

facebook

DNS Infrastructure @ FB

Manu Bretelle

Production Engineer

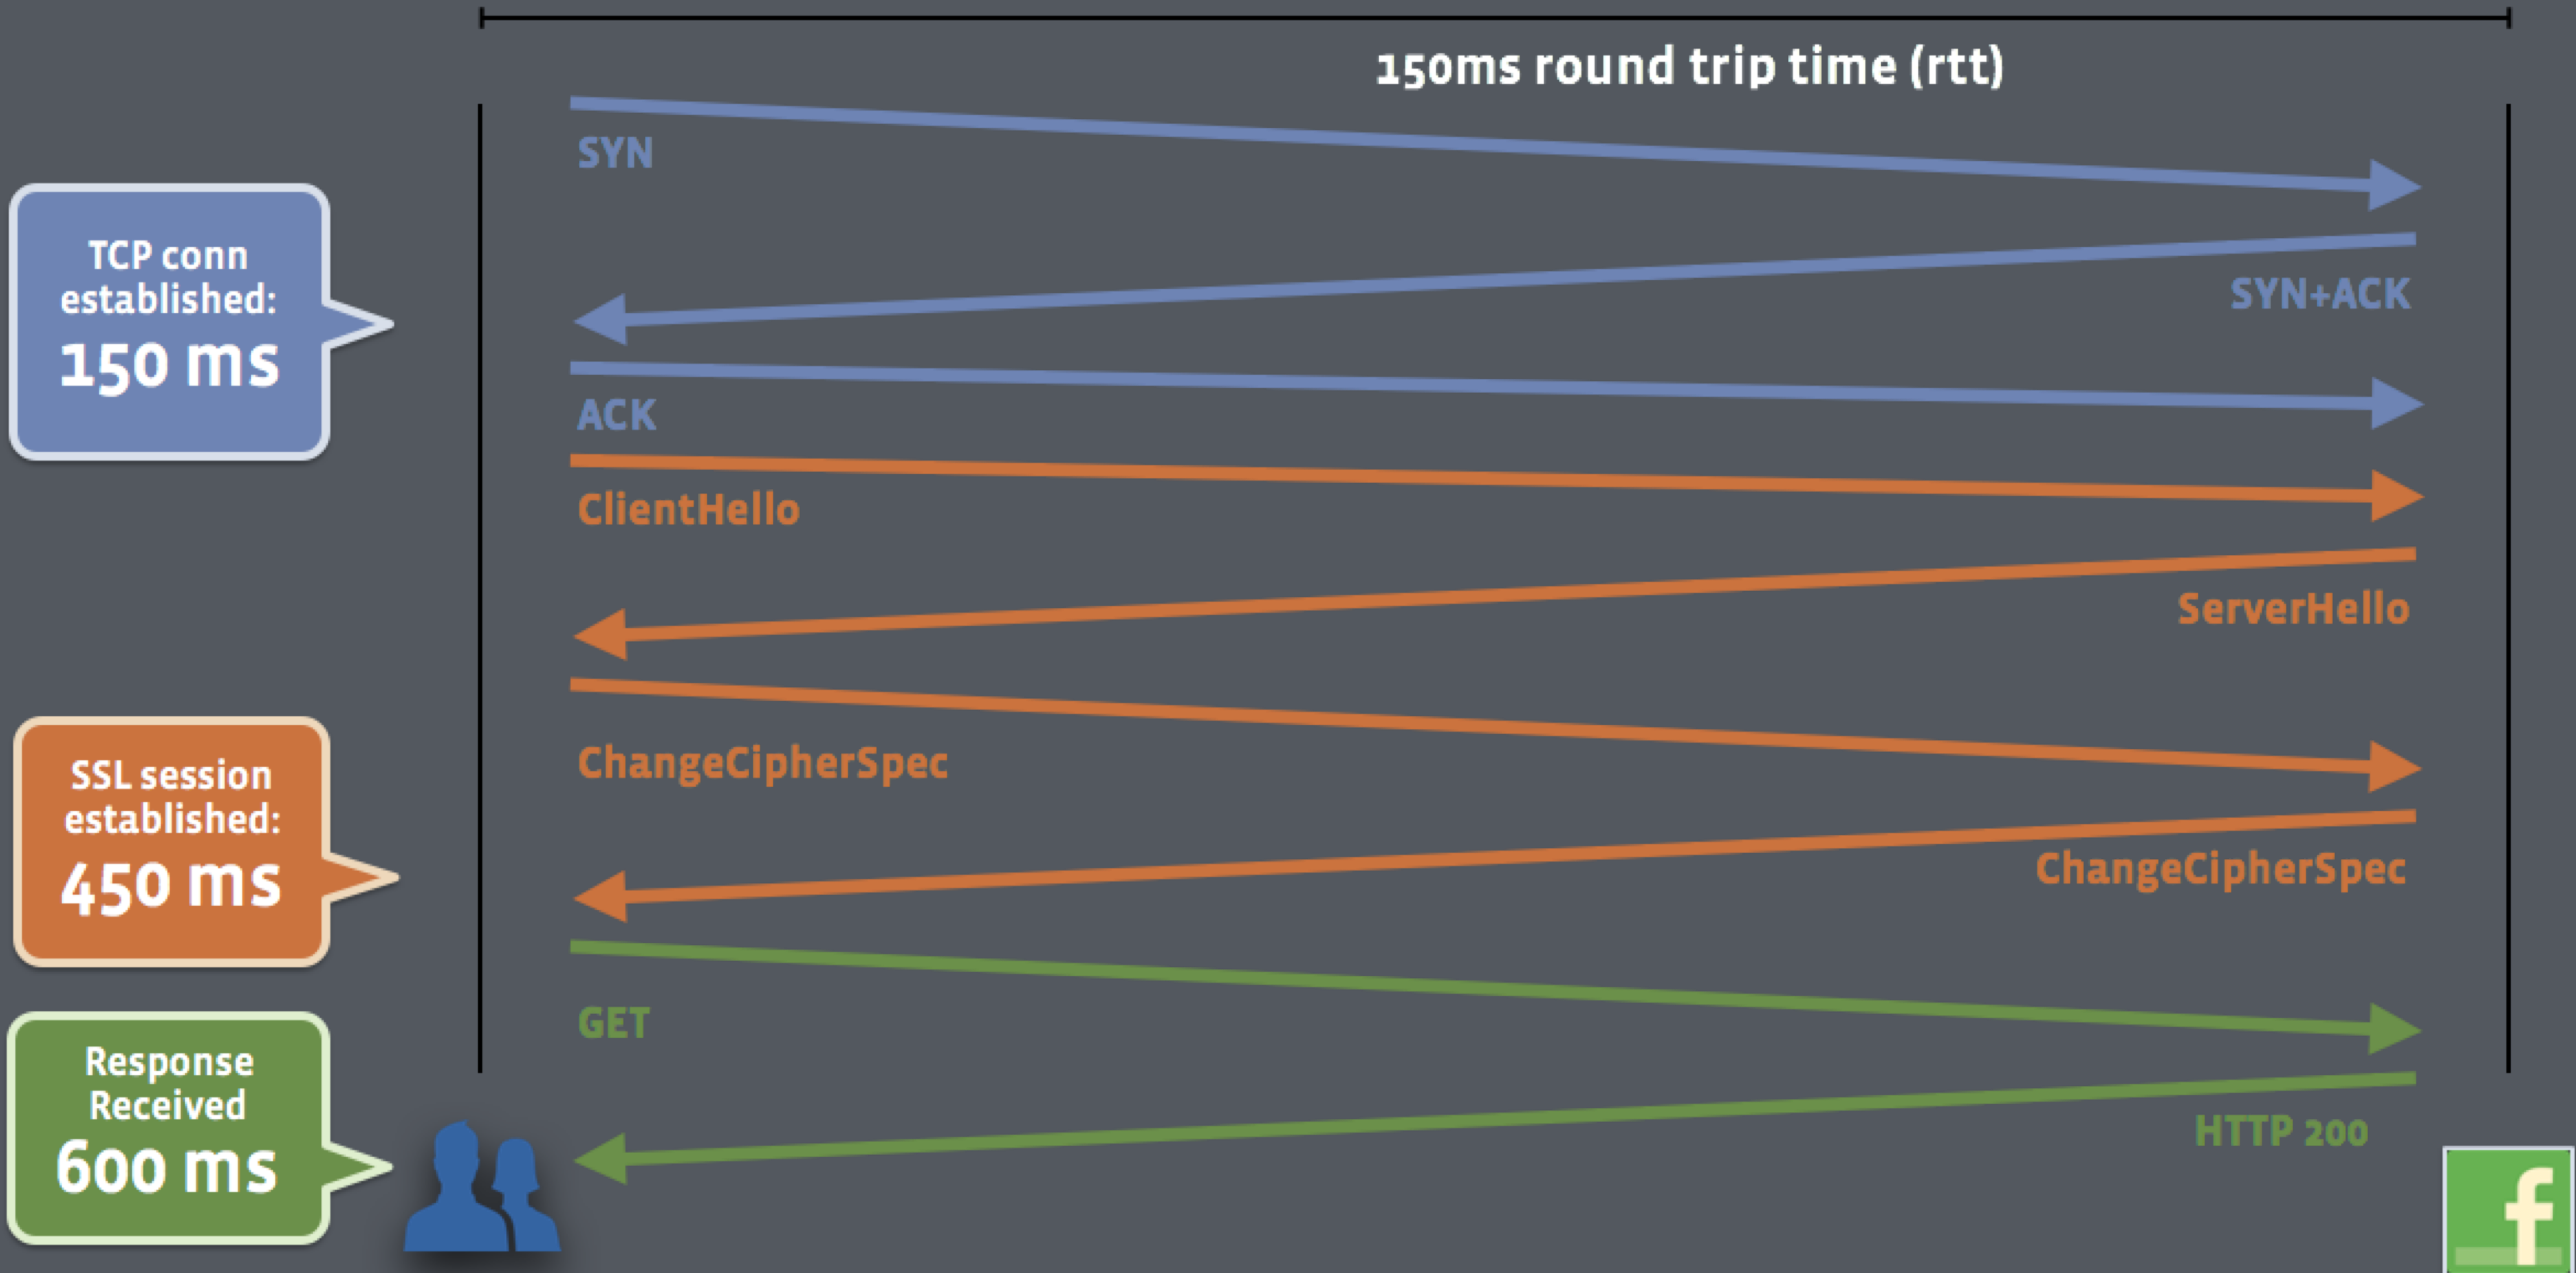
Who am I

- Traffic Team
- At FB for 5.380821917808219 years

A stylized world map is centered in the background, rendered in a light yellowish-gold color. Overlaid on the map is a complex network of thin, light blue lines that connect various points across the continents, suggesting a global communication or data network. The entire scene is set against a dark blue background with a subtle, darker blue grid pattern.

