iTAP: In-network Traffic Analysis Prevention using Software-Defined Networks

Roland Meier, David Gugelmann, Laurent Vanbever

https://itap.ethz.ch

ACM SOSR 2017. Santa Clara, CA, USA (April 2017).

ETH zürich

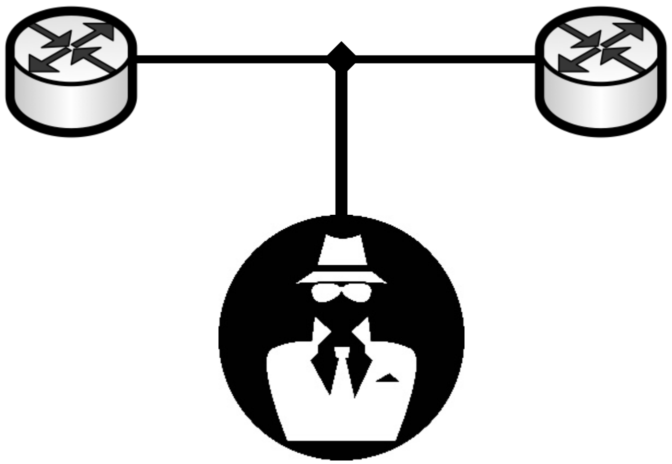

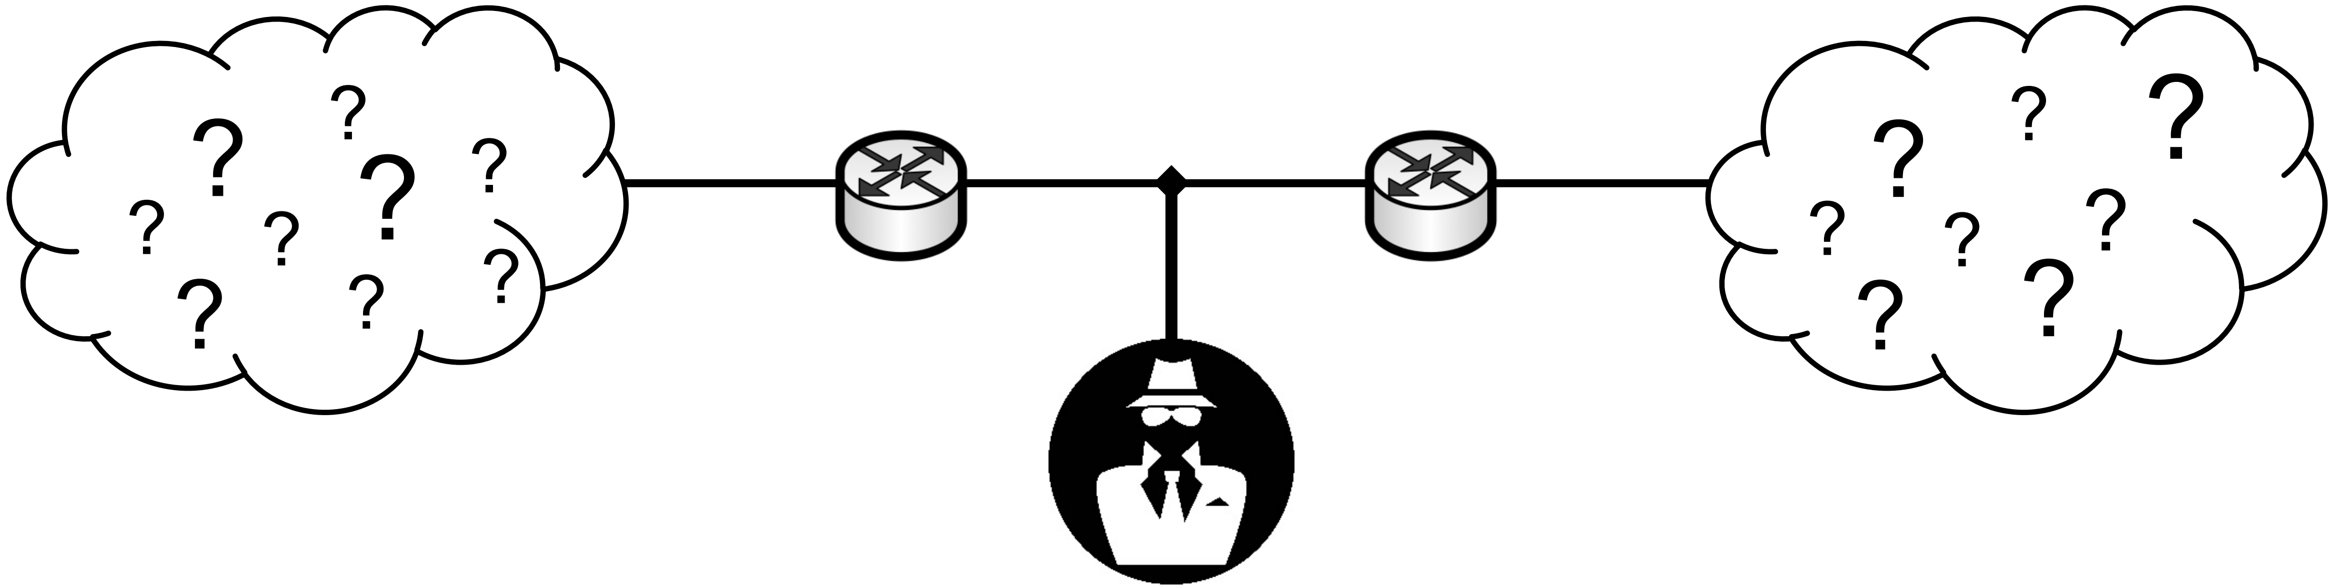

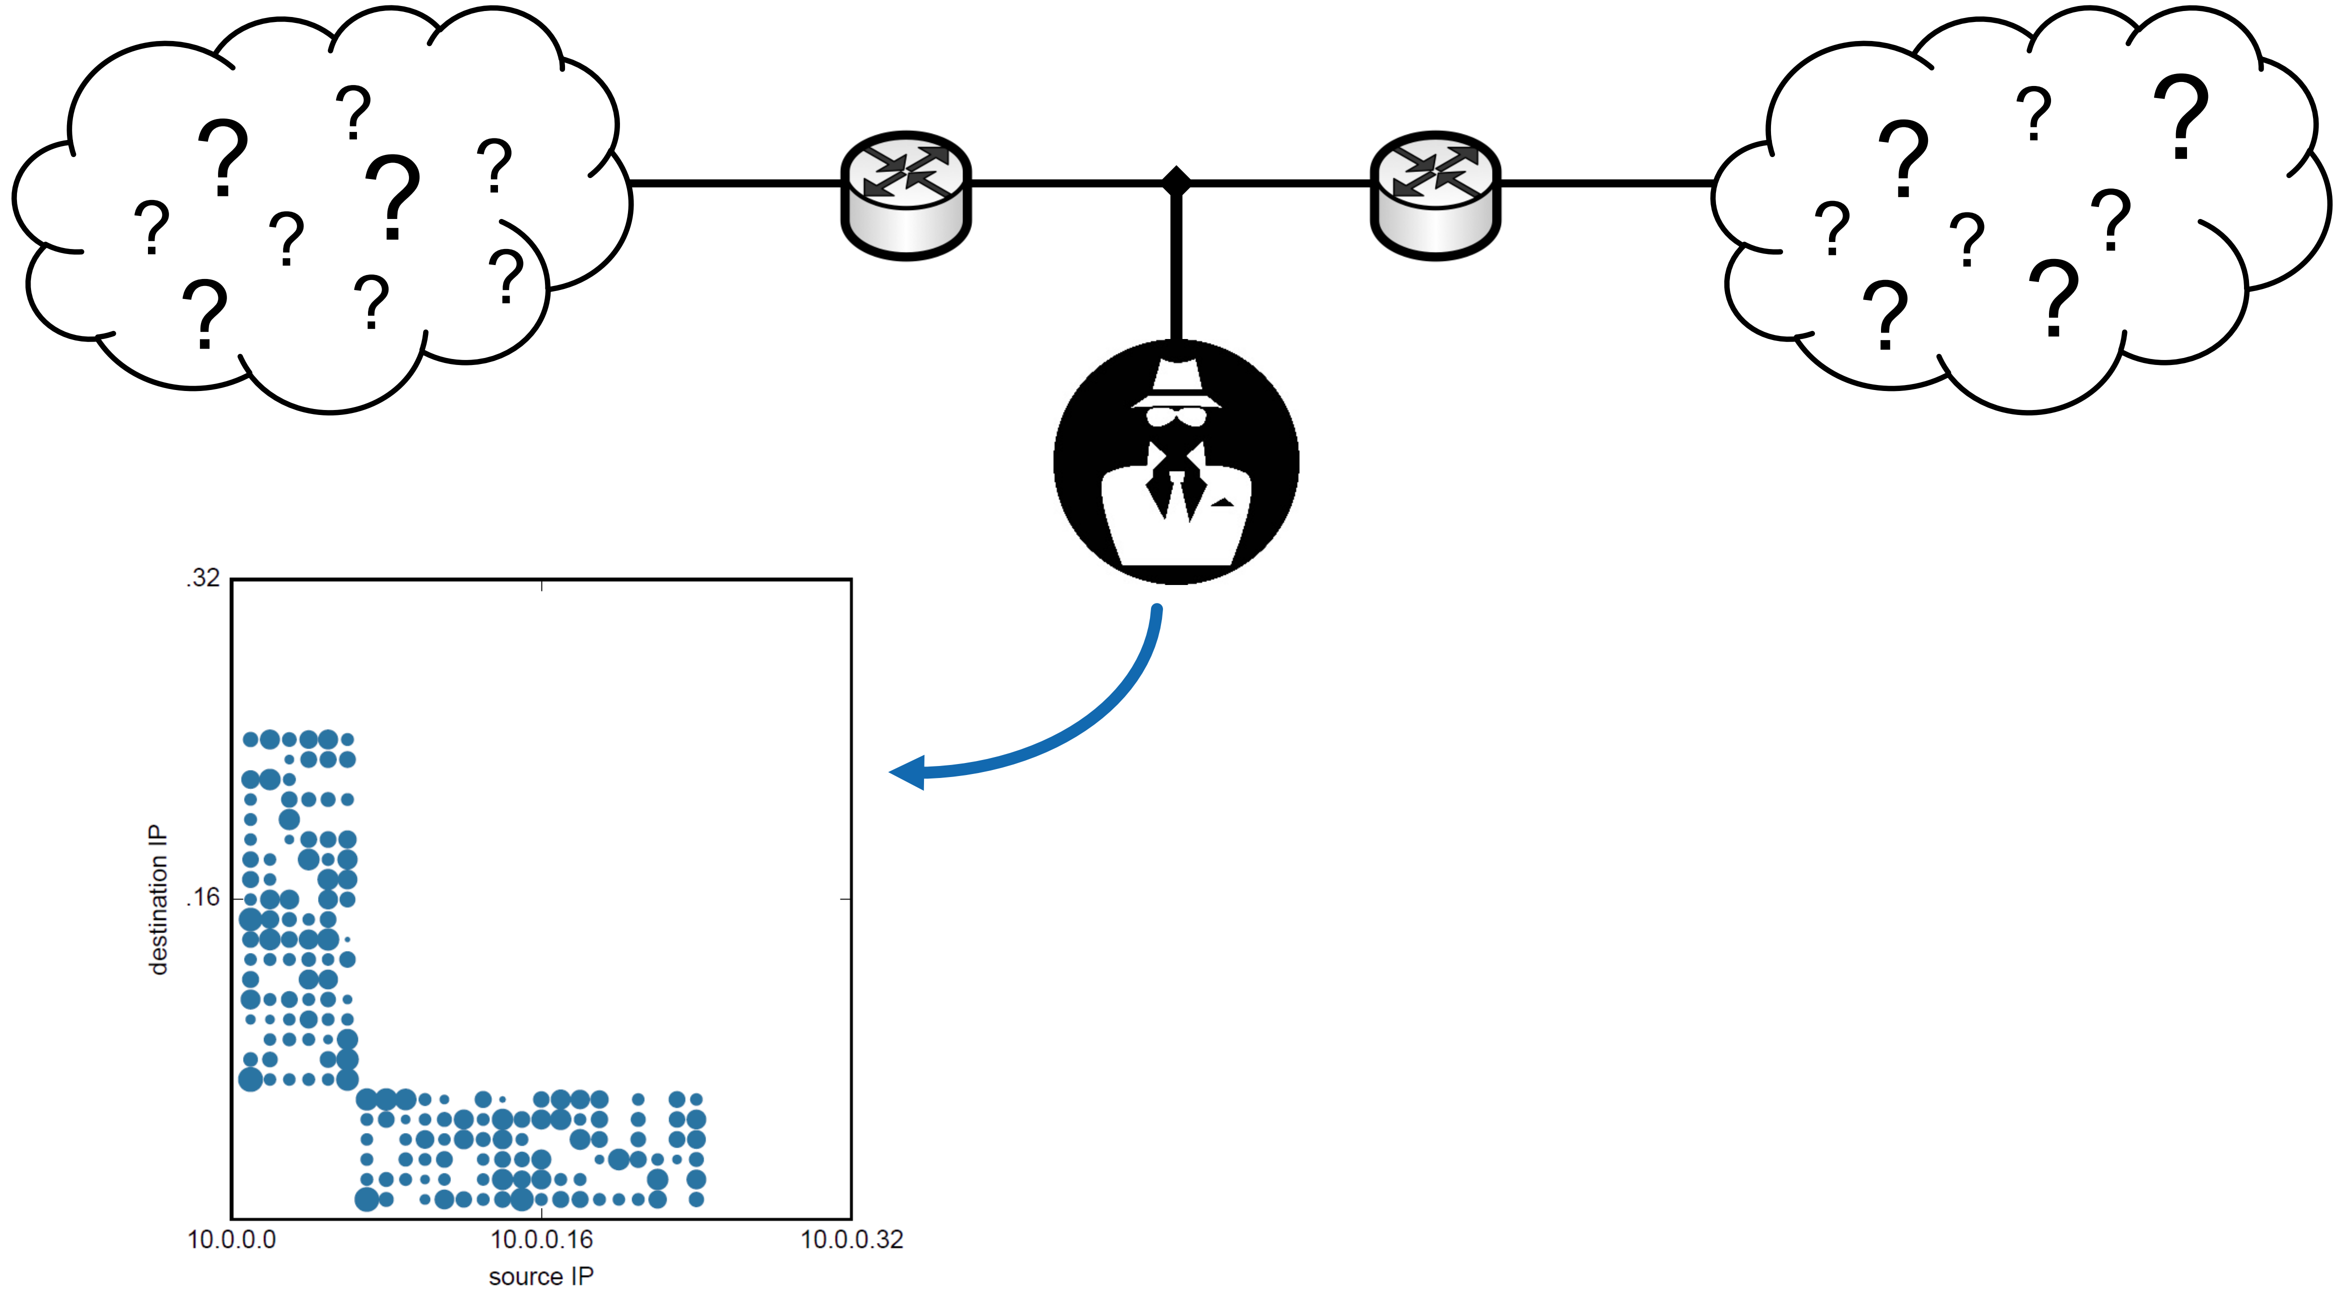

.32
.16
destination IP
10.0.0.0
10.0.0.16
10.0.0.32
source IP

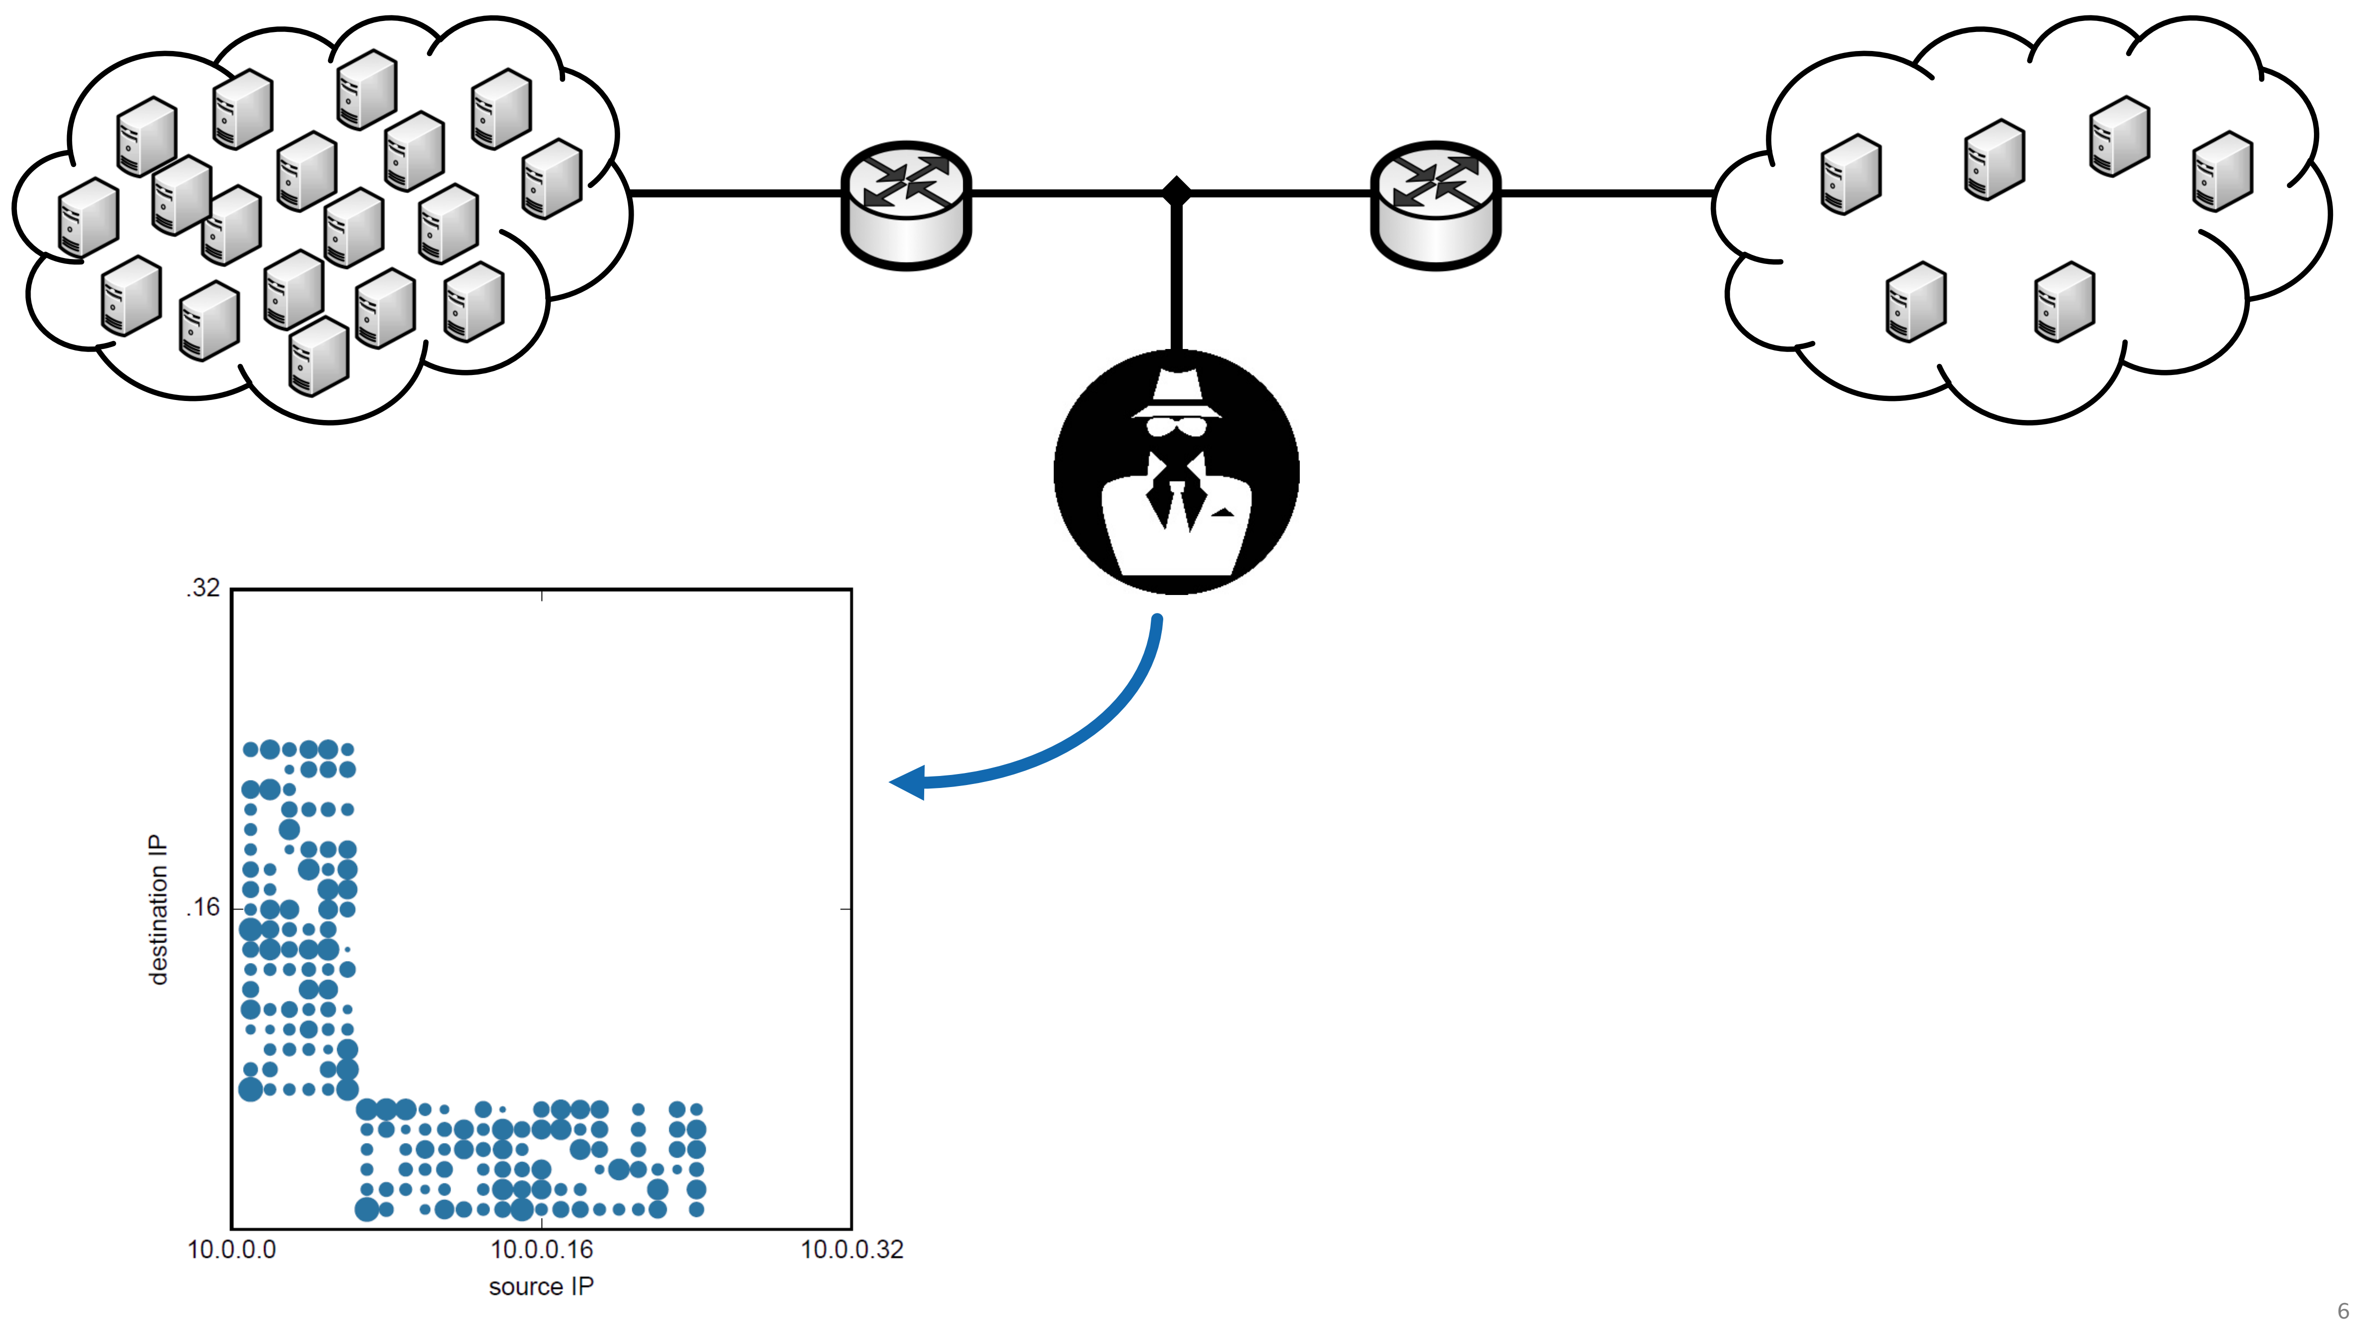

.32
.16
destination IP
10.0.0.0
