Universität Münster

IIJ Seminar

A brief look at cloud IBR

Ongoing work

Nils Kempen

8 July 2025

University of Münster - NetSec Research Group

living.knowledge

Who am I

- Hello, I'm Nils 👋
- First year PhD Student @ University of Münster, Germany
- Interested in network telescopes and new approaches to them
- Advised by Ralph Holz

Network telescope (mis)-adventures

Internet Background Radiation (IBR) & network telescopes

- IBR describes all unsolicited traffic a host receives
- Mainly made up of three types:
  - Scans
  - DDoS backscatter
  - Misconfigurations
- Captured mainly by the use of darknets / Internet telescopes
  - Typically deployed in unused IP ranges of university networks → UCSD-NT, MERIT-NT
  - May introduce bias, other notable approaches:
  - Deployment in company networks (Bailey et al. 2005)
  - Deployment at IXPs (Wagner et al. 2023)
  - Deployment in CDNs (Richter & Berger 2019)
  - Deployment in public clouds (Pauley et al. 2023)

Idea(s)

- Different telescopes/ vantage-points provide different views
- Understanding which is best for specific observations
- Cloud-based approaches seem promising
  - Still unclear what the best way to operate them is
  - e.g. Holding time of an IP Address,
  - VM configuration,
  - economic perspective

When do we need which lense?

- *current literature™* dosn't provide clear answers yet

Approach - data collection

- **Idea:** Build a distributed, multi-cloud network telescope
  - configurable lifetime
  - provider agnostic
  - variable size
- **Current state:**
  - Build with Terraform as Infrastructure as Code (IaC)
  - 5 cloud providers supported: AWS, GCP, Azure, Vultr, DigitalOcean
  - Create 1 VM/ IP per “availability zone”
  - approx. 300 VMs total

Approach - data collection

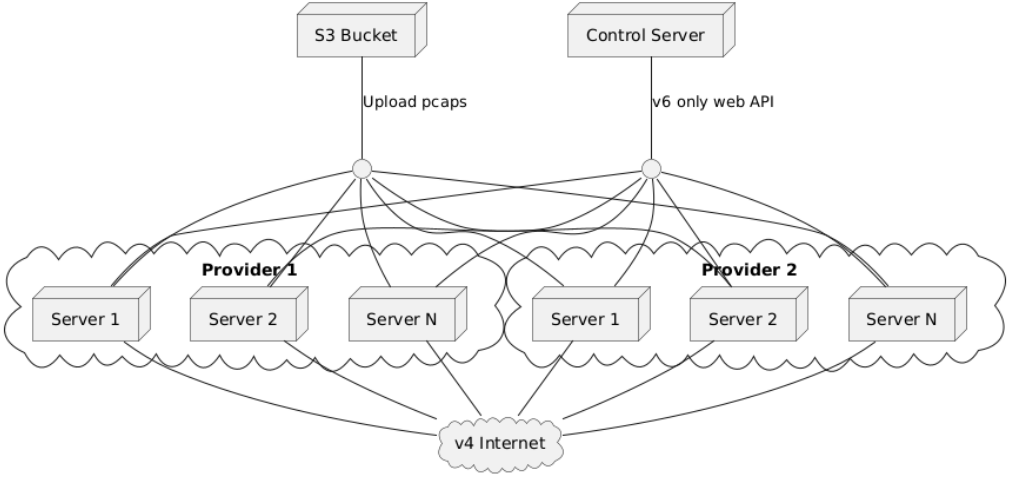

Cost consideration

Provider	Cost (IP/M)	Approach
DigitalOcean	4.0$	One VM per IP
OVH	1.8$	Leasing subnet
AWS	7.5$	One VM per IP
Azure	9.0$	One VM per IP
Azure	4.8$	Load balancer
GCP	8.5$	One VM per IP
GCP	5.4$	Load balancer
Alibaba	3.8$	VM with multiple IPs
Vultr	3.5$	One VM per IP

Deployment - learnings

- Deploying a cloud-telescope is hard
  - All cloud-providers work a bit different
  - Destination IPs are often not directly linked to the interface (NAT)
  - Old software
  - Cloud-internal traffic

Deployment results

- Data collected over two weeks (April 3–17, 2025)
- Approx. 115k PCAP files
- Approx. 130M connections from 985k sources
- Stable Baseline of packets with notable variance

