DNS Coffee ☕

Ian Foster

$ whoami

Ian Foster

UCSD Graduate B.S./M.S. (2014/2015)
https://ian.ucsd.edu

DNS Researcher
https://dns.coffee

DNS Coffee

Collect, archive, and analyze TLD zone files daily

Provide current & historical data for researchers

Current web UI is minimal and not fully functional, still WIP

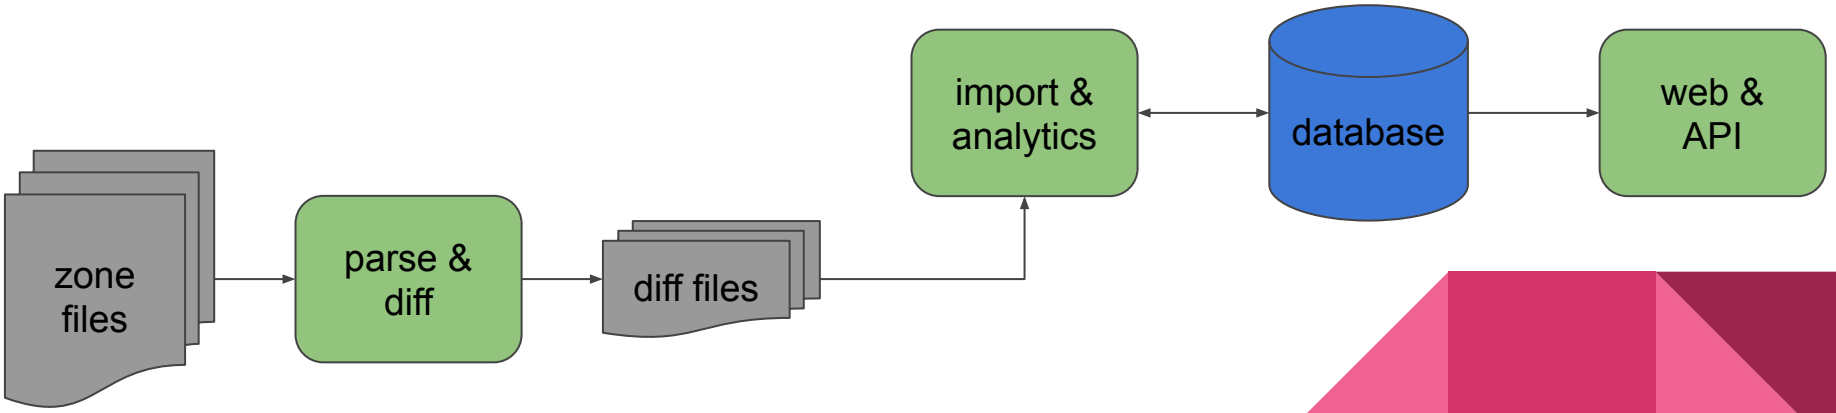

Brief History

- Started with 4 TLDs in April 2011, now index over 1258 TLDs daily.
  - > 500M unique domains
  - > 20M nameservers
  - > 5M IPs (90k IPv6)
  - > 2M zone files
  - > 3B individual DNS records
- Watched the ROOT zone grow from 306 TLDs in 2011 to 1516 today
  - Peaked at 1547 in mid 2017.
- Watched the birth and death of 62 TLDs
  - 46 fully indexed before EOL

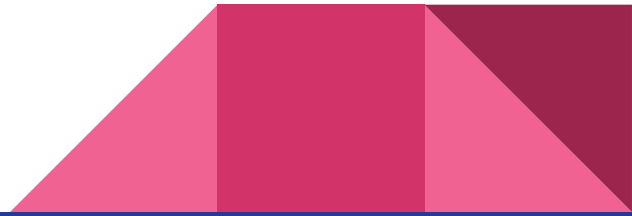

Data Sources

- Zone File Access Agreements (FTP)
- Centralized Zone Data Service (CZDS)
- Zone Transfers (AXFR)

- Currently index NS, A, AAAA records
  - Recently added support for partial SOA
  - Additional records on roadmap

DNS Coffee Data

- Each NS/A/AAAA record is stored in a relational table
  - Each has a first seen and last seen marker
- Allows for easy querying the data at a particular timestamp
  - Also saves a LOT of space
- Easily track changes over time

domain	nameserver	first_seen	last_seen	zone
COMMONLENS.NET	NS1.CAIDA.ORG	[NULL]	2012-10-19	NET
COMMONLENS.ORG	NS1.CAIDA.ORG	[NULL]	2012-10-19	ORG
NLANR.NET	NS1.CAIDA.ORG	[NULL]	2013-04-11	NET
DATCAT.ORG	NS1.CAIDA.ORG	[NULL]	2019-06-16	ORG
CAIDA.ORG	NS1.CAIDA.ORG	[NULL]	[NULL]	ORG

DNS Coffee Data

- Using DNS, you can query "up" the DNS "tree"
  - DNS Coffee allows you to also query "down" the tree.
    - IE: See all the domains a nameserver is responsible for
  - Can also match multiple NS sharing the same IP together
    - "Branded nameservers"
- Historical records of changes in zone files
- Find unregistered nameserver's domains
  - > 10k domains currently have unregistered name server's domains