10.0.0.16
10.0.0.32
source IP

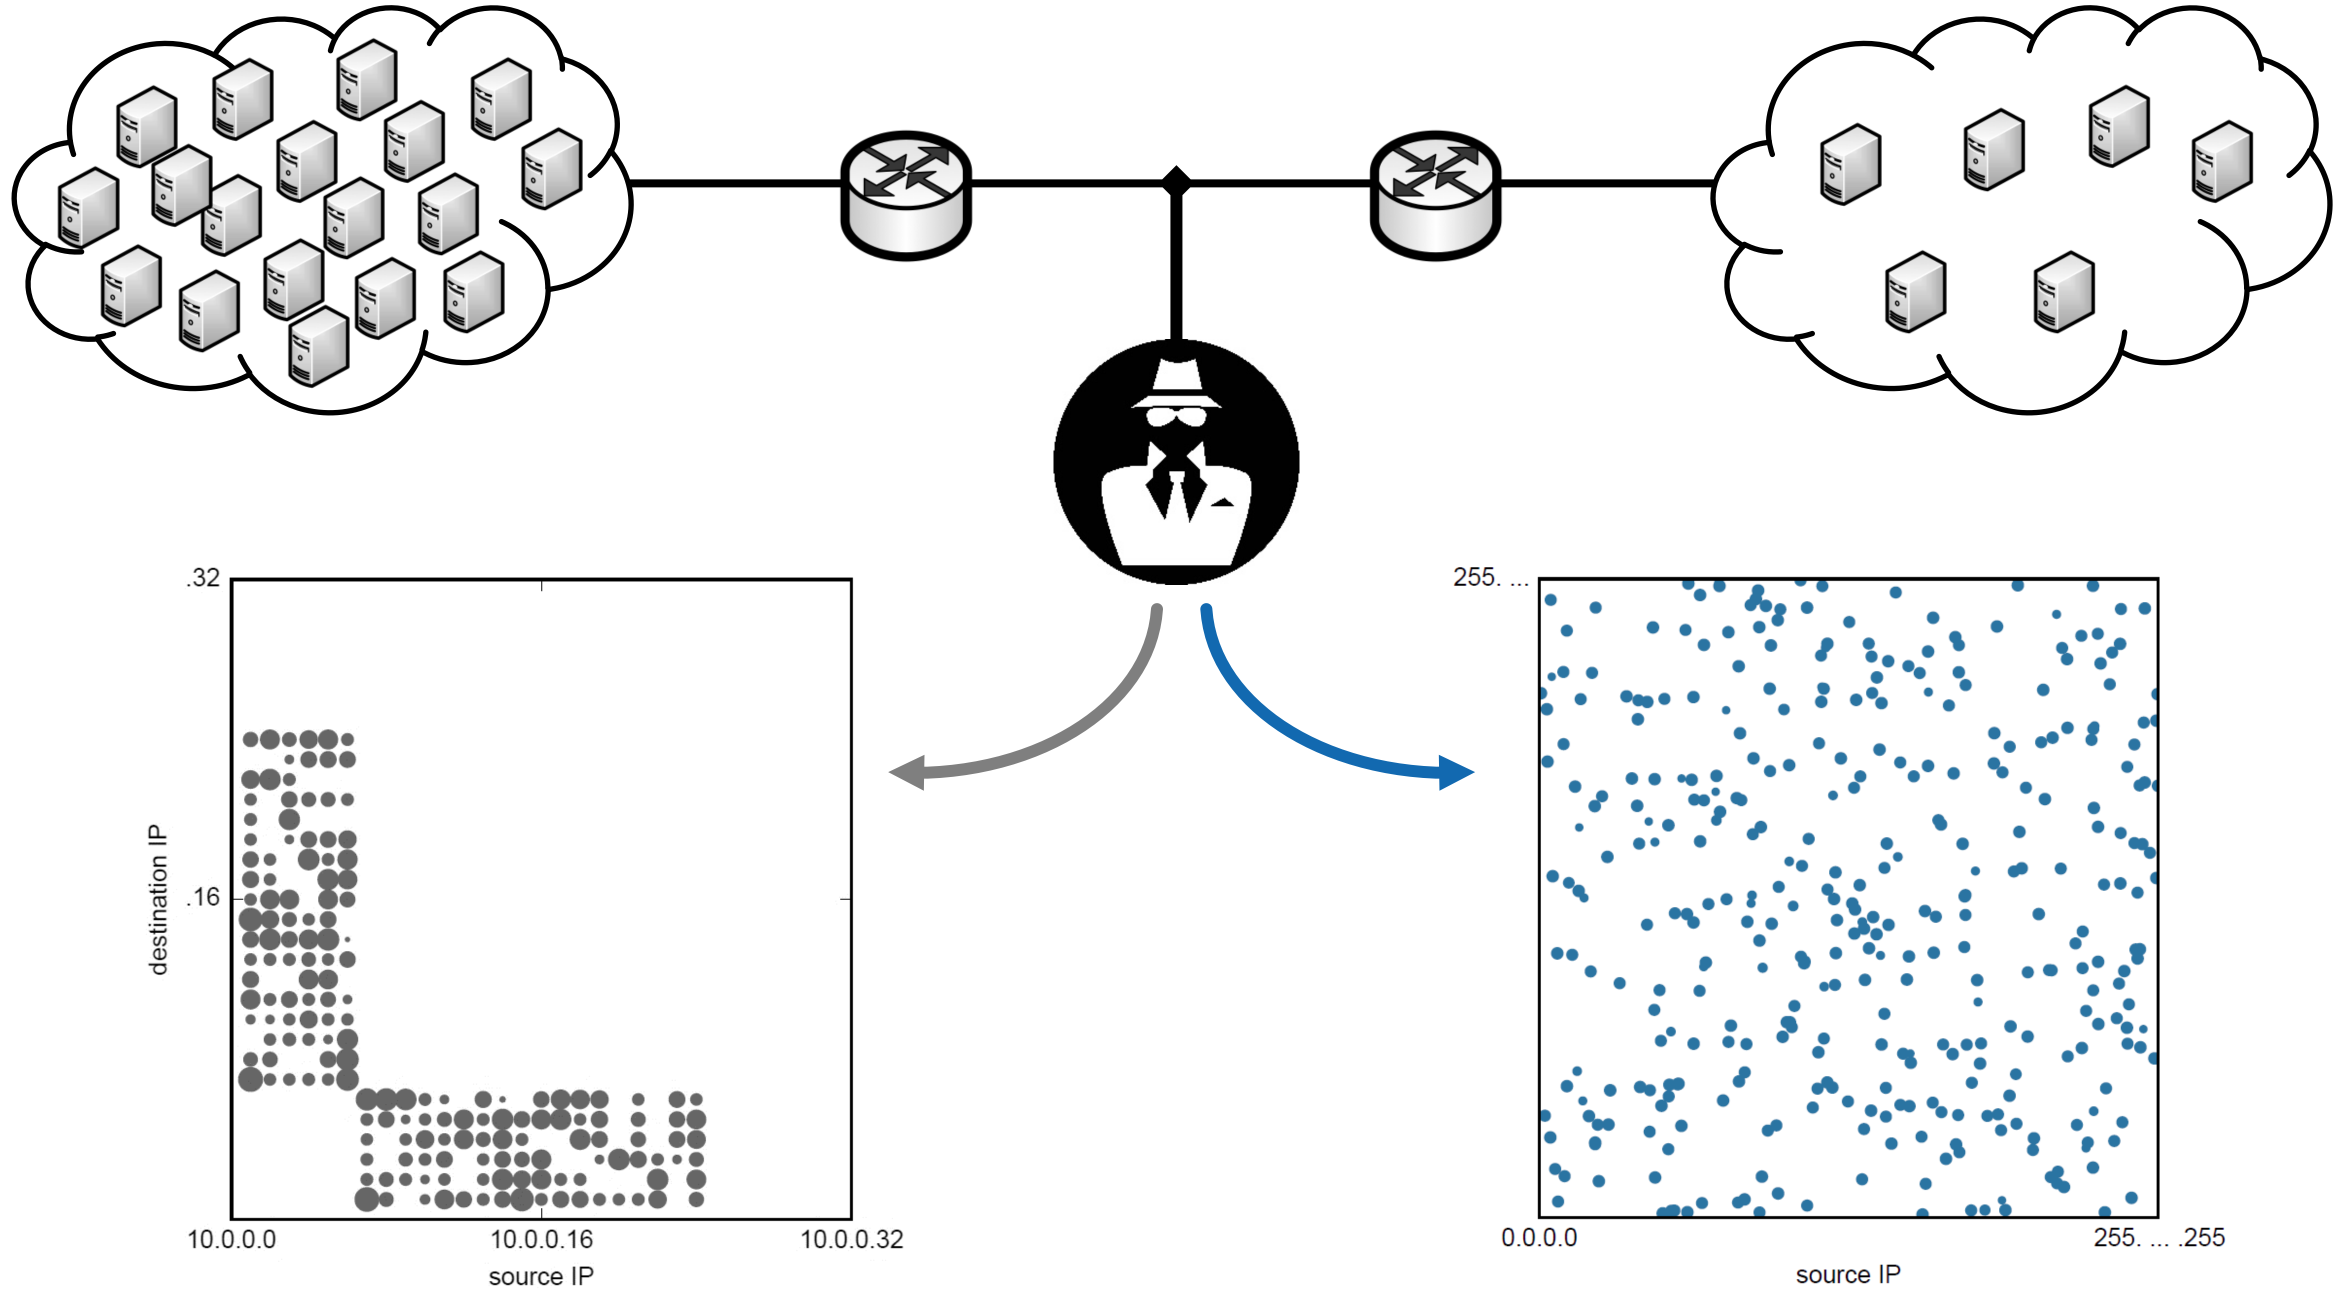

.32
.16
destination IP
10.0.0.0
10.0.0.16
10.0.0.32
source IP
255. ...
0.0.0.0
255. ... .255
source IP

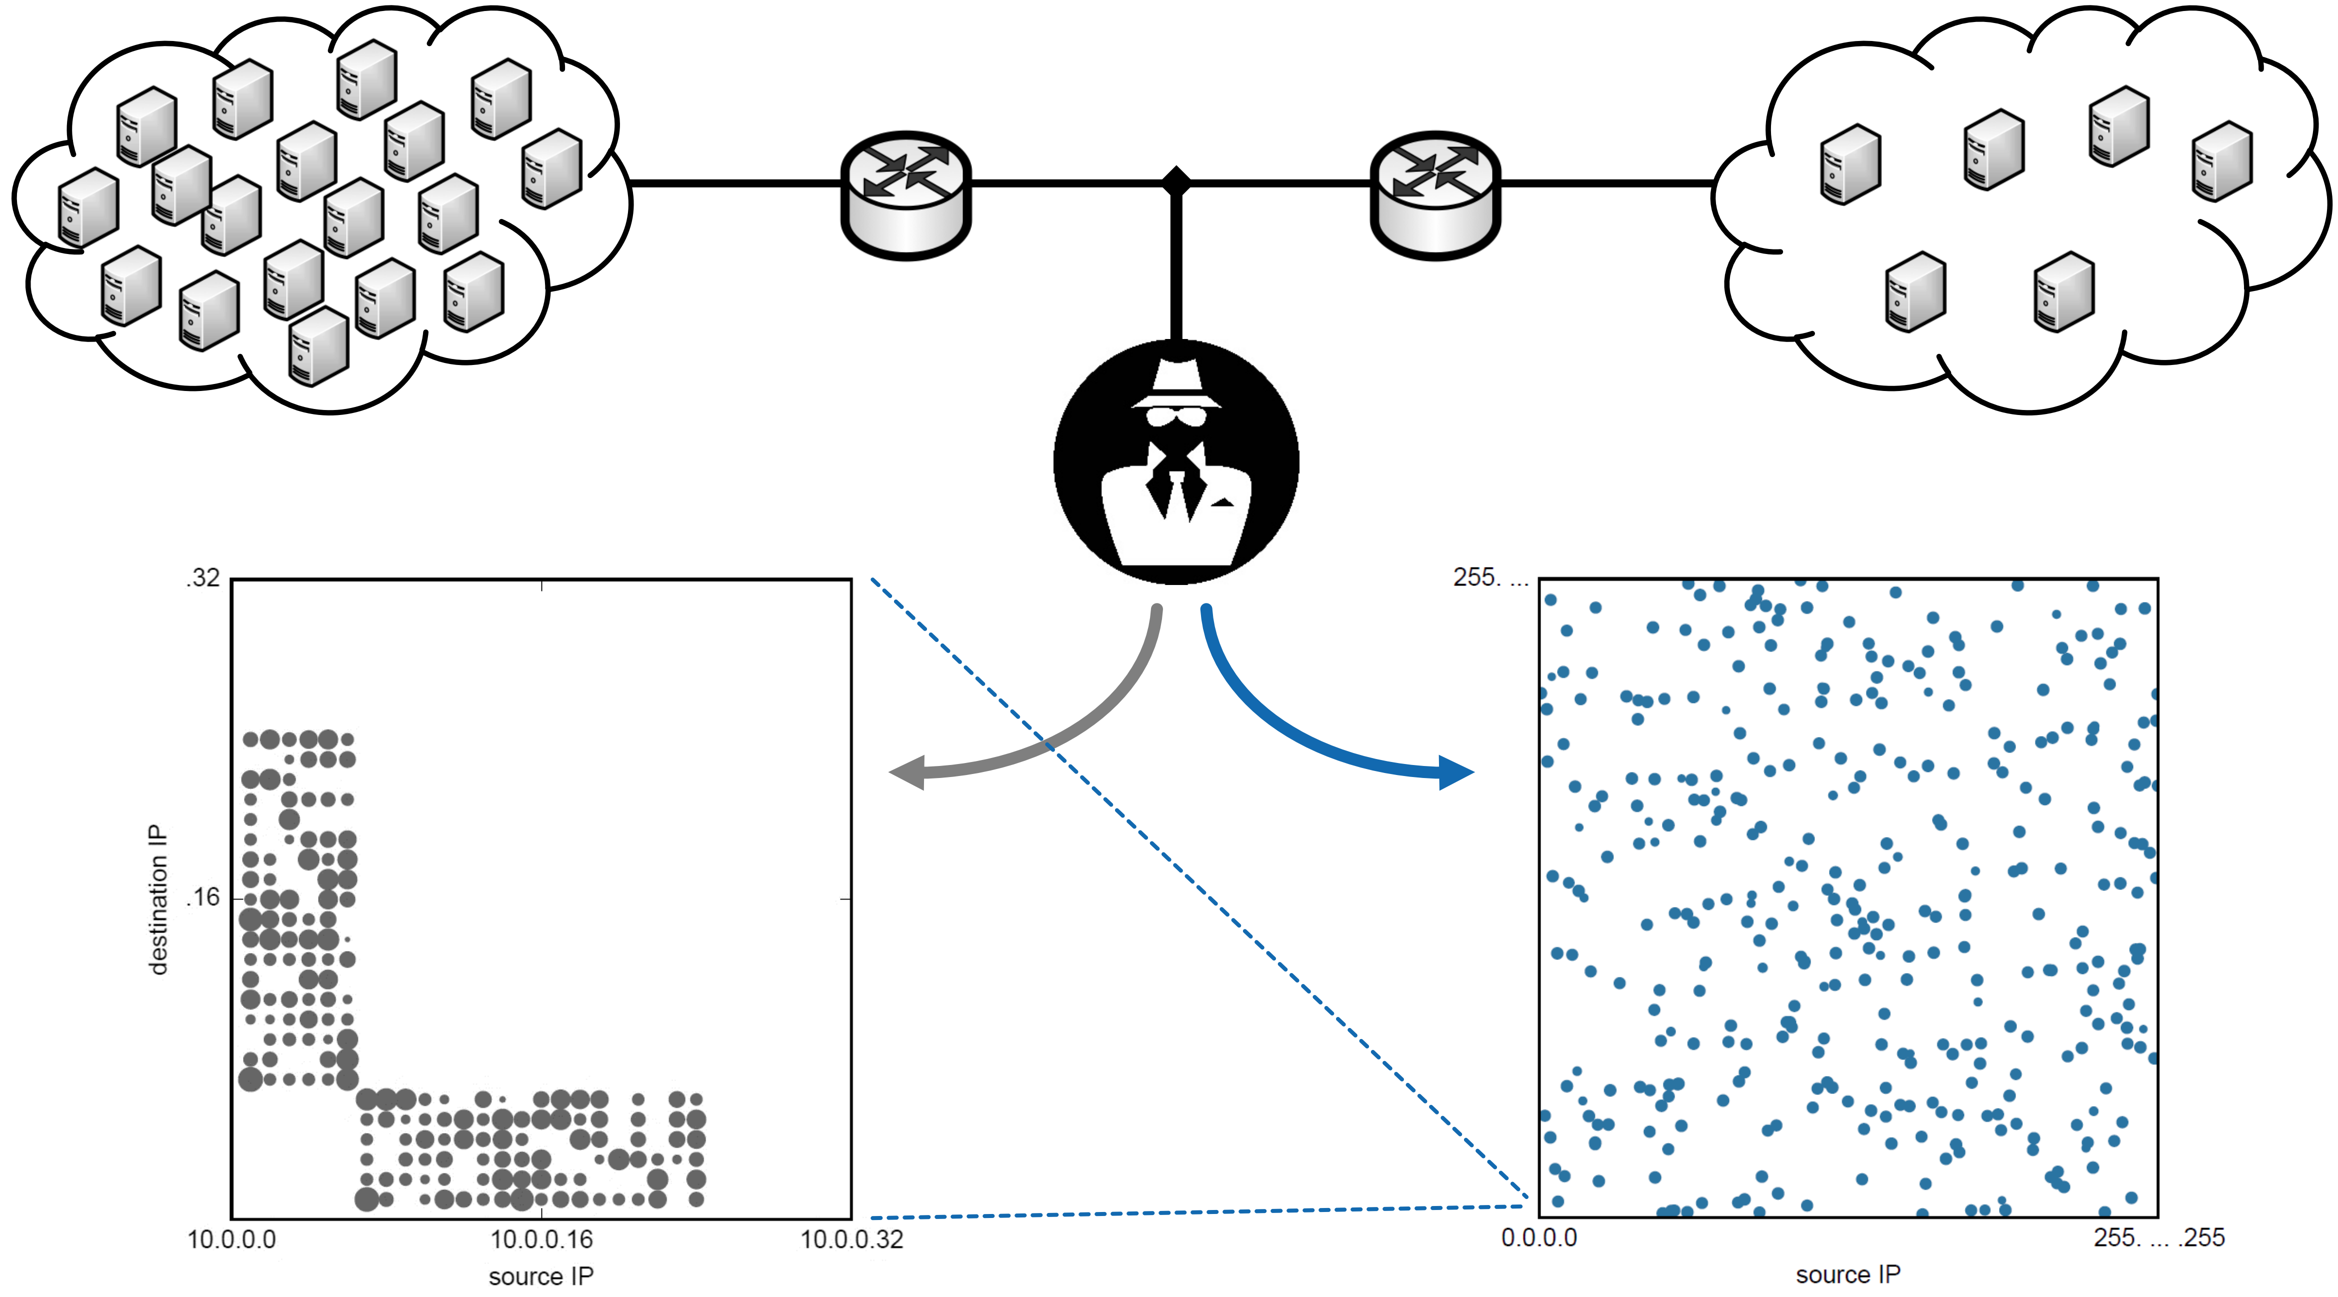
.32
.16
destination IP
10.0.0.0
10.0.0.16
10.0.0.32
source IP
255. ...
0.0.0.0
255. ... .255
source IP

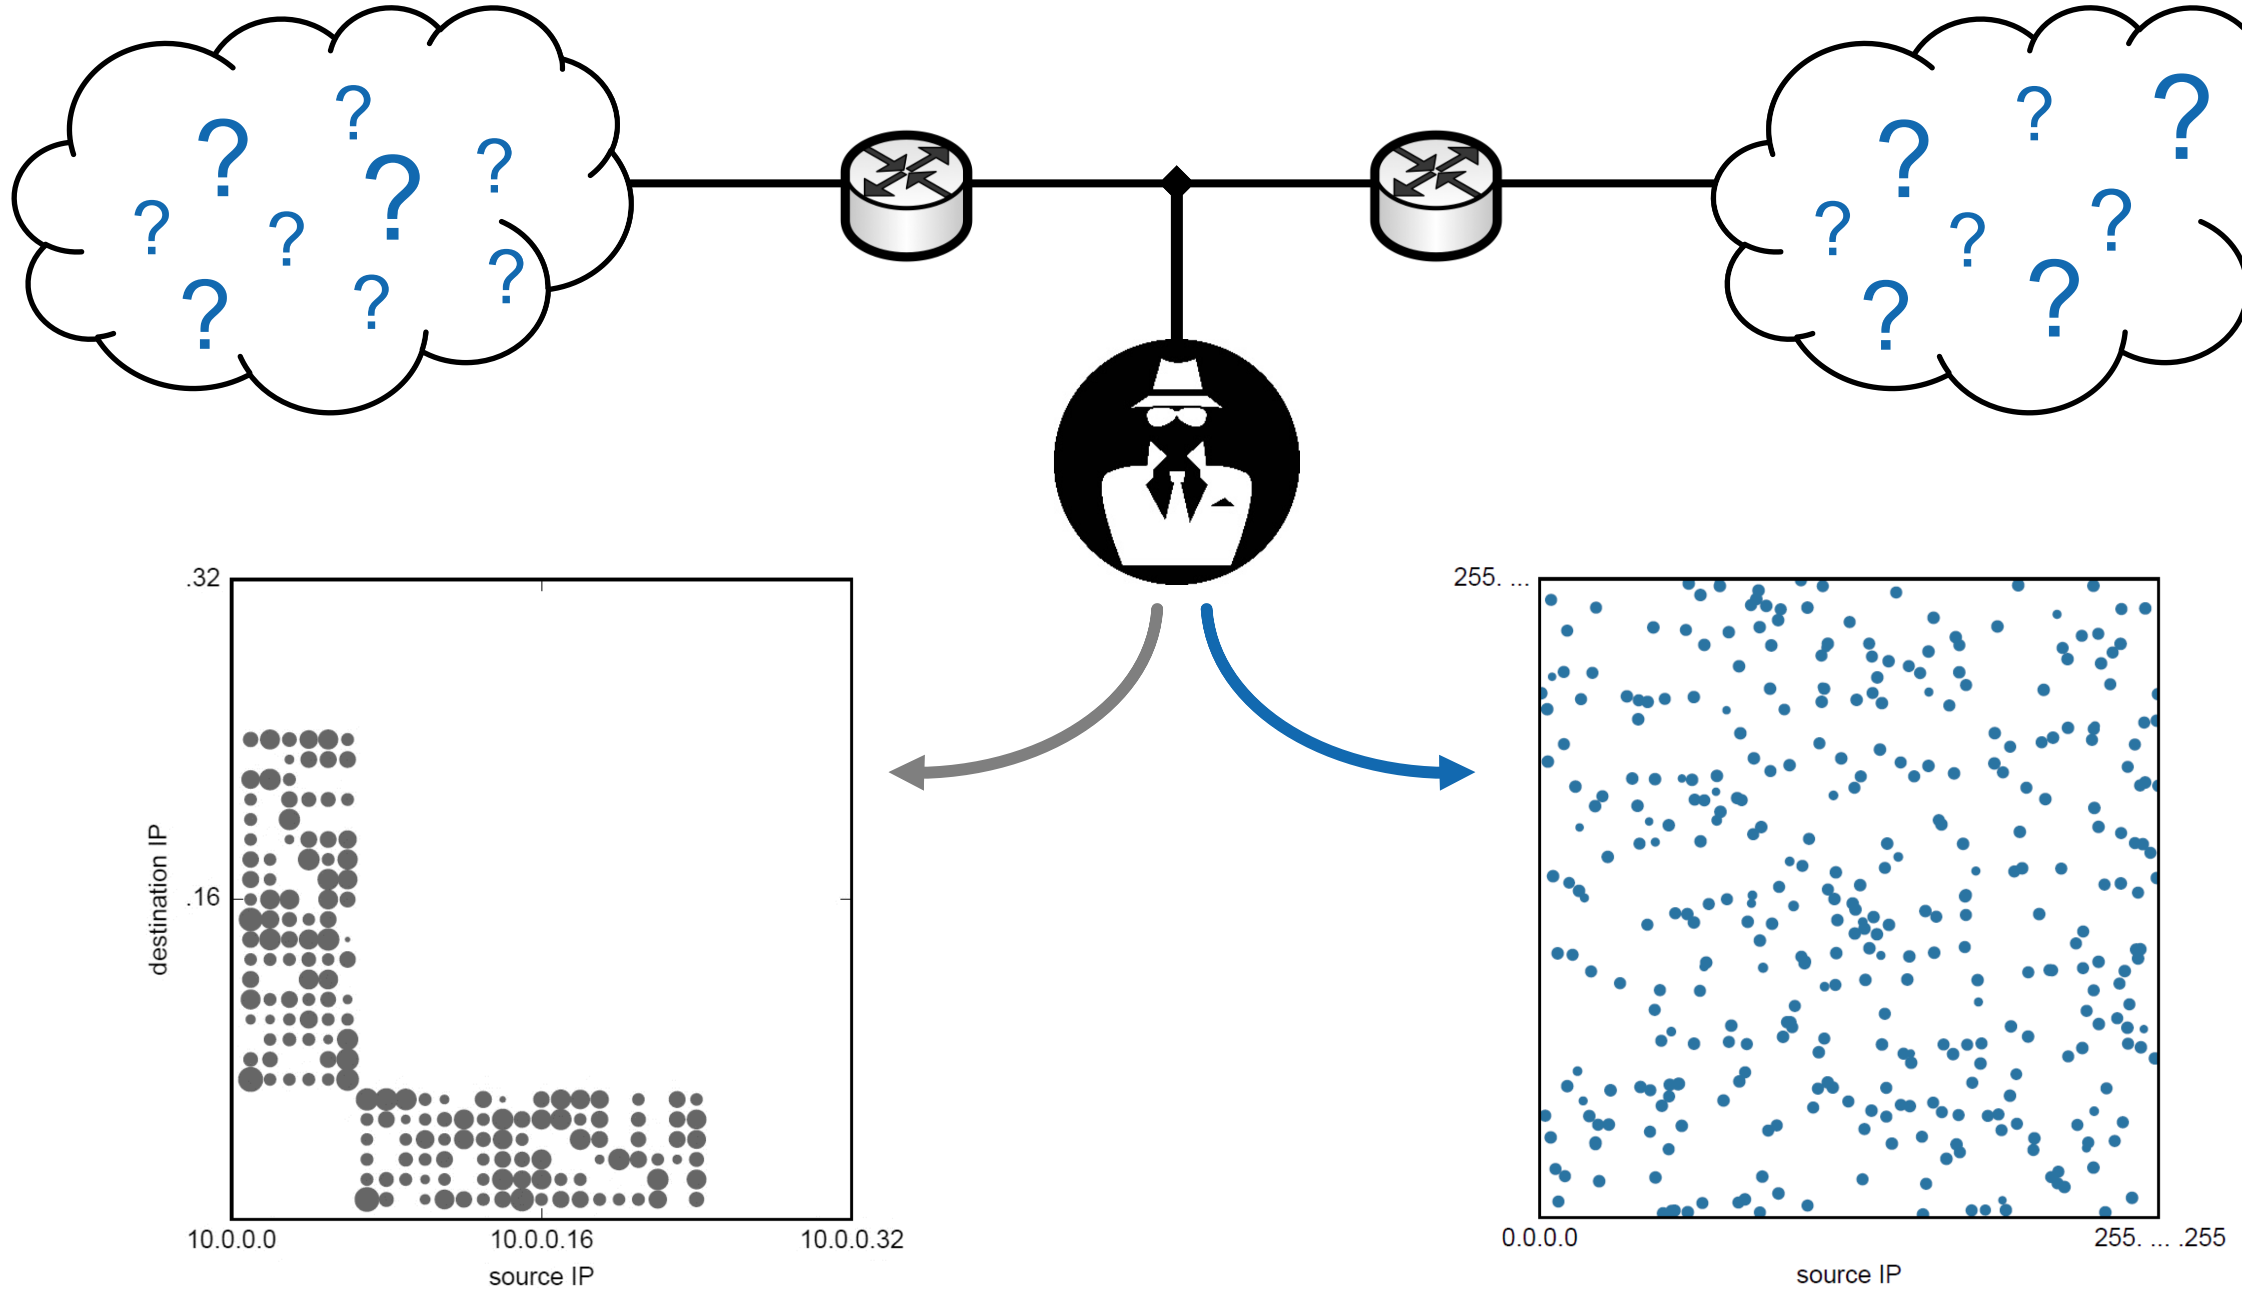

?
.32
.16
destination IP
10.0.0.0
10.0.0.16
10.0.0.32
source IP
255. ...
0.0.0.0
255. ... .255
source IP

