Systems and Internet Infrastructure Security Laboratory

PennState

Measuring and Mitigating the Risk of IP Reuse on Public Clouds

Eric Pauley, Ryan Sheatsley, Blaine Hoak, Quinn Burke, Yohan Beugin, Patrick McDaniel

Pennsylvania State University

Contact: epauley@psu.edu

Public Clouds: Disruption at Scale

Amazon Web Services posts record $13.5B in *profits* for 2020 in Andy Jassy's AWS swan song

BY TODD BISHOP on February 2, 2021 at 4:29 pm

Public clouds leverage resource sharing and reuse to improve performance.

Vulnerabilities due to Resource Sharing

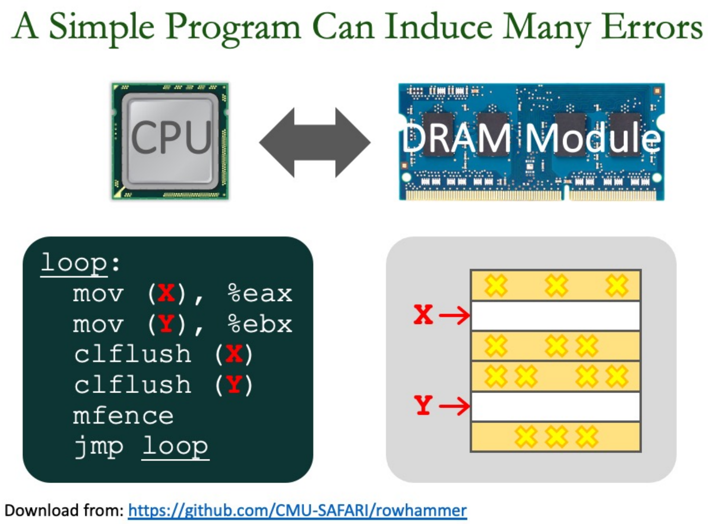

Row Hammer (Kim et al. 2014)

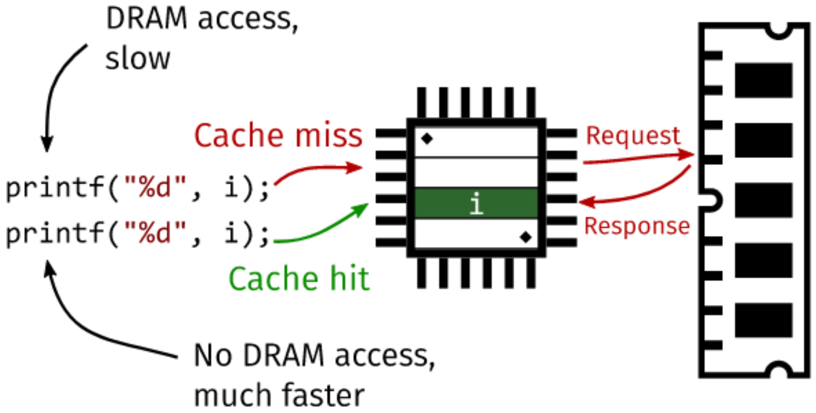

Meltdown (Lipp et al. 2018)

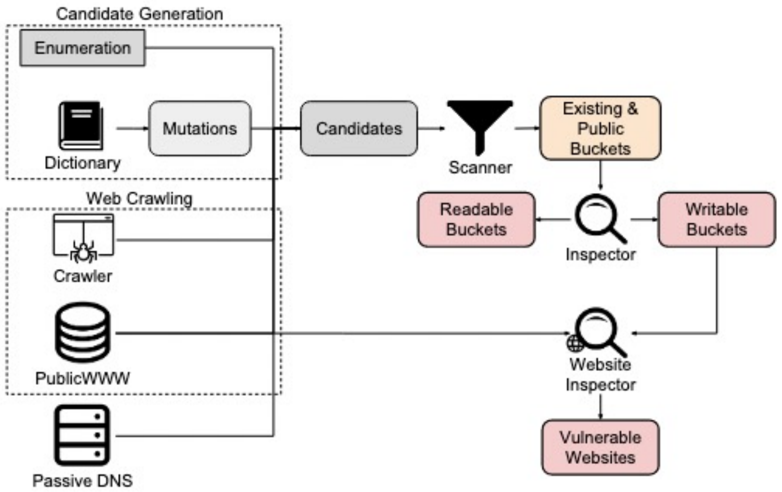

Storage Policies (Continella et al. 2018)

How does the resource lifecycle of public clouds affect security?

