Accountable Internet Protocol

David Andersen, Hari Balakrishnan, Nick Feamster,
Teemu Koponen, Daekyeong Moon, Scott Shenker

http://www.aip-arch.net/

Drawbacks (a sampler)

- **Complicated Mechanisms**
  - Many details to circumvent IP weaknesses
- **External Sources of Trust**
  - Trusted certificate authorities (e.g., SBGP)
- **Operator Vigilance**
  - Semi-manual configuration (e.g., filters, registries)

IP Layer Names Don't Have Secure Bindings

- Three kinds of IP layer names:
  IP address, IP prefix, AS number
- No secure binding of host to its IP addresses
- No secure binding of AS number to its IP prefixes

Accountability

- Many problems easier to solve with *network-layer accountability:*

  Ability to associate a principal with a message

- There's a way to make accountability intrinsic

How?

- Key idea: New addressing scheme for networks and hosts
- Addresses are self-certifying
- Simple protocols that use properties of addressing scheme as foundation
  - Anti-spoofing, secure routing, DDoS shut-off, etc.

AIP Addressing

Autonomous domains, each with unique ID

An AD...
Would fail together
Single administrative domain

Key Idea:

AD and EID are *self-certifying flat names*

- AD = hash(public_key_of_AD)
- Self-certification binds name to named entity

a glo

AD1:EID,AD2:EID,AD3:EID

AIP Forwarding and Routing

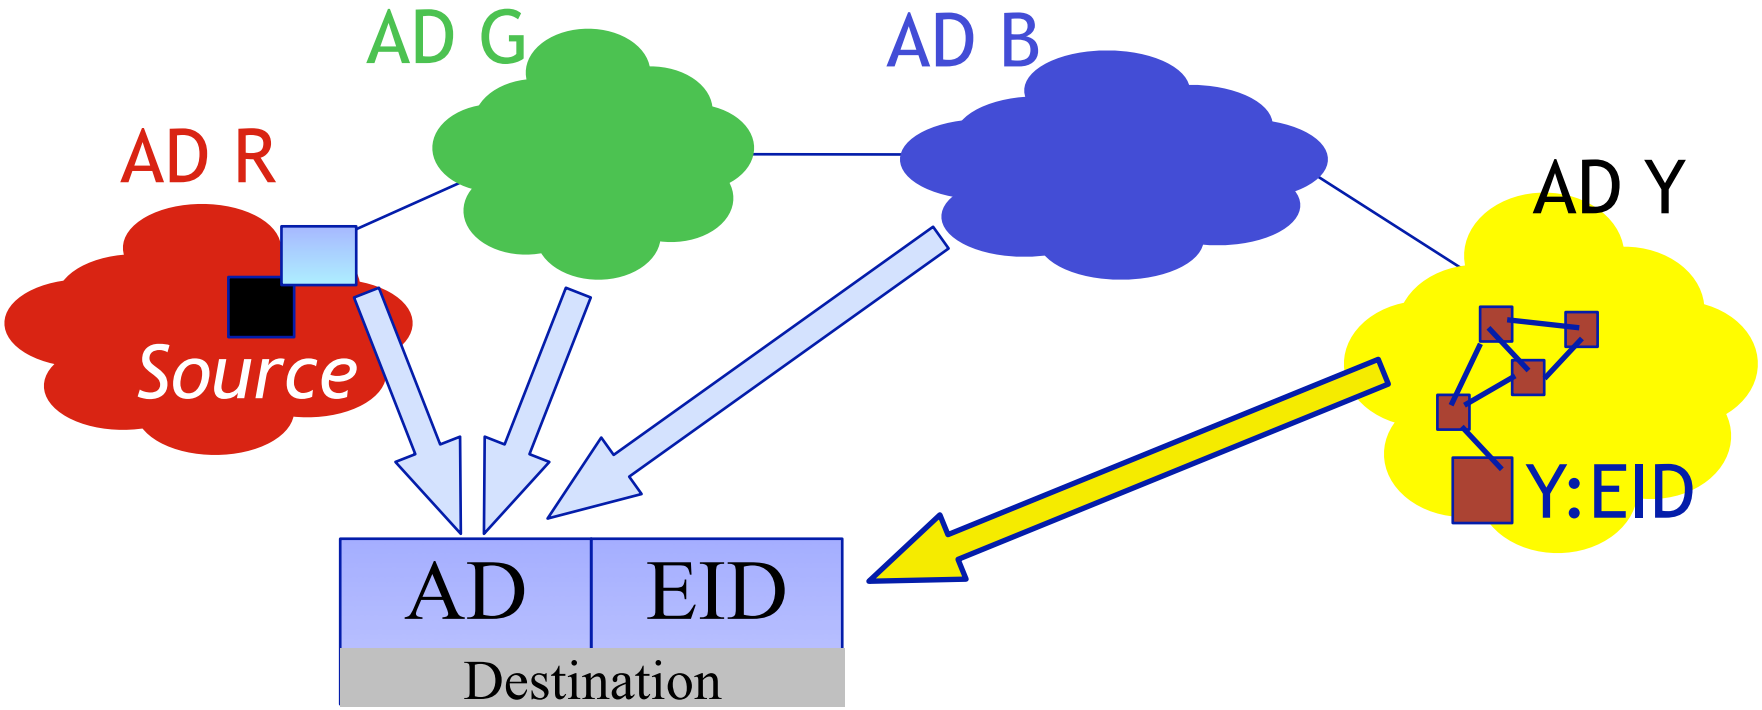

Inter-AD routing & forwarding: AD #s only.