**To give people the power to build community
and bring the world closer together.**

HTTPS Seoul -> Oregon



Edge POPs




 Datacenter

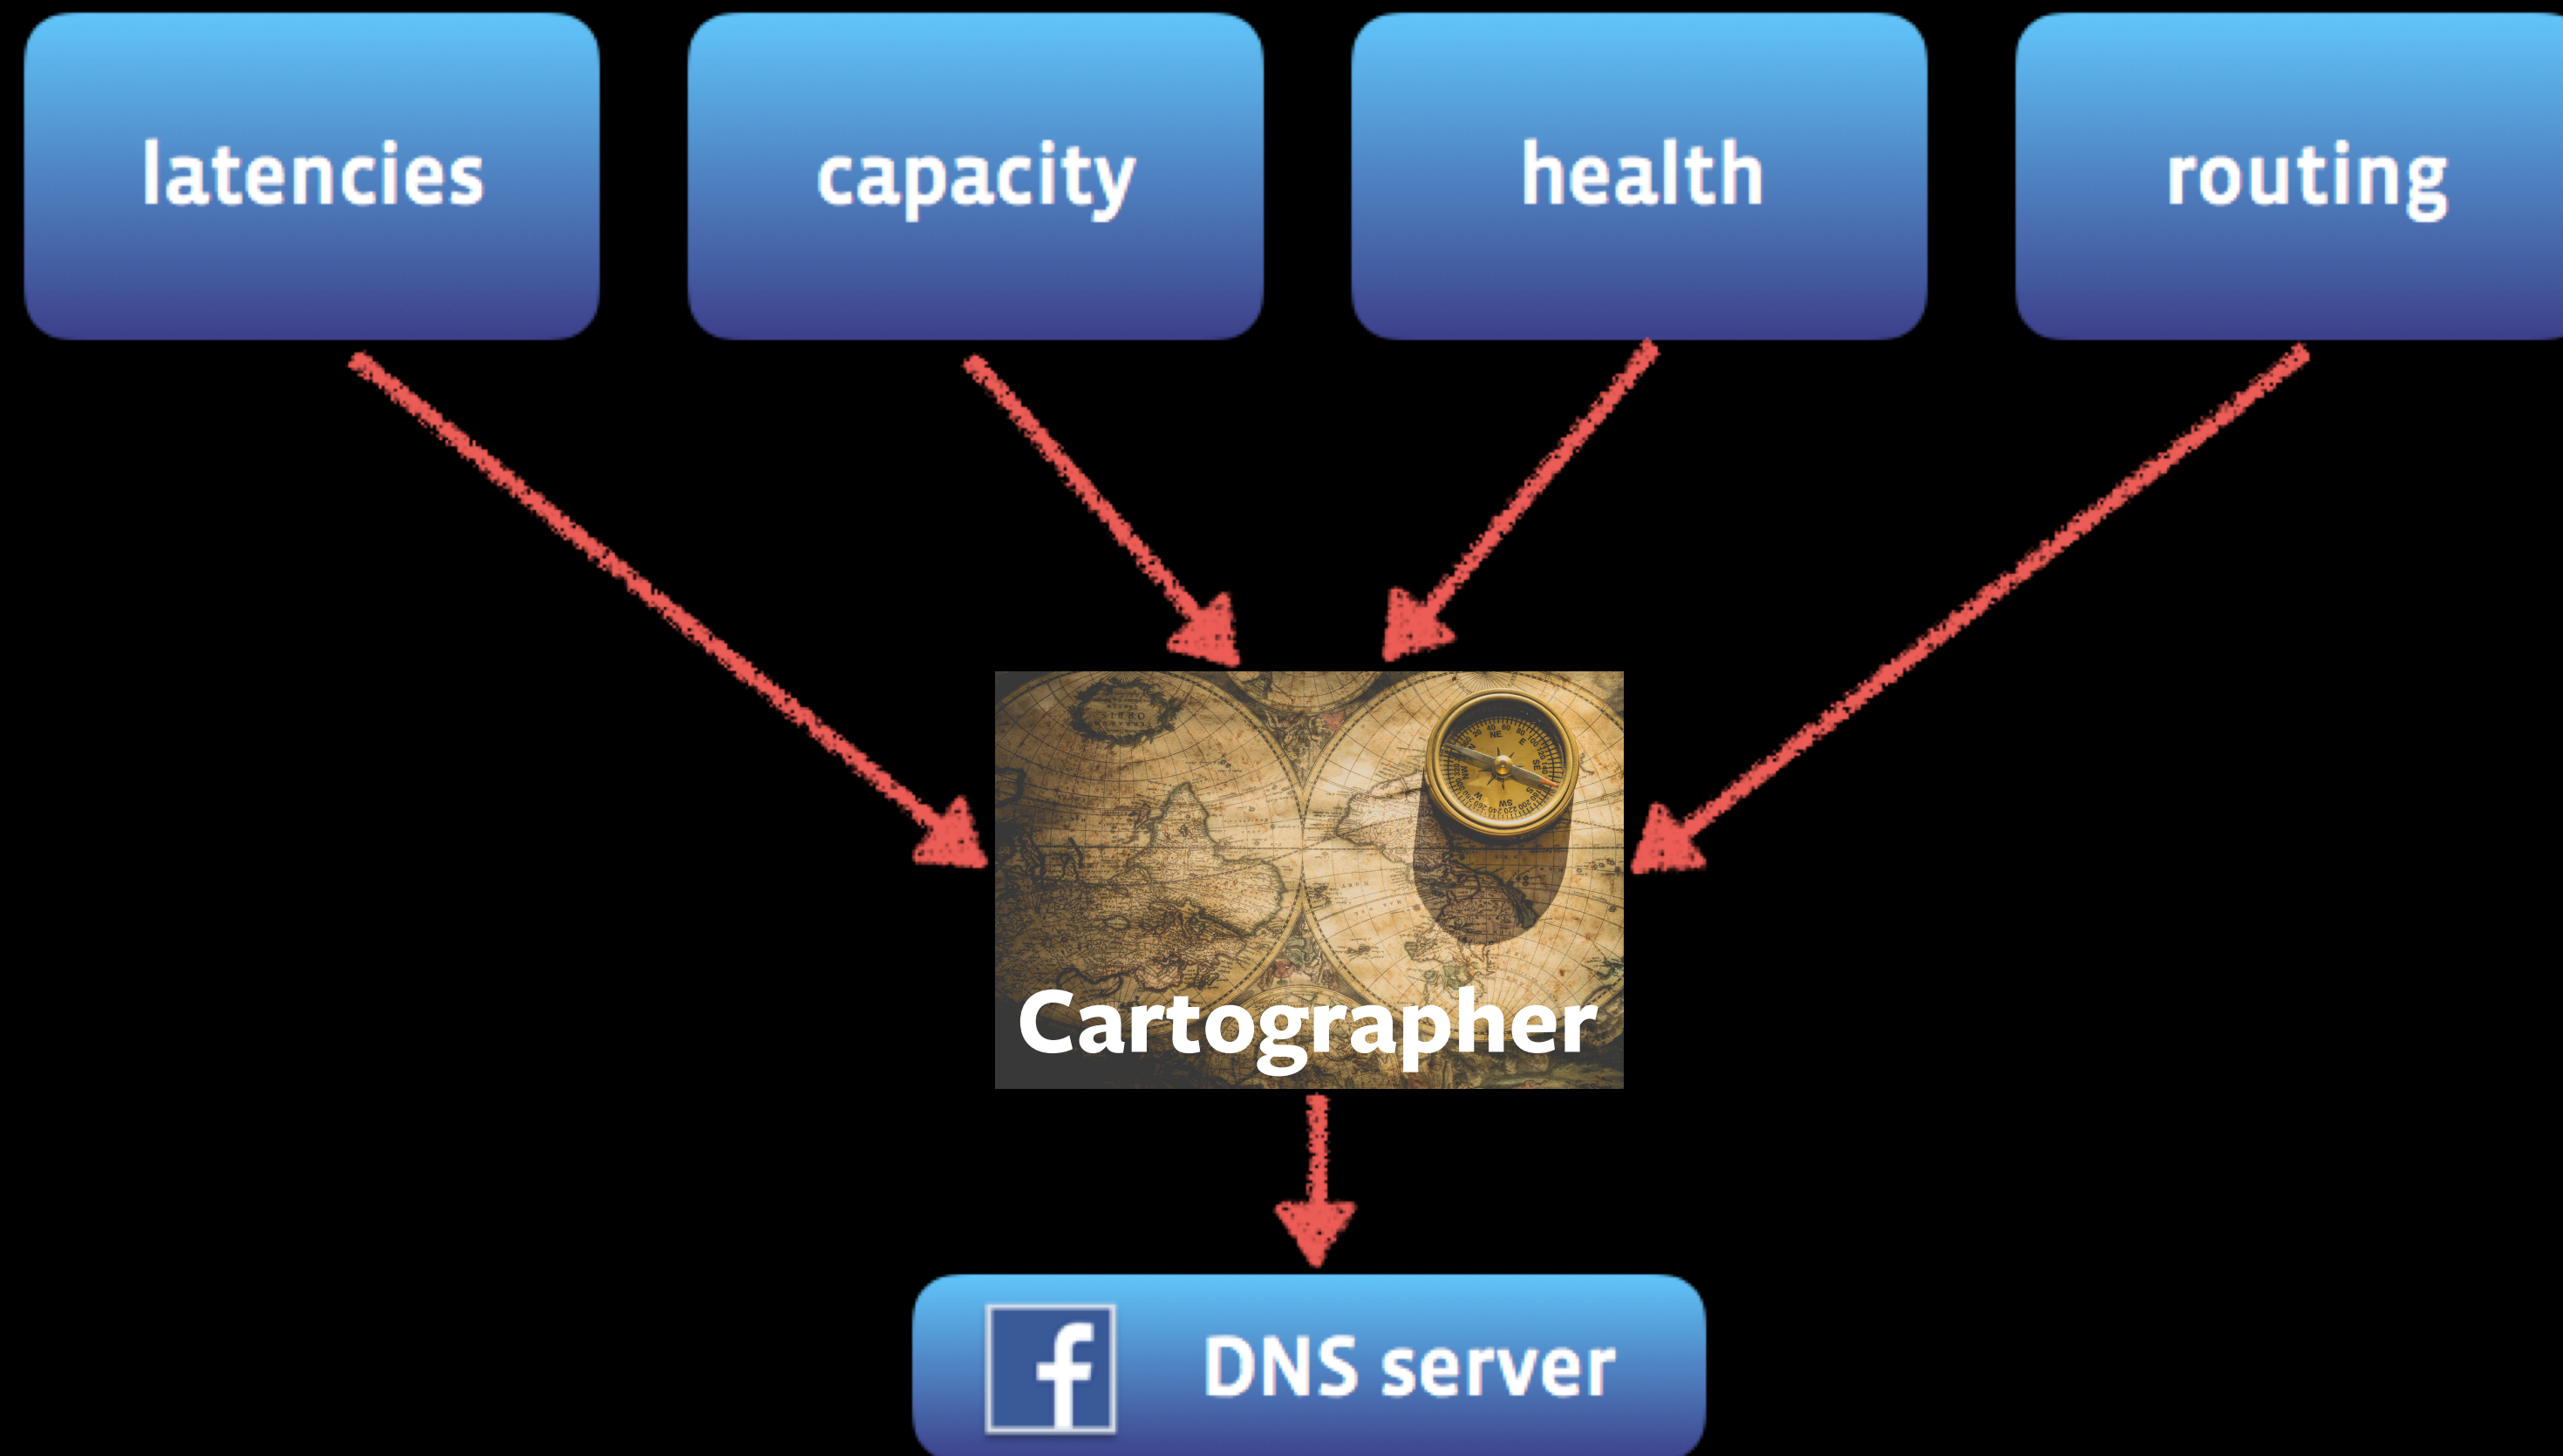
 POP

*POP = points of presence. Map doesn't reflect actual POP locations.

HTTPS Seoul->Tokyo->Oregon



Global Load Balancing



Resolver Allocation

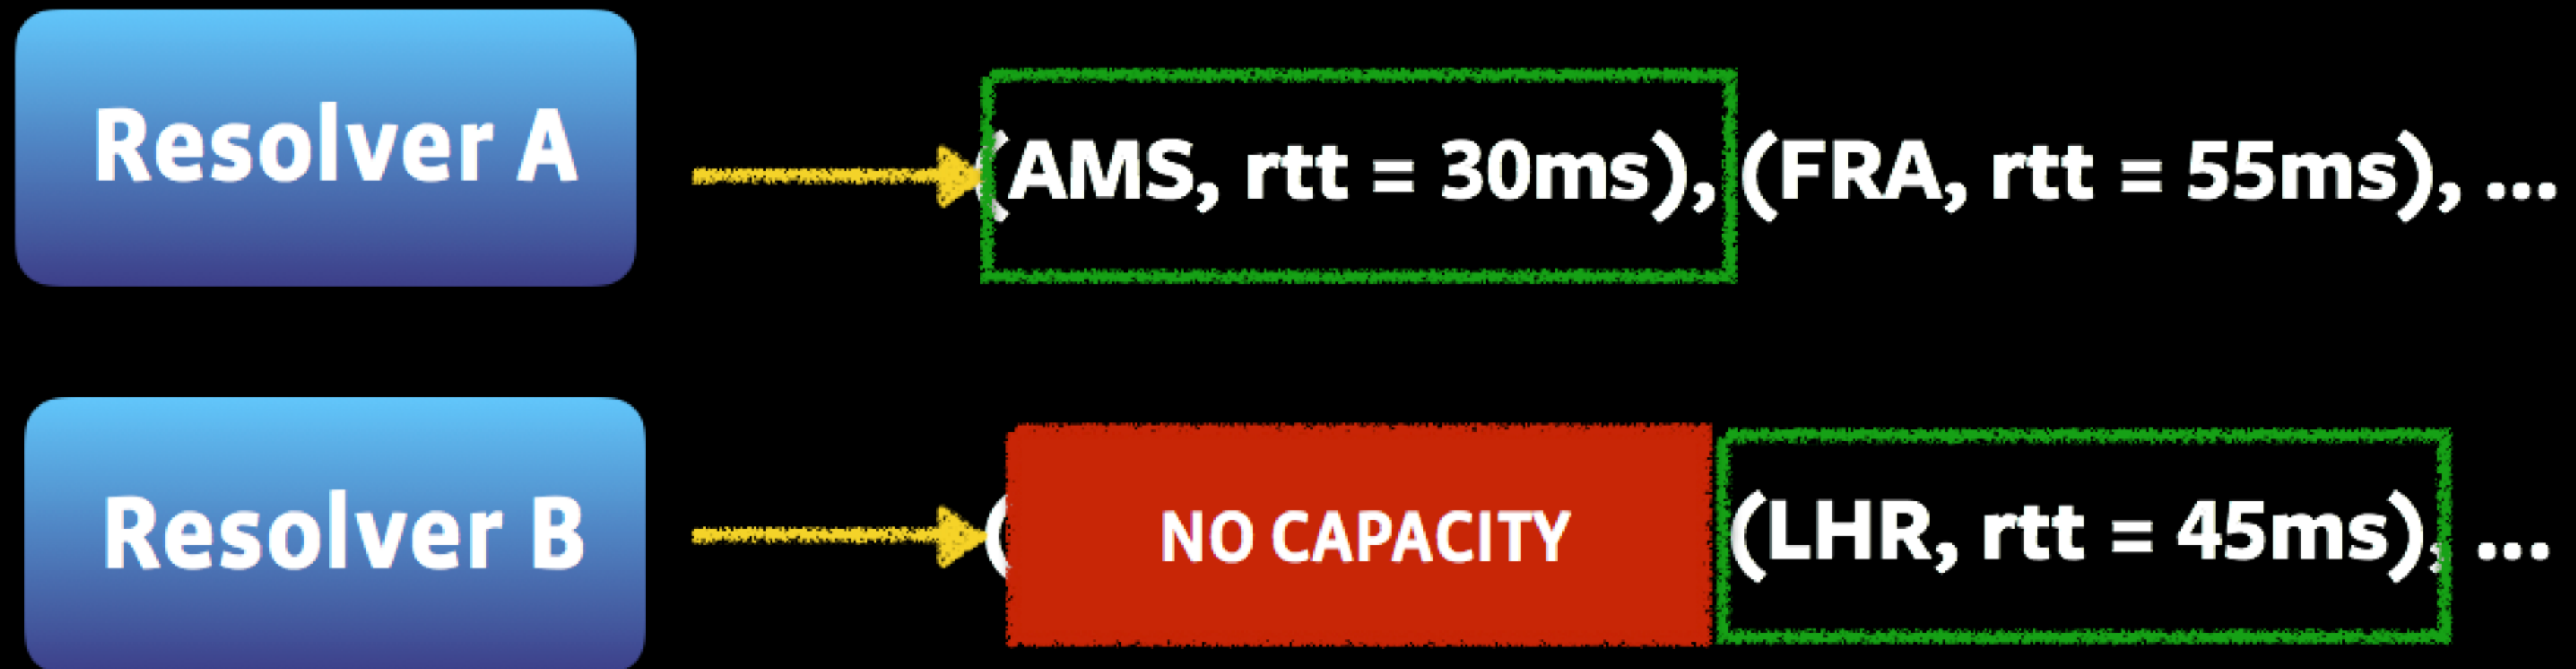
Resolver A

→ (AMS, rtt = 30ms), (FRA, rtt = 55ms), ...