Issue: Resource Reuse

1. Tenants create configuration that refers to cloud resources (e.g., IP addresses):
   - Causes clients to use resources
   - Establishes a trust relationship

2. Cloud resources reused by other tenants
   - Configuration is now *latent*

3. Previous tenant’s clients send data
   - Adversary listens (*cloud squatting*)

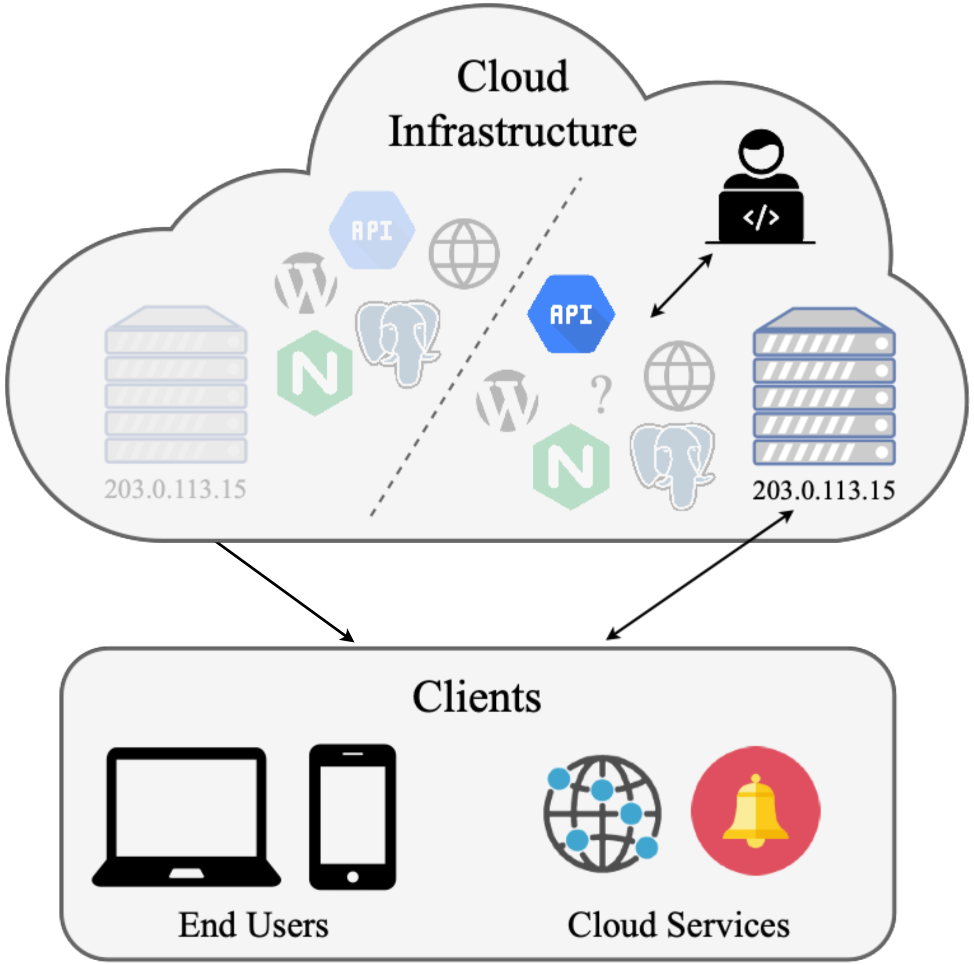

Cloud Internet Telescope

PennState

Experiment (March 8 – May 15, 2021)

- 3M servers allocated on AWS `us-east1`
- ~500M network sessions
- ~½ TB of raw network traffic data
- 1.5M unique IP addresses
  - 56% of total available in pool

