Authenticating Internet Routing Using Zero-Knowledge Proofs

Jacob D. White
A-4: Advanced Research in Cyber Systems
Email: `jdwhite@lanl.gov`

Mentor: Michael Dixon
Co-Mentor: Boris Gelfand
Collaborator: Zachary DeStefano

August 9th, 2023

LA-UR-23-29806

Internet Routing: The Sky-High View

The **Internet** is a complex “network of networks”, allowing computers to **route** messages to each other across the globe.

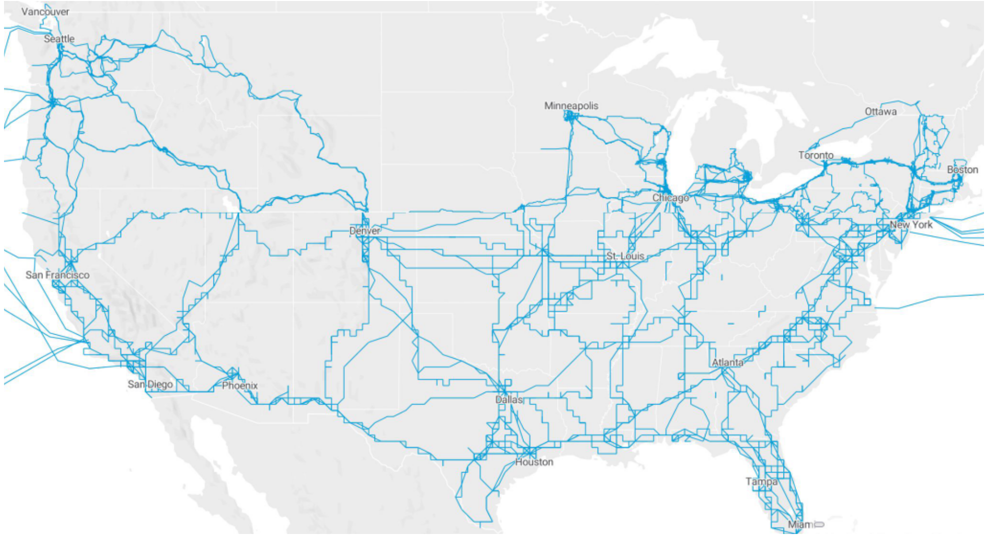

Figure: US high-speed fiber optic connections (Lumen 2023)

Internet Routing: The Sky-High View

- Computers send data across **routers**
- Organizations form a **network** of routers
- Routers use **policies** and **protocols** to find and communicate with each other on the Internet

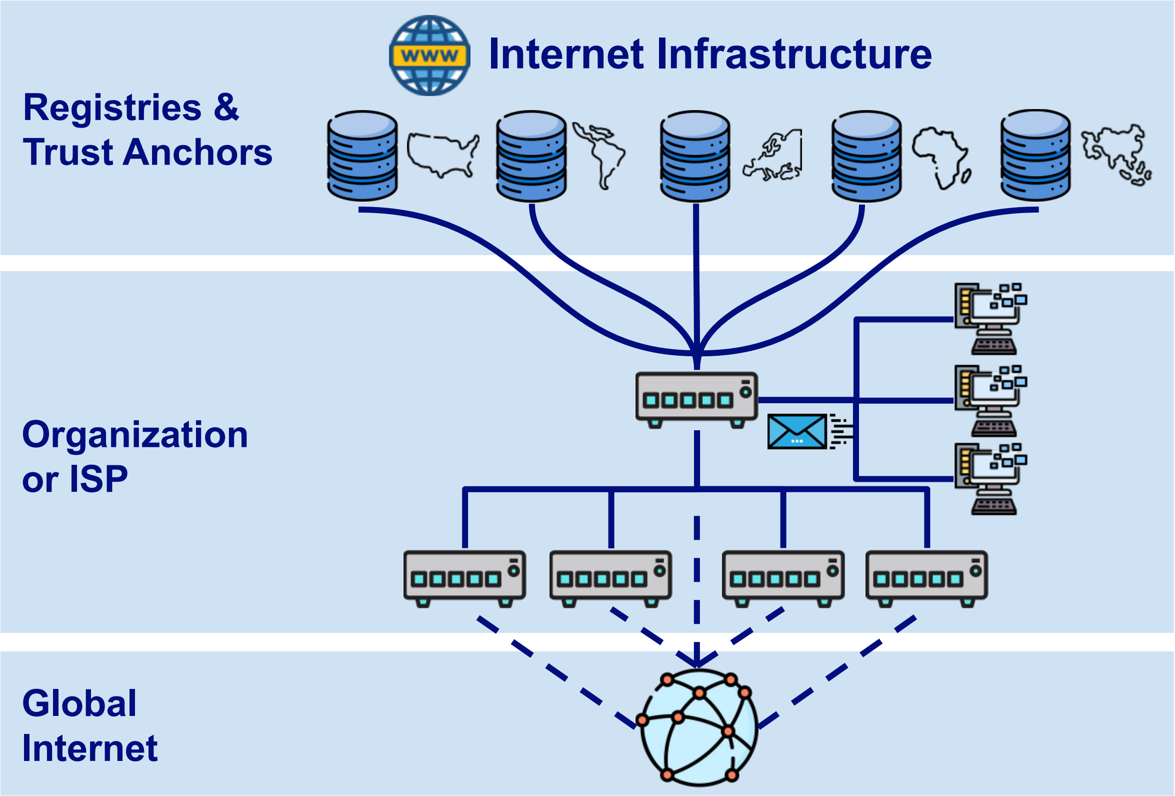

Internet Routing for ISPs and Organizations

An organization manages an **Autonomous System** (AS) or domain of routers. Routing policies are defined internal or external to the domain.

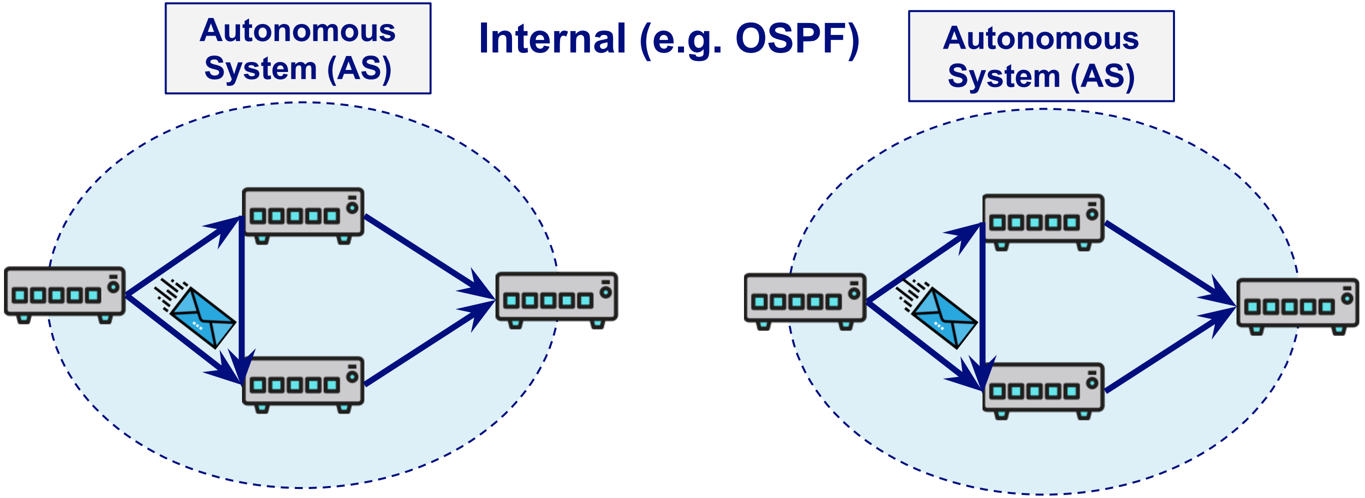

Internet Routing for ISPs and Organizations

An orga… of
routers… in.

Problem: Routers who *lie* about how they route network data can cause serious disruptions and privacy issues!

