Carnegie Mellon University

BALANCING ACCOUNTABILITY & PRIVACY IN THE NETWORK

David Naylor

Matt Mukerjee

Peter Steenkiste

ACCOUNTABILITY

operators want to **know who sends each packet**

so they can stop malicious senders

VS

PRIVACY

users want to **hide who sends certain packets**

so they can do stuff without the whole world knowing

ACCOUNTABILITY

Accountable Internet Protocol

[Andersen et al., SIGCOMM 2008]

cryptographic addresses

anti-spoofing mechanism
+ shutoff protocol

No Privacy

Shutoff is Stop-Gap Fix

Requires “Smart NIC”

VS

PRIVACY

Tor Instead of IP

[Liu et al., HotNets 2011]

routers act as onion nodes

No Accountability

Heavyweight

ACCOUNTABILITY

Accountable Internet Protocol

[Andersen et al., SIGCOMM 2008]

unforgeable **source addresses**

PRIVACY

Tor Instead of IP

[Liu et al., HotNets 2011]

hidden **source addresses**

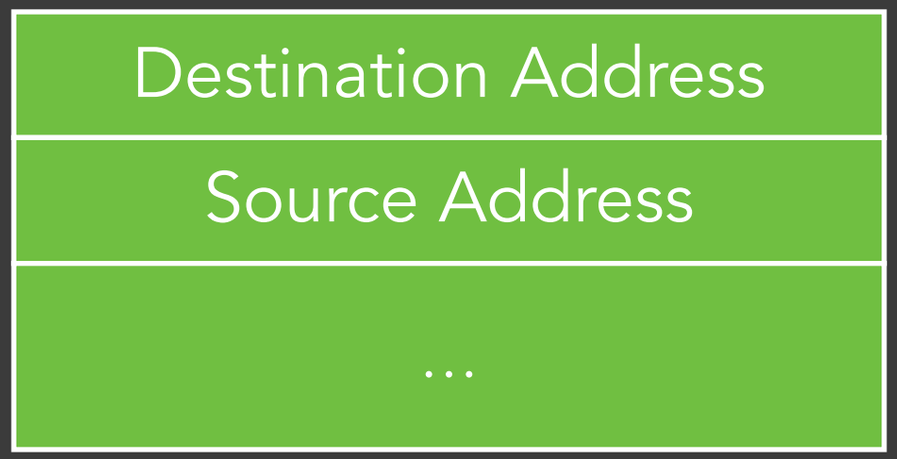
Destination Address
Source Address
. . .

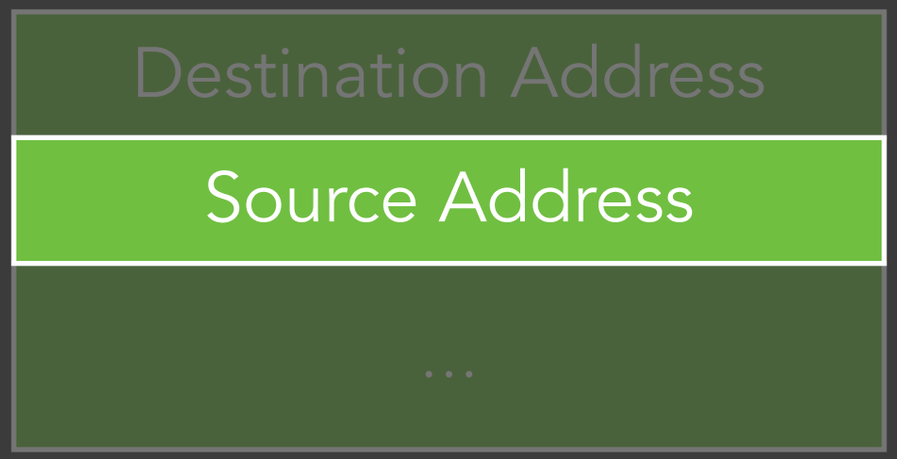

return address

accountability

sender identity

error reporting

flow ID


Destination Address
Source Address
...

