



How Spotify Migrated HTTP Ingress Systems to Envoy

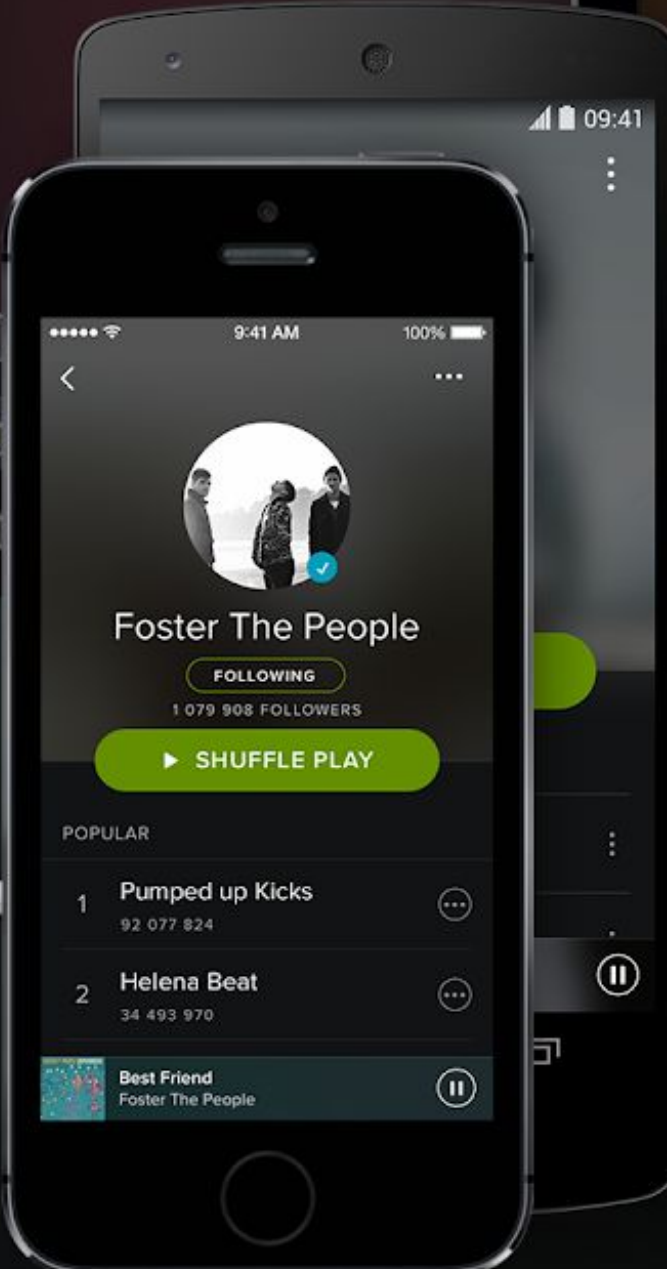