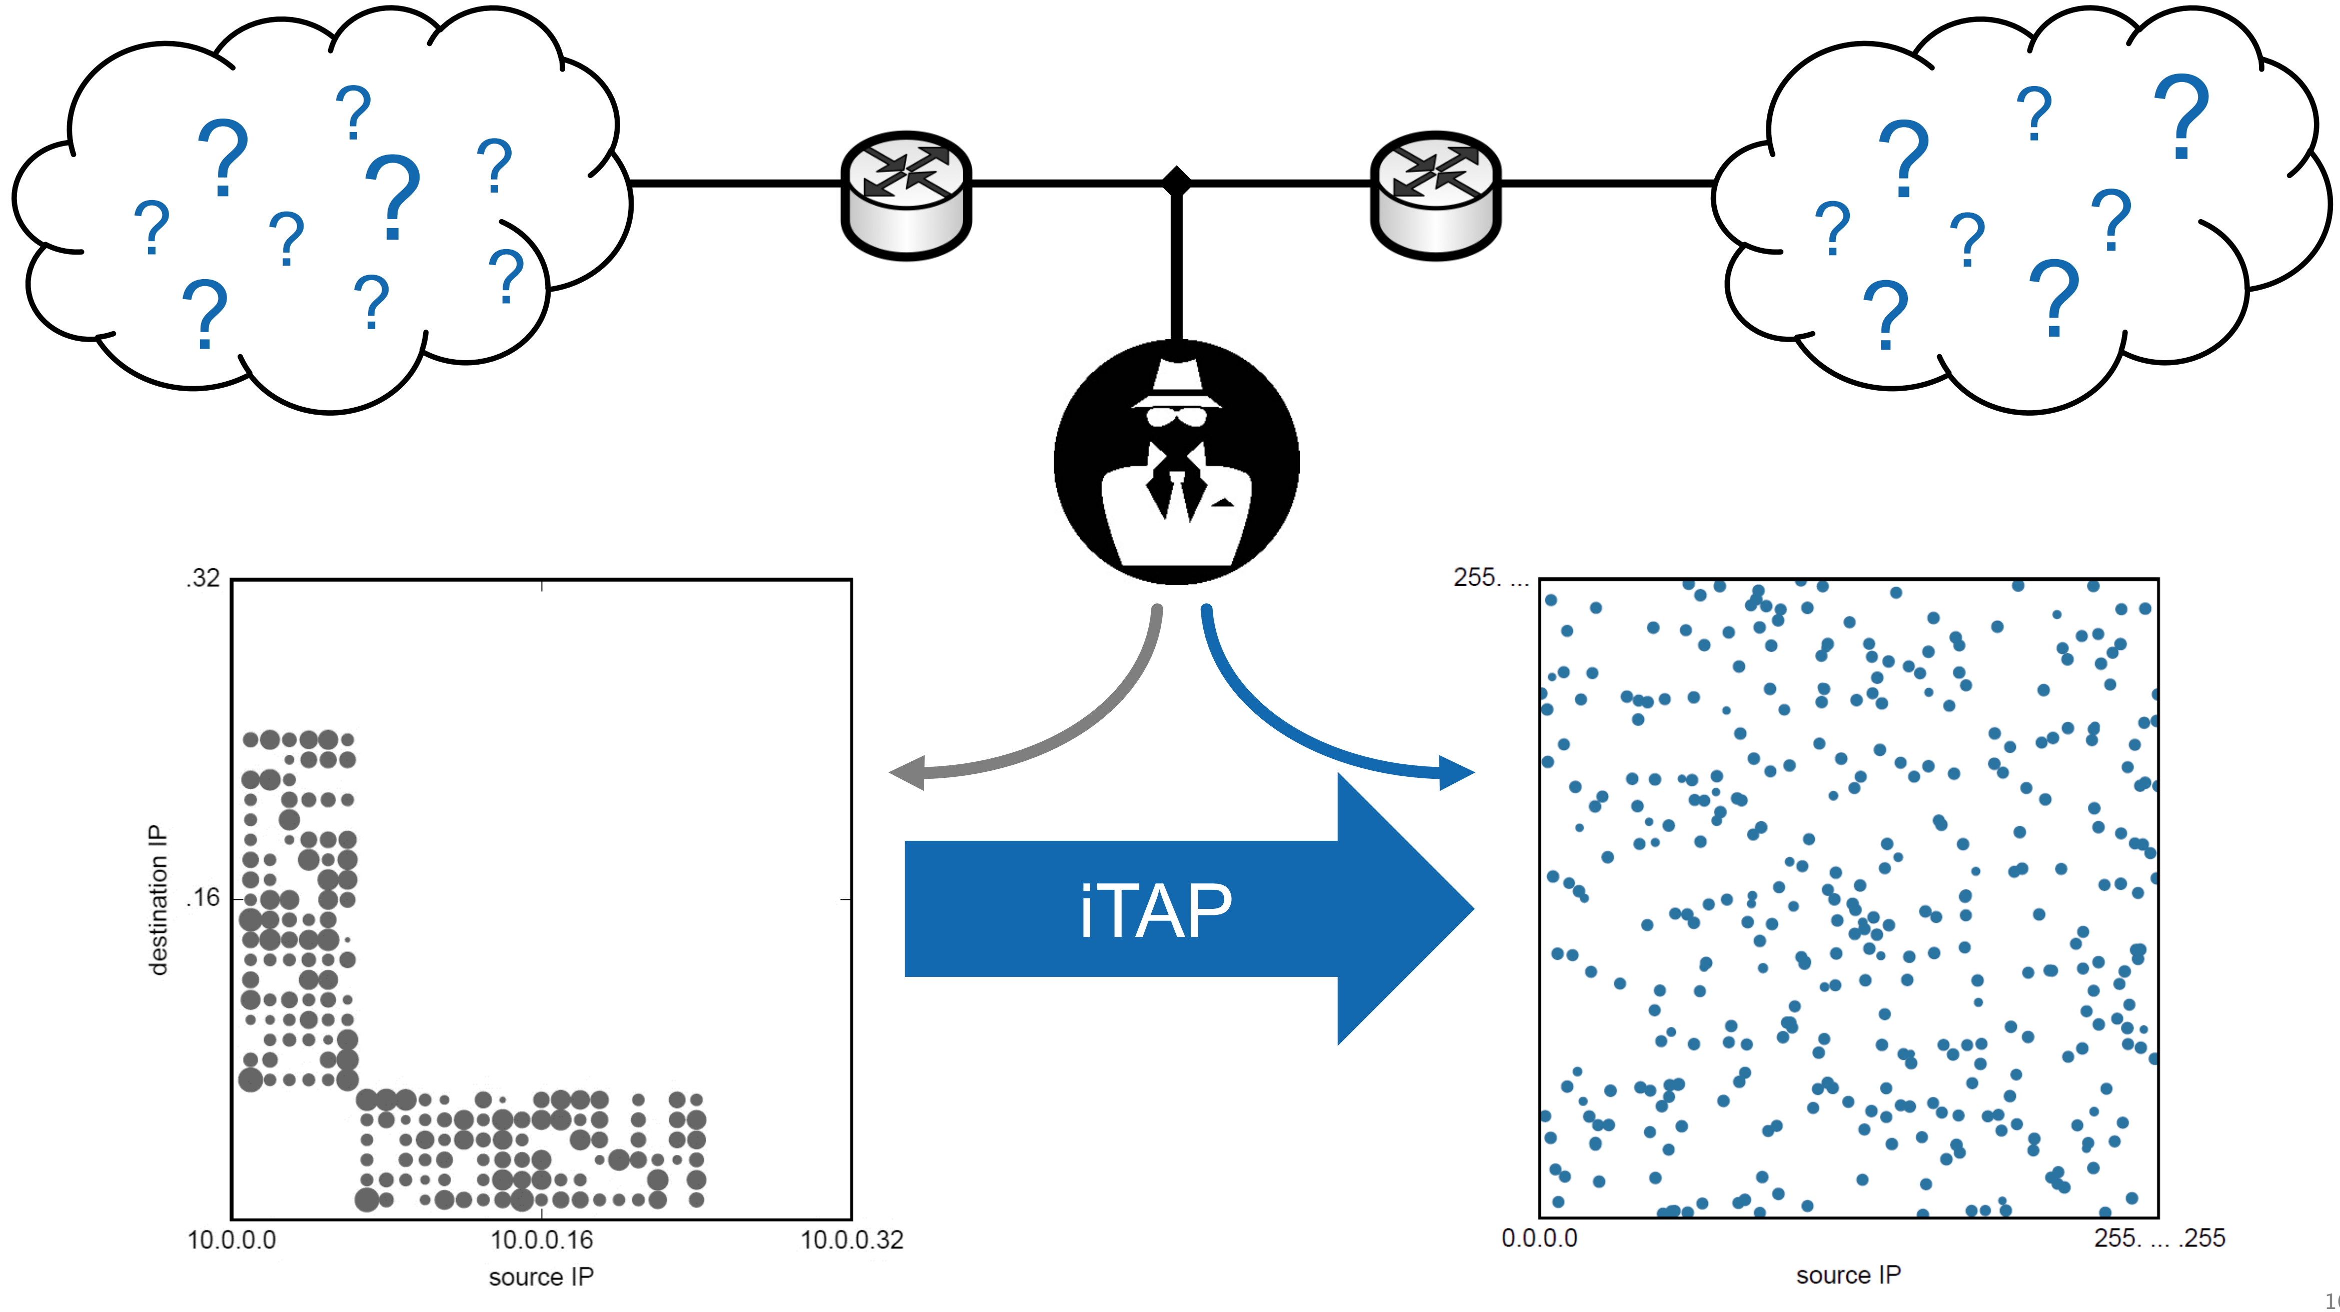

.32
.16
destination IP
10.0.0.0
10.0.0.16
10.0.0.32
source IP
iTAP
255. ...
0.0.0.0
255. ... .255
source IP

N.S.A. May Have Hit Internet Companies at a Weak Spot

The Internet companies' data centers are locked down with full-time security […]. But *between the data centers […] information was unencrypted and an easier target for government intercept efforts*, according to three people with knowledge of Google's and Yahoo's systems who spoke on the condition of anonymity.

- The New York Times, Nov. 25, 2013

The Washington Post

Technology

Google encrypts data amid backlash against NSA spying

By Craig Timberg September 6, 2013

Google is racing to encrypt the torrents of information that flow among its data centers around the world in a bid to thwart snooping by t
foreign governments

The move by Google
that recent revelati
sweeping surveillan
backlash within an

Post and the G
from American technology companies,
der various legal authorities.

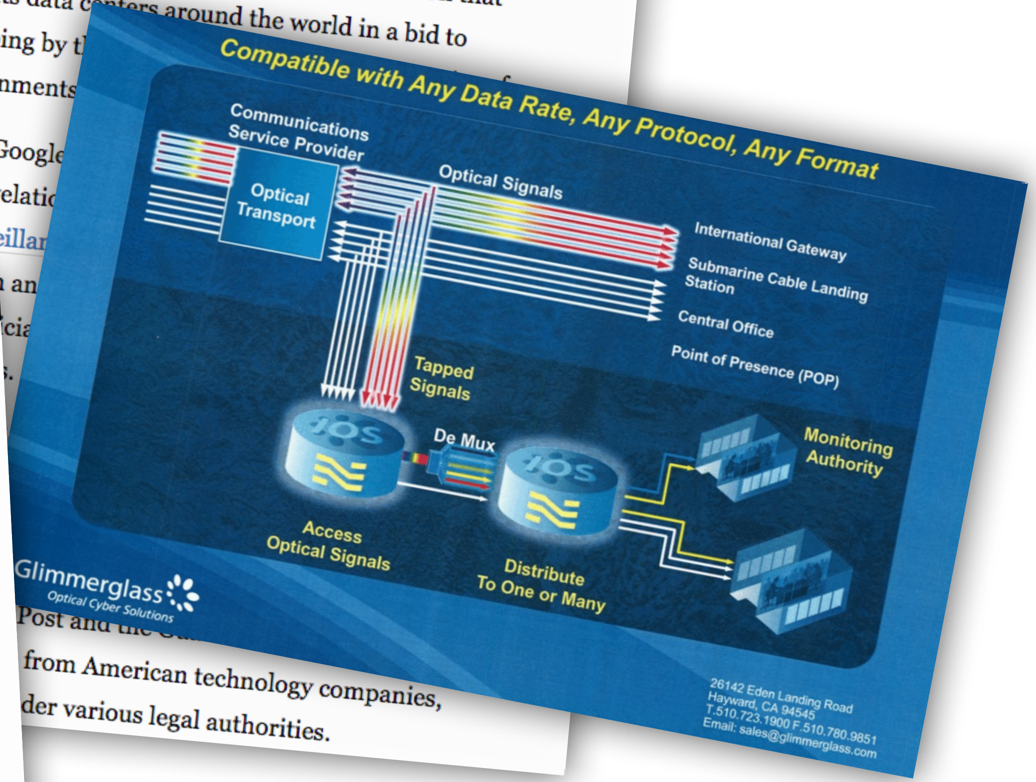

The Atlantic SUBSCRIBE SEARCH MENU

GLOBAL

The Creepy, Long-Standing Practice of Undersea Cable Tapping

The newest NSA leaks reveal that governments are probing "the Internet's backbone." How does that work?

1.8k 808

OLGA KHAZAN | JUL 16, 2013

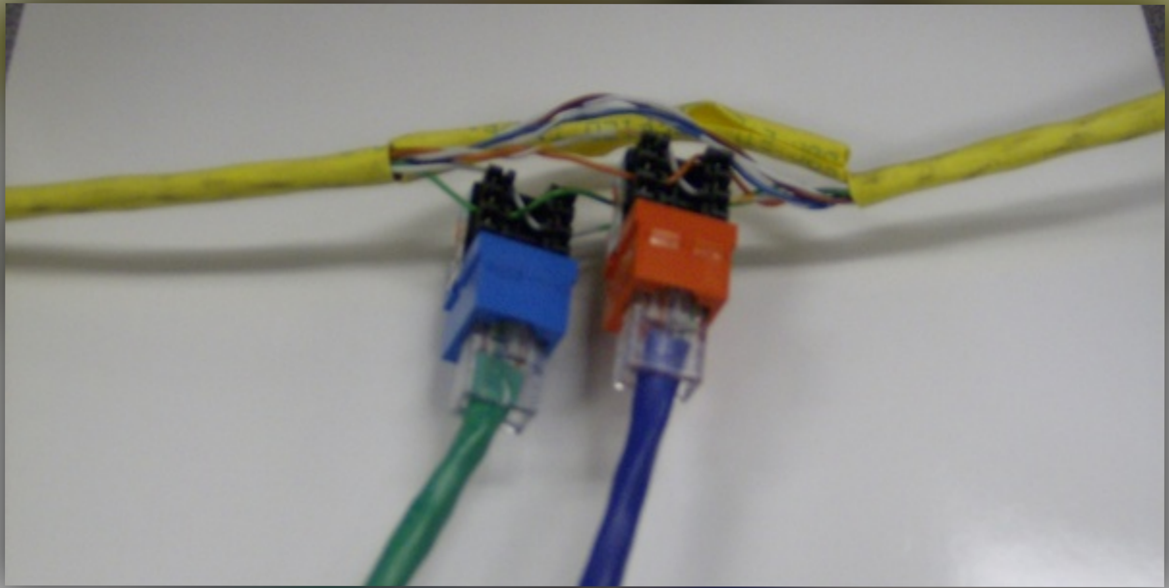

[janitha.com]


SharkTap 10/100
www.midBitTech.com
NETWORK
TAP

Existing solutions

Do not protect communicating parties
[SSL/TLS, IPsec Transport, MACsec]

Require modifications at end-hosts or additional middleboxes
[APOD, CONTRA]

Do not support partial deployment or have scalability problems
[MACsec, PHEAR]

More references provided in the paper

iTAP

In-network Traffic Analysis Prevention using Software-Defined Networks

Roland Meier, David Gugelmann, Laurent Vanbever

iTAP

In-network Traffic Analysis Prevention
using Software-Defined Networks

iTAP

In-network Traffic Analysis Prevention
using Software-Defined Networks

- Communication anonymity
  who is communicating with whom?

iTAP

In-network Traffic Analysis Prevention using Software-Defined Networks

- Communication anonymity
  who is communicating with whom?
- Volume anonymity
  how much traffic flows between X and Y?

iTAP

In-network Traffic Analysis Prevention
using Software-Defined Networks

- Communication anonymity
  who is communicating with whom?
- Volume anonymity
  how much traffic flows between X and Y?
- Topology anonymity
  how many hosts are in the network?