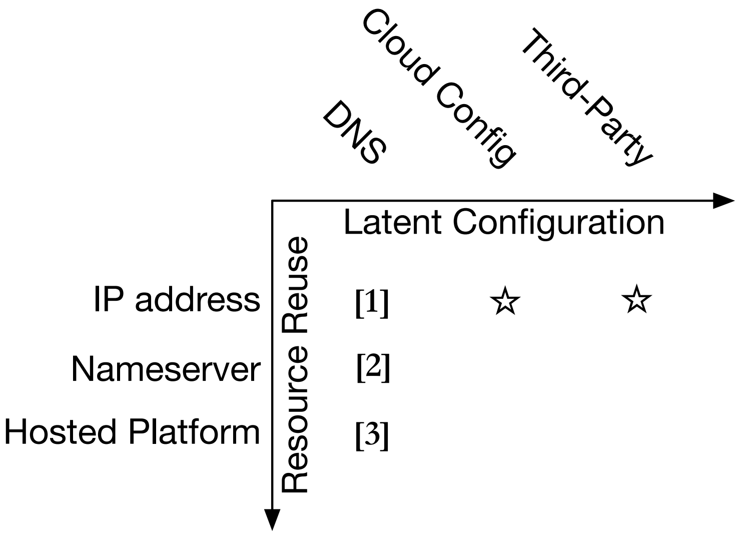

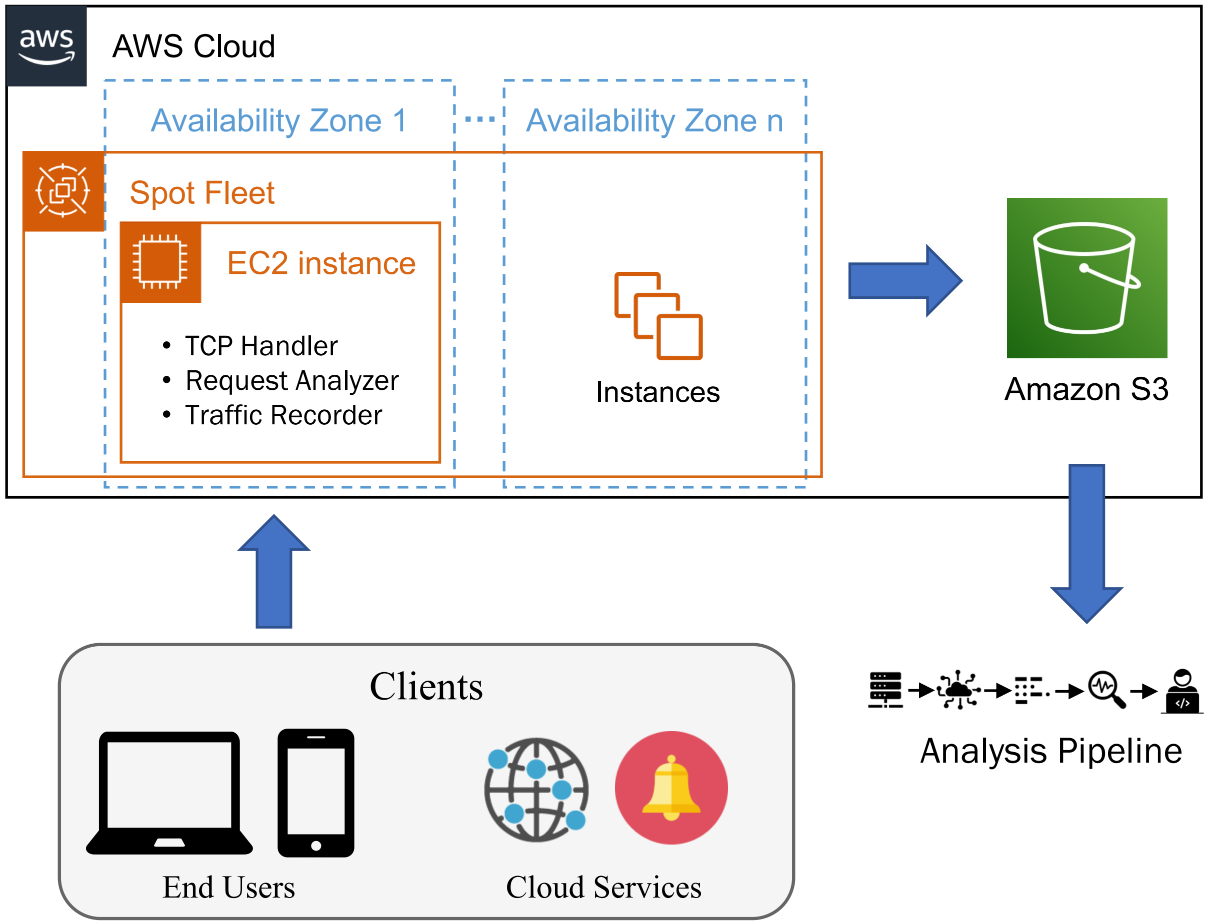

[1] Borgolte et al. 2018
[2] Alowaisheq et al. 2020
[3] Liu et al. 2016

Cloud Squatting: Vulnerability at Scale

PennState

Cloud Services

- >5M messages
- 4 cloud services

Third-Party Services

- >3M messages
- Numerous Services

