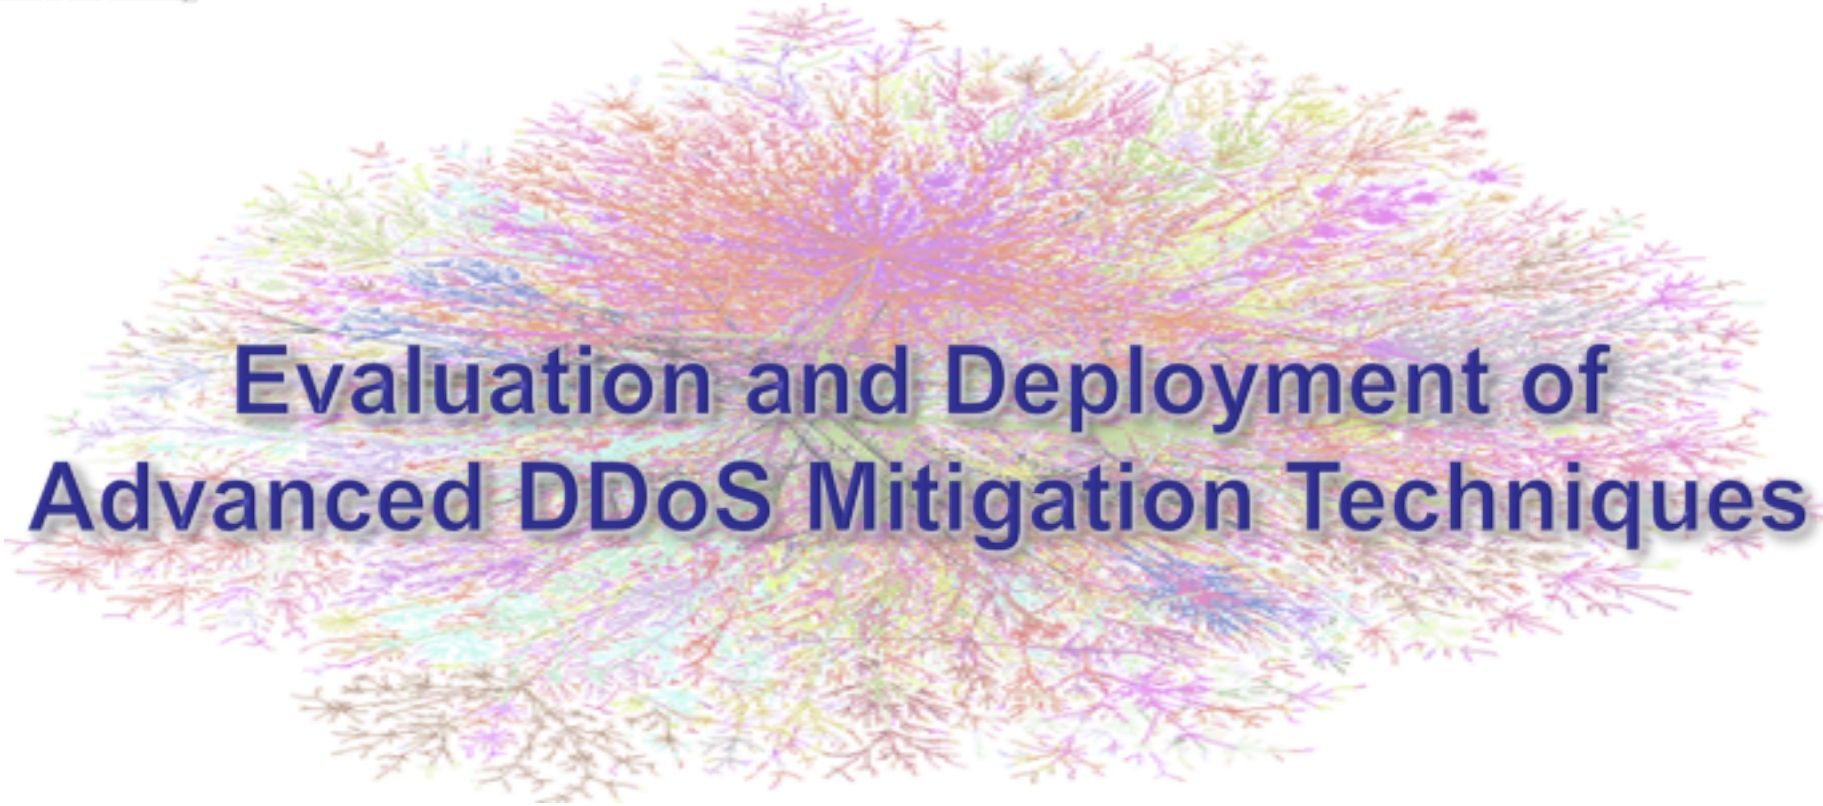

Evaluation and Deployment of Advanced DDoS Mitigation Techniques

Doug Montgomery, Kotikalapudi Sriram ({dougm | ksriram}@nist.gov)
Joint work with Mark Carson and Okhee Kim {carson | okim}@nist.gov
Advanced Network Technologies Division
Information Technology Laboratory
http://www.antd.nist.gov/

DDoS in the Headlines

All sites are vulnerable.

Average DDoS Attacks Now Large Enough to Take Most Organizations Completely Offline

Jul 19, 2016 6:19 PM PDT | Comments: 0 | Views: 1,567

— Average is large enough: "A 1 Gbps DDoS attack is large enough to take most organizations completely offline. Average attack size in 1H 2016 was 986Mbps, a 30% increase over 2015. Average attack size is projected to be 1.15Gbps by end of 2016."

http://www.circleid.com/posts/20160719_average_ddos_attacks_large_enough_to_take_most_organizations_down/

Attacks hit institutions and infrastructure.

US Congress Website Recovers from a Crippling 3-Day DNS Attack

Jul 20, 2016 12:11 PM PDT | Comments: 0 | Views: 1,464

By CircleID Reporter Comment | Print

A number of websites owned and operated by the United States Congress are recovering from a three-day DNS attack. Adam Mazmanian reporting in FCW: "The Library of Congress was the target of a denial-of-service attack that has knocked out Congress.gov and the U.S. Copyright Office website, and caused outages at other sites hosted by the library. Library spokesperson Gayle Osterberg told FCW that the DNS attack was launched July 17 and continues to affect library operations, including internal websites and employee email."

http://www.circleid.com/posts/20160720_us_congress_website_recovers_from_a_crippling_3_day_dns_attack/

Attacks are financially motivated.

DDoS is most common cyber attack on financial institutions

http://www.computerweekly.com/news/4500272230/DDoS-is-most-common-cyber-attack-on-financial-institutions

Bitcoin Hit By 'Massive' DDoS Attack As Tensions Rise

http://www.forbes.com/sites/leoking/2014/02/12/bitcoin-hit-by-massive-ddos-attack-as-tensions-rise/

Low barrier to entry for attackers.

Lizard Squad launches DDoS tool that lets anyone take down online services, starting at $6 per month

EMIL PROTALINSKI DECEMBER 30, 2014 8:37 AM

TAGS: DDOS, LIZARD SQUAD, MICROSOFT, PLAYSTATION 3, PLAYSTATION 4, PLAYSTATION NETWORK, SONY, TOP-STORIES, XBOX 360, XBOX LIVE, XBOX ONE

http://venturebeat.com/2014/12/30/lizard-squad-launches-ddos-tool-that-lets-anyone-take-down-online-services-starting-at-5-99-per-month/

DDoS Industry Views

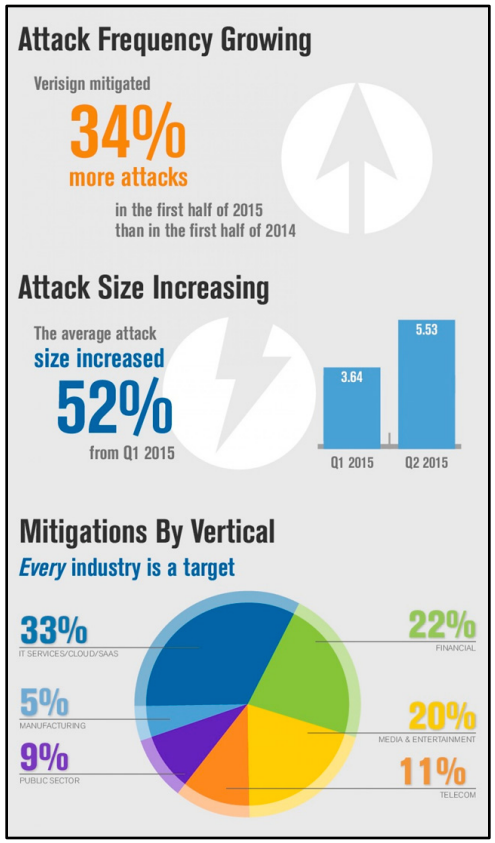

http://www.circleid.com/posts/20160525_ddos_trends_attack_activity_increases_111_percent_year_over_year/

http://www.circleid.com/posts/20150827_ddos_for_bitcoin_increasingly_targets_financial_industry/

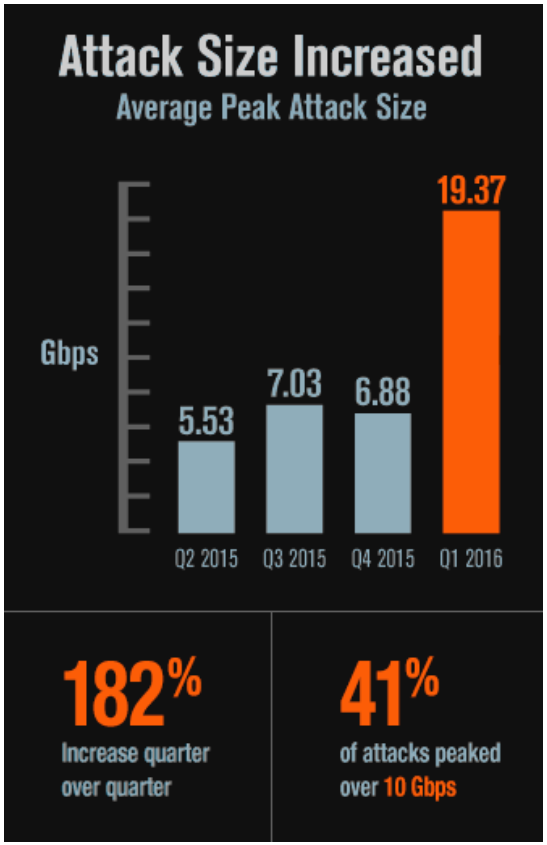

- Abor Networks: Worldwide Infrastructure Security Report.
  - https://www.arbornetworks.com/images/documents/WISR2016_EN_Web.pdf
- Verisign Distributed Denial of Service Trends Report
  - https://www.verisign.com/assets/report-ddos-trends-Q12016.pdf

A Quick Review of Our Goals

DDoS Mitigation Techniques

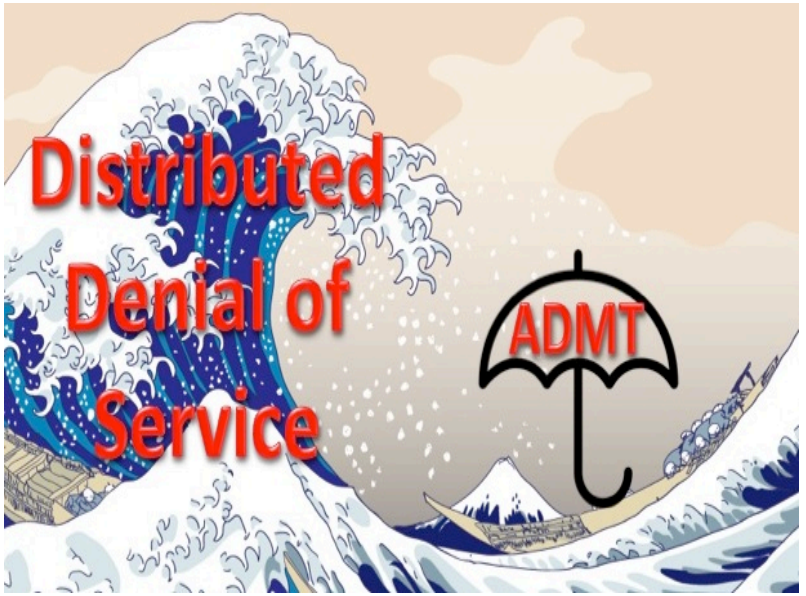

- **Advanced Techniques**
  - Focus of much research and development.
  - Detection, information sharing, rate limiting, distributed trace back and packet filtering.
- **Source Address Spoofing**
  - Enabler of reflection attacks.
  - Disabler of trace-back techniques.
  - **Focus on Source Address Validation (SAV)**.
- **SAV Existing Techniques**
  - Are / why aren't existing techniques (BCP-38,84) deployed?
  - What are operational impacts of deployment &management of SAV?
  - Are existing SAV techniques applicable to modern environments?

Isn't This a Solved Problem?

- **IETF / NANOG / RIPE Deployment Guidance**
  - BCP38 - *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing.* 2000.
    - https://tools.ietf.org/html/bcp38
  - BCP84 - *Ingress Filtering for Multihomed Networks.* 2004.
    - https://tools.ietf.org/html/bcp84
- **Why Aren't These Techniques Widely Deployed**
  - Maybe they are?
  - Much more needs to be known about the source of spoofed traffic, where it originates from and how.
  - Much more needs to be known about the cost / complexities of implementing SAV.

How To Implement SAV

- **Access Control Lists**
  - Only scalable for simple fixed subsets.
- **Reverse Path Checks (RPF)**
  - Leverages dynamic BGP control plane to maintain view of valid sources
  - **Strict RPF**
    - SA in the FIB and reachable by the interface that receives the packet.
  - **Loose RPF**
    - SA in the FIB and reachable over any interface.
  - **Feasible RPF**
    - SA in FIB and was announced over interface that receives the packet.
- **Link Level SAV**
  - IP Source Guard, DHCP Linked mechanism.
- **Other Approaches?**
  - SDX Filter Rules

Where to Implement SAV

- **BCP38 is only a vision..**
  - Reality is much more complicated.
  - Goal: SAV that is
    - As strict as possible.
    - As close to the source as possible.
    - Minimize performance impact.
    - Minimize management burden.
    - Minimize failure scenarios.
- **No single answer**
  - Each scenario is different.
  - Each mechanism is different.
  - Ingress / egress is different.
  - Each platform is different.

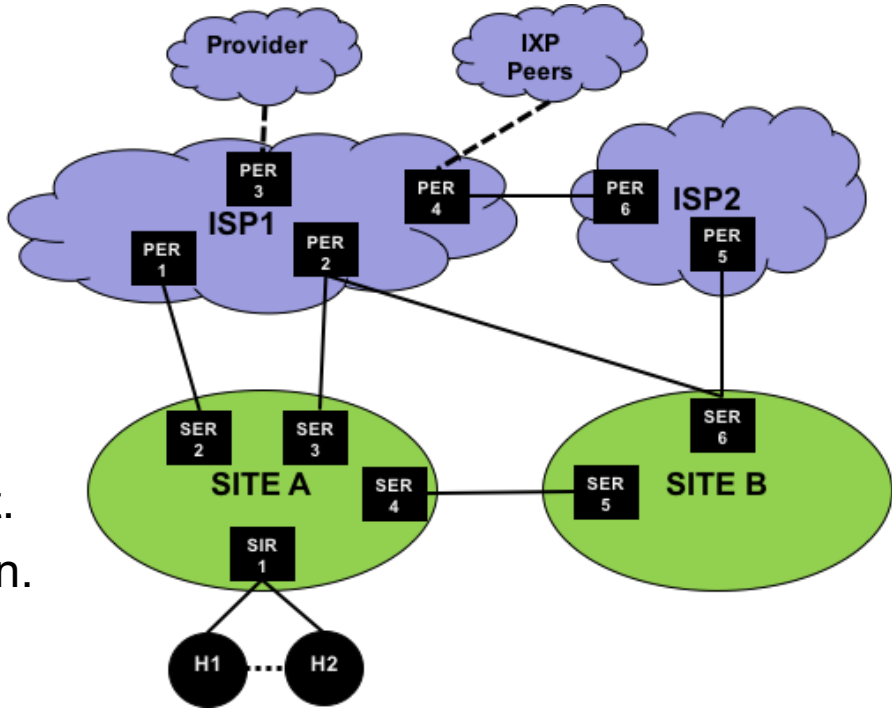

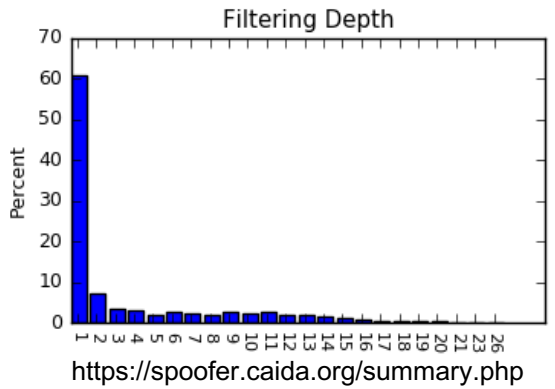

https://spoofer.caida.org/summary.php

Our Objectives

- **Independent Technical Evaluation of SAV**
  - **Quantitatively characterize SAV mechanisms**
    - Focus on those in current commercial products.
    - Applicability, effectiveness, and operational impacts.
    - Different deployment scenarios: enterprise (multi-homed), stub ISP, small transit ISP, large transit ISP.
    - Characterize complexity to manage filtering, performance impact on data plane.
  - **Focus on SAV at Domain Boundaries**
    - Other deployment scenarios (subnet, multi-tenant) later.
    - Develop workload models of inter-domain SAV.
- **SAV Deployment Guidance**

Projected Impact

- **Measurements:**
  - Set of NIST tools, techniques and data sets that can form the basis for repeatable quantitative measurements of SAV mechanisms in commercial networking products.
  - Set of vendor independent characterizations of the performance impact, management complexity and robustness of SAV filtering mechanisms.
- **Deployment Guidance:**
  - Published NIST technical guidance for the deployment of inter-domain SAV filtering.