iTAP

In-network Traffic Analysis Prevention
using Software-Defined Networks

- No modifications at end-hosts

iTAP

In-network Traffic Analysis Prevention
using Software-Defined Networks

- Central controller
- Rewriting capabilities of switches

An iTAP-protected network

An iTAP-protected network

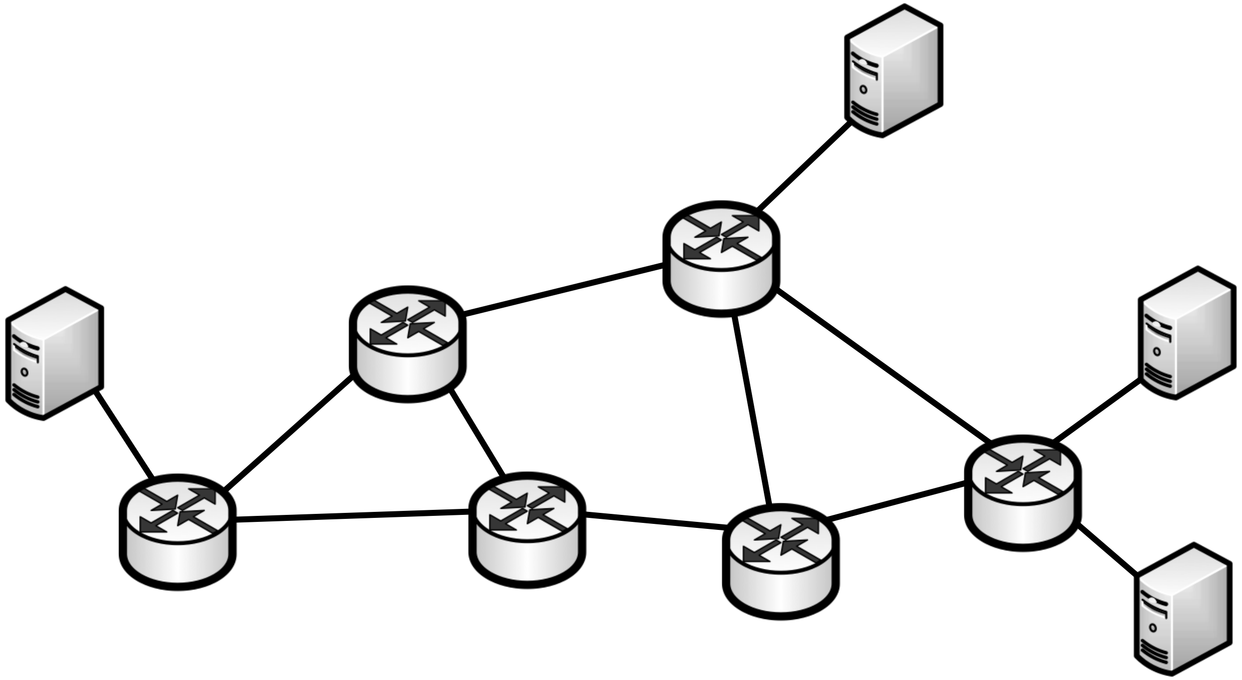

Layer 2 network

An iTAP-protected network

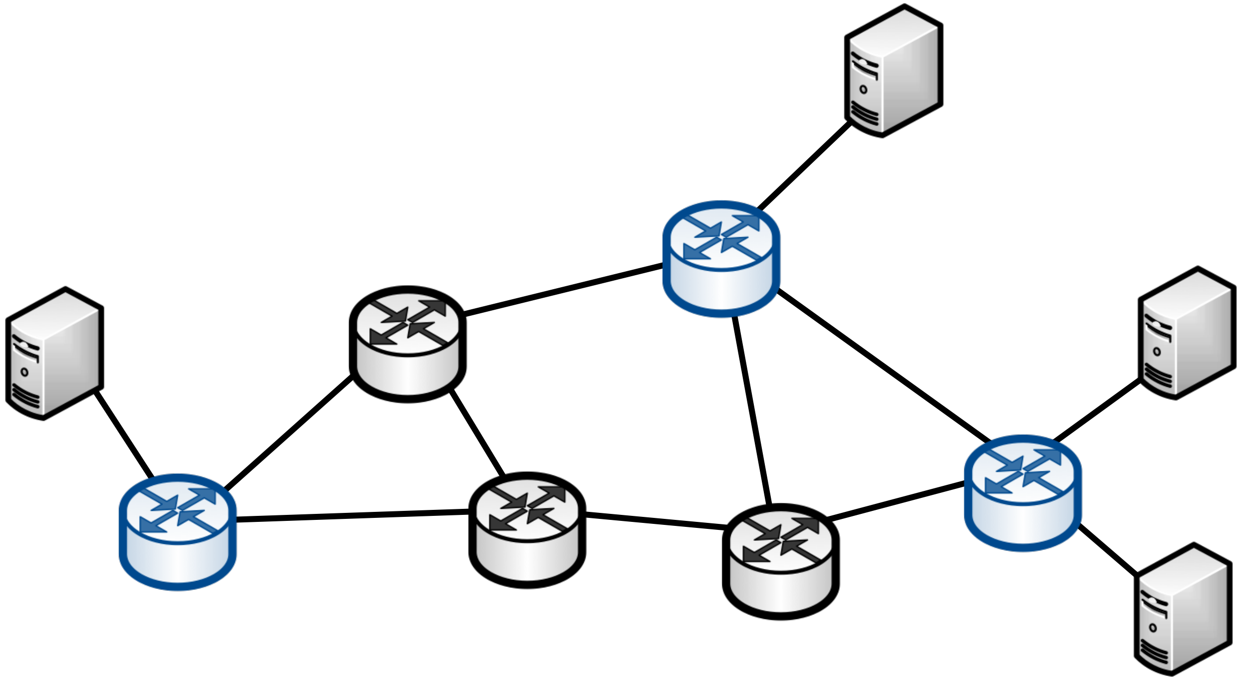

 Layer 2 network

 With some SDN switches

An iTAP-protected network

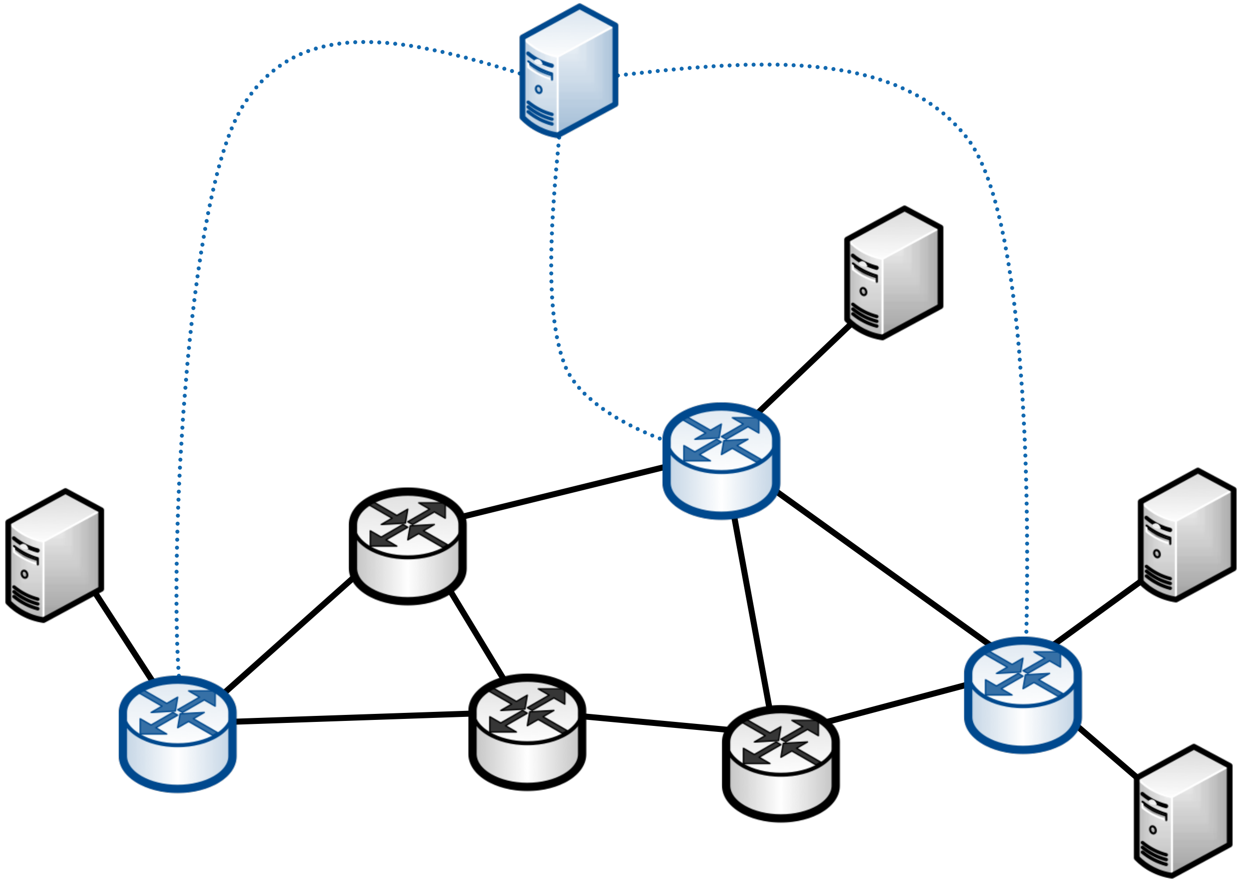

Layer 2 network

With some SDN switches

And a central controller

An iTAP-protected network

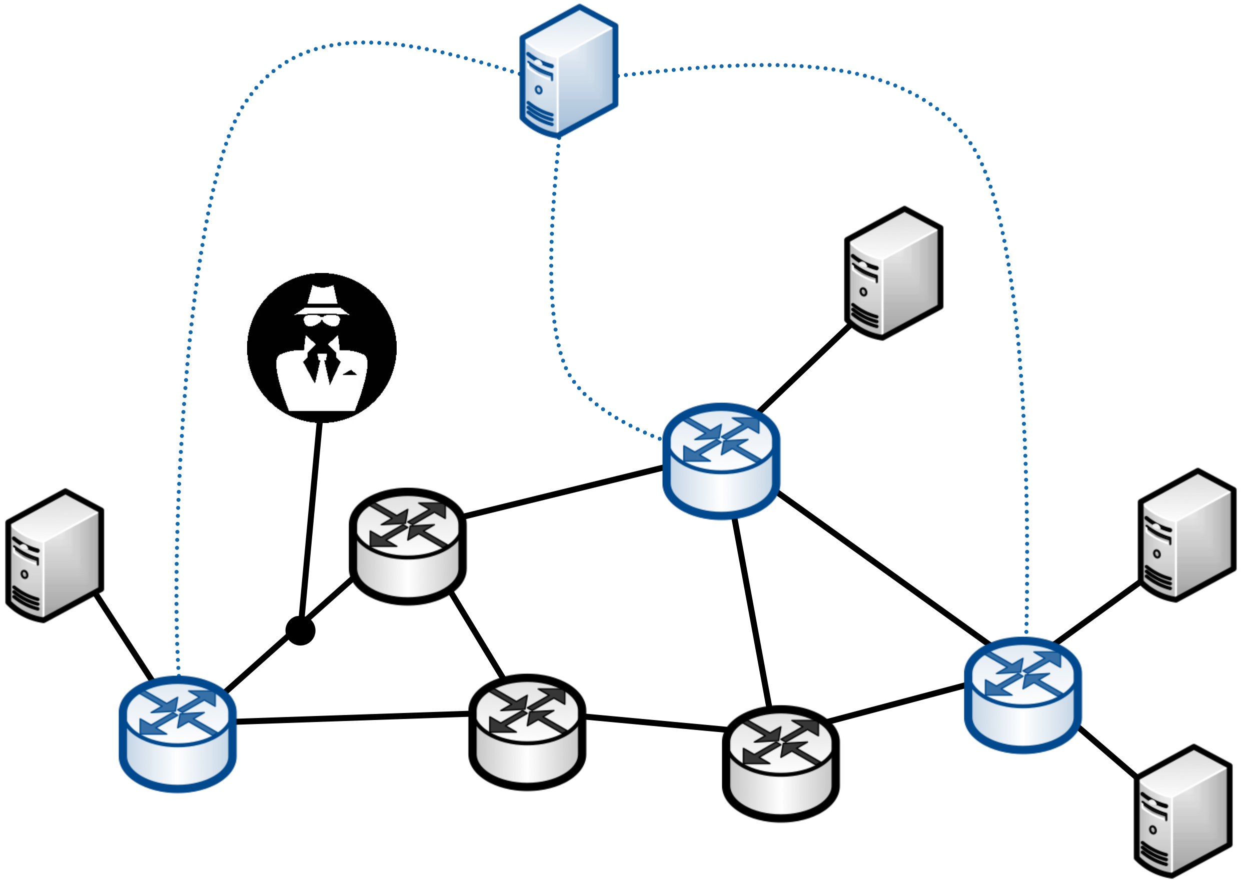

Layer 2 network

With some SDN switches

And a central controller

Attacked by an eavesdropper

An iTAP-protected network

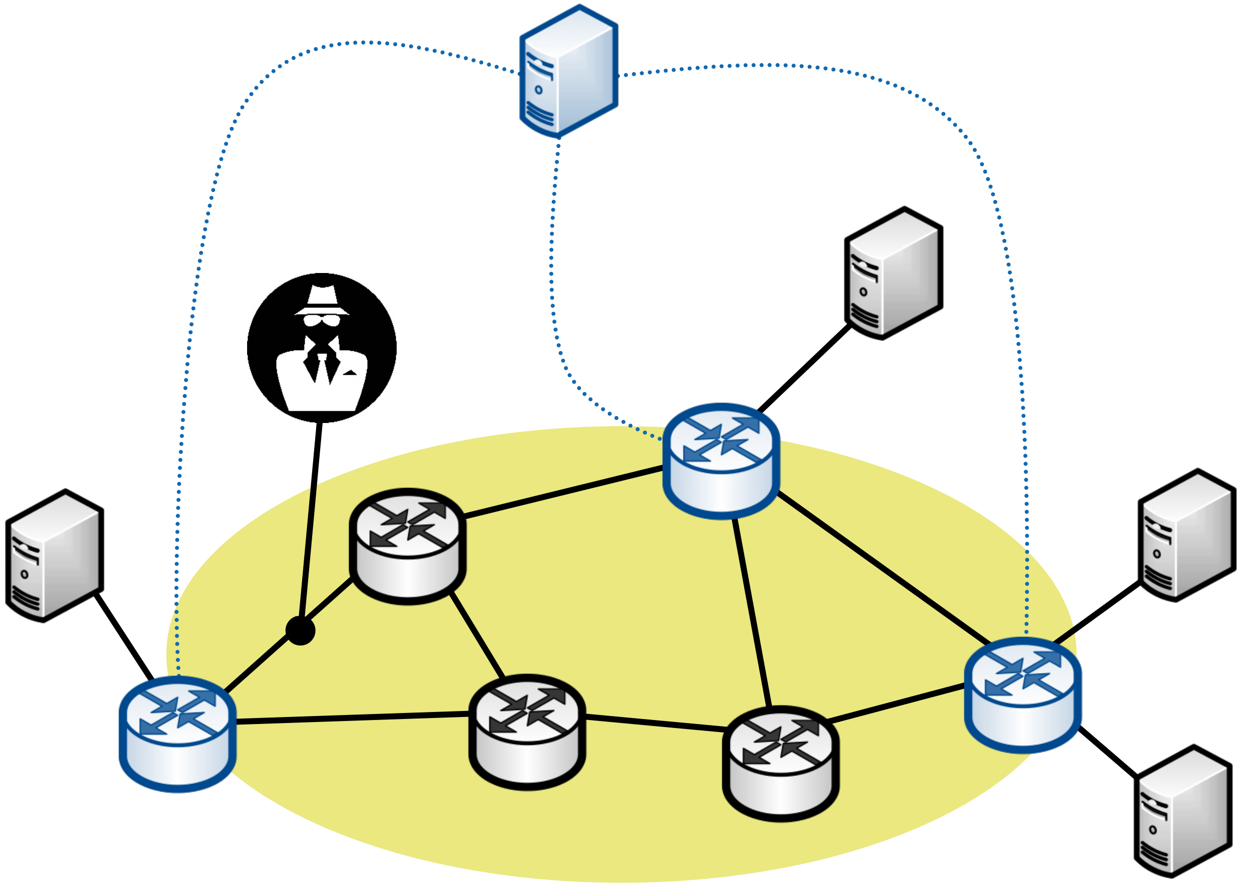

- Layer 2 network
- With some SDN switches
- And a central controller
- Attacked by an eavesdropper
- Protected by iTAP