Intra-AD routing disseminates EIDs.

Many routing protocols possible - derive security from AIP self-certification

Roadmap

- Uses
  - **Secure Routing**
  - **Anti-Spoofing**
  - **Shut-Off Packets**
- Concerns
  - **Scalability**
  - Key Management
  - Traffic Engineering

Secure Routing with AIP (for BGP)

- ***Origin authentication***:
  prefix originated by AS X actually belongs to X
- ***Path authentication***: accuracy of AS path

– S-BGP requires external infrastructures

Routing Registry	
Prefix	Pub Key

AS PKI	
AS	Pub Key

  – In past, registries notoriously inaccurate

✓ With AIP: ADs exchange pub keys via BGP messages

  ✓ Origin auth automatic: ADs *are* keys!

  ✓ Path auth: Just like S-BGP, but no PKI

Detecting & Preventing Spoofing

- Self-certified entity can *prove* it sent message:

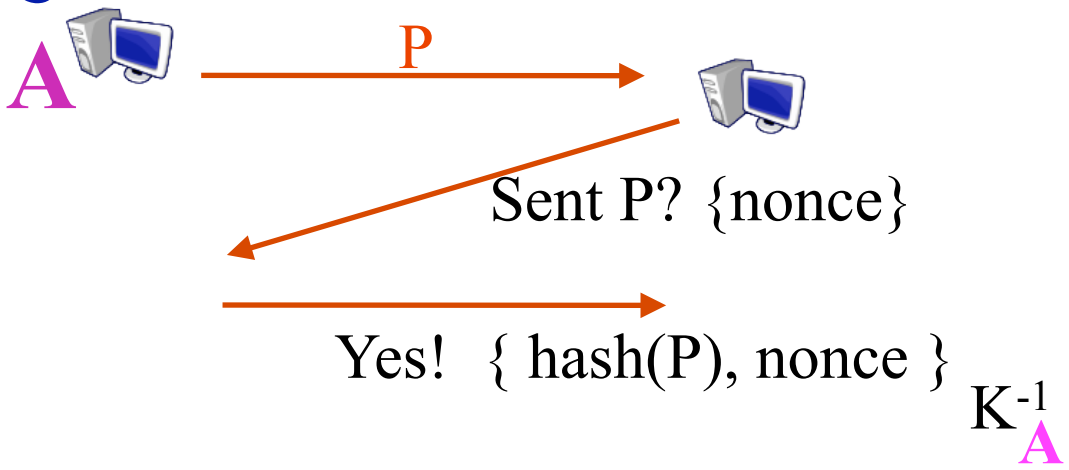

- Routers or hosts seeing packet can check the AD or EID using a challenge-response protocol

Spoofing vs. Minting

- AIP guarantee:
  - Nobody but X can claim to be X
- However:
  - X could invent a new identity (minting)

Mitigating Minting

- Peering ADs:
  - Today: List which ASes/Prefixes A can use (painful for clients and ISPs)
  - AIP: Configure reasonable limit on number of ADs can announce
- Edge ADs can limit EIDs similarly

AIP Enables Secure Shut-Off

- Problem: Compromised zombie sending stream of unwanted traffic to victim
- Zombie is "well-intentioned", owner benign [Shaw]