SAV Metrics

- **Data Plane**
  - Throughput, latency, jitter, packet loss.
  - SAV computational load, power consumption.
- **Control Plane**
  - SAV configuration space / time complexity.
  - SAV volatility: rate and volume of change.
- **Robustness**
  - Correctness / completeness of SAV filtering data.
  - Risks of incorrect data.
  - Threats / vulnerabilities of accidental / malicious corruption of data.

SAV Workload Models

- **SA Workload Models**
  - Inter-domain SAV requires computing "customer cone" of operational ISPs.

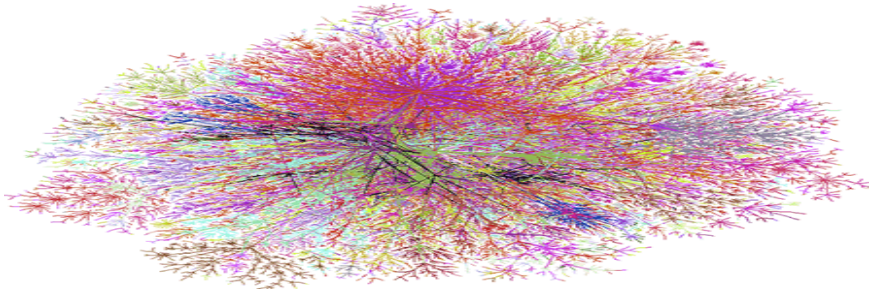

  - *SAV Customer Cone* – set of prefixes that describe the set of valid source addresses that may exit a given domain.
  - Questions of the source, accuracy, security, scale and volatility of information used to compute customer cones.
  - Develop a set of models representative of stub, simple transit, transit and tier 1 ISPs.
  - Develop tools to parameterize models from common information source.
- **See backup slides**

SAV Benchmark Methodology and Performance Measurements

SAV Implementation

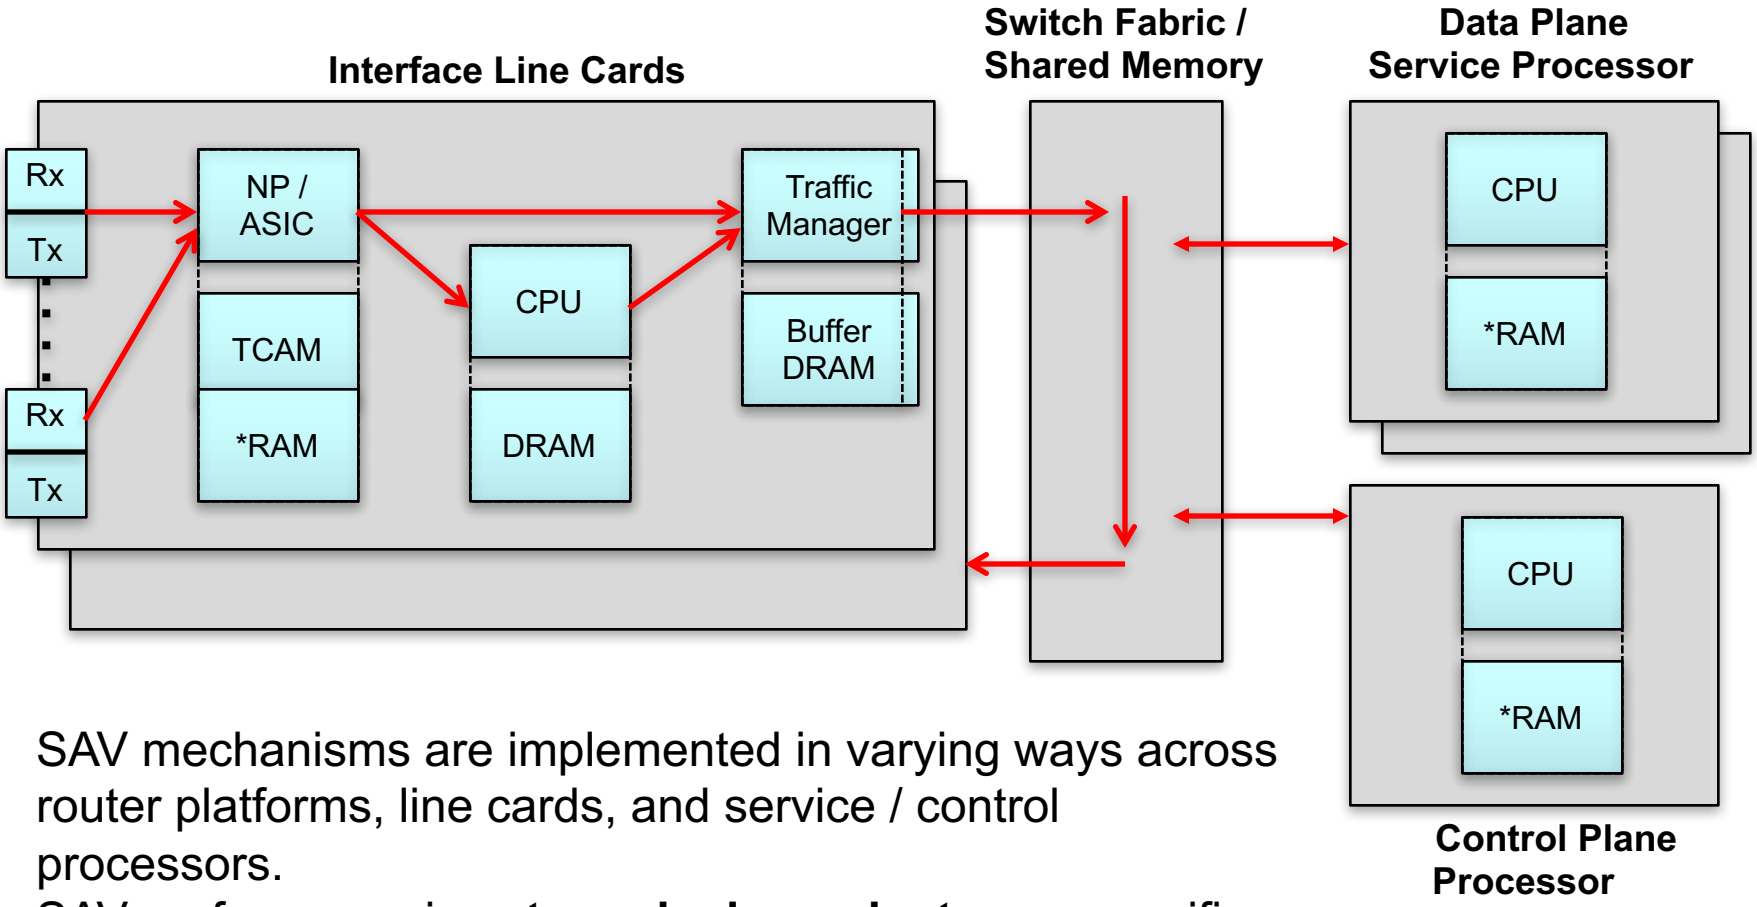

- SAV mechanisms are implemented in varying ways across router platforms, line cards, and service / control processors.
- SAV performance is **extremely dependent** upon specific components, port densities, memory sizes, etc.
- Even with a single hardware configuration, specific loading / traffic scenarios can produce vastly different results.

Initial (Poor Man's) Approach

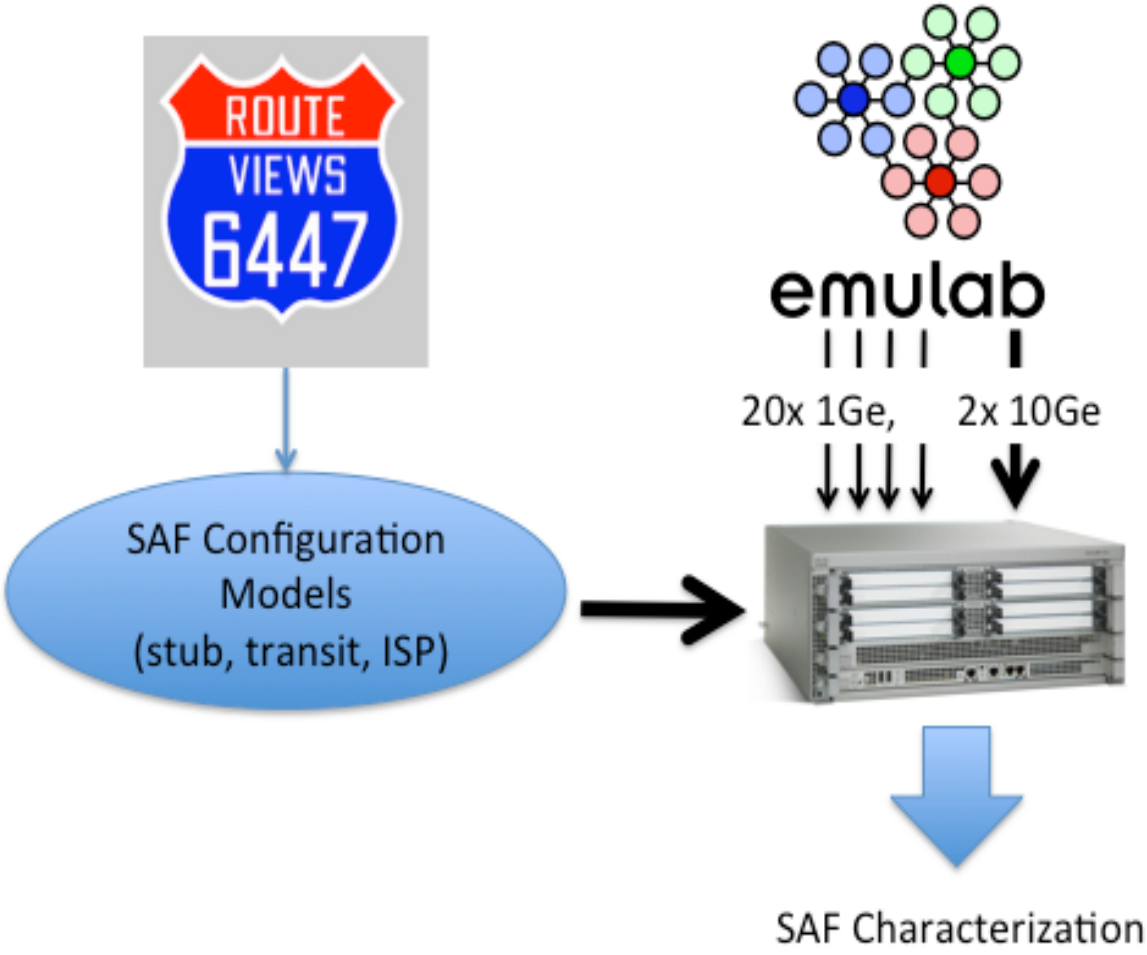

Emulab-Based Testbed

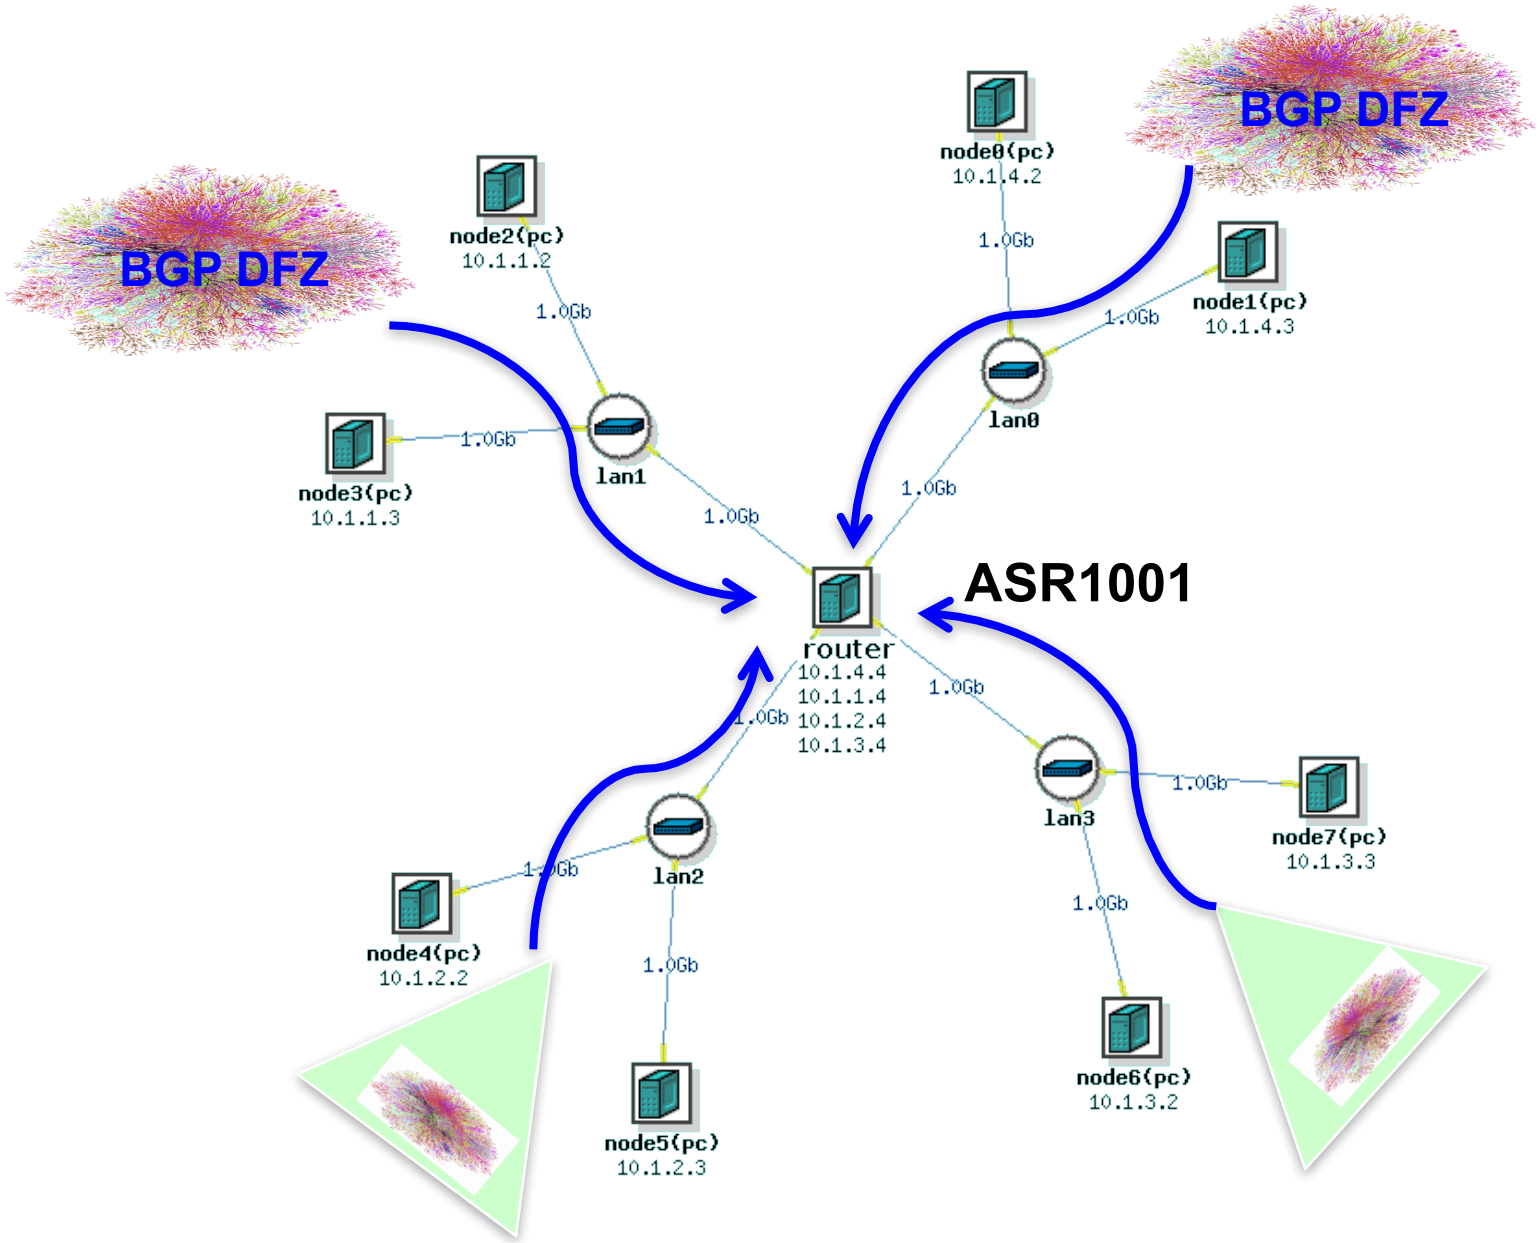

Scaling Up & Reaching Out

- **Poor men can't test fast routers**
  - Ixia XGS-12 - $300K
    - 4 x 40Ge
    - 16 x 10Ge
    - 8 x 1Ge

- **Collaboration with Router Vendors**
  - Develop an open source set of test scenarios, metrics, and reporting formats for SAV benchmarking.
  - NIST will develop test scripts (Ixia / open source) to execute tests.
  - Vendors execute tests across their product line.