Example

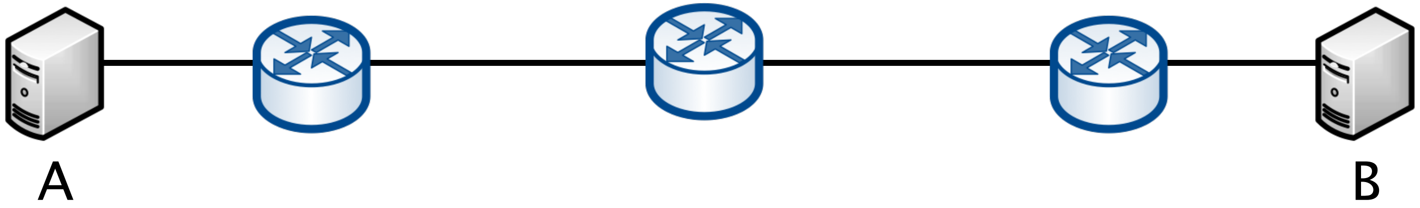

Example

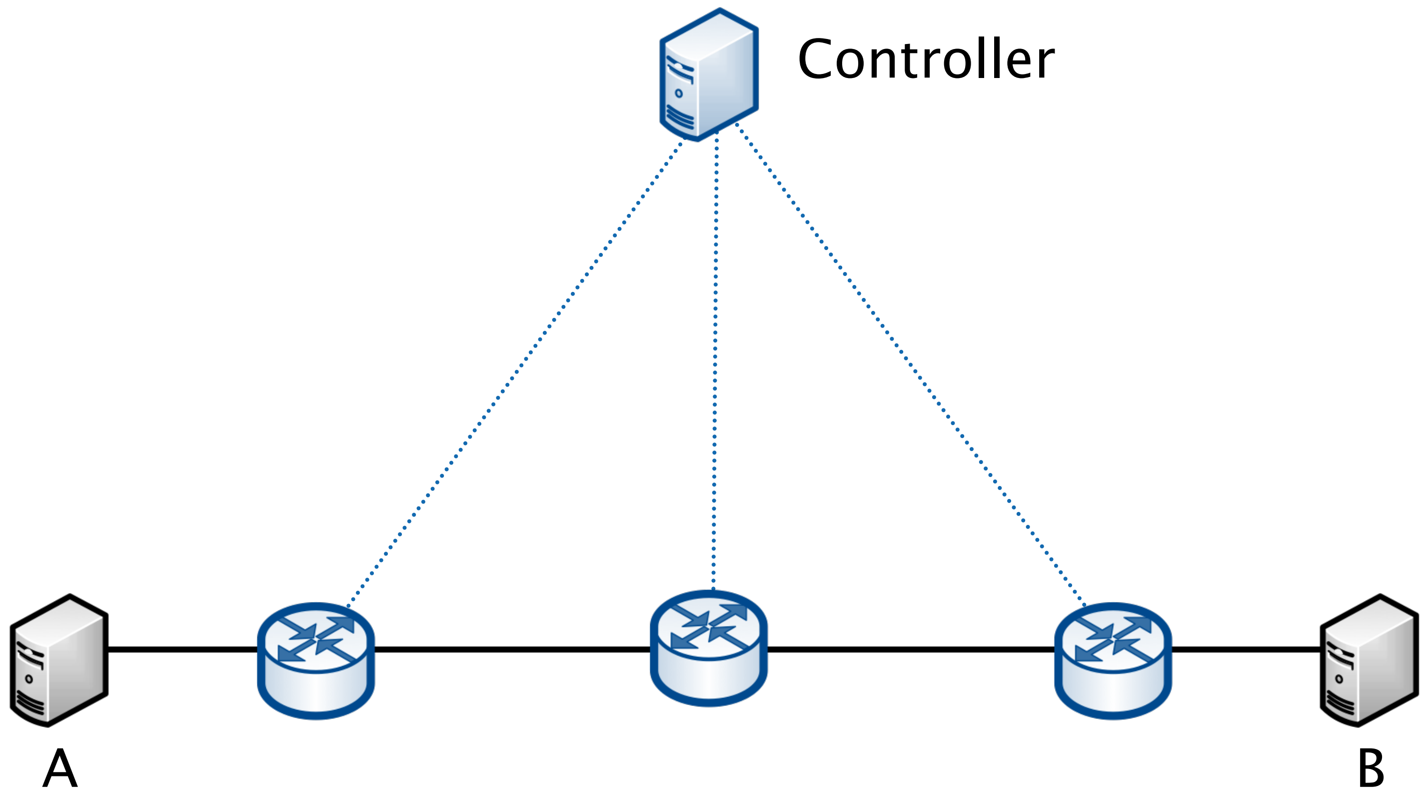

Packet from A to B enters the network

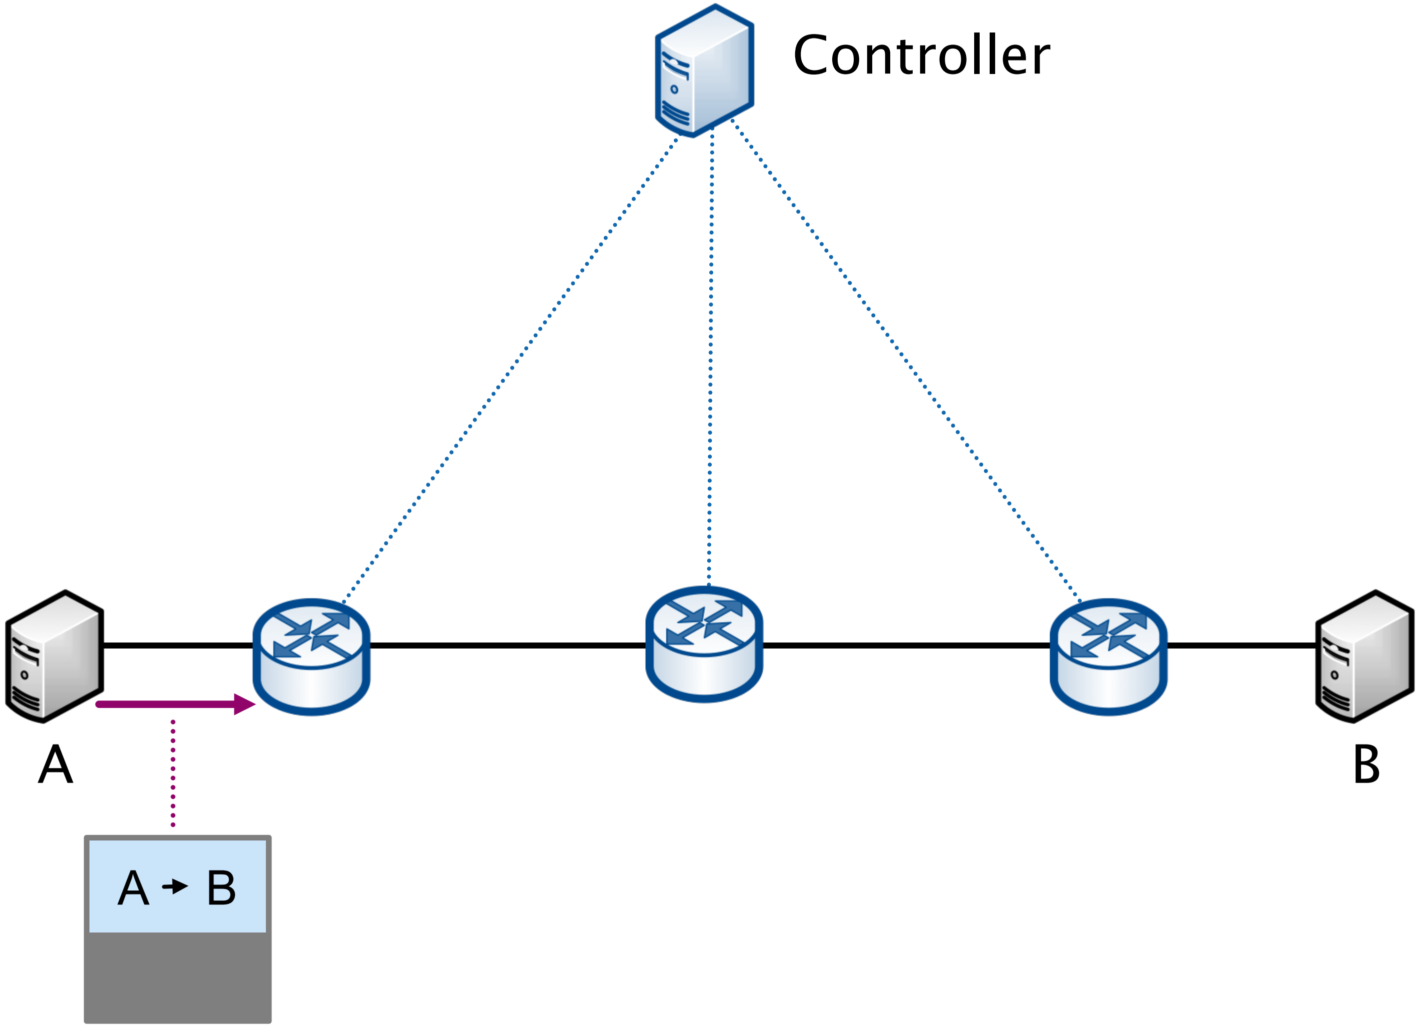

Ingress switch notifies controller

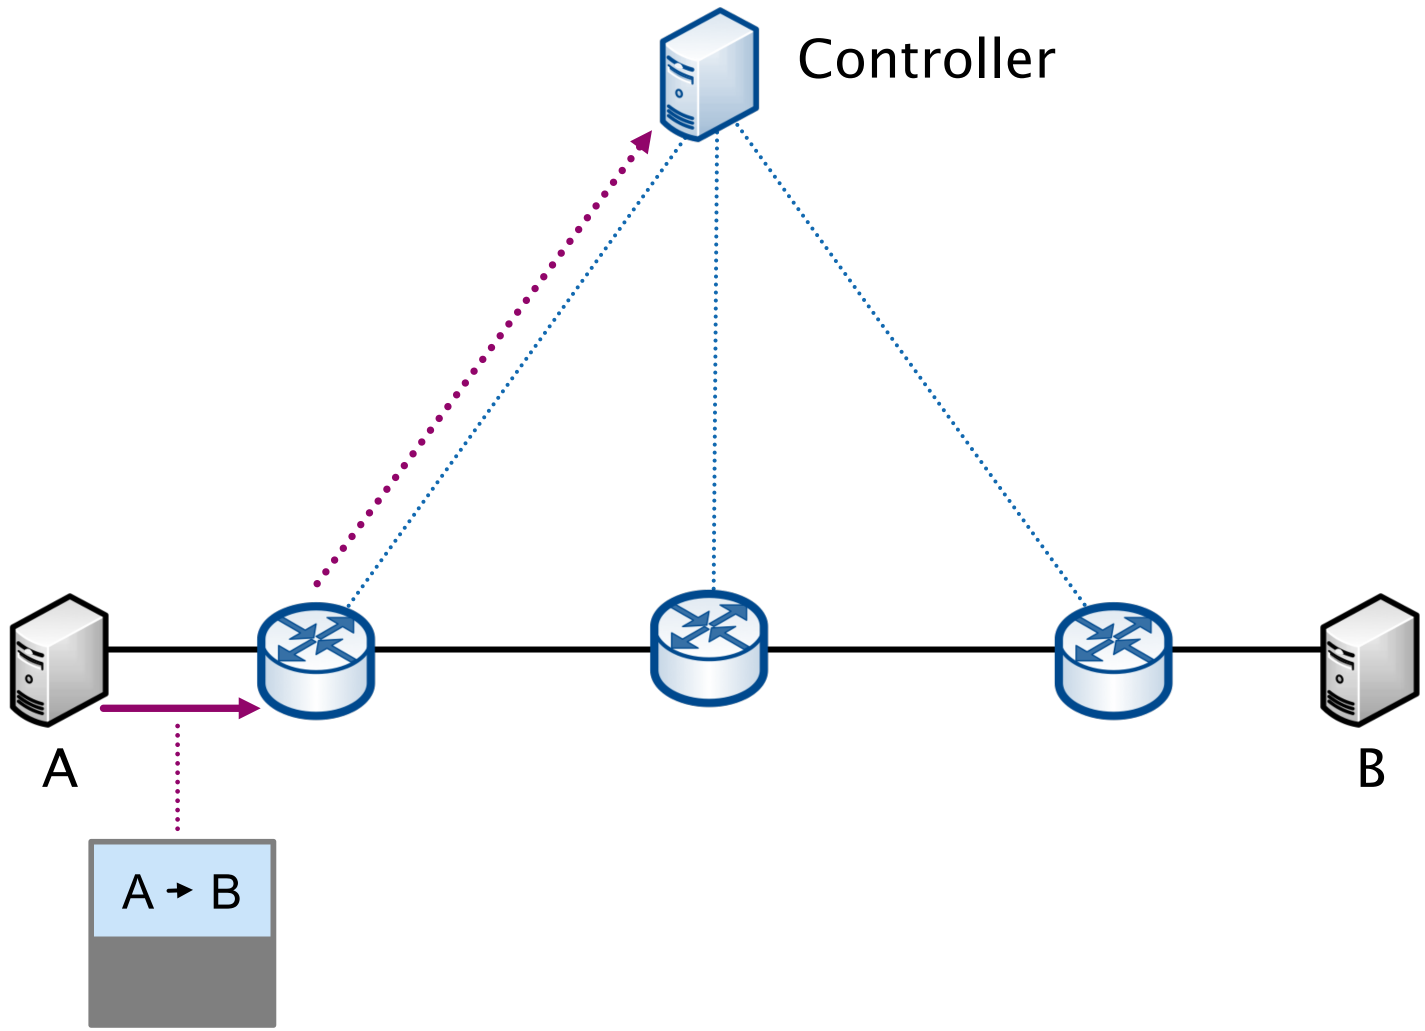

Controller computes & installs flow rules

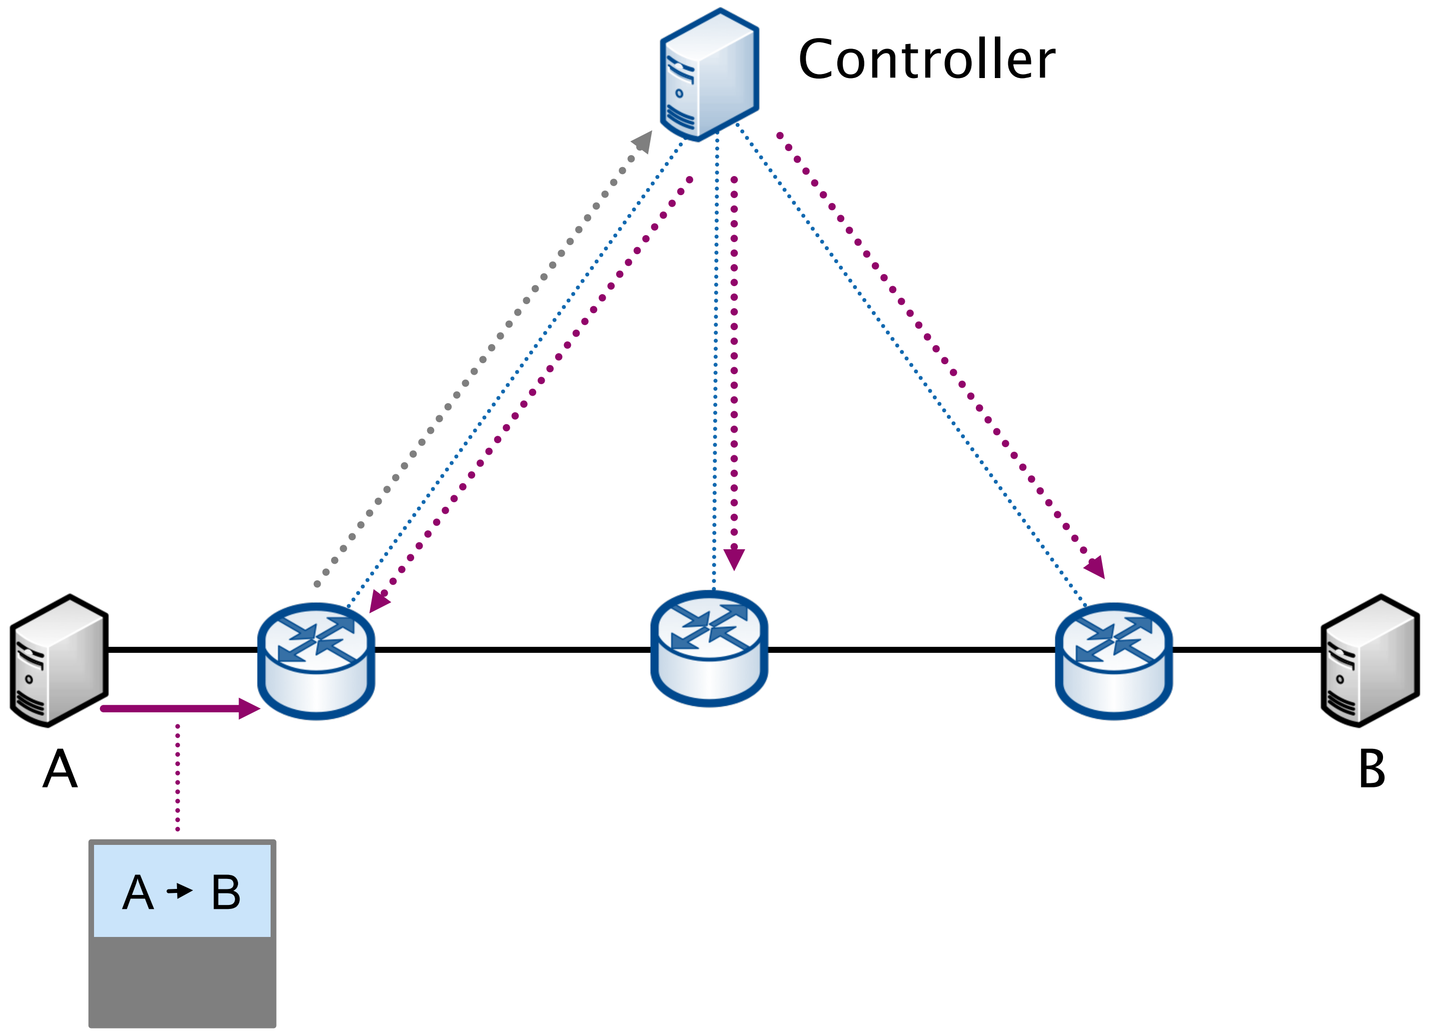

Ingress switch obfuscates source and destination

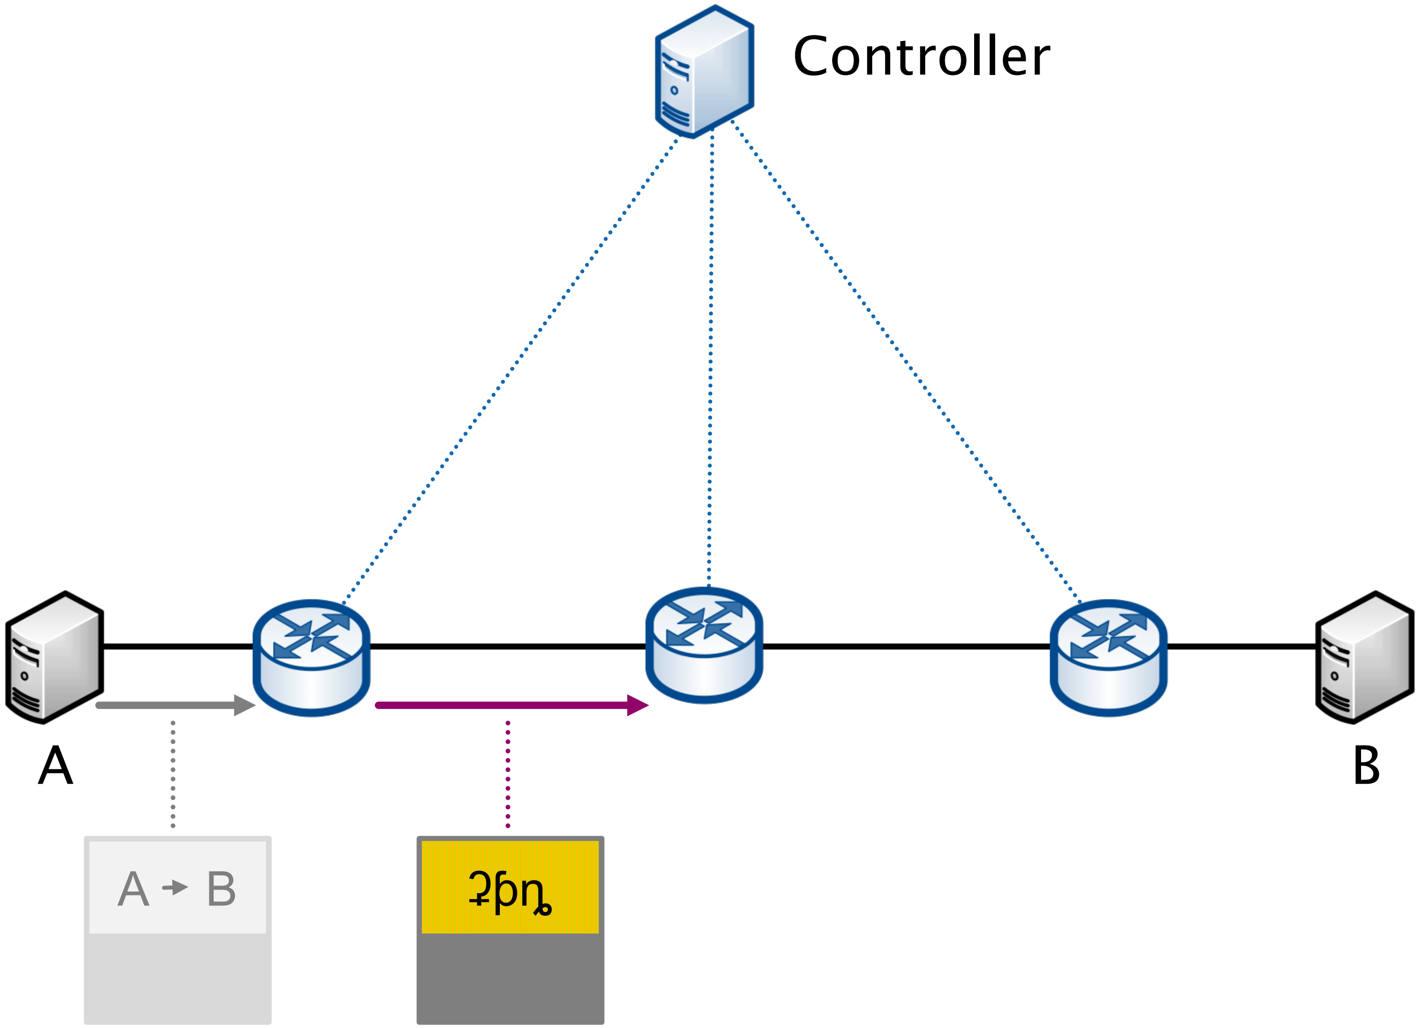

Core switch forwards obfuscated packet

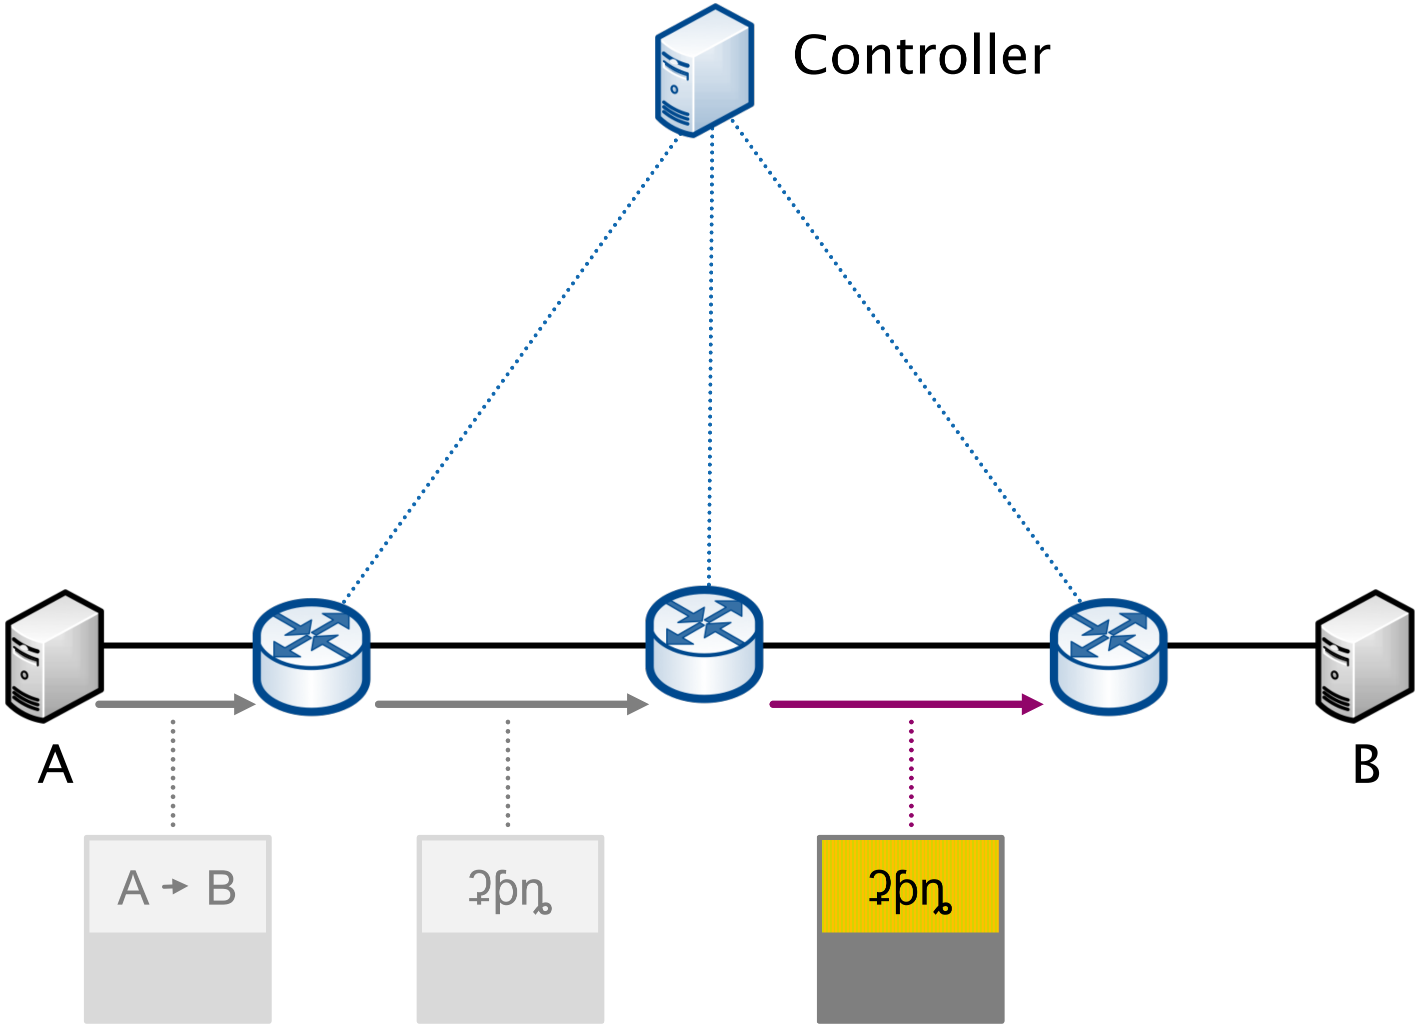

Egress switch de-obfuscates source and destination

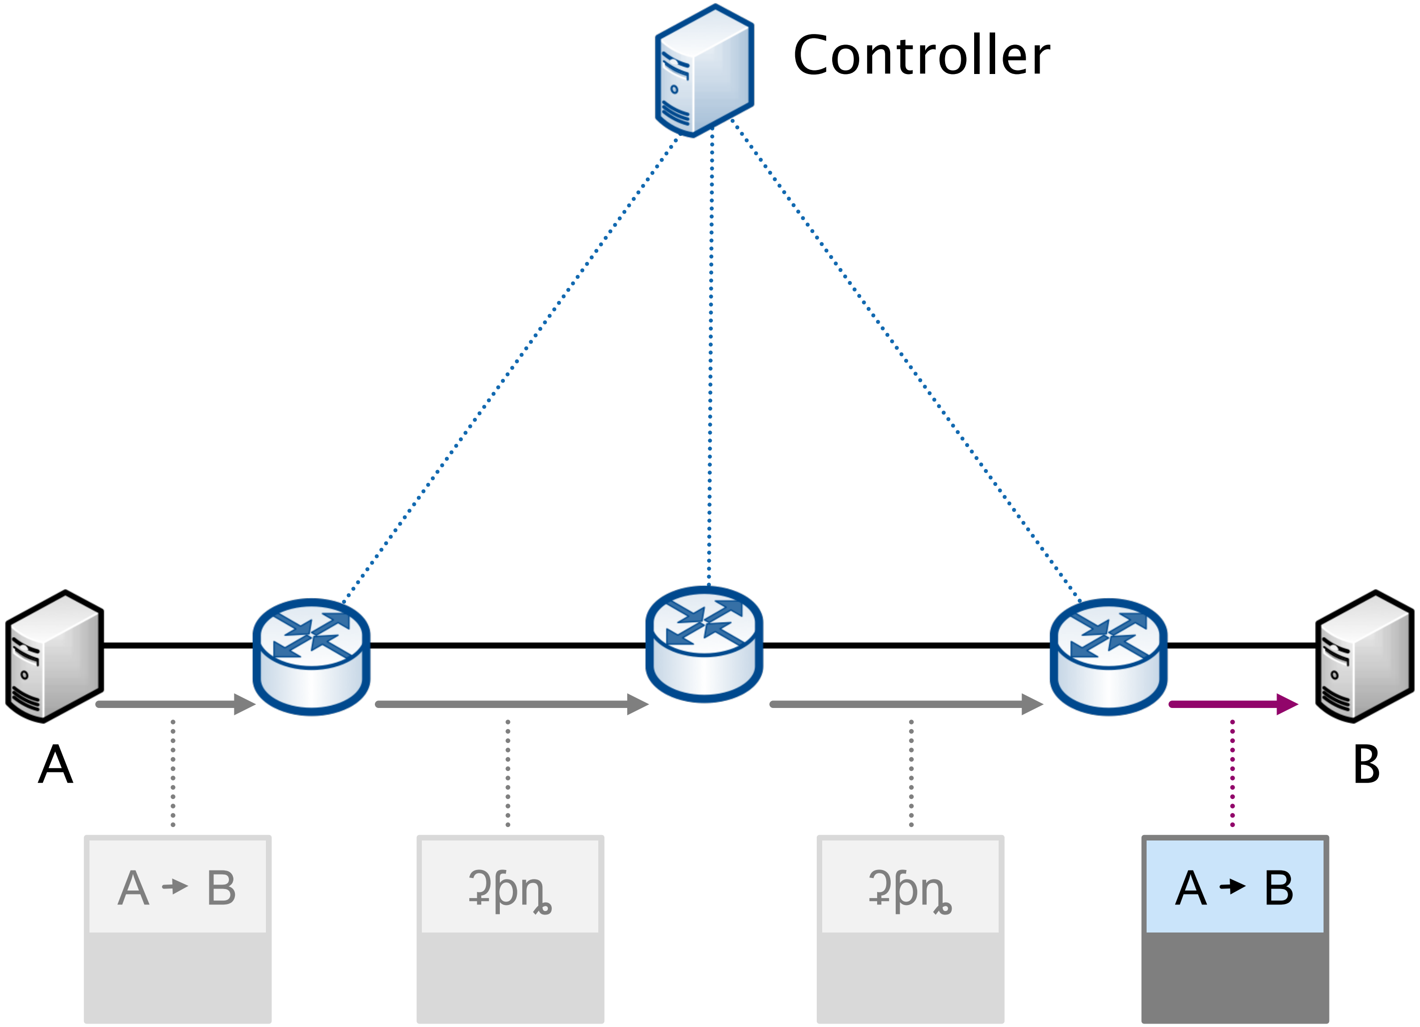

How does the rewriting work?