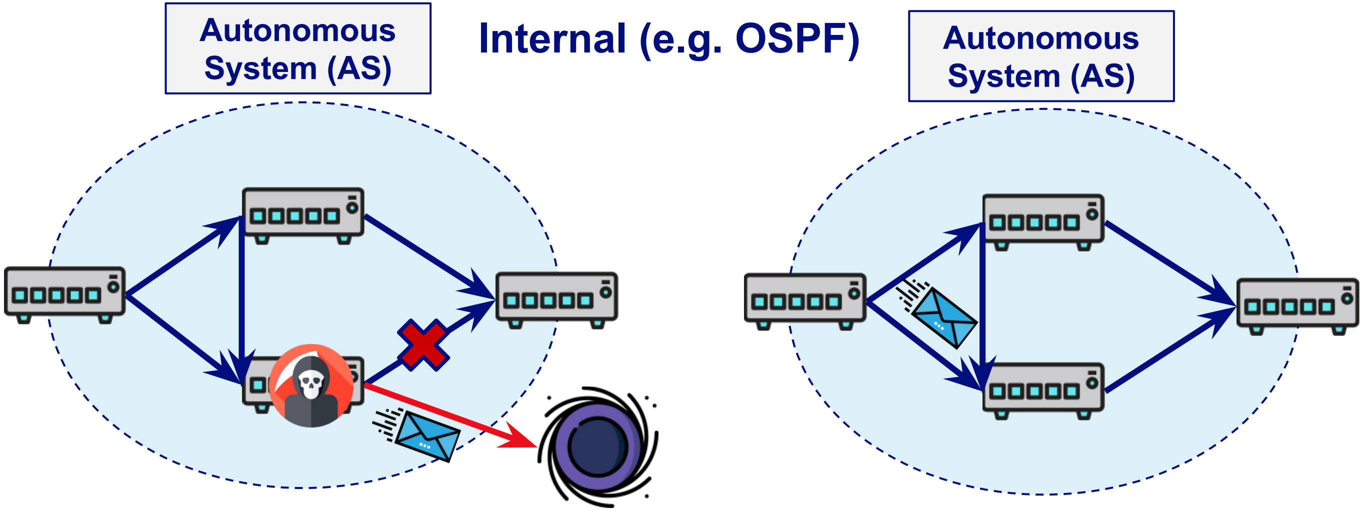

Traffic Diversion (Blackholing)

Internet Routing for ISPs and Organizations

An orga … of
routers … in.

Problem: Routers who *lie* about how they route network data can cause serious disruptions and privacy issues!

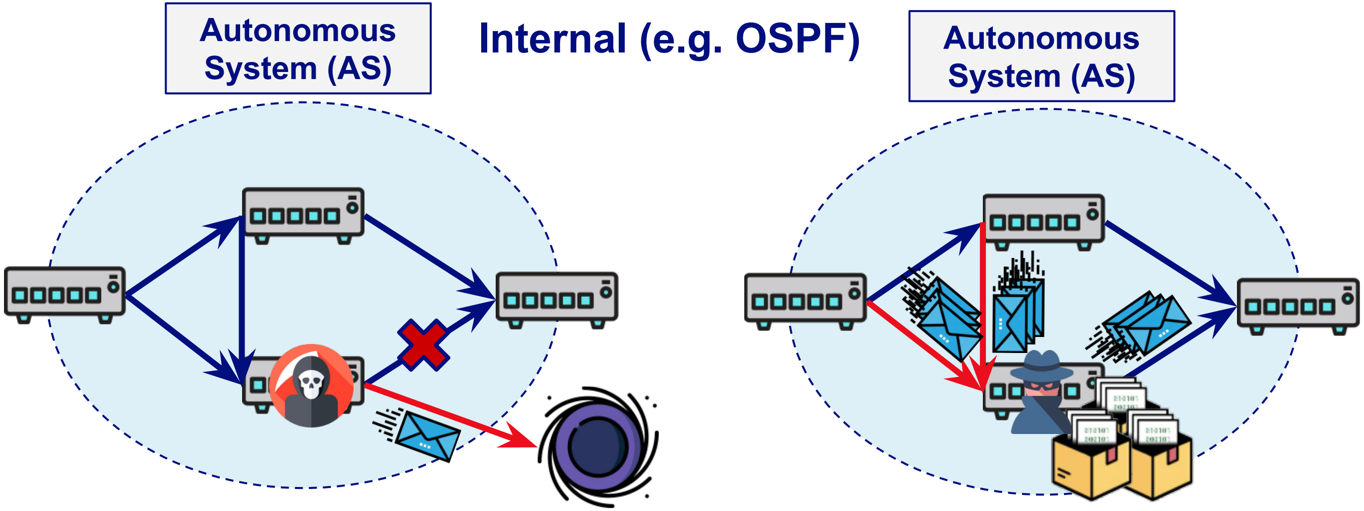

Traffic Diversion (Blackholing)

Interception Attack

Internet Routing for ISPs and Organizations

An organization manages an **Autonomous System** (AS) or domain of routers. Routing policies are defined internal or external to the domain.

External (BGP)

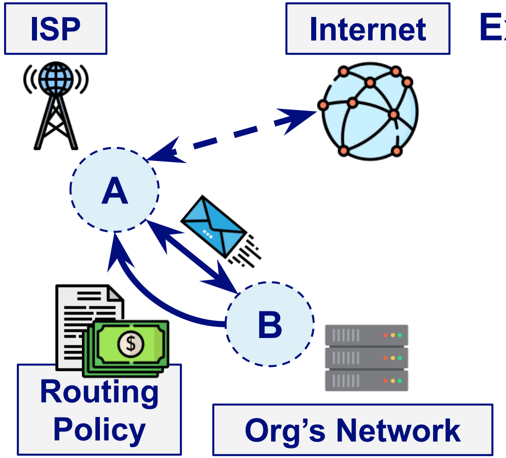

Transit (Customer-Provider)

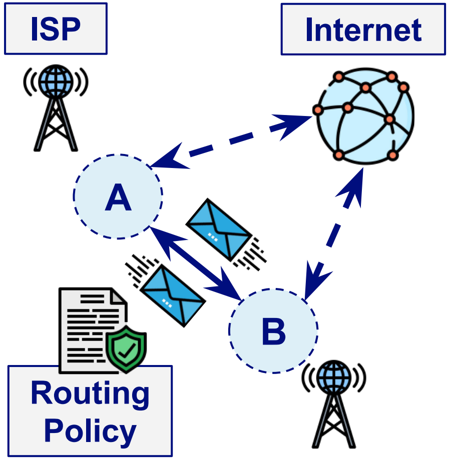

Peers

Internet Routing for ISPs and Organizations

An orga … of
routers … in.

Problem: Routers who *lie* about their routing network data can cause serious disruptions and privacy issues!

External (BGP)

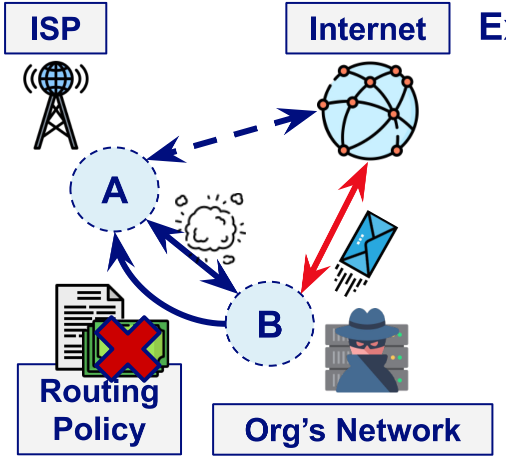

Transit (Customer-Provider)

Policy Violation

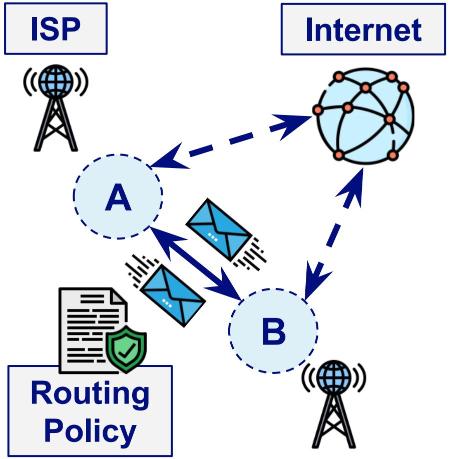

Peers

Internet Routing for ISPs and Organizations

An orga **Problem:** Routers who *lie* about their routing network data of
routers can cause serious disruptions and privacy issues! in.

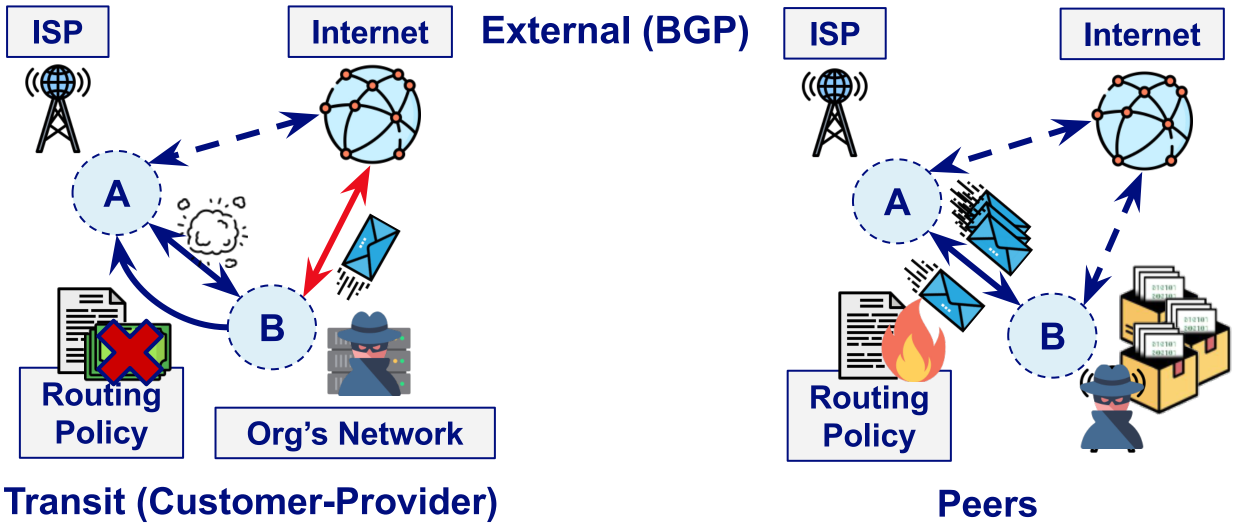

Policy Violation

Interception Attack

Internet Routing for ISPs and Organizations

Key Questions:

1. How can we **authenticate network operations**, having routers learn from the global network, and behave according to policy?