Resolver B


→ (FRA, rtt = 35ms), (LHR, rtt = 45ms), ...

Resolver Allocation




+ rules!

Sonar



<https://fc7b506e240KEYQAAAA47YABRIEzD42UETHK7Y-szxlfecf1403573393-sonar.xx.fbcdn.net/thumb.jpg>

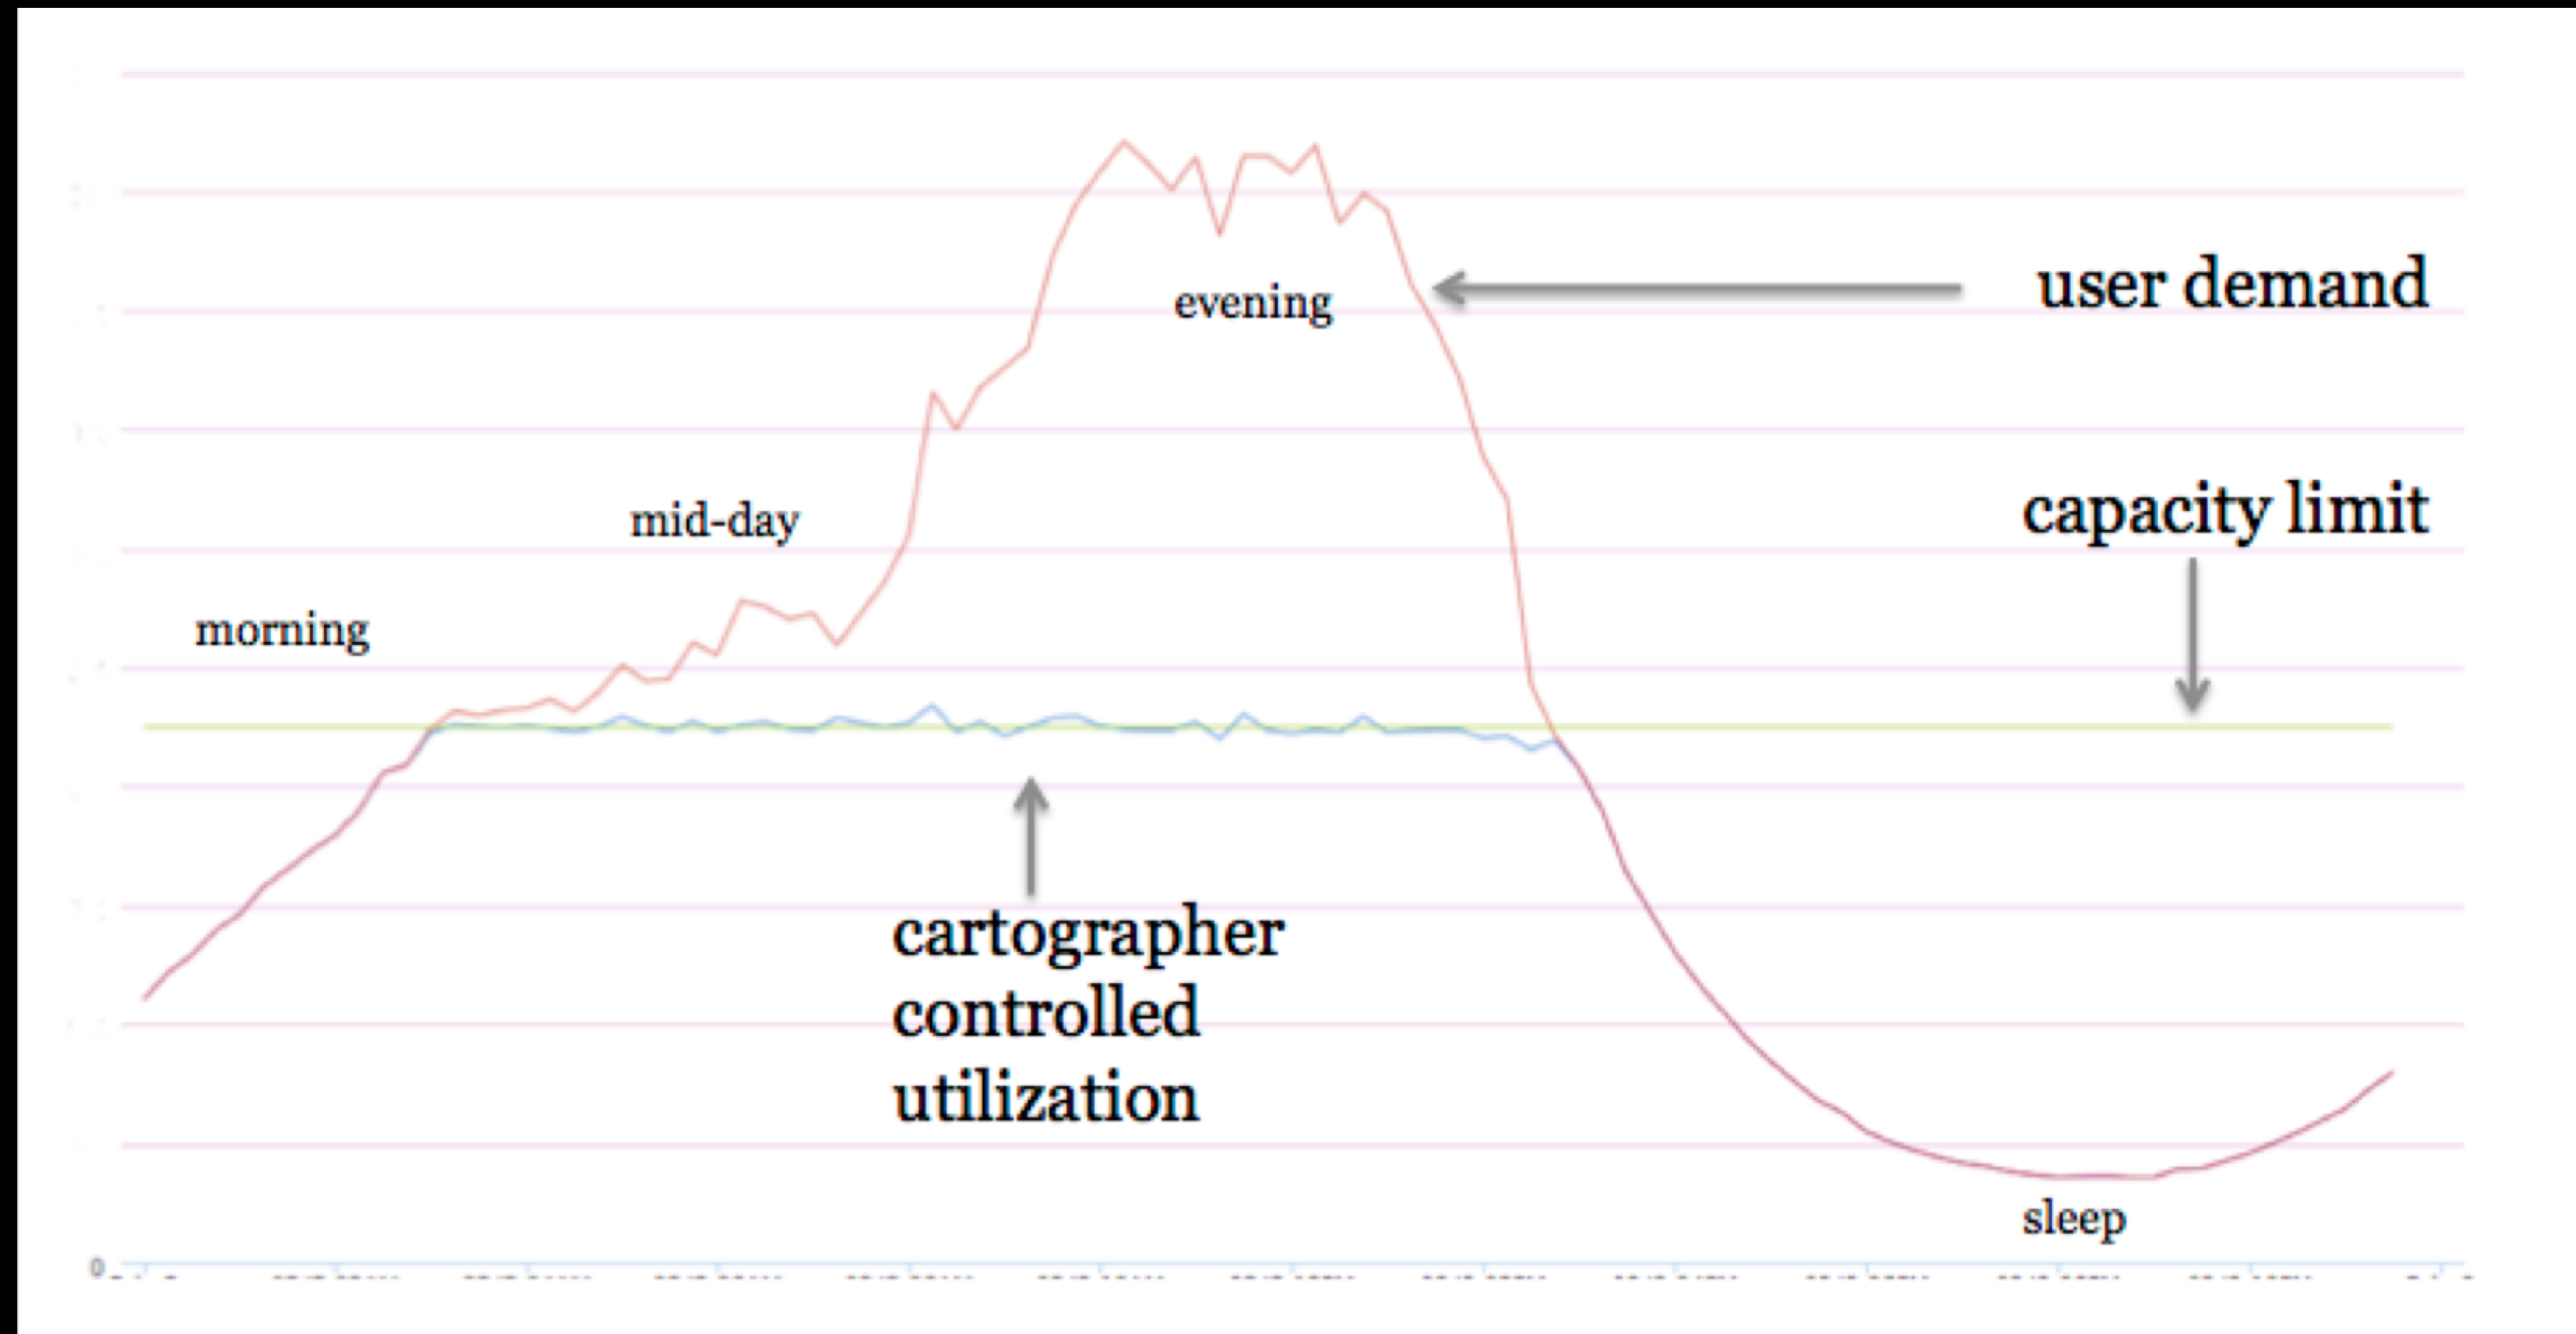


DNS log: (unique hostname, resolver IP)
proxygen log: (unique hostname, client IP, edge location, RTT)



