DNS Baseline Dynamics

Gautam Akiwate[1] **Mattijs Jonker**[2] **Ian Foster**[3]
Geoffrey Voelker[1] **Stefan Savage**[1]
[1]*University of California, San Diego* [2]*University of Twente* [3]*DNS Coffee*

RESEARCH QUESTIONS

Q. How **consistent** are configurations of different portions of the **DNS ecosystem**?

Implicit assumptions about relationships between domains, nameservers, and glue records that are not explicitly enforced lead to inconsistencies that create opportunities for attackers.

Q. How **distributed** is the **DNS ecosystem** across the **IPv4** address space?

Reliance on concentrated portions of the IPv4 address space can hinder reliability, robustness, and availability.

RESULTS

CONSISTENCY

NAMESERVERS WITH NO OR INVALID GLUE RR(s)

Category	+/-	Balance
1) All NS	+18,910,222	18,910,222
2) Fully Resolved	−16,689,058	2,221,164
3) Other TLDs	−1,319,388	901,776
4) SLD Resolution	−628,244	273,532
5) Hardcoded IP	−5,519	268,013
6) Invalid TLD/Malformed	−10,309	257,704
7) Dangling Delegations	−	257,704

- Out of 19M nameservers, **257k** have **no corresponding glue record** over 8 years. Of these, **45k** are **still active** as of April 2019. **99.4%** of the latter can **be registered today**, potentially **compromising** nearly 75k domains.

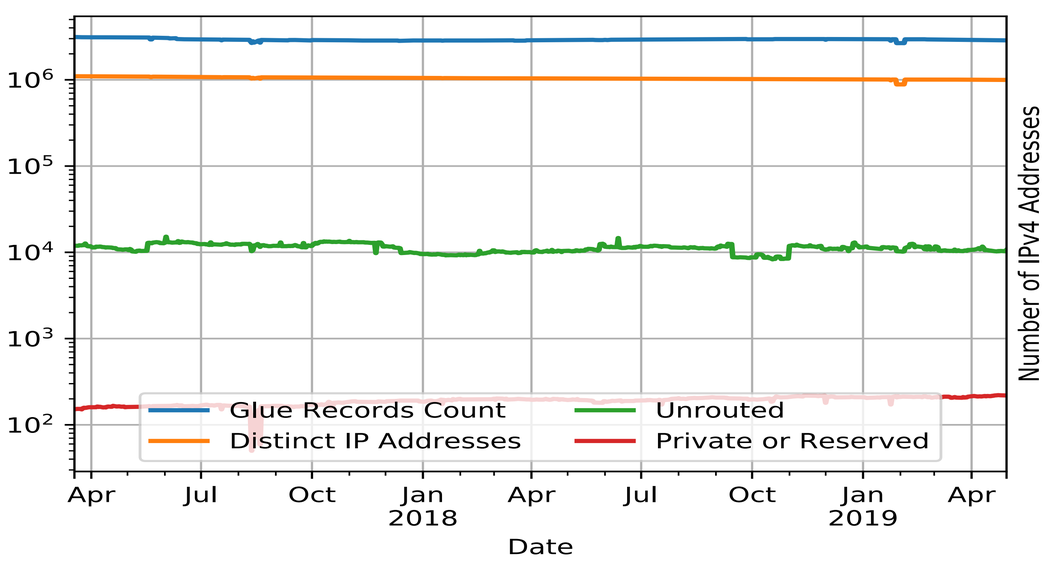

- Nearly 1% of glue RRs are unroutable, in private/ reserved address space, or do not follow RFCs.

DIVERSITY

CYCLES OF CONCENTRATION & DIVERSIFICATION

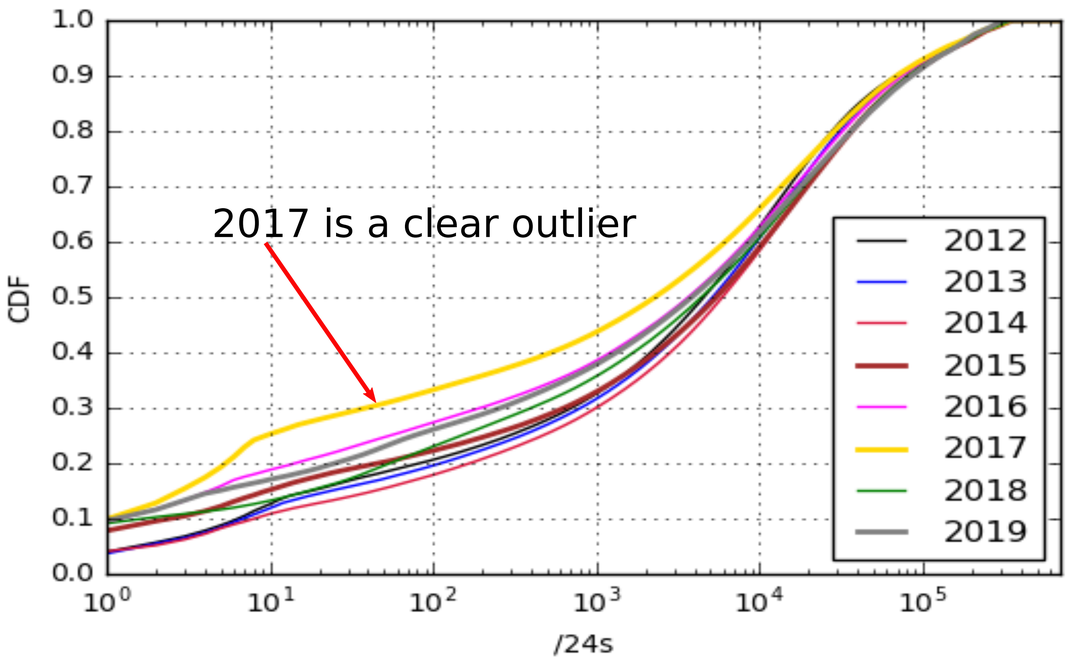

The **large skew** in **nameserver concentration** in **2017** is a result of **three ASNs:**

- **Bitcanal-AS**, hijacked dormant address space and sold it to malicious actors. In 2018, Bitcanal-AS was kicked off the Internet.
- Nearly all nameservers that point to the three ASNs belong to the **.US TLD** and look like **machine-generated domains.**
- The IP ranges routed by the three ASNs show up in **multiple blacklists**.
- **Handoff of nameservers** between the three ASNs as Bitcanal-AS gets shutdown.

DATASETS

DNS Coffee: Longitudinal Dataset of TLD Zone Files for legacy gTLDs, new gTLDs, and some ccTLDs collected over the last **8 years**.

Dataset	Domains	NS (NS)	IPv4 (A)	IPv6 (AAAA)
DNS Coffee	456.3 M	1879 M	4.8 M	8 K

OpenINTEL: Longitudinal Dataset of active DNS measurements for legacy gTLDs, new gTLDs, and some ccTLDs collected over the last **3 years**.

Dataset	Domains	NS (NS)	IPv (A)	IPv6 (AAAA)
OpenINTEL	276.6 M	14.1 M	2.1 M	153 K

CHALLENGES AND FUTURE WORK

Anycast Addressing: Use of anycast addressing can lead to overestimation of concentration.

Access to TLD zone files: while significantly improved as a result of ICANN CZDS, access is still not uniform.

UC San Diego | DNS Coffee | UNIVERSITY OF TWENTE.