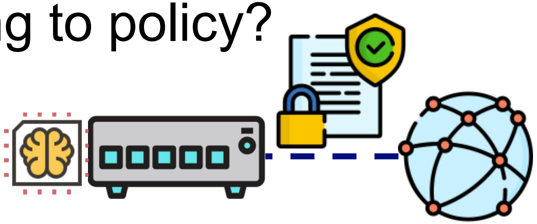

2. How can we **avoid leaking private information** about organizations' networks and relationships?

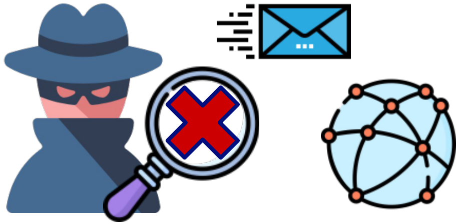

(Non-interactive) Zero-knowledge Proofs

Zero-knowledge proofs (ZKPs) allow us to prove that a claim **IS** true without revealing **WHY** it is true, even if the prover is considered untrusted and **malicious**.

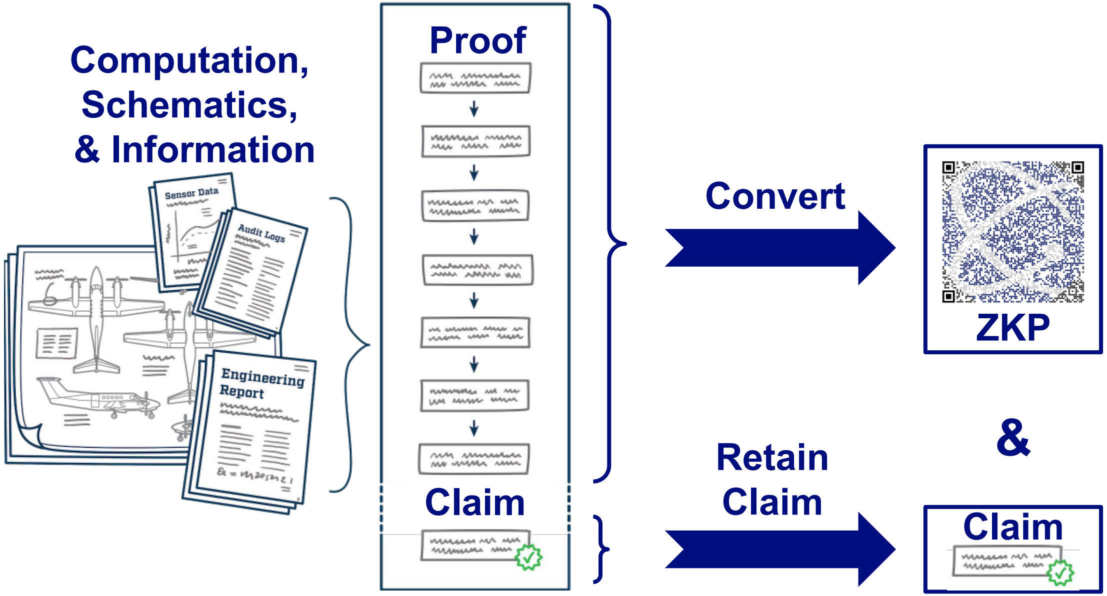

Example Claim.
"The packet can reach Router Y from X, even if Router Z goes offline"

Features of (Non-interactive) Zero-knowledge Proofs

Zero-knowledge proofs (ZKPs) allow us to prove that a claim **IS** true without revealing **WHY** it is true, even if the prover is considered untrusted and **malicious**.

Miniscule Footprint

~3k Bits

Some ZKP variants are tiny, often only slightly larger than a regular QR code

Fine-grained Control

ZKPs give fine-grained control over secret information, yet allows trustless verification

Composable

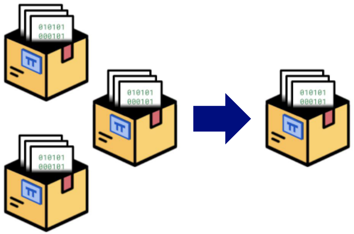

ZKPs can be collected and combined into new ZKPs without growing in size

Expanded Trust

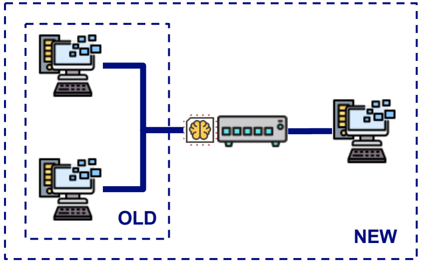

Portable proofs extend our trusted view beyond that of our own system

ZKPNet: An Overview

Developed a prototype Rust library which provides useful gadgets for authenticating network routing attestations using ZKPs

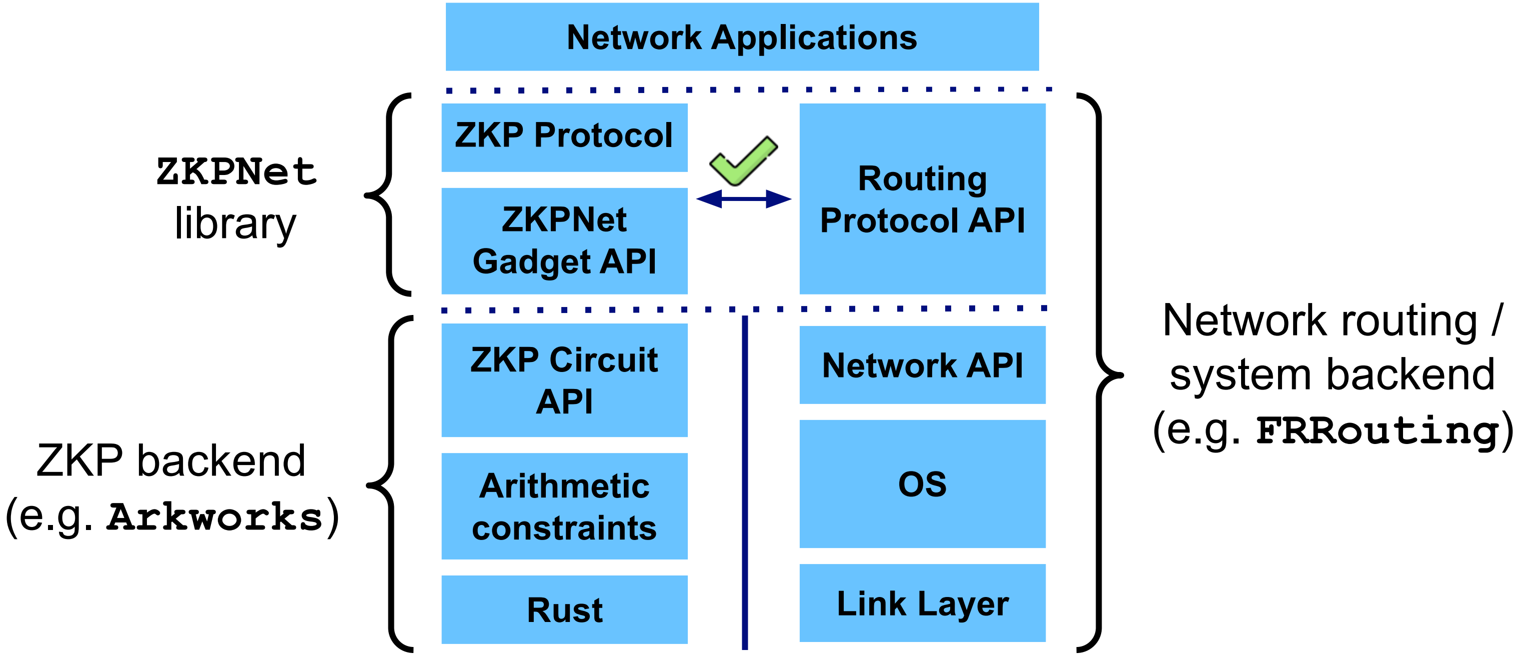

Demo: Proving Route Reachability

Group A wants to send important data to Group C, but will need to go through Group B first. A and C first want to verify that B can deliver the data, but Group B is unwilling to reveal details about the network for security reasons. *How does Bob prove this?*