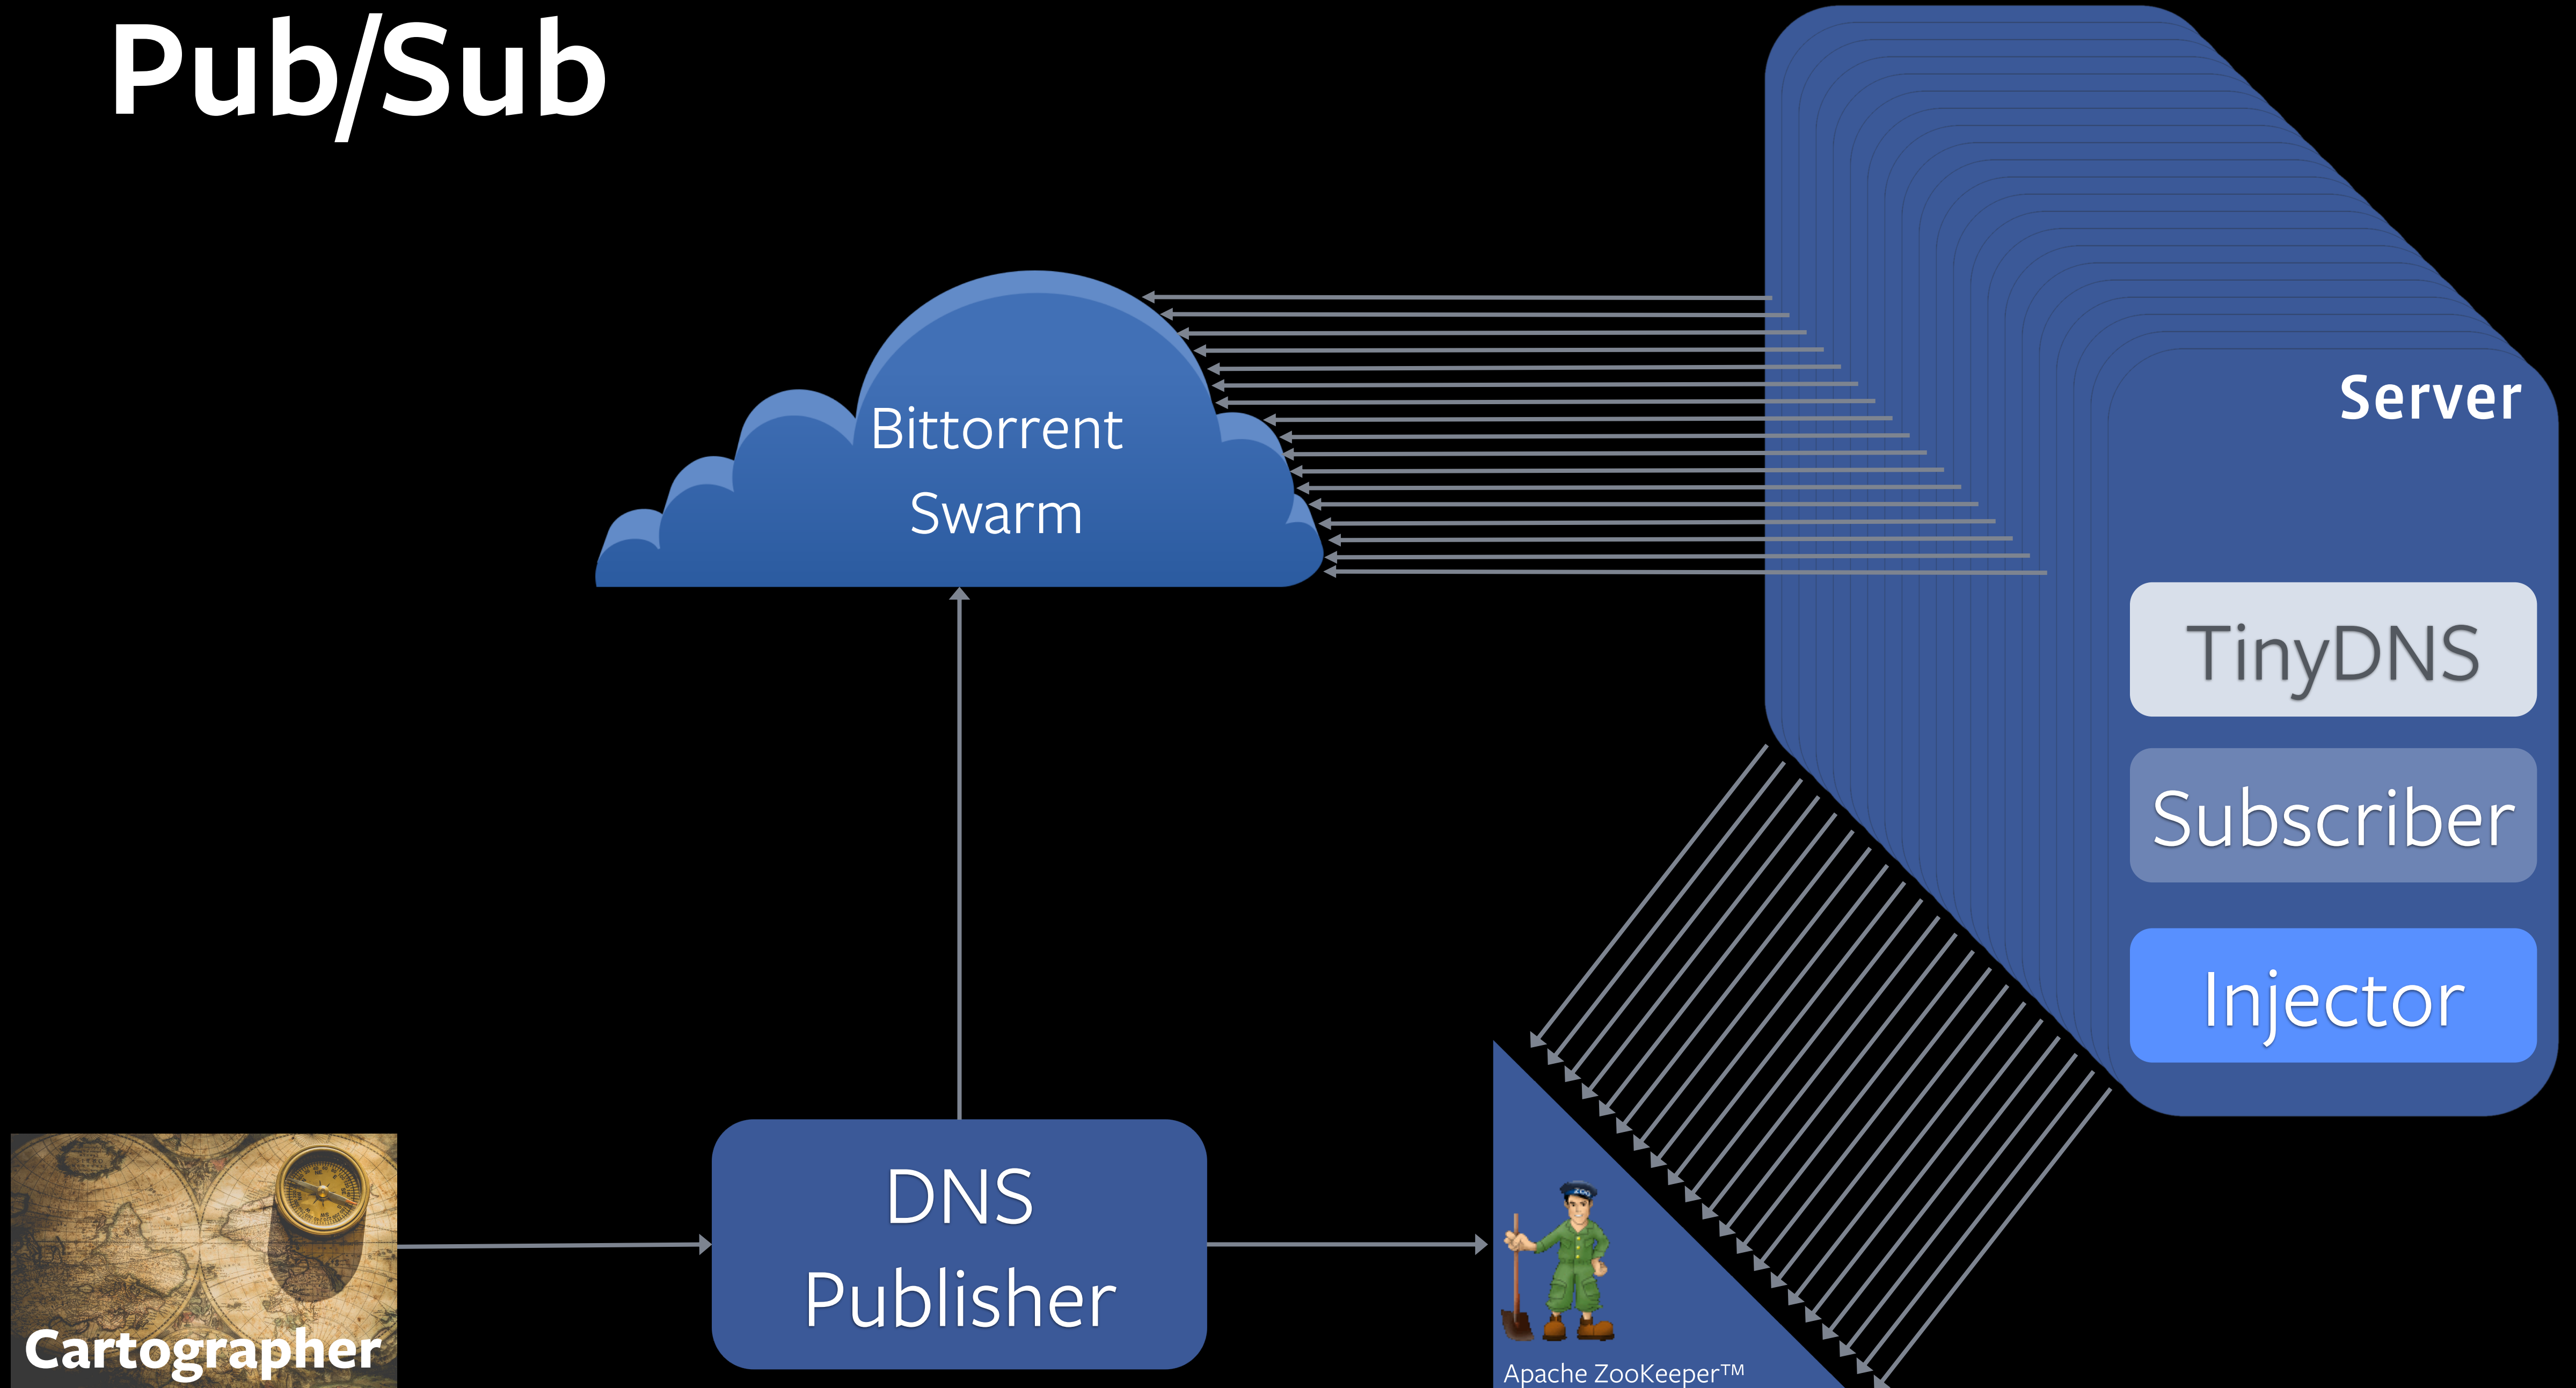
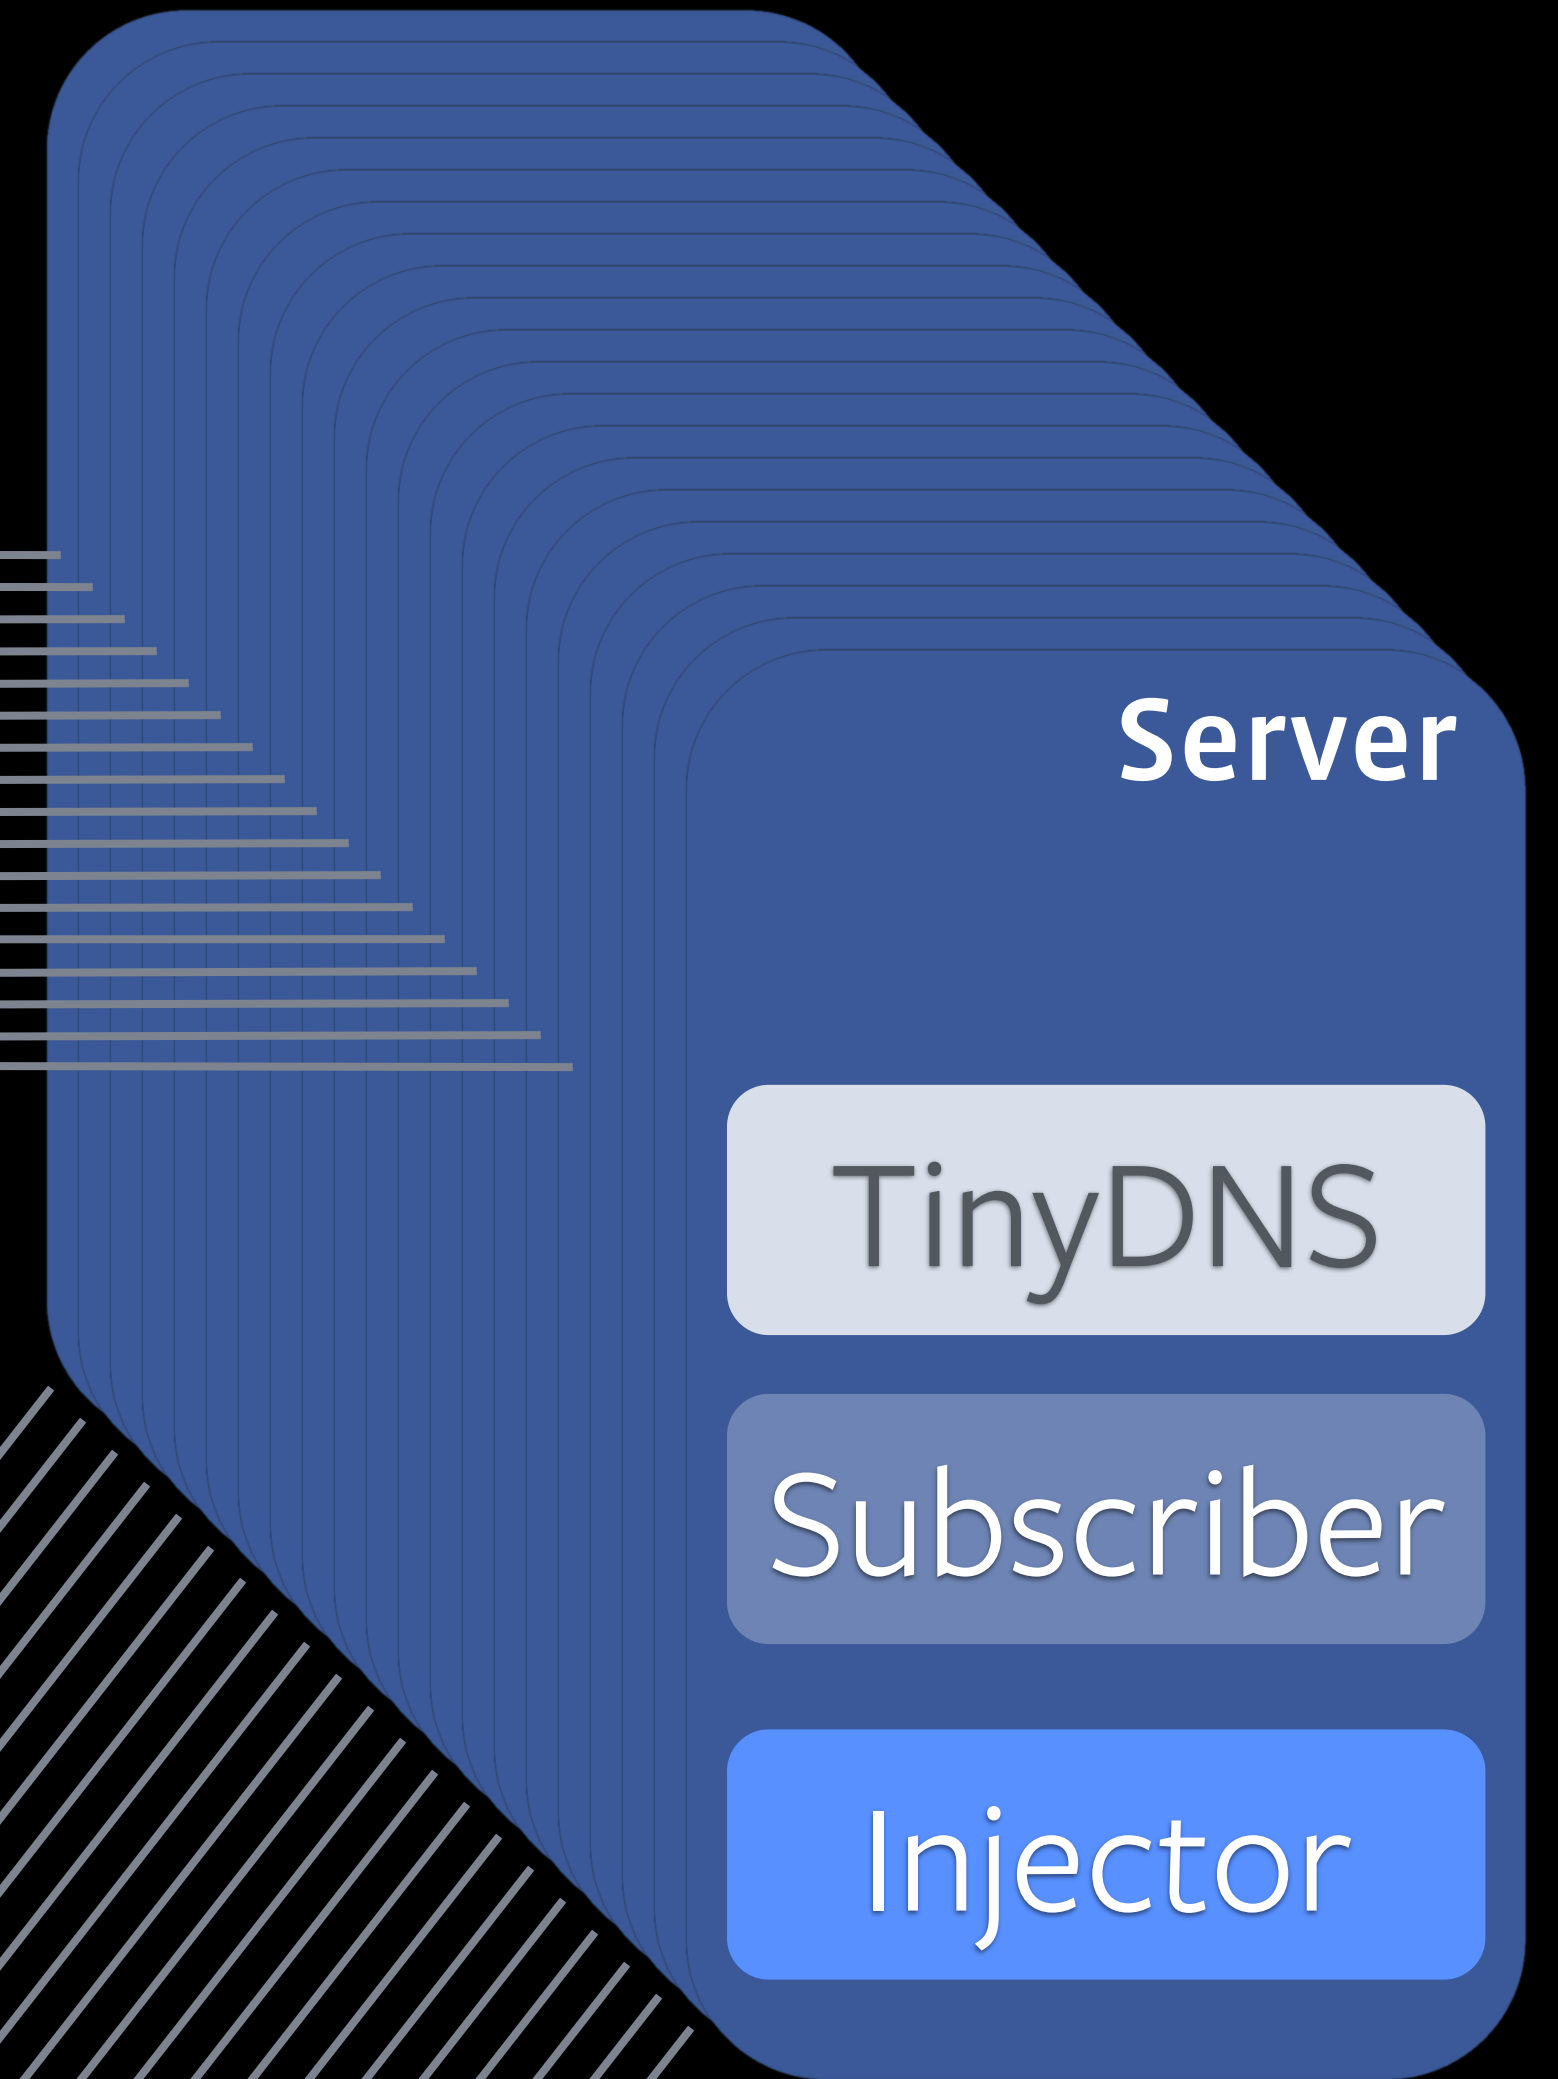