Separate Accountability and Return Addresses

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

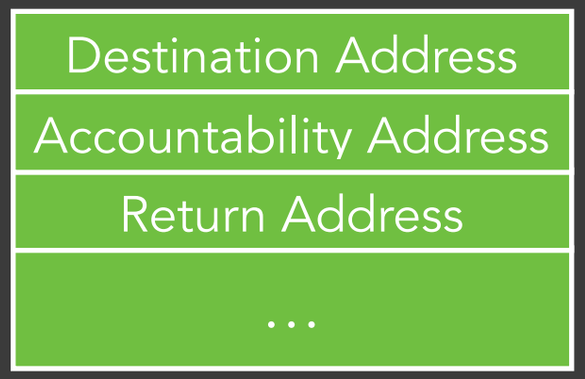

Separate Accountability and Return Addresses

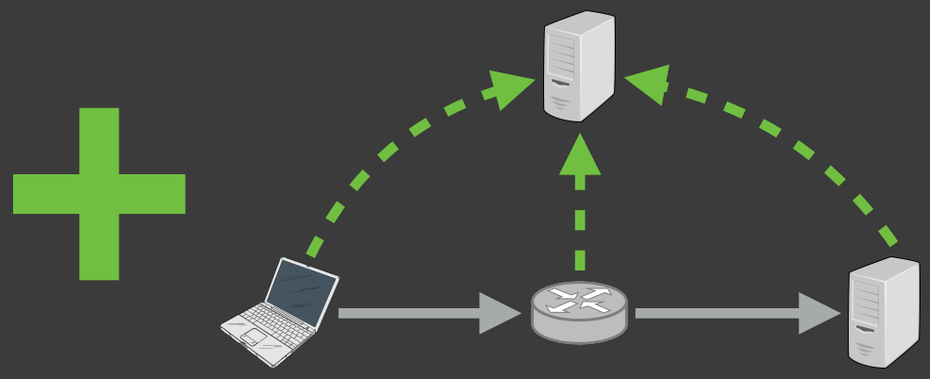

Delegated Accountability

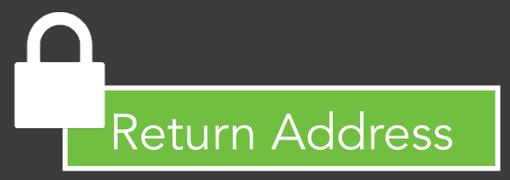

Hidden Return Addresses

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

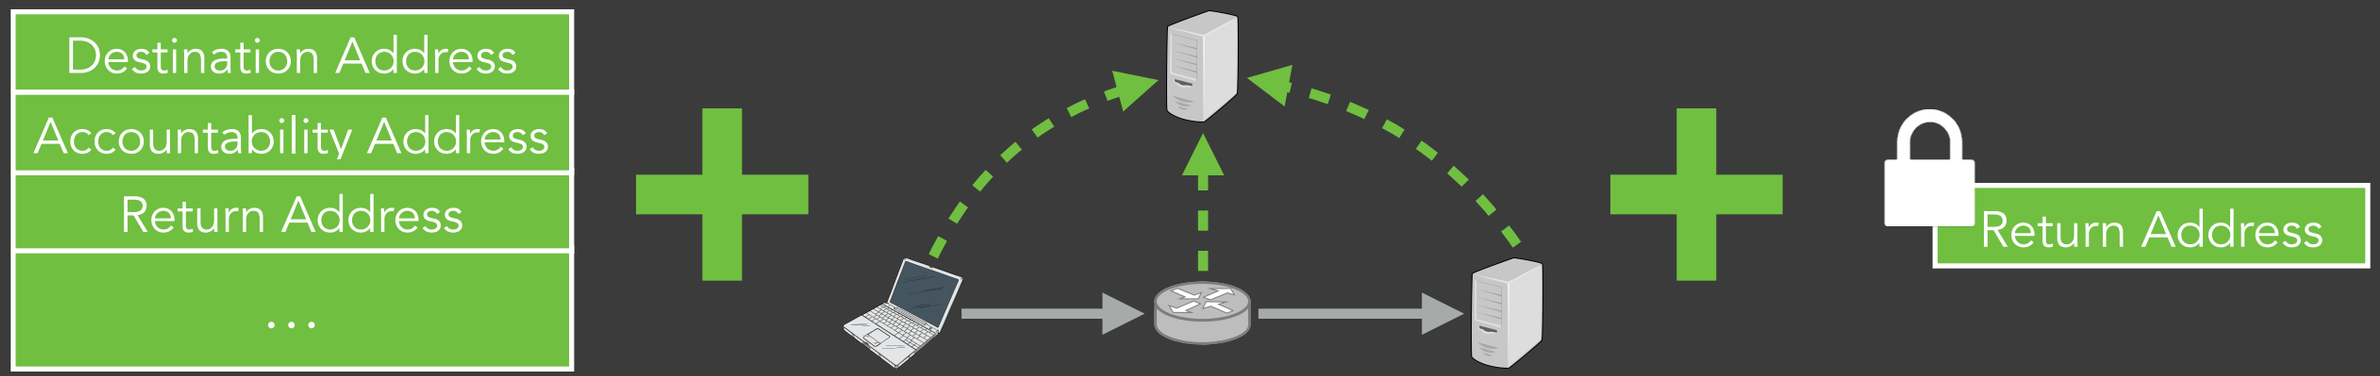

Separate Accountability and Return Addresses

Delegated Accountability

Hidden Return Addresses

Real-World Deployment

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

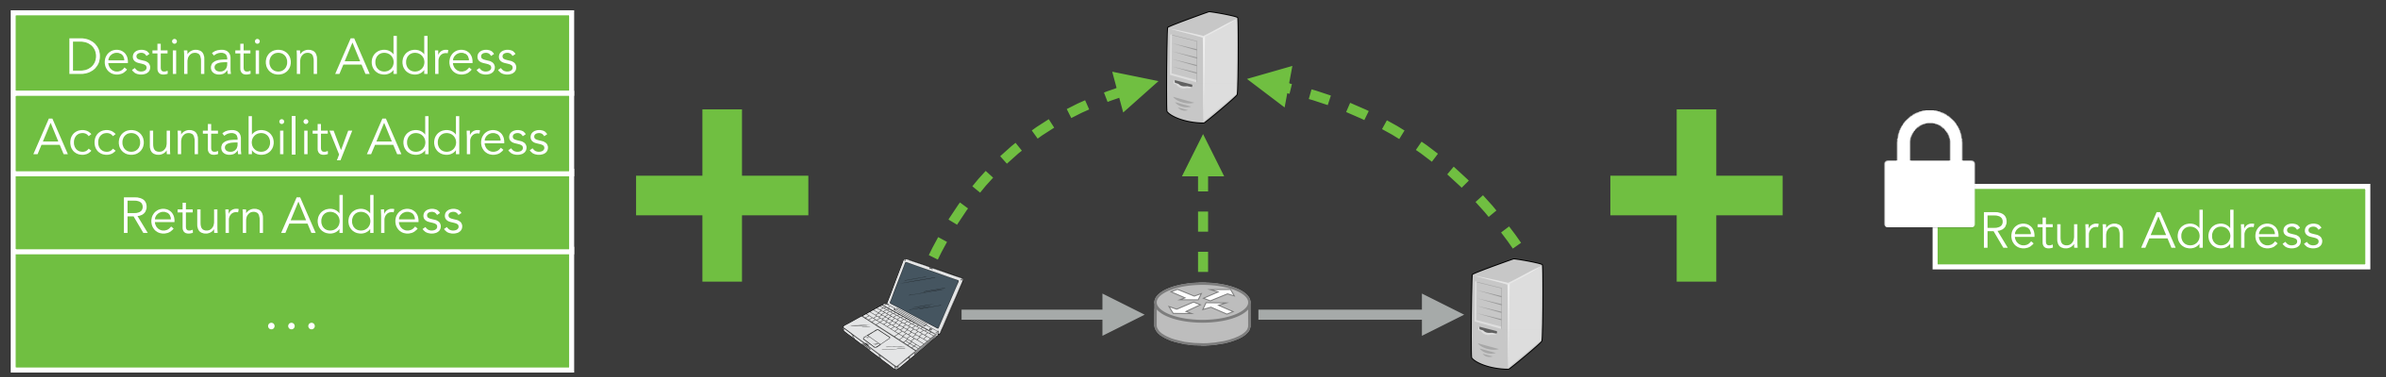

Separate Accountability and Return Addresses

Delegated Accountability

How it Works

Feasibility

Flow Granularity

Hidden Return Addresses

Real-World Deployment

DELEGATED ACCOUNTABILITY

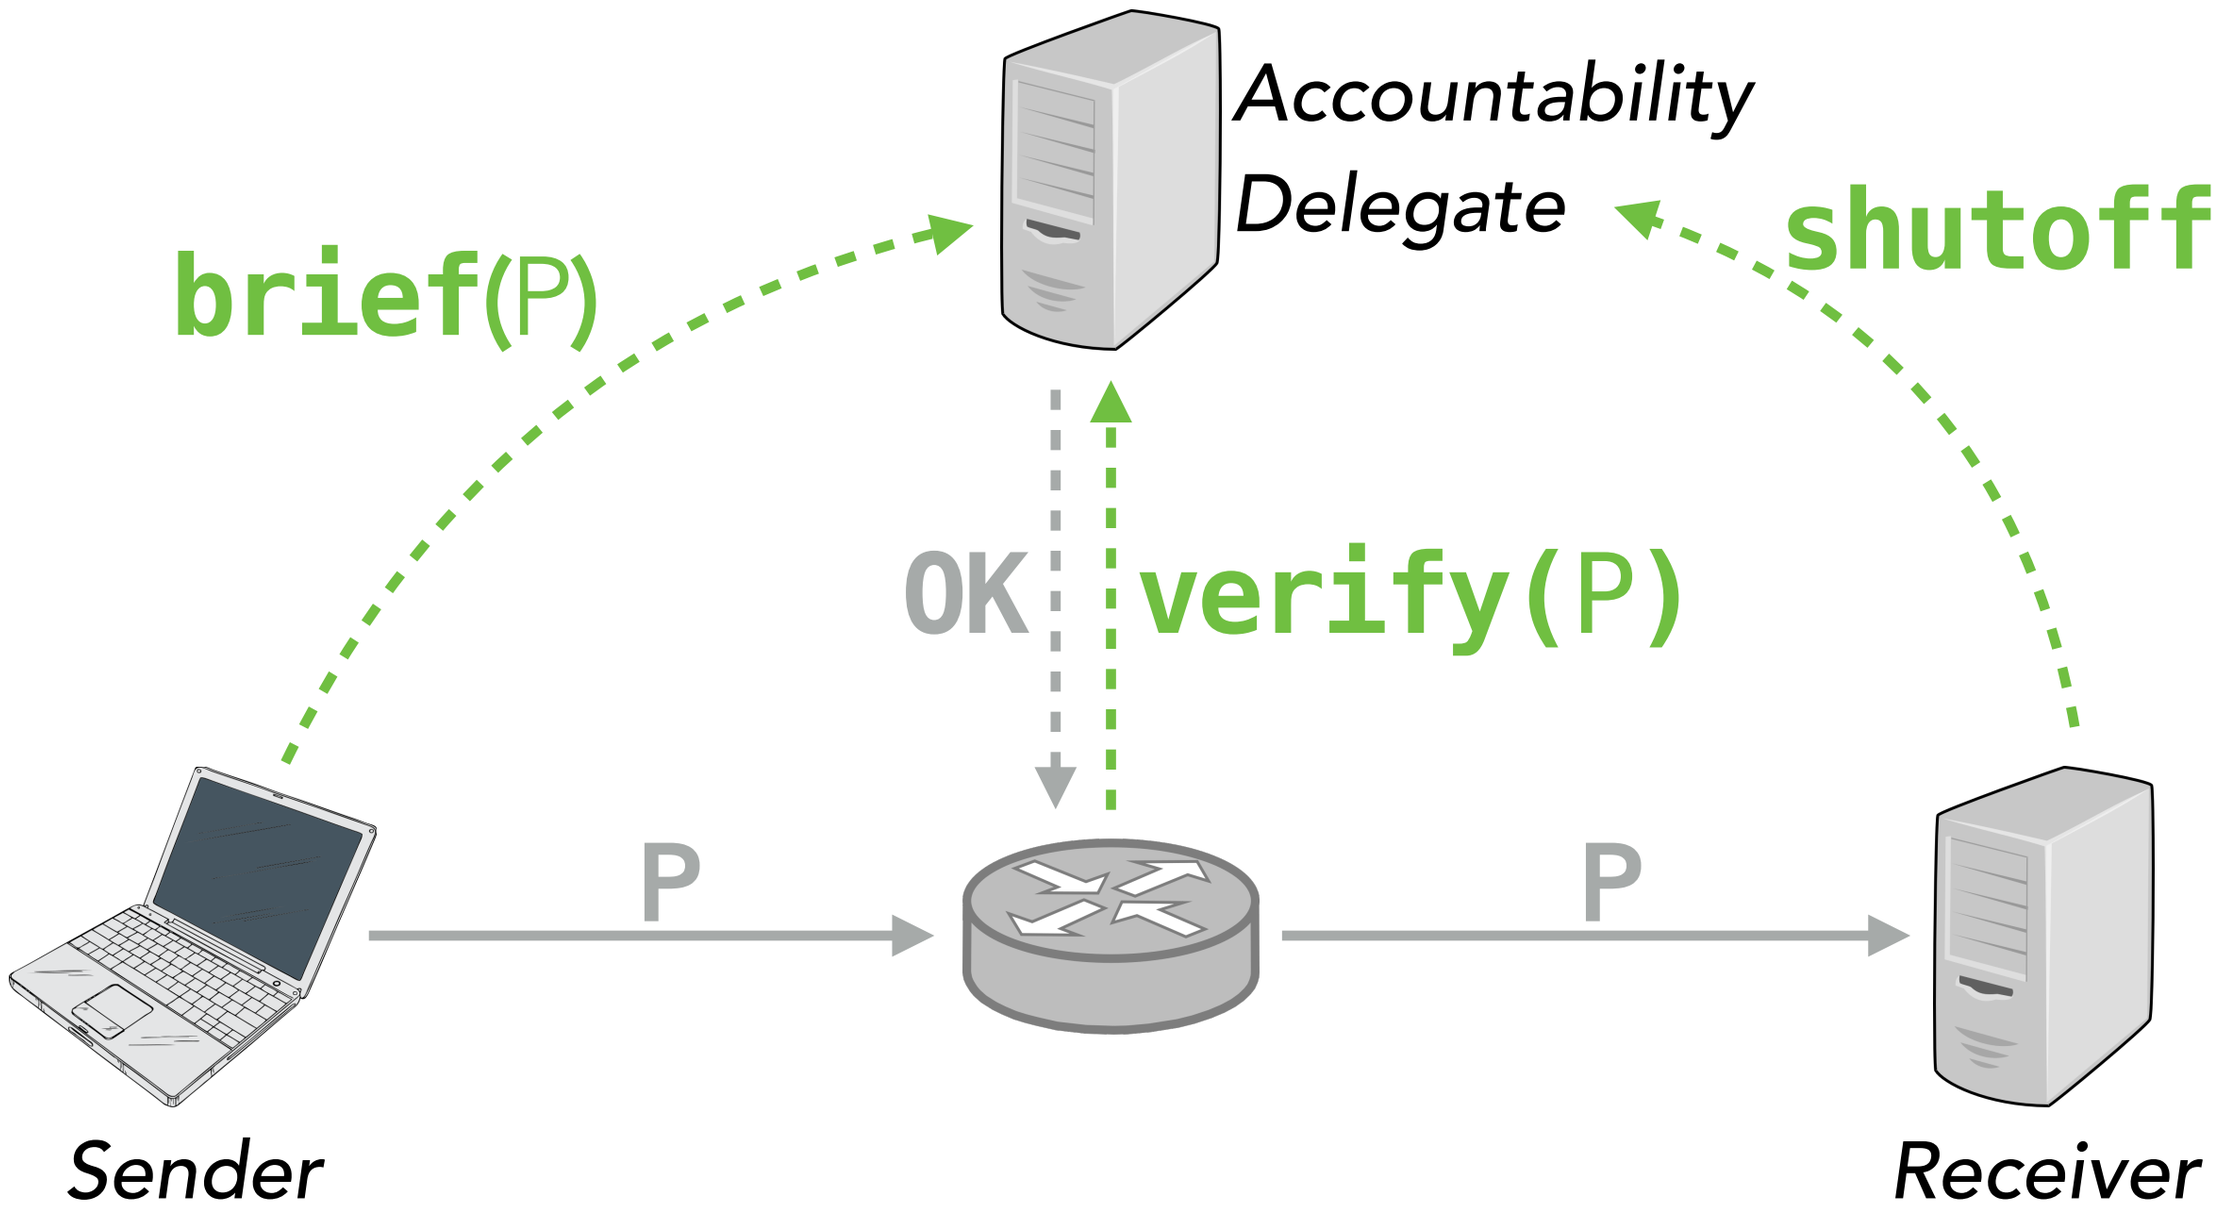

```
brief(P)
```

Sender to Delegate:

“I sent this packet.”

brief(P)

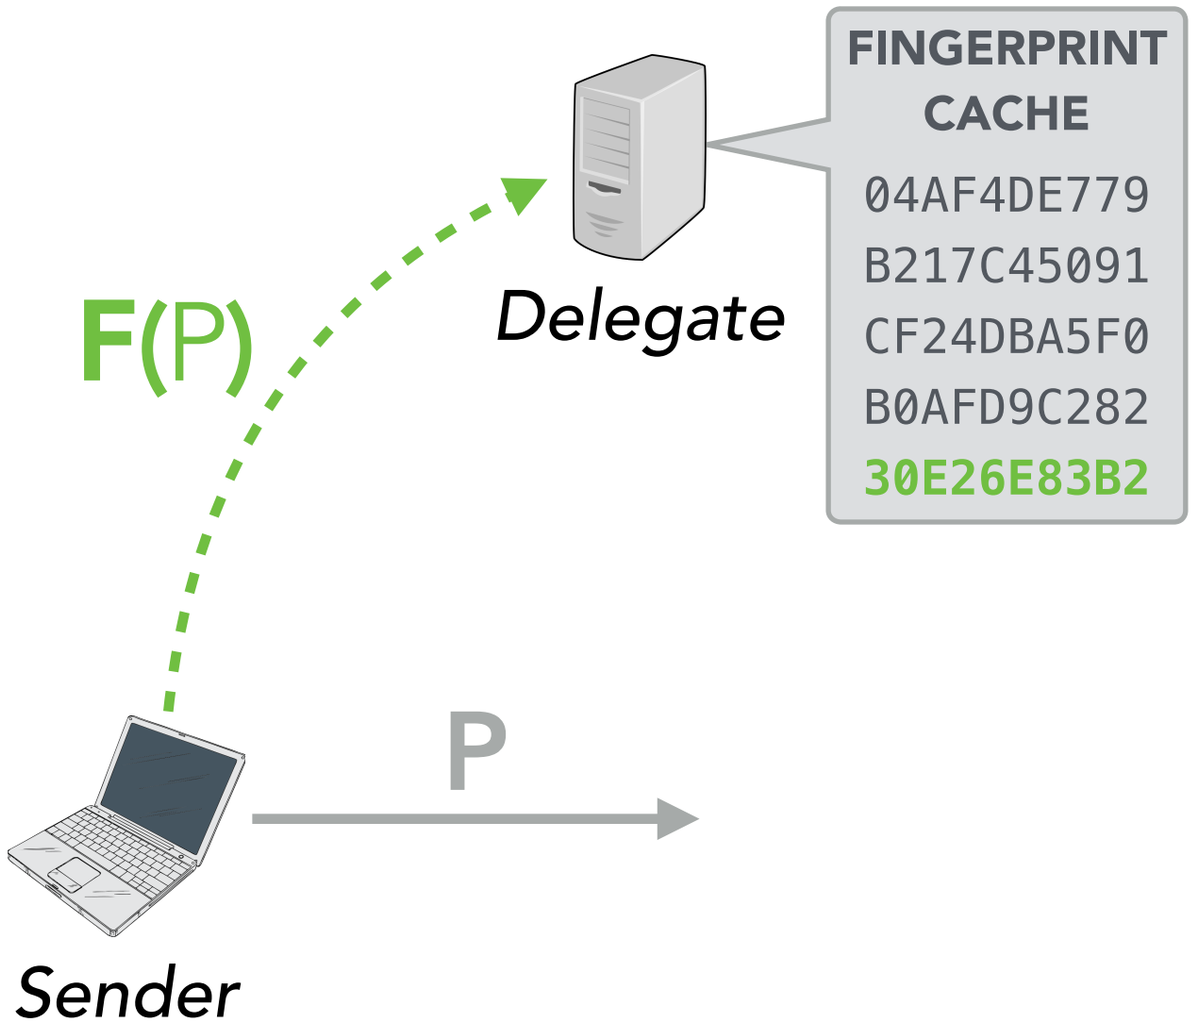

Batch fingerprints in Bloom filter

Delegate does not learn packet contents

DELEGATED ACCOUNTABILITY

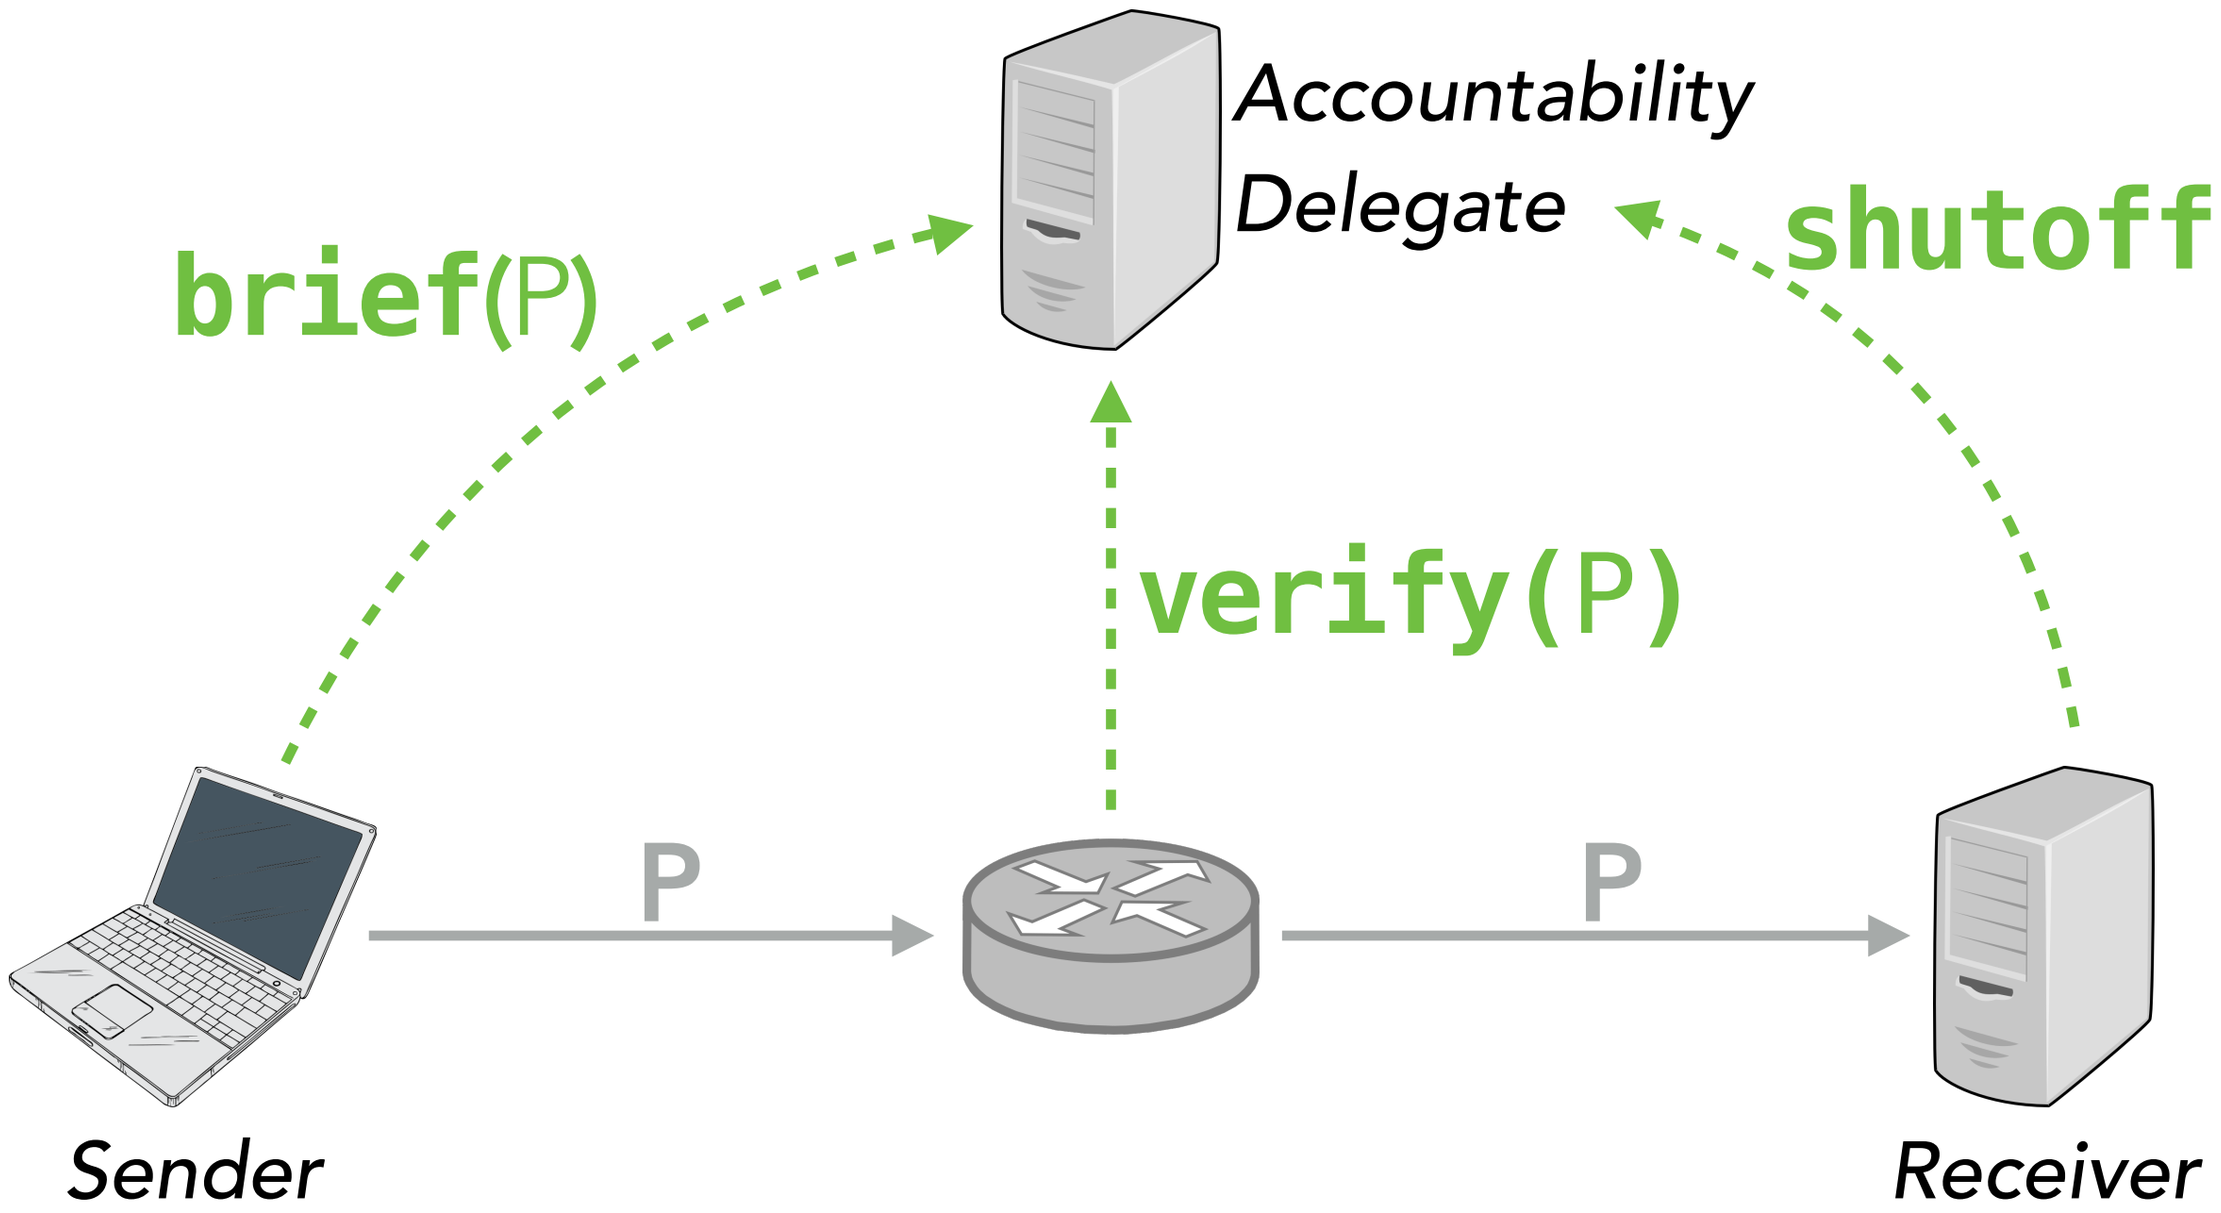

```
verify(P)
```

Verifier to Delegate:

"Do you vouch for this packet?"