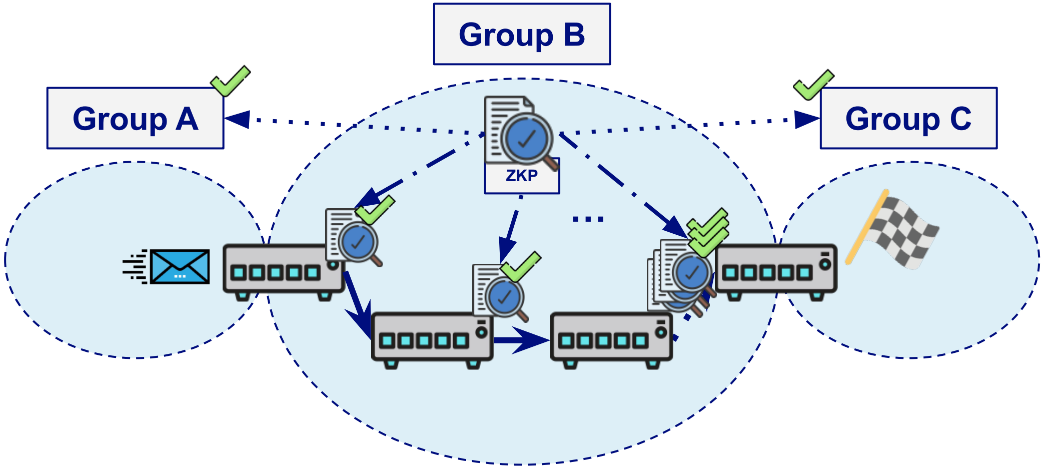

Results: ZKPNet Demo Benchmarks

Using realistic OSPF entries for internal routing, we have constructed zero-knowledge proof for route reachability for a single hop. Benchmarks were performed on a Apple M1 Max CPU with 32 GB of memory.

ZKP Technique	# of constraints	Proof Size (Bandwidth)	Proving Time (Latency / Delay)	Verification Time (Latency)
Single Proof	104	224 B*	468.03 ms	2.7165 ms
Depth-2 Recursion on Proof	13976	299 B*	TBD	TBD

*Estimated from Groth16 proof sizes with MNT4&6 curves

Looking Ahead: Feature Support for Routing Auth.

	ZKPNet	BGPSec (best auth)	S-BGP	so-BGP	RPKI (deployed)	IRR (worst auth)
Network Route Integrity	✓	✓	✓	⚠	⚠	✗
Comm/Bandwidth Efficiency	?	✗	✗	⚠	⚠	✓
Dynamic/Adaptive Recovery	?	✗	✗	✗	✓	✓
Trustless Authentication	✓	✗	✗	✗	✗	✗
Privacy Preservation	✓	✗	✗	✗	✗	✗

Legend	
✓	Complete
⚠	Partial
✗	Missing/Bad
?	Varies

Future Work: ZKP Compiler for Verifiable Routing

```
as-num: 64496
import: {
      from AS64497 at 192.0.2.1
            action pref=0;
            accept community
                  .contains(GRACEFUL-SHUTDOWN);
      from AS64497 action pref=10 accept ANY;
      from AS64496:AS-SECRET # ...
} except {
      from AS64497 at 192.0.2.1 accept RS-BOGONS-V4;
      # ...
}
```

```
router bgp 64496 # ...
      neighbor 192.0.2.1 route-map AS64497-in in
      neighbor 192.0.2.1 route-map AS64497-out out
!
route-map AS64497-in permit 10
      set local-preference 0
      match community graceful-shutdown
route-map AS64497-in deny 10
      match ip address prefix-list bogons-v4
!
# ...
```

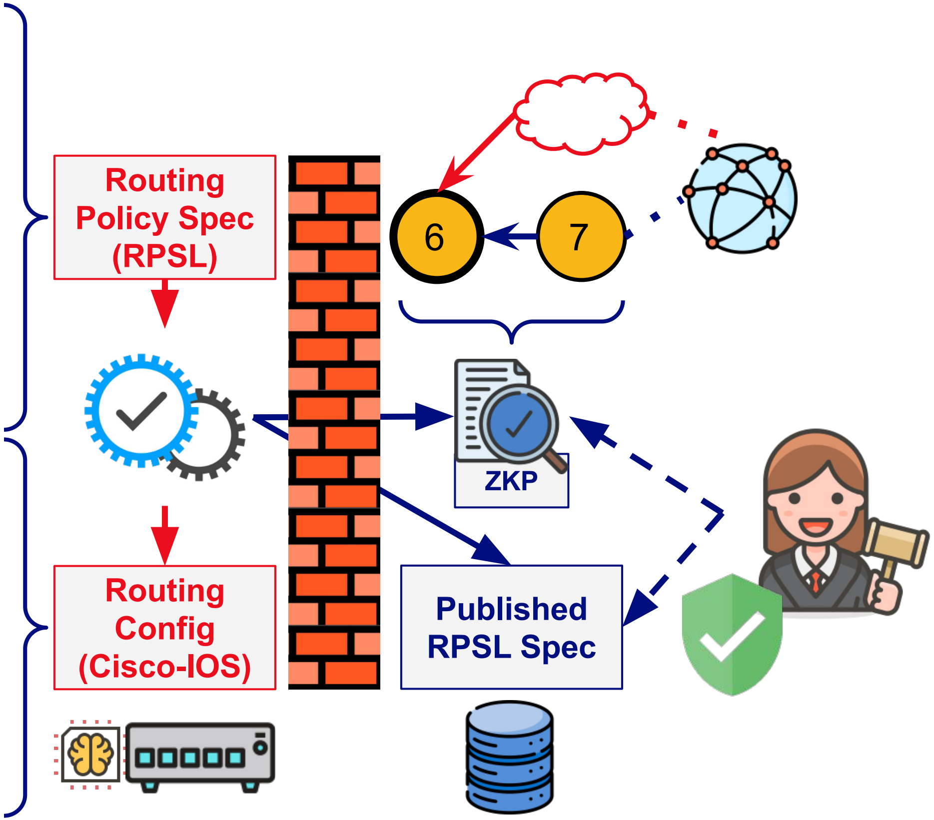

Future Work: Using ZKPs to Inform RL-based SDNs

Software Defined Networking (SDN) routers take a different approach: adopt **Reinforcement Learning** (RL) techniques to decide optimal routing policies.

SDN requires much more data (often sensitive!) to inform routers.

Human-on-the-loop approach gives verifiable ZKP claims, allowing RL-based routers to reason about secret info as well!

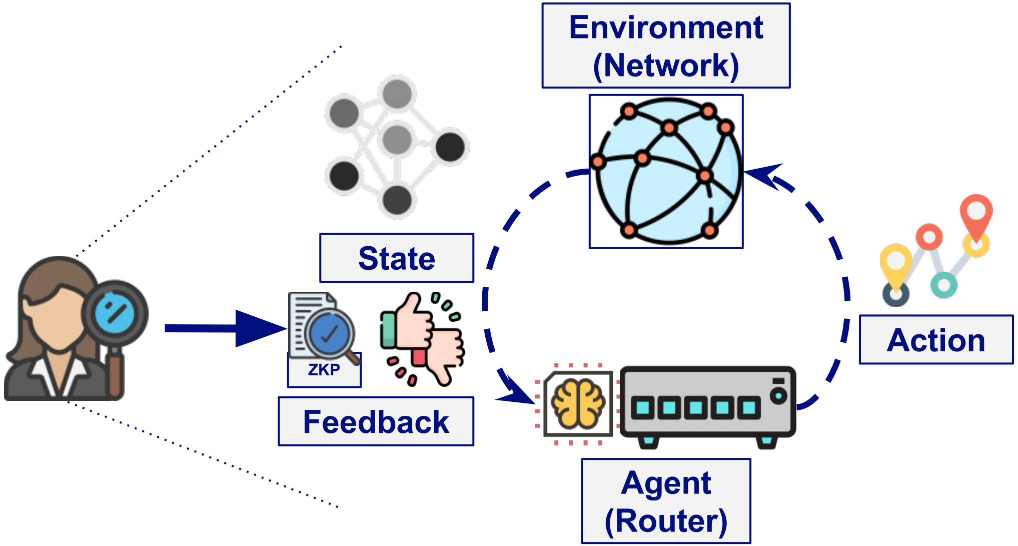

Conclusion

- ZKPs can provide both **privacy** and **authenticated routing** guarantees, ensuring conformance to both protocol and policy specifications.

- Since ZKPs do **NOT** rely on key infrastructure, they are a promising tool for authenticating routing in a distributed environment.

- ZKPs will likely **increase proving and verification times**, with many overhead and maintenance challenges to consider before widespread adoption.