Joined: (resolver IP, edge location, client IP, RTT)

C10R managed load



Deploying DNS maps

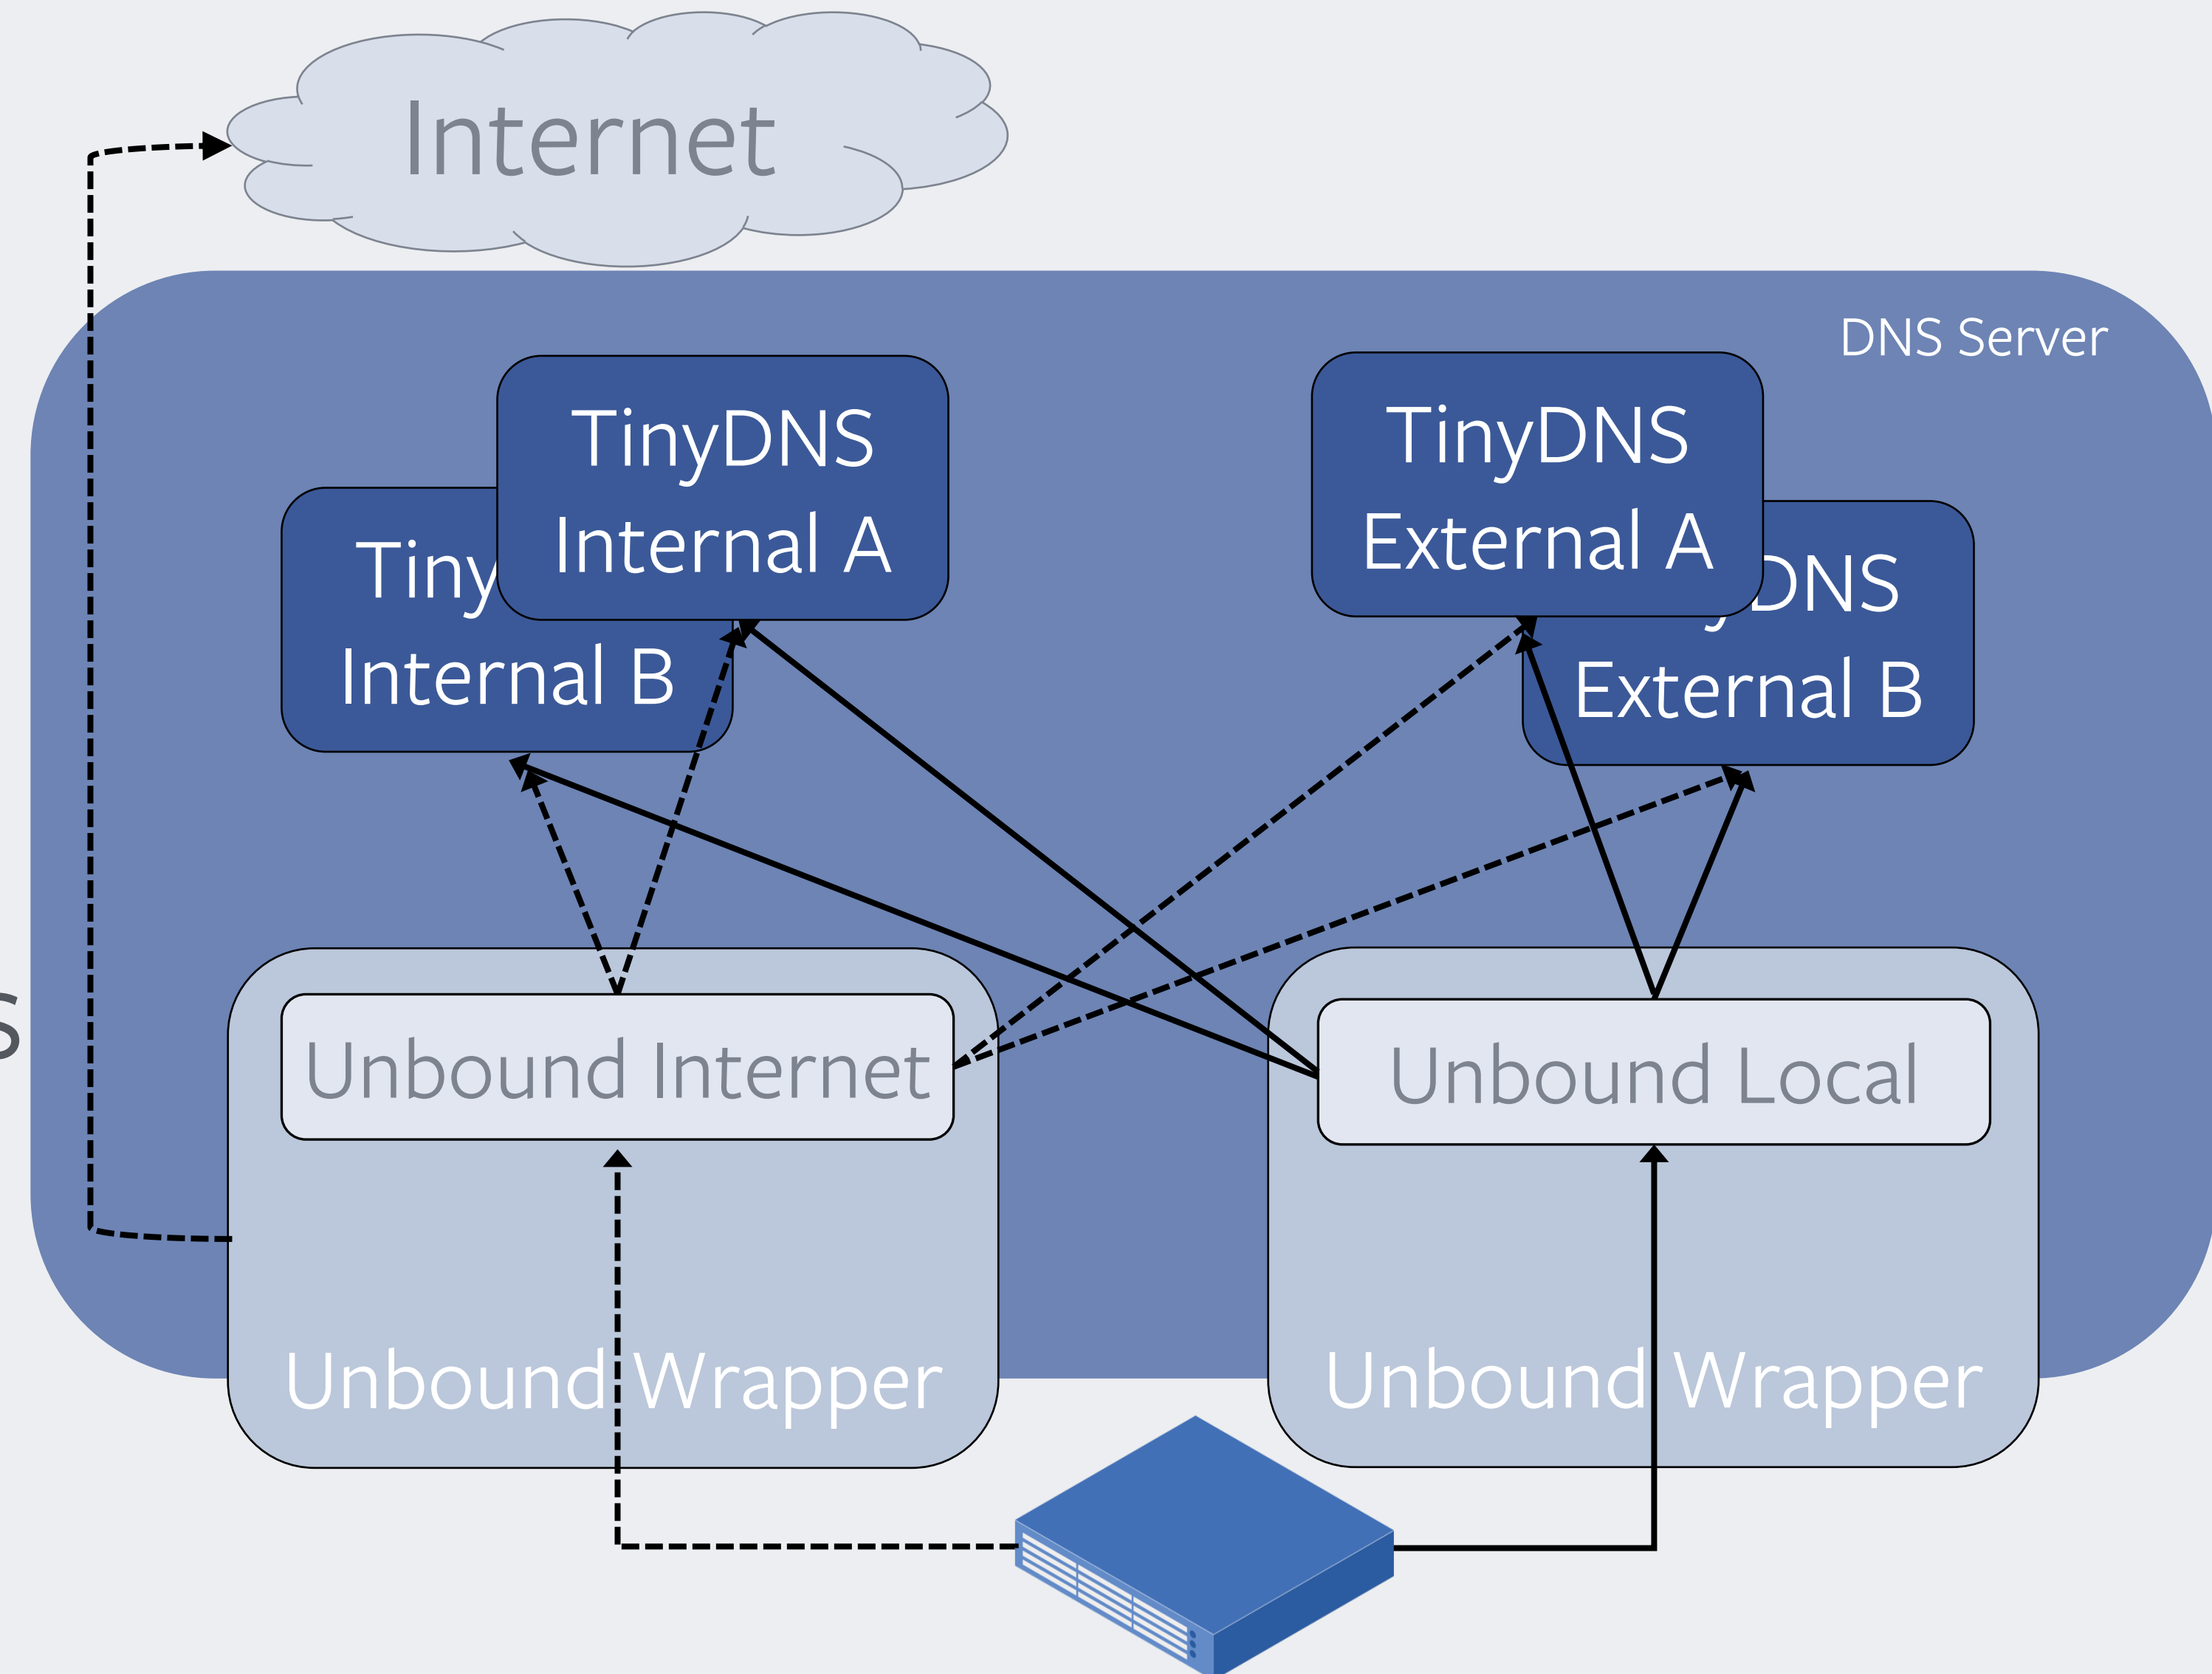
Pub/Sub



Internal DNS Infra


Recursive resolver

- Stub-zones
- Healthchecks
- Advertise VIP
- Monitor configs
- Rolled restart




Anycast VIP Advertisement