Ixia / Emulab Hybrid Testbed

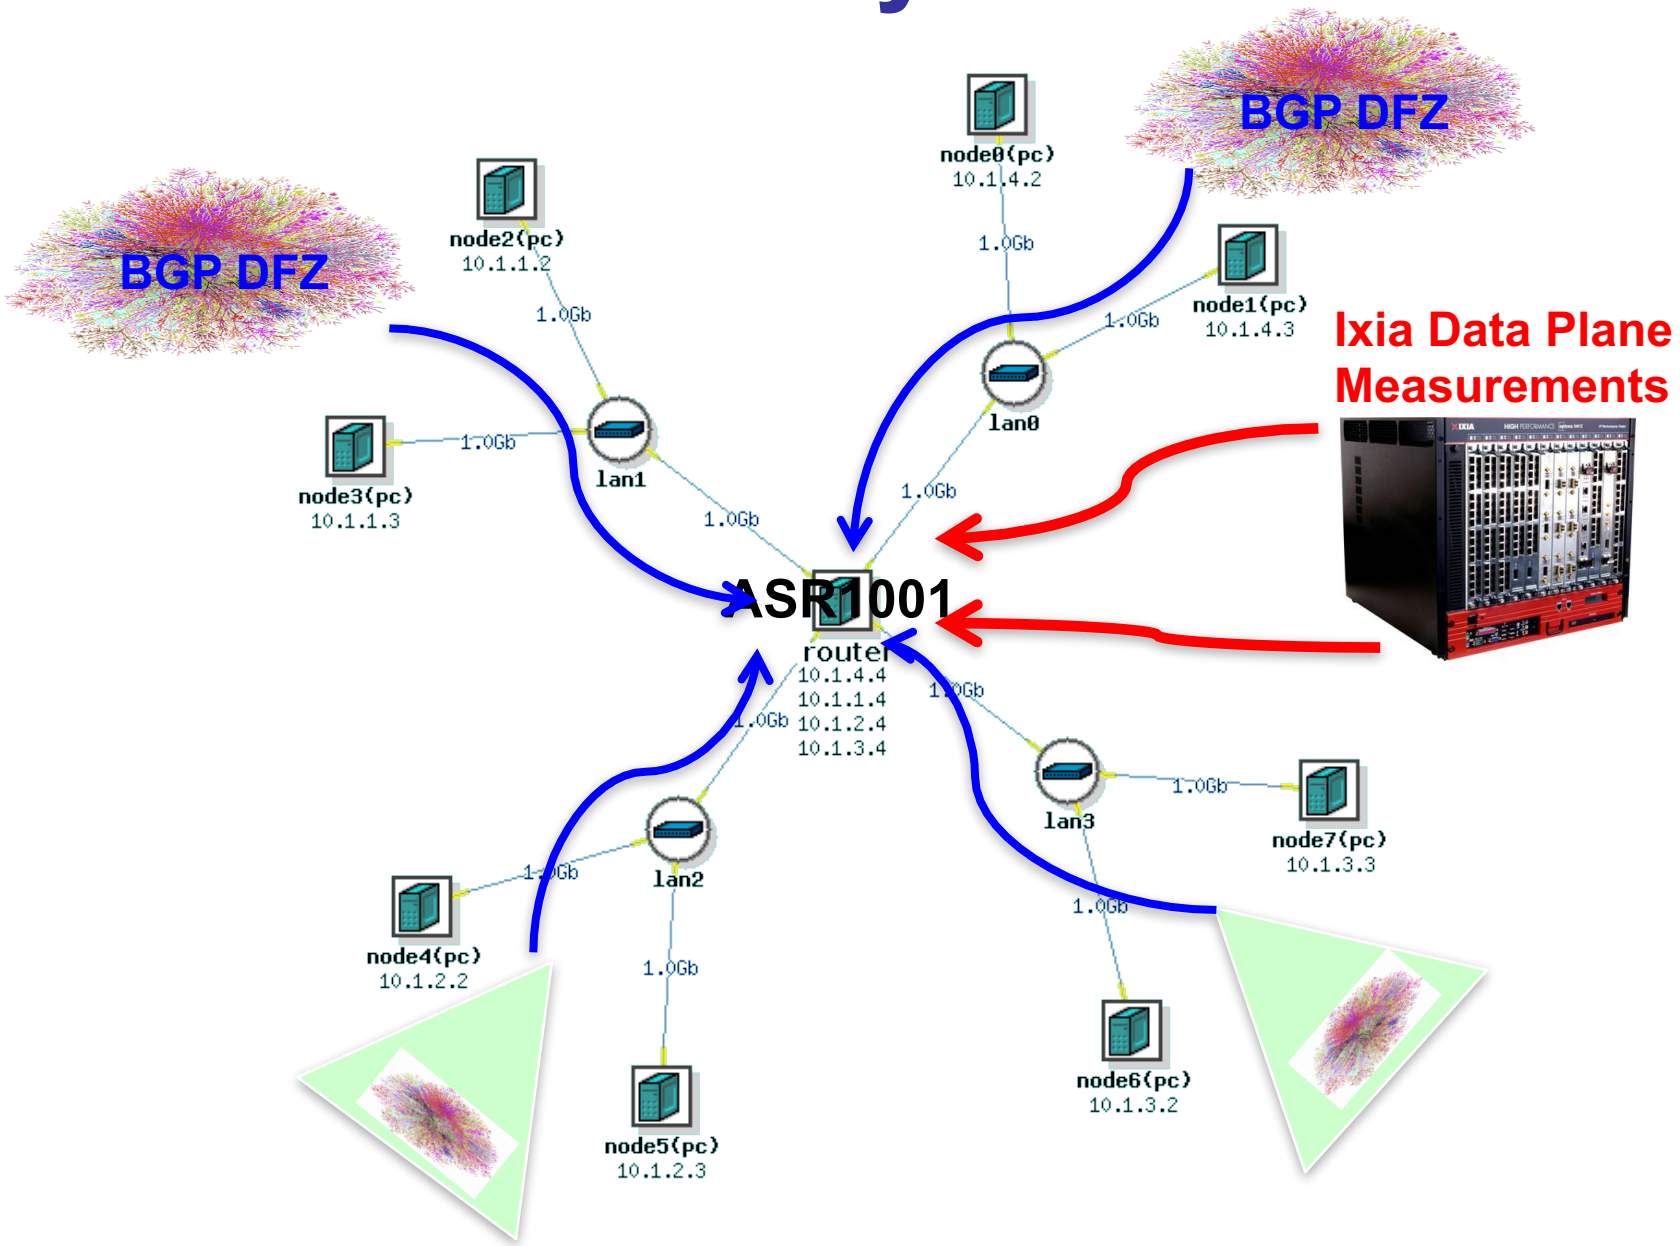

Initial SAV Benchmark Tests

- **Router Performance Baseline**
  - Focus on PPS/Mbps forwarding rates.
- **Performance impact of uRPF mechanisms?**
  - Strict mode; feasible (semi-loose) mode; loose mode
- **Impact of of BGP routing table on uRPF?**
  - Test range of FIB sizes
    - Including those larger than TCAM.
    - "Spray" source and destination addresses over varying address ranges.
- **Impact of varying % of spoofed packets?**
  - Cost of SAV filtering of spoofed traffic.
- **Impact of varying packet size?**

Initial DUT

- Cisco ASR1001
  - 5 Gbps ESP throughput
  - 4 x 1G ports
  - 5Mb TCAM
  - 65K IPv4 routes

SAV Measurements: ASR1001 (5 GHz ESP) with Load from Ixia box (Part 1 of 3)

- uRPF turned OFF
- IPv4
- Packet size: 74B (plus 20B PHY overhead)
- ACL, QoS, Netflow OFF

	BGP table size	Tx Rate (Mbps)	Rx Rate (Mbps)	Loss (%)	Rx (Tx) Bandwidth utilization
no BGP (static routes)	0	2636	2636	0.00	83.71%
random SAs, fixed DAs	32,000	2636	2636	0.00	83.71%
random SAs, fixed DAs	64,000	2636	2636	0.00	83.71%
random SAs, fixed DAs	128,000	2636	2636	0.00	83.71%
random SAs, fixed DAs	256,000	2636	2636	0.00	83.71%

SAV Measurements: ASR1001 (5 GHz ESP) with Load from Ixia box (Part 2 of 3)

- uRPF turned ON (strict mode)
- IPv4
- Packet size: 74B (plus 20B PHY overhead)
- ACL, QoS, Netflow OFF

	BGP table size	Tx Rate (Mbps)	Rx Rate (Mbps)	Loss (%)	Rx (Tx) Bandwidth utilization	Penalty due to uRPF
random SAs, fixed DAs	32,000	2394	2394	0.00	76.03%	9.18%
random SAs, fixed DAs	64,000	2393	2393	0.00	75.99%	9.22%
random SAs, fixed DAs	128,000	2393	2393	0.00	75.99%	9.22%
random SAs, fixed DAs	256,000	2392	2392	0.00	75.96%	9.26%

SAV Measurements: ASR1001 (5 GHz ESP) with Load from Emulab (Part 3 of 3)

- uRPF turned ON (strict mode)
- IPv4
- Packet size: 46B
- ACL, QoS, Netflow OFF

	Throughput rate (Mpps)
Without uRPF	5.38
With uRPF	4.44
% drop due to uRPF	17%

Initial Observations

- **Based on the measurements done so far:**
  - Throughput penalty due to uRPF observed!
    - 9% (74B, Ixia setup) to
    - 17% (46B, Emulab) for ASR1001.
  - Expected TCAM Caching Effects Not Observed!
    - BGP/FIB table size does not seem to impact performance.
    - Distribution of source / destination addresses used does not seem to effect performance.
- **Compare Initial Results with Industry Tests**
  - Verify we are on the right track before exhaustive tests.

IPv4 Measurements by Router Analysis, Inc.

ASR1004 Router with 40 GHz ESP & Load Generated by Ixia box

64B packets (minimum size)

	IPv4 Throughput (MPPS)
Forwarding only (Base case)	26.60
Forwarding with uRPF	23.45
Forwarding with ACL	21.77
Forwarding with ACL & uRPF	18.14

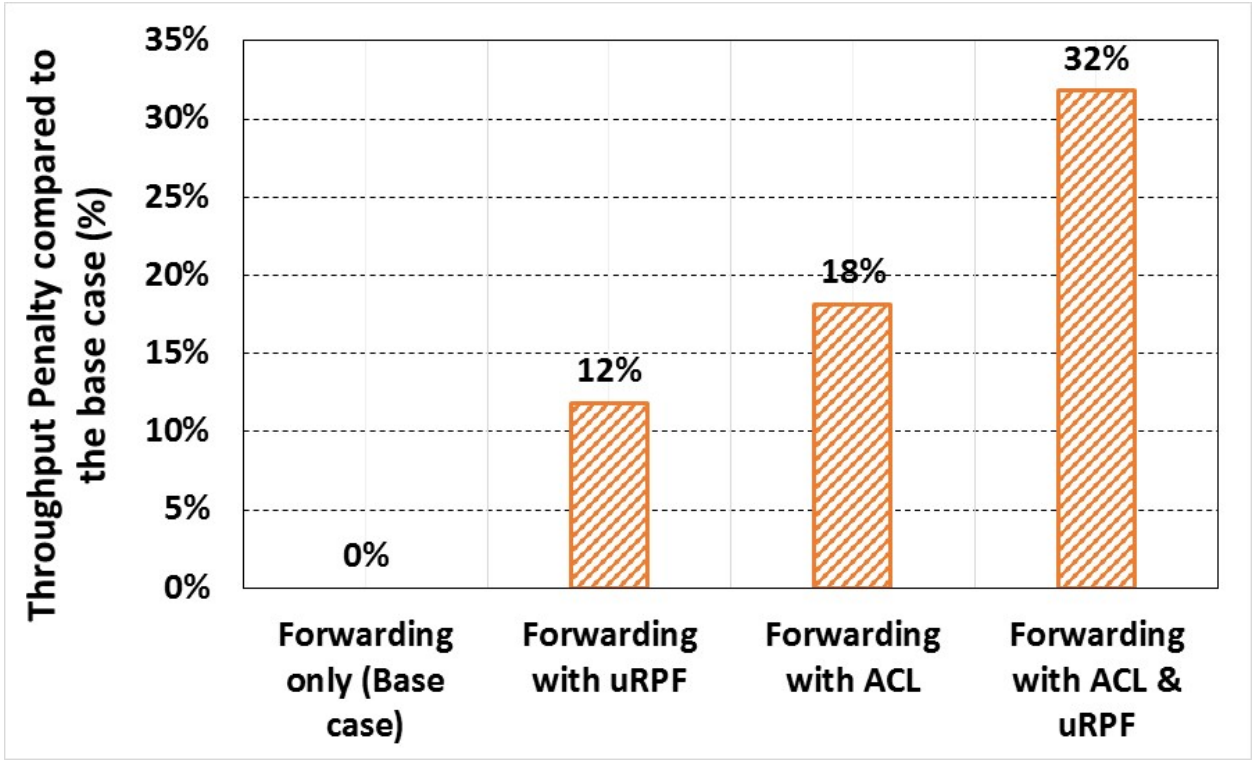

NIST plot using data from Router Analysis, Inc.
Ref: http://www.slideshare.net/RouterAnalysis/cisco-asr-1000-series-testing-results-and-analysis

NIST National Institute of Standards and Technology

INFORMATION TECHNOLOGY LABORATORY

IPv6 Measurements by Router Analysis, Inc.