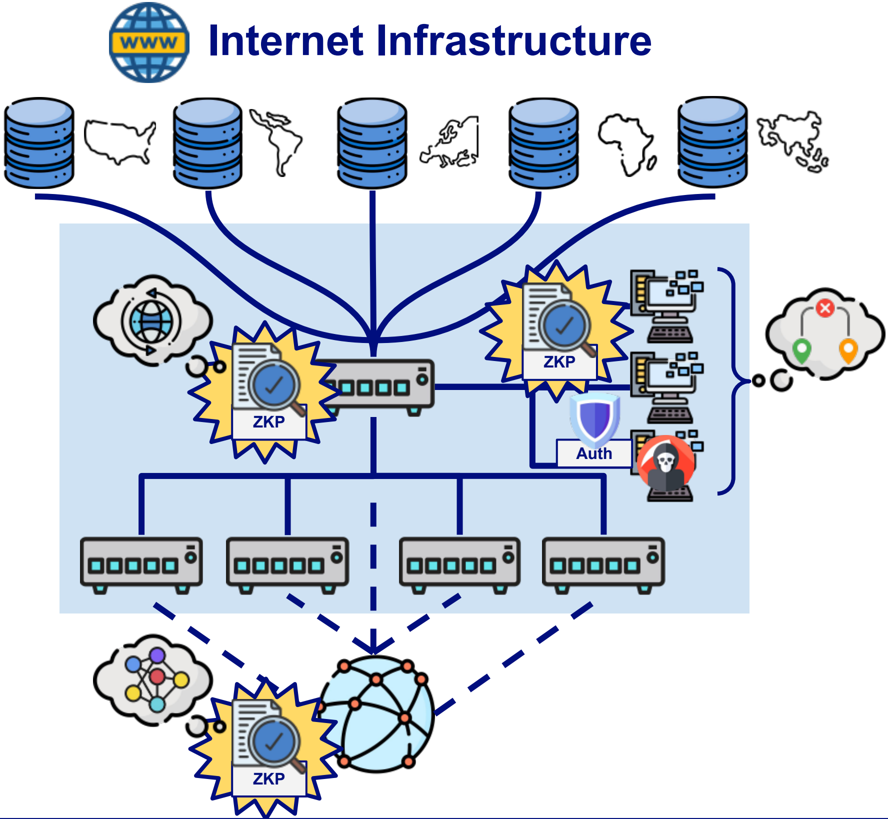

Backup

Integrating ZKP Information into RL-based SDNs (Backup)

Software-Defined Networking (SDN) Overview

Hardware routing not very complex – “on-chip” accelerators to perform specialized routing tasks very quickly

… also not very flexible

SDNs allow software itself to decide how to best route incoming packets / react to changing scenarios

Reinforcement Learning (RL) Overview

Agent (here, router) performs action given current state, environment (here, ML model/network sim) impacted, new state produced with reward/punishment for said action, back to agent.

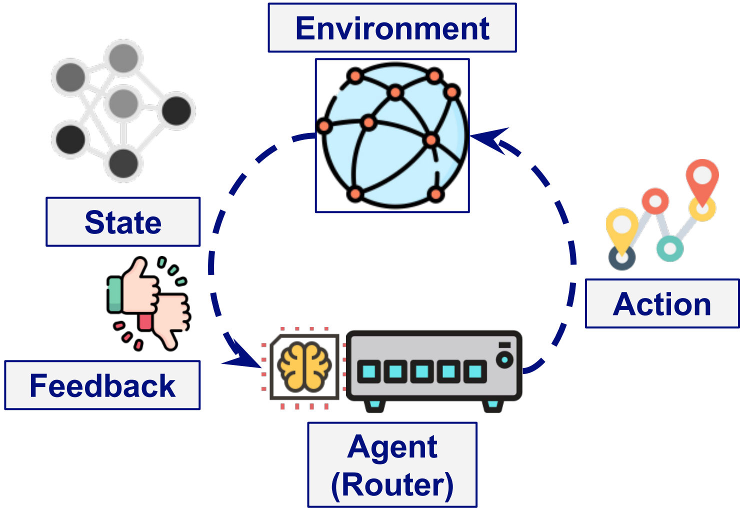

Background: Resilient and Secure Cyber Networks

Traditional routers use heuristic networking protocols, such as BGP, to route and deliver messages between clients.

Traditional protocols are not resilient to drastic changes injected by adversaries.

Recent research has focused on learning-based software defined networking (SDN) routers that use reinforcement learning to ingest network state data and optimally route.

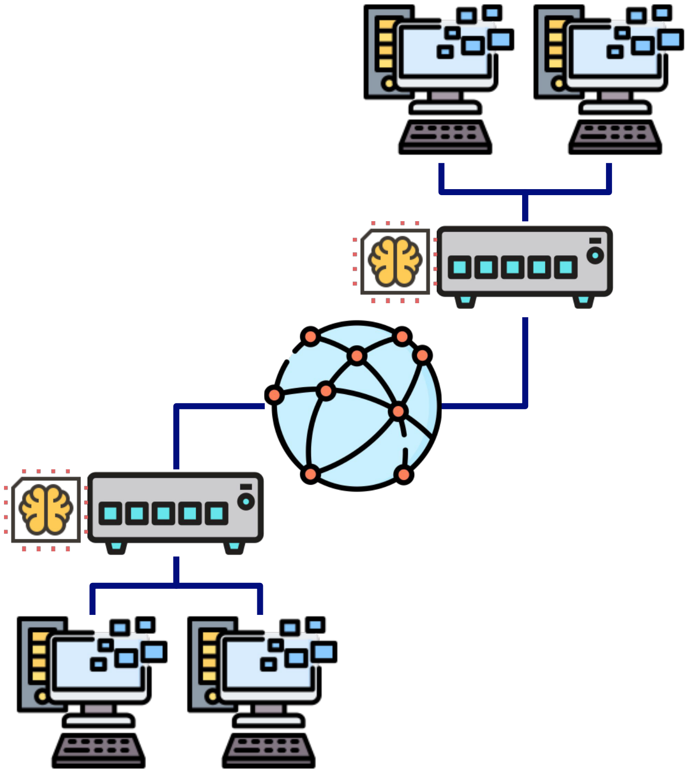

Challenge: Input Validation for Learning-based Routers

SDN (e.g. AI-based) routers require extra information about the network from other hosts to quickly adapt to new changes.

Problem I: Some hosts, including neighbors, may be malicious.

Problem II: Network details and/or messages may contain sensitive or proprietary information.

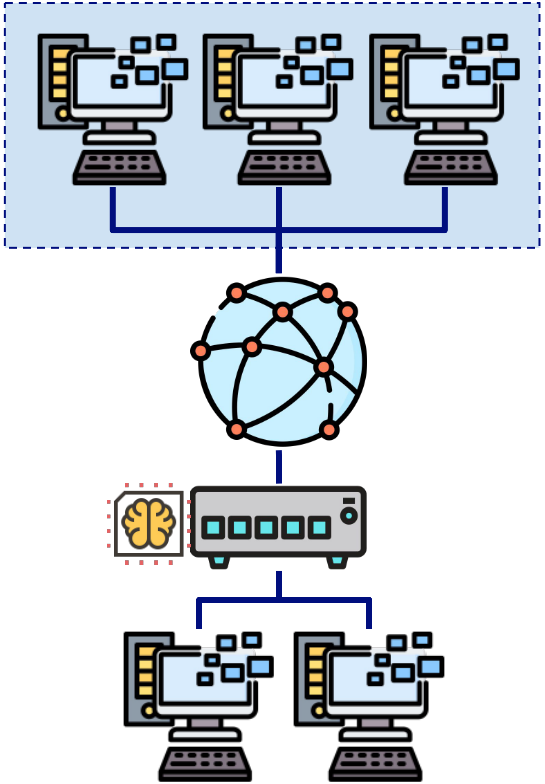

Challenge: Input Validation for Learning-based Routers

SDN (e.g. AI-based) routers require extra information about the network from other hosts to quickly adapt to new changes.

Problem I: Some hosts, including neighbors, may be malicious.

Problem II: Network details and/or messages may contain sensitive or proprietary information.

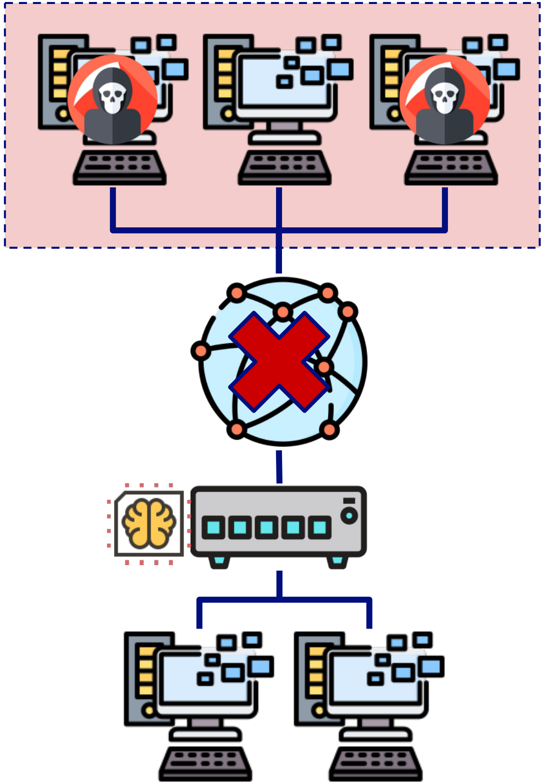

Solution: ZKPs for Network Security Properties

We can authenticate relevant peer-provided information used by smart routers using zero-knowledge proofs.