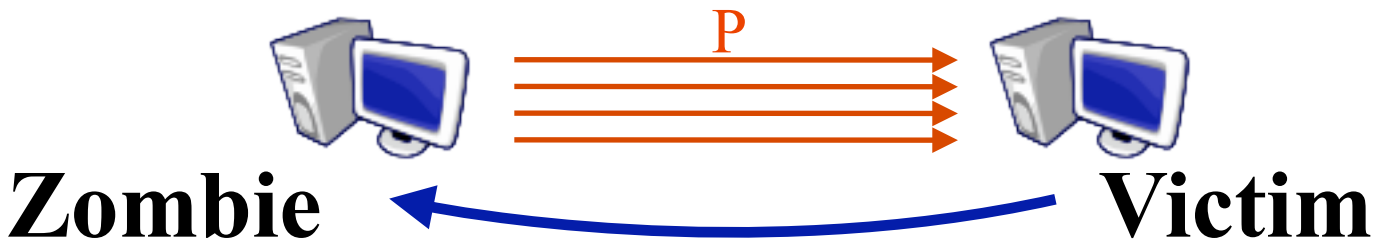

Shut-off packet

$\{ \text{key} = K_{victim}, \text{TTL}, \text{hash}=H(P) \}_{K^{-1}_{victim}}$

- Shut-off scheme implemented in NIC
  (NIC firmware update requires physical access)
- Hardware requirements practical
  - Bloom filter for replay prevention (8MB SRAM)

Can AIP Scale?

- How big will the routing tables be?
  - # of entries: Scale from IP
    (ASes vs. prefixes vs. ADs)
  - Diameter: Shrinking in IP
    AIP: more ADs on path
  - Size of entries: Larger AIP addresses