`verify(P)`

TWO CHECKS:

1. $P_{A \rightarrow B}$ in fingerprint cache
2. Flow A→B not shut off

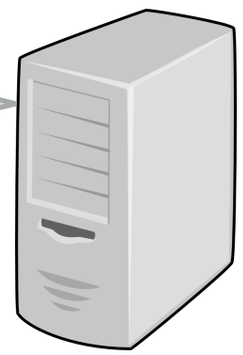

A's Delegate

`verify(P)`

OK

`verify(P)`

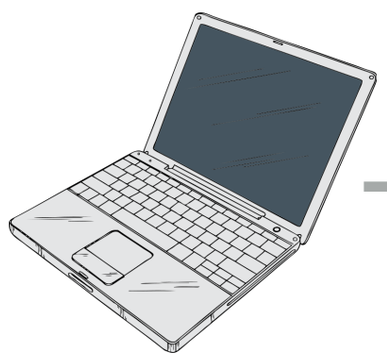

A

$P_{A \rightarrow B}$

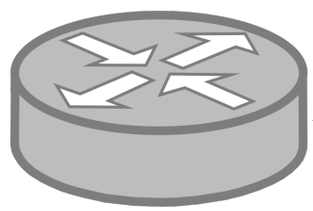

VERIFIED FLOWS

`A → B`

verify(P)

TWO CHECKS:

1. $P_{A \rightarrow B}$ in fingerprint cache
2. Flow A→B not shut off

A's Delegate

Most effective at first hop

Verified flow entries periodically expire

Routers keep no state during verification

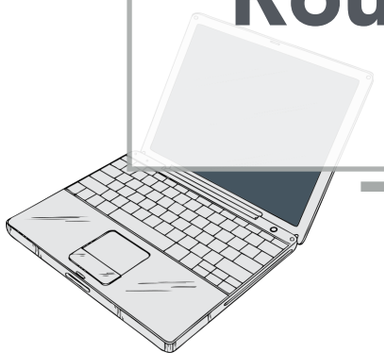

A

VERIFIED FLOWS

A → B

DELEGATED ACCOUNTABILITY

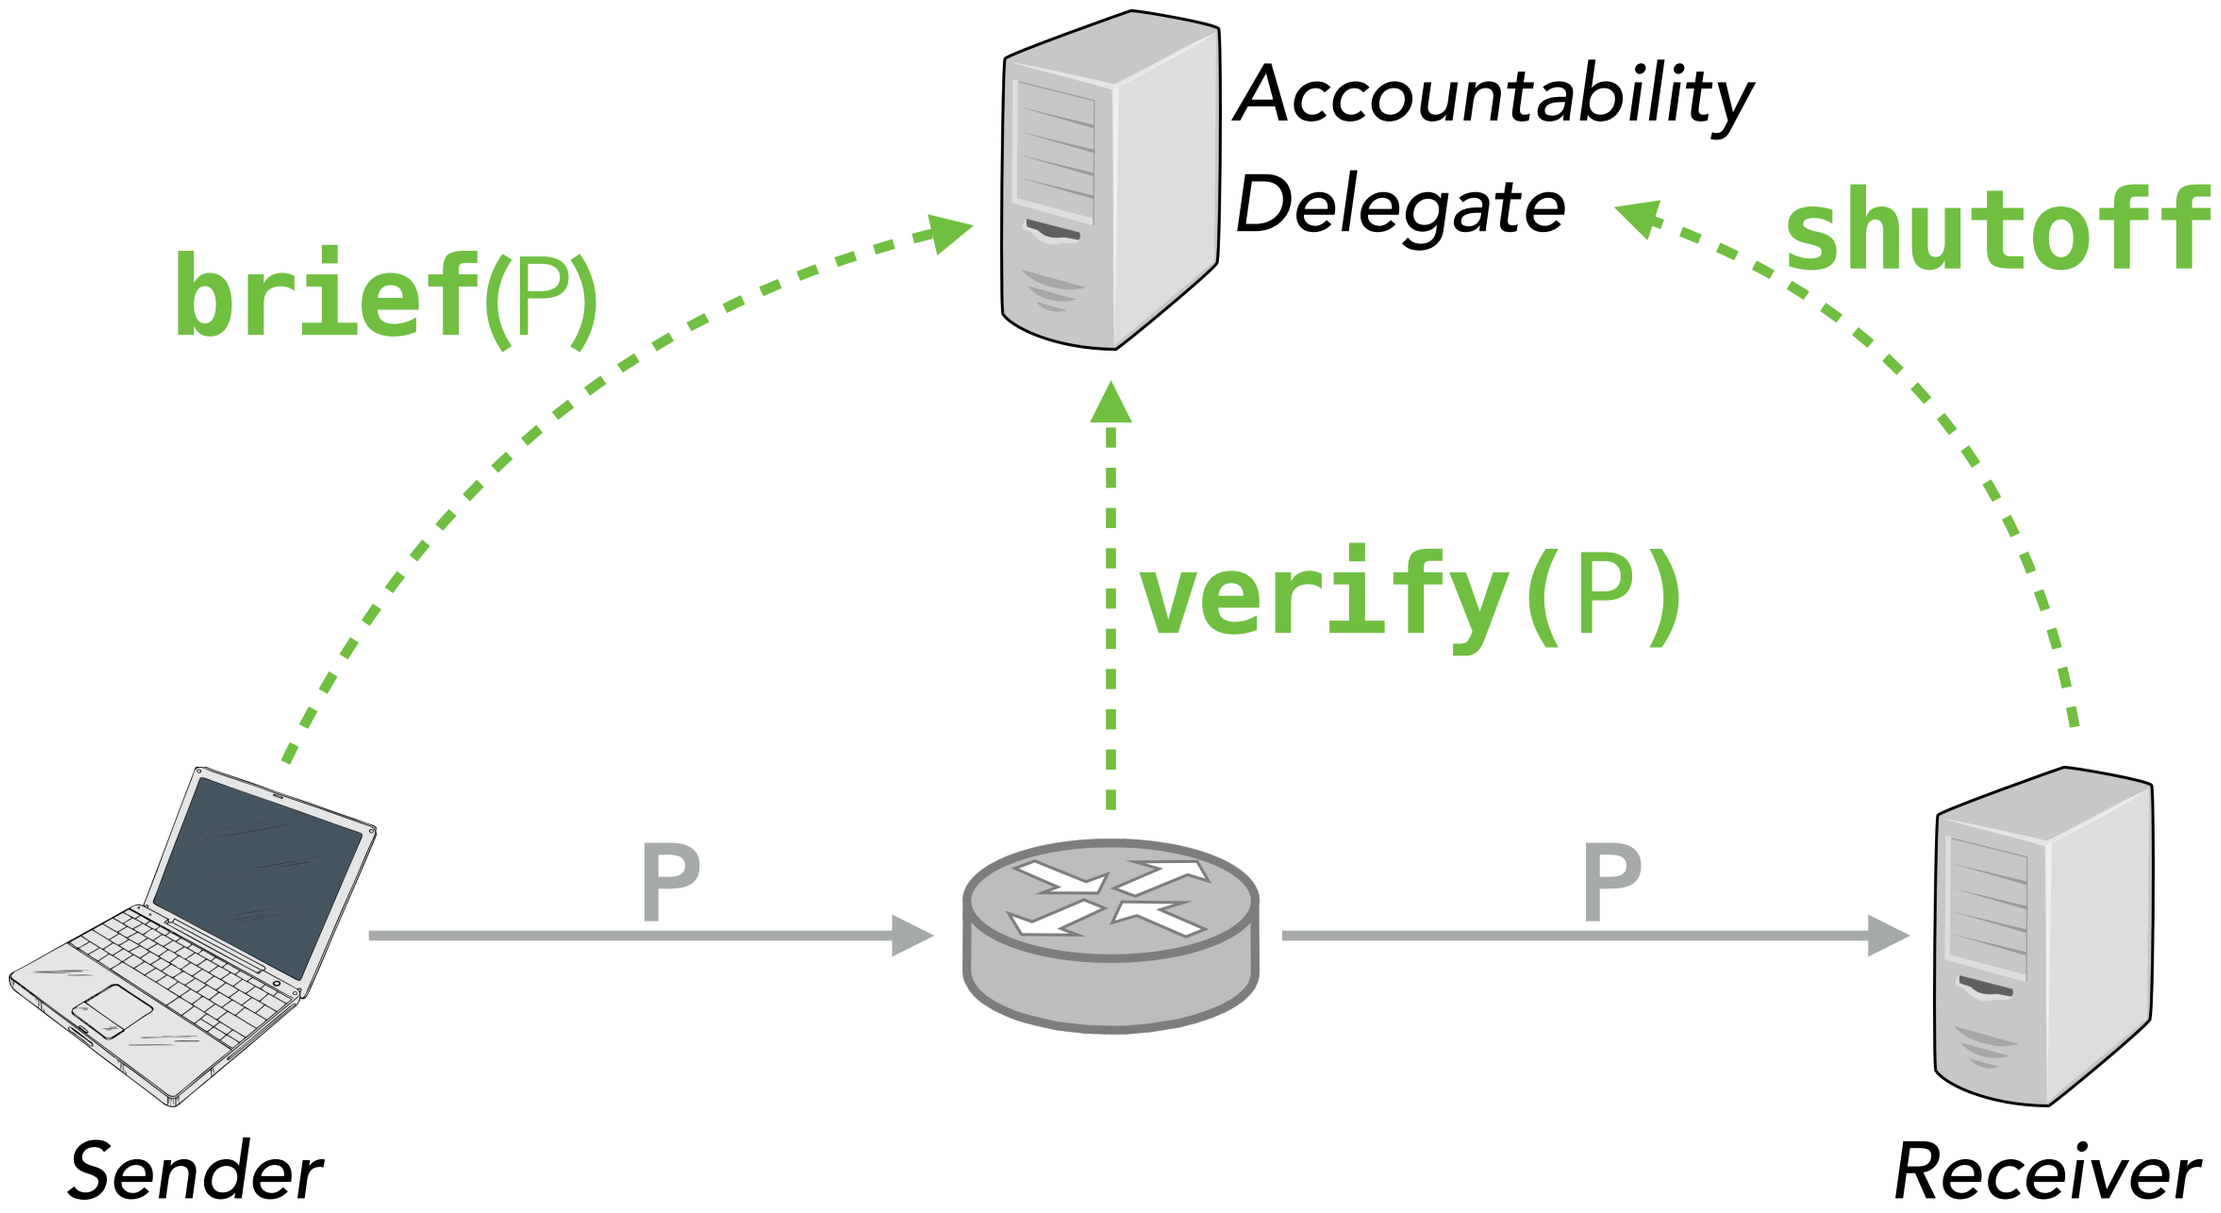

```
shutoff(P)
```

Receiver to Delegate:

"Stop this flow."

shutoff(P)