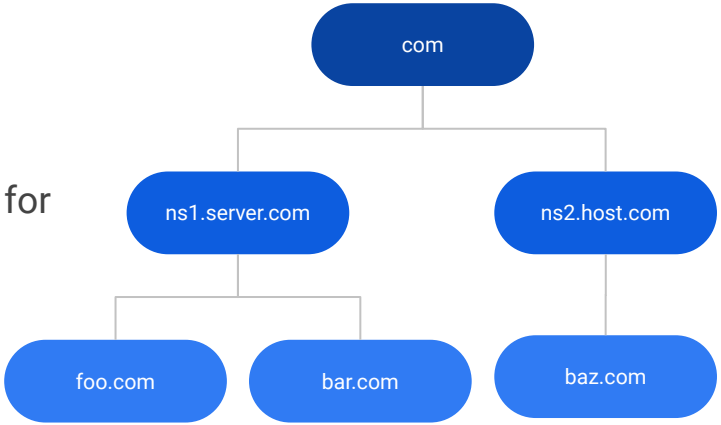

Zone File Weirdness

- Not all zone files conform to RFC Spec (US, BIZ, TEL)
- Misconfigurations
  - NS listed as IP (PSL)
  - IPs in private subnet or loopback
- Some zone files have different NS/A/AAAA records than the ROOT
- Some AXFR zones return different information than querying them for the same record
- SOA records do not always increase
- Some zone files contain records belonging to other zones
- Some zone files contain subdomains
- The same zone from multiple sources (CZDS/FTP/AXFR) may contain dramatically different records

ICANN Centralized Zone Data Service (CZDS)

- ICANN's portal to get zone file access for the new gTLDs (1200 zones)
  - Now being used by other TLDs such as NET, BIZ, CAT
  - Refreshed in January to provide a REST API
    - Only download endpoints documented
- Process
  - Users request access to zone files
  - Registries receive and review requests, may approve or deny
    - Approvals can have an expiration date
  - Once approved, user can download zone files until expiration date
  - After expiration users need to request again
  - Can file a complaint to ICANN for unresponsive registries
- Users must wait for access to expire before submitting a new request
  - Forces users to lose access while waiting on the regiesteries
- Registries may take as long as they like to respond to requests

Zone Transfers (AXFR)

- Some TLD name servers allow AXFR of their zones intentionally, most don't.
- Some are misconfigured and may allow AXFR on just one of their nameservers. Sometimes just on a single IP for a single server
- Often times this is a temporary misconfiguration and is fixed shortly.
- On a given day, ~ 30 different TLDs allow AXFR on one of their servers
  - mostly ccTLDs
- The same methodology can be used on all the nameservers in a zone
  - ~1% success rate per domain in a given zone file

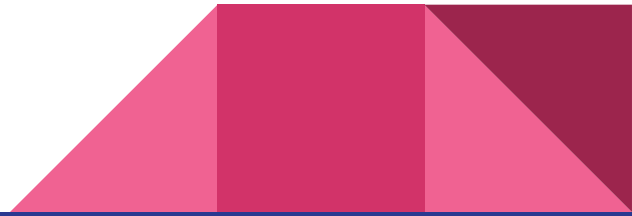

Tooling

CZDS Client https://github.com/lanrat/czds

A utility and golang library implementing a client to the CZDS REST API using both the documented and undocumented API endpoints

All XFR https://github.com/lanrat/allxfr

Tool to perform opportunistic zone transfers (AXFR) requests against every domain/nameserver/IP in a zone

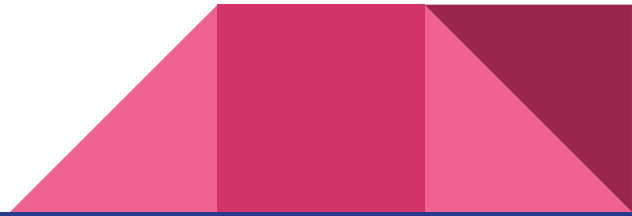

Prior Research Utilizing DNS Coffee

- DNS Baseline Dynamics
  - ACM IMC 2019 Poster: Gautam Akiwate, Mattijs Jonker, Ian Foster, Stefan Savage, Geoffrey M. Voelker
- BygoneSSL: dealing with residual certificates for pre-owned domains
  - DEFCON 26 & ToorCon 20: Ian Foster
- CertGraph: Crawling the Graph of SSL Certificate Alternate Names using CT
  - ShmooCon 2018: Ian Foster
- From .academy to .zone: An Analysis of the New TLD Land Rush
  - ACM IMC 2015: Tristan Halvorson, Matthew F. Der, Ian Foster, Stefan Savage, Lawrence K. Saul, Geoffrey M. Voelker
- Who is .com? Learning to Parse WHOIS Records
  - ACM IMC 2015: Suqi Liu, Ian Foster, Stefan Savage, Geoffrey M. Voelker, Lawrence K. Saul
- Pillaging DVCS Repos for Fun and Profit
  - Defcon 19 & ToorCon 13: Adam Baldwin