ASR1004 Router with 40 GHz ESP & Load Generated by Ixia box

64B packets (minimum size)

	IPv6 Throughput (MPPS)
Forwarding only (Base case)	21.00
Forwarding with uRPF	17.50
Forwarding with ACL	14.90
Forwarding with ACL & uRPF	11.50

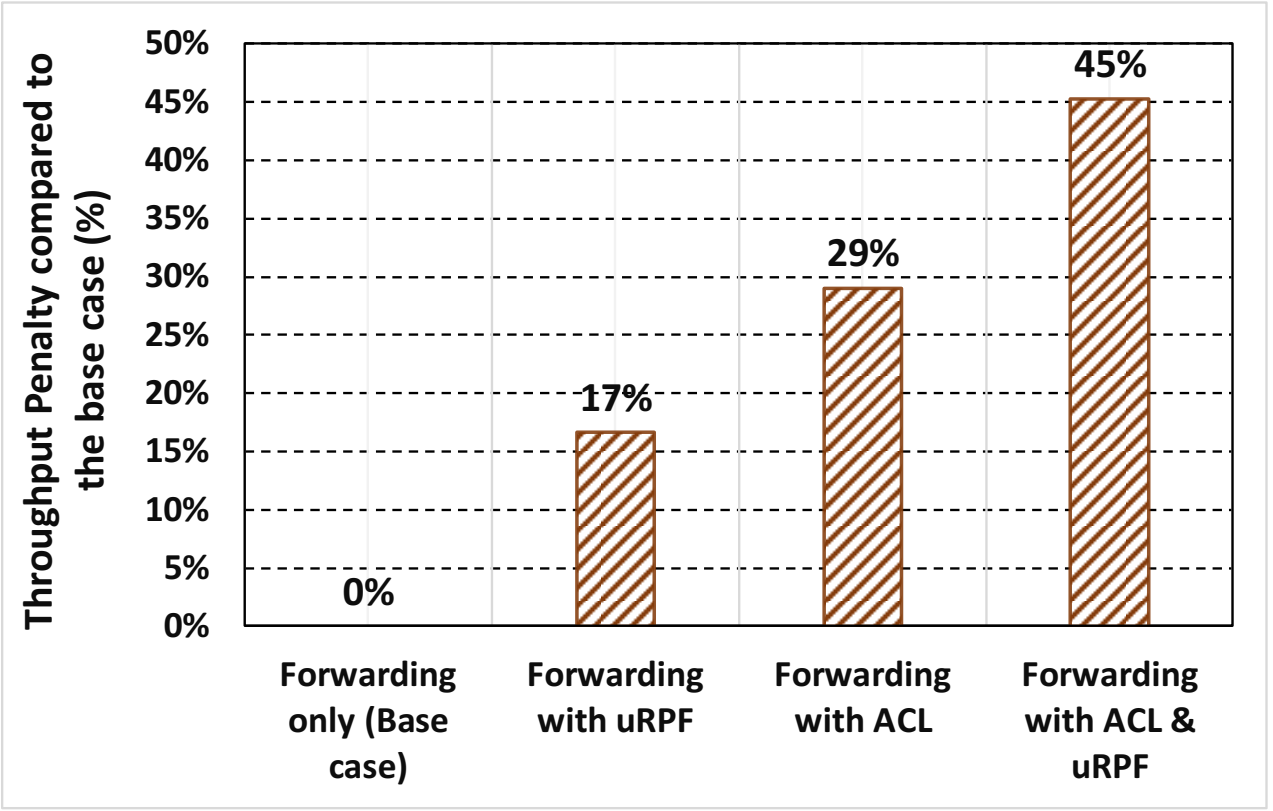

NIST plot using data from Router Analysis, Inc.
Ref: http://www.slideshare.net/RouterAnalysis/cisco-asr-1000-series-testing-results-and-analysis

Comparison of IPv4 & IPv6

Measurements by Router Analysis, Inc.

ASR1004 Router with 40 GHz ESP & Load Generated by Ixia box

64B packets (minimum size)

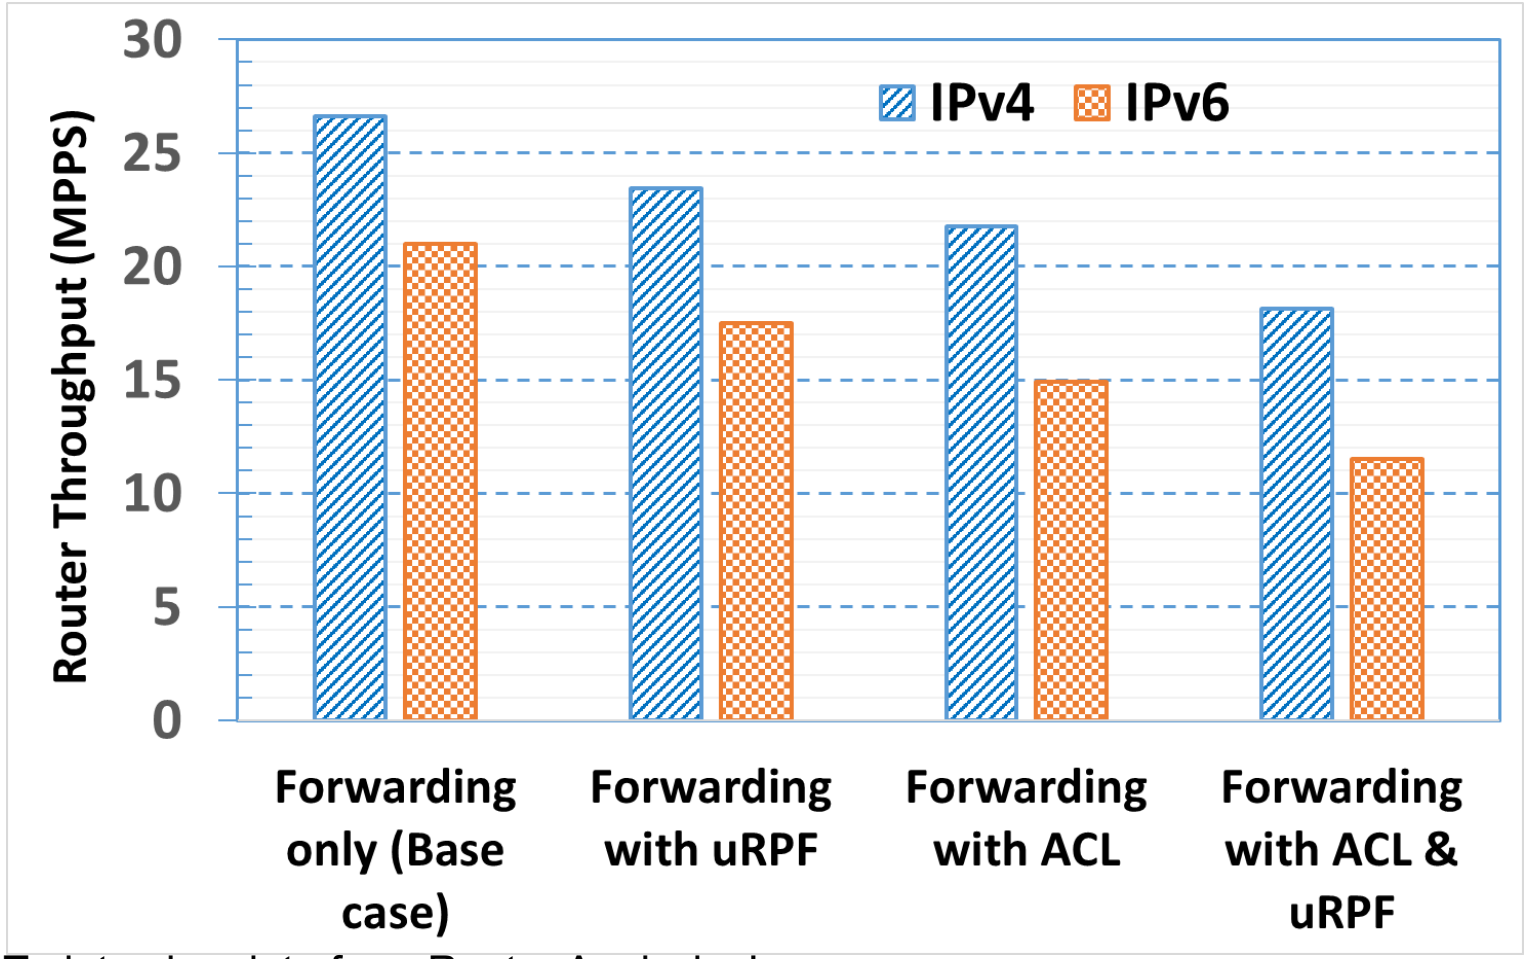

NIST plot using data from Router Analysis, Inc.

Ref: http://www.slideshare.net/RouterAnalysis/cisco-asr-1000-series-testing-results-and-analysis

Cisco Advertised Router Performance

Packet size = 64B

	Cisco ASR 1000 Series Routers		
ESP speed (GHz)	Forwarding only (MPPS)	Forwarding plus ACL, uRPF, QoS (MPPS)	Penalty due to ACL, uRPF, QoS (MPPS)
2.5	4	2	50.00%
5	7.5	4	46.67%
10	15	8	46.67%
20	19	6.7	64.74%
36	30	19	36.67%
40	23	10.4	54.78%
100	58	26	55.17%

ESP = Embedded Services Processor

- Working with assistance from Cisco to fill in missing performance results (Forwarding plus uRPF, Forwarding plus ACL).

Ref:
http://www.cisco.com/c/en/us/products/collateral/routers/asr-1000-series-aggregation-services-routers/datasheet-c78-731640.html

Benchmarking Next Steps

- Complete SAV Benchmarks Scripts
  - Pure Ixia test environment
- Test Range of Platforms
  - Cisco ASR1000
    - 5 Gbps throughput / 1G ports / 5Mb TCAM
    - 65K IPv4 routes.

  - Cisco ASR1001X
    - 20 Gbps throughput / 10 G ports / 80Mb TCAM
    - 1M IPv4 routes
  - Brocade MLXe-8
    - 3.2Tbps throughput / 100G ports
    - 1M IPv4 routes

  - Software / Virtual Routers
    - CloudRouter https://cloudrouter.org/

Benchmarking Impact

- **NIST Publication**
  - Methodologies for SAV Benchmarking.
  - Workload models.
  - Metrics for SAV Evaluation.
  - Open source scripts for Ixia driven SAV tests.
- **Performance Insight**
  - Factor into deployment guidance.

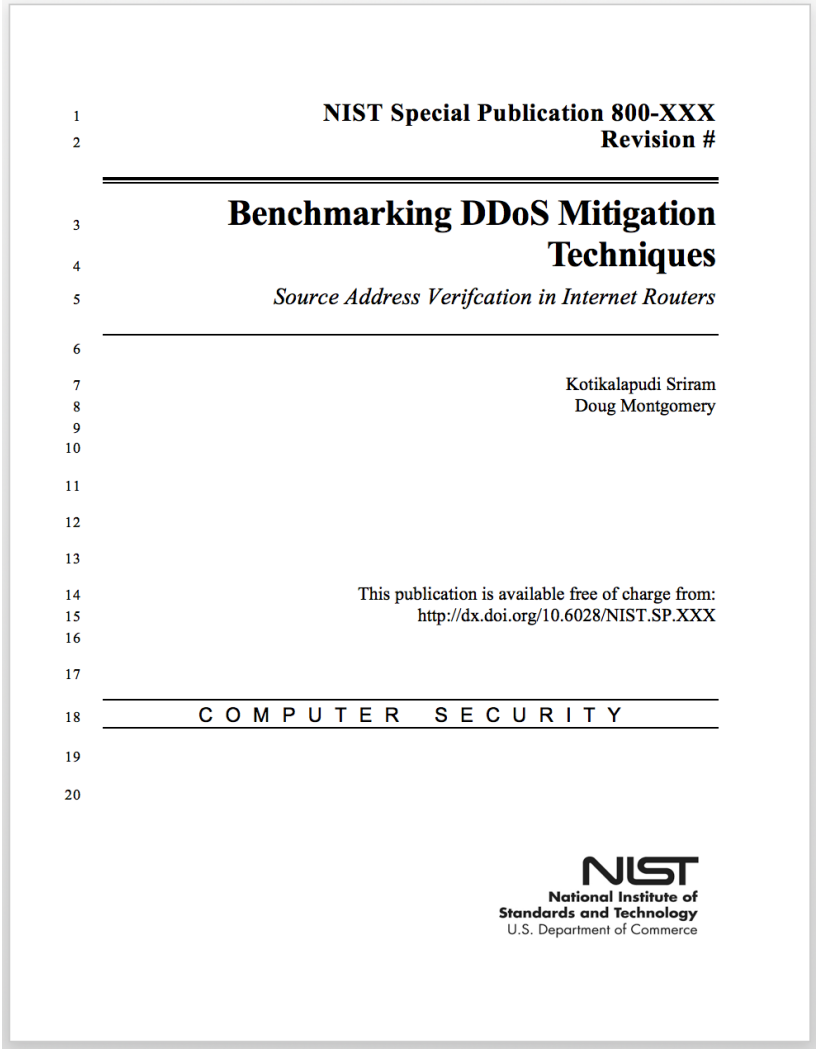

NIST Special Publication 800-XXX
Revision #

Benchmarking DDoS Mitigation Techniques

Source Address Verifcation in Internet Routers

Kotikalapudi Sriram
Doug Montgomery

This publication is available free of charge from:
http://dx.doi.org/10.6028/NIST.SP.XXX

COMPUTER SECURITY

NIST National Institute of Standards and Technology
U.S. Department of Commerce

Deployment Guidance

- **Need for detailed guidance**
  - Within a site.
  - Single homed sites.
  - Multi homed sites.
  - Intra sites / VPNs.
  - Could / outsourced services.
  - ISP customer, provider, peer, IXP connections.
- **No single answer**
  - Each scenario is different.
  - Each mechanism is different.
  - Ingress / egress is different.
  - Each platform is different.

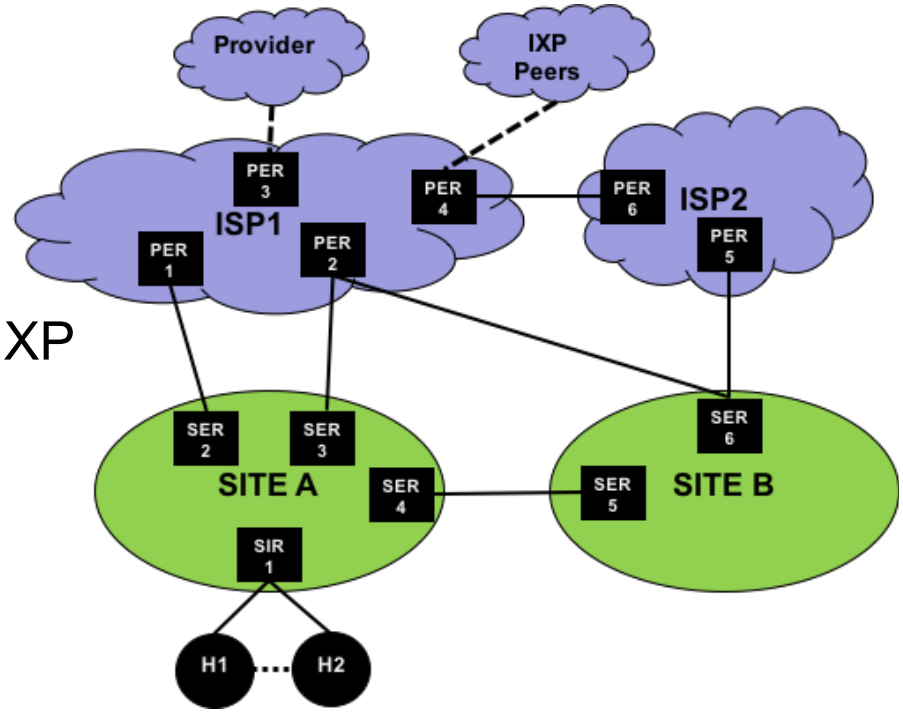

https://spoofer.caida.org/summary.php

Obstacles to SAV Deployment

Examples of Service Disruption due to uRPF:

- Various multi-homing scenarios: ISP may shy away from uRPF due to risk to multi-homed customer in case customer is not savvy about their announcements.
- PMTUD problem faced at IXP and its clients due to uRPF.

Concerns about Multi-homed Customer Service

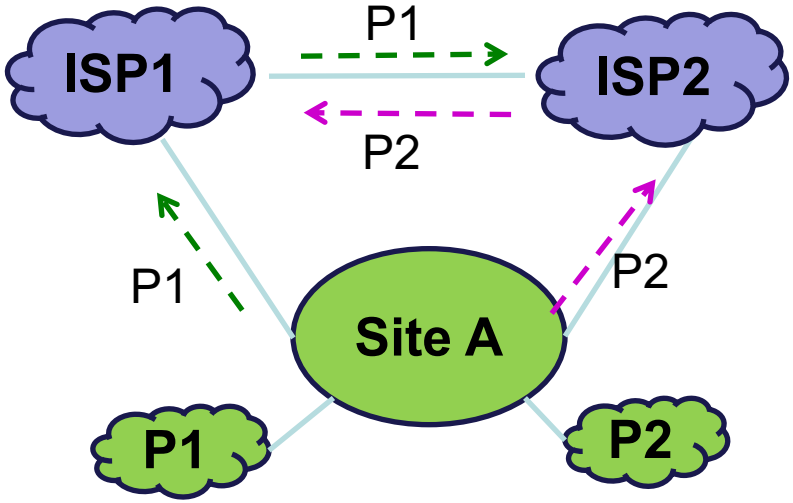

✗ Strict uRPF fails

✗ Feasible uRPF fails

✔ Loose uRPF works

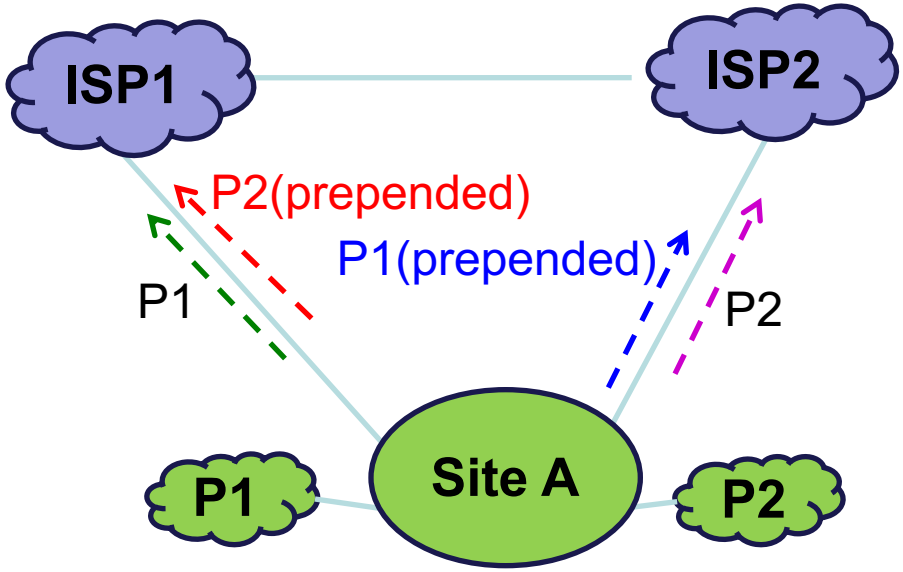

✔ Strict uRPF works

✔ Feasible uRPF works

✔ Loose uRPF works

Assume: Customer wants to engineer traffic for P1 via ISP1 and for P2 via ISP2. However, uses addresses from either P1 or P2 in SA fields towards either transit ISP.