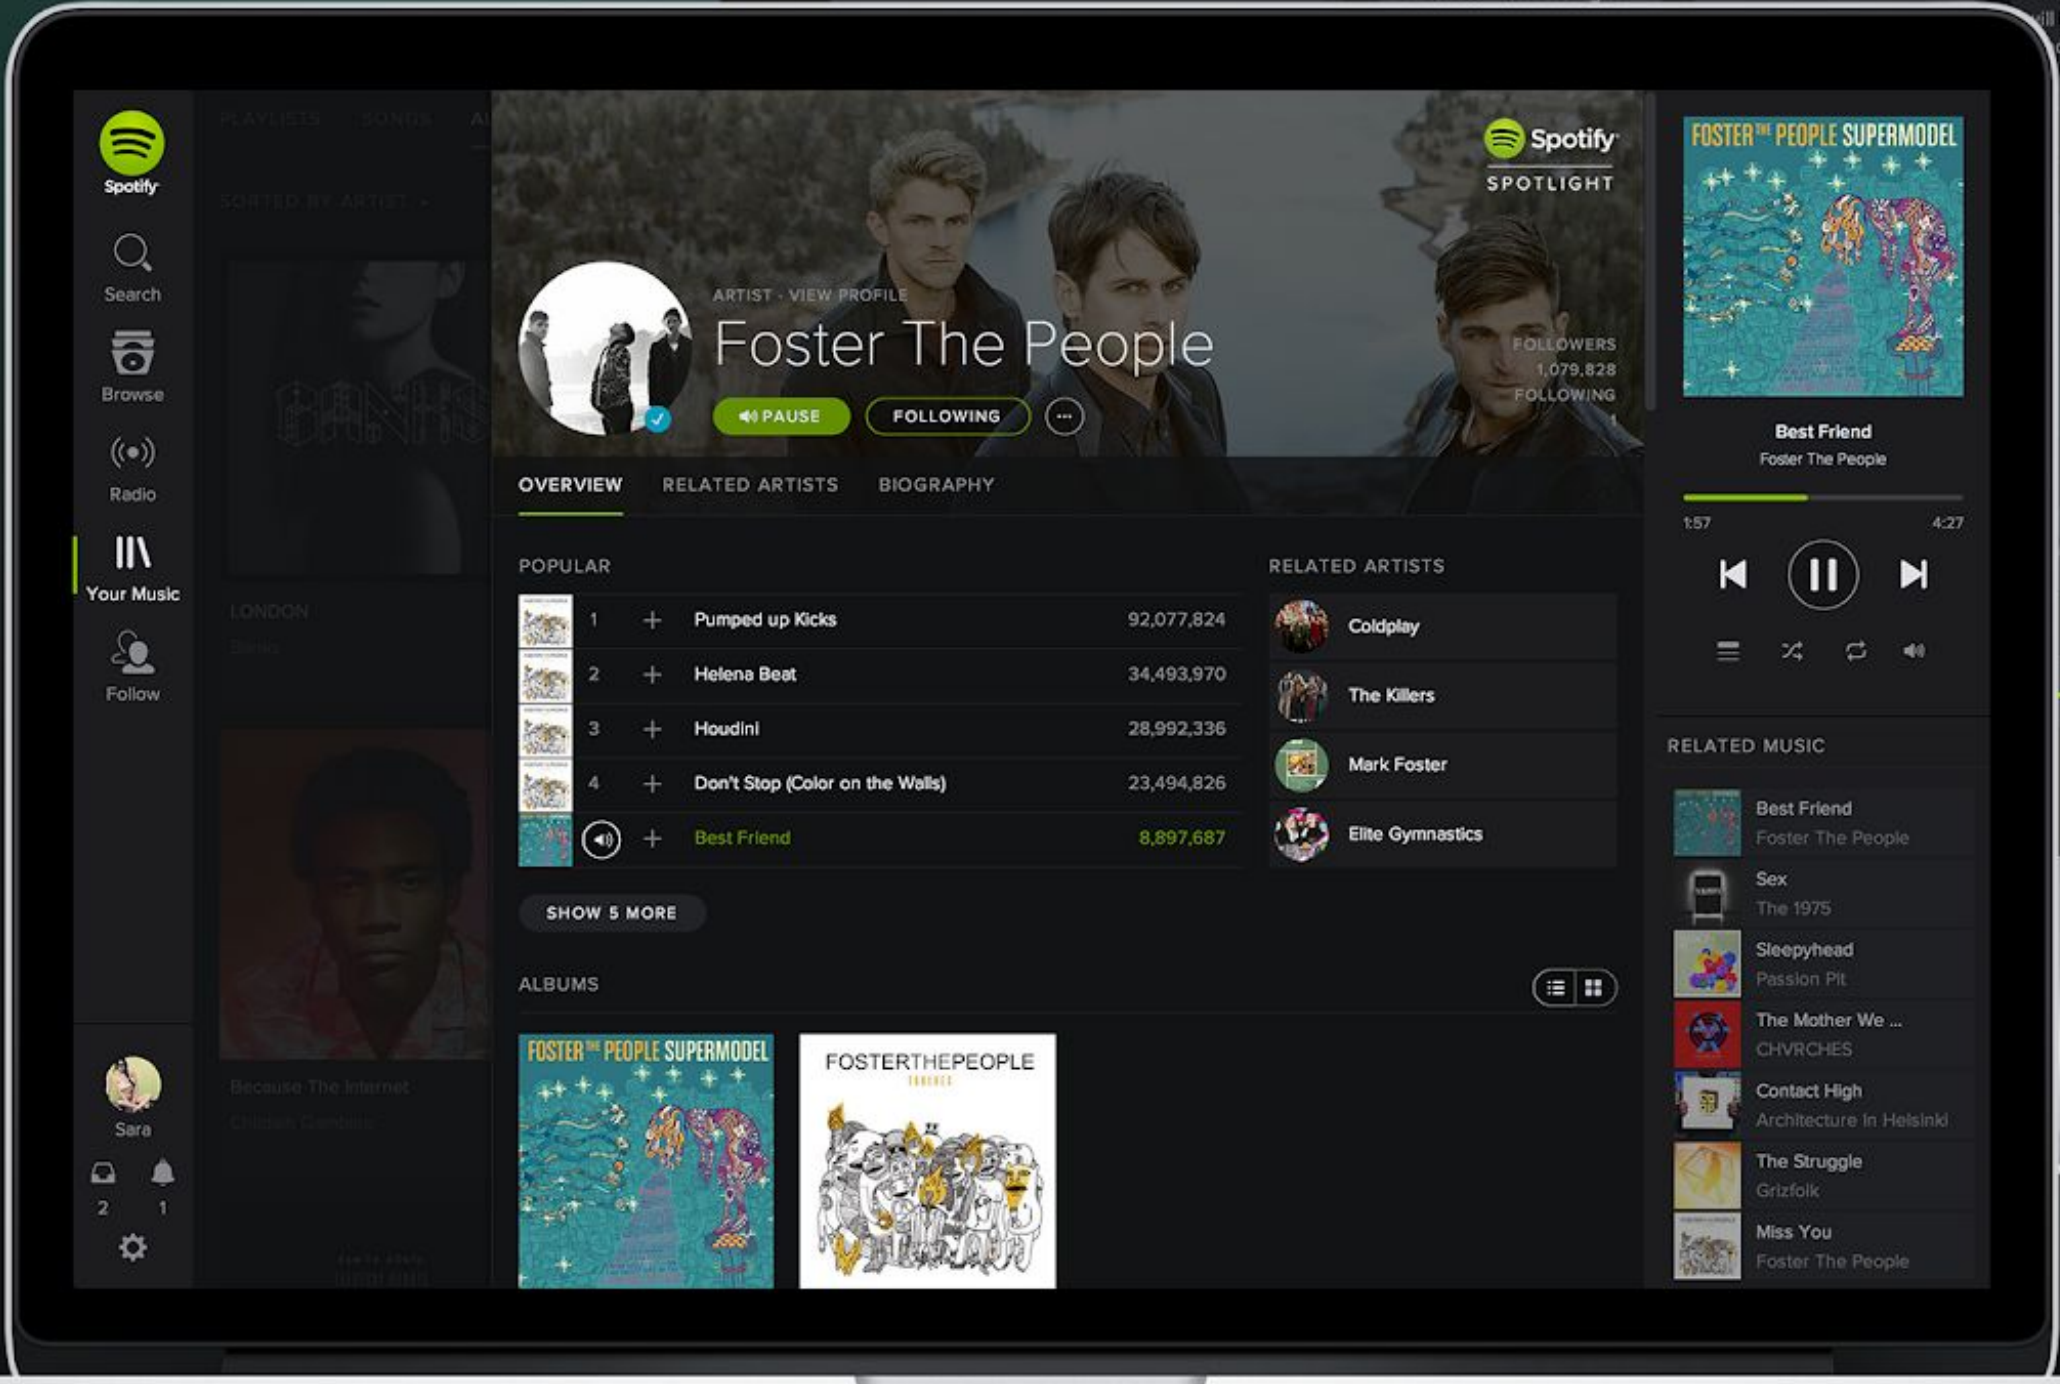
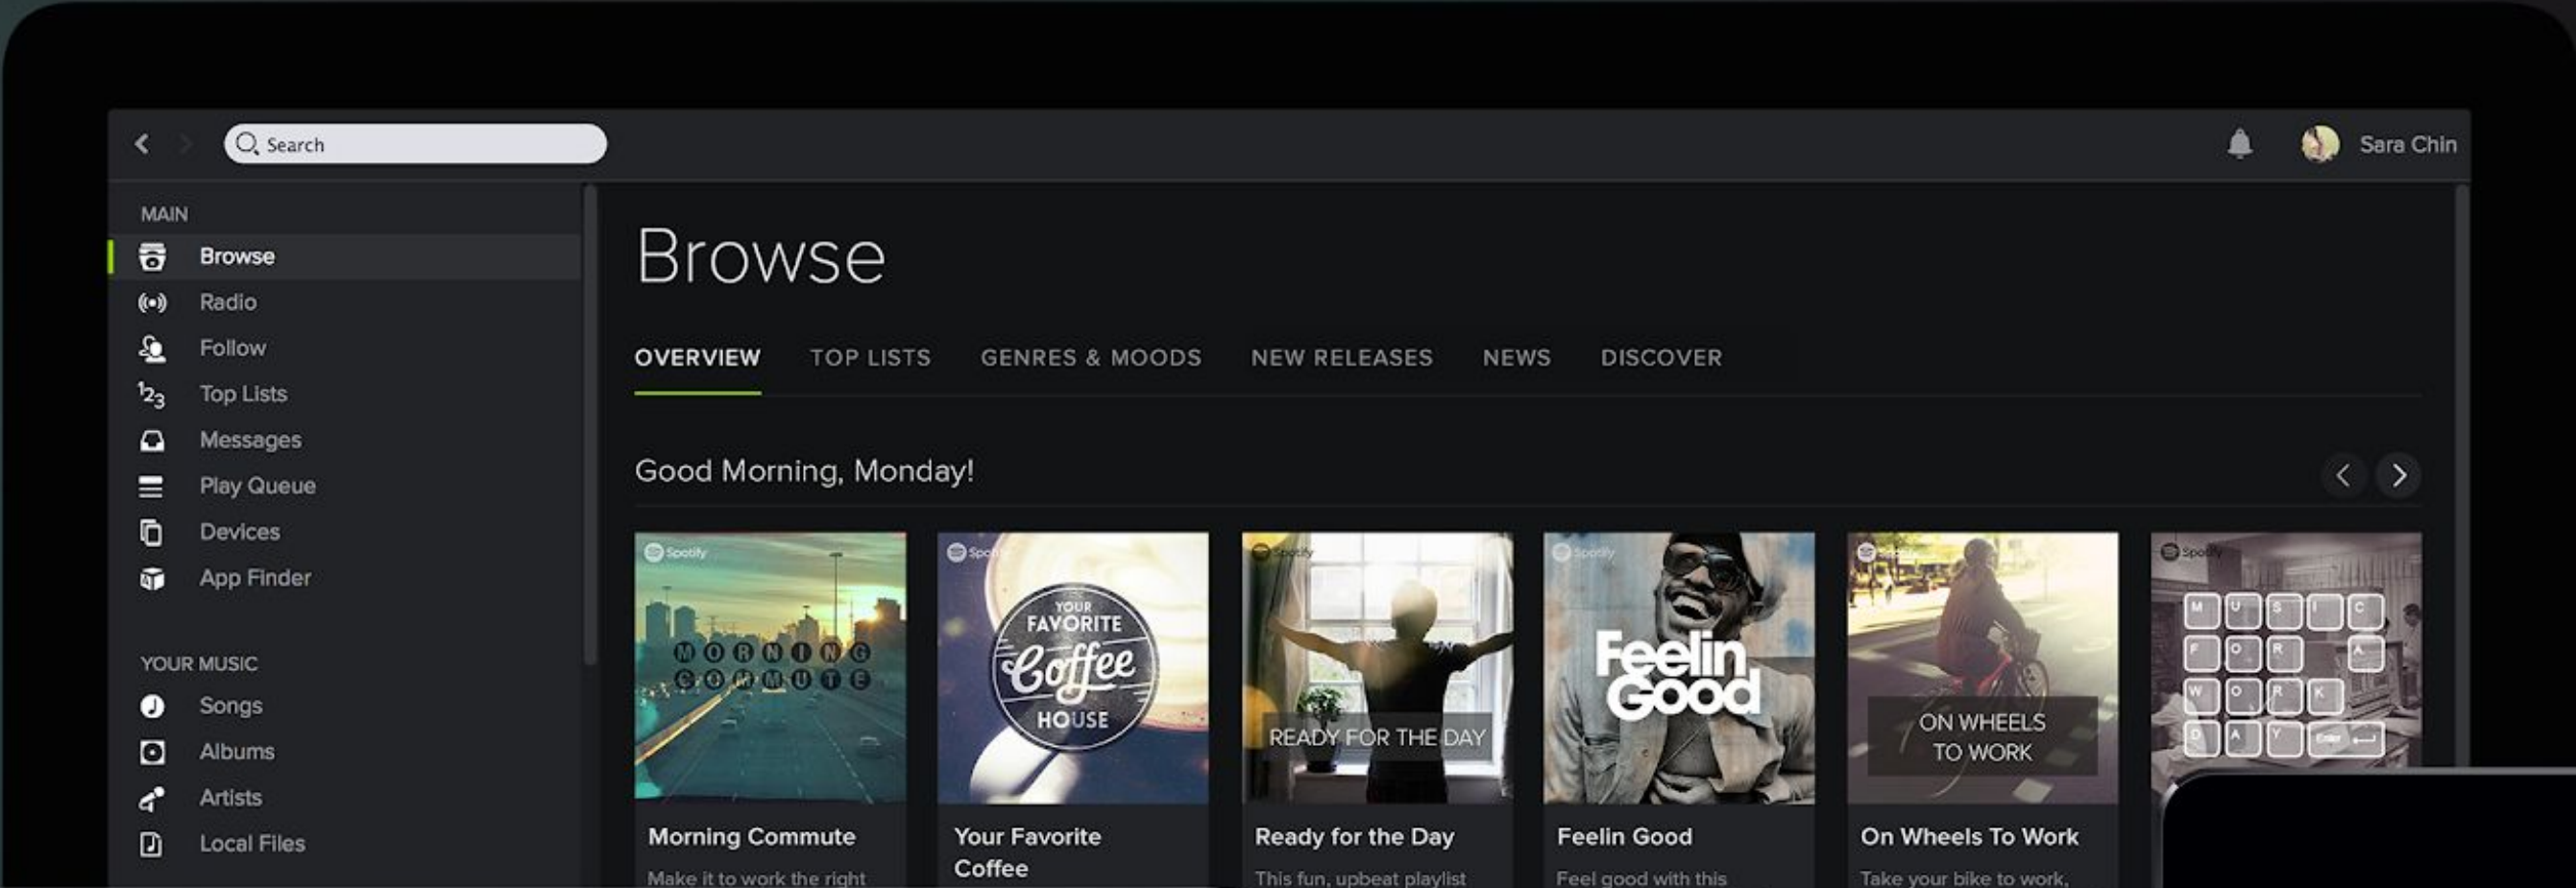
Erica Manno - ericam@spotify.com