Properties that are true on one end of the network can be communicated to the other side with little-to-no trust.

We will use succinct ZKPs, so they will be small enough to add minimal overhead to the network.

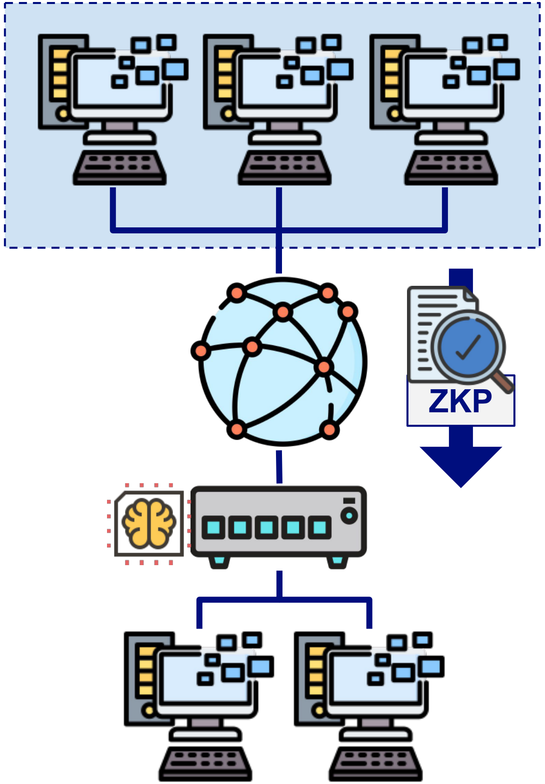

Single-Prover ZKPs
(Backup)

Zero-Knowledge Proof for *Where's Waldo?*

Example. Proving that you know the solution to *Where's Waldo?*

Zero-Knowledge Proof for *Where's Waldo?*

Example. Proving that you know the solution to *Where's Waldo?*

Traditional Proof: Circle Waldo's location

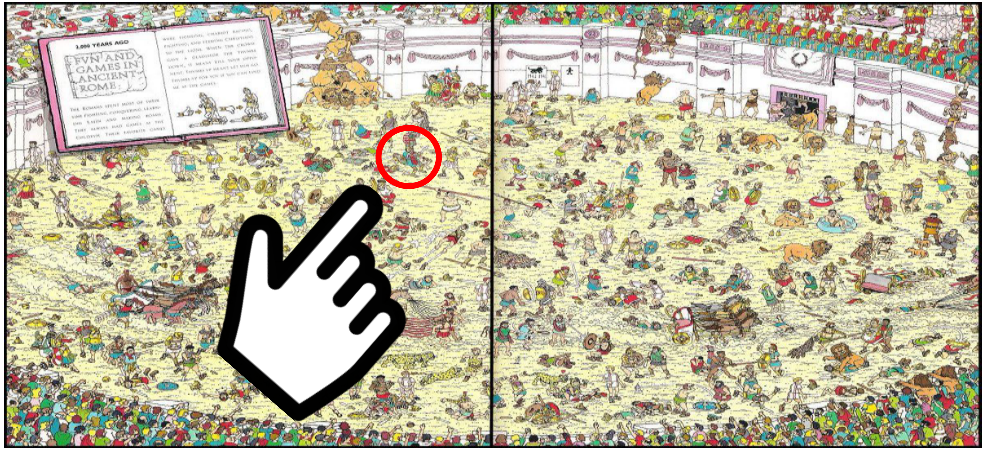

Zero-Knowledge Proof for *Where's Waldo?*

Example. Proving that you know the solution to *Where's Waldo?*

Traditional Proof: Circle Waldo's location

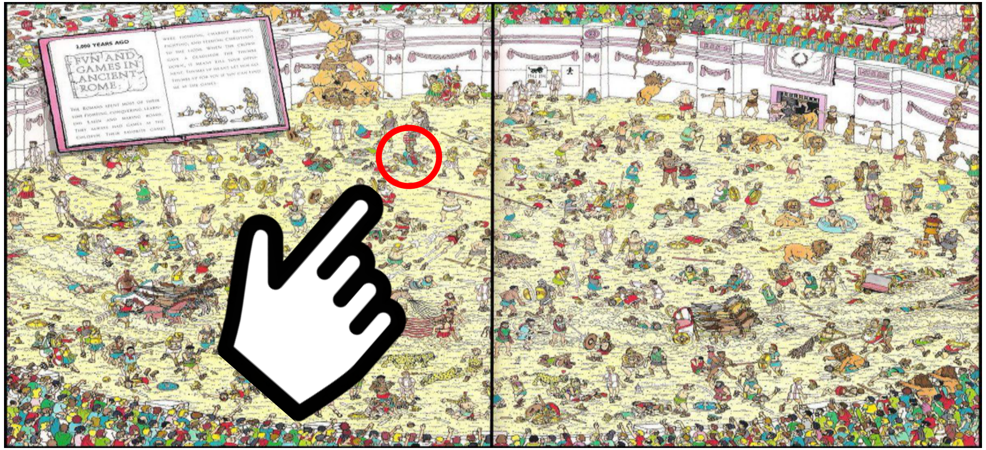

Problem

This kind of proof leaks all information about Waldo's location, much more than simply that you have *knowledge* of the location (not zero-knowledge)!

Zero-Knowledge Proof for *Where's Waldo?*

Zero-knowledge Protocol

1. Cut out a Waldo shaped hole in a much larger piece of paper
2. Position the hole over Waldo's location

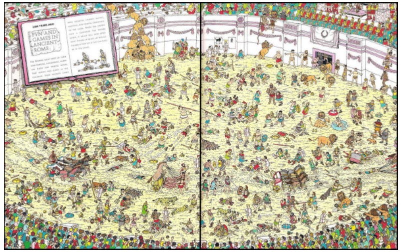

Slide under paper

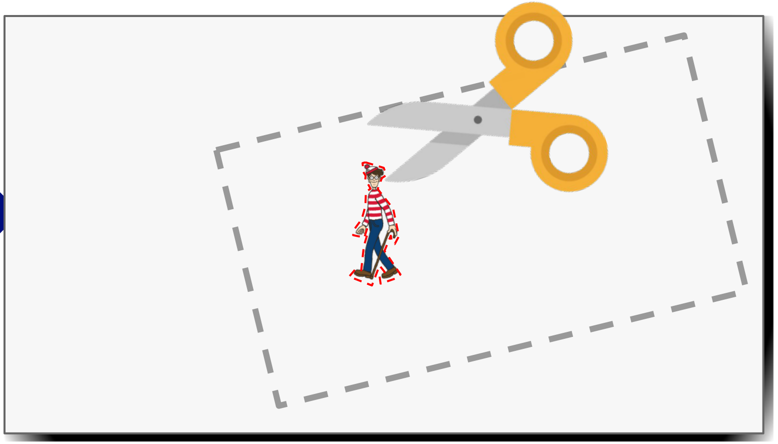

The sheet acts as an **obfuscating mask** for Waldo's location

To verifiers, the book underneath could hypothetically be in any random orientation

