

The Cloud that Runs the Mobile Internet

H. Doroud^{1,2}, F. Michelinakis^{1,2}, A. Razaghpanah³, A. Lutu⁴, N. Vallina Rodriguez^{2,5}, P. Gill⁶, J. Widmer²
¹UC3M, ²IMDEA Networks Institute, ³Stony Brook University, ⁴Simula, ⁵ICSI, ⁶UMass_Amherst

Introduction

Mobile applications and on-line services use third-party Cloud Service Providers (CSPs) (e.g., Amazon and CDNs) to outsource their on-line infrastructure.

Motivation:

No previous systematic study analyzed the complex relationships between mobile apps, CSPs and ISPs

Main Objectives:

- Identifying the main CSPs supporting mobile apps and services and the strategies followed by app developers and on-line services (e.g., Multi-CDN).
- Measuring empirically the interconnections between ISPs and CSPs in real mobile ISPs.

METHODOLOGY AND DATASET

Datasets:


Lumen Privacy Monitor Dataset¹

- Mobile app that intercepts and analyzes mobile traffic in user space with real user- and network-stimuli.
- Available for free in Google Play
- UI provides transparency about app's hidden communication patterns and privacy leaks.
- By running on local host, it correlates rich contextual information such as process IDs, with flows, DNS lookups and TTL values.


Release Date	Users	Num- App	Num-Flow
Oct, 2015	4,200	7,600	5M

CSP Identification:

- CSPs hide the direct relationships between IP addresses and hostnames.
- IP addresses and IP metadata are not sufficient to classify all CSPs.
- CDNFinder² is the largest, yet incomplete, public list of PTR records associated with CDN-related domains.
- Custom classification method³:**
 - Reverse DNS lookup: Mapping FQDNs to CSPs via their PTR records




- Map Autonomous Systems(AS) to CSPs: In case a AS entirely belongs to one CSP



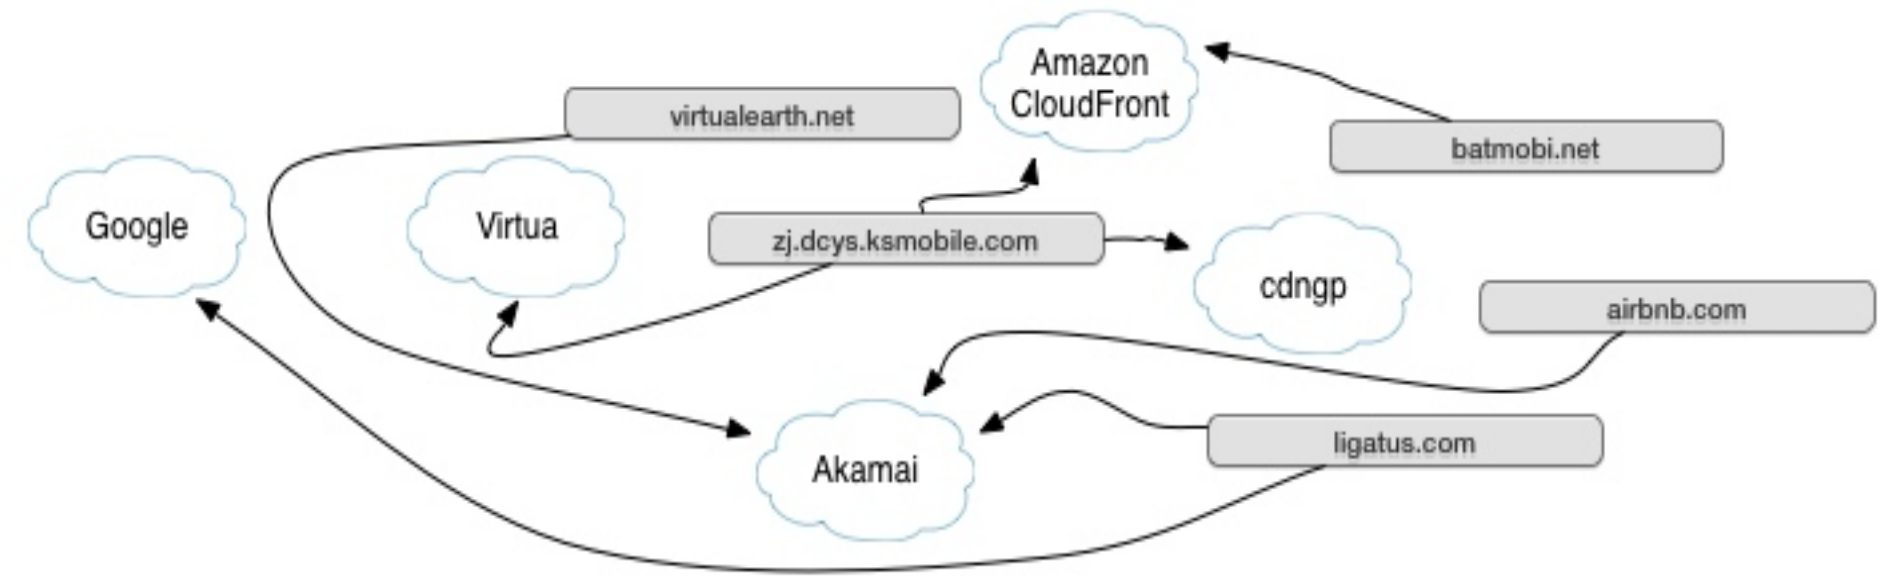
Results

- CSP app penetration: 83,7%
- Five large CSPs dominate the market



- Multi-CSP strategy detection:

Num-CSP	1	2	3	4	5
SLD(%)	85.15	13.5	1.14	0.17	0.04



Ongoing Work

- Measurements of ISP-CSPs inter-connection, replica selection and performance empirically
- Find peering and replica selection inefficiencies
- Identify cache deployments within ISP's internal address space

MONROE MobileBroadband Testbed:³

- Open access platform for mobile performance testing
- Suitable for independent, multihomed, and large-scale experimentation

Implemented Tests:

- DNS lookups for each CSP-hosted mobile domains from several DNS resolvers
- Traceroutes for IP addresses
- Measure the TCP connection time and UDP latency tests

Conclusions

- 5 CSPs provide support to 83.8% of apps and 44.9% domains.
- Most apps and on-line services use a single CSP provider.
- Observation of multi-CSPs strategy in several domains
- Future work: identify replica selection inefficiencies.

References

- Haystack project website: <https://www.haystack.mobi>
- CNDFinder website: <https://www.cdnplanet.com/tools/cdnfinder/>
- cdnDetector: Git repository, 2017. <https://github.com/Hossein-Doroud/cdn-detector.git>.
- Monroe's website: <https://www.monroe-project.eu/>