- BGP
- ECMP
- ExaBGP between servers and TOR



Who makes changes

Engineers






Servers Themselves

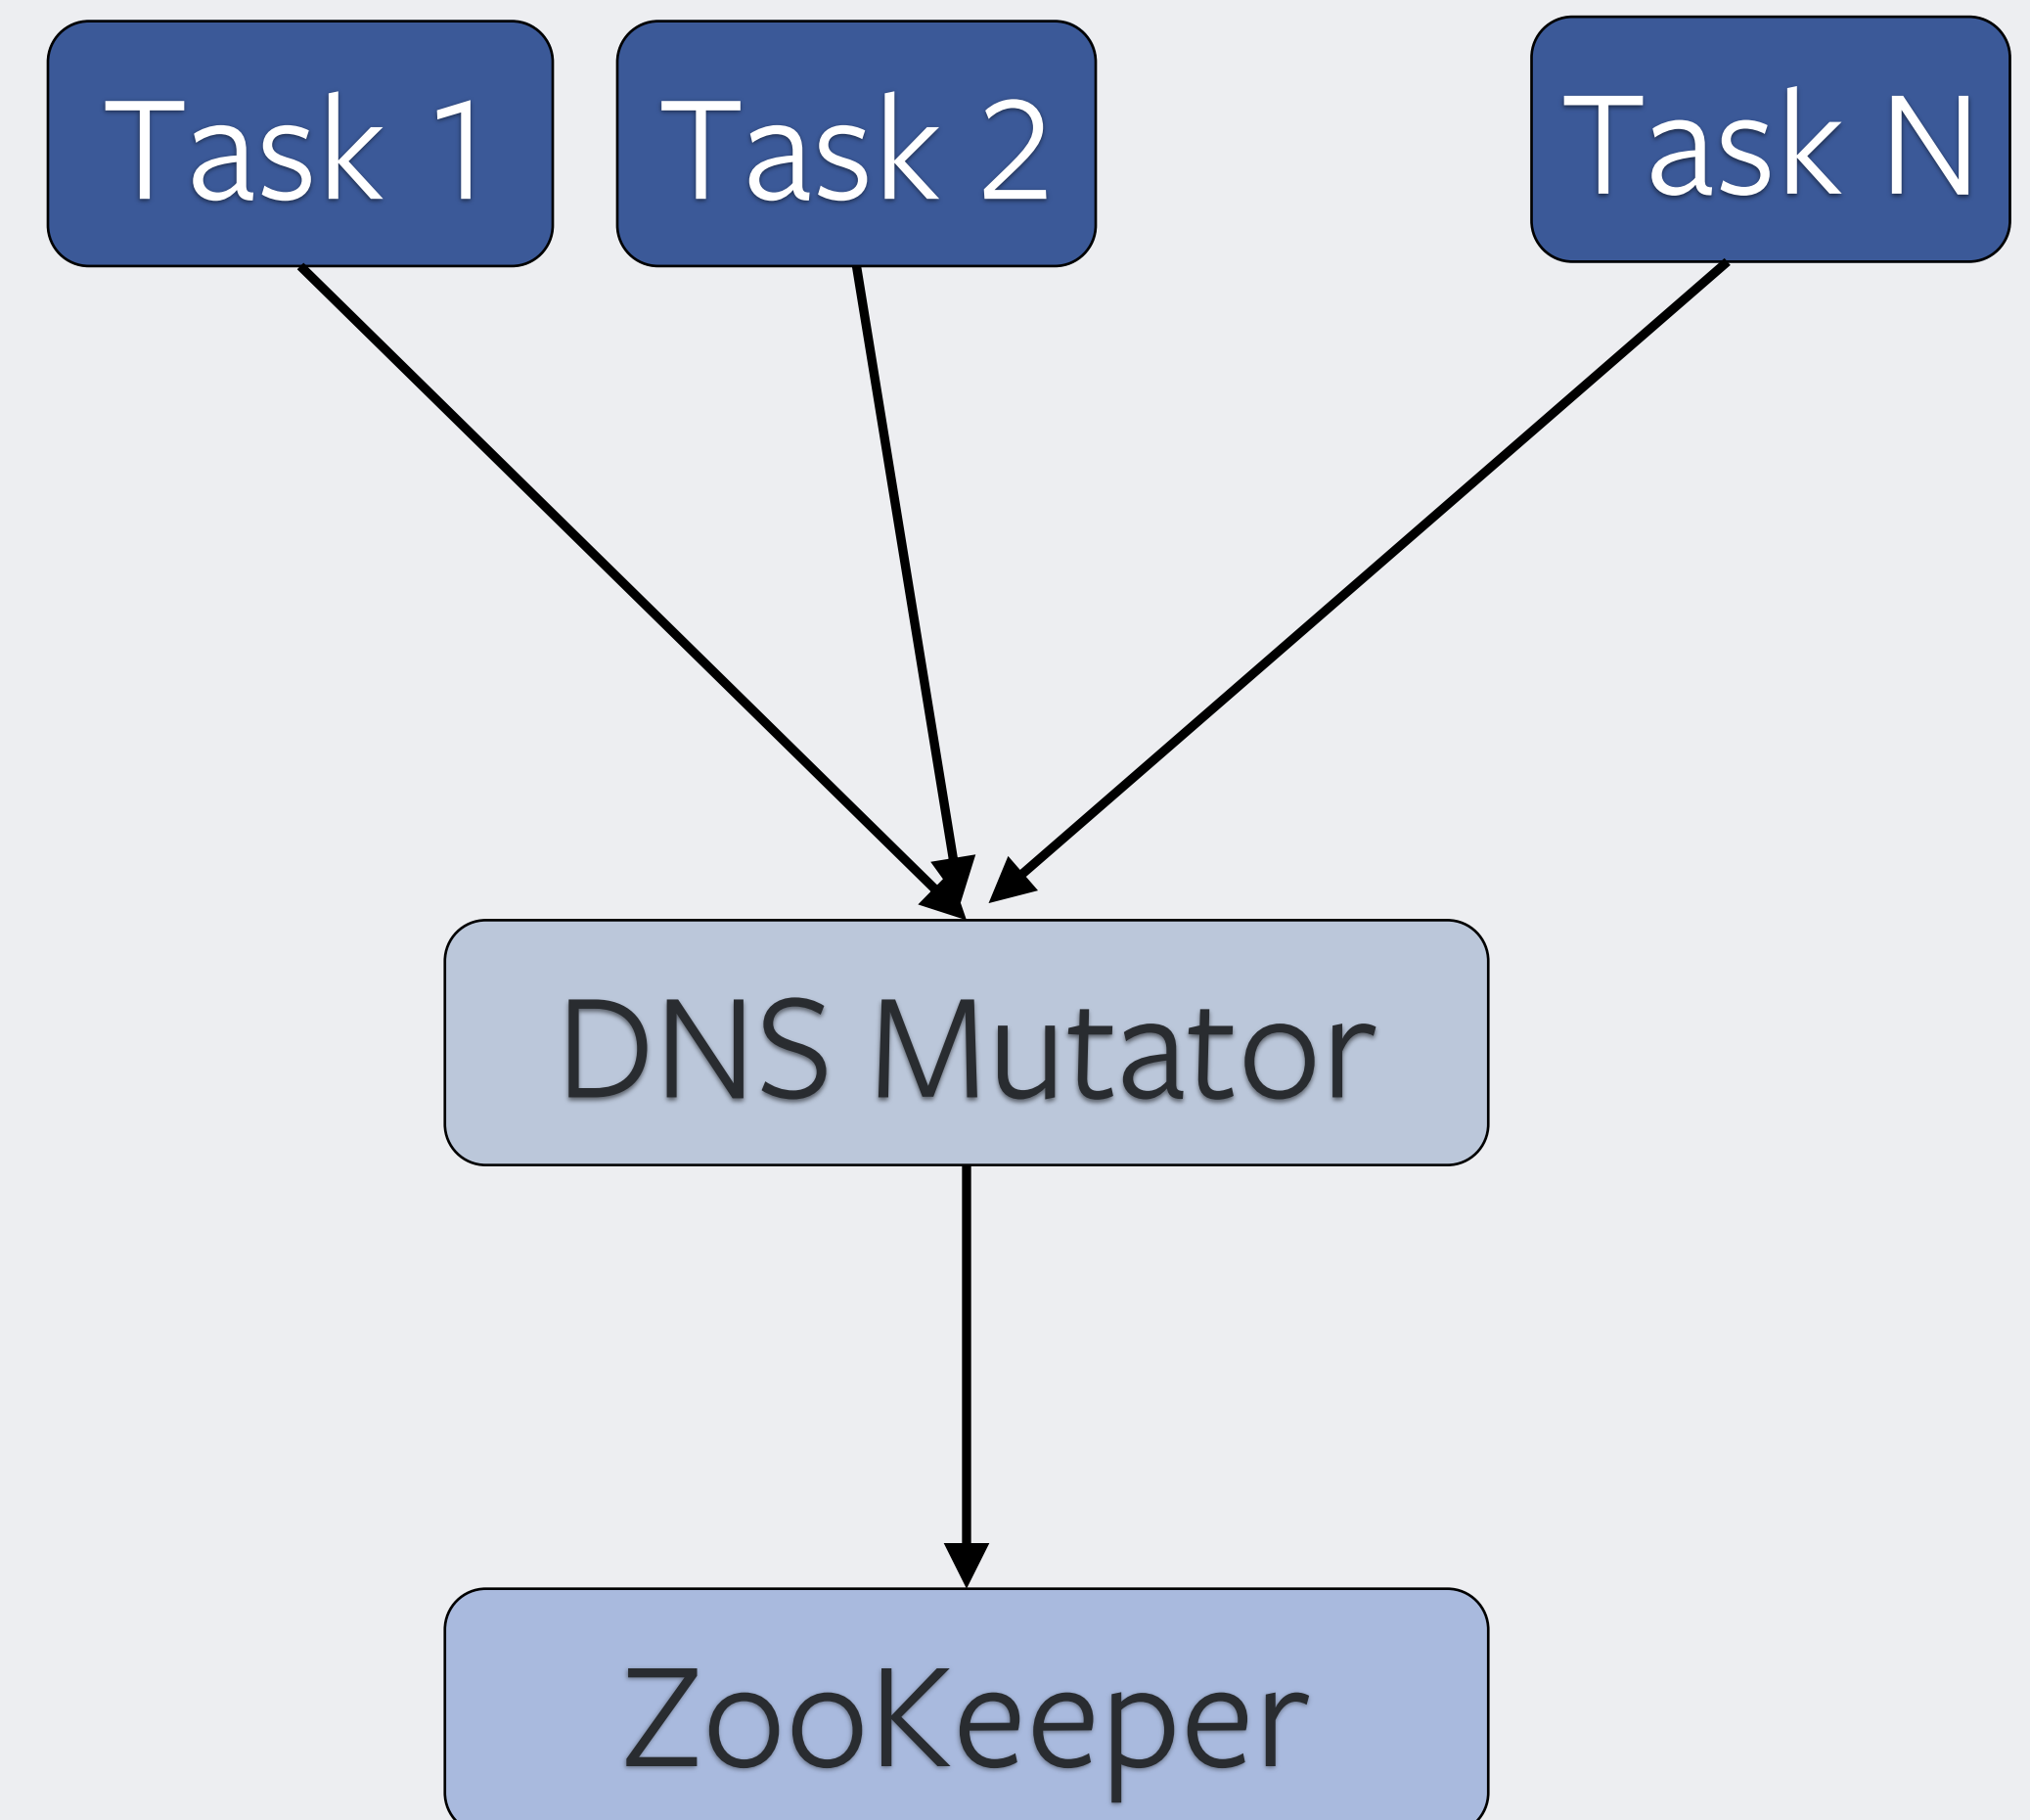
Human generated changes

- Version controlled/Code review
- Linter/CI to catch errors
 - Duplicates, cname cycles, unknown zone