Zero-knowledge Proofs: High Level View

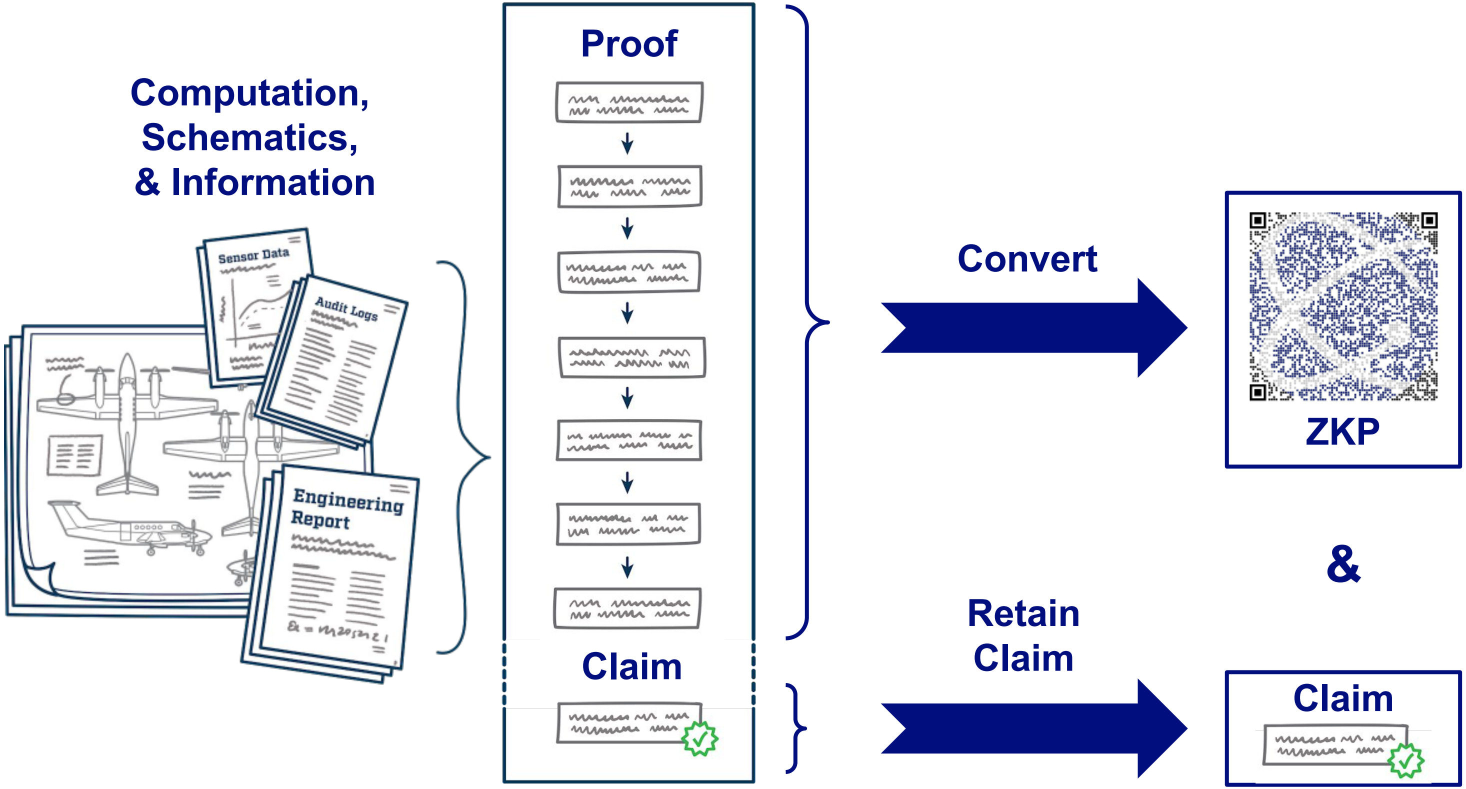

Example Claim.
“The packet can reach Router Y from X, even if Router Z is offline”

Zero-knowledge Proofs: High Level View

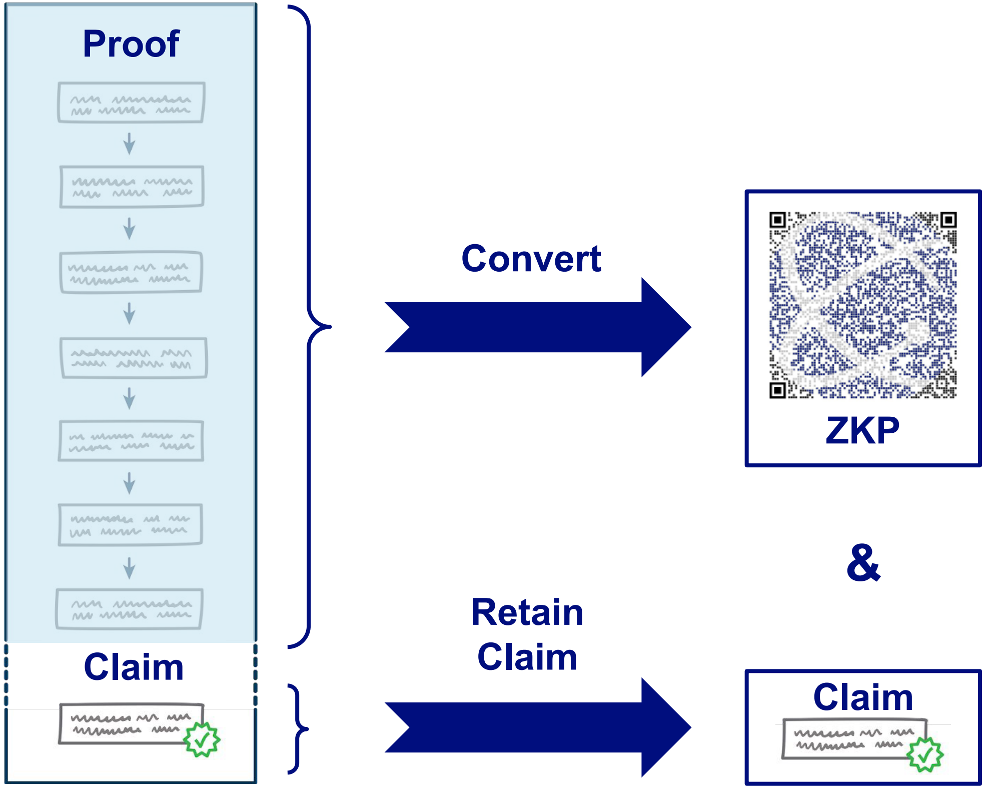

The Zero Knowledge Proof replaces the need for sensitive proof information

(effectively completely redacting the original proof)

The Claim can be quickly verified without any knowledge of the original proof.

Zero-Knowledge Proofs and Network Authentication

Zero-knowledge proofs (ZKPs) allow us to prove that a claim **IS** true without revealing **WHY** it is true, even if the prover is considered untrusted and **malicious**.

zkSNARKs are special ZKPs that are ***tiny*** and ***non-interactive***

Inputs:
Audit Logs
Schematics
Policies
Signal Analysis
Encryption Keys
Attestations
etc.

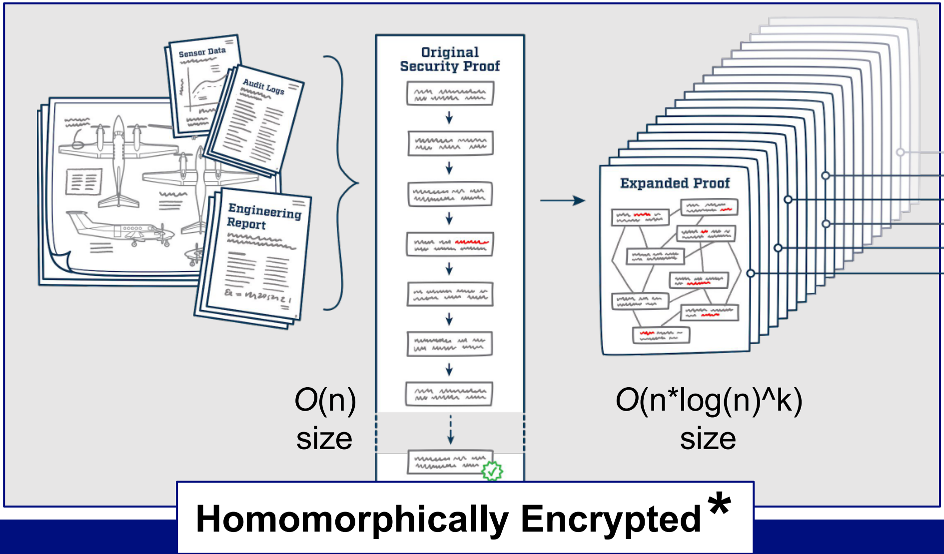

Outputs:
Results
&
ZK Proof of
Network
Integrity

***with tweaks**

Features of (Non-interactive) Zero-knowledge Proofs

Zero-knowledge proofs (ZKPs) allow us to prove that a claim **IS** true without revealing **WHY** it is true, even if the prover is considered untrusted and **malicious**.

Ideal Secrecy

Secrets are *NOT* revealed even if the cryptography is completely broken

Miniscule Footprint

~3k Bits

Proofs are tiny, often only slightly larger than a regular QR code

Fine-grained Control

Exacting control over need-to-know while enabling trustless verification

Composable

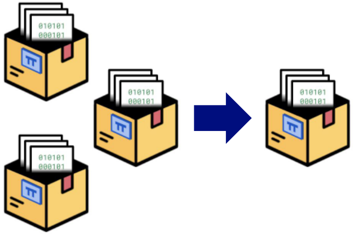

ZKPs can be collected and combined into new ZKPs without growing in size

Cryptographic Proof Systems

Cryptographic proof systems have variable completeness and soundness. For non-interactive zero-knowledge proofs we care about:

(Completeness)	$\mathbb{P}$[true statement AND verifier accepts] = 1 "Everything true is provable"
(Soundness)	$\mathbb{P}$[false statement AND verifier rejects] = 1 - ε "Low chance that a proof of a false statement is encountered"

We sacrifice minimal amount of soundness (have to break crypto to produce counter-example) in order to get valuable proof properties

zkSNARK Construction for Verified Computation [BCGTV13]

```
int myFunction(int a) {
     int b=a*a-4;
     return 3*b+a;
}
```

Rank-1 Constraint System (R1CS):

S • A * S • B = S • C

S	A	S	B	S	C
1	0	1	0	1	0
a	1	a	1	a	1
t0	0	t0	0	t0	0
b	0	b	0	b	0

Arkworks

Computation

Arithmetic Circuit

R1CS

Proof Representation Of Network Robustness

QAP

Zero Knowledge Added

LPCP

Succinctness Added

LIP

Interactivity Removed

zkSNARK

Arkworks backend

Verifier Net View

Prover View

π

zkSNARK for Network Integrity

Spare (Ignored/Skipped)

Alert: RPKI is Vulnerable and Risky!