DNS

- 5400 Websites
- 23 top-1000

Example Sensitive Data Received

Financial

Personal

Location

Remote Code Execution

Passwords

Images

Measuring IP Reuse: Bottom-Up

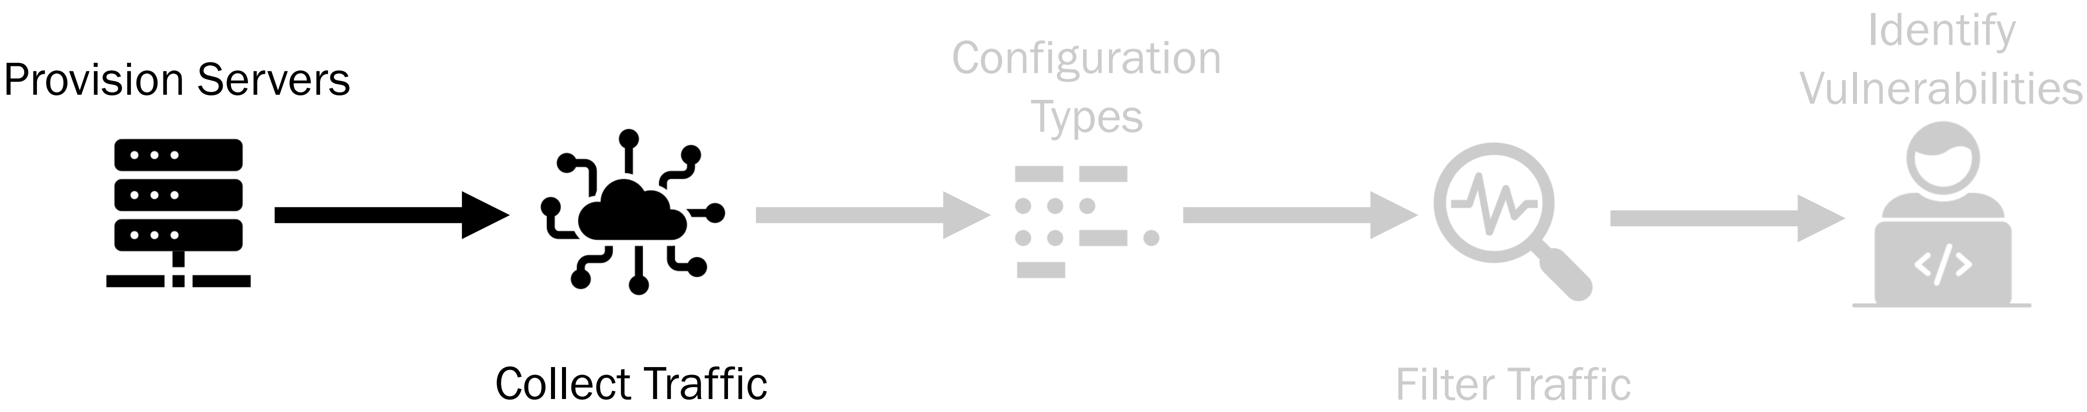

Characterizing Cloud IP Reuse

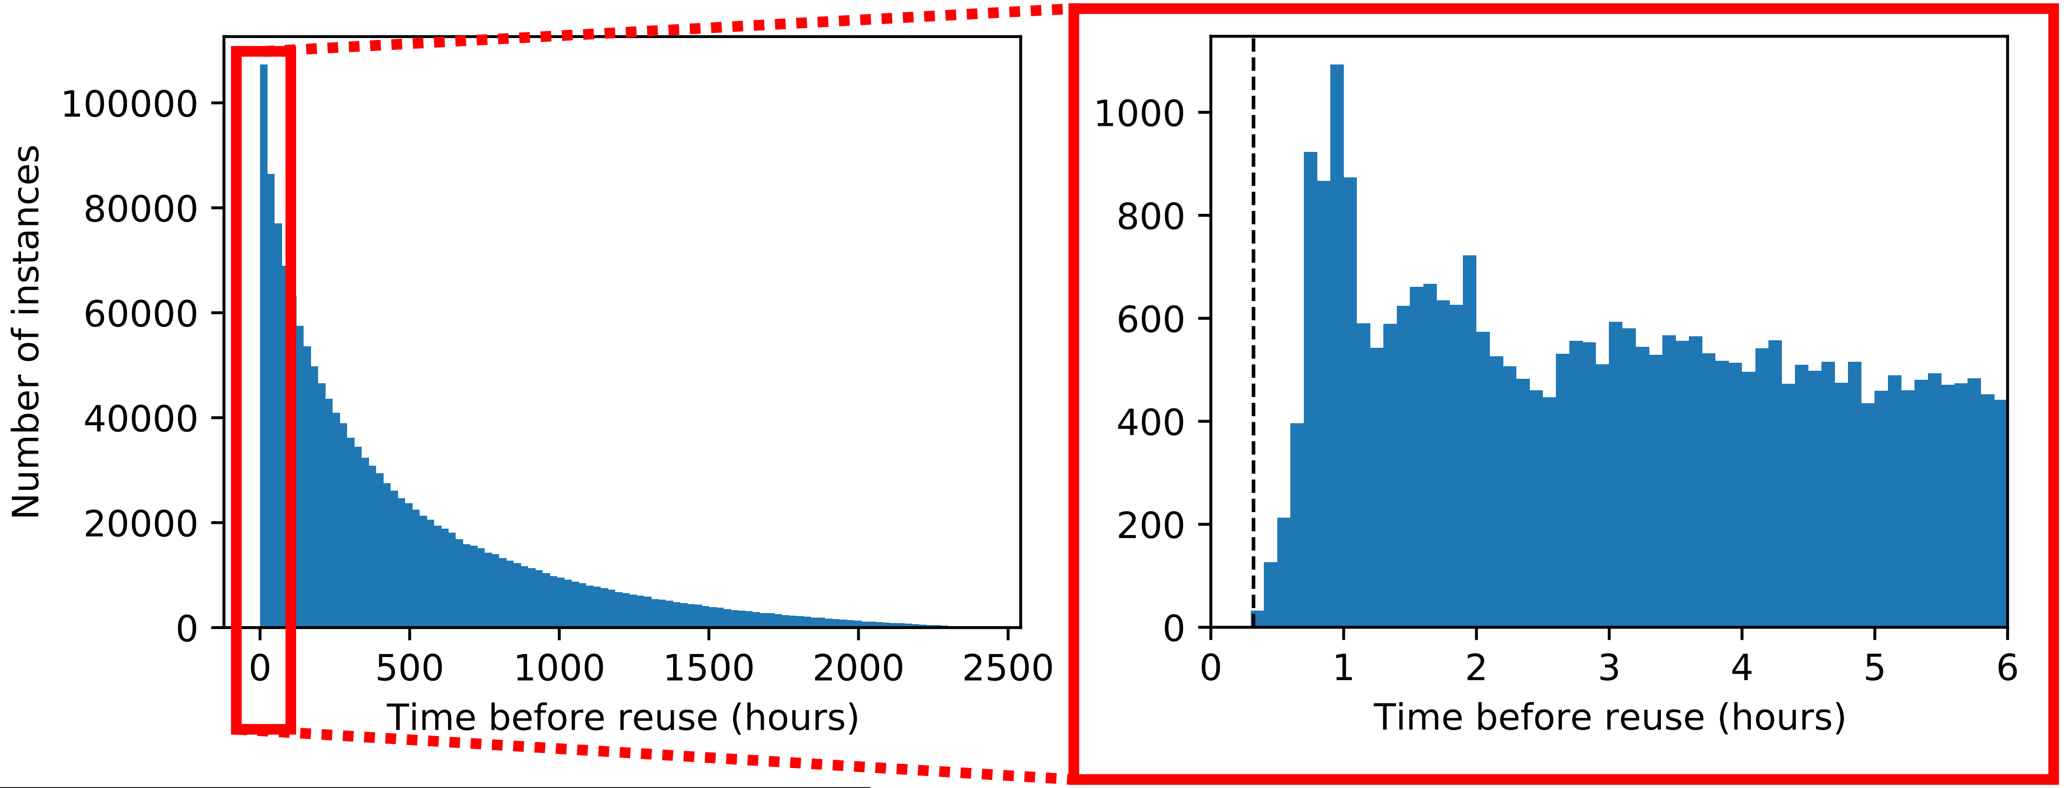

Zone	Servers	Unique IPs	Estimated IPs	Capture Rate
us-east-1a	581 k	383 k	789 k	49 %
us-east-1b	607 k	389 k	762 k	51 %
us-east-1c	630 k	236 k	313 k	76 %
us-east-1d	573 k	360 k	700 k	51 %
us-east-1f	647 k	171 k	198 k	87 %
Total	3039 k	1540 k	2762 k	56 %

Pseudorandom IP allocation allows adversaries to easily explore the IP space with high coverage.