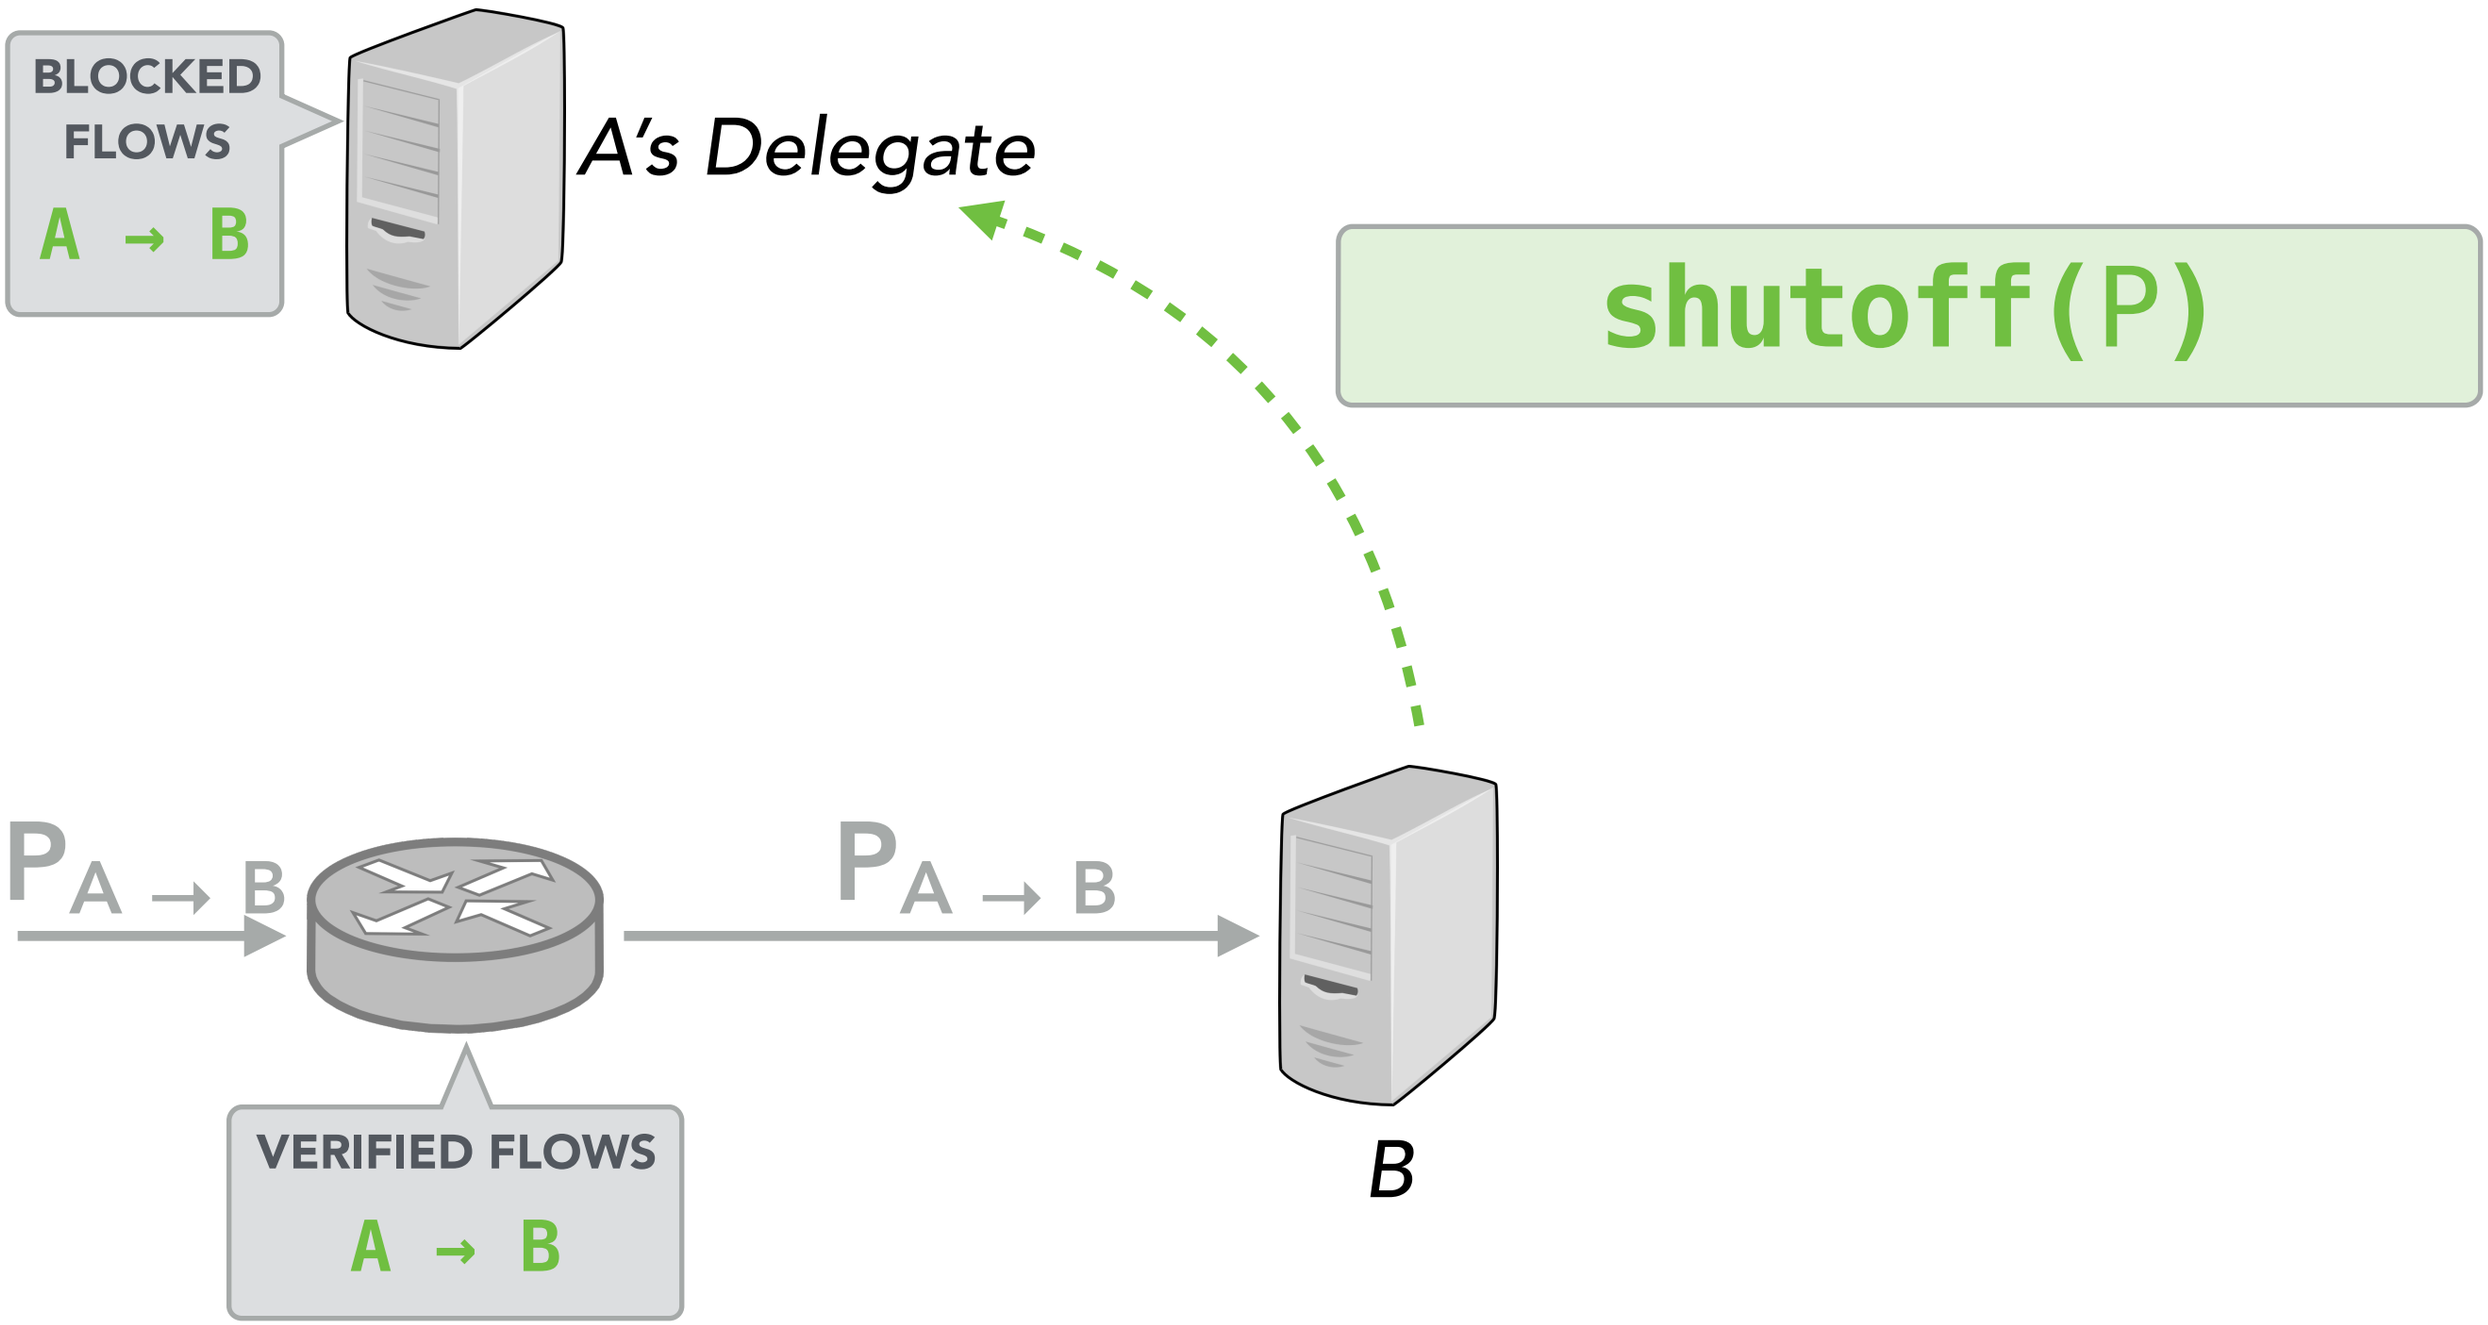
BLOCKED FLOWS
A → B
A's Delegate
shutoff(P)
P A → B
P A → B
B
VERIFIED FLOWS
A → B

shutoff(P)

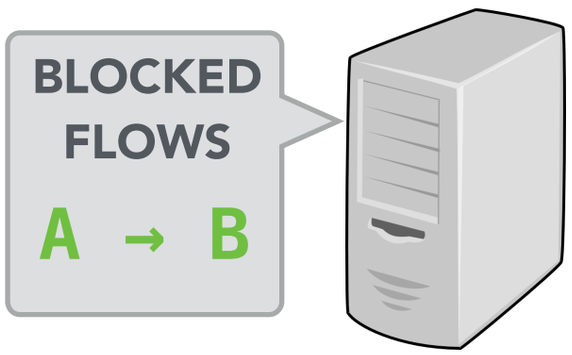

A's Delegate

shutoff(P)

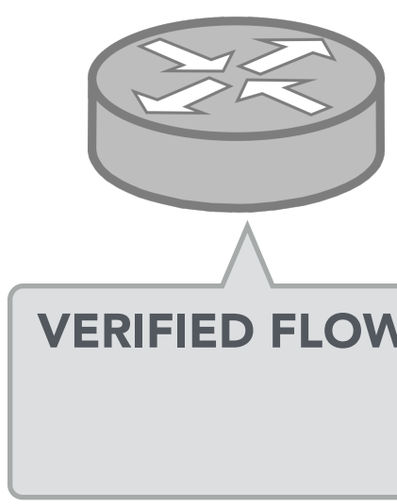

B

shutoff(P)

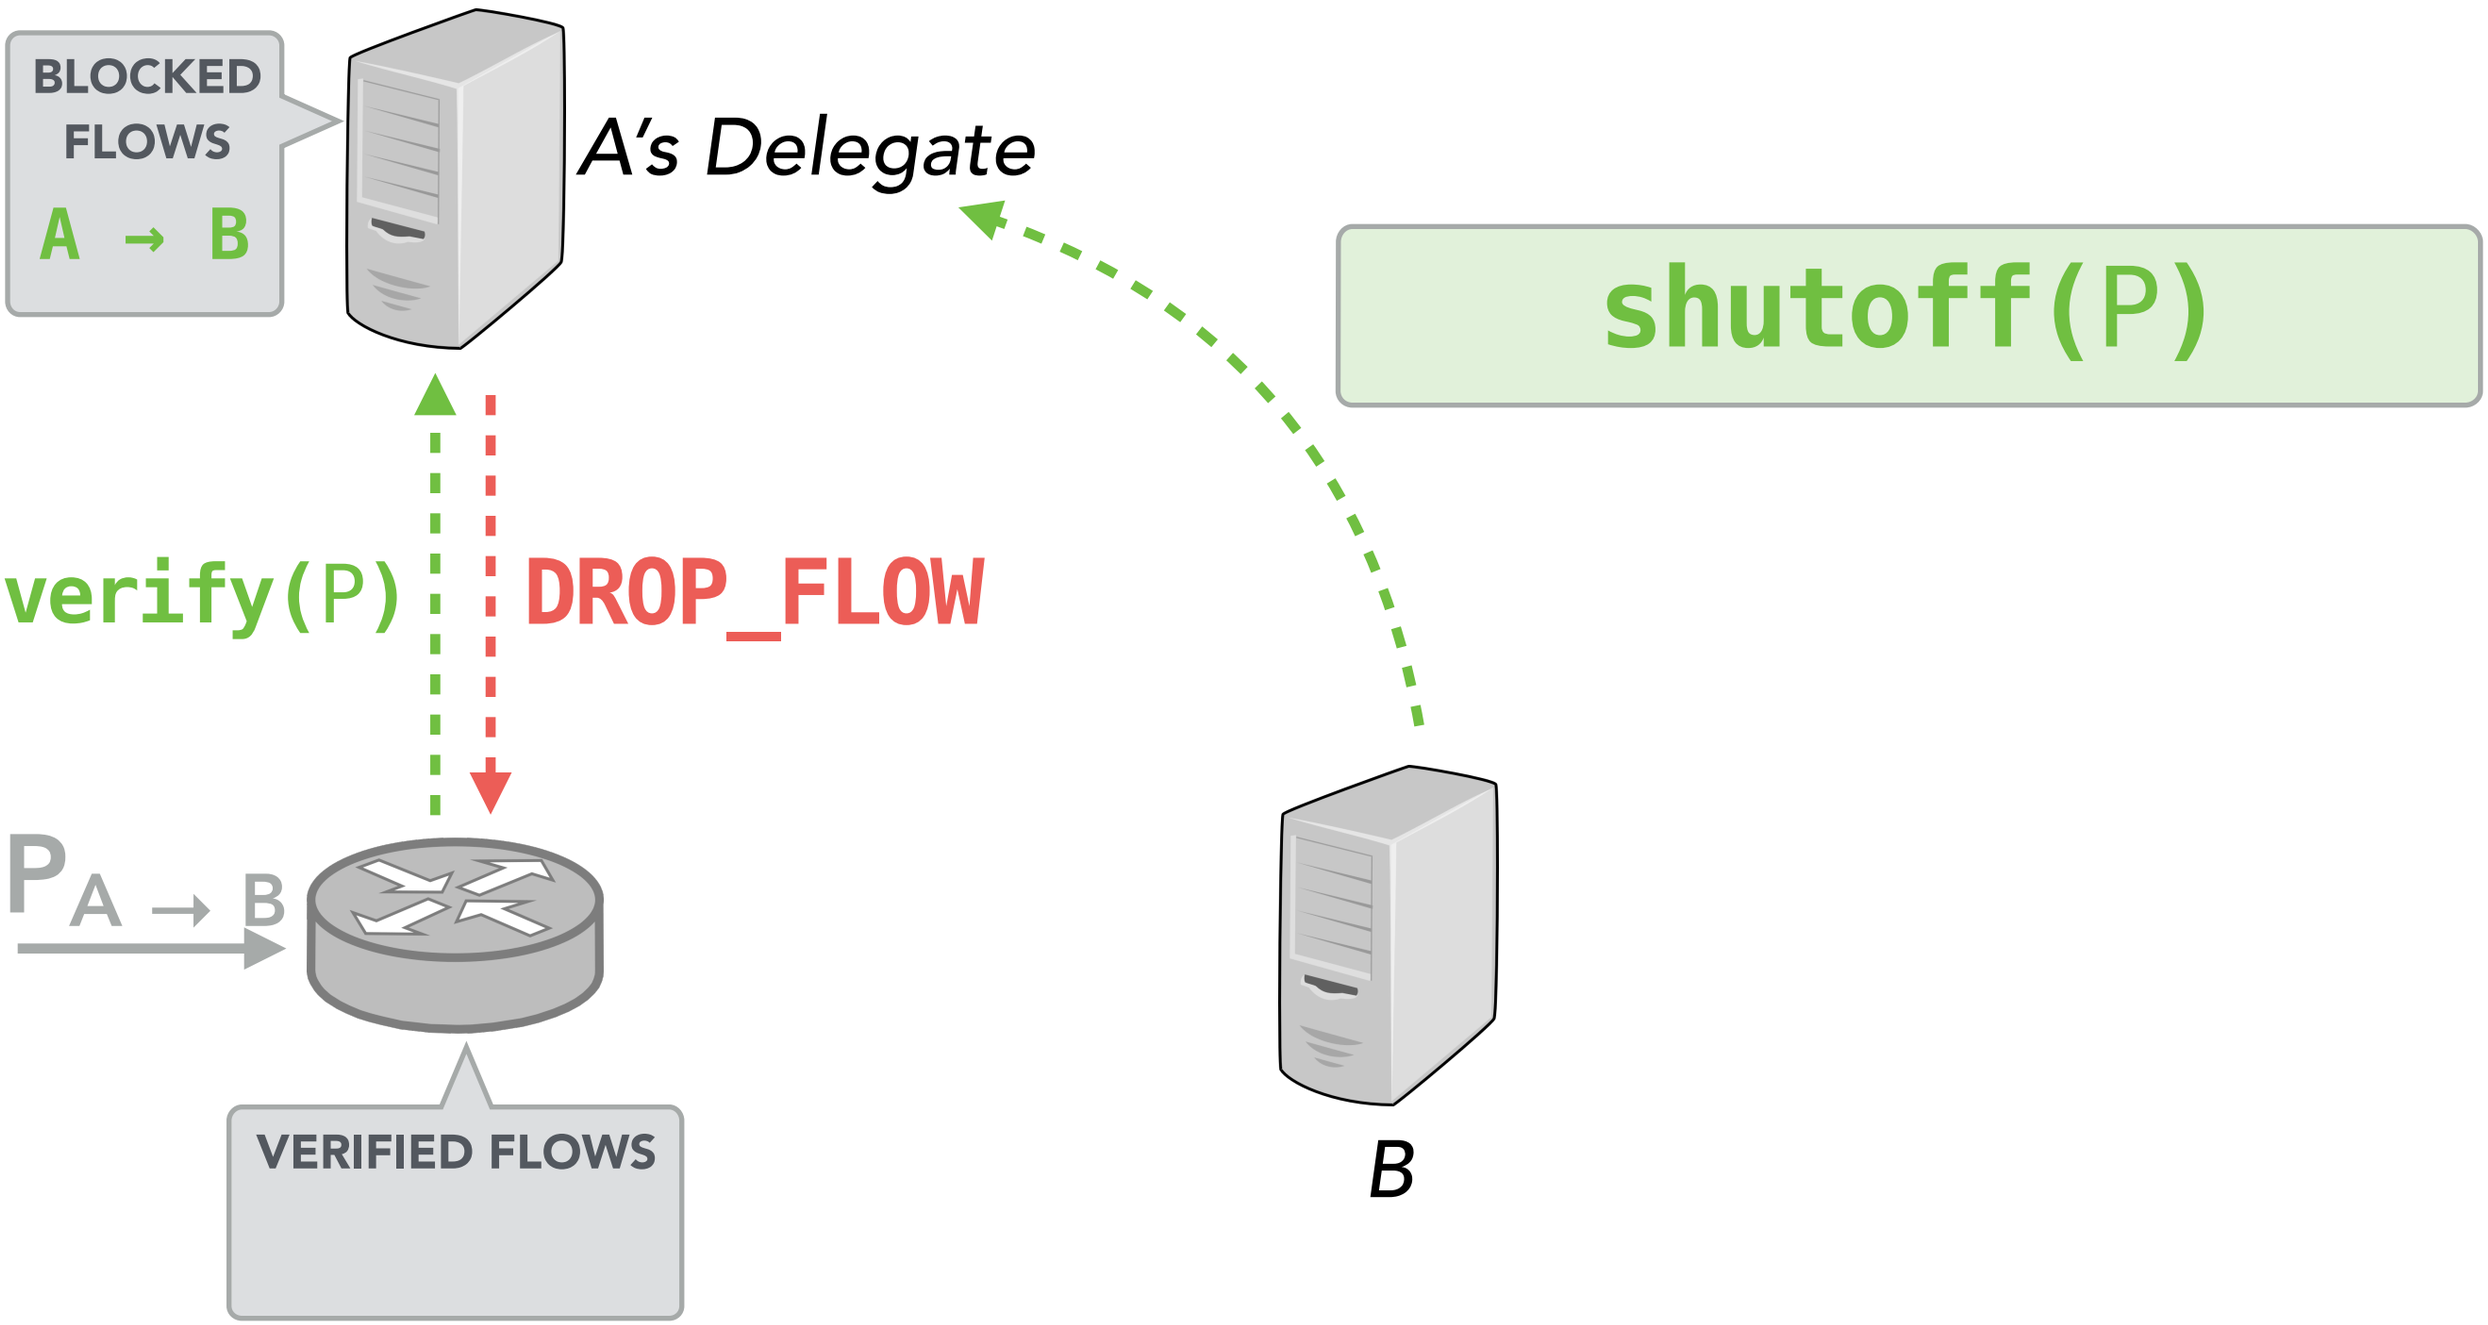
BLOCKED FLOWS
A → B
A's Delegate
shutoff(P)
verify(P)
DROP_FLOW
P A → B
VERIFIED FLOWS
B