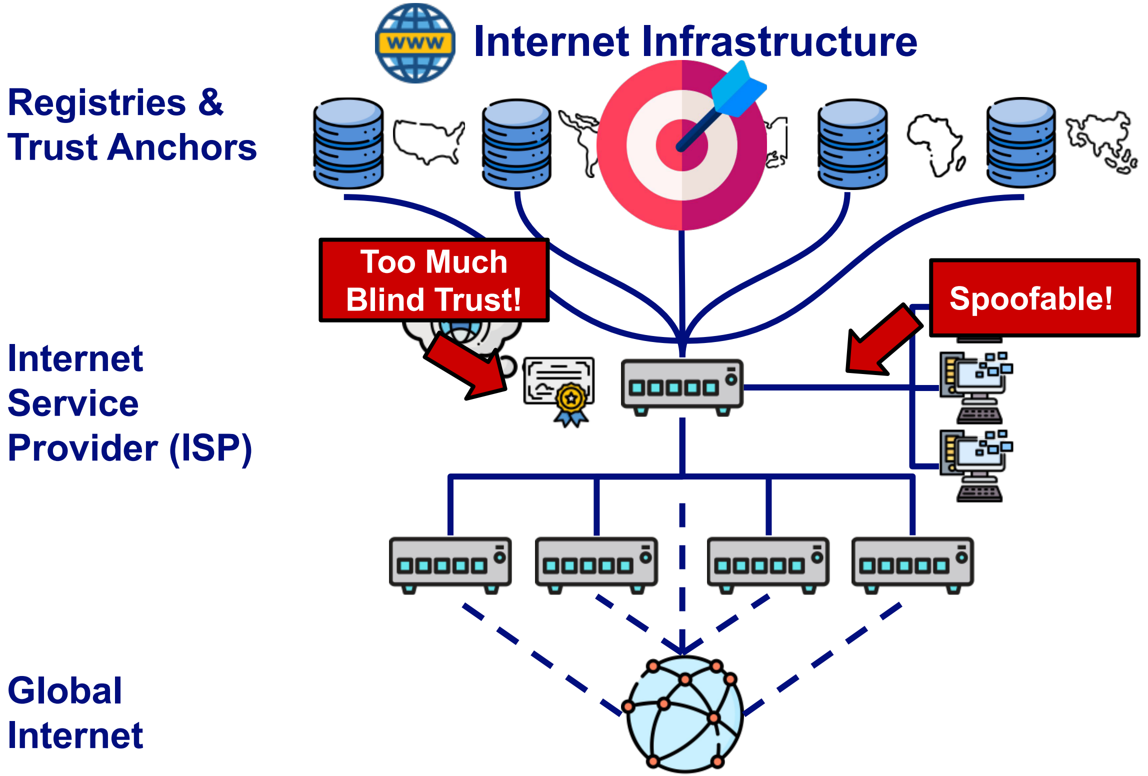

Weaknesses:

- Centralized trust is a **point of failure**
- Can’t certify entire **route / network**
- Keys are a **target** and **hard to manage**

Secure and Robust ISP Network Routing

Existing: RPKI

- Trust sources are **points of failure**
- Only certifies info of **route's origin**
- Authentication needs **centralized keys**
- Can only decide route **public (often local) info**

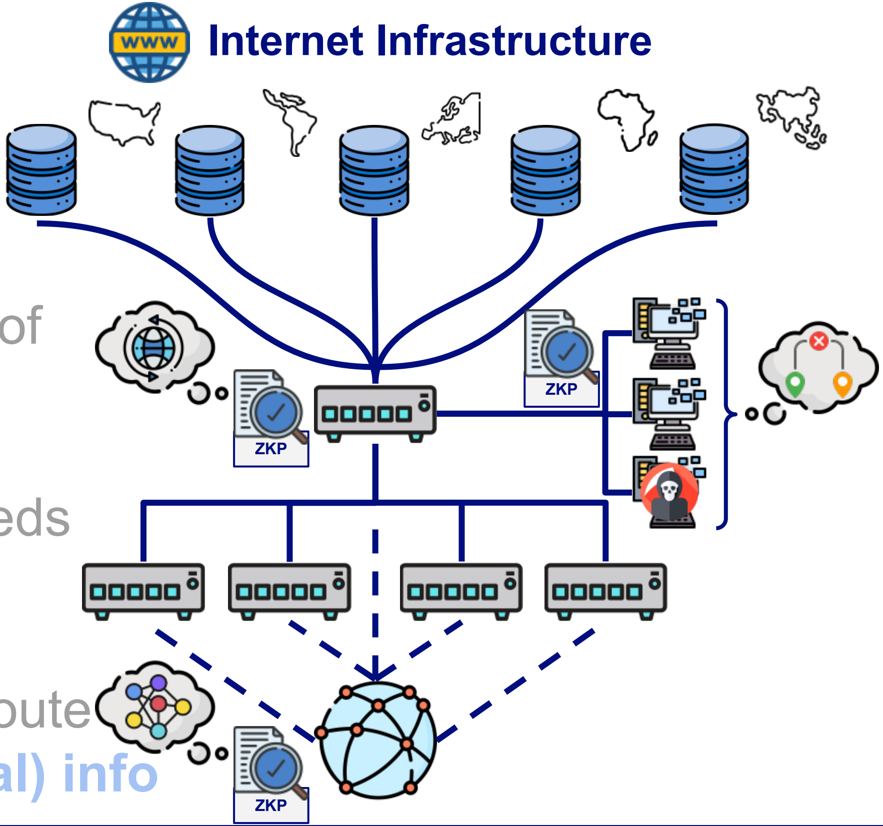

Secure and Robust ISP Network Routing

Existing: RPKI

- Trust sources are **points of failure**
- Only certifies info of **route's origin**
- Authentication needs **centralized keys**
- Can only decide route **public (often local) info**

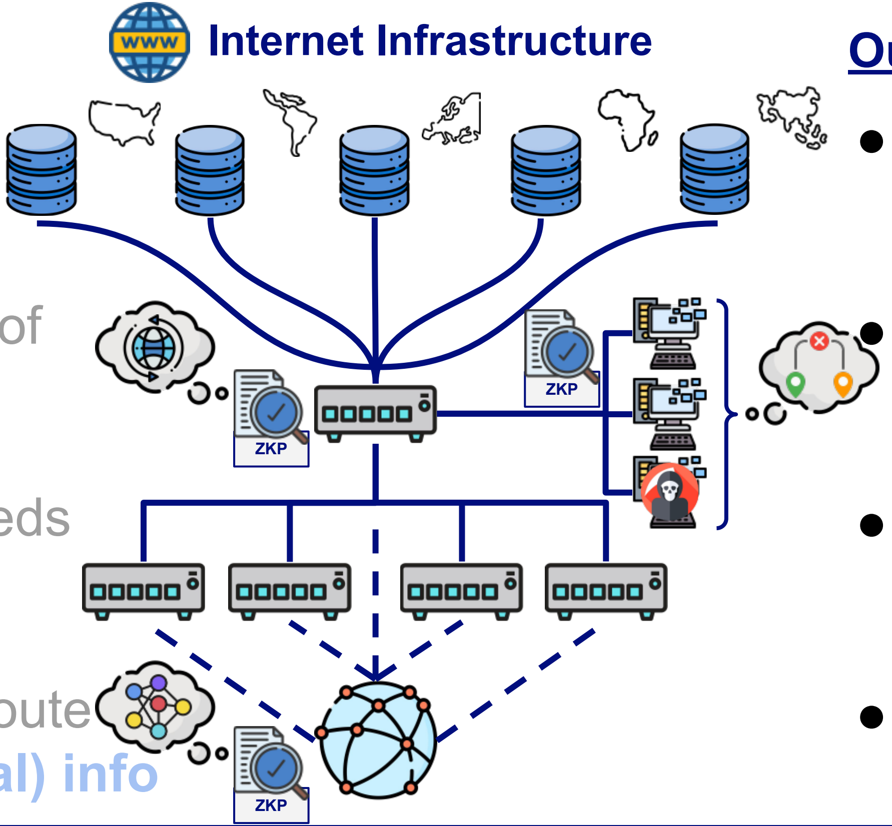

Our Solution: ZKPNet

- Trust sources are **distributed**
- Correctly verifies **arbitrary info**
- No auth keys, only **trusted setup**
- Also decides with **secret global info**