Measuring IP Reuse: Bottom-Up

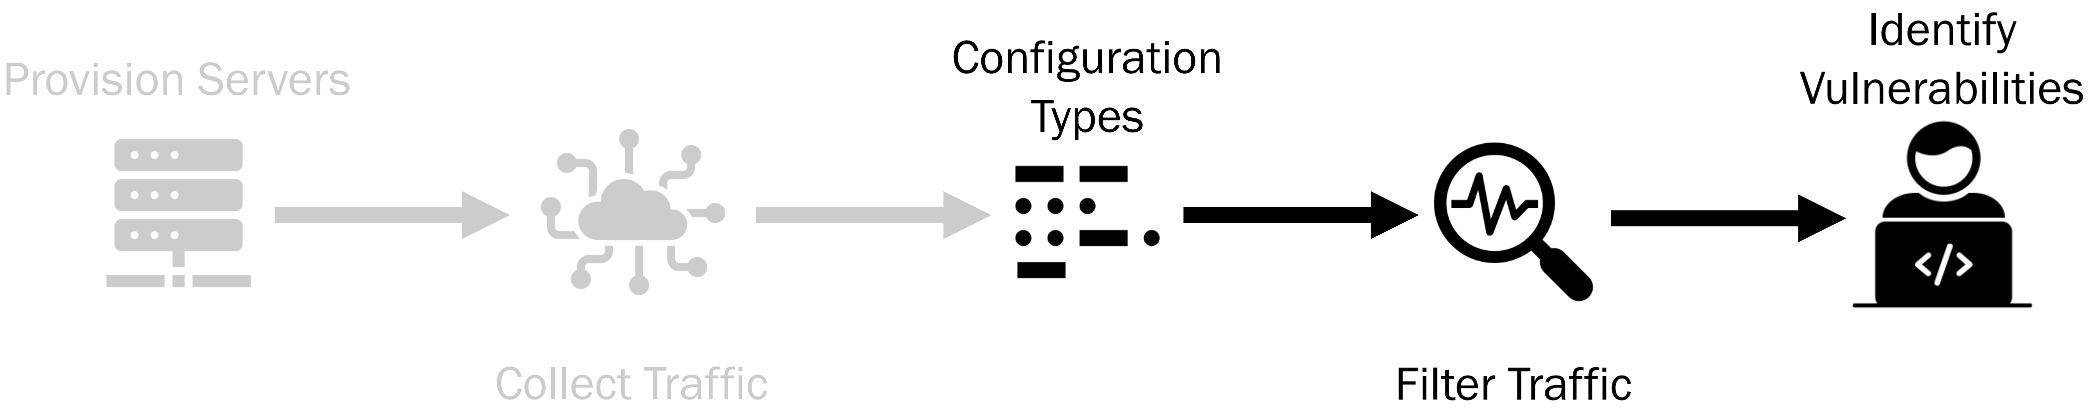

Types of Latent Configuration

Cloud Services
- Managed by cloud provider
- Configured to connect to IP addresses
- E.g., SNS, Route53

Third-Party Services
- Client software referencing reused IPs
- E.g., Databases, APIs