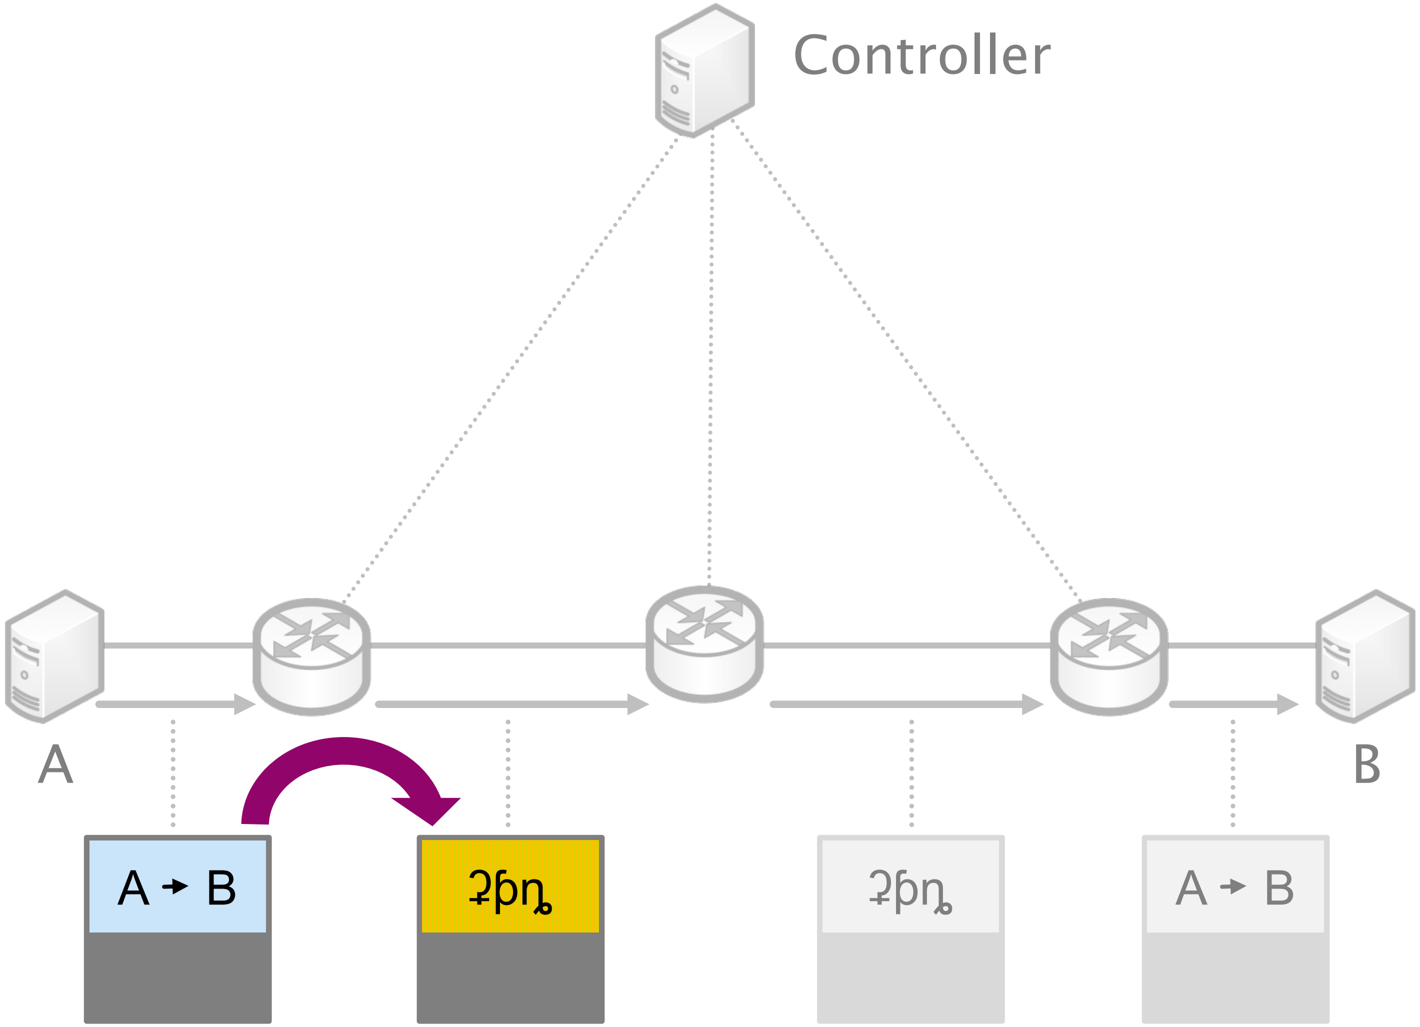

Rewriting packet headers

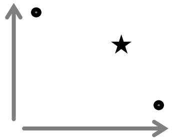

Trade-off between anonymity and scalability

iTAP approach: Mixing per-host IDs and random bits

Measure information leakage & counteract attacker

Rewriting packet headers as a trade-off between anonymity and scalability

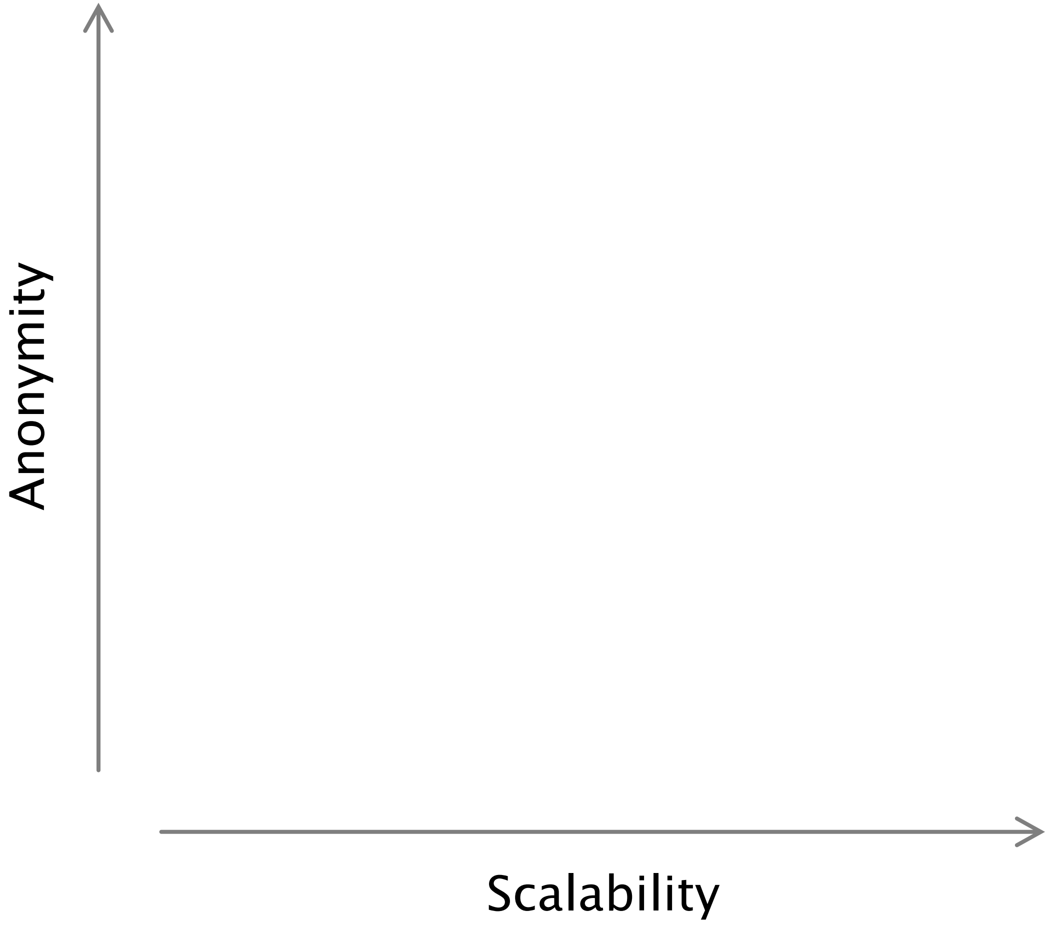

Rewriting packet headers as a trade-off between anonymity and scalability

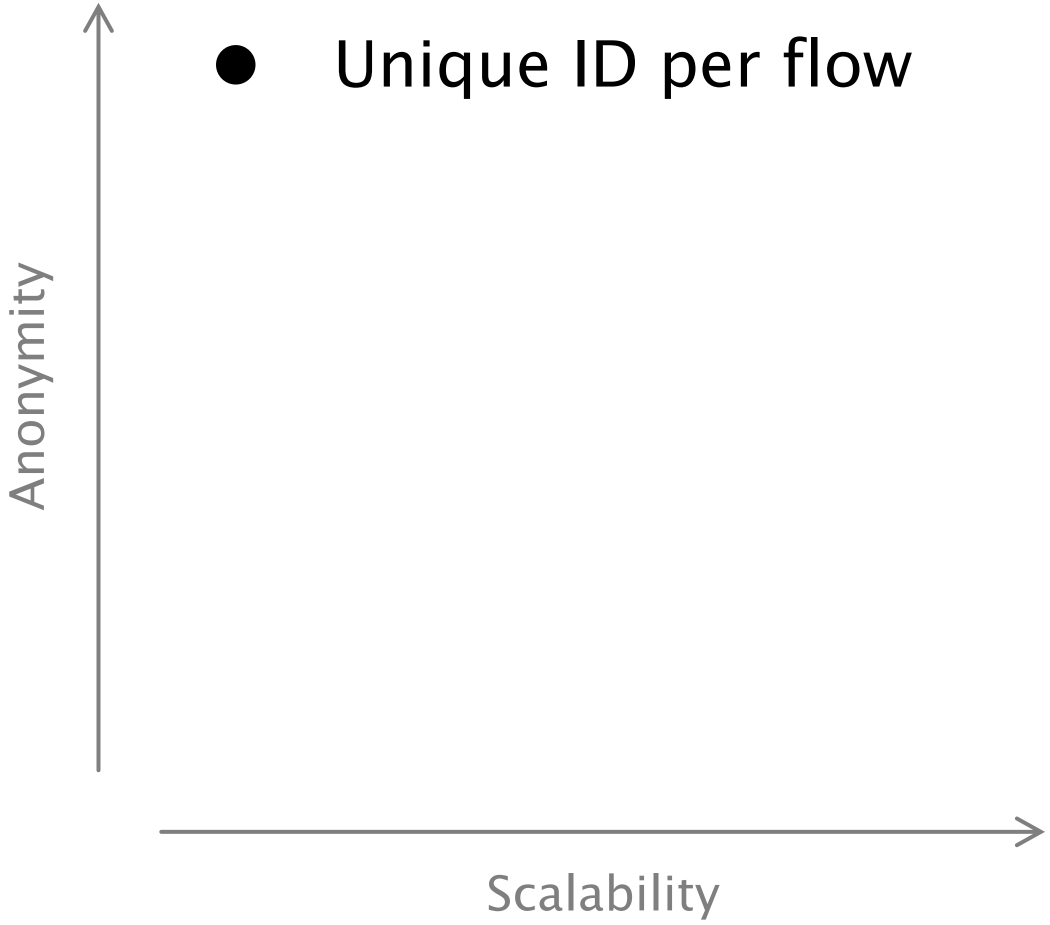

Rewriting packet headers as a trade-off between anonymity and scalability

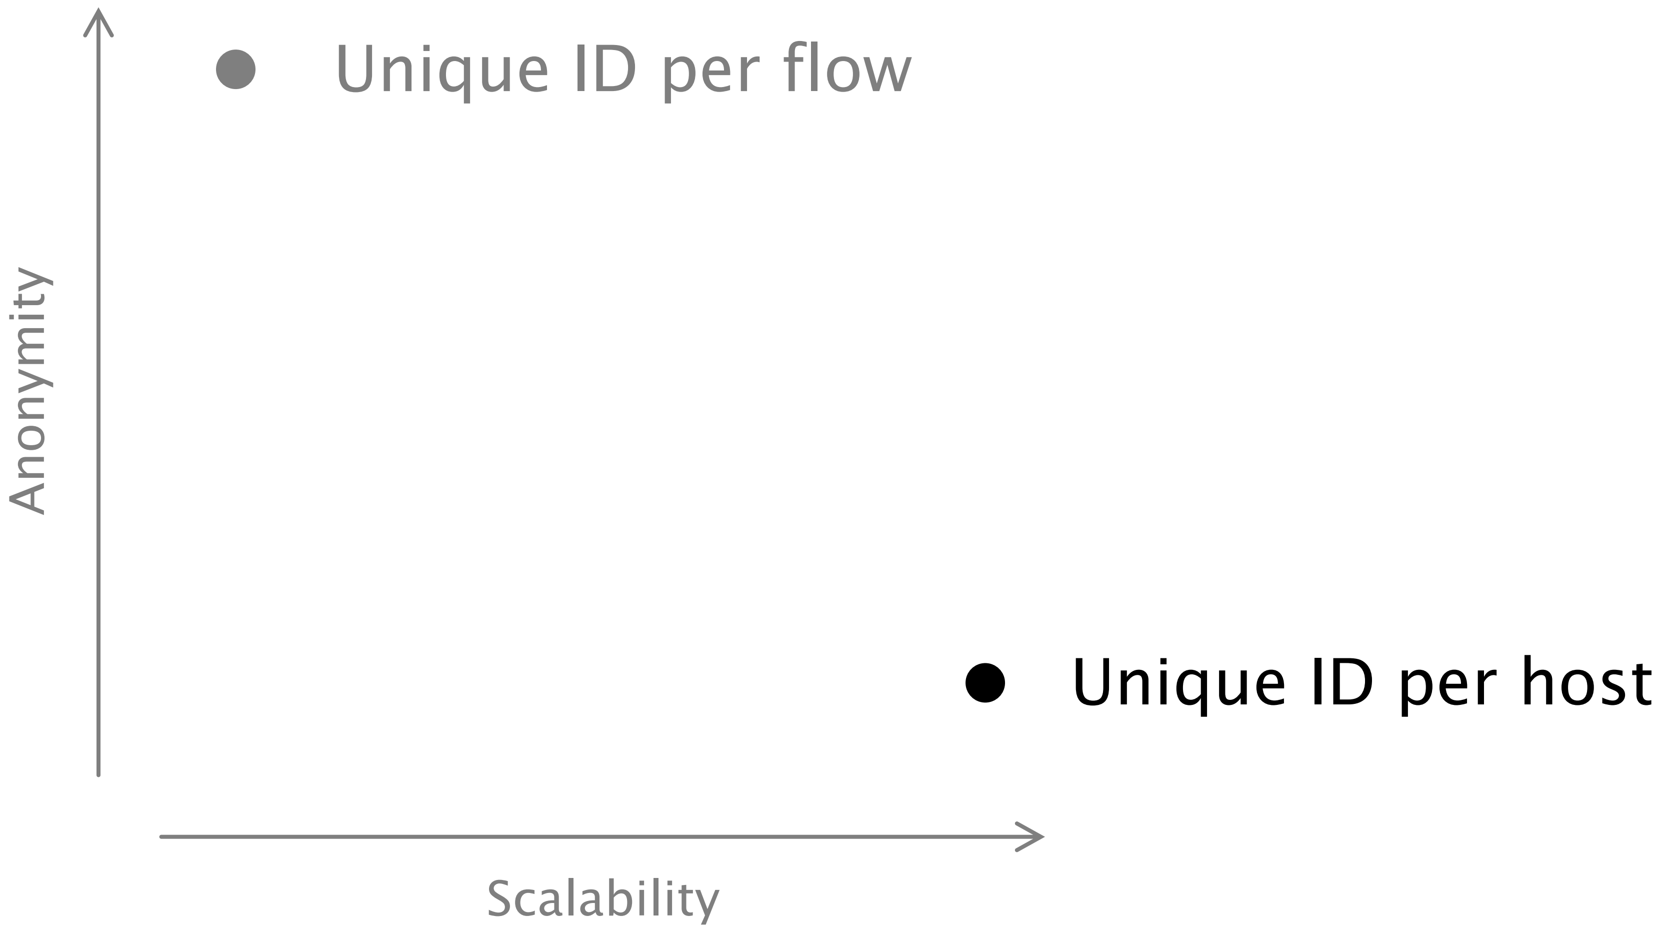

Rewriting packet headers as a trade-off between anonymity and scalability

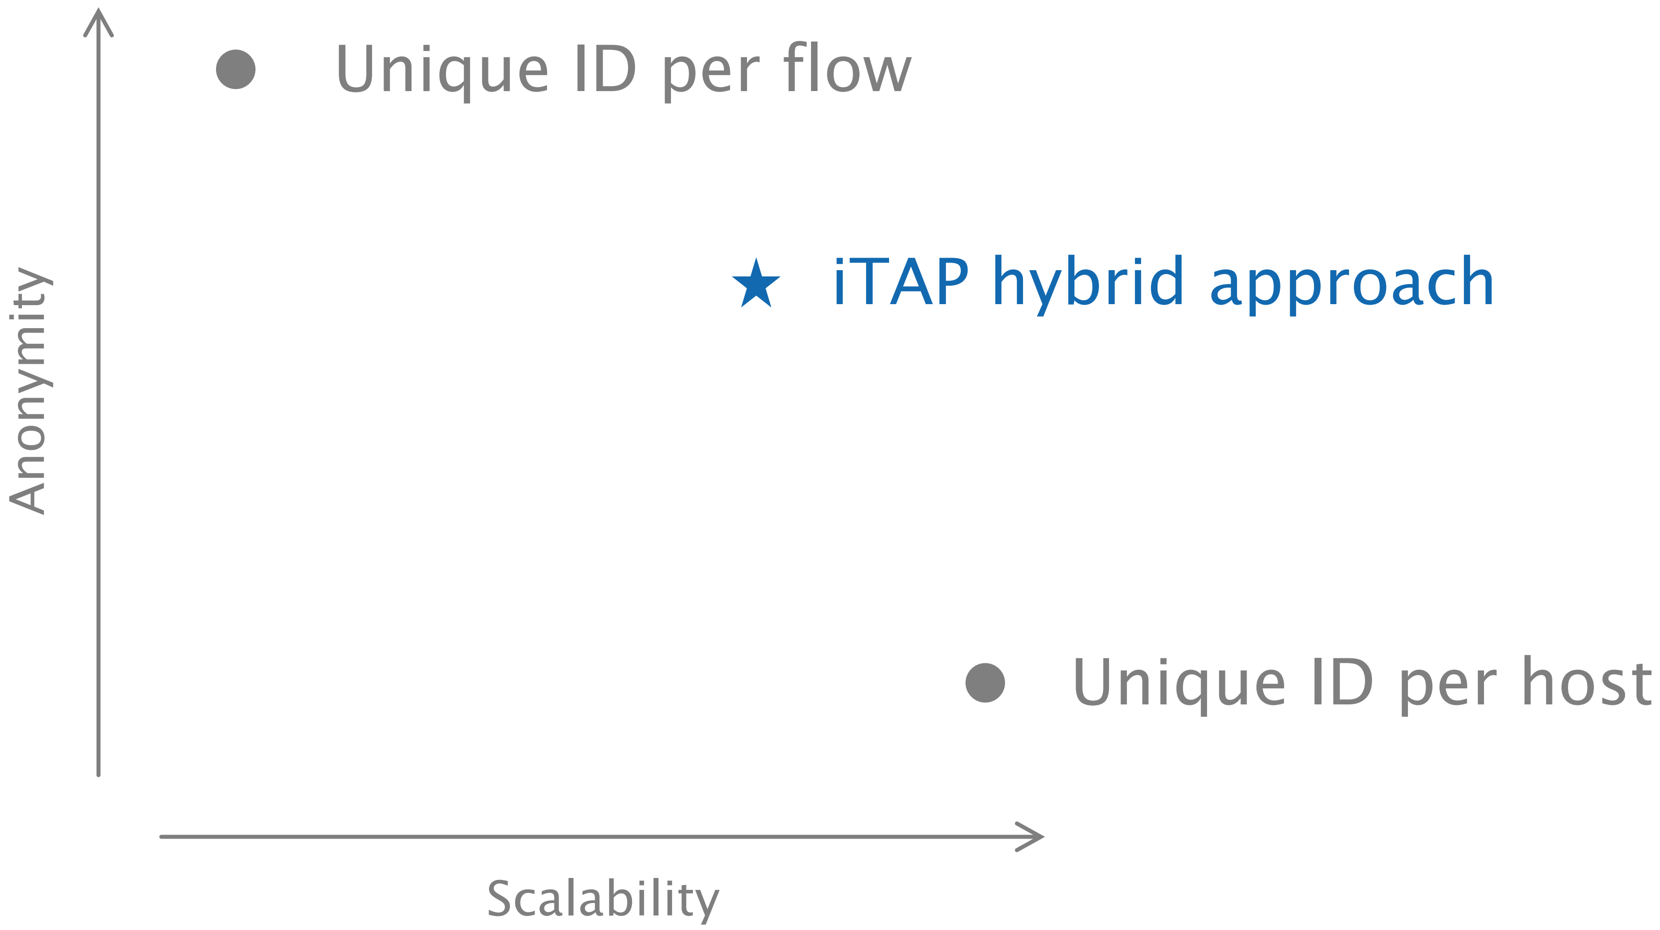

iTAP hybrid obfuscation scheme

iTAP hybrid obfuscation scheme

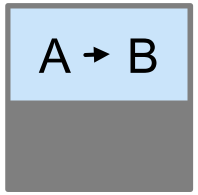

iTAP hybrid obfuscation scheme

Map source and destination to IDs

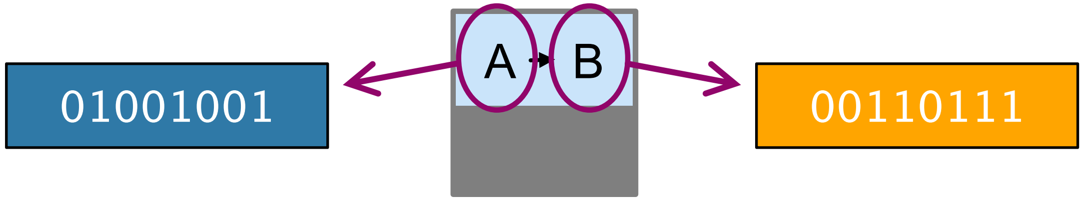

iTAP hybrid obfuscation scheme

Map source and destination to IDs

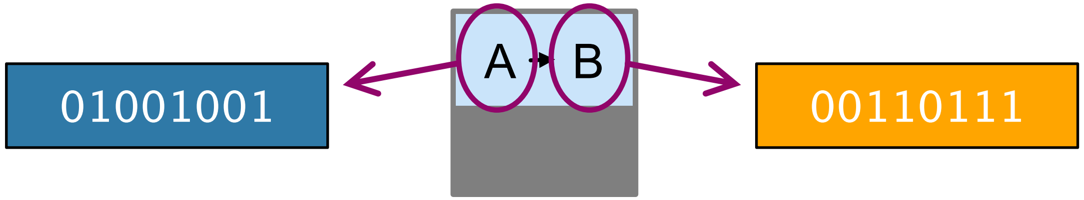

Match-fields with arbitrary bitmasks

iTAP hybrid obfuscation scheme

Map source and destination to IDs

Match-fields with arbitrary bitmasks

Interpret as bit-string of 160 bits

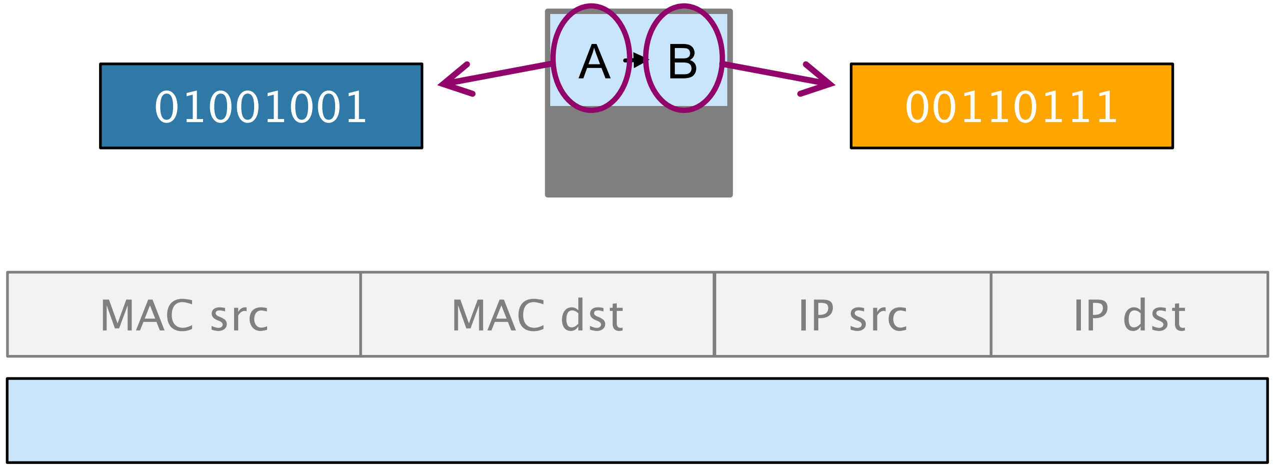

iTAP hybrid obfuscation scheme

Map source and destination to IDs

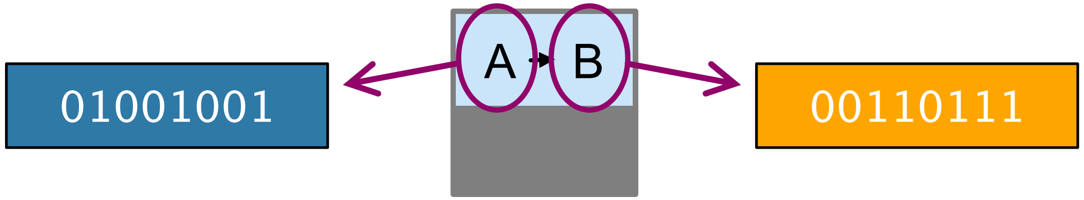

Match-fields with arbitrary bitmasks

Interpret as bit-string of 160 bits

Randomly select bits that are used for source and destination ID

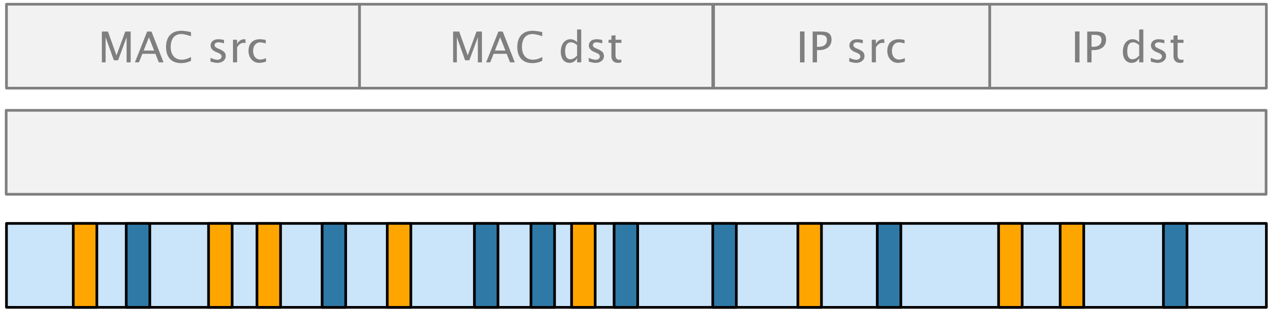

iTAP hybrid obfuscation scheme

Map source and destination to IDs