- BCP recommendation: Multi-homed customer AS must announce all its routed prefixes or more specifics to each of its transit ISPs and should depref routes by prepending as necessary.

PMTUD Problem Faced by IXP Clients

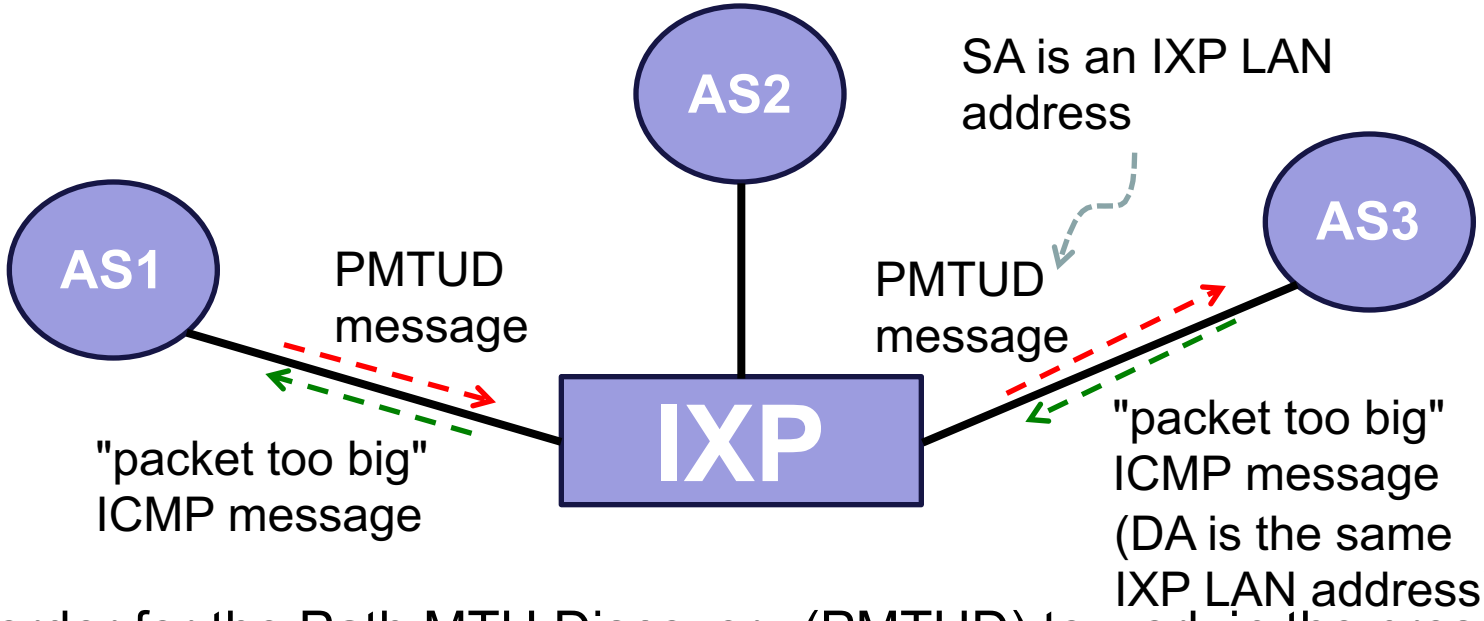

In order for the Path MTU Discovery (PMTUD) to work in the presence of uRPF:

- IXP should announce in BGP its own LAN prefix to all clients
- All AS clients at the IXP must install the IXP LAN prefix in their routing tables
- All AS clients must not accept the IXP LAN prefix or a more specific prefix from other peers

BGP Hijacking Can Subvert SAV

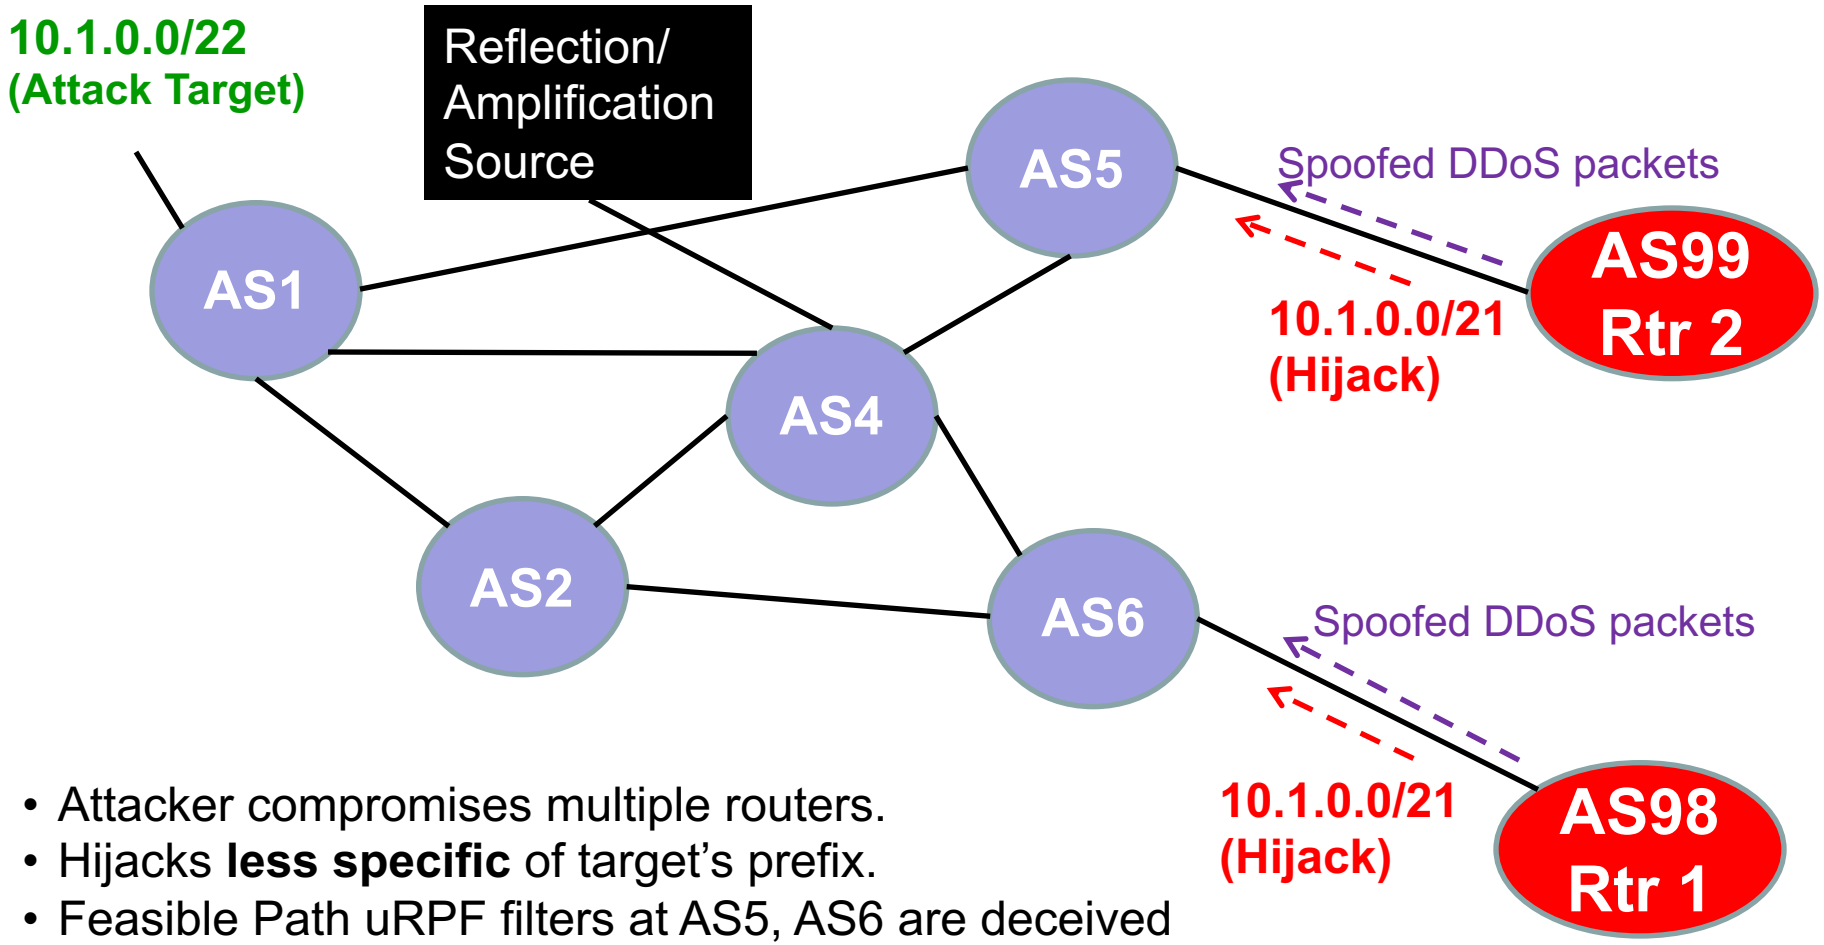

- Attacker compromises multiple routers.
- Hijacks **less specific** of target's prefix.
- Feasible Path uRPF filters at AS5, AS6 are deceived
- Attacker can launch DDoS reflection attack.

NIST Deployment Guidance

- **NIST Security Guidance**
  - Addressing both Control and Data plane issues.
  - Recommendations for BGP prefix filtering.
  - Recommendations for IP SAV filtering.
  - Addressing multiple deployment scenarios
- **Coordination with**
  - USG TIC architectures
  - FedRAMP guidance
  - FISMA requirements

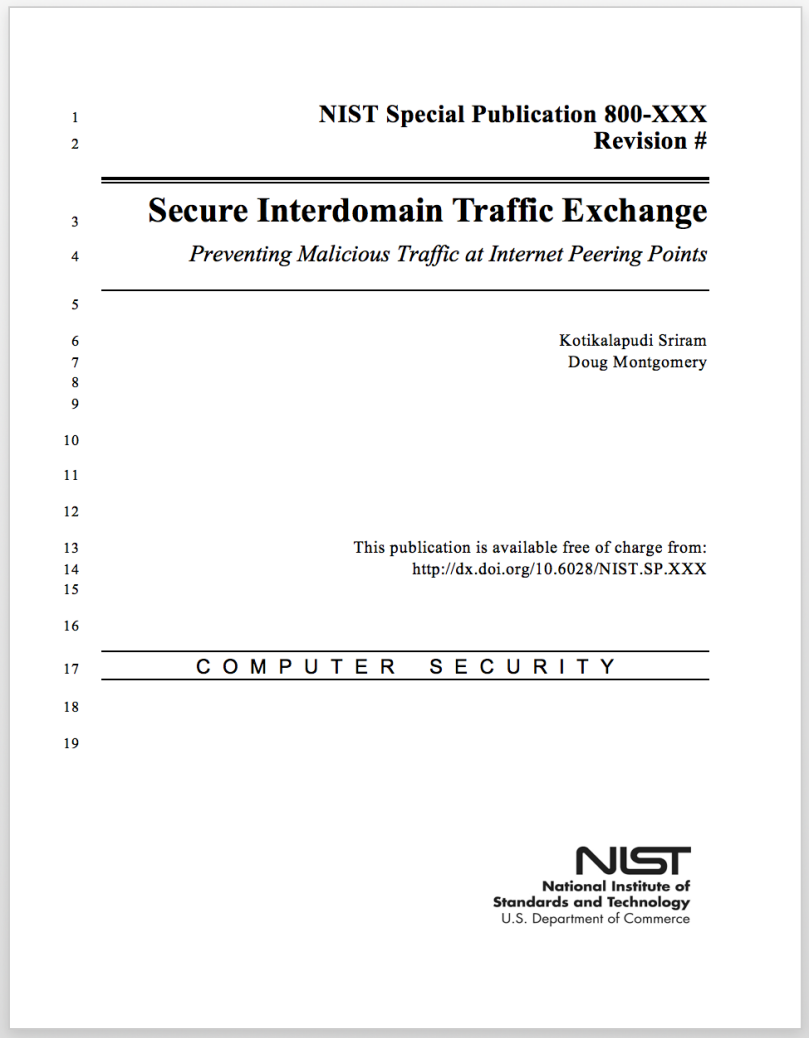

NIST Special Publication 800-XXX
Revision #

Secure Interdomain Traffic Exchange

Preventing Malicious Traffic at Internet Peering Points

Kotikalapudi Sriram
Doug Montgomery

This publication is available free of charge from:
http://dx.doi.org/10.6028/NIST.SP.XXX

COMPUTER SECURITY

NIST National Institute of Standards and Technology
U.S. Department of Commerce

References

- M. Luckie, B. Huffaker, A.Dhamdhere, V. Giotsas, and kc claffy, “AS Relationships, Customer Cones, and Validation.”
  - http://www.caida.org/publications/papers/2013/asrank/asrank.pdf

- **Cisco router performance measurements:**
  - http://www.cisco.com/c/en/us/products/collateral/routers/asr-1000-series-aggregation-services-routers/datasheet-c78-731640.html

- **Router Analysis, Inc. router performance measurements:**
  - http://www.slideshare.net/RouterAnalysis/cisco-asr-1000-series-testing-results-and-analysis

Questions?

{dougm | ksriram}@nist.gov

Acknowledgements / Disclaimers

- This research was supported by the Department of Homeland Security under the DDoSD program and the NIST Information Technology Laboratory under the Internet Infrastructure Protection program.
- Certain commercial equipment, instruments, or materials are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.
- The views and conclusions contained in this document are those of the authors, and should not be interpreted as representing the official policies, either expressed or implied of the NIST, or the U.S. Government.