Vladimir Shakhov - vladimir@spotify.com

About us

- ❑ software engineers in **ATC** squad (**A**ir **T**raffic **C**ontr**o**l)
- ❑ owners of Spotify's perimeter
- ❑ "ops in teams"

Agenda

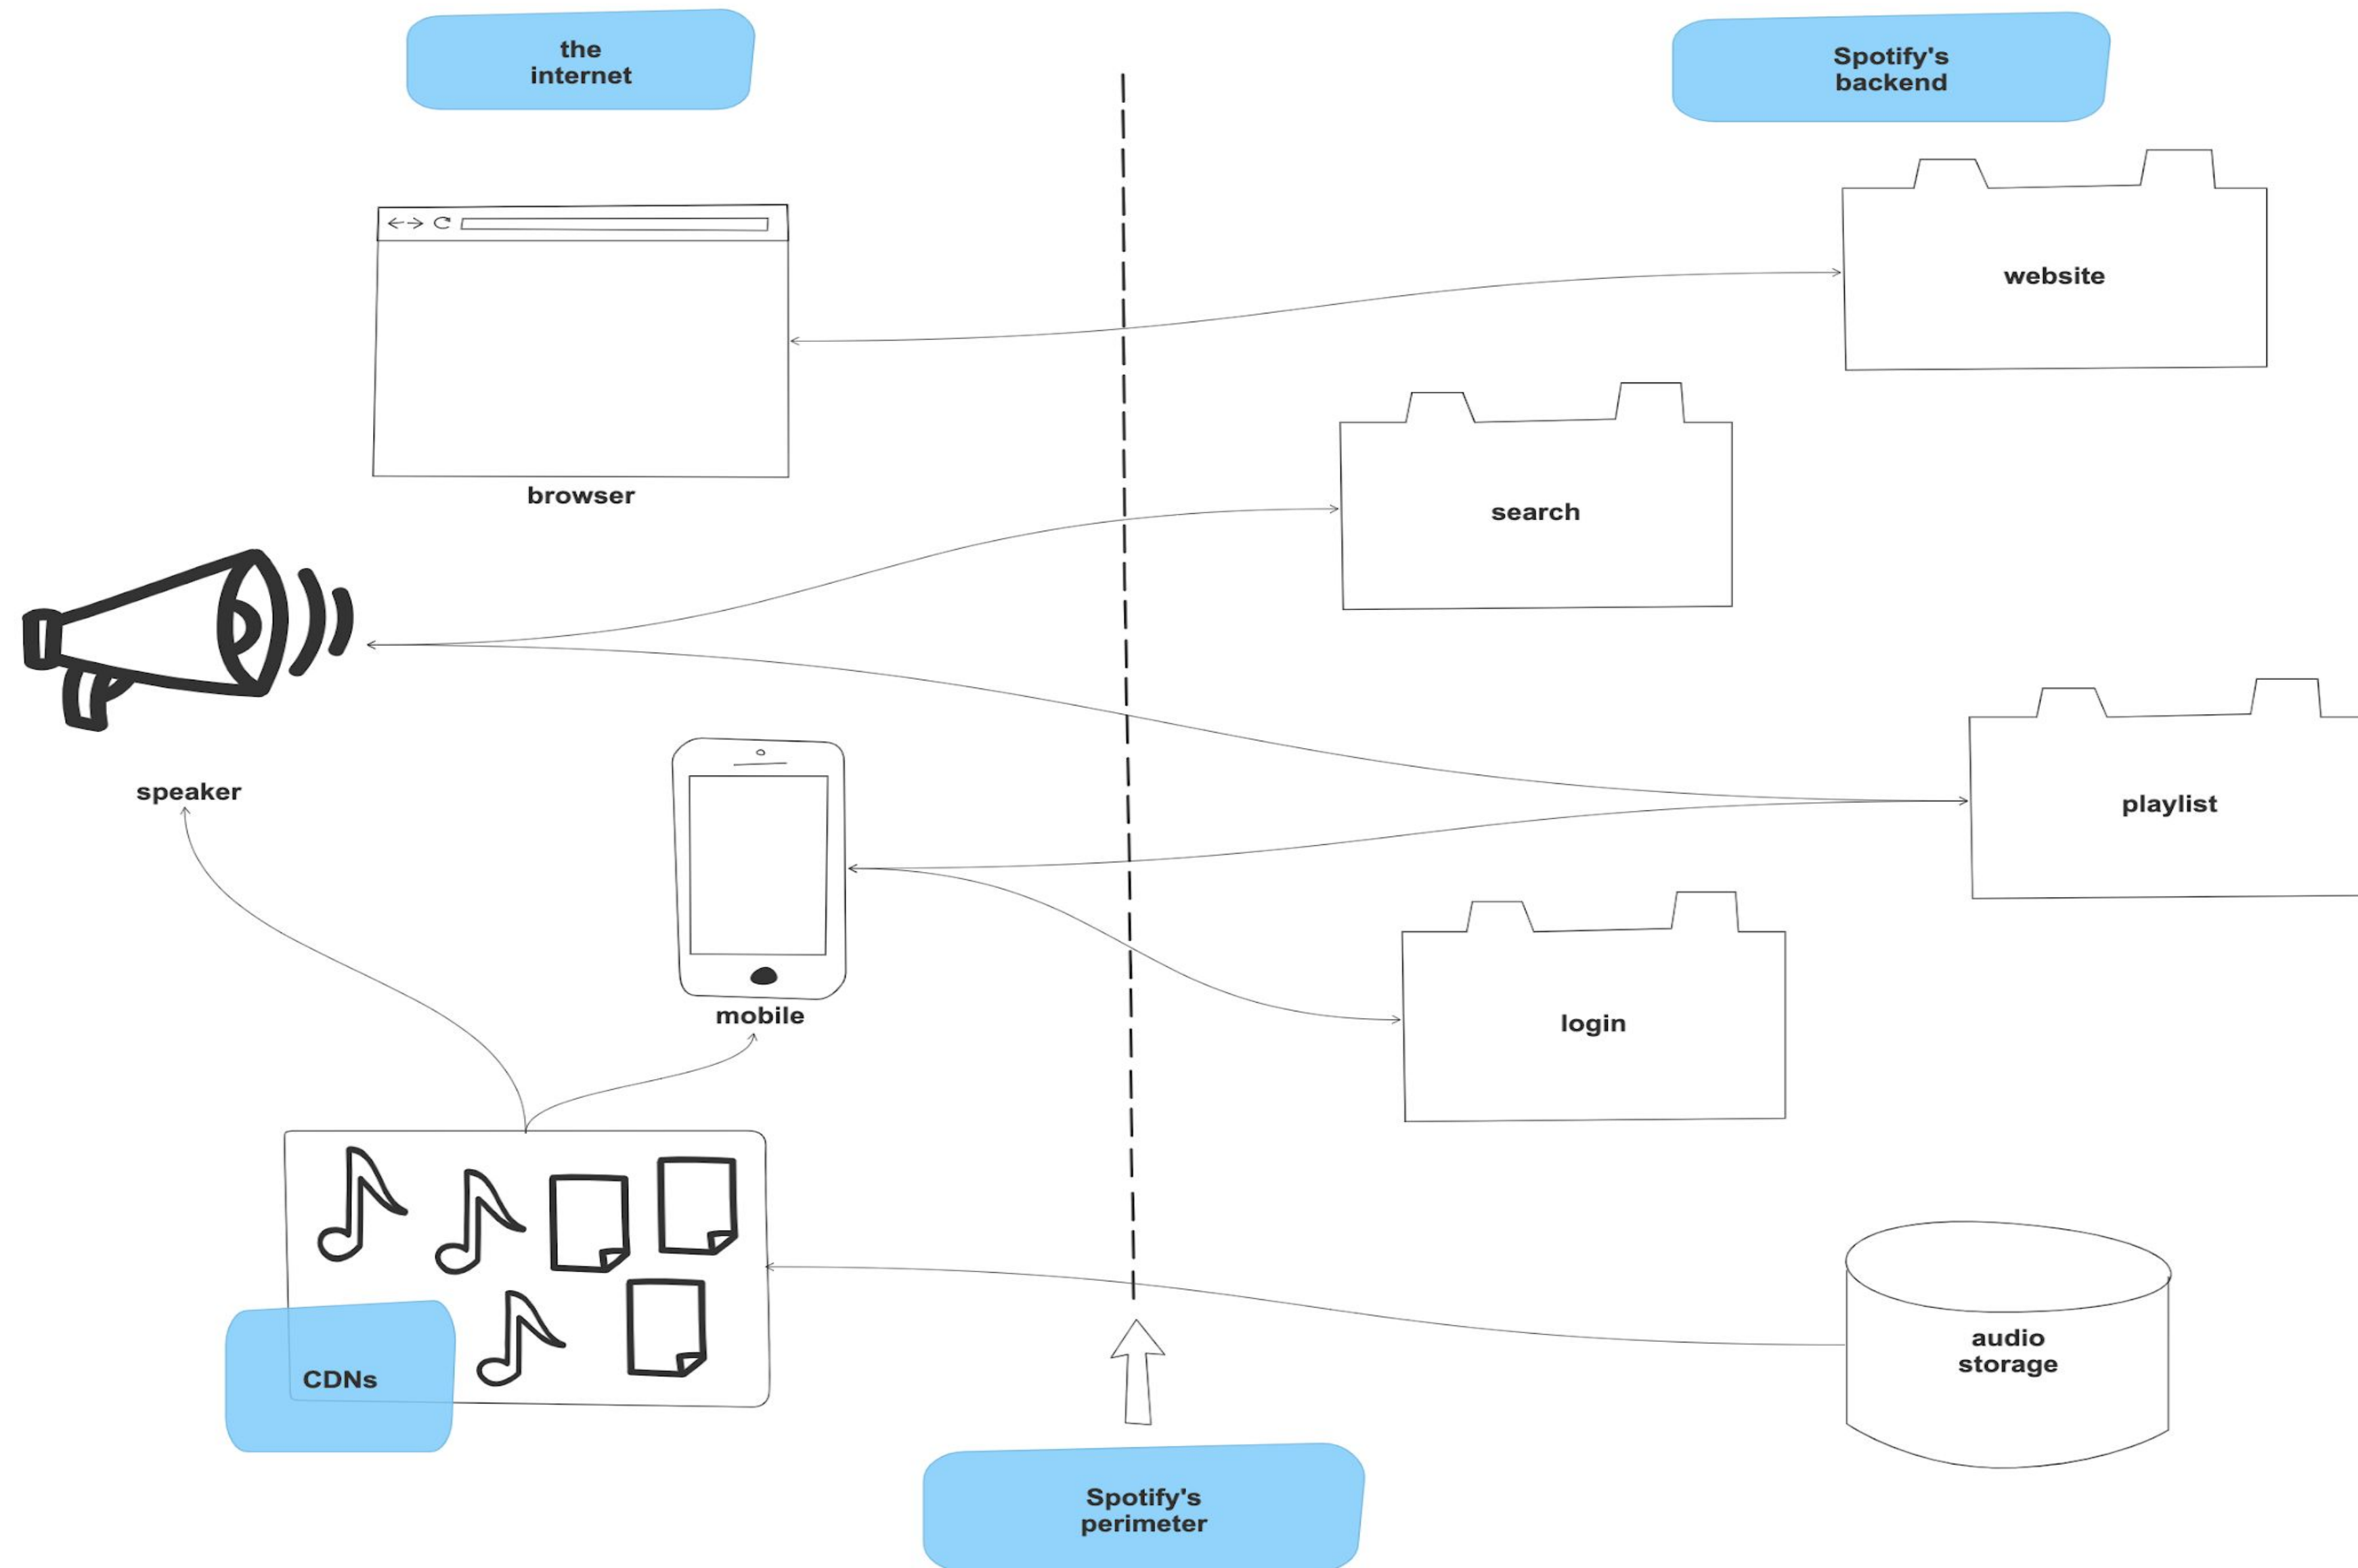