shutoff(P)

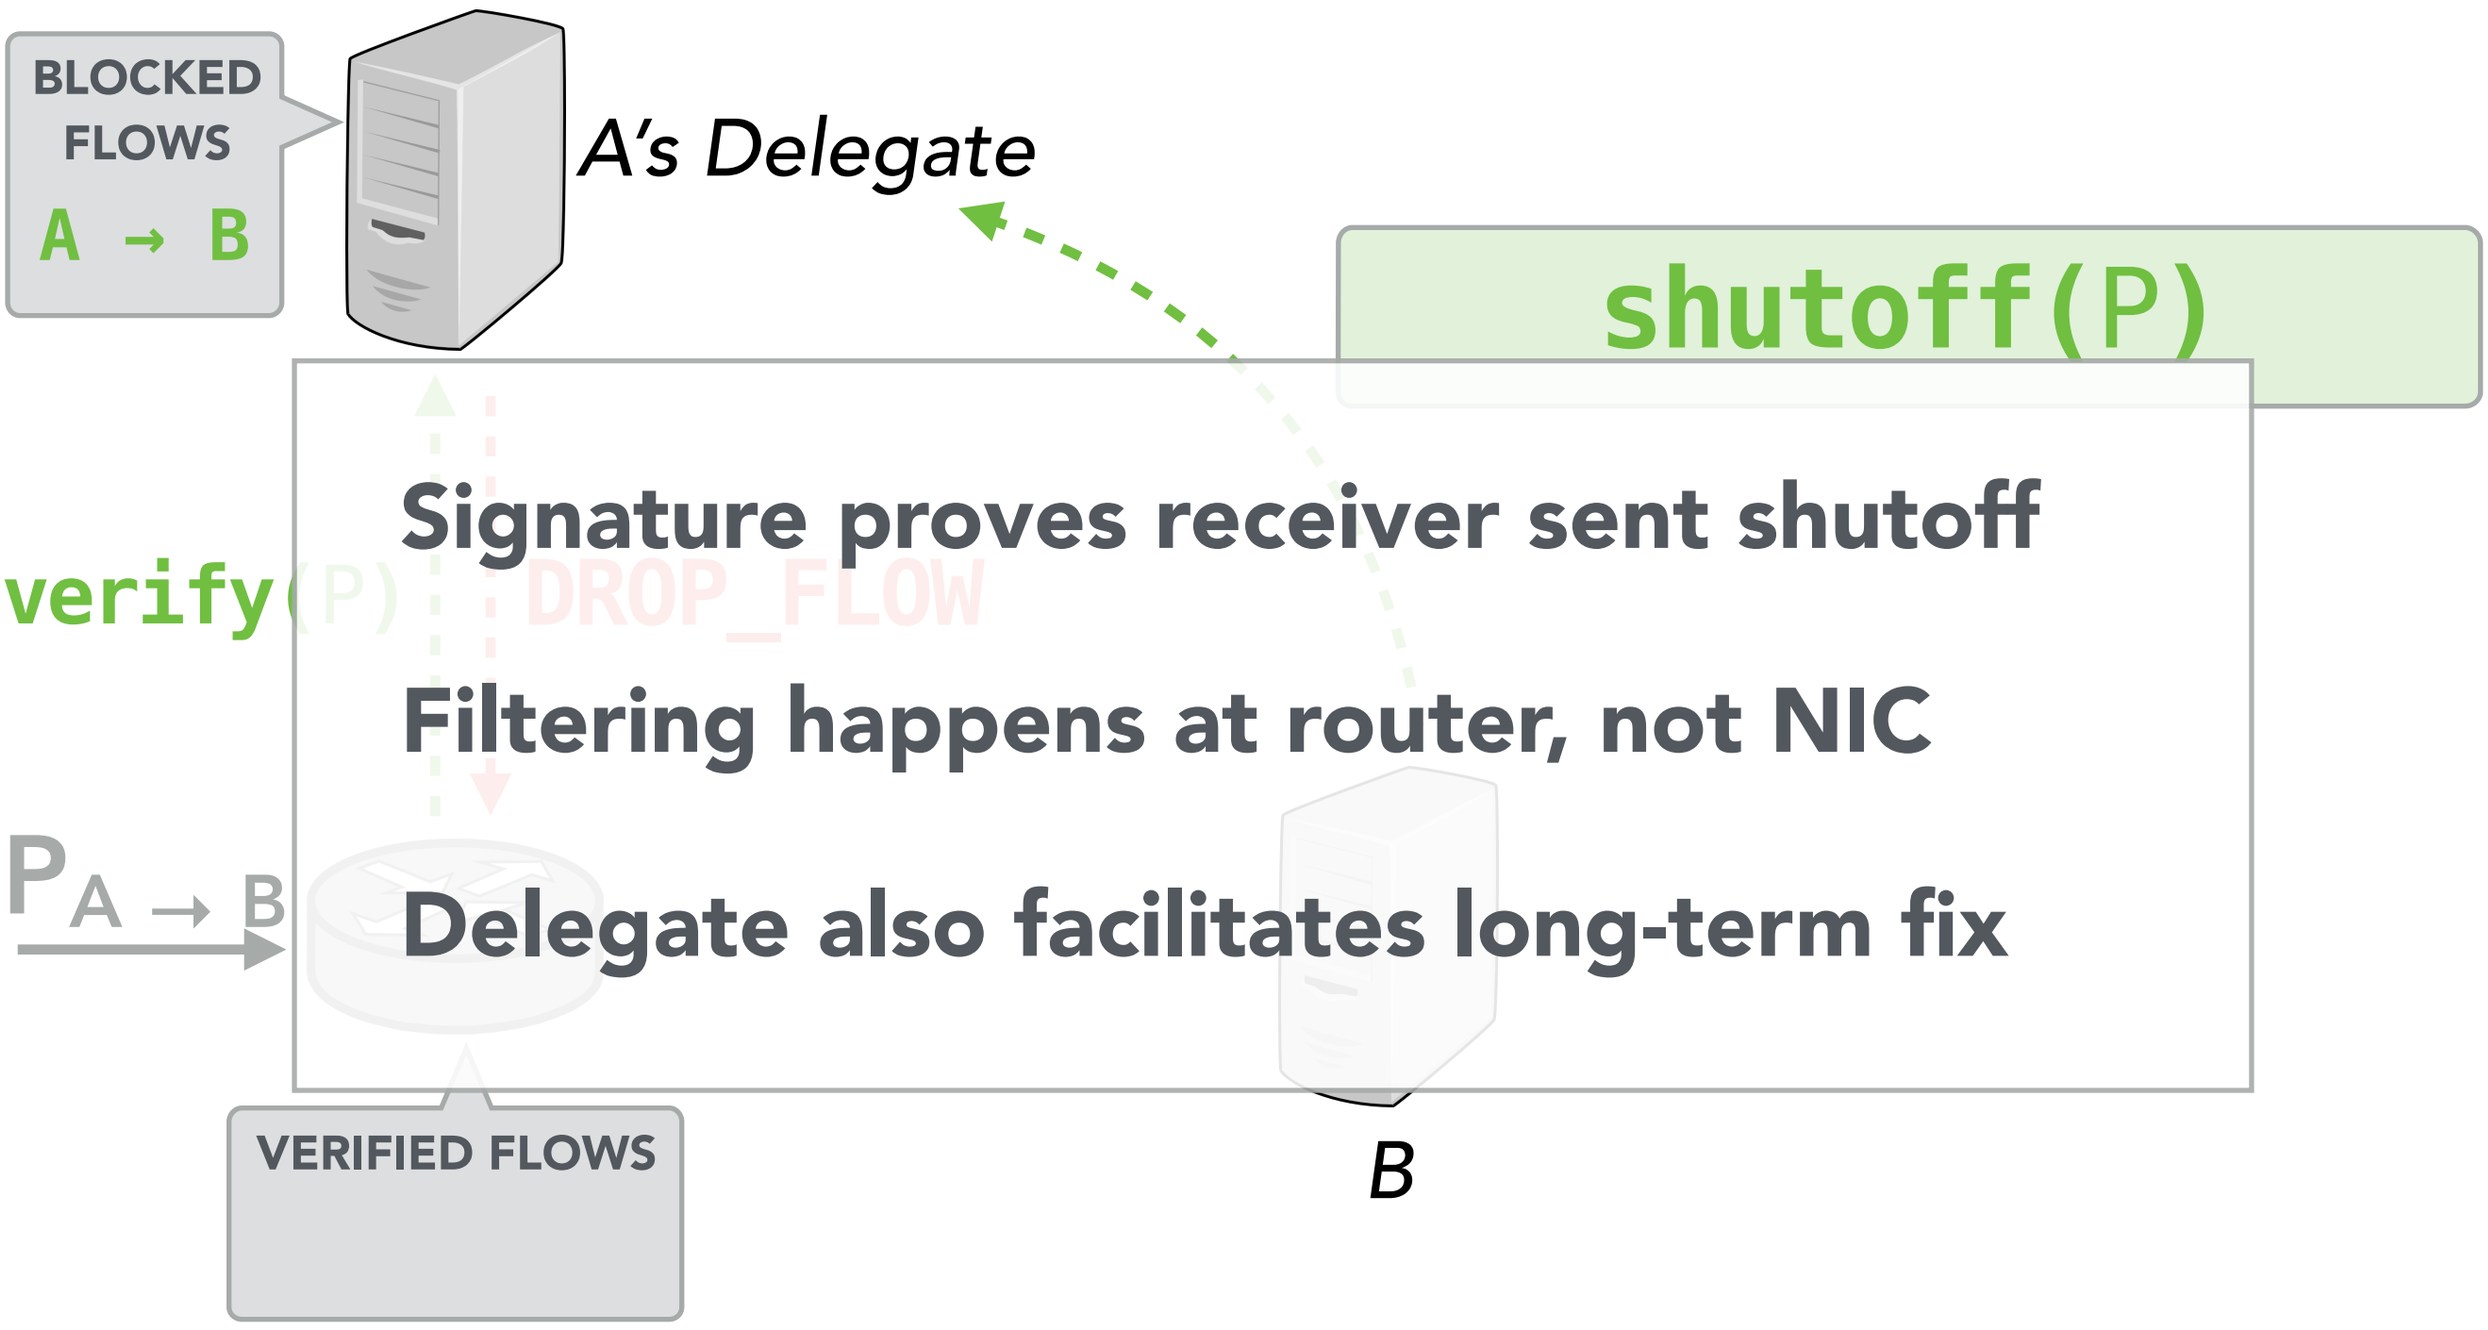
BLOCKED FLOWS
A → B
A's Delegate
shutoff(P)
Signature proves receiver sent shutoff
verify(P) : DROP_FLOW
Filtering happens at router, not NIC
P A → B
Delegate also facilitates long-term fix
VERIFIED FLOWS
B