Backup Slides

SAV Workload Models

Computing "Customer Cones" for SAV

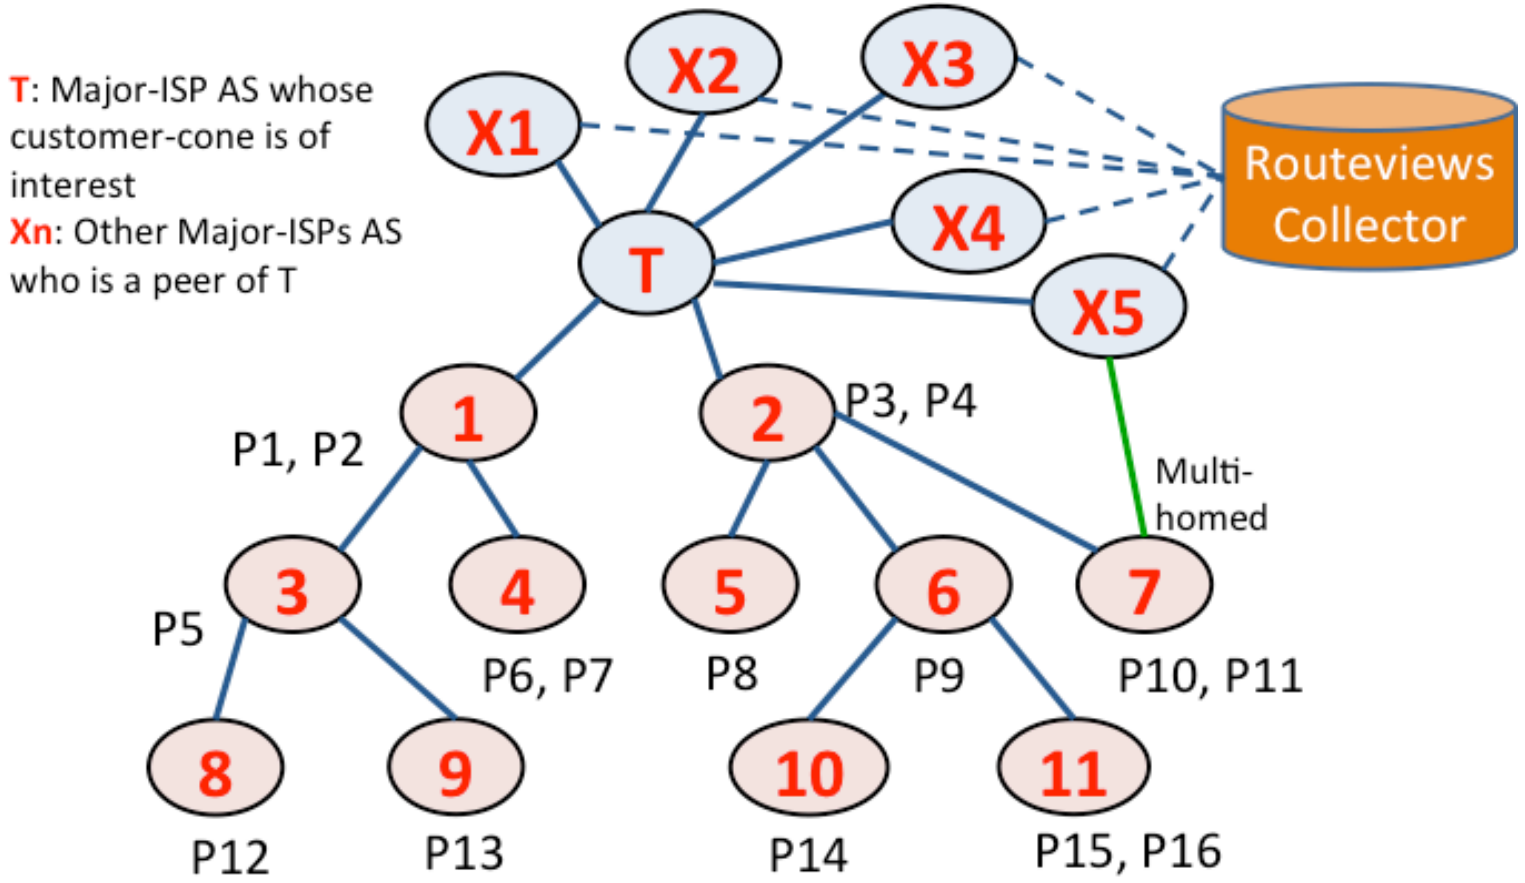

- Gather all the RIB entries of the kind: P { Xn, T, AS-i, AS-j, AS-k }; n = 1,2,3,4,5
- Unlikely that a prefix or an AS in the customer cone of T would be multi-homed to all five of the peers of T (i.e., X1, X2, X3, X4, X5)

Prefixes Observed at Level3 (AS3356)

Prefixes observed at Level3 (AS3356) -- as viewed by Routeviews	
Total # prefixes in Customer Cone of AS3356	287,841
Total # prefixes at AS3356 but learned from Peers/Providers	282,242

Additionally (for uRPF), # unannounced /8's in the Internet	57

- The number of customer prefixes per customer facing PE router will vary, and they all add up to 287,841 (for AS3356)

Characteristics of Customer Cone of Level3 (AS3356)

Prefixes

AS path length (depth)	prefix count
1	2612
2	67569
3	135516
4	65642
5	18609
6	4156
7	575
8	47
9	3
10	2
Total	294731
Unique	287841

ASes

AS path length (depth)	#ASes
1	1
2	4008
3	18069
4	10797
5	3069
6	1057
7	144
8	17
9	3
10	2
Total	37167
Unique	33415

Walking the Tiers

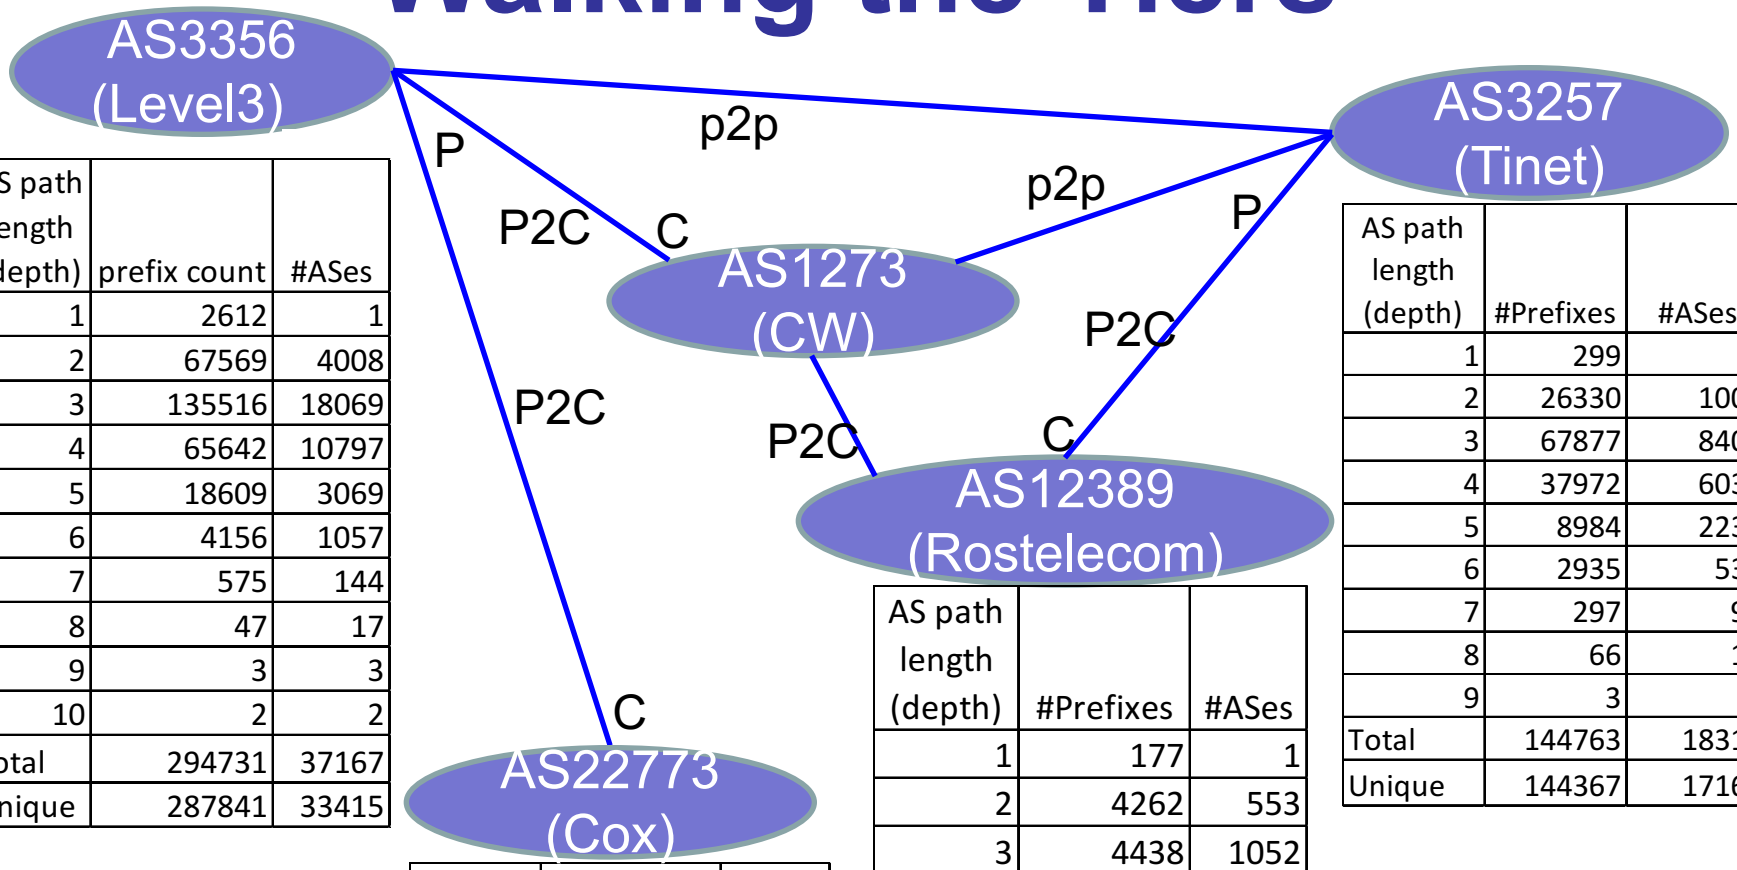

AS path length (depth)	prefix count	#ASes
1	2612	1
2	67569	4008
3	135516	18069
4	65642	10797
5	18609	3069
6	4156	1057
7	575	144
8	47	17
9	3	3
10	2	2
Total	294731	37167
Unique	287841	33415

AS path length (depth)	#Prefixes	#ASes
1	299	1
2	26330	1004
3	67877	8405
4	37972	6031
5	8984	2231
6	2935	538
7	297	93
8	66	14
9	3	2
Total	144763	18319
Unique	144367	17164

AS path length (depth)	#Prefixes	#ASes
1	177	1
2	4262	553
3	4438	1052
4	1734	510
5	1063	201
6	129	35
7	60	8
8	1	1
Total	11864	2361
Unique	11855	

AS path length (depth)	prefix count	#ASes
1	2931	1
2	1330	412
3	120	46
4	1	1
Total	4382	460
Unique	4382	

P: provider
C: Customer
p2p: peer to peer

Only 65 prefixes are seen at AS3356 as originated by 12389 ! Why not 177?