Deployment results

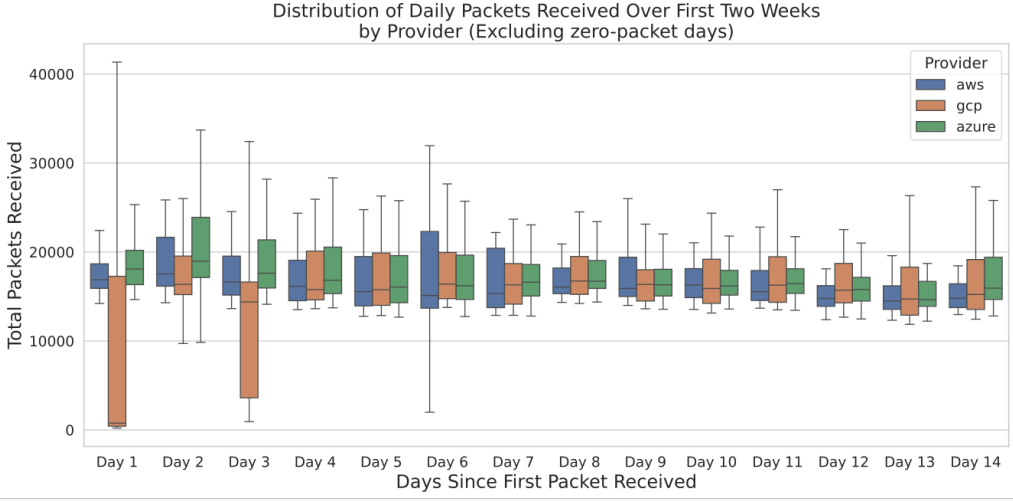

Deployment results

sld	count
googleusercontent.com	16245123
amazonaws.com	2920831
linodeusercontent.com	2537175
modat.io	2080634
coop.net	2010113
censys-scanner.com	1872200
hinet.net	1777494
shadowserver.org	1466999
4cloud.mobi	1326969
onyphe.net	1119946
bufferover.run	990792
internet-measurement.com	822108
recyber.net	675361
4vendeta.com	564725
stretchoid.com	515102
deepfield.net	499842
infornetnetwork.net.br	432905
pacesettersports.com	392724
tube-hosting.com	385923
criminalip.com	357202

Figure: Top source SLDs over all

provider	sld	count
DigitalOcean	hinet.net	397806
DigitalOcean	amazonaws.com	193928
DigitalOcean	linodeusercontent.com	120140
aws	googleusercontent.com	3725330
aws	amazonaws.com	1063459
aws	hinet.net	745764
azure	googleusercontent.com	6210778
azure	4cloud.mobi	1212207
azure	linodeusercontent.com	898888
gcp	googleusercontent.com	5944100
gcp	amazonaws.com	852464
gcp	linodeusercontent.com	820235
vultr	hinet.net	475813
vultr	ip-94-23-87.eu	260690
vultr	googleusercontent.com	246134

Figure: Top 3 source SLDs per provider

Future work - open questions

- Evaluate importance of cloud-telescope IP's history for IBR
- Evaluate traffic patterns across different regions and/or cloud-providers
- Evaluate best approach for cloud setup (e.g. responsiveness)
- Evaluate scanner behavior in cloud vs. "normal" telescopes

Current/ Future work

- Data enrichment:
  - OpenIntel reverse DNS data for source IPs
  - CAIDAs Hoiho for rDNS based geolocation
  - IPInfo geolocation
  - Routeviews prefix to ASN
- IDS scanner detection
- Filter cloud internal traffic
- Compare to other network telescopes

Future work - inspiration

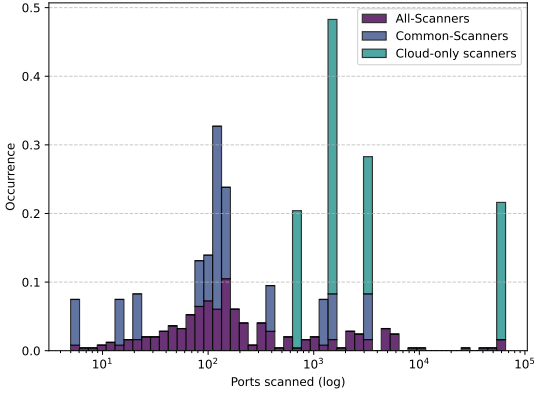

- If you scan cloud address space you are likely to hit something
- Resource intensive scans could be more focussed and may not be seen in “normal” telescopes.
- Further investigation of cloud-scanner behavior is needed.

Questions?