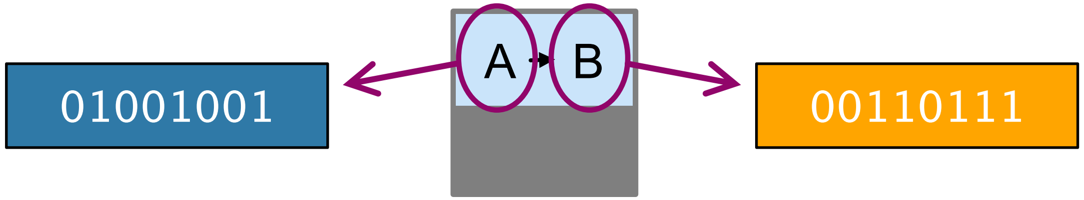

Match-fields with arbitrary bitmasks

Interpret as bit-string of 160 bits

Randomly select bits that are used for source and destination ID

Add source and destination ID

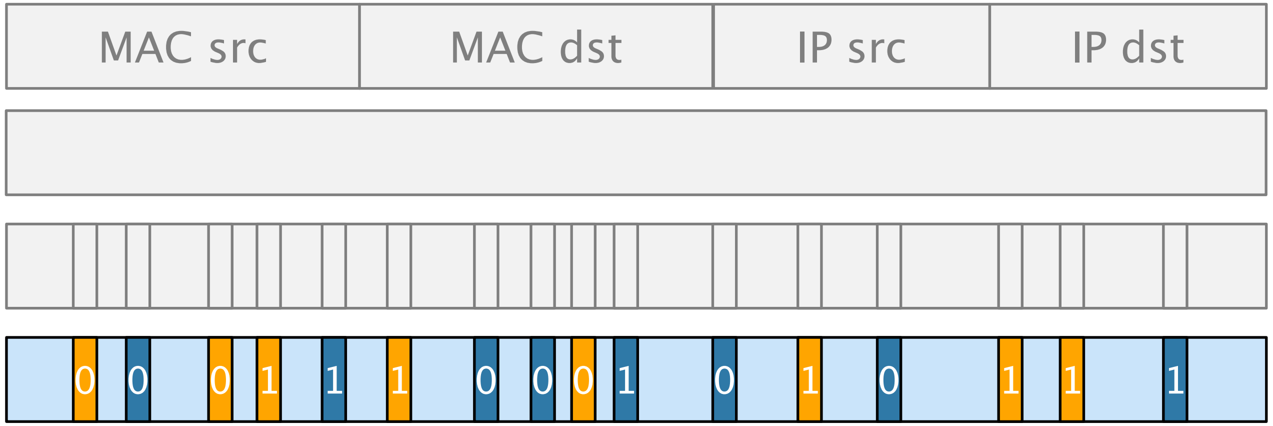

iTAP hybrid obfuscation scheme

Map source and destination to IDs

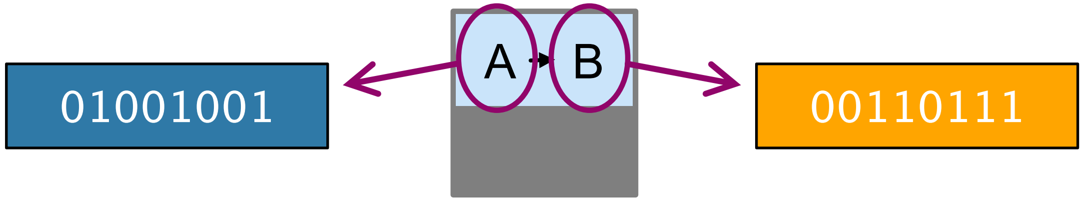

Match-fields with arbitrary bitmasks

Interpret as bit-string of 160 bits

Randomly select bits that are used for source and destination ID

Add source and destination ID

Set other bits to random values

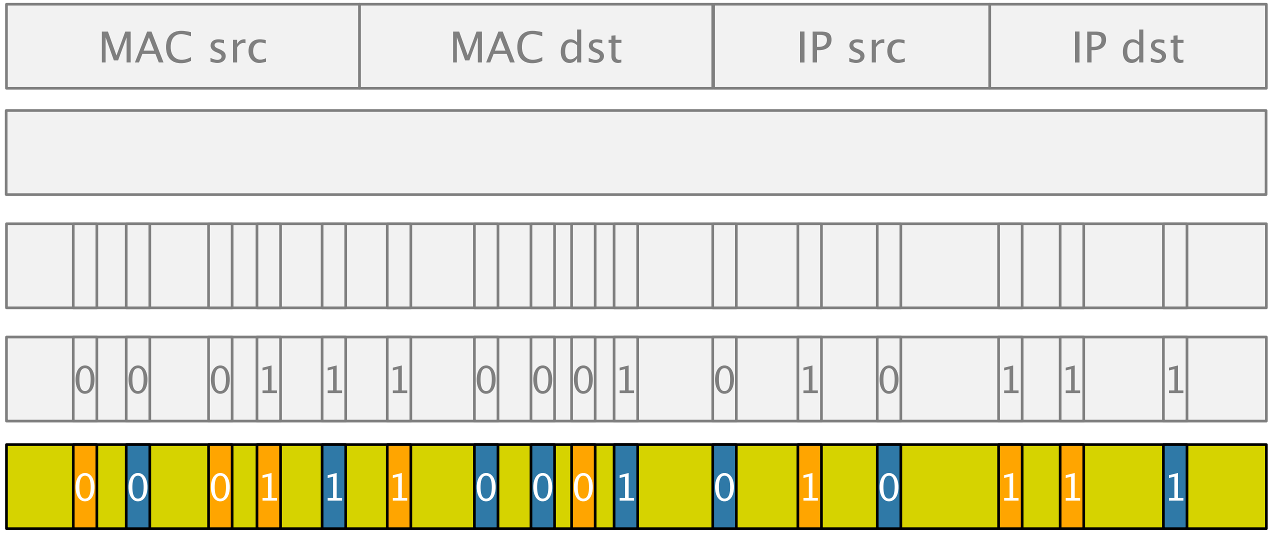

iTAP hybrid obfuscation scheme

Map source and destination to IDs

Match-fields with arbitrary bitmasks

Interpret as bit-string of 160 bits

Randomly select bits that are used for source and destination ID

Add source and destination ID

Set other bits to random values

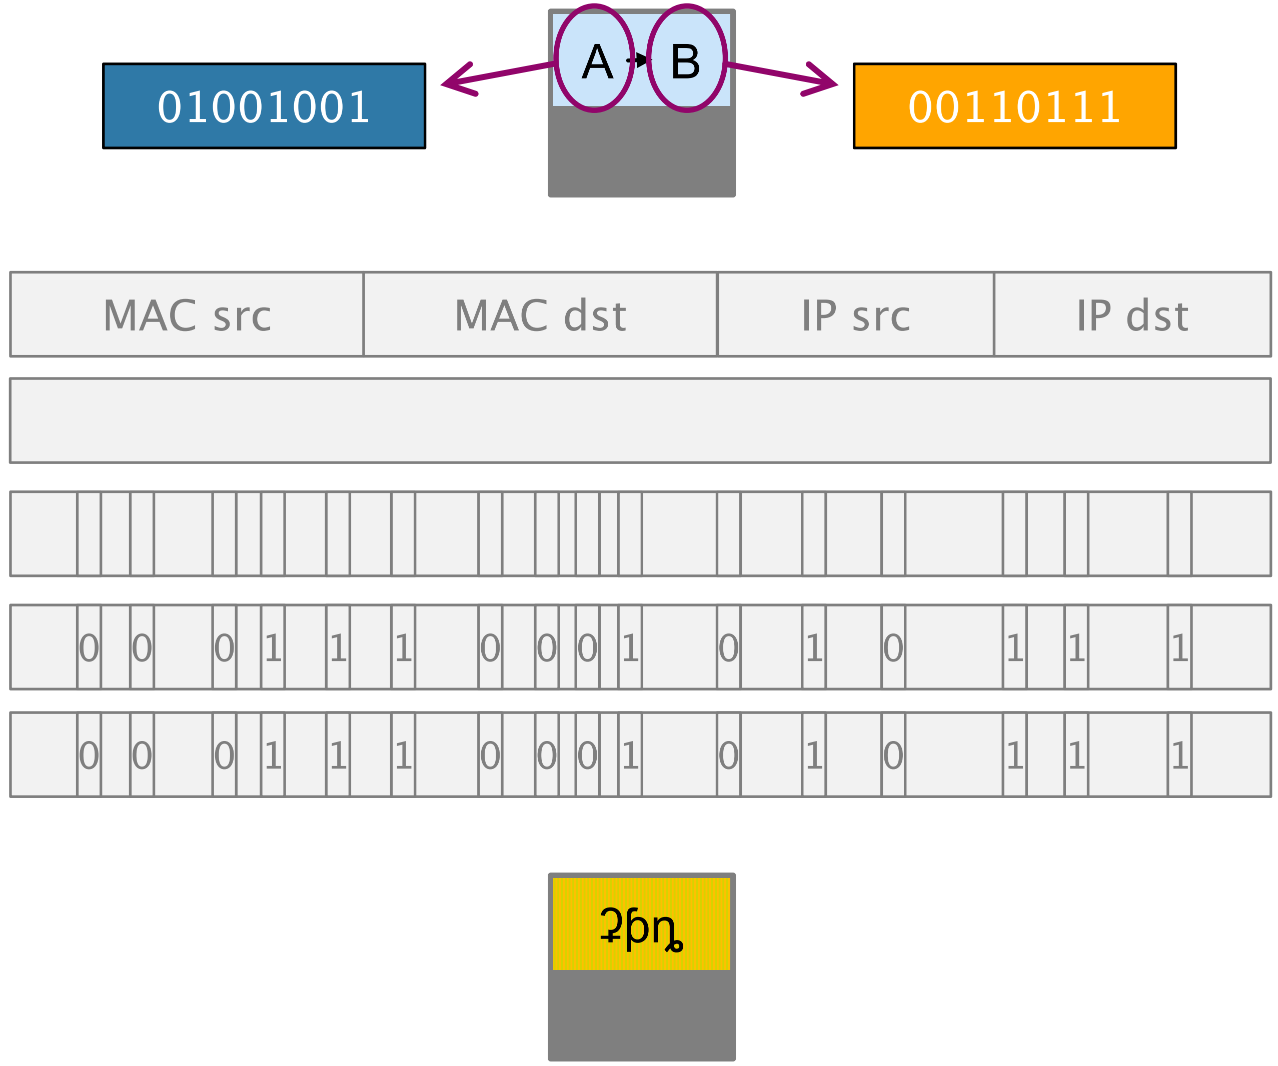

iTAP hybrid obfuscation scheme

Forwarding based on the destination ID

→ good scalability

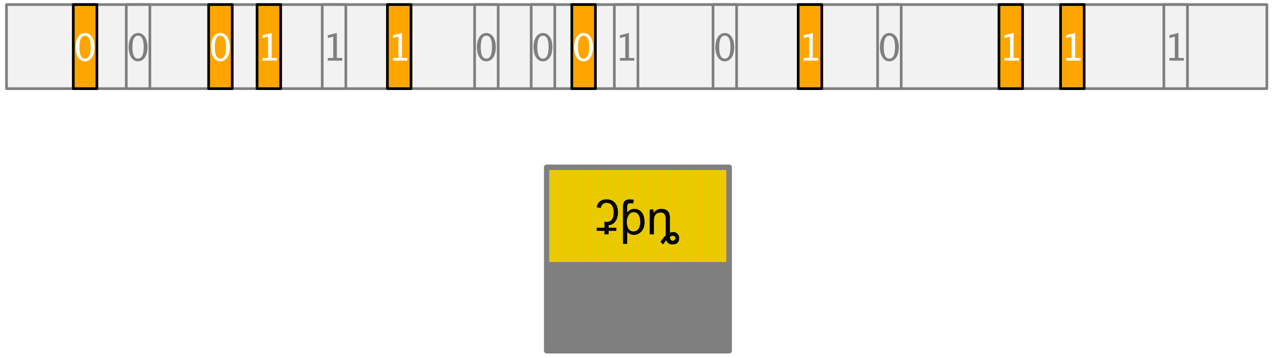

iTAP hybrid obfuscation scheme

Eavesdropper cannot distinguish between random and non-random bits

→ good anonymity

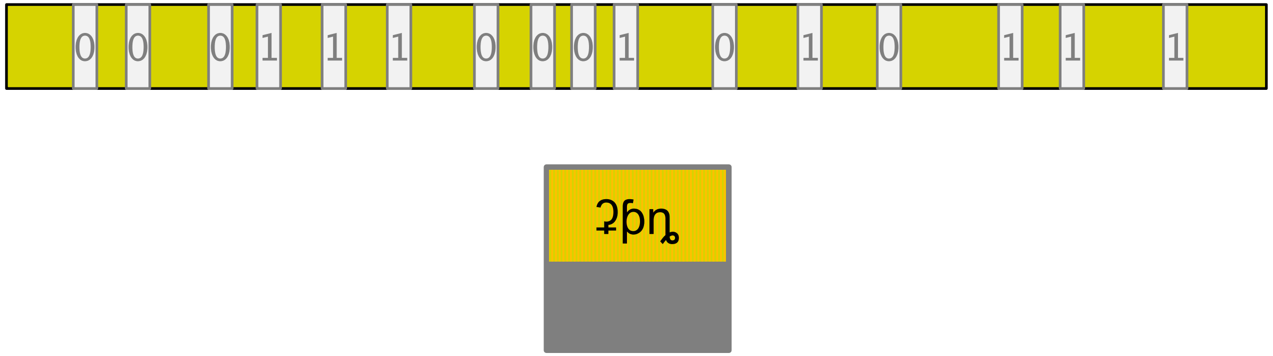

What if an attacker analyzes multiple flows?

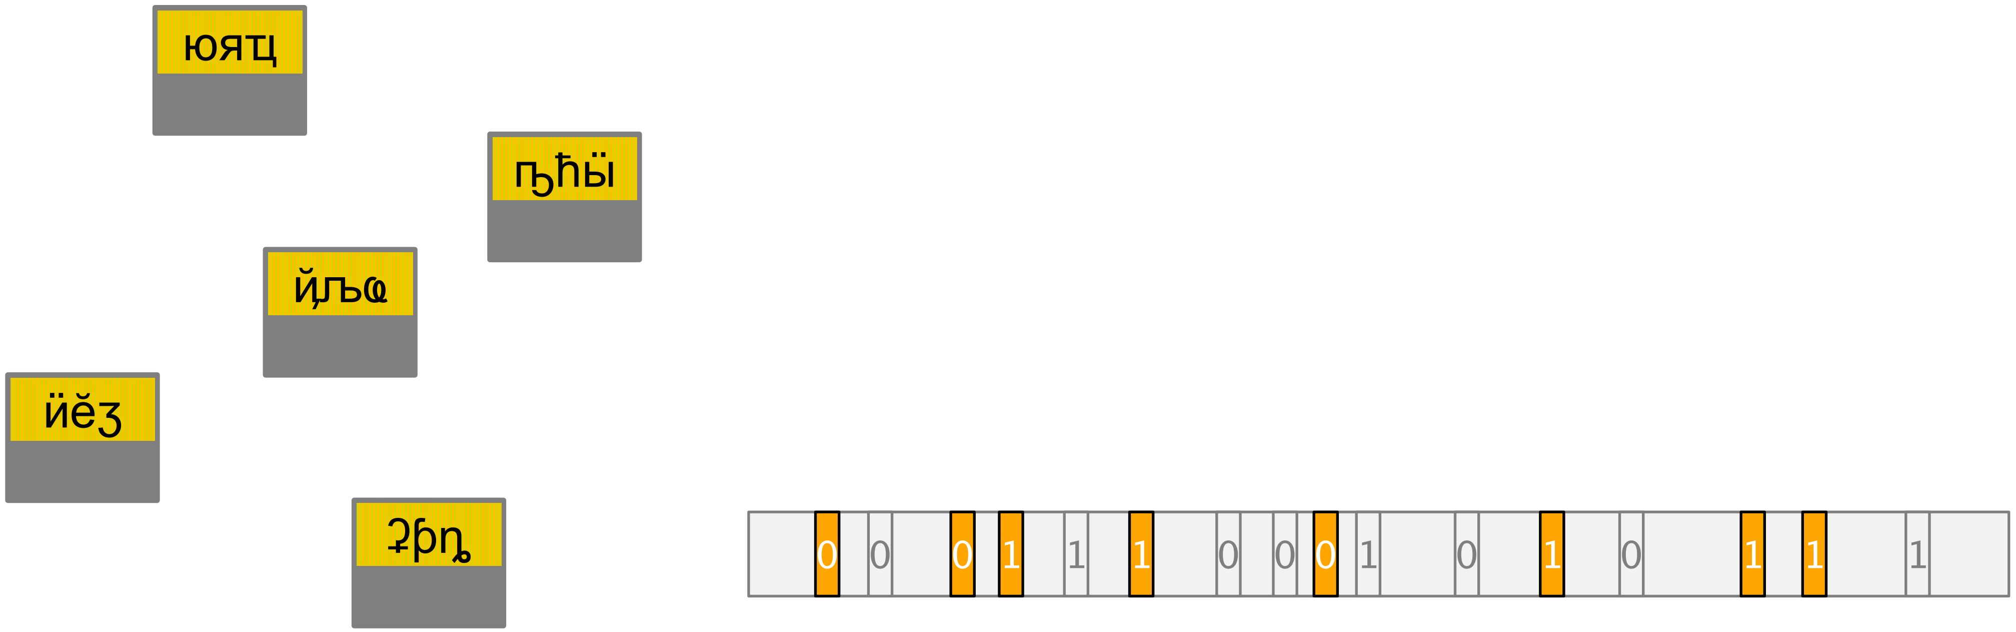

What if an attacker analyzes multiple flows?

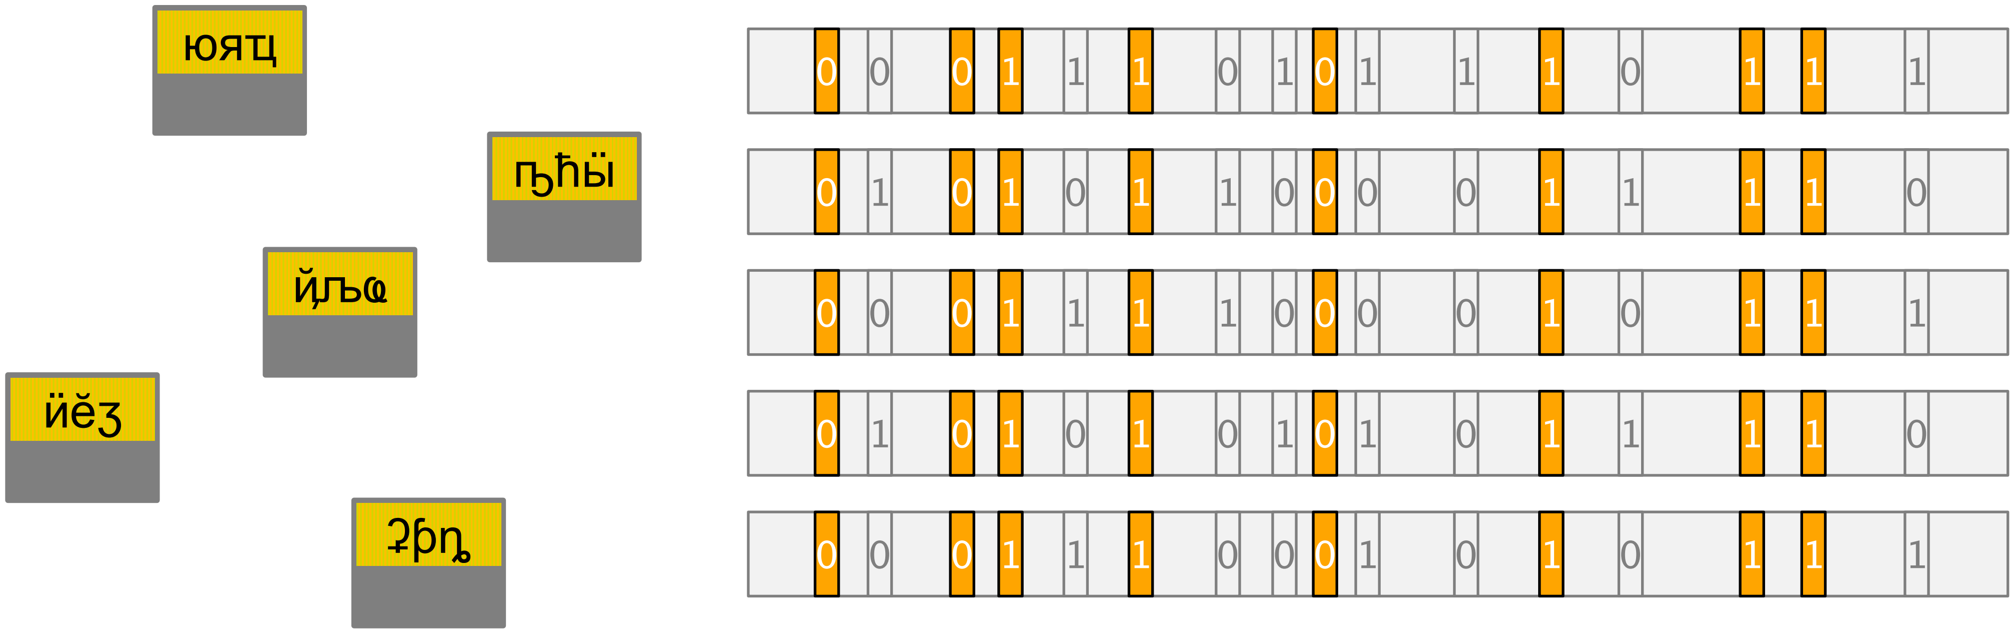

What if an attacker analyzes multiple flows?

iTAP controls information leakage and proactively adapts the encoding

The controller monitors the observed entropy for each link…

… and changes the encoding before an eavesdropper is able to break it.