Domain Names

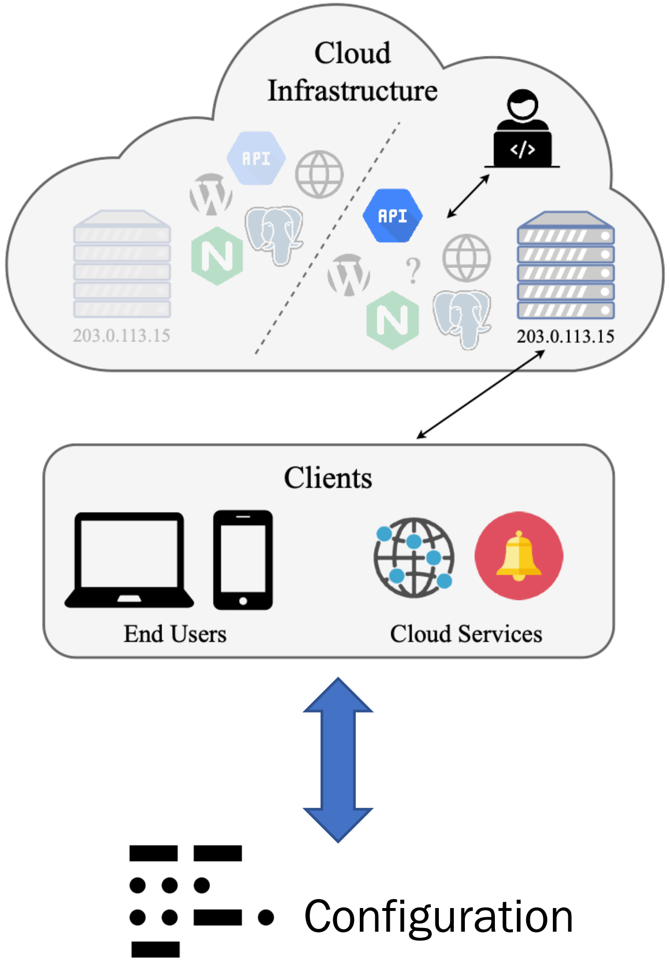

Measuring Cloud Services

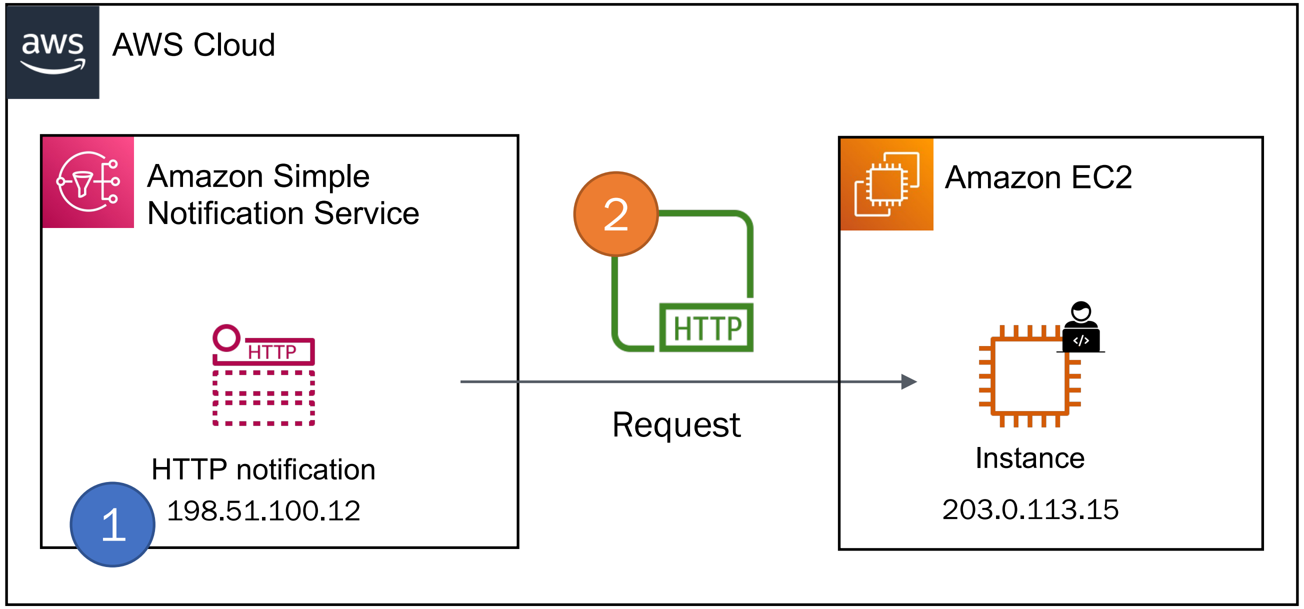

1. AWS-verified IP Address
2. Cloud service identified in HTTP headers

Cloud Services are Vulnerable

Service	SNS	Route53	Cloudfront	API Gateway
IPs	24.9 k	2.8 k	65	3
Sessions	1.6 M	3.6 M	1.7 k	10
Sessions w/ DNS	25	567 k	767	2
Unique Tenants	78	3.1 k	64	3

Leaked data:

Financial

Personal

Location

Remote Code Execution

Filtering Third-Party Services

- Main idea: filter out likely bot/scanner traffic to analyze remaining share
- Method: series of filters at various levels of protocol stack:
  1. Network filtering (Blocklists)
  2. Transport filtering (IP/Port scanning)
  3. Session filtering
  4. Application Filtering

Step	IPs	TCP Sessions	Size
Initial	3.13 M	596 M	410 GB
Network	3.03 M	280 M	148 GB
Transport	1.70 M	10.2 M	11 GB
Session	1.14 M	4.89 M	9.3 GB
Application	340 k	2.95 M	6.3 GB