- How much work to process updates?
  - Crypto overhead

BGP Table Size Trends

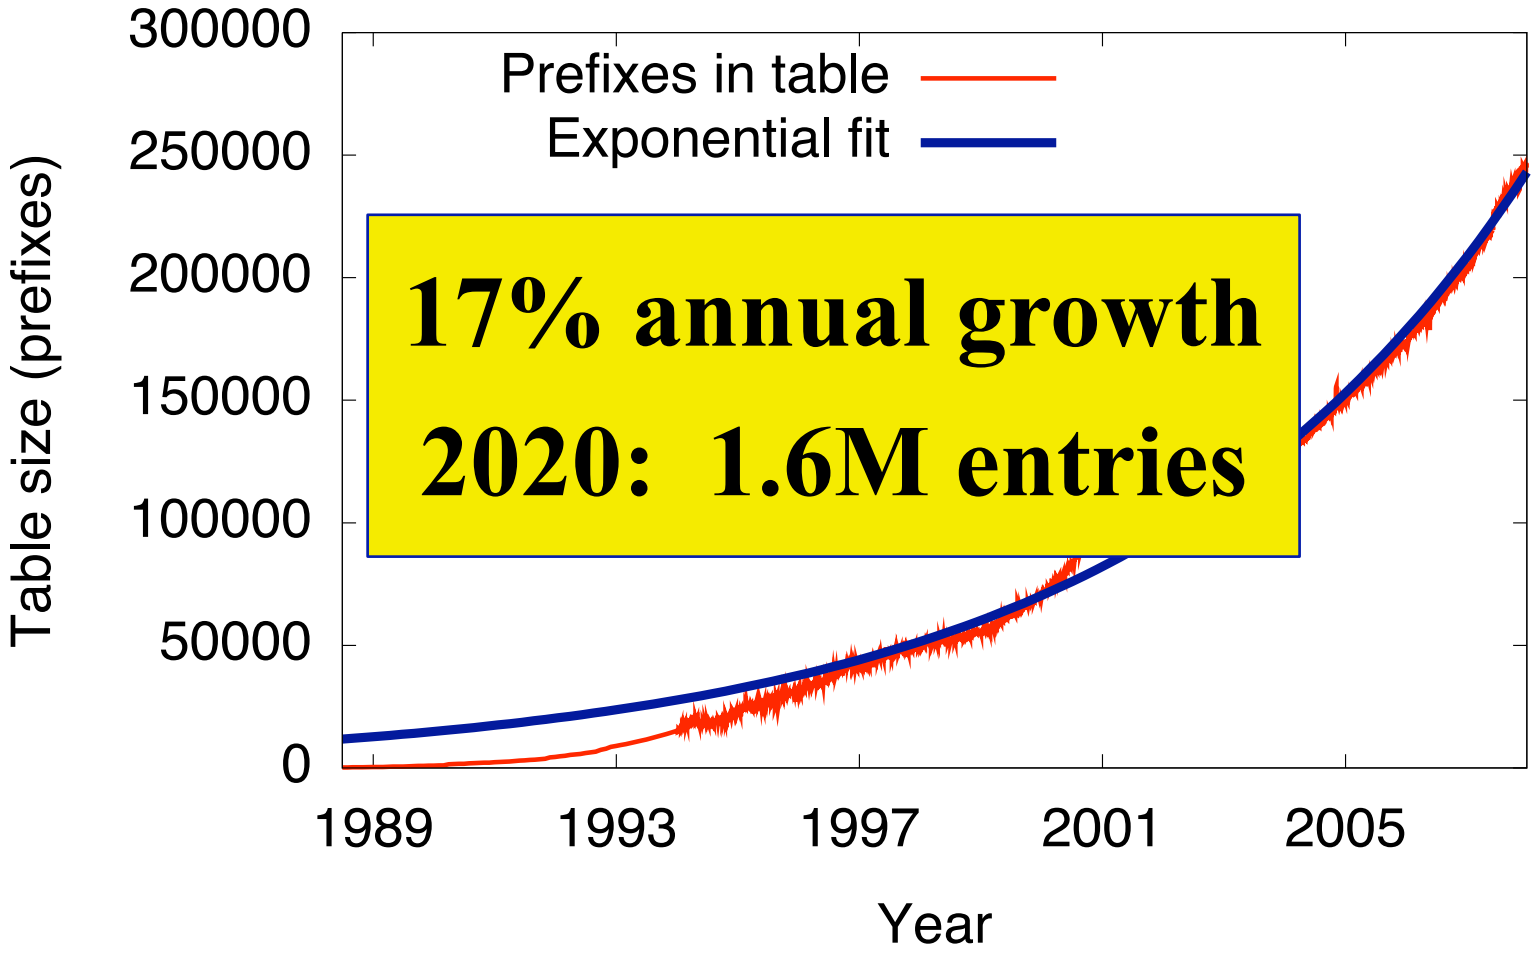

Growth vs. Hardware

- Semiconductor industry roadmap projects doubling in ~3 years
  - 50% >> 17%. But let's look at some #s...

- In 2020, can we build a cost-effective router for AIP traffic?

RIB Memory (20 full-table peers, core)

Gigabytes (2007 Dollars)

	2007	2011	2020
IP	0.4 ($30)	0.7 ($14)	2.9 ($7)
AIP	1.3 ($103)	2.0 ($40)	8.2 ($21)

- By 2020…
  - FIB: Will grow 5-9x
  - DRAM, SRAM, TCAM: 16x growth per $

"I/O Data Rates on commodity DRAM devices will increase to over 8 GB/s by 2022"
ITRS 2007 roadmap

But what about speed?

- Scariest challenge: Update processing
  - Load ~20 full tables on boot, *fast*.
  - ... And do S-BGP style crypto verification

- Limitations: Memory bandwidth, crypto CPU
  - Memory bandwidth: 8.2GB of memory; *today's* memory can handle 1.7GB/sec.
    - Without AIP/S-BGP future router could load in ~30 seconds.
  - With crypto, however...

Crypto overhead still hurts

- Process update: Validate RSA signature
- Trivially parallelized

	2008 (2.8Ghz quad-core)	2020
RSA Validate	35k/sec	480k/sec
AIP/S-BGP Table Load	~141 seconds	~66 seconds

- Worst-case result - crypto acceleration or clever BGP tricks reduce time

Scaling summary

- Assuming continued network growth and semiconductor trends...

✓ An AIP router in 2020 will be cheaper than an IP router in 2007

(From RIB/FIB perspective)

Things I haven't talked about

- AIP still requires DNS to go from name->AIP
- Traffic engineering
- Detecting key compromise
- Key management (2 level hierarchy)
- Hierarchical AIP addresses
  - beyond the 2-level flat hierarchy presented here
- AIP's benefits to mobility (HIP/TCP Migrate)

Conclusion

- Q: How to achieve network-layer accountability in an internetwork?

- A: **Self-certifying internetwork addresses**
  - AD:EID (AIP)
  - Each field derived from public keys
- **Accountability *intrinsic*** - has many uses
- We believe AIP will scale
  AIP composes well with mechanisms for mobility, DoS mitigation, availability, etc.

Cryptographic Evolution

- Each crypto version: 1 combination of algorithm and parameters
- To move to new one:
  - Add support in *all* routers
  - Once reasonably global, start using
  - Begin phase-out of old version
- We anticipate ~5+ year cycle for this
- (Must pre-deploy one alternate version)

What is an AD?