Range of #Prefixes in Customer Cone

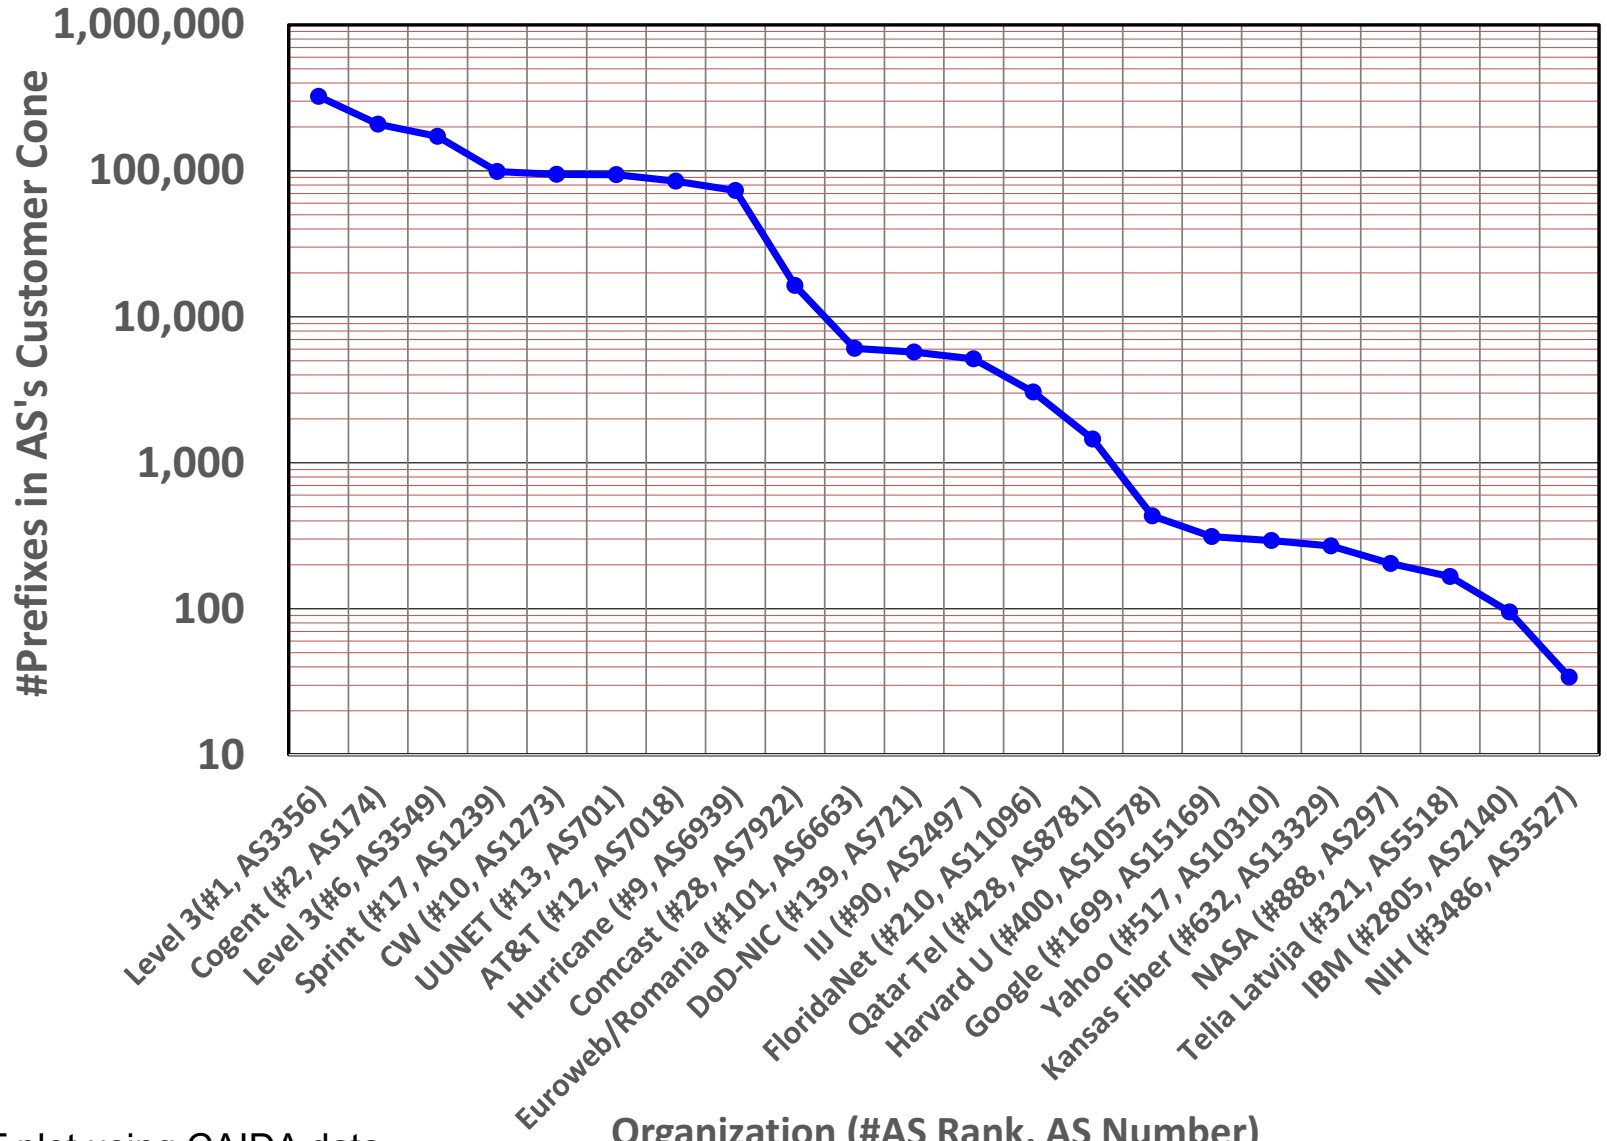

NIST plot using CAIDA data

Range of #ASes in Customer Cone

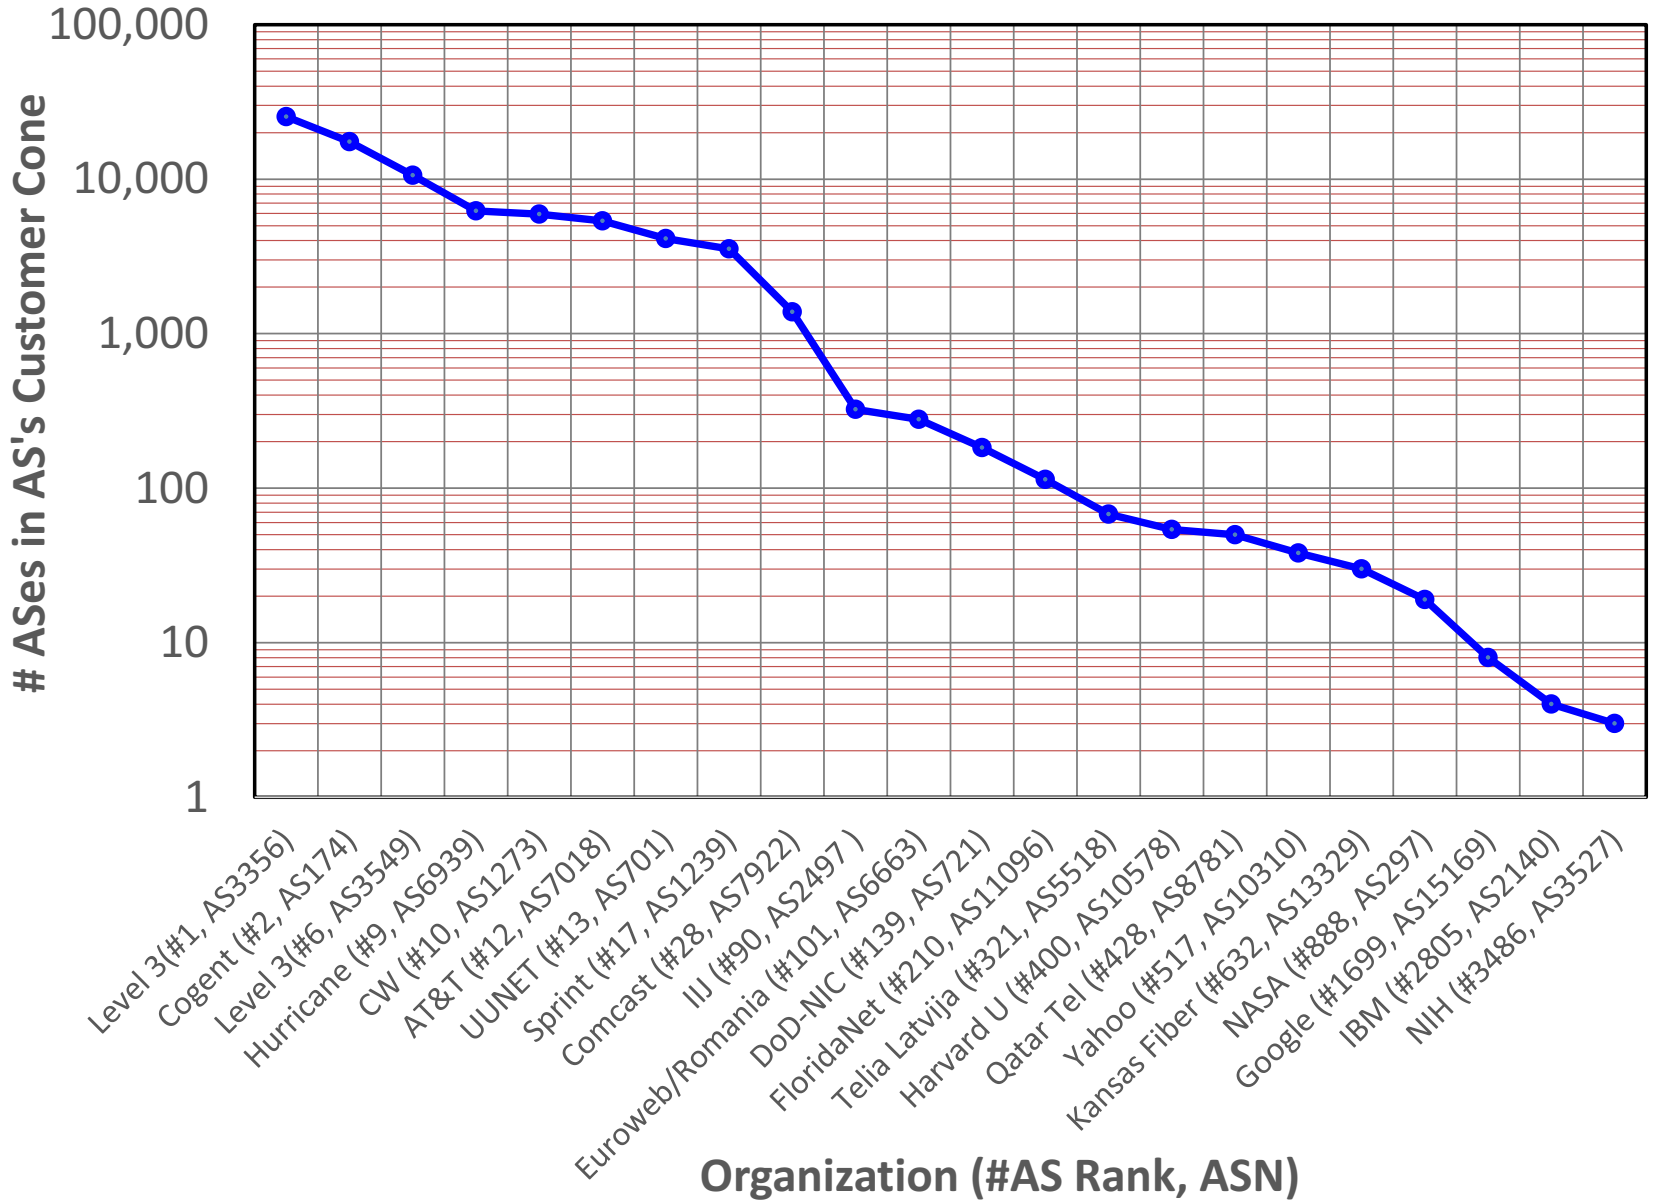

Distribution of #Prefixes Originated by an AS

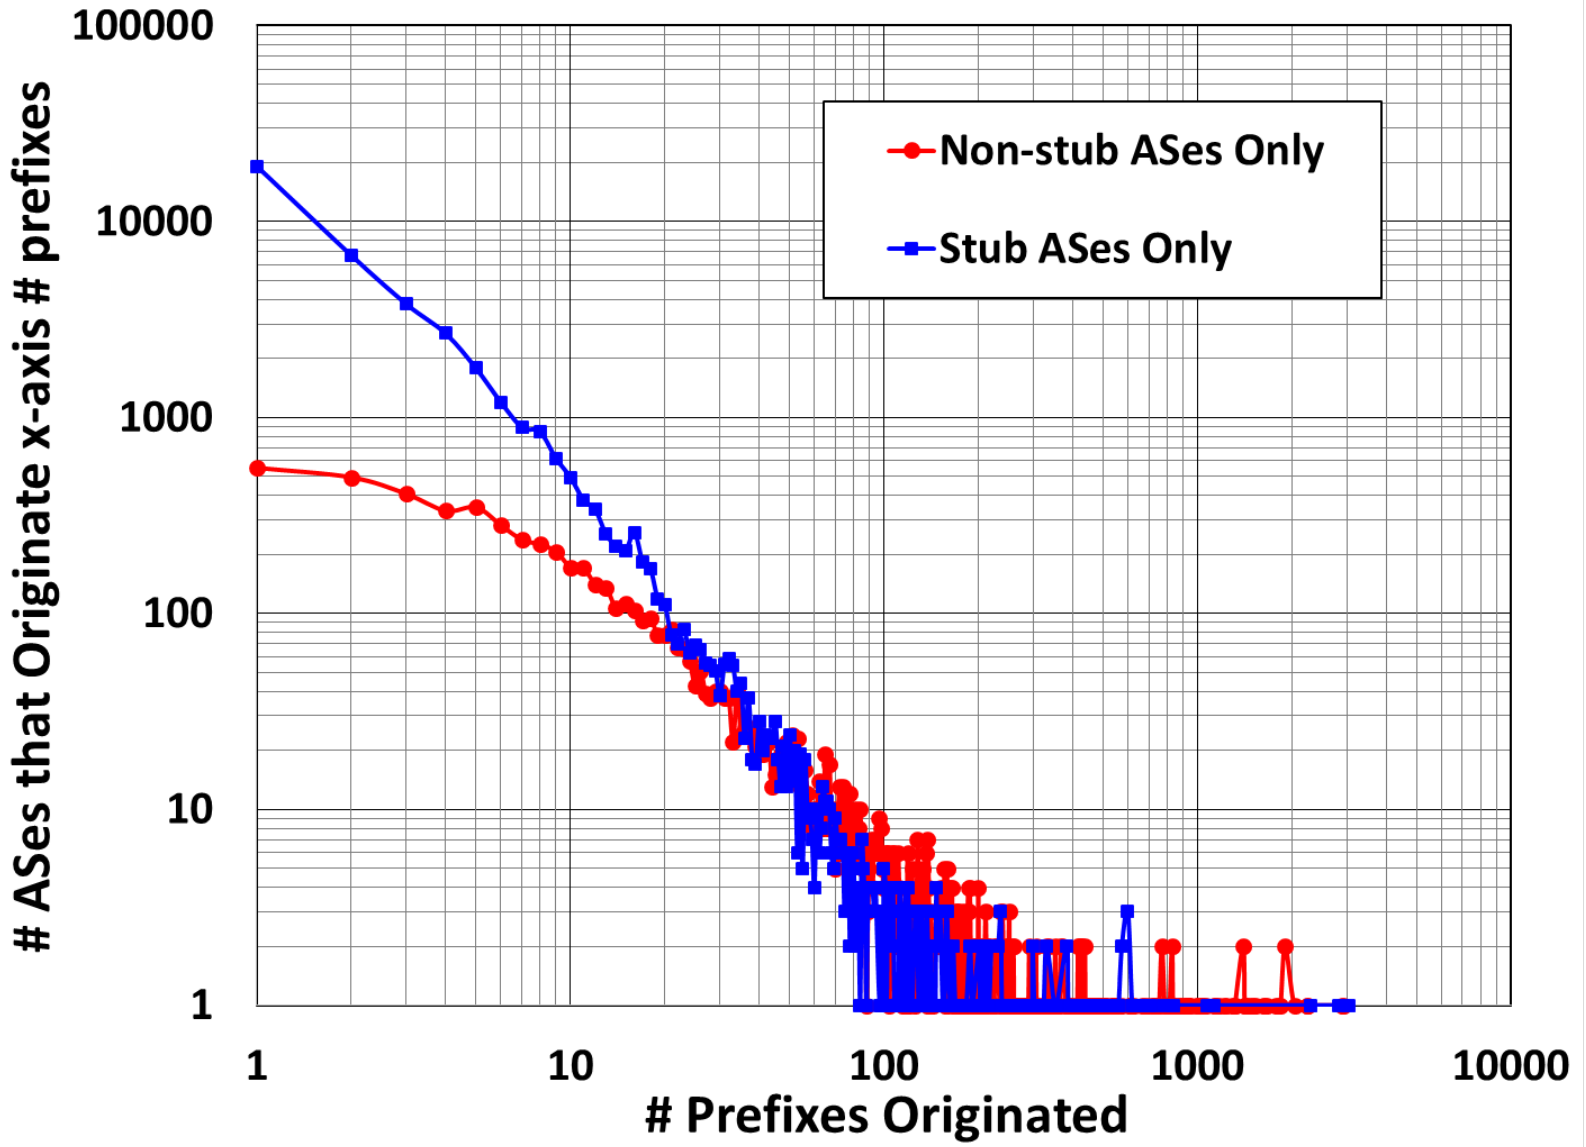

IP Spoofing – Defeat Trace-Back

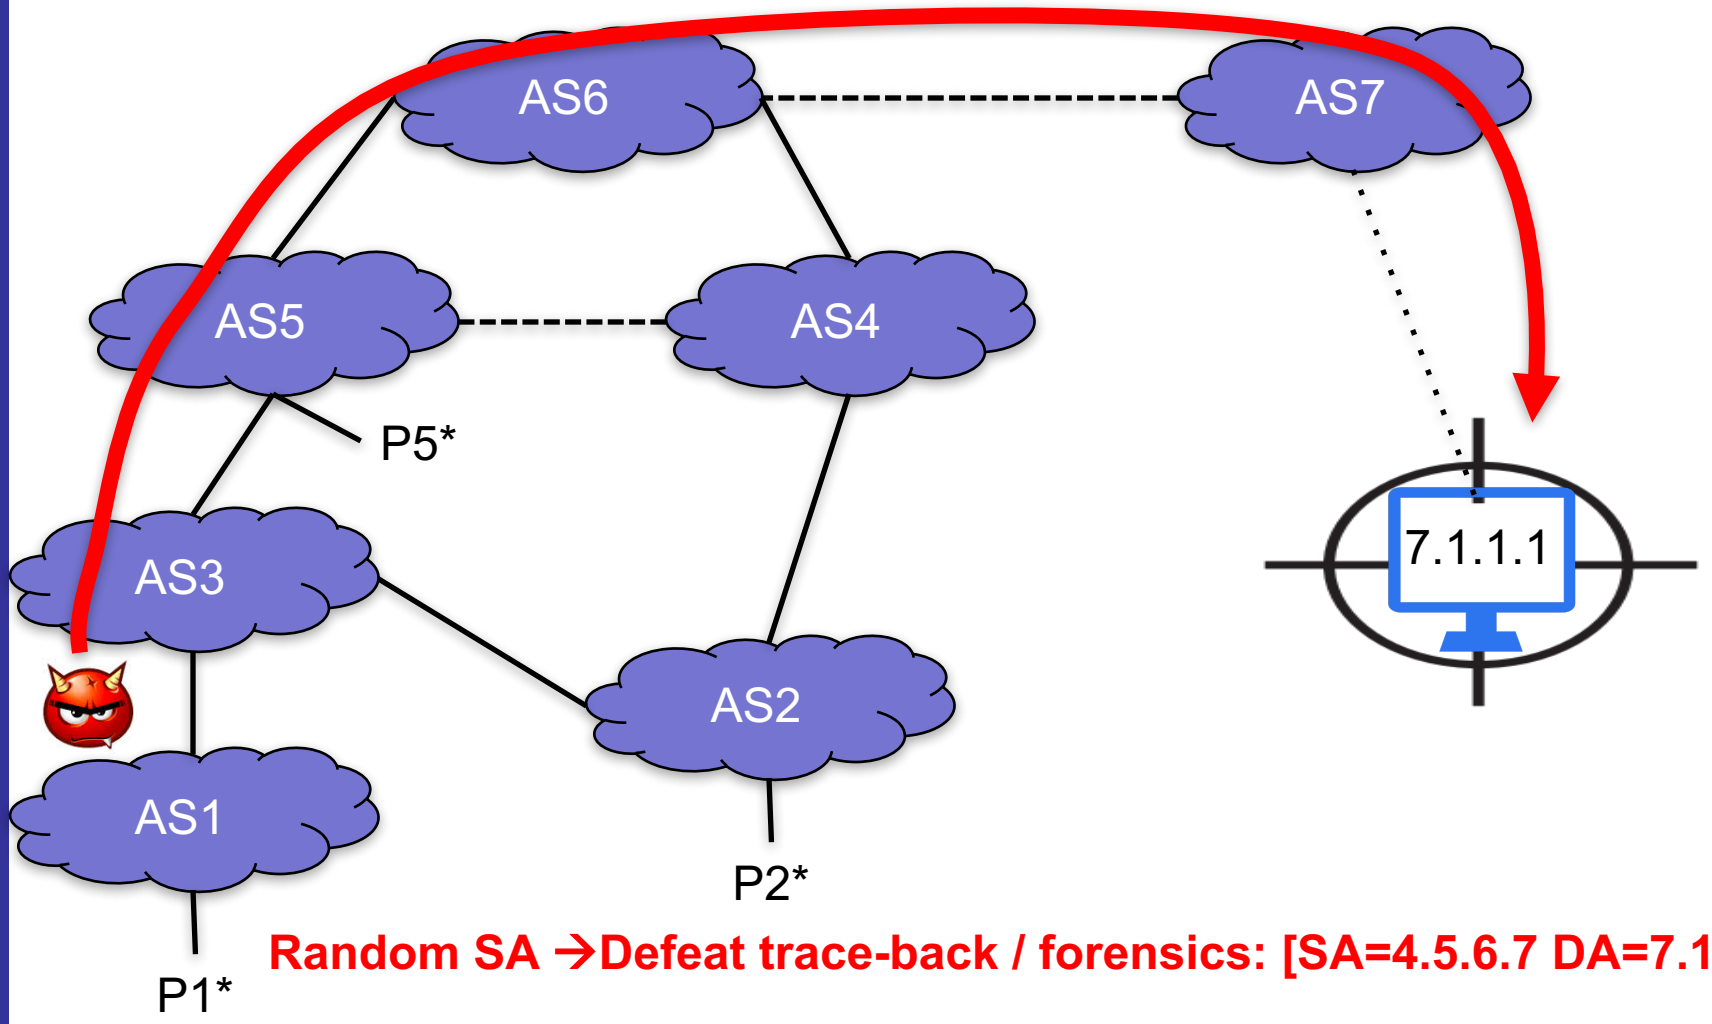

Random SA →Defeat trace-back / forensics: [SA=4.5.6.7 DA=7.1.1.1]