DELEGATED ACCOUNTABILITY

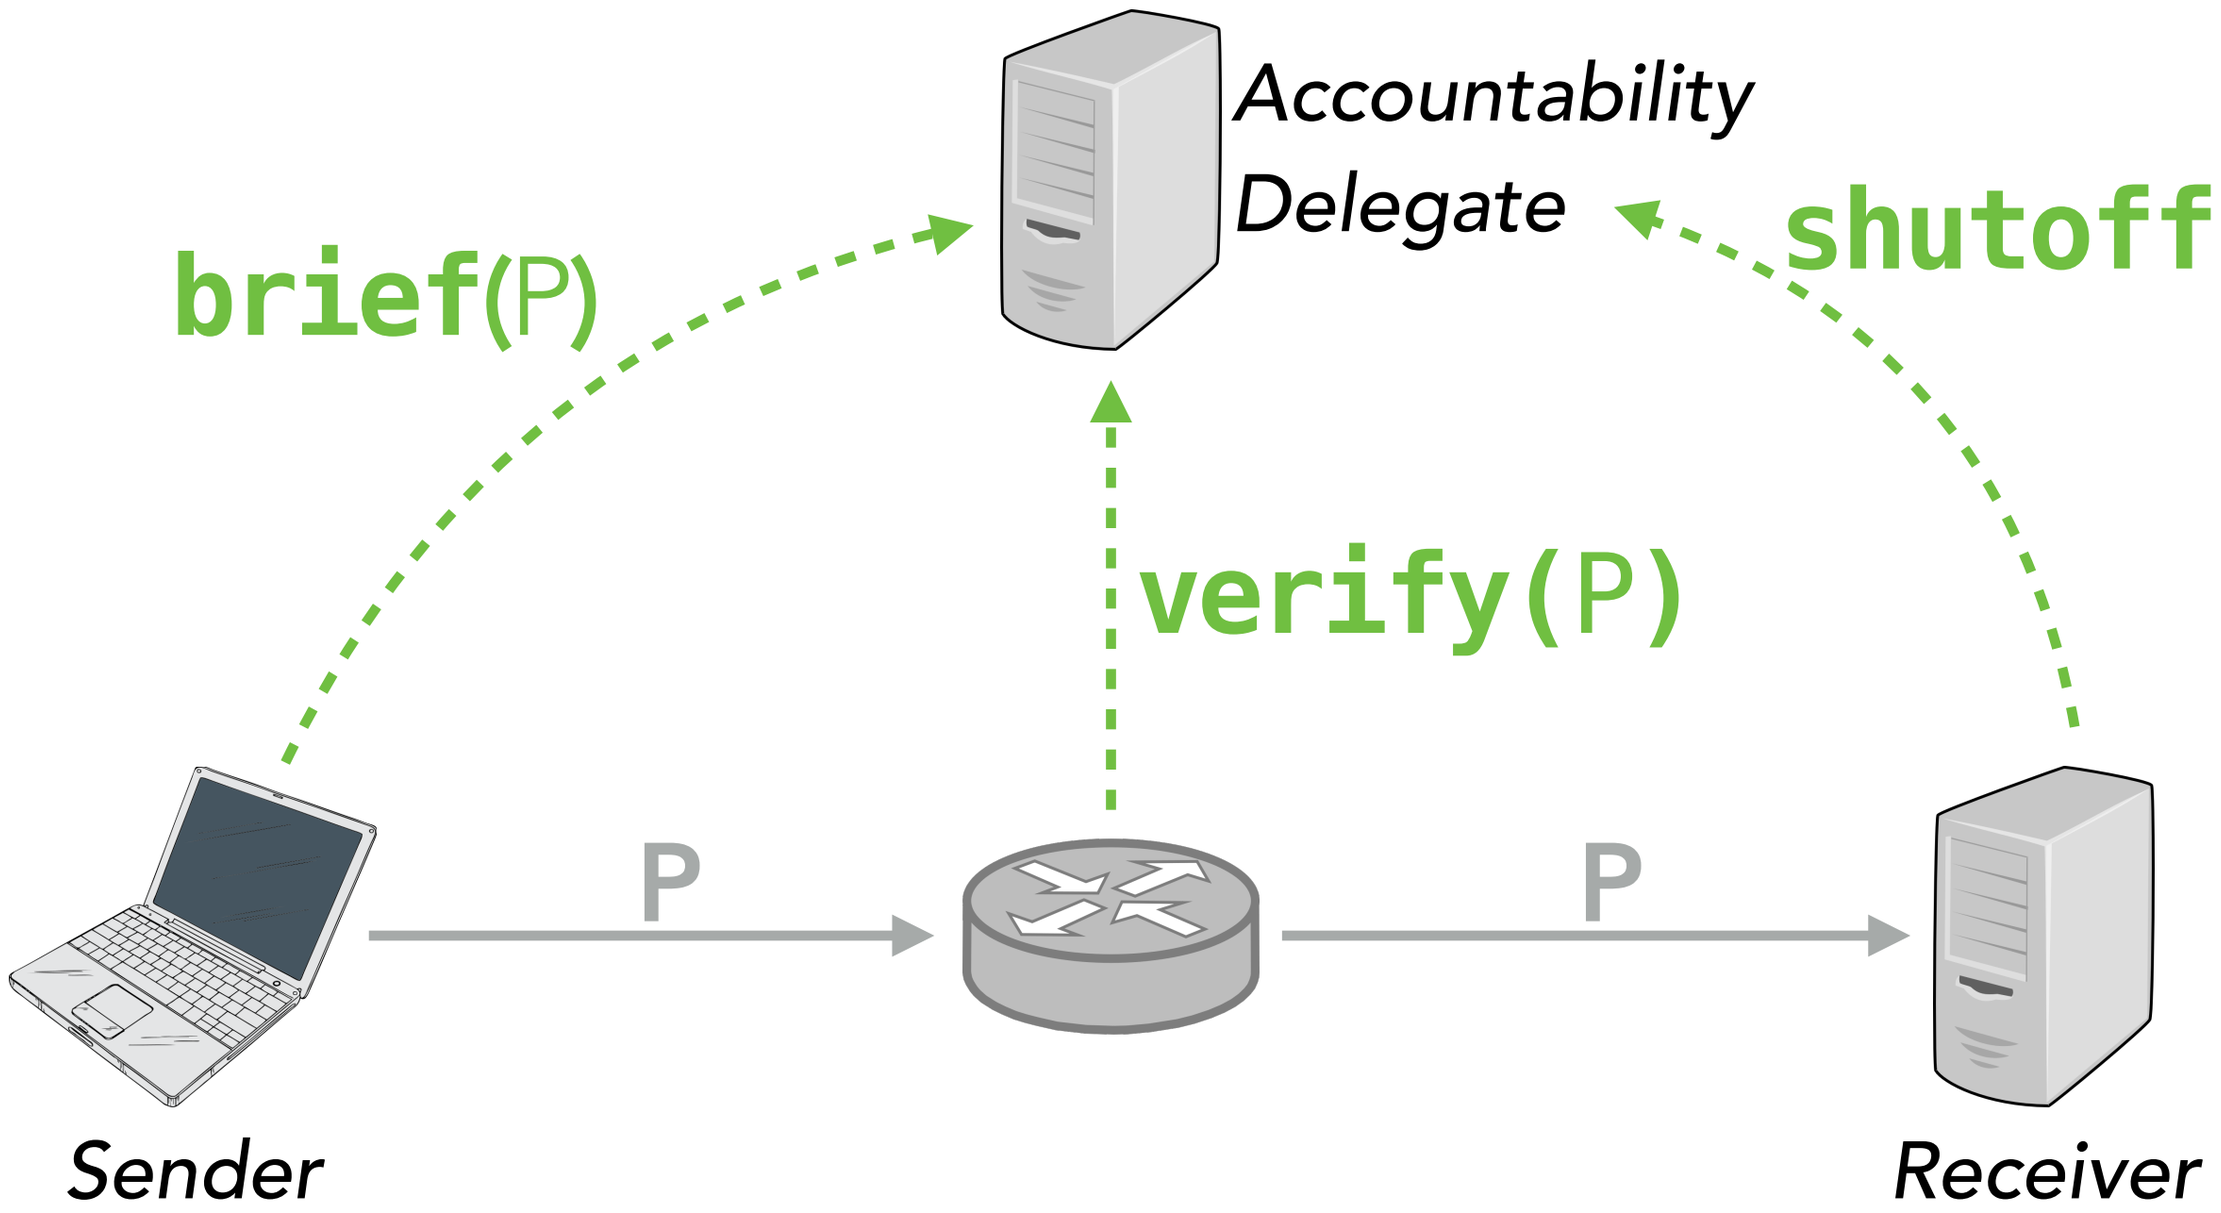

IS THIS TECHNICALLY FEASIBLE?

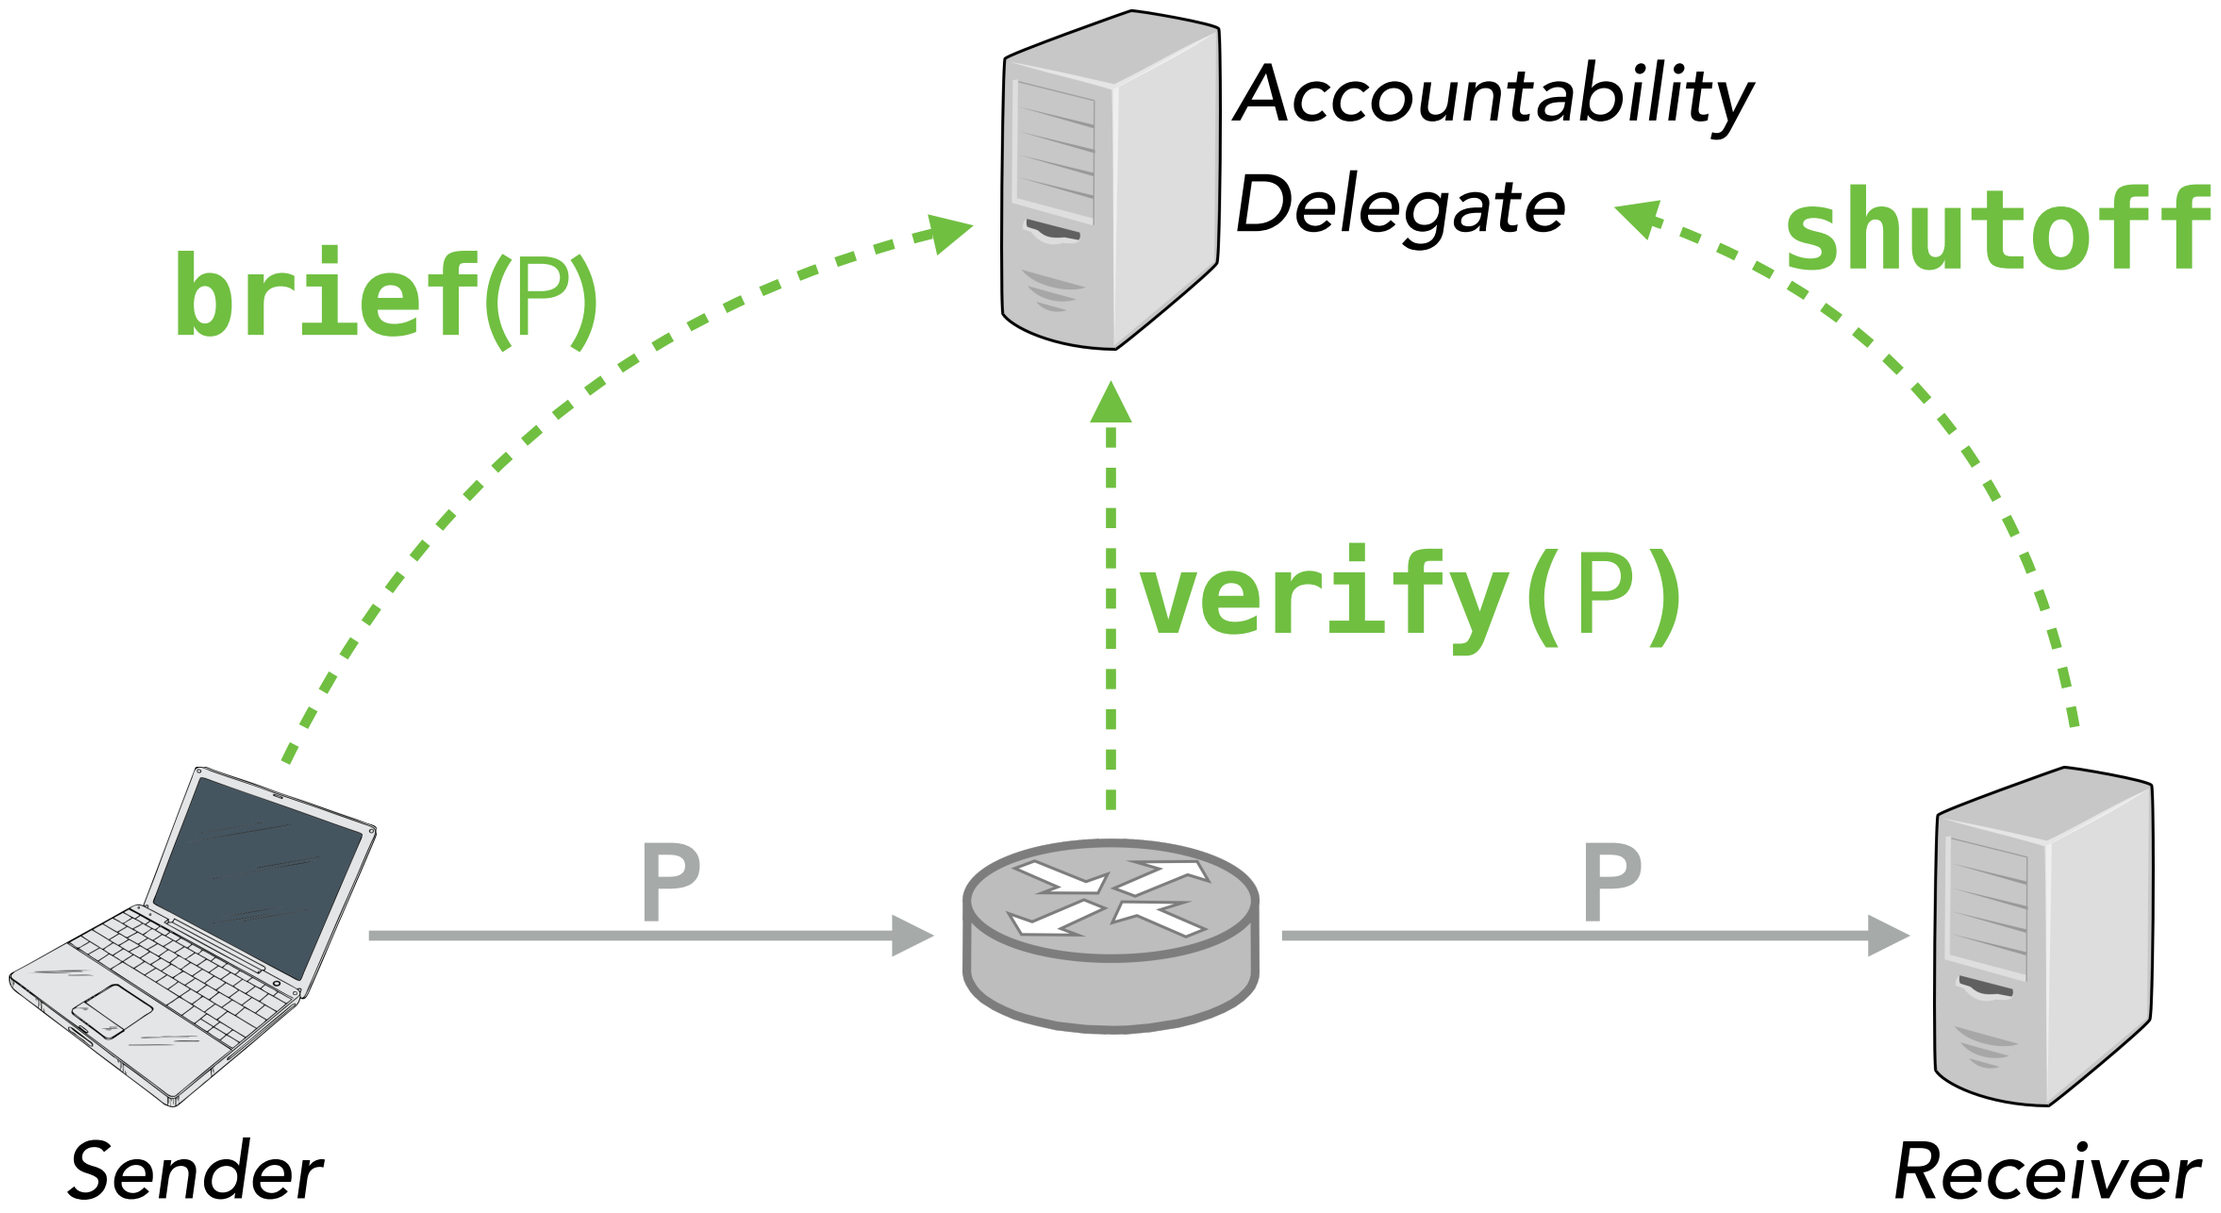

IS THIS TECHNICALLY FEASIBLE?

brief(P)

Storage Overhead
fingerprints at delegate

< 1GB

Network Overhead
sending fingerprints

0.5%

IS THIS TECHNICALLY FEASIBLE?

verify(P)

Computational Overhead
at delegate

78K
verifies per sec

Storage Overhead
verified flow list at router

94MB

CuckooFilter: [Zhou et al., CoNEXT 2013]

ed25519: [Bernstein et al., 2012]