 - Target IP, TinyDNS format
- Propagate to all hosts



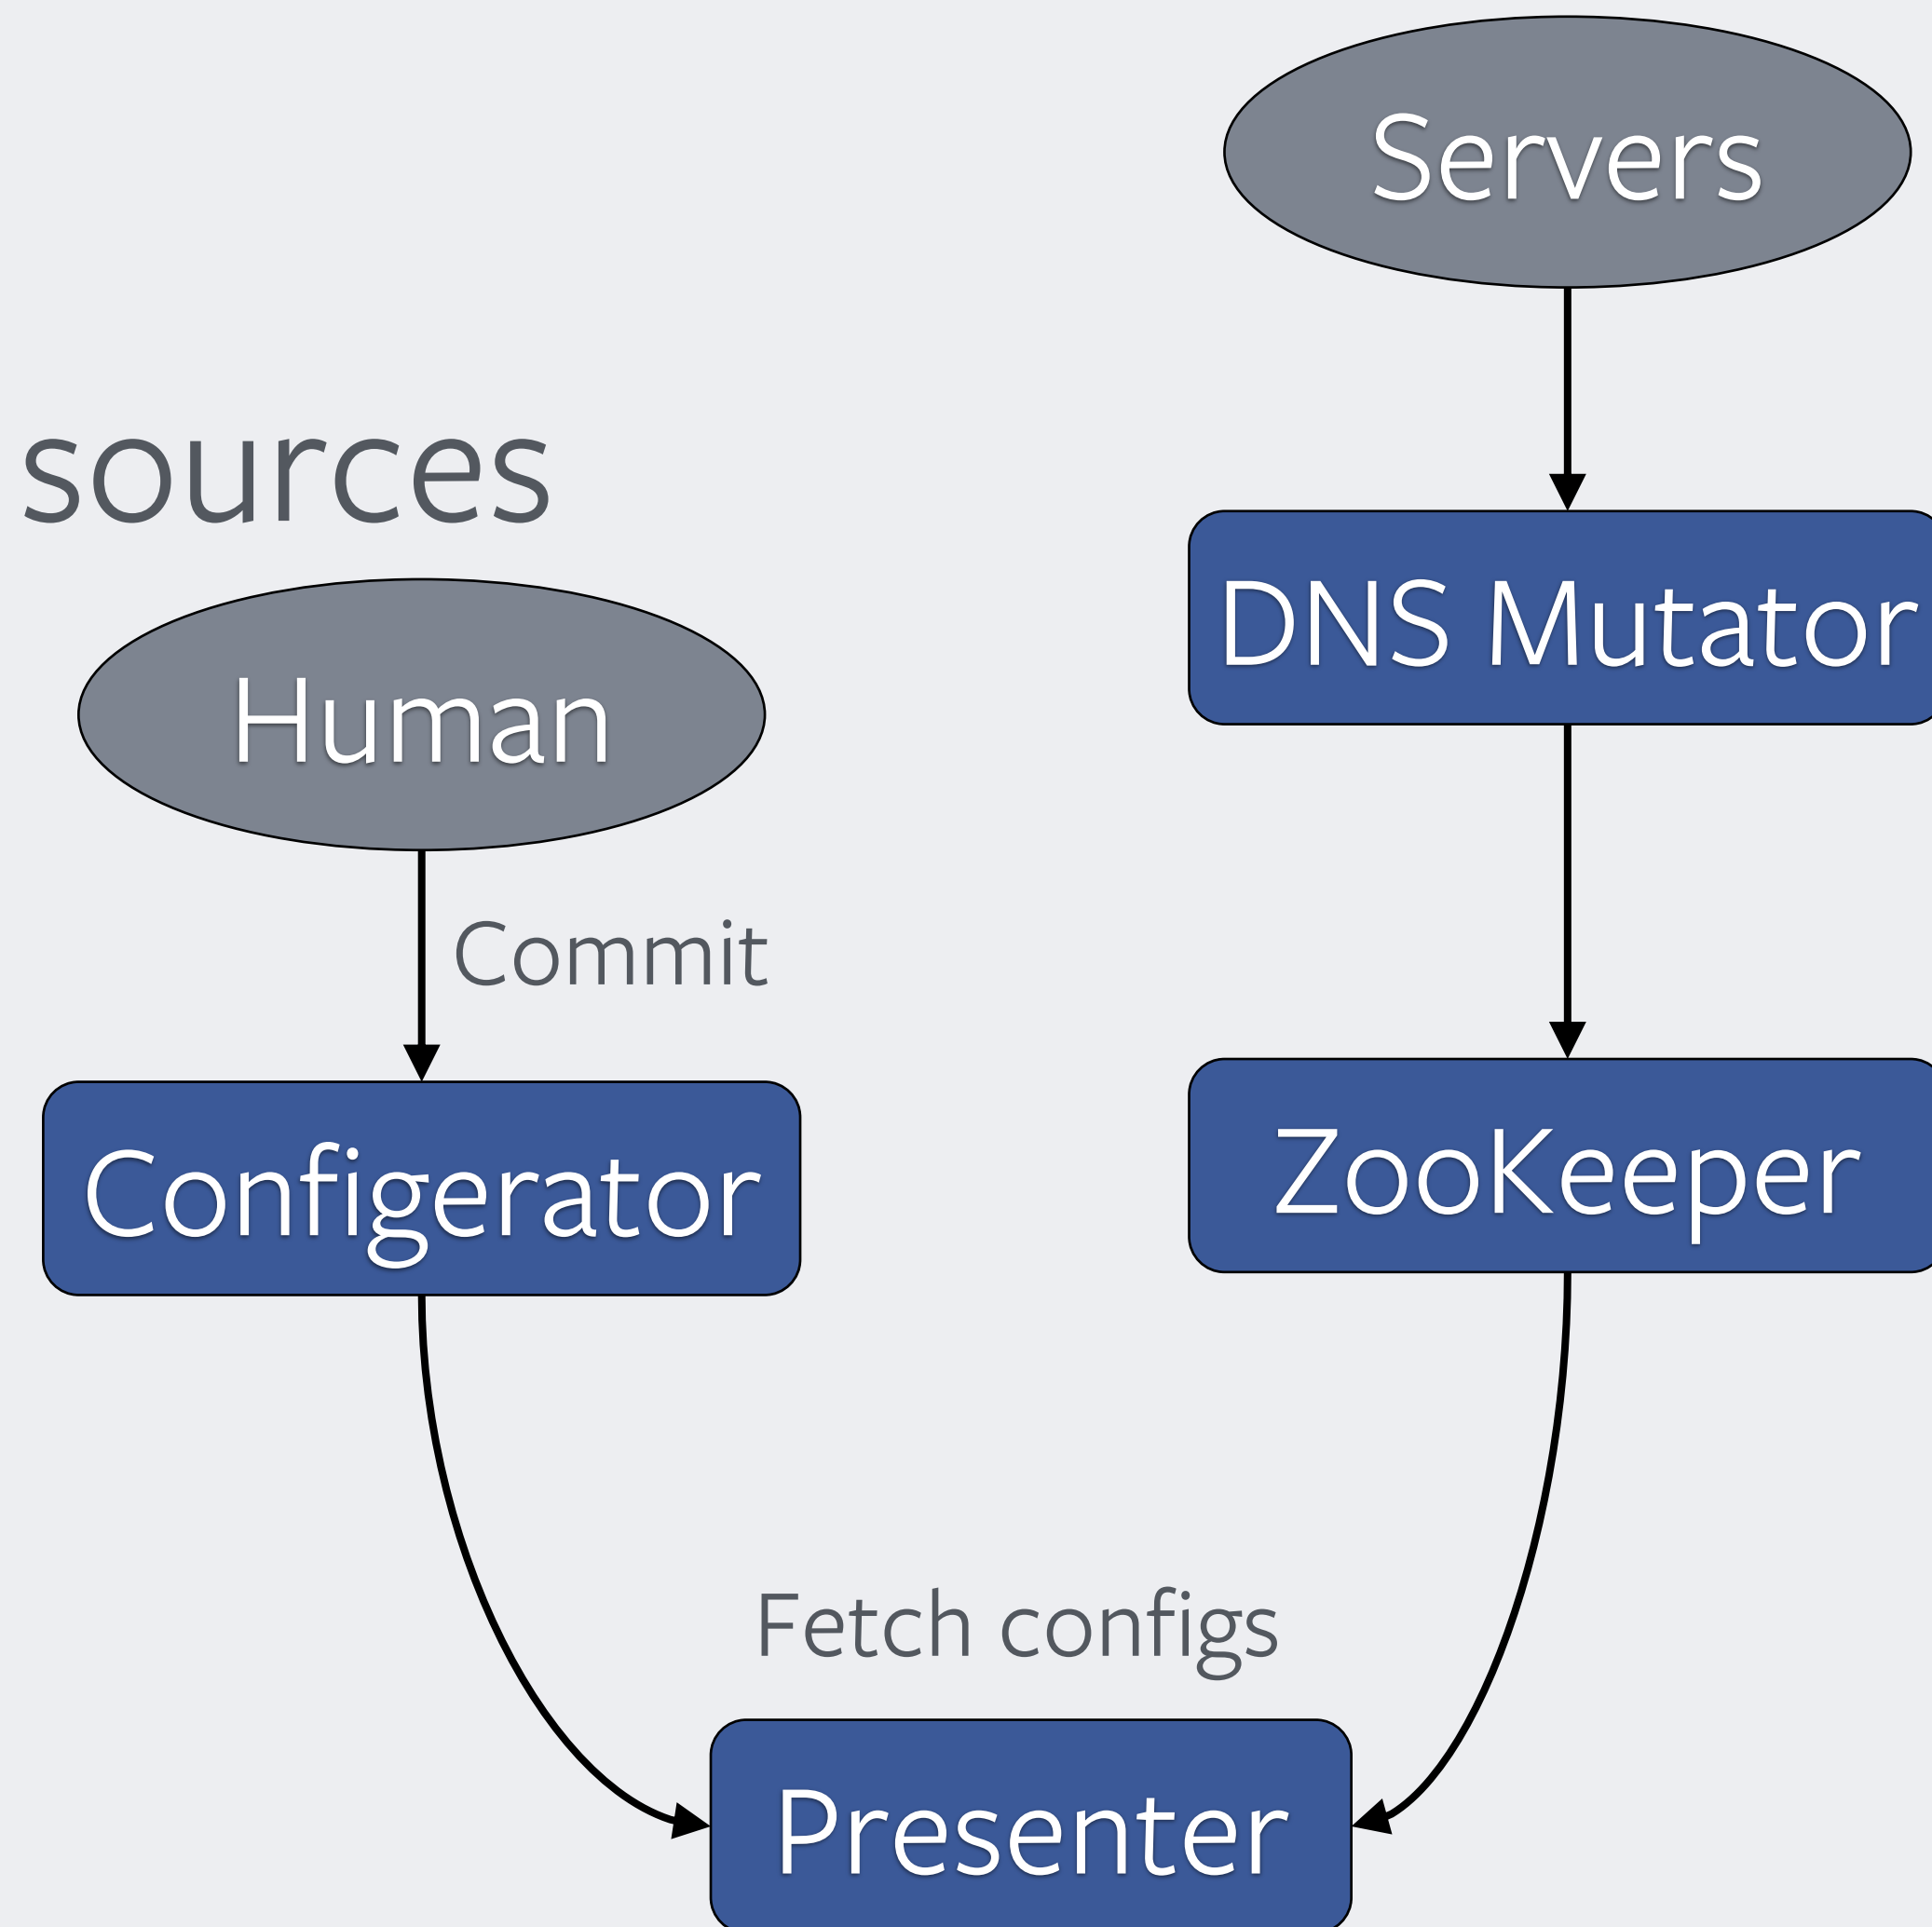
Automatically generated configs

- Python Framework
 - Only need to provide DNS data
 - abstract storage details, retrying...
- Periodically re-generate config
- Mutator checks generated data