IP Spoofing – Reflection Attack

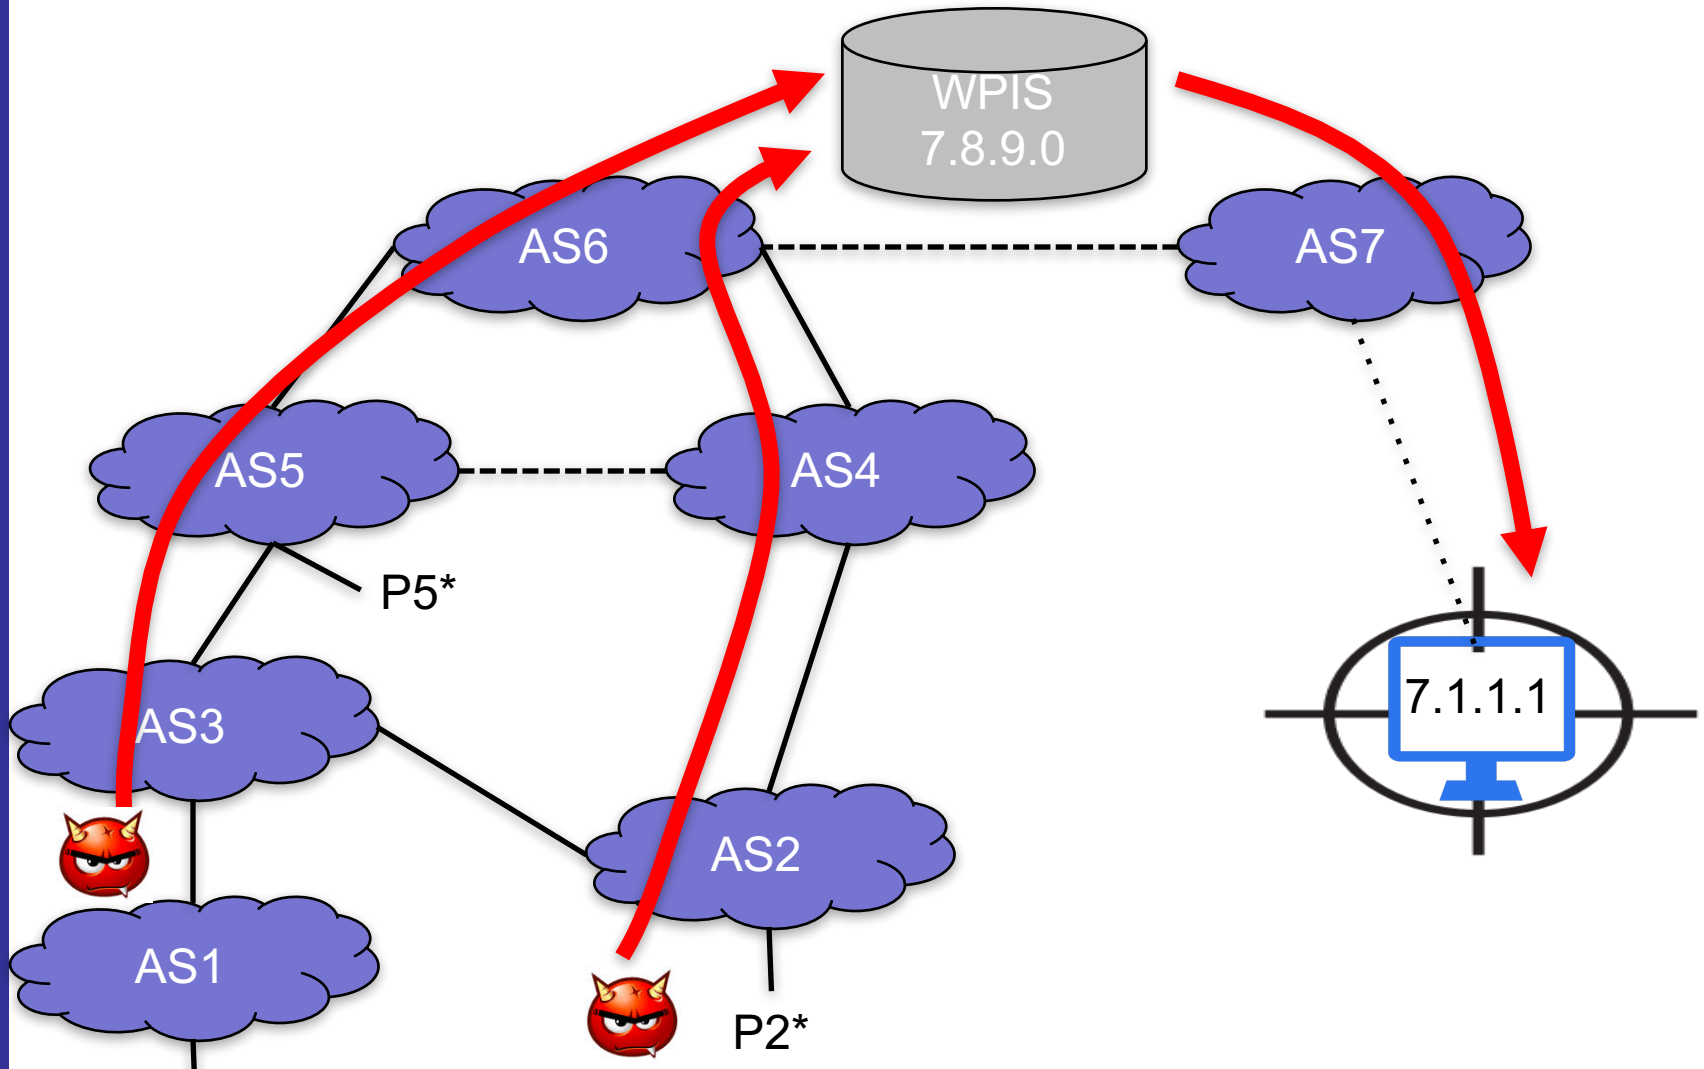

Target = SA →Reflection Attack: Query [SA=7.1.1.1 DA=7.8.9.0]

Source Address Verification (SAV)

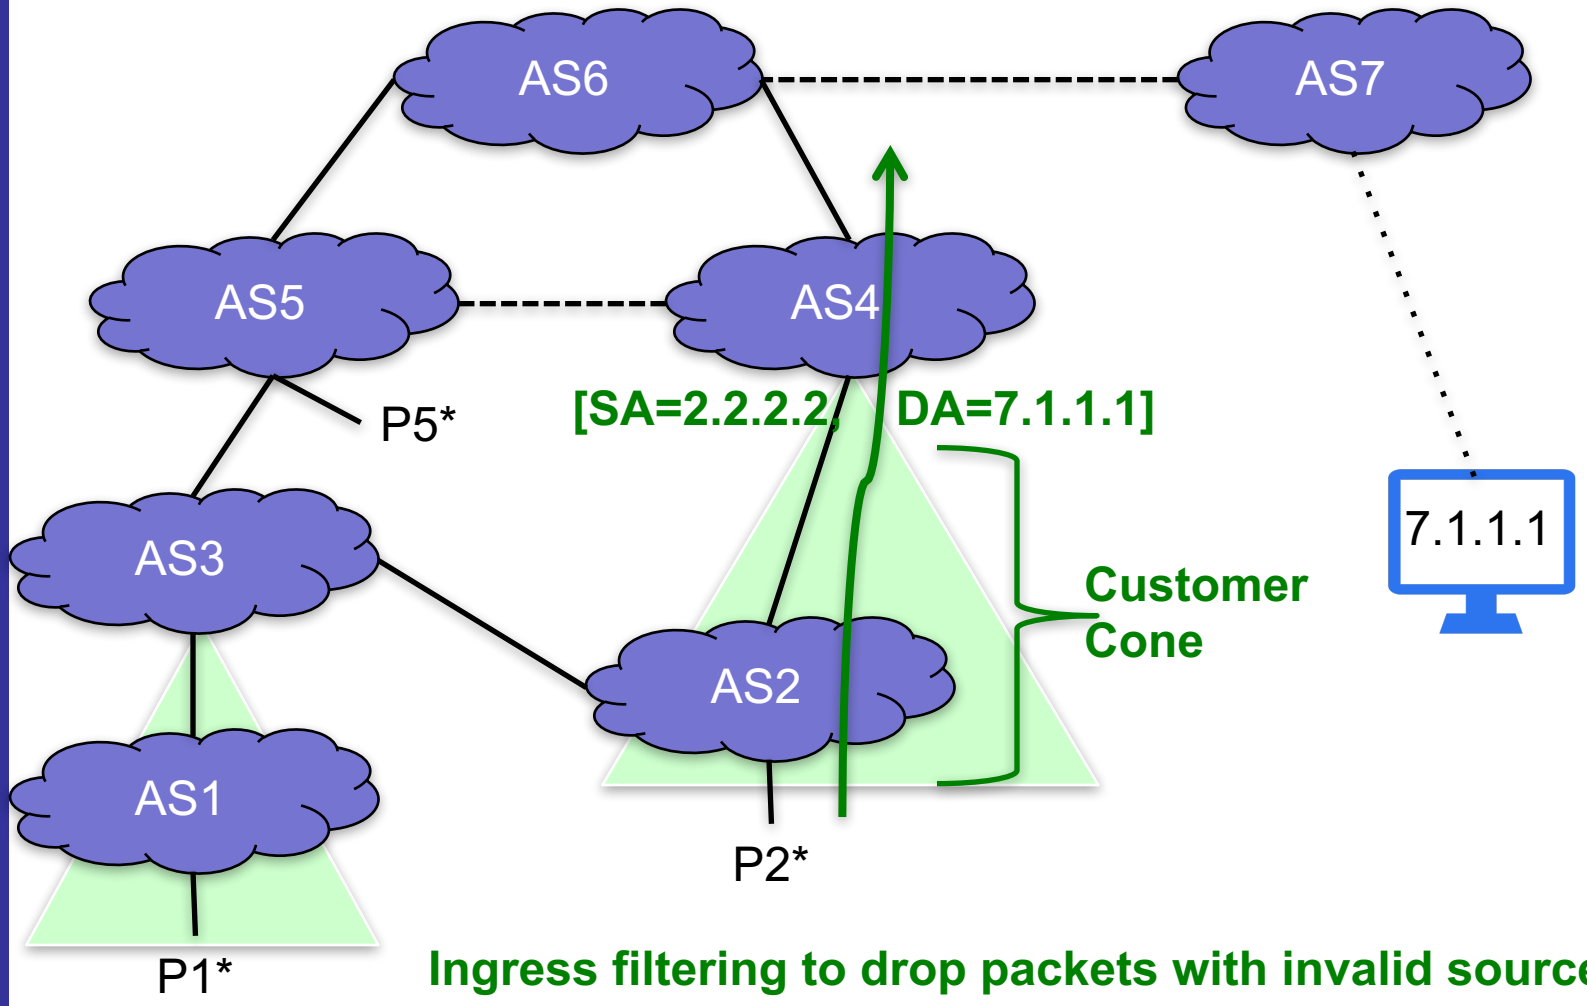

Ingress filtering to drop packets with invalid source address

Source Address Verification (SAV)

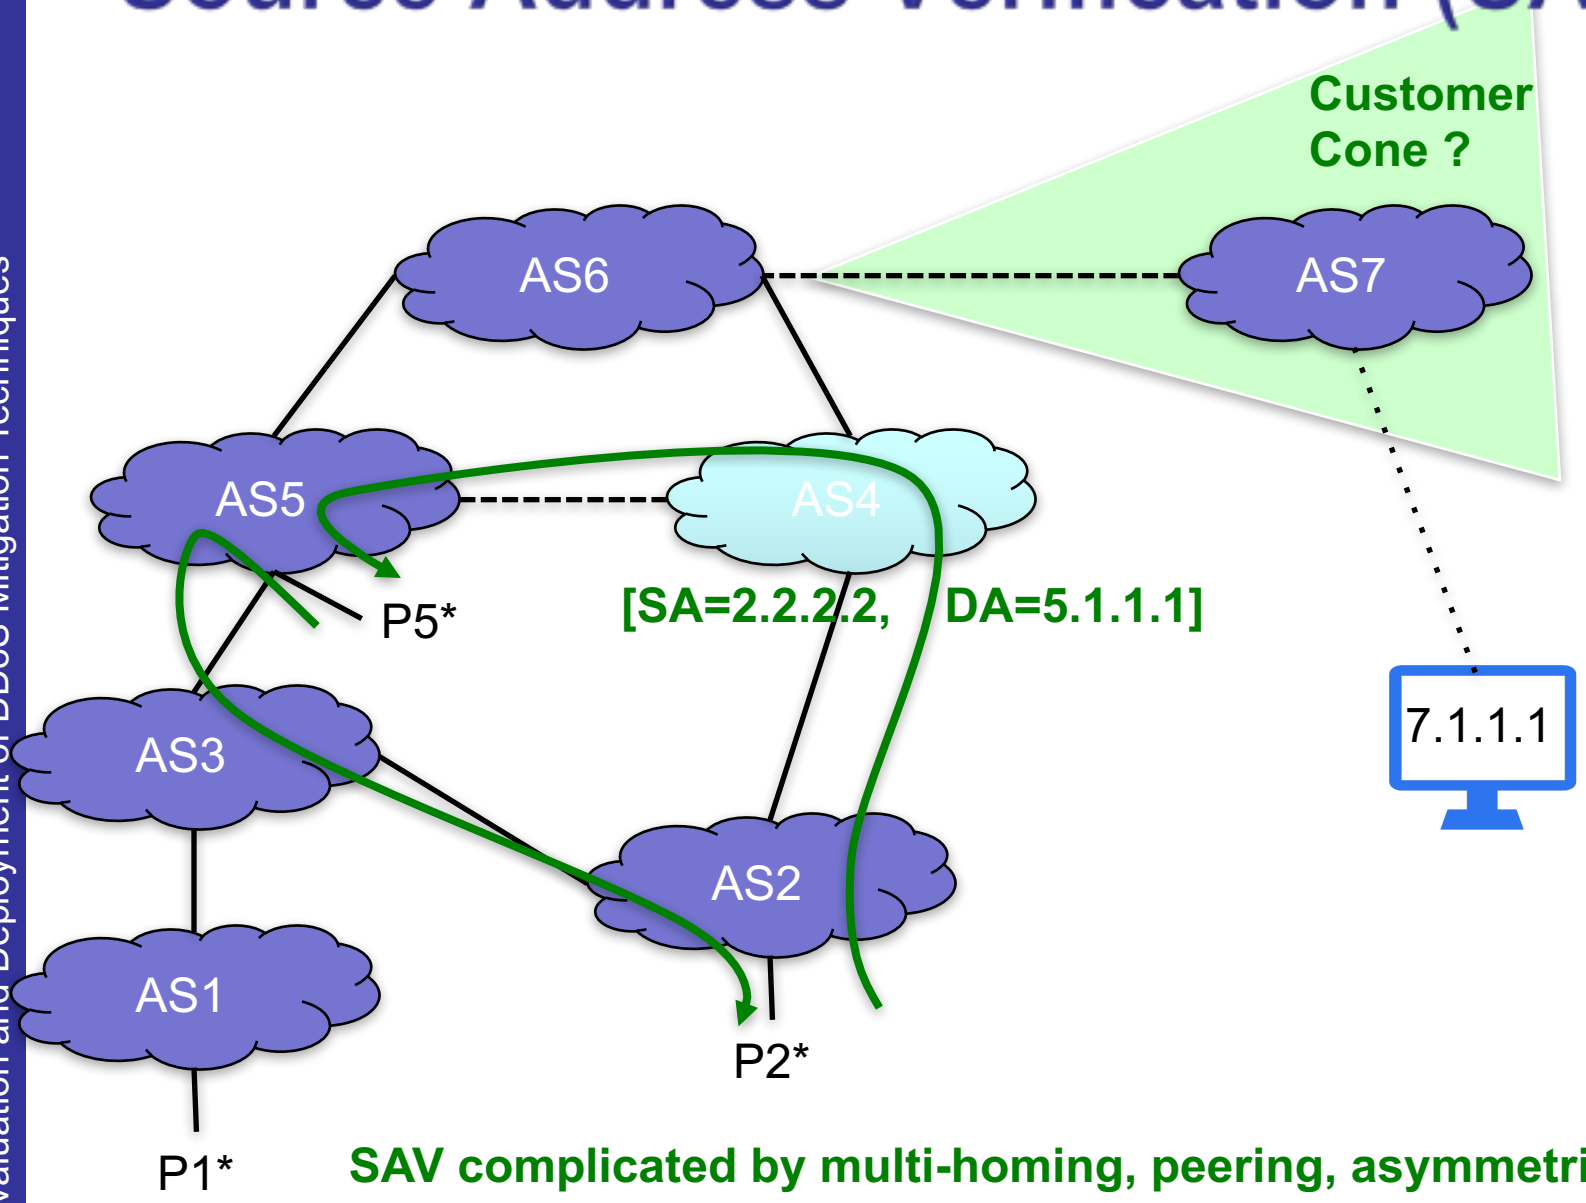

SAV complicated by multi-homing, peering, asymmetric routes.