FLOW GRANULARITY

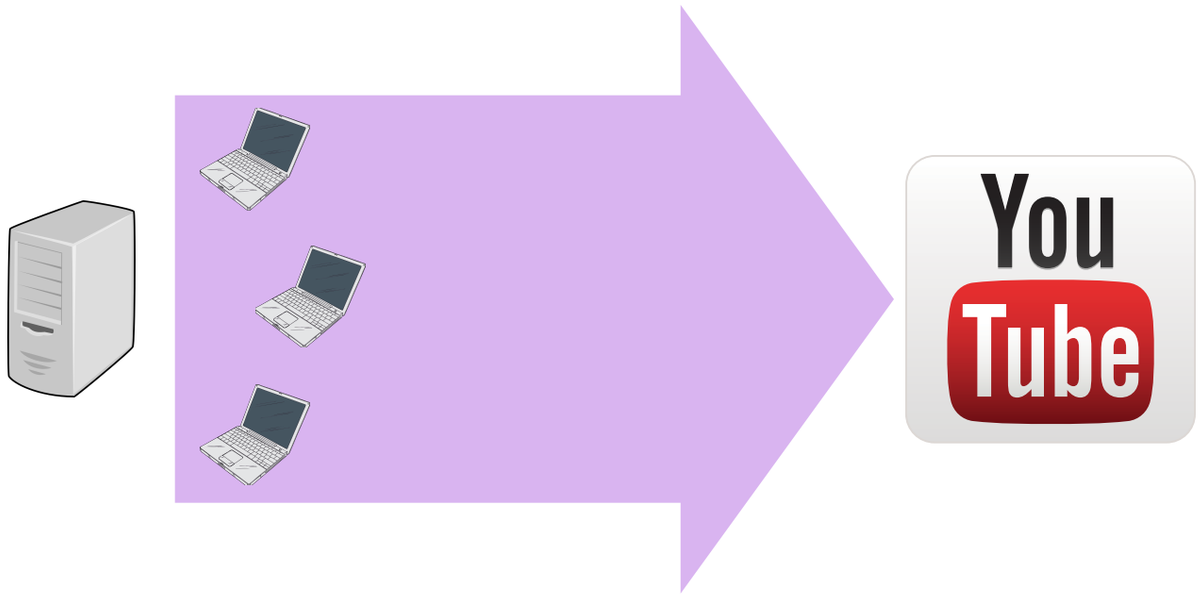

One flow ID for all clients

GRANULARITY: DELEGATE ↔ DESTINATION

Large Anonymity Set

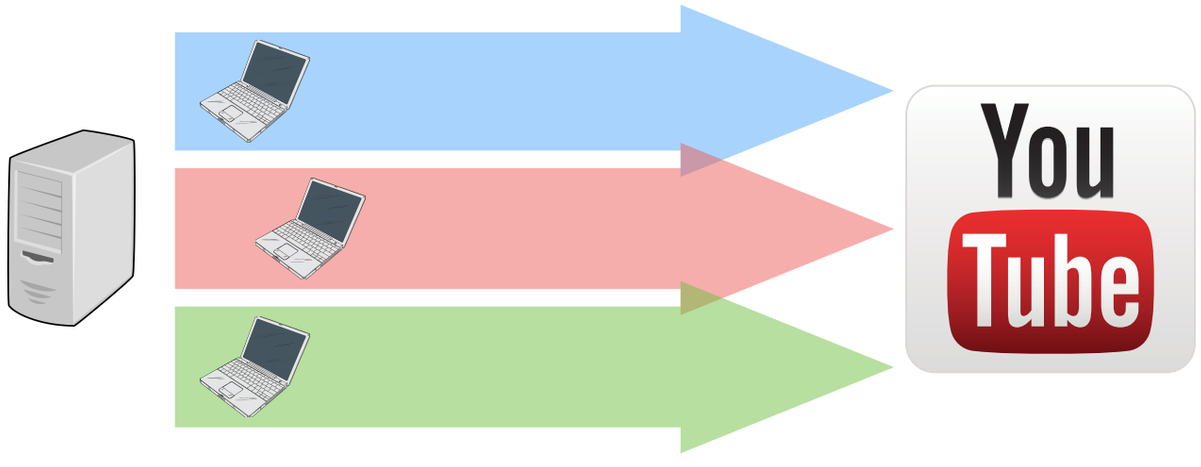

One flow ID per connection

GRANULARITY: TCP FLOW

No Collateral Damage for Shutoff

ASSIGNING FLOW IDS

DELEGATE

DELEGATE'S CLIENTS

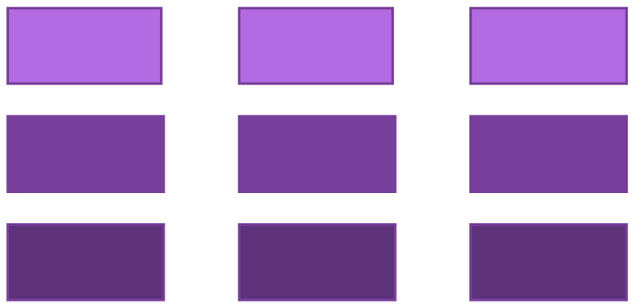

FLOW IDS

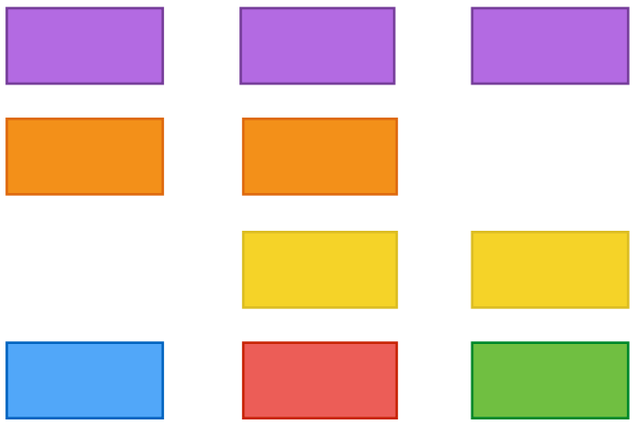

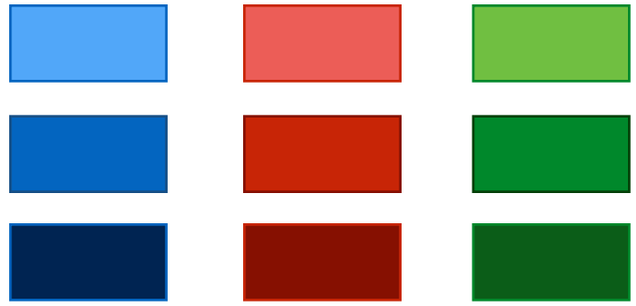

SHARED

Large Anonymity Set

VARIETY OF CLASSES

Flexible

UNIQUE

No Collateral Damage

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

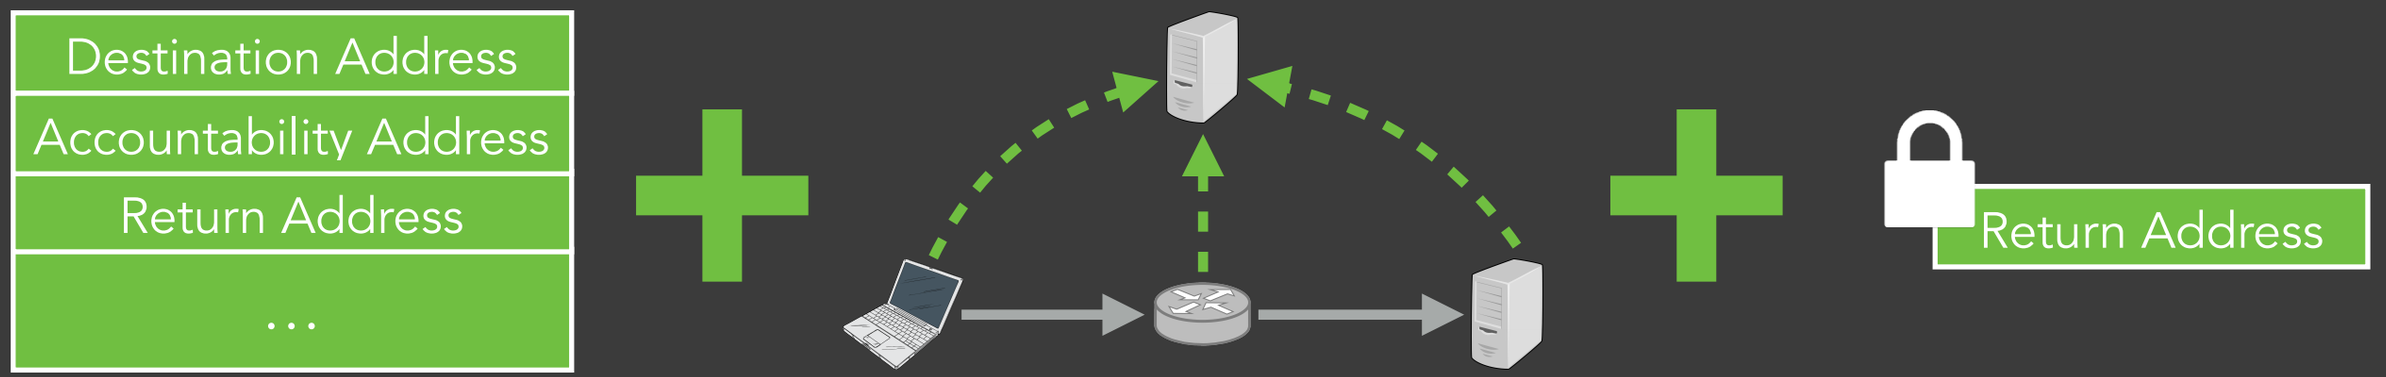

Separate Accountability and Return Addresses

Delegated Accountability

How it Works

Feasibility

Flow Granularity

Hidden Return Addresses

Real-World Deployment

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

Destination Address
Accountability Address
Return Address
...

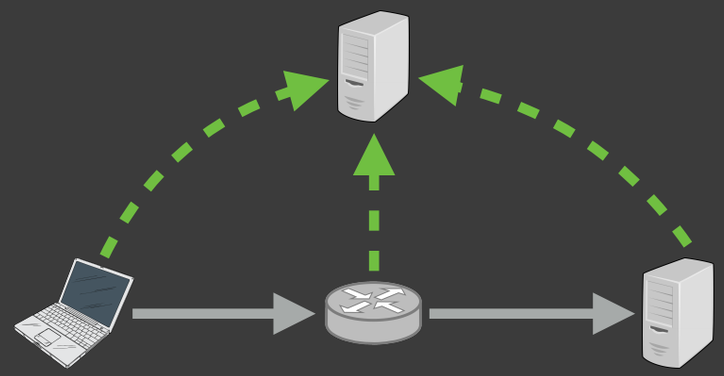

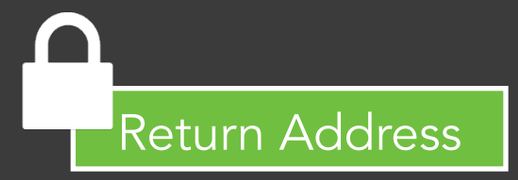

Separate Accountability and Return Addresses

Delegated Accountability

Hidden Return Addresses

Real-World Deployment

HIDING RETURN ADDRESSES

1 END-TO-END ENCRYPTION

Protection From:

- *Source Domain*
- ✓ *Local Observers*
- ✓ *Transit Networks*
- *Receiver*

2 ADDRESS TRANSLATION

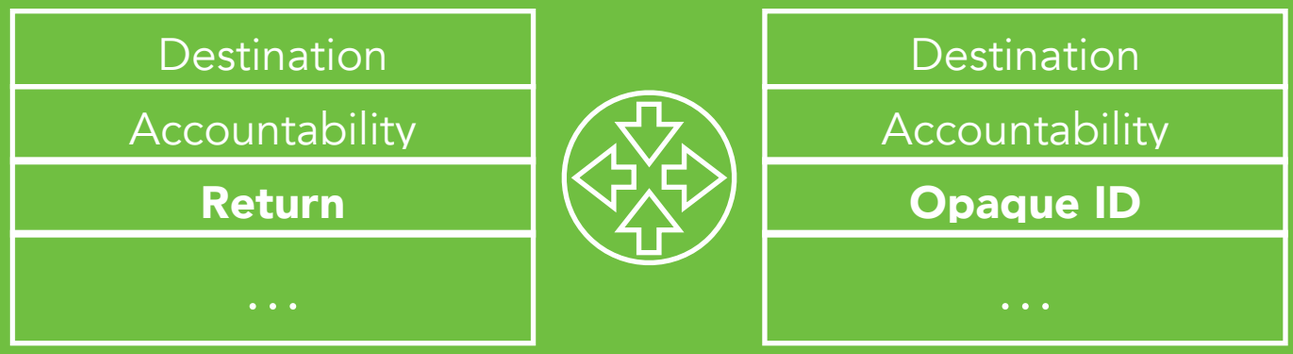

Protection From:

- *Source Domain*
- *Local Observers*
- ✓ *Transit Networks*
- ✓ *Receiver*

Stateless and secure: [Raghavan 2009]

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

Destination Address
Accountability Address
Return Address
...

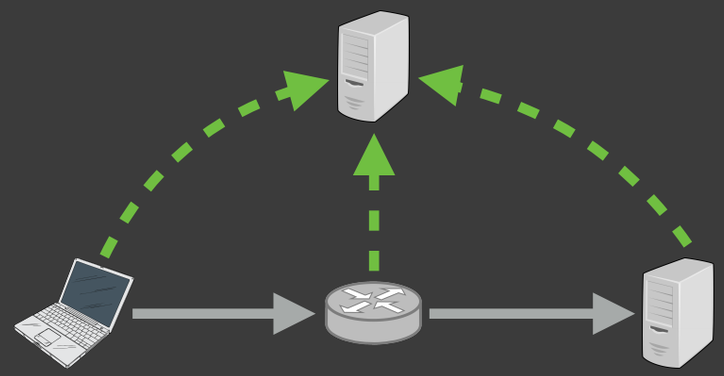

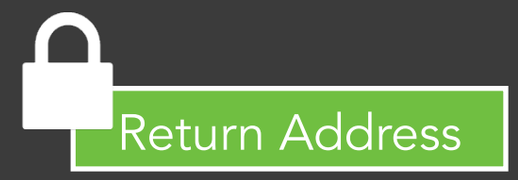

Separate Accountability and Return Addresses

Delegated Accountability

Hidden Return Addresses

Real-World Deployment

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

Destination Address
Accountability Address
Return Address
...