- ❑ how we migrated a part of Spotify's perimeter to Envoy
- ❑ future of Spotify's perimeter
- ❑ key takeaways from migration



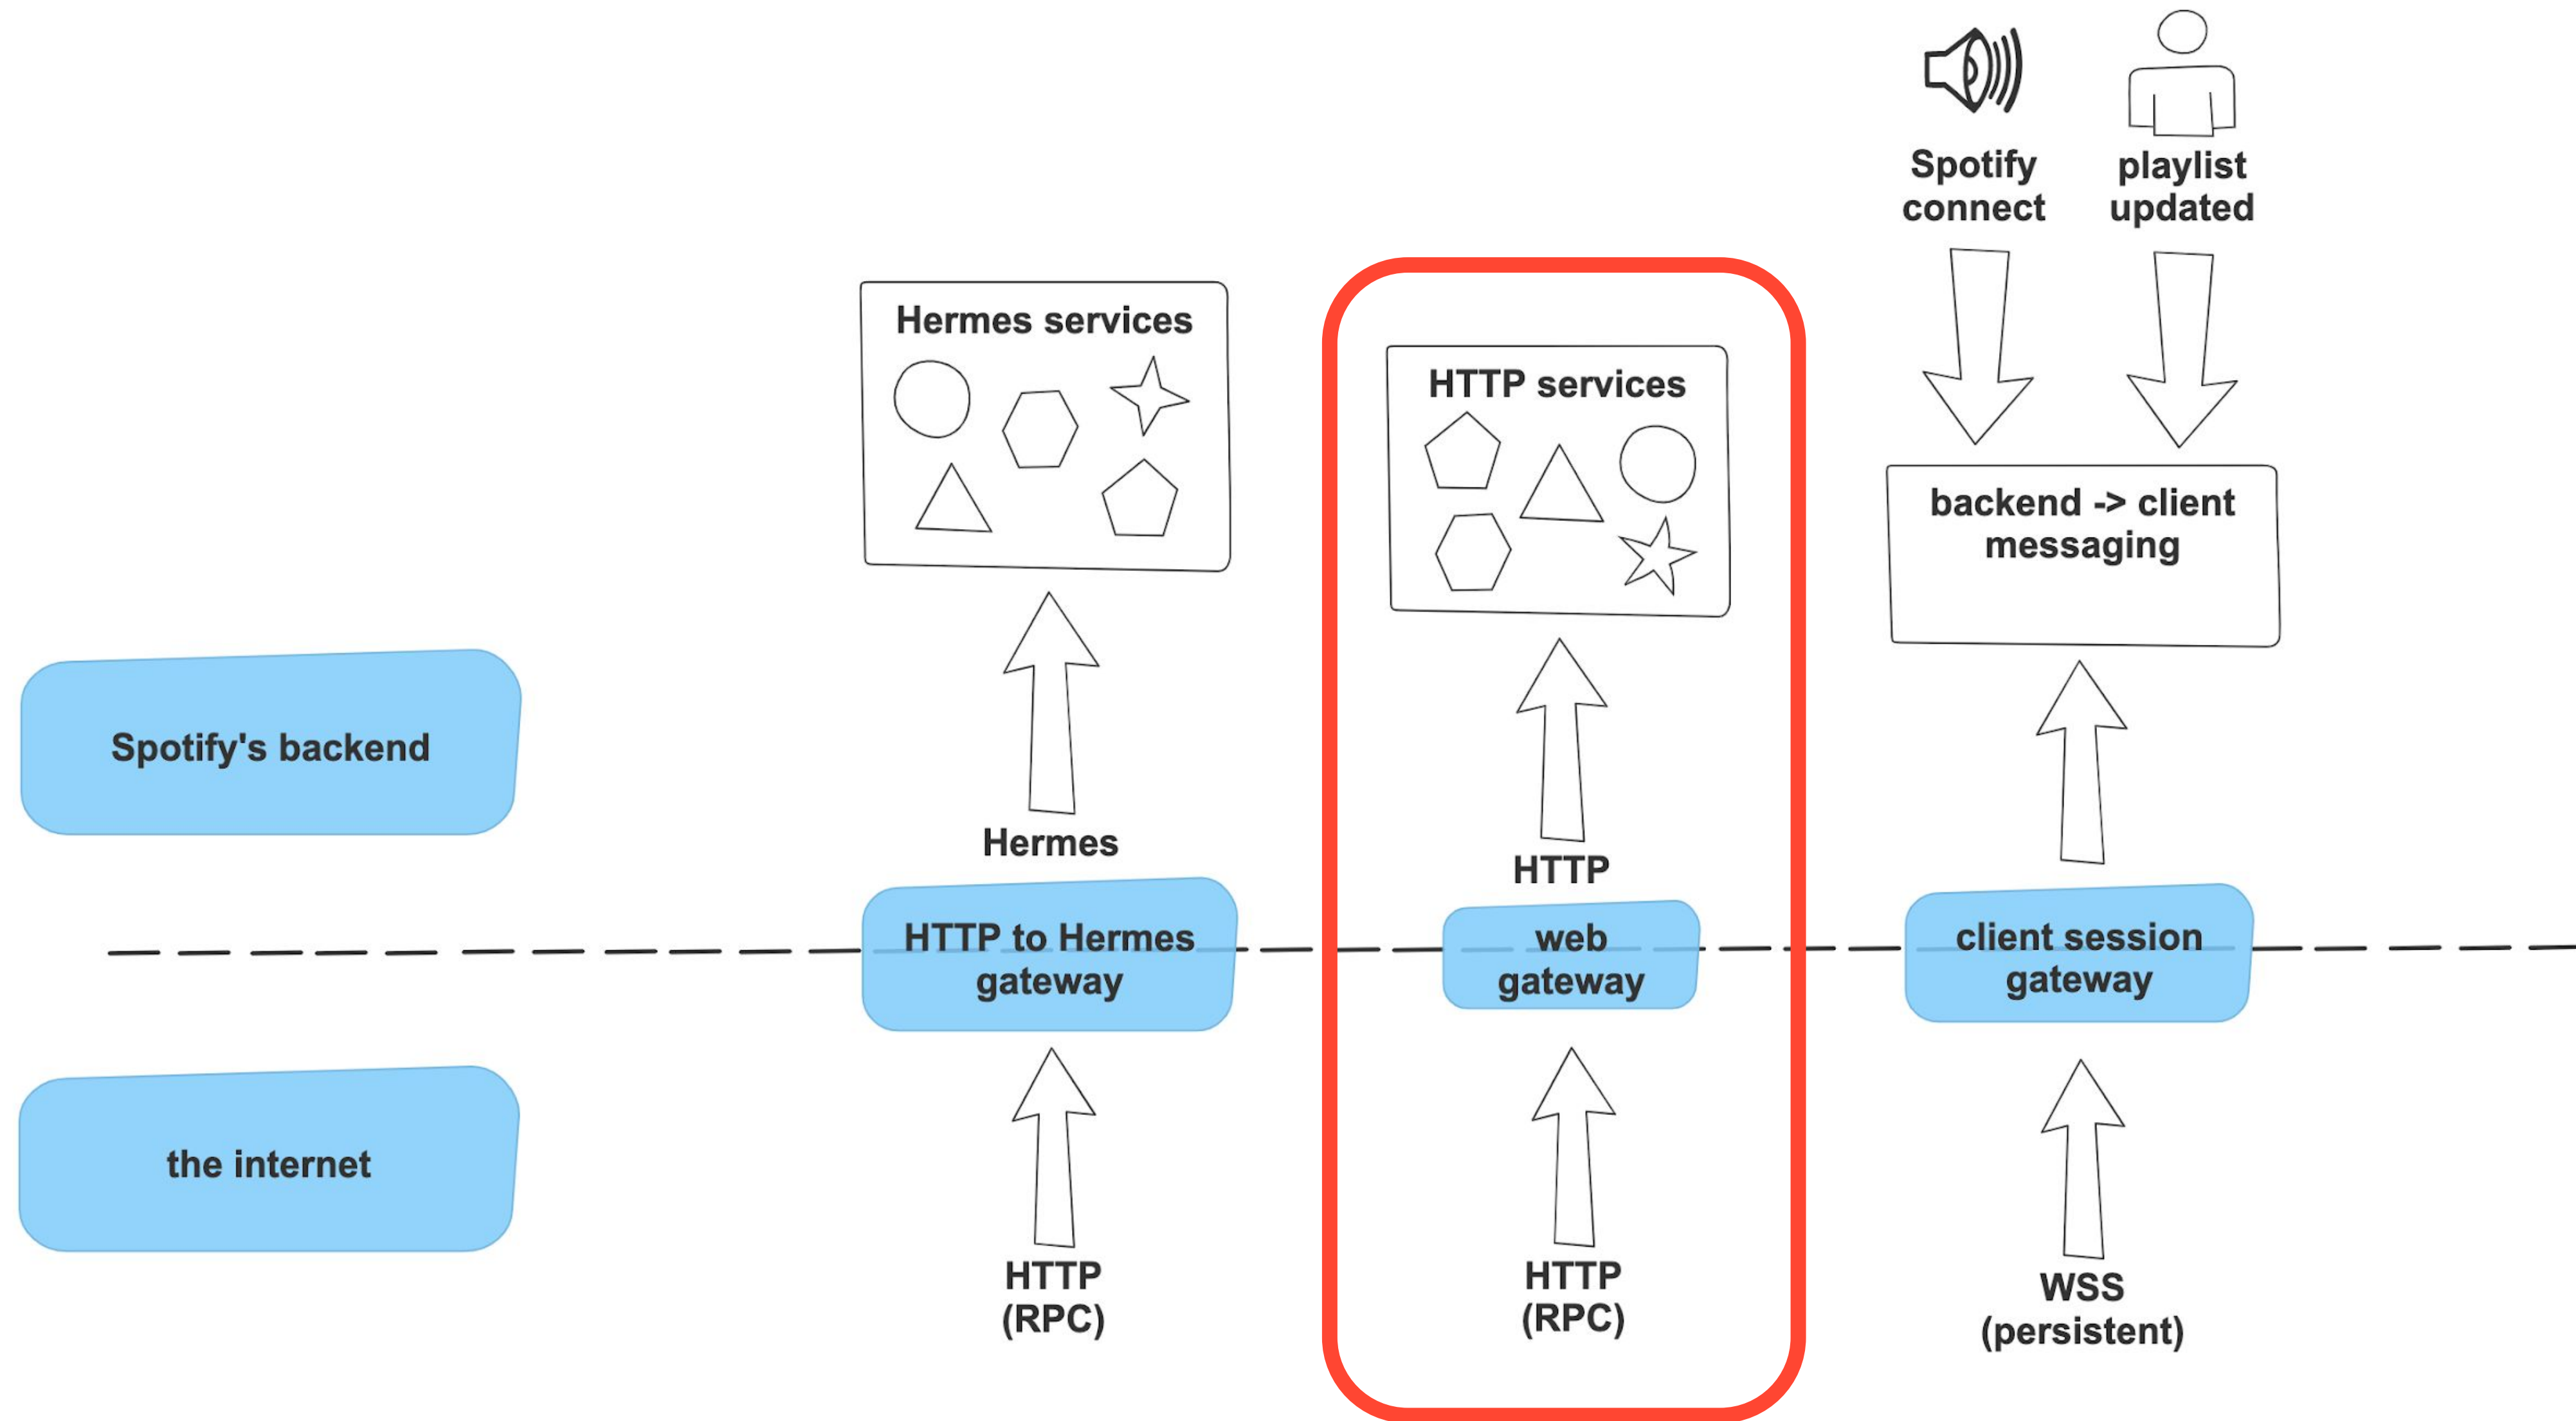
About Spotify

- ❑ 248M MAUs
- ❑ 79 markets