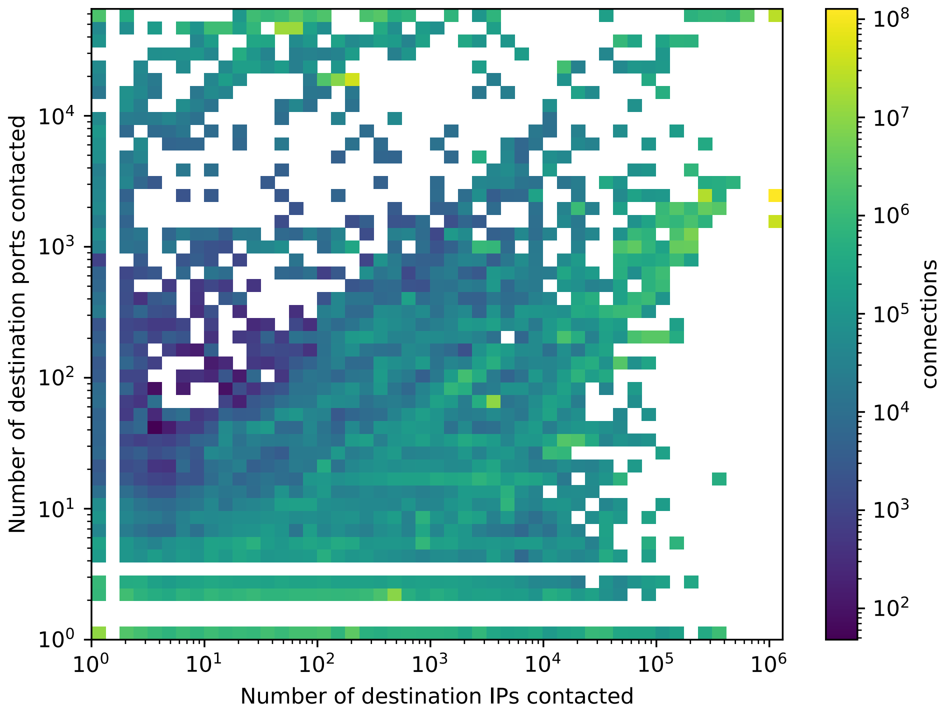

Exploitable third-party services

PennState

Databases

Financial Traffic

Webhooks

Caches

Logging

Custom APIs

Vulnerable Domain Names

Site rank	Domain
31	amazonaws.com
68	akadns.net
76	cnn.com
129	wix.com
146	harvard.edu
164	go.com
177	usatoday.com
284	intuit.com
298	cornell.edu
300	intel.com
302	slack.com
434	vice.com
450	redhat.com
470	trafficmanager.net
495	upenn.edu
497	elsevier.com
535	ieee.org
578	jhu.edu
588	nvidia.com
618	lenovo.com
767	ea.com
782	hhs.gov
957	justice.gov

From banner info: Over 5,400 domains found vulnerable

- 23 in top-1000
- Many domains had several vulnerable subdomains

Wide variety of associated organizations:

- Industry
- Academic
- Government

Disclosures and Root Causes

Direct tenant disclosures and surveys reveal root causes

Integration: lift-and-shift

- Transfers assumptions from private data center
- No consideration for service decommissioning

Failure to manage configuration

- No centralized view of cloud configuration
- Failure to follow best practices

Insufficient/broken automation

- No automated DevOps (e.g., CloudFormation)
- Bespoke deployments without decommissioning

Defenses and Mitigations

PennState

Resource (IP) reuse

- Reserved IP ranges
- Private networking
- IPv6
- IP allocation policy (e.g., **IP Tagging**)

Policy	Unique IPs	Mean Prev. Tenants	Median Reuse Time
RANDOM	377 596	228.2	5.7×10^3 s
LRU	385 774	209.6	9.2×10^3 s
TAGGING	240	2.387	2.9×10^6 s

Latent configuration

- Centralized configuration (DNS)
- Configuration auditing
- Provider scanning for vulnerabilities
- Policy Enforcement
- Resource-based Naming

Cloud Tenant Cloud Provider

Amazon Actions

PennState

Cloud Configuration

New alerts in console for dangerous configuration

Scanning & Disclosure

Analysis of control-plane across all regions

Automated Policy Enforcement

Managed Config rules to enforce best practices

Updated Best Practices

New documentation on IP hygiene and latent configuration

Takeaways

PennState

Public clouds bring new security concerns

- Latent configuration is widespread and dangerous
- Cloud services may not sufficiently protect tenants

Adversaries can discover and exploit vulnerabilities

- IP addresses are pseudo-random, and allow sampling of pool

Cloud squatting can be prevented

- Reducing IP address reuse
- Preventing latent configurations

Systems and Internet
Infrastructure Security Laboratory

PennState

Thank You!

@EricPauley_

epauley@psu.edu

pauley.me/cloudsec