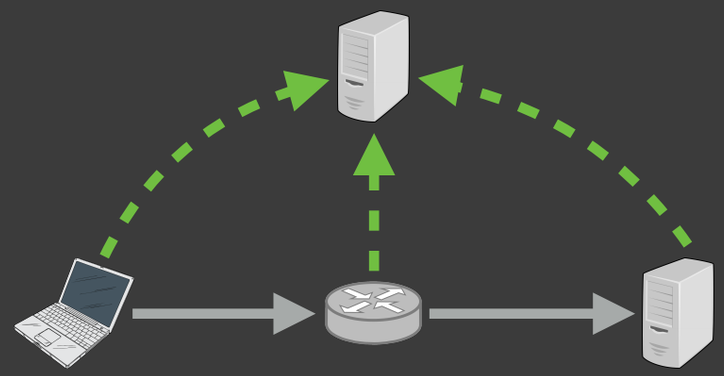

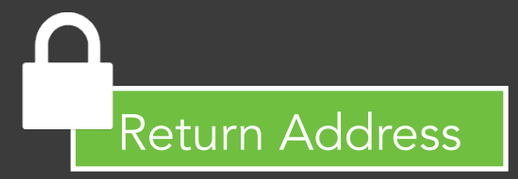

Separate Accountability and Return Addresses

Delegated Accountability

Hidden Return Addresses

Real-World Deployment

EXAMPLE DEPLOYMENTS

Specialized Companies as Delegates

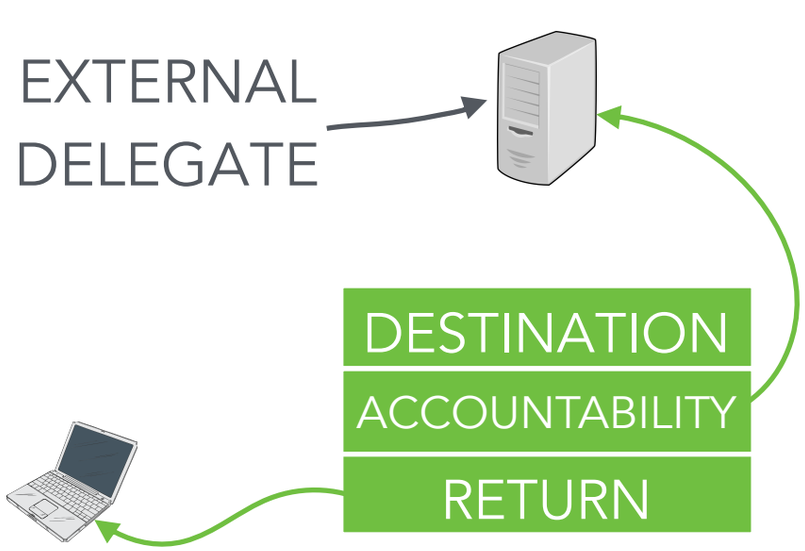

Source Domains as Delegates

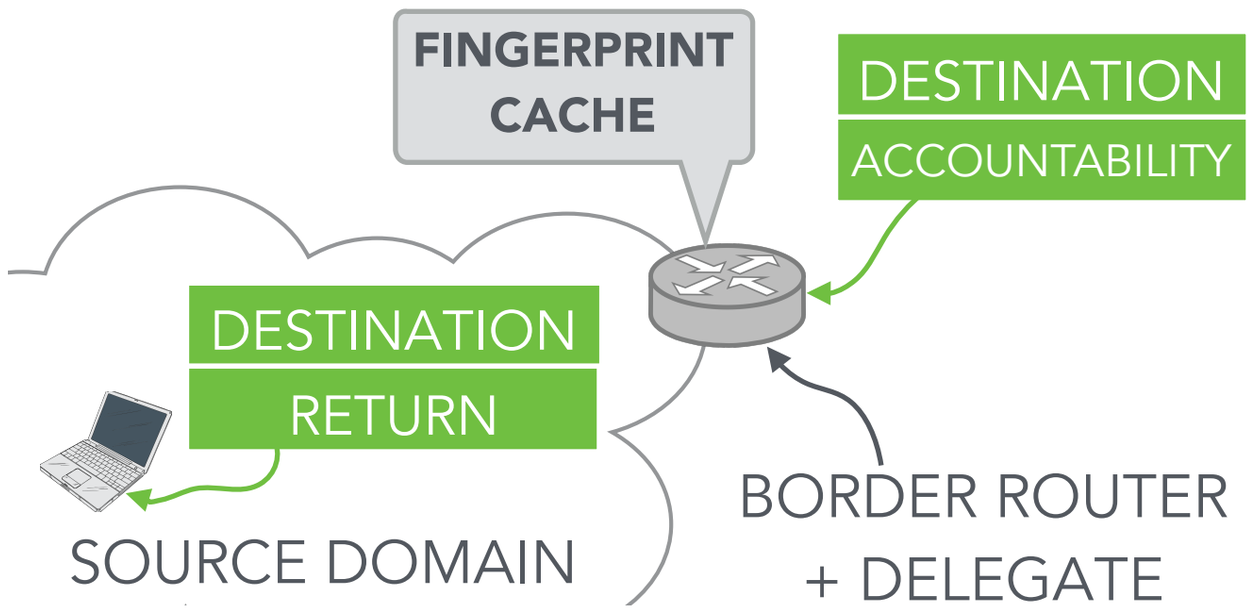

No burden on source domains

Larger anonymity set

No briefing overhead

Lower verification latency

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

Destination Address
Accountability Address
Return Address
...

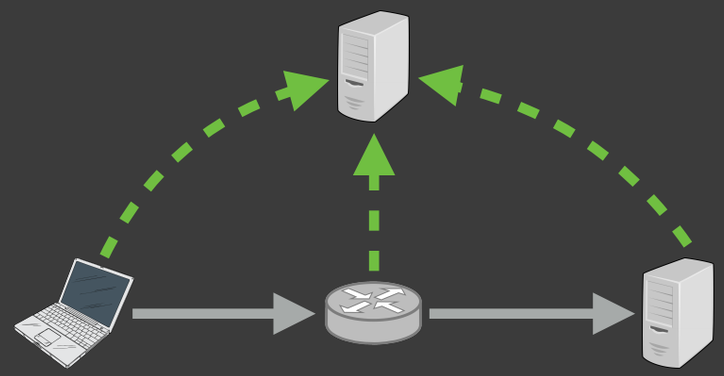

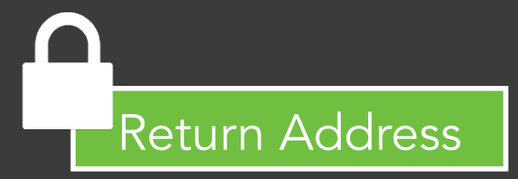

Separate Accountability and Return Addresses

Delegated Accountability

Hidden Return Addresses

Real-World Deployment

APIP:

ACCOUNTABLE AND PRIVATE INTERNET PROTOCOL

Destination Address
Accountability Address
Return Address
…

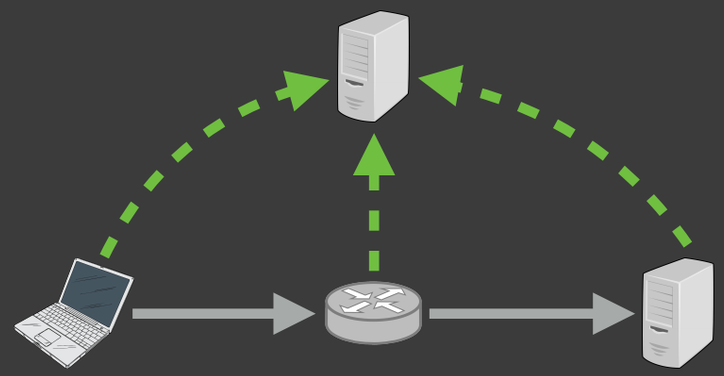

Return Address

Separate Accountability and Return Addresses

Delegated Accountability

Hidden Return Addresses

Real-World Deployment

IN THE PAPER

Source address roles

Who can be a delegate?

Anonymity set analysis

Attacking APIP

Trust/key management

Protocol details

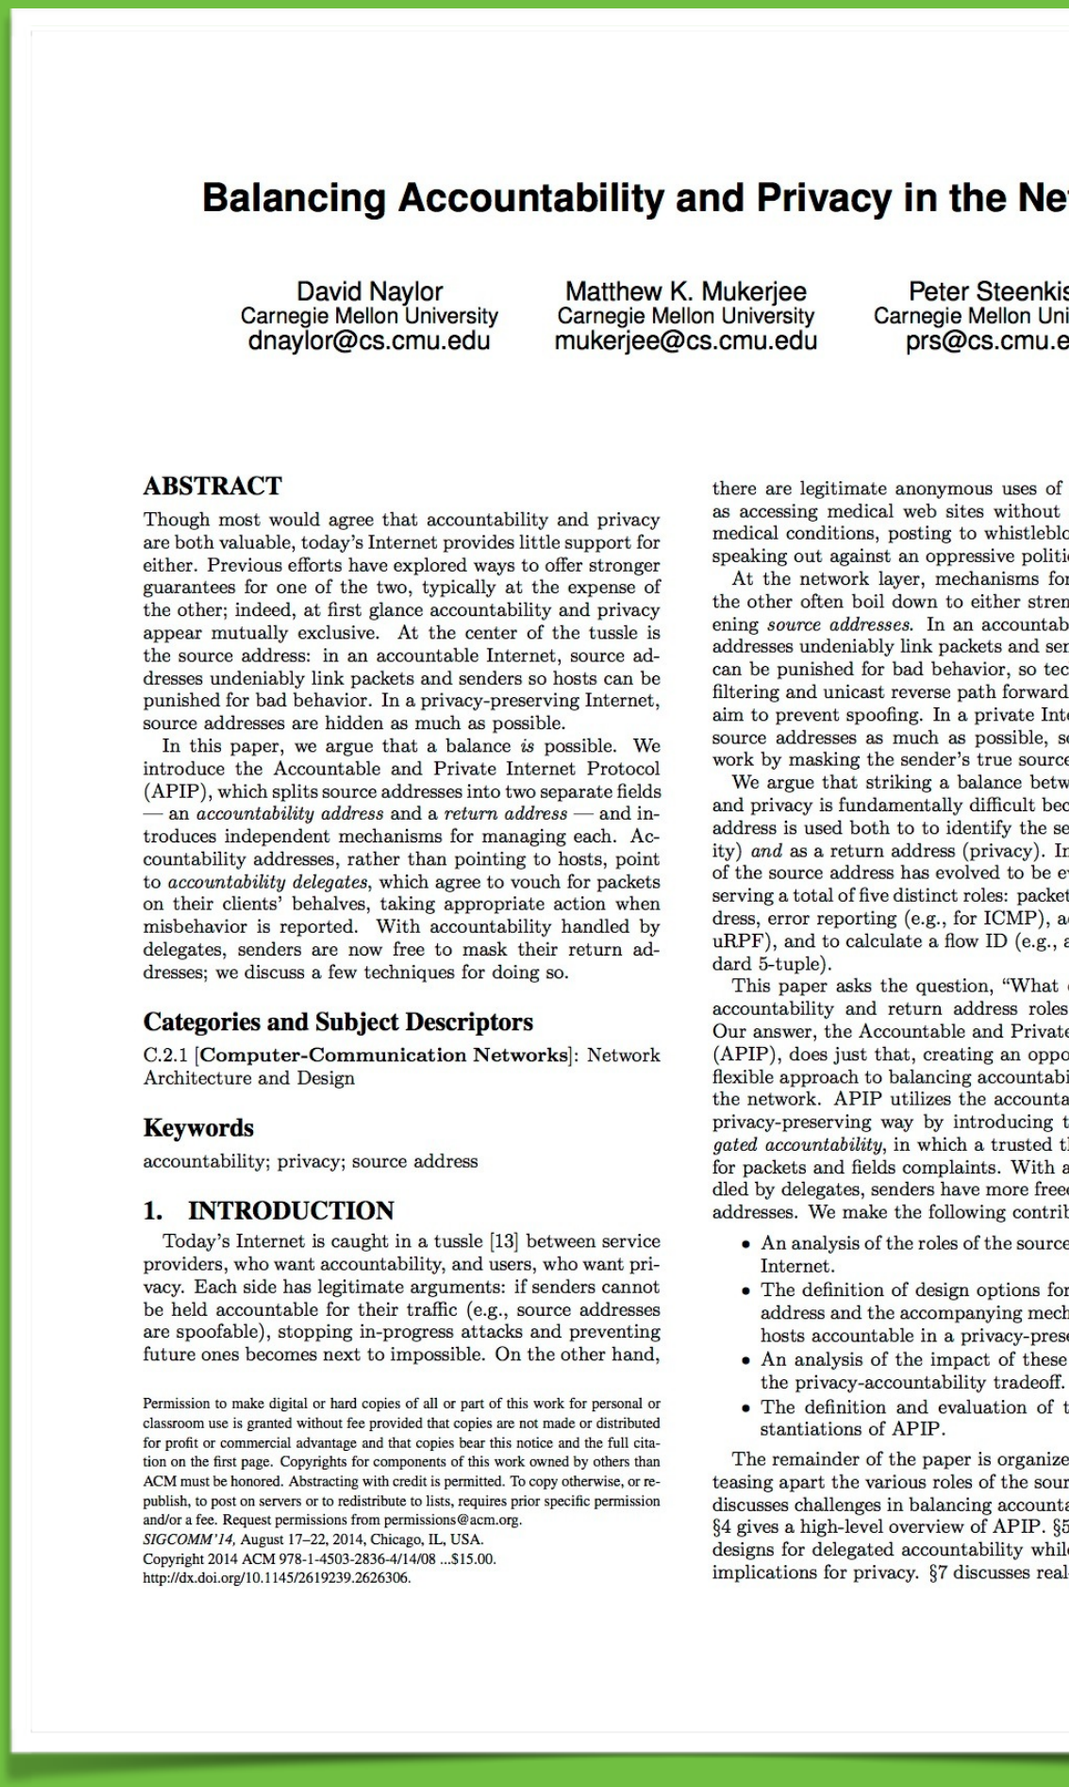

Balancing Accountability and Privacy in the Ne

David Naylor
Carnegie Mellon University
dnaylor@cs.cmu.edu

Matthew K. Mukerjee
Carnegie Mellon University
mukerjee@cs.cmu.edu

Peter Steenkis
Carnegie Mellon Uni
prs@cs.cmu.e

ABSTRACT

Though most would agree that accountability and privacy are both valuable, today's Internet provides little support for either. Previous efforts have explored ways to offer stronger guarantees for one of the two, typically at the expense of the other; indeed, at first glance accountability and privacy appear mutually exclusive. At the center of the tussle is the source address: in an accountable Internet, source addresses undeniably link packets and senders so hosts can be punished for bad behavior. In a privacy-preserving Internet, source addresses are hidden as much as possible.

In this paper, we argue that a balance *is* possible. We introduce the Accountable and Private Internet Protocol (APIP), which splits source addresses into two separate fields — an *accountability address* and a *return address* — and introduces independent mechanisms for managing each. Accountability addresses, rather than pointing to hosts, point to *accountability delegates*, which agree to vouch for packets on their clients' behalves, taking appropriate action when misbehavior is reported. With accountability handled by delegates, senders are now free to mask their return addresses; we discuss a few techniques for doing so.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

Keywords

accountability; privacy; source address

1. INTRODUCTION

Today's Internet is caught in a tussle [13] between service providers, who want accountability, and users, who want privacy. Each side has legitimate arguments: if senders cannot be held accountable for their traffic (e.g., source addresses are spoofable), stopping in-progress attacks and preventing future ones becomes next to impossible. On the other hand,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGCOMM'14, August 17–22, 2014, Chicago, IL, USA.
Copyright 2014 ACM 978-1-4503-2836-4/14/08 ...$15.00.
http://dx.doi.org/10.1145/2619239.2626306.

there are legitimate anonymous uses of
as accessing medical web sites without
medical conditions, posting to whistleblo
speaking out against an oppressive politi

At the network layer, mechanisms for
the other often boil down to either stren
ening *source addresses.* In an accountab
addresses undeniably link packets and se
can be punished for bad behavior, so tec
filtering and unicast reverse path forward
aim to prevent spoofing. In a private Int
source addresses as much as possible, s
work by masking the sender's true source

We argue that striking a balance betw
and privacy is fundamentally difficult bec
address is used both to to identify the se
ity) *and* as a return address (privacy). In
of the source address has evolved to be e
serving a total of five distinct roles: packet
dress, error reporting (e.g., for ICMP), a
uRPF), and to calculate a flow ID (e.g., a
dard 5-tuple).

This paper asks the question, "What
accountability and return address roles
Our answer, the Accountable and Private
(APIP), does just that, creating an oppo
flexible approach to balancing accountabi
the network. APIP utilizes the accounta
privacy-preserving way by introducing t
gated accountability, in which a trusted t
for packets and fields complaints. With a
dled by delegates, senders have more free
addresses. We make the following contrib

- An analysis of the roles of the source
  Internet.
- The definition of design options for
  address and the accompanying mech
  hosts accountable in a privacy-pres
- An analysis of the impact of these
  the privacy-accountability tradeoff.
- The definition and evaluation of t
  stantiations of APIP.

The remainder of the paper is organize
teasing apart the various roles of the sour
discusses challenges in balancing accounta
§4 gives a high-level overview of APIP. §5
designs for delegated accountability while
implications for privacy. §7 discusses real

ACCOUNTABILITY

unforgeable **source addresses**

PRIVACY

hidden **source addresses**

ACCOUNTABILITY

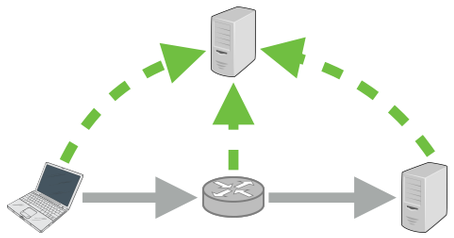

Delegated Accountability

every packet carries an accountability address

for reporting misbehavior

&

PRIVACY

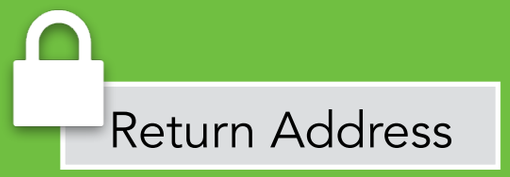

Hidden Return Addresses

return address can be hidden

since network just needs accountability address

Carnegie Mellon University

BALANCING ACCOUNTABILITY & PRIVACY IN THE NETWORK

David Naylor

Matt Mukerjee

Peter Steenkiste