- Group of addresses that
  - Are administered together
  - Would fail together under common failures
- Examples:
  - A campus, a local organization
- Non-examples:
  - CMU Pittsburgh / CMU Qatar
    - (Each would be different AD)

Traffic Engineering

- ADs are good match for inbound TE techniques - granularity of campus/ customer/reachable subnet
- If need finer-grained:
  - Note ECMP unchanged;
  - Note DNS load-balancing unchanged;

  - AIP address interface bits to sub-divide AD
  - 8 bits of interface space
  - partition to up to 255 “paths” to a domain

Handling Key Compromise

- ***Preventing***:
  - Two-level key hierarchy (master signs offline; routers have temporary key)
- ***Detecting***:
  - Registry of addresses used
    - e.g., AD registers "EID X is connecting through me"
    - Registries simple: entirely self-certifying
- ***Recovering***:
  - Renumber + (self-certifying) revocation registry

Shut-Off Replay Prevention

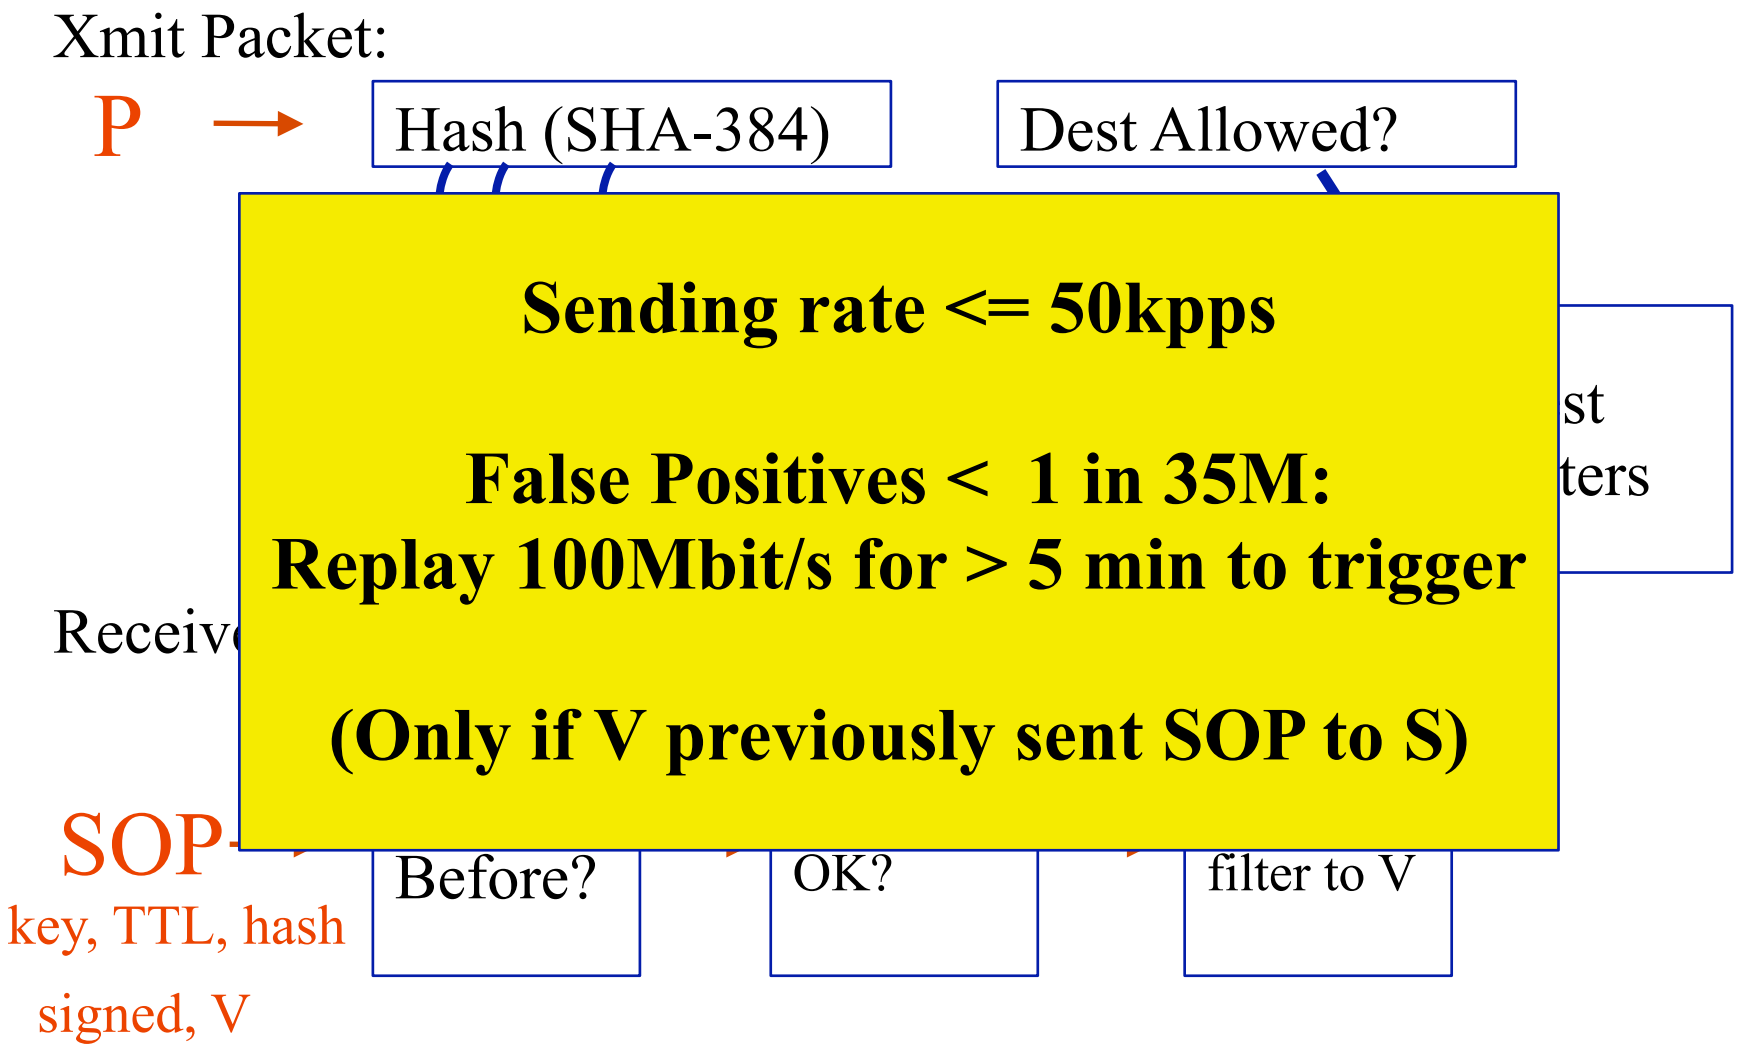

Mutual Shut-Off

- Attack:
  - Zombie Z wants to flood victim V
    - First, Z pings V. Gets response back.
    - Z sends Shut-Off packet to V.
    - Z floods V.

- Resolution:
  - Smart-NIC allows V to send SOPs at very low rate (1 per 30 seconds) even though filtered

➡Hosts can mutually shut-off...