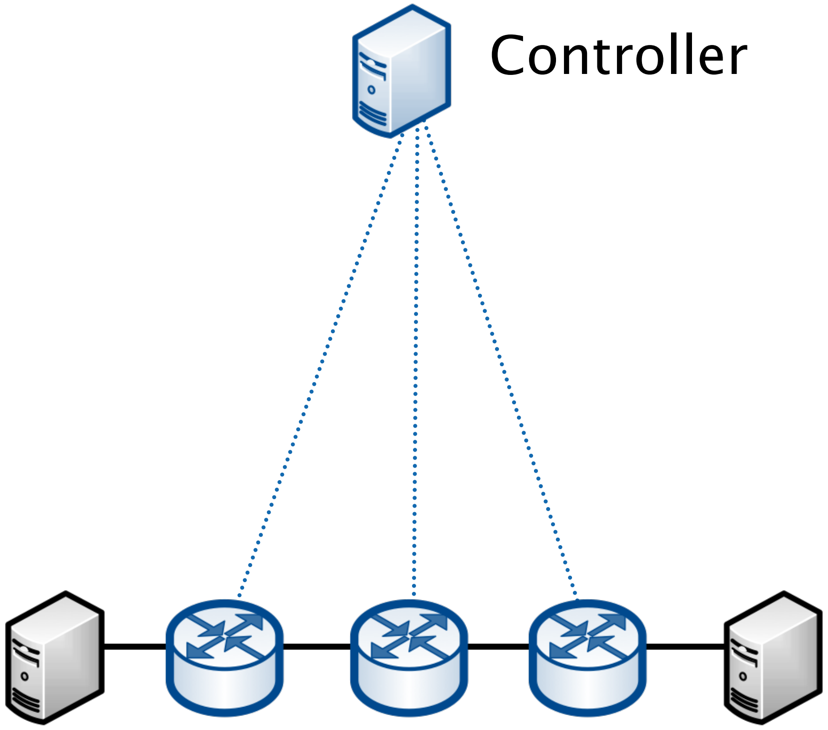

iTAP controls information leakage and proactively adapts the encoding

The controller monitors the observed entropy for each link…

… and changes the encoding before an eavesdropper is able to break it.*

* According to the Unicity Distance

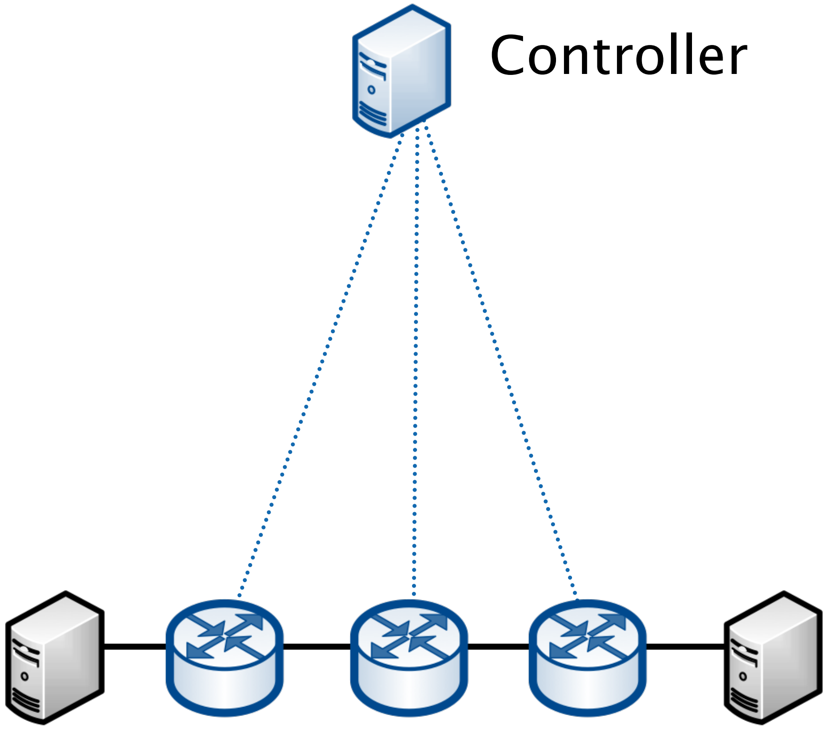

iTAP controls information leakage and proactively adapts the encoding

iTAP controls information leakage and proactively adapts the encoding

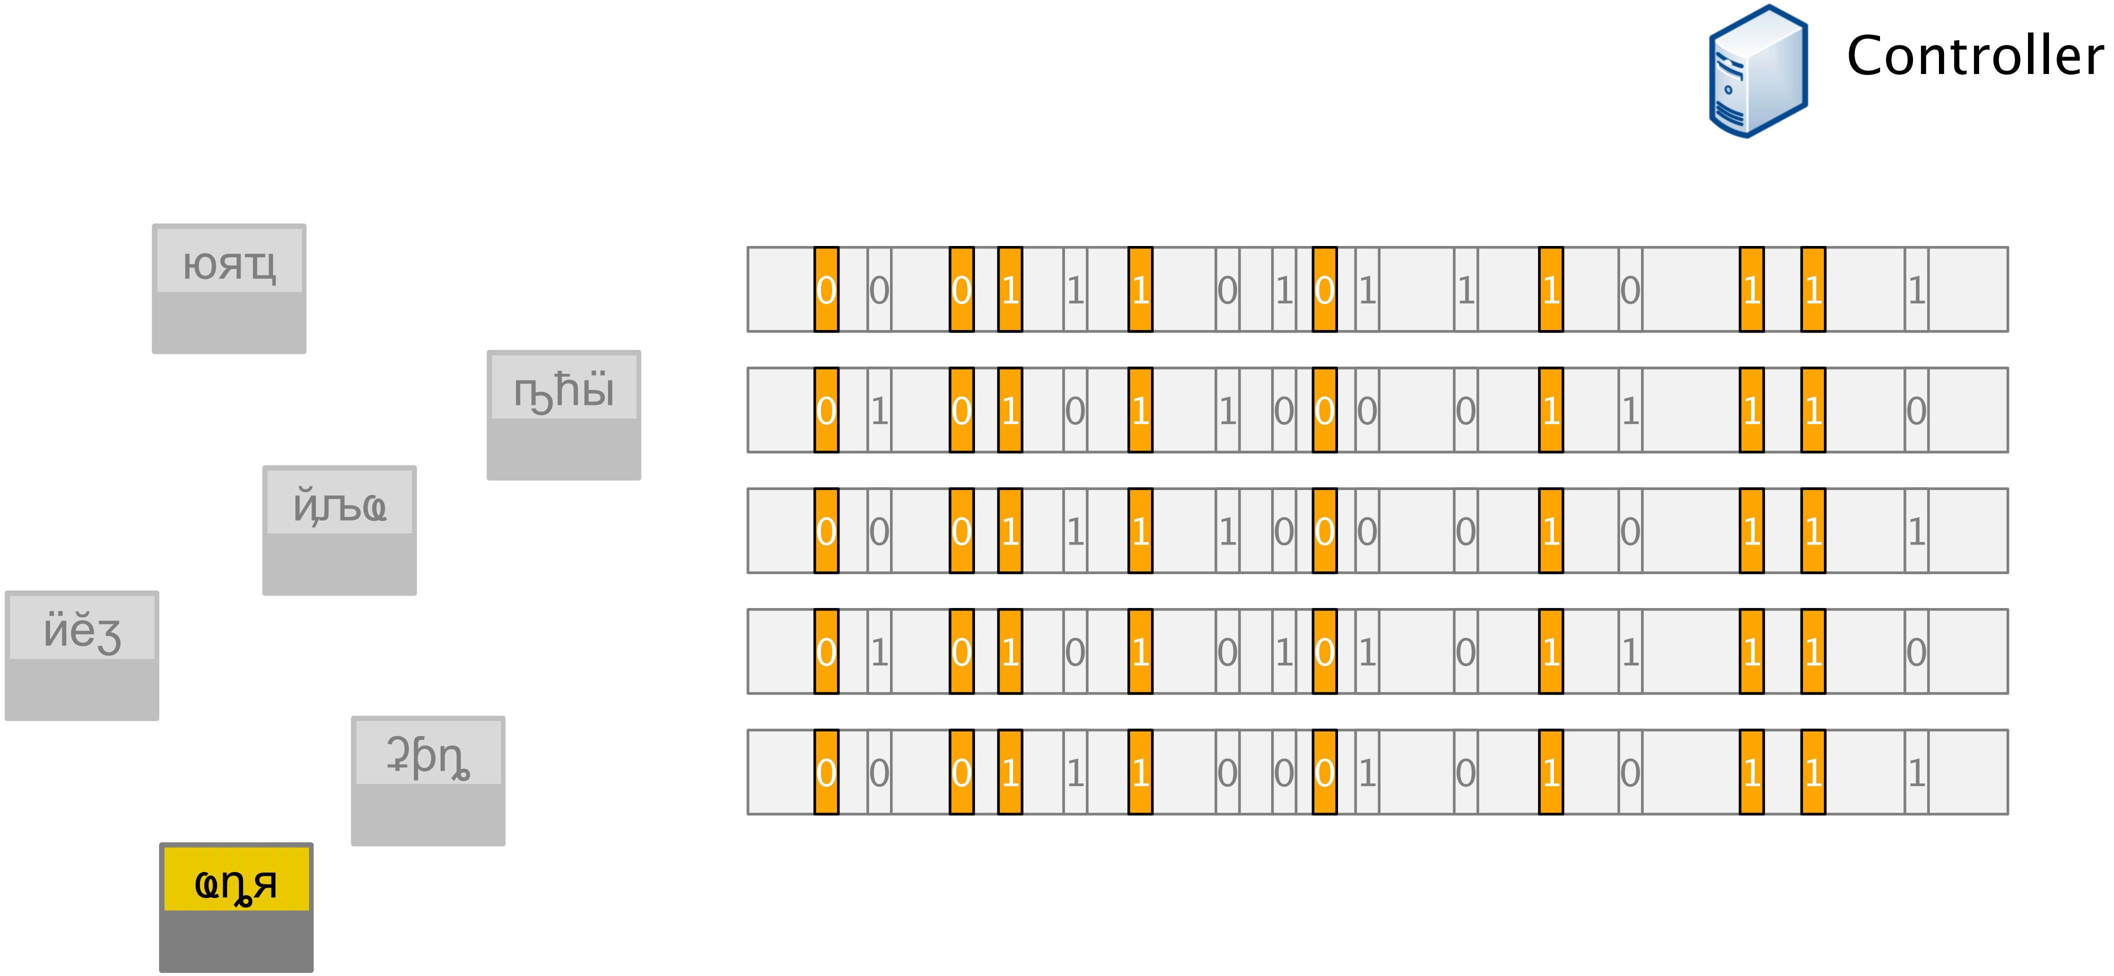

iTAP controls information leakage and proactively adapts the encoding

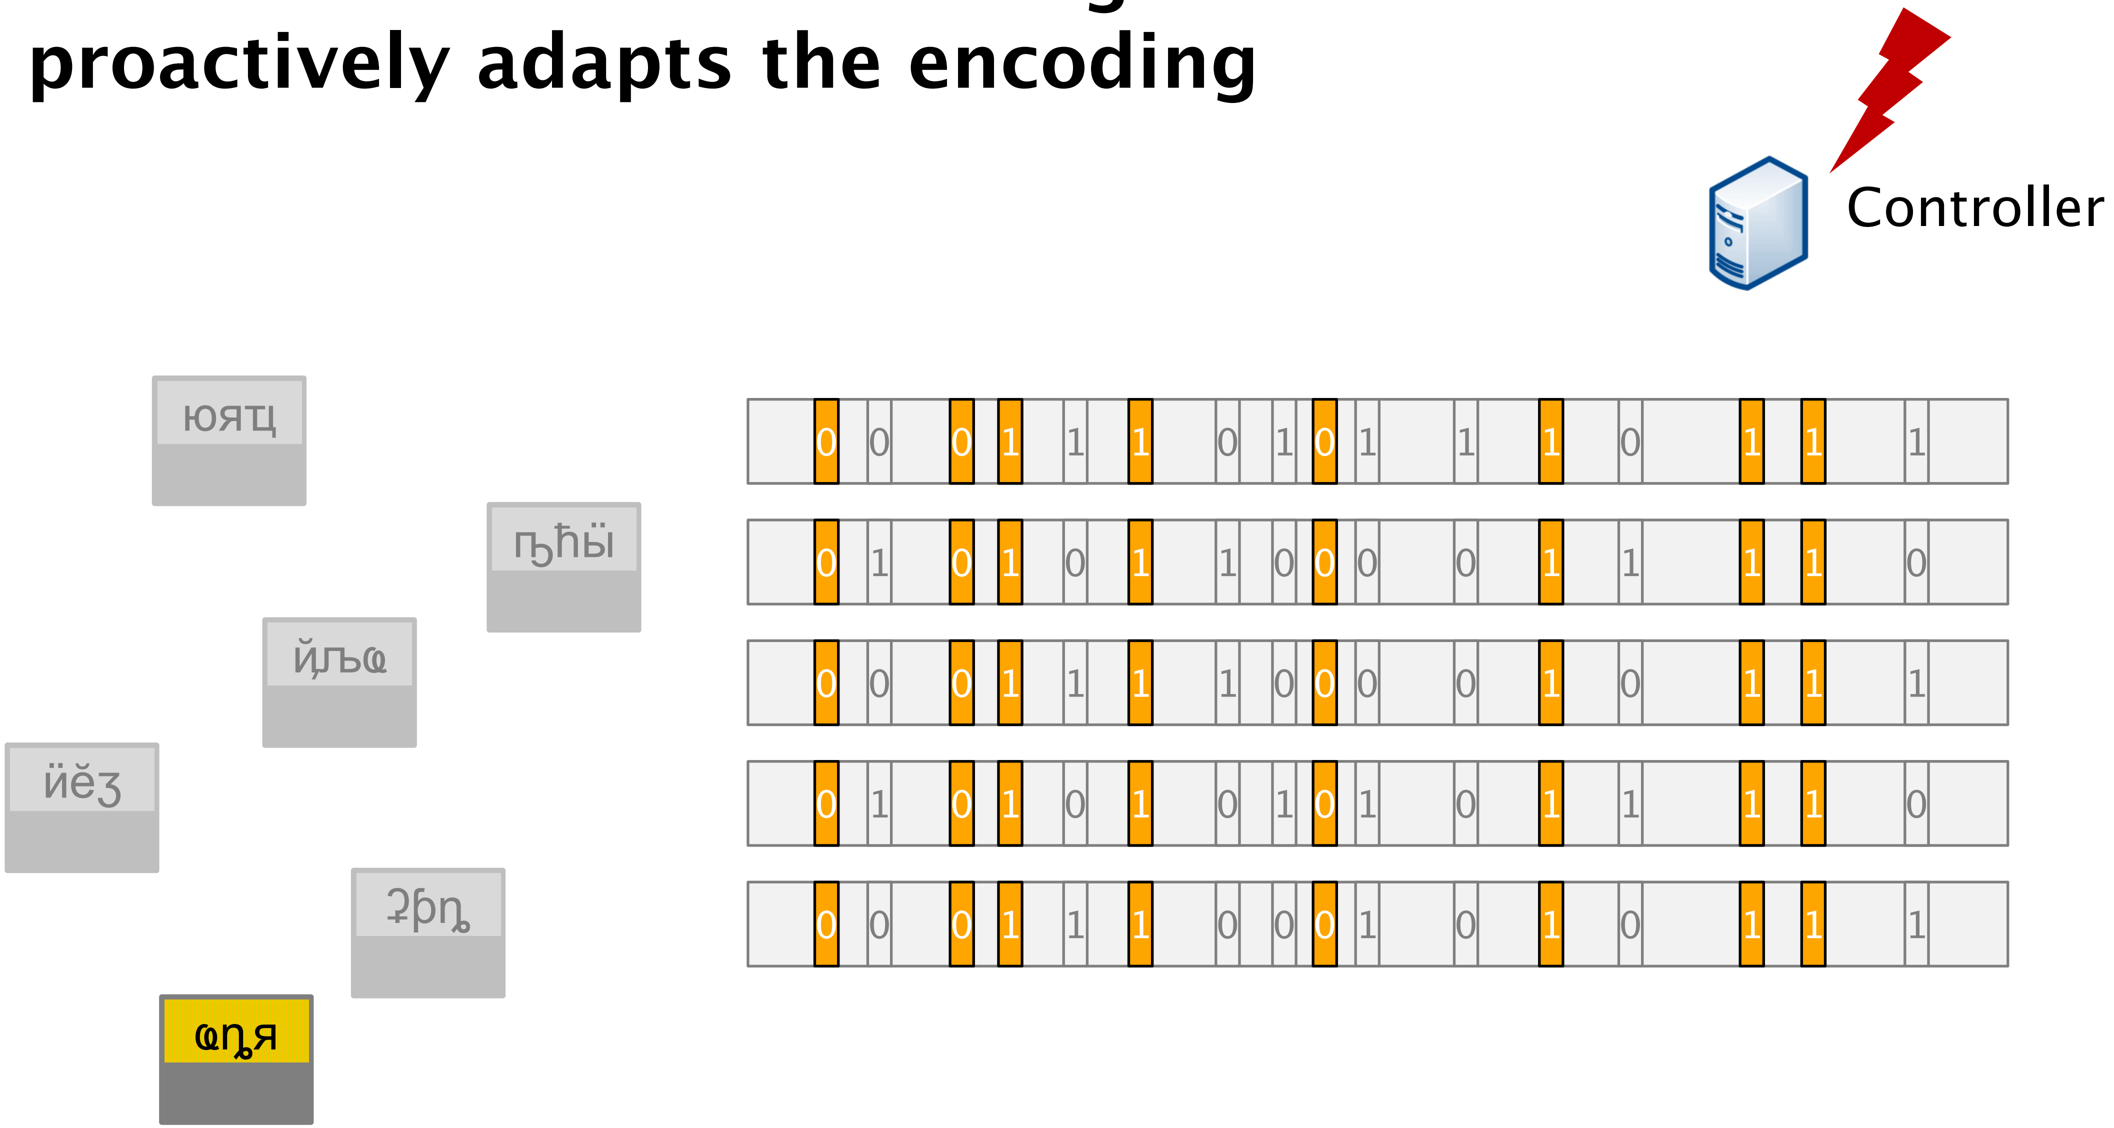

iTAP controls information leakage and proactively adapts the encoding

iTAP evaluation based on real network traffic

iTAP evaluation based on

7 days	of network traffic
400	hosts
128 M	flows

iTAP evaluation based on real network traffic

7 days of network traffic | 400 hosts | 128 M flows

Indicators: controller actions / s

flow table updates / s

forwarding rules

iTAP works in practice

7 days of network traffic   400 hosts   128 M flows

avg	max	
200	700	controller actions / s
50	250	flow table updates / s
600	2.5 k	forwarding rules

Only a small share of SDN switches is sufficient to protect a large share of the network traffic

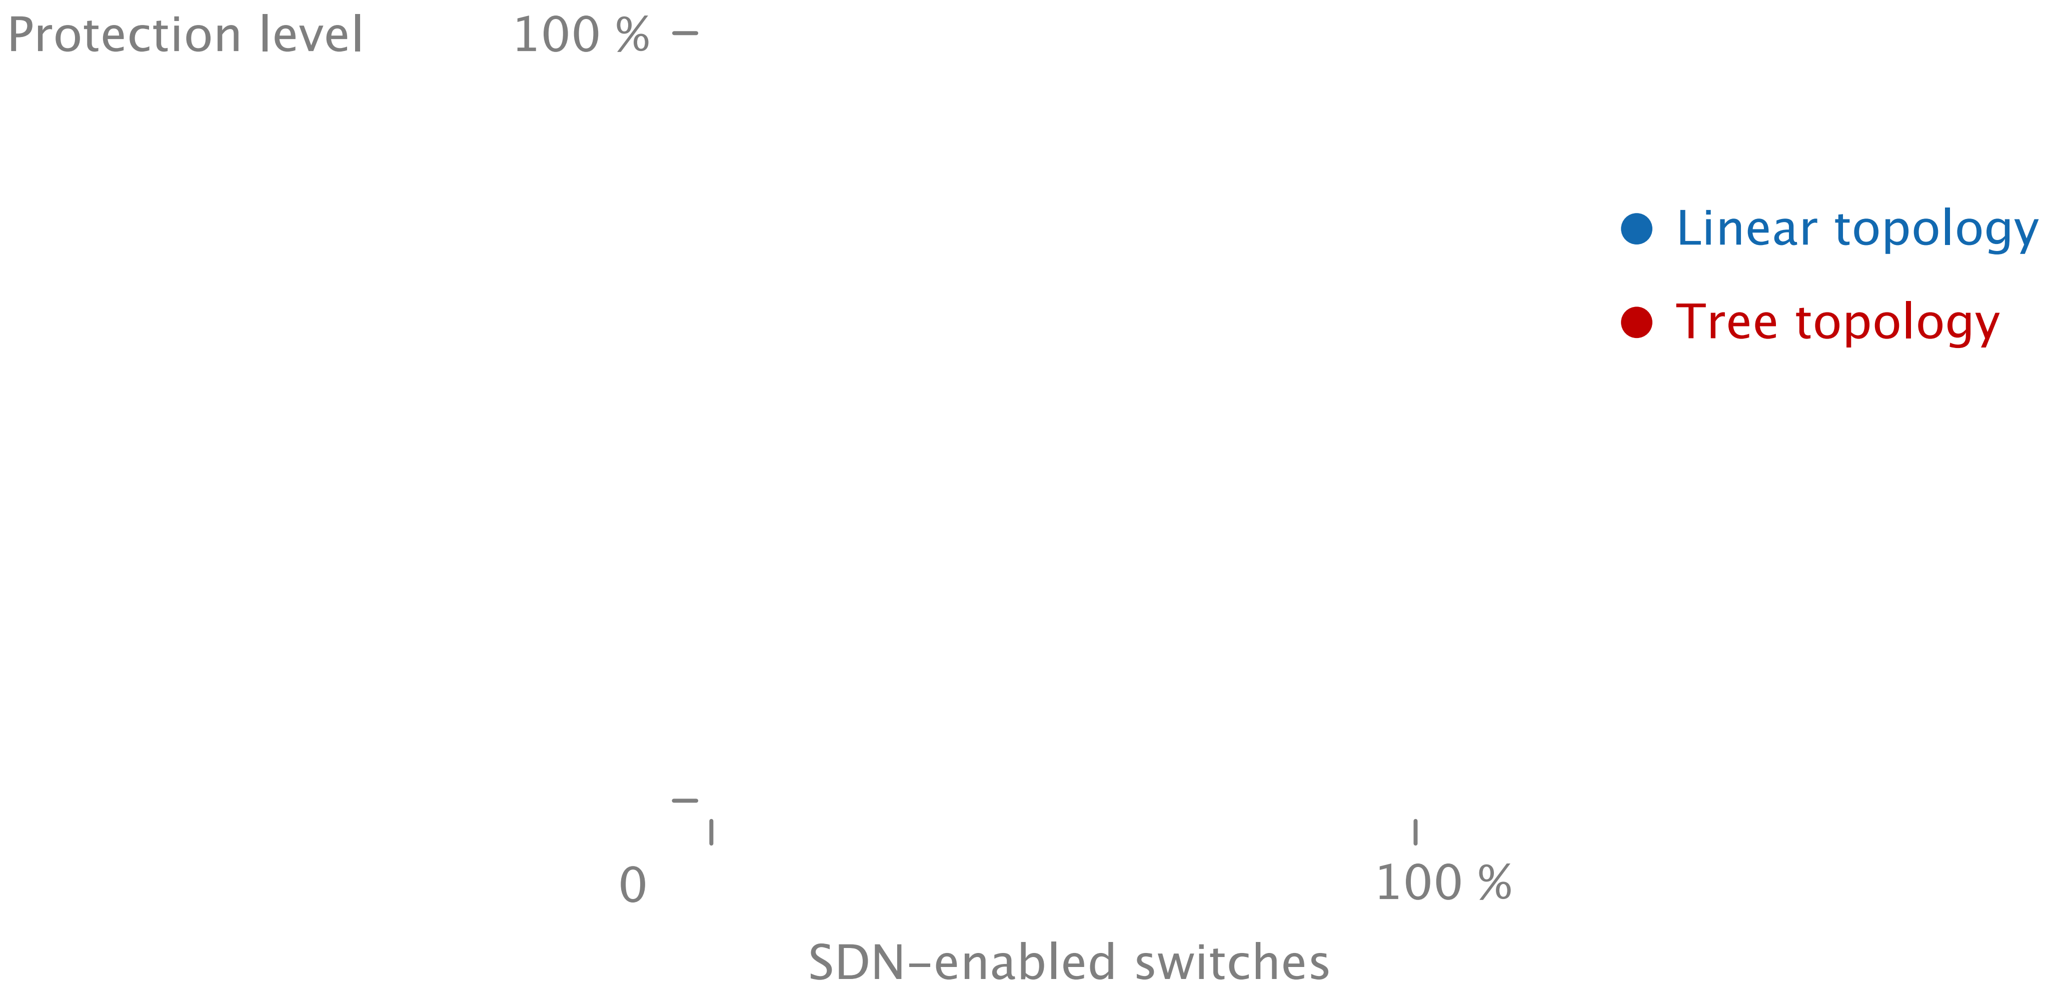

Only a small share of SDN switches is sufficient to protect a large share of the network traffic

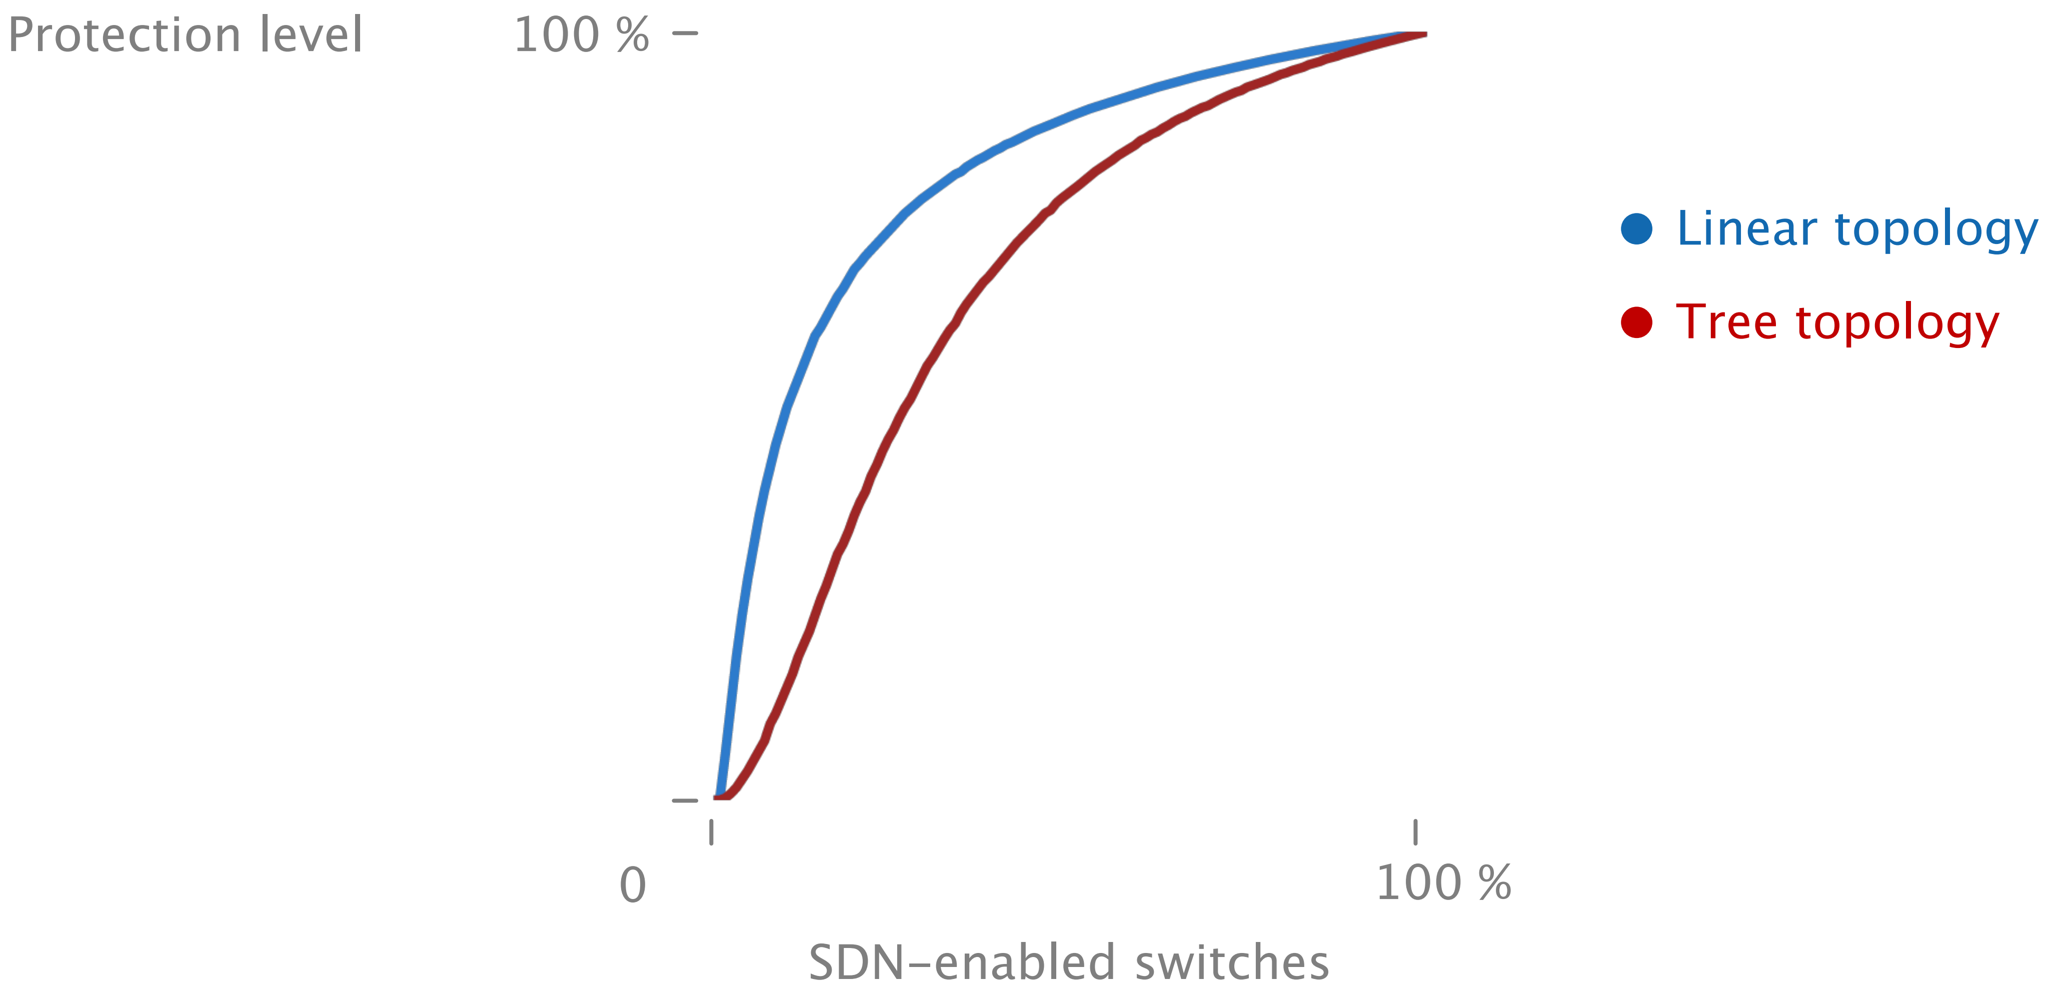

Only a small share of SDN switches is sufficient to protect a large share of the network traffic

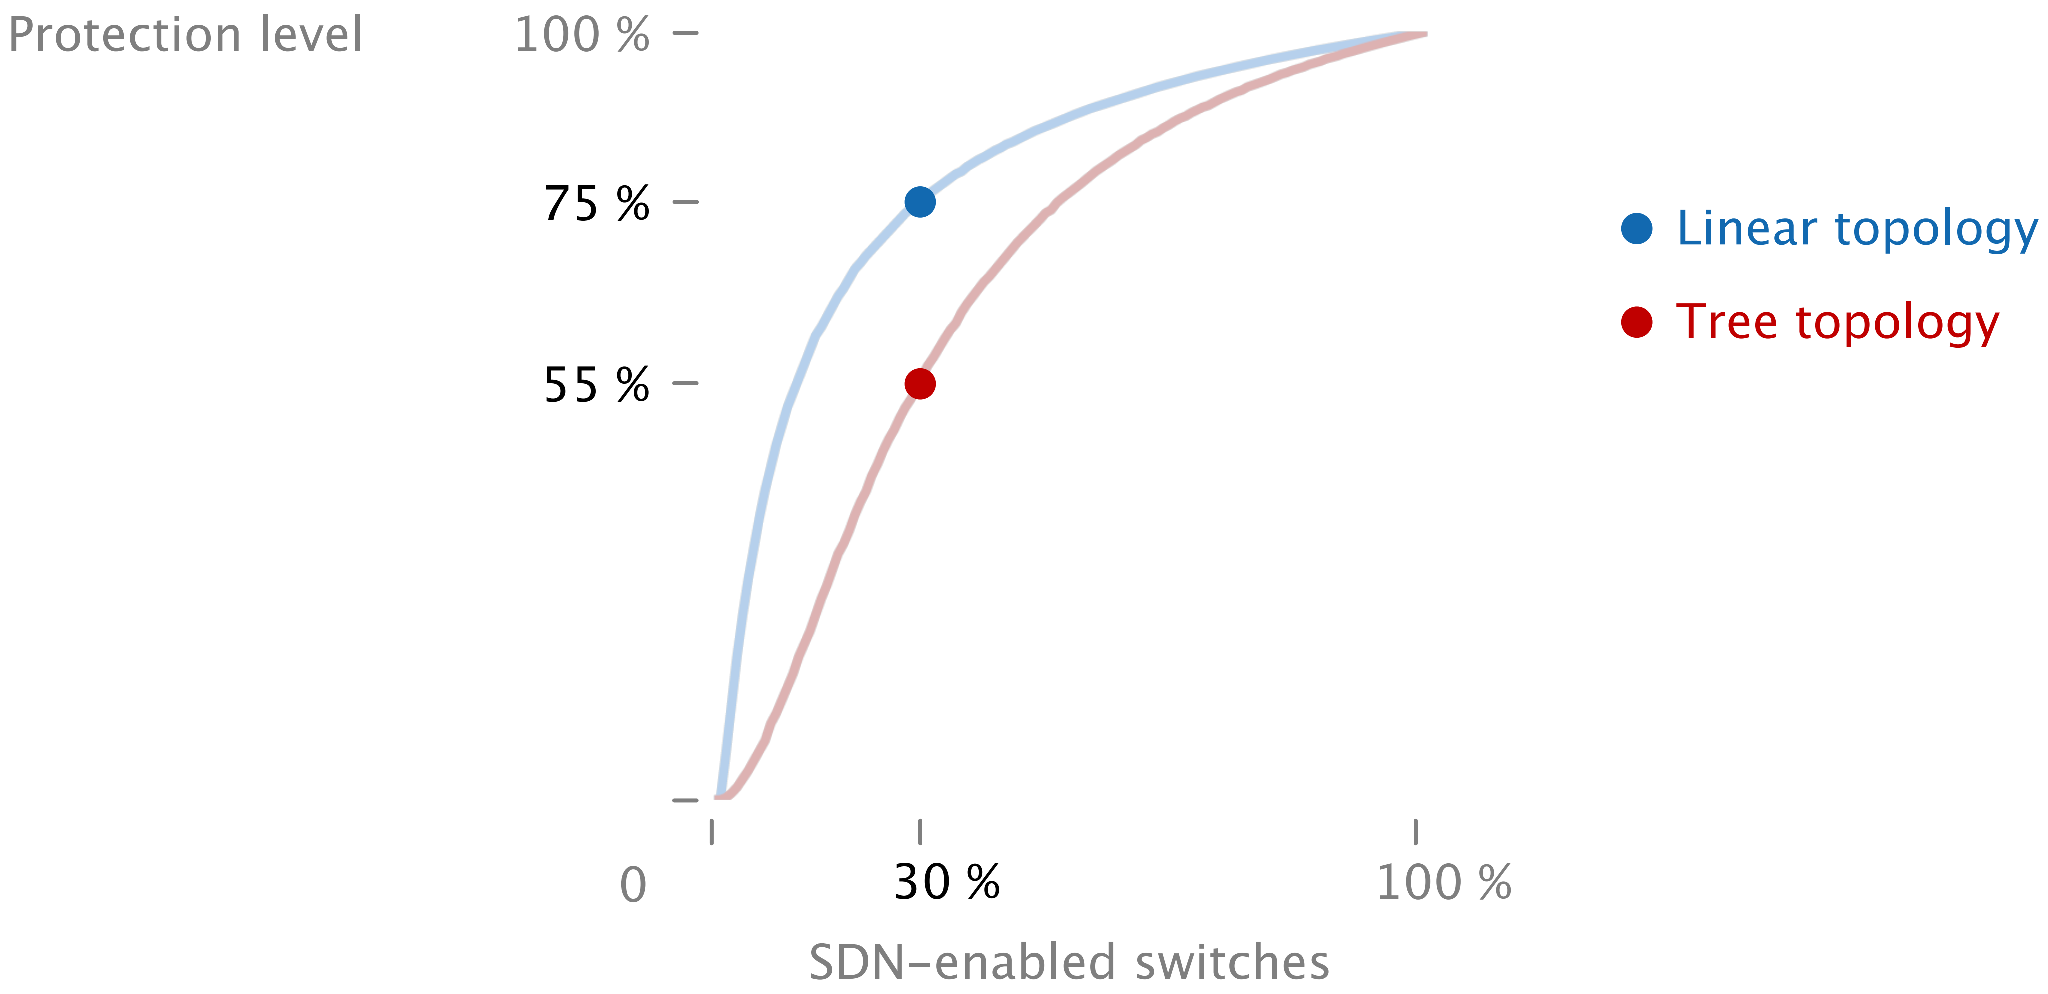

Contributions

https://itap.ethz.ch

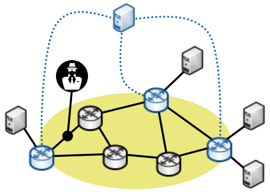

iTAP design

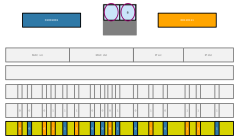

Scalable & anonymity-providing header rewriting scheme

iTAP prototype implementation

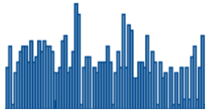

Evaluation based on real user traffic