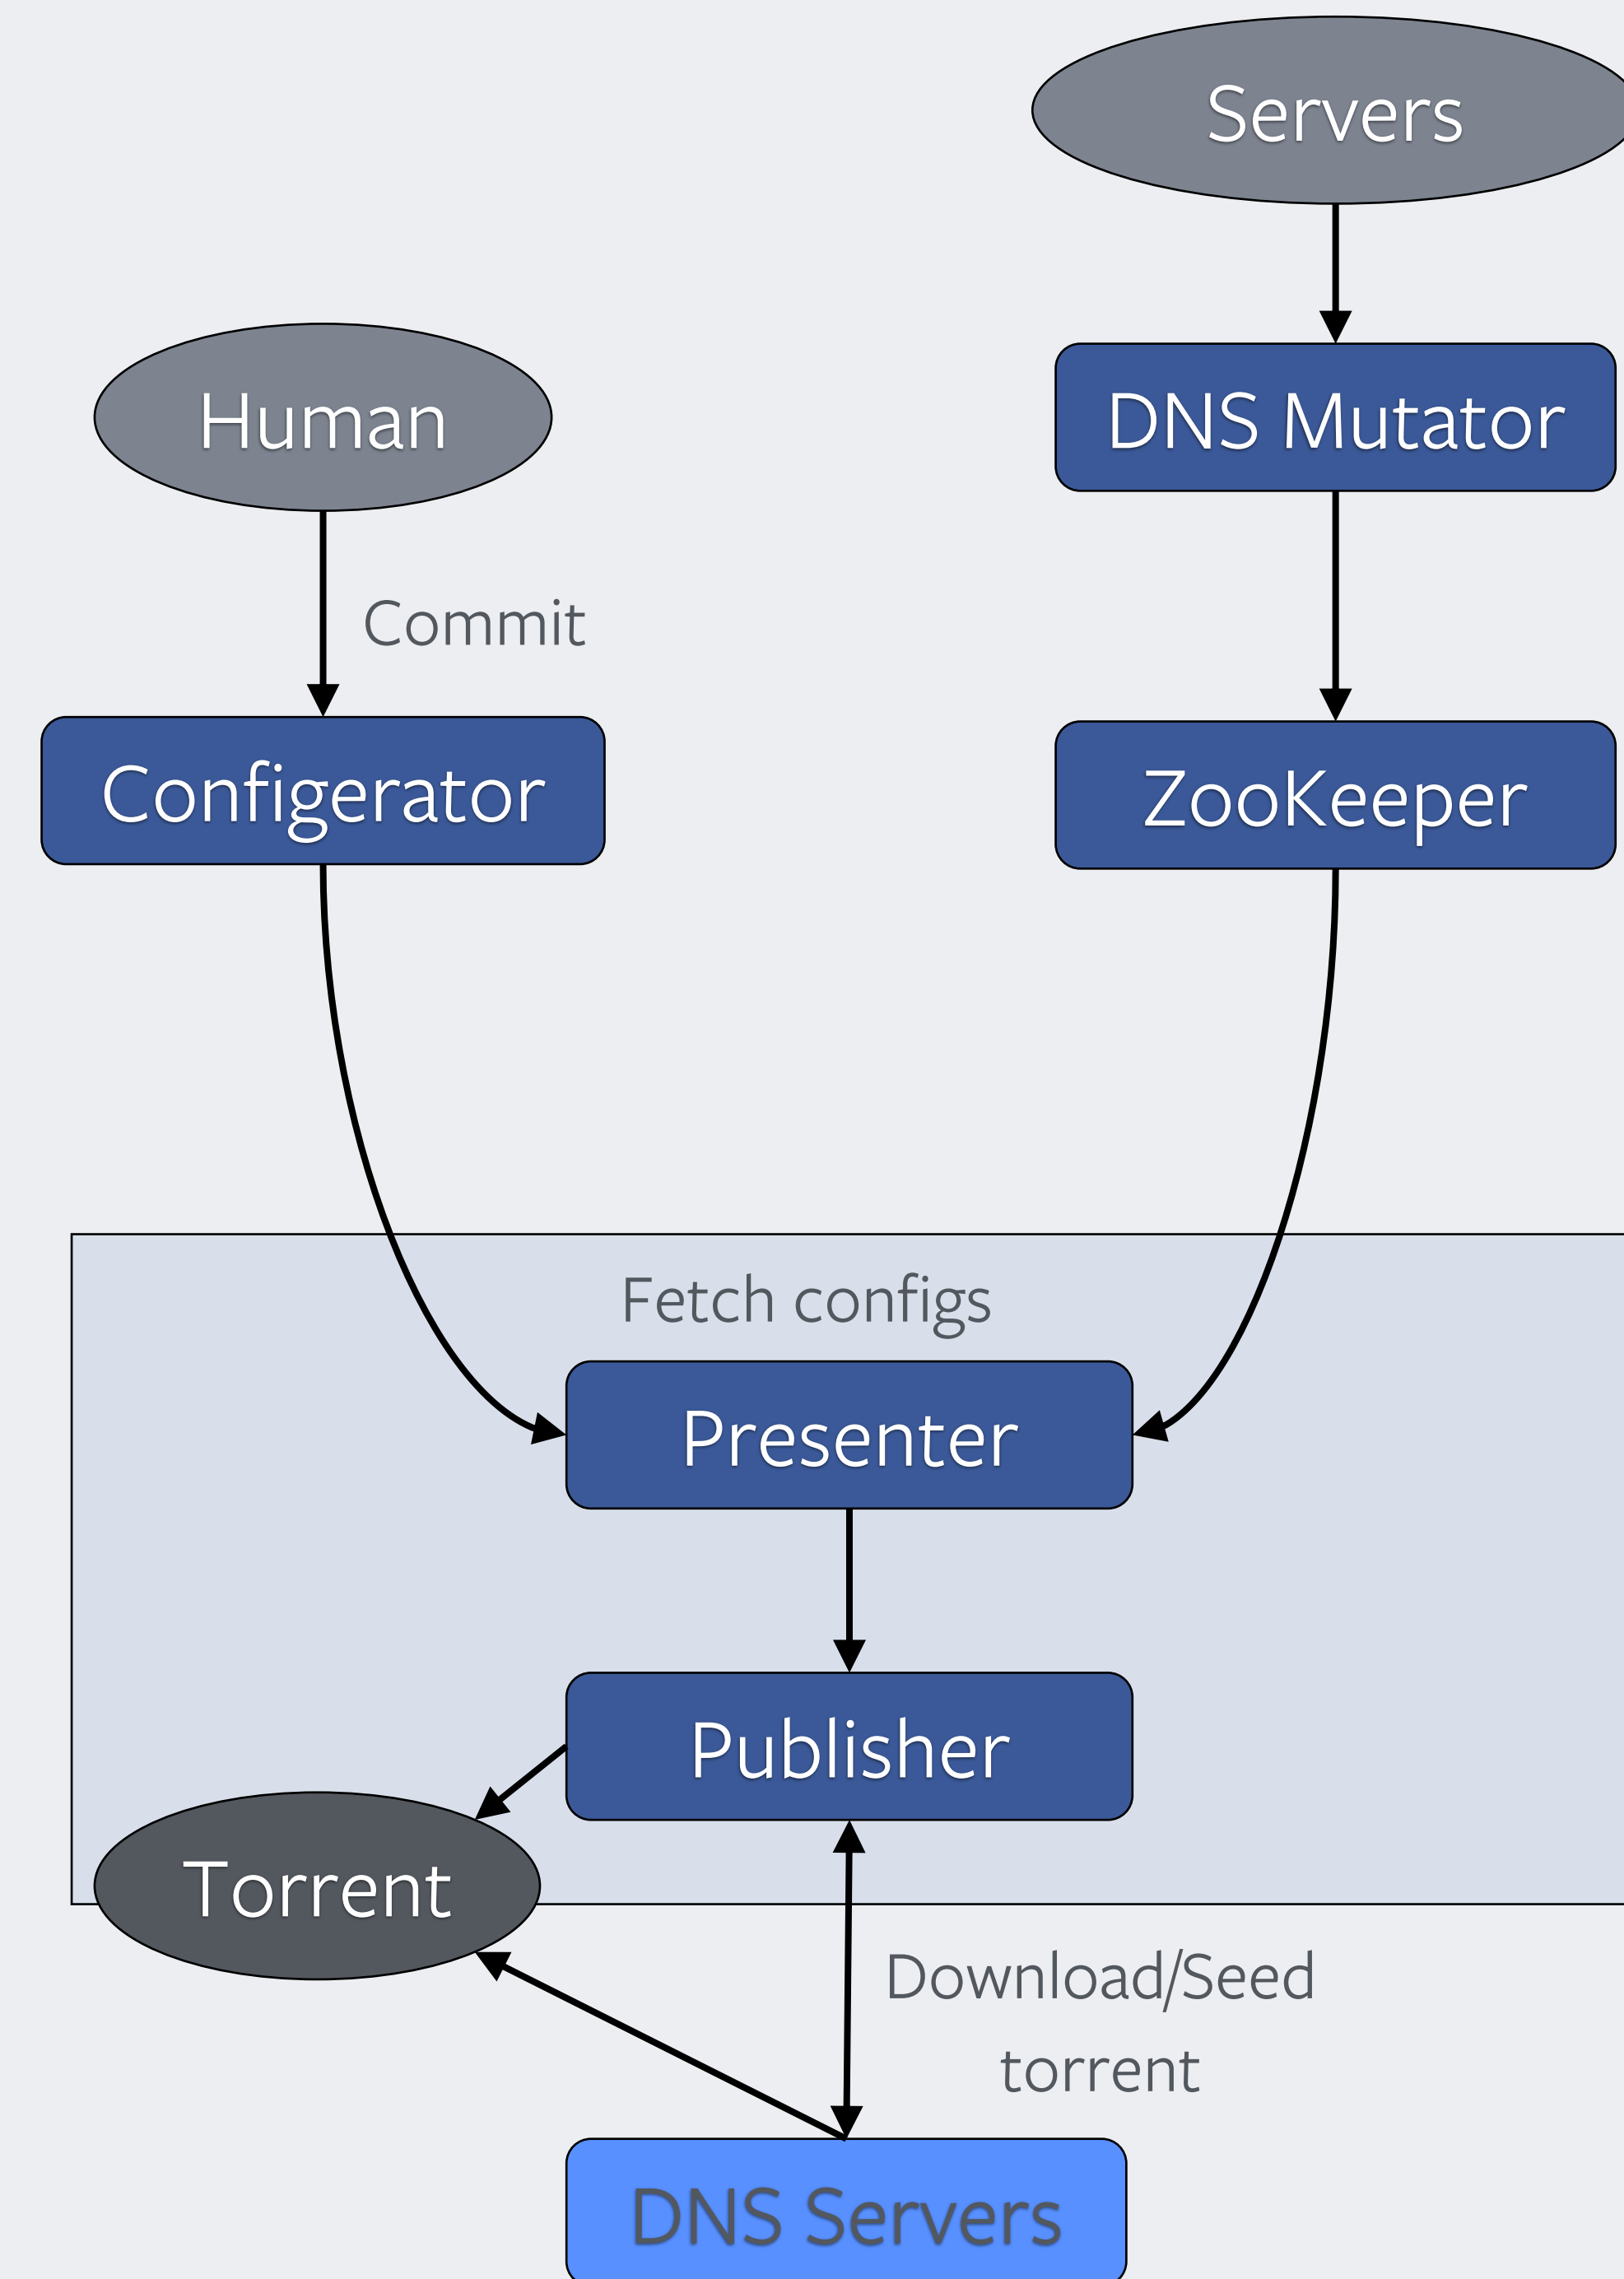
DNS Presenter

- Collect records from multiple data sources
- Record and Logical Validation
- Global view of DNS map
 - Perform final sanity checks
 - Collision/Resolution
 - Post processing



DNS Pipeline Overview

- Complexity broken into:
 - Small and simple services
 - Isolation
- Checks at every levels
 - Catch issues early
 - Easy to trace
 - Limit failure blast



Thank You

facebook