AIP Address

AIP Header

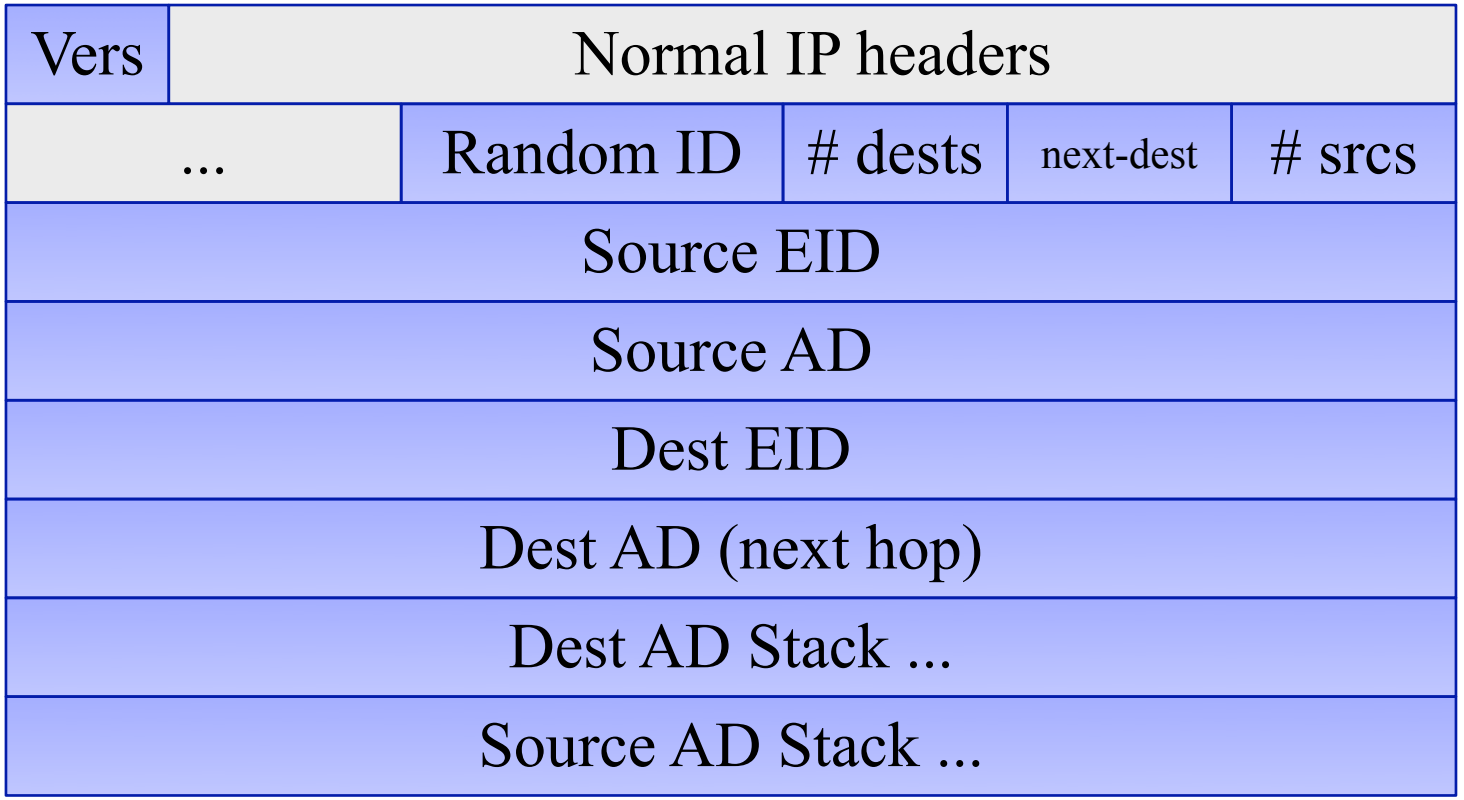

AIP Verification Protocol

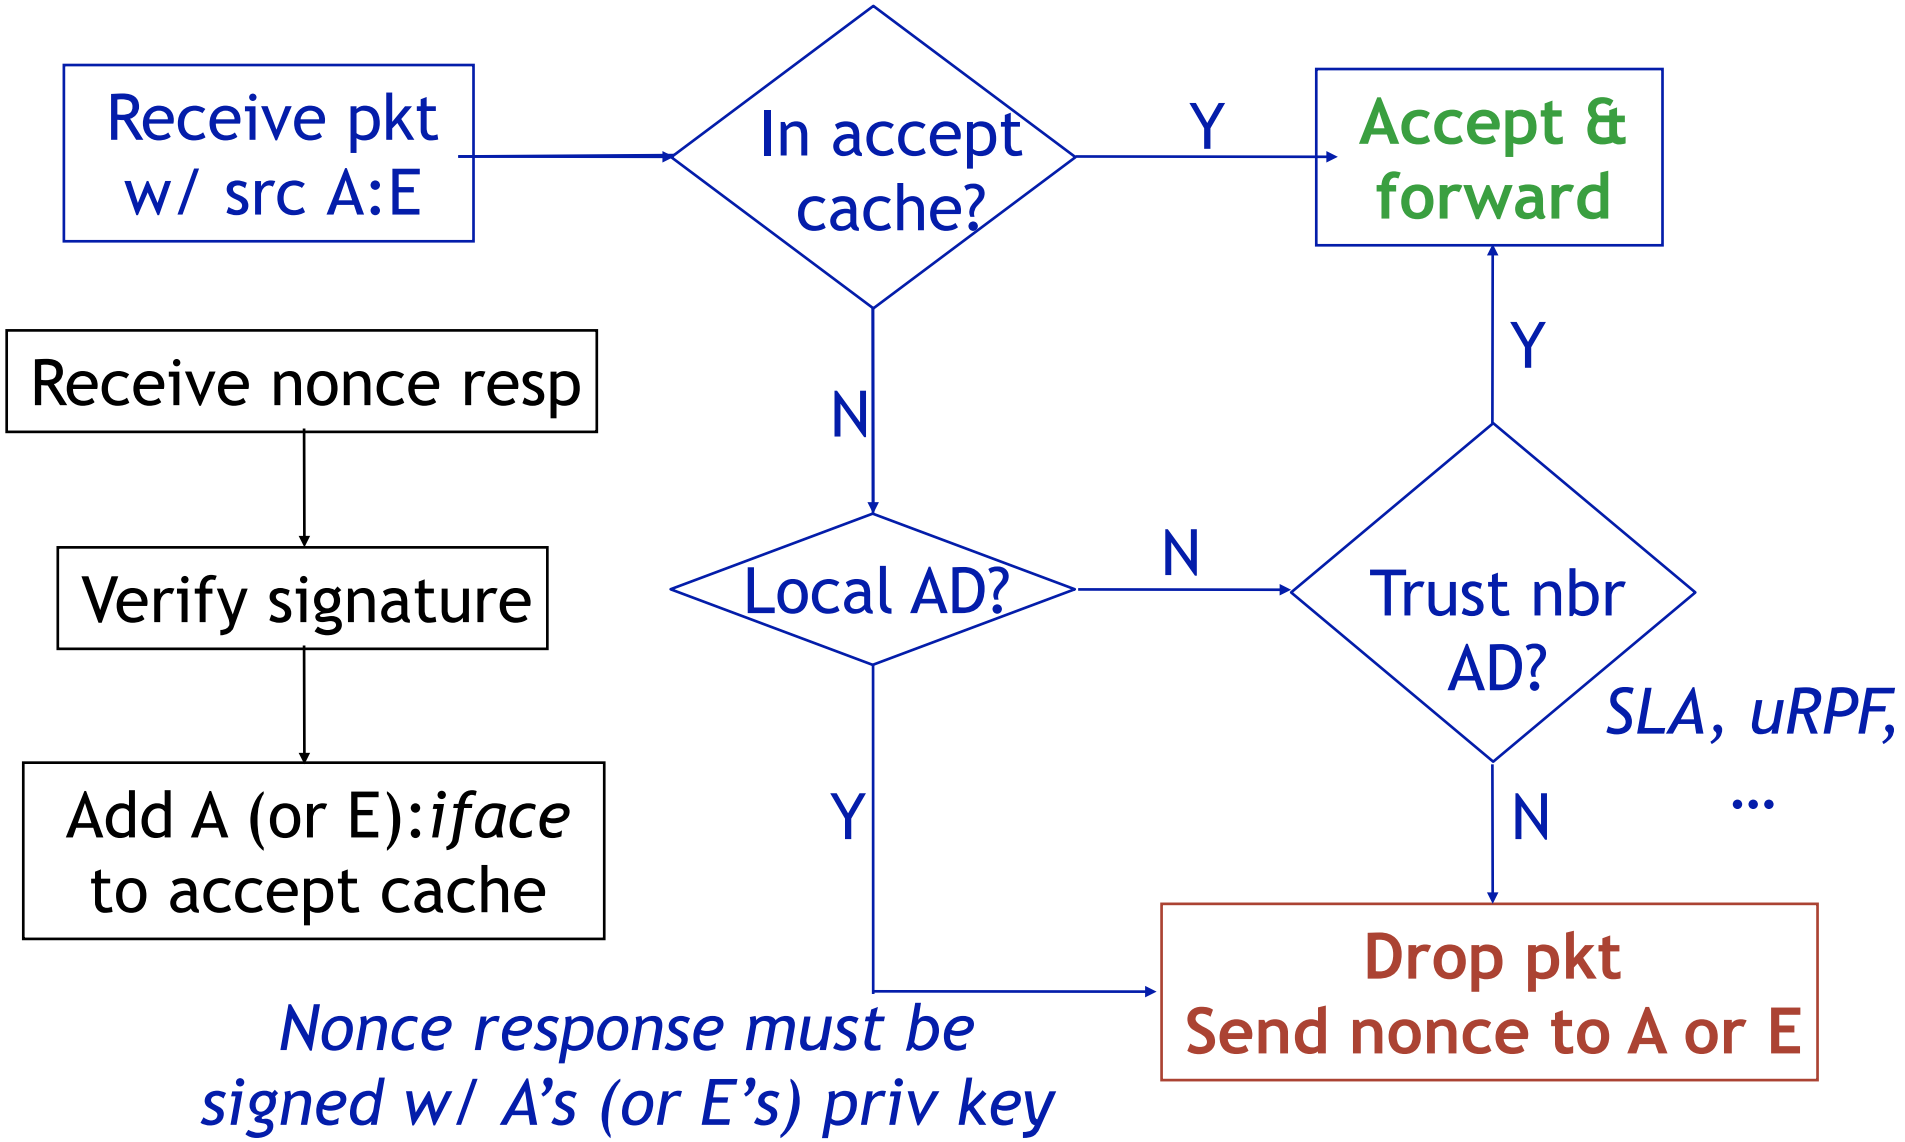

Protecting Those who Protect Themselves

- To bound size of accept cache,
  - if too many entries of AD:x, AD:x2, ...
  - Upgrade to “wildcard”: AD:*
- If many compromised hots in AD, they can allow others to spoof AD
- If AD secure, nobody can spoof it

Table Size Projections

Year	17% Growth	Fuller/Huston
2008	Observed: 247K	
2011	396K	600K-1M
2020	1.6M	1.3-2.3M

- 17% growth and predictions from Fuller & Huston; rough agreement for 2020