- ❑ 8M+ RPS
- ❑ 1500+ engineers, 280+ squads
- ❑ 1200+ microservices in production

About Spotify (cont.)

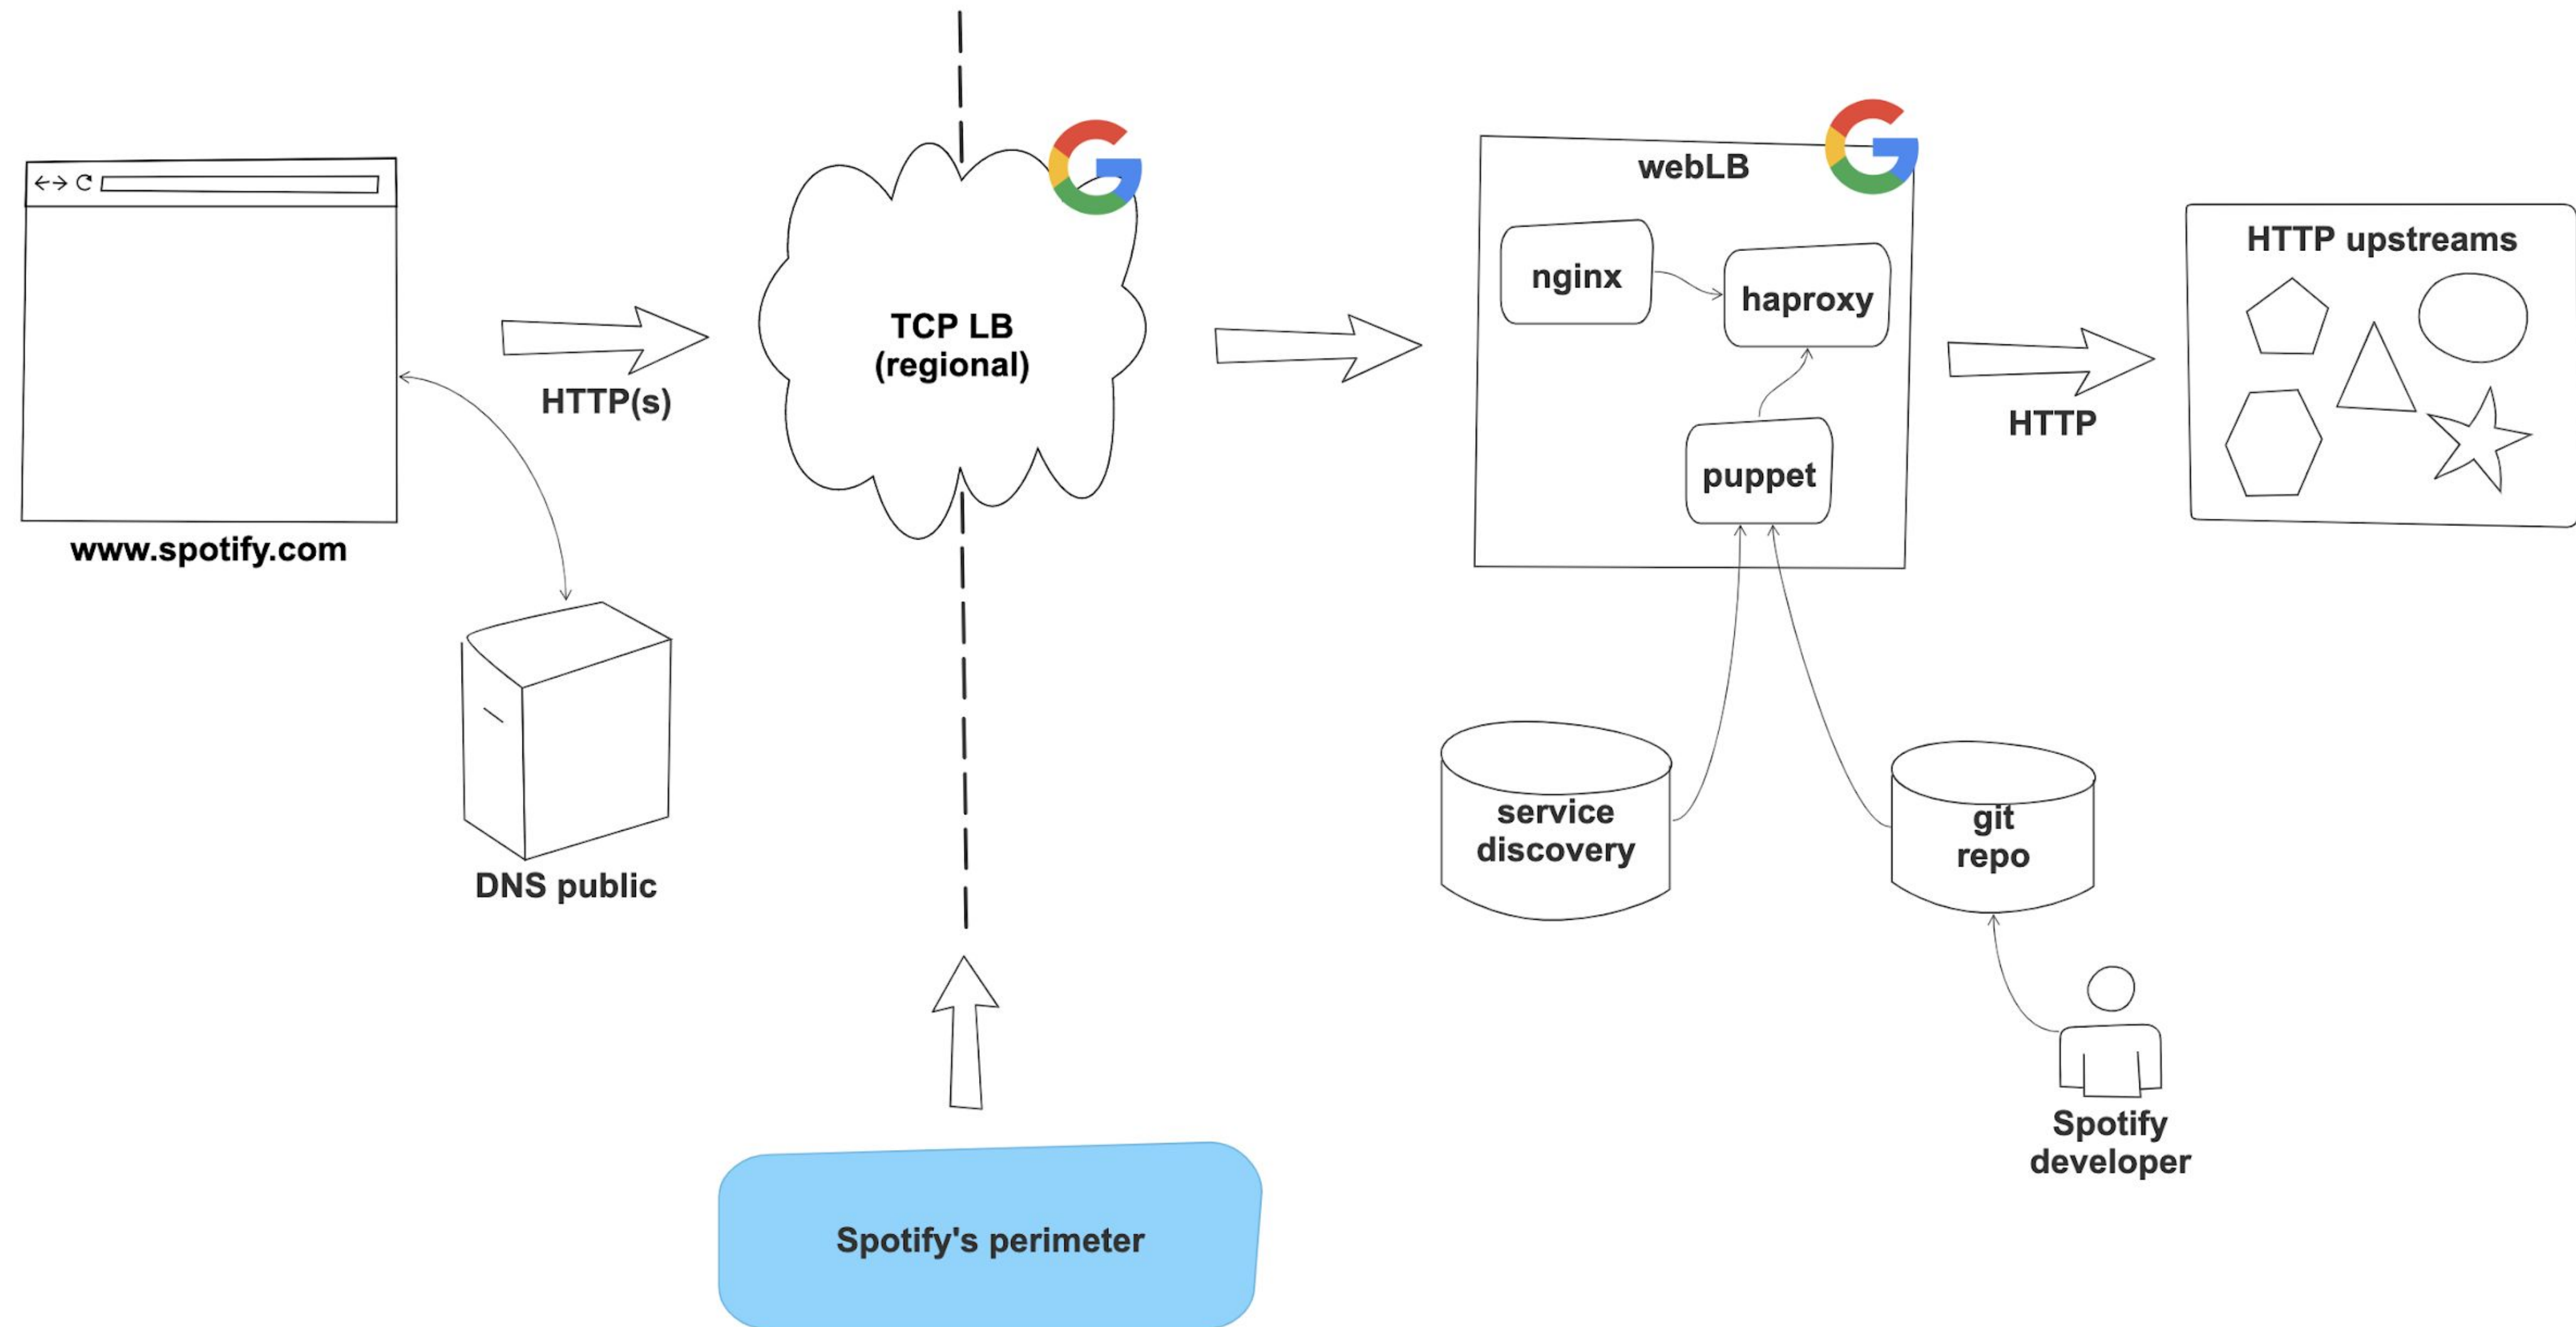


Spotify's perimeter



Hermes: proprietary protocol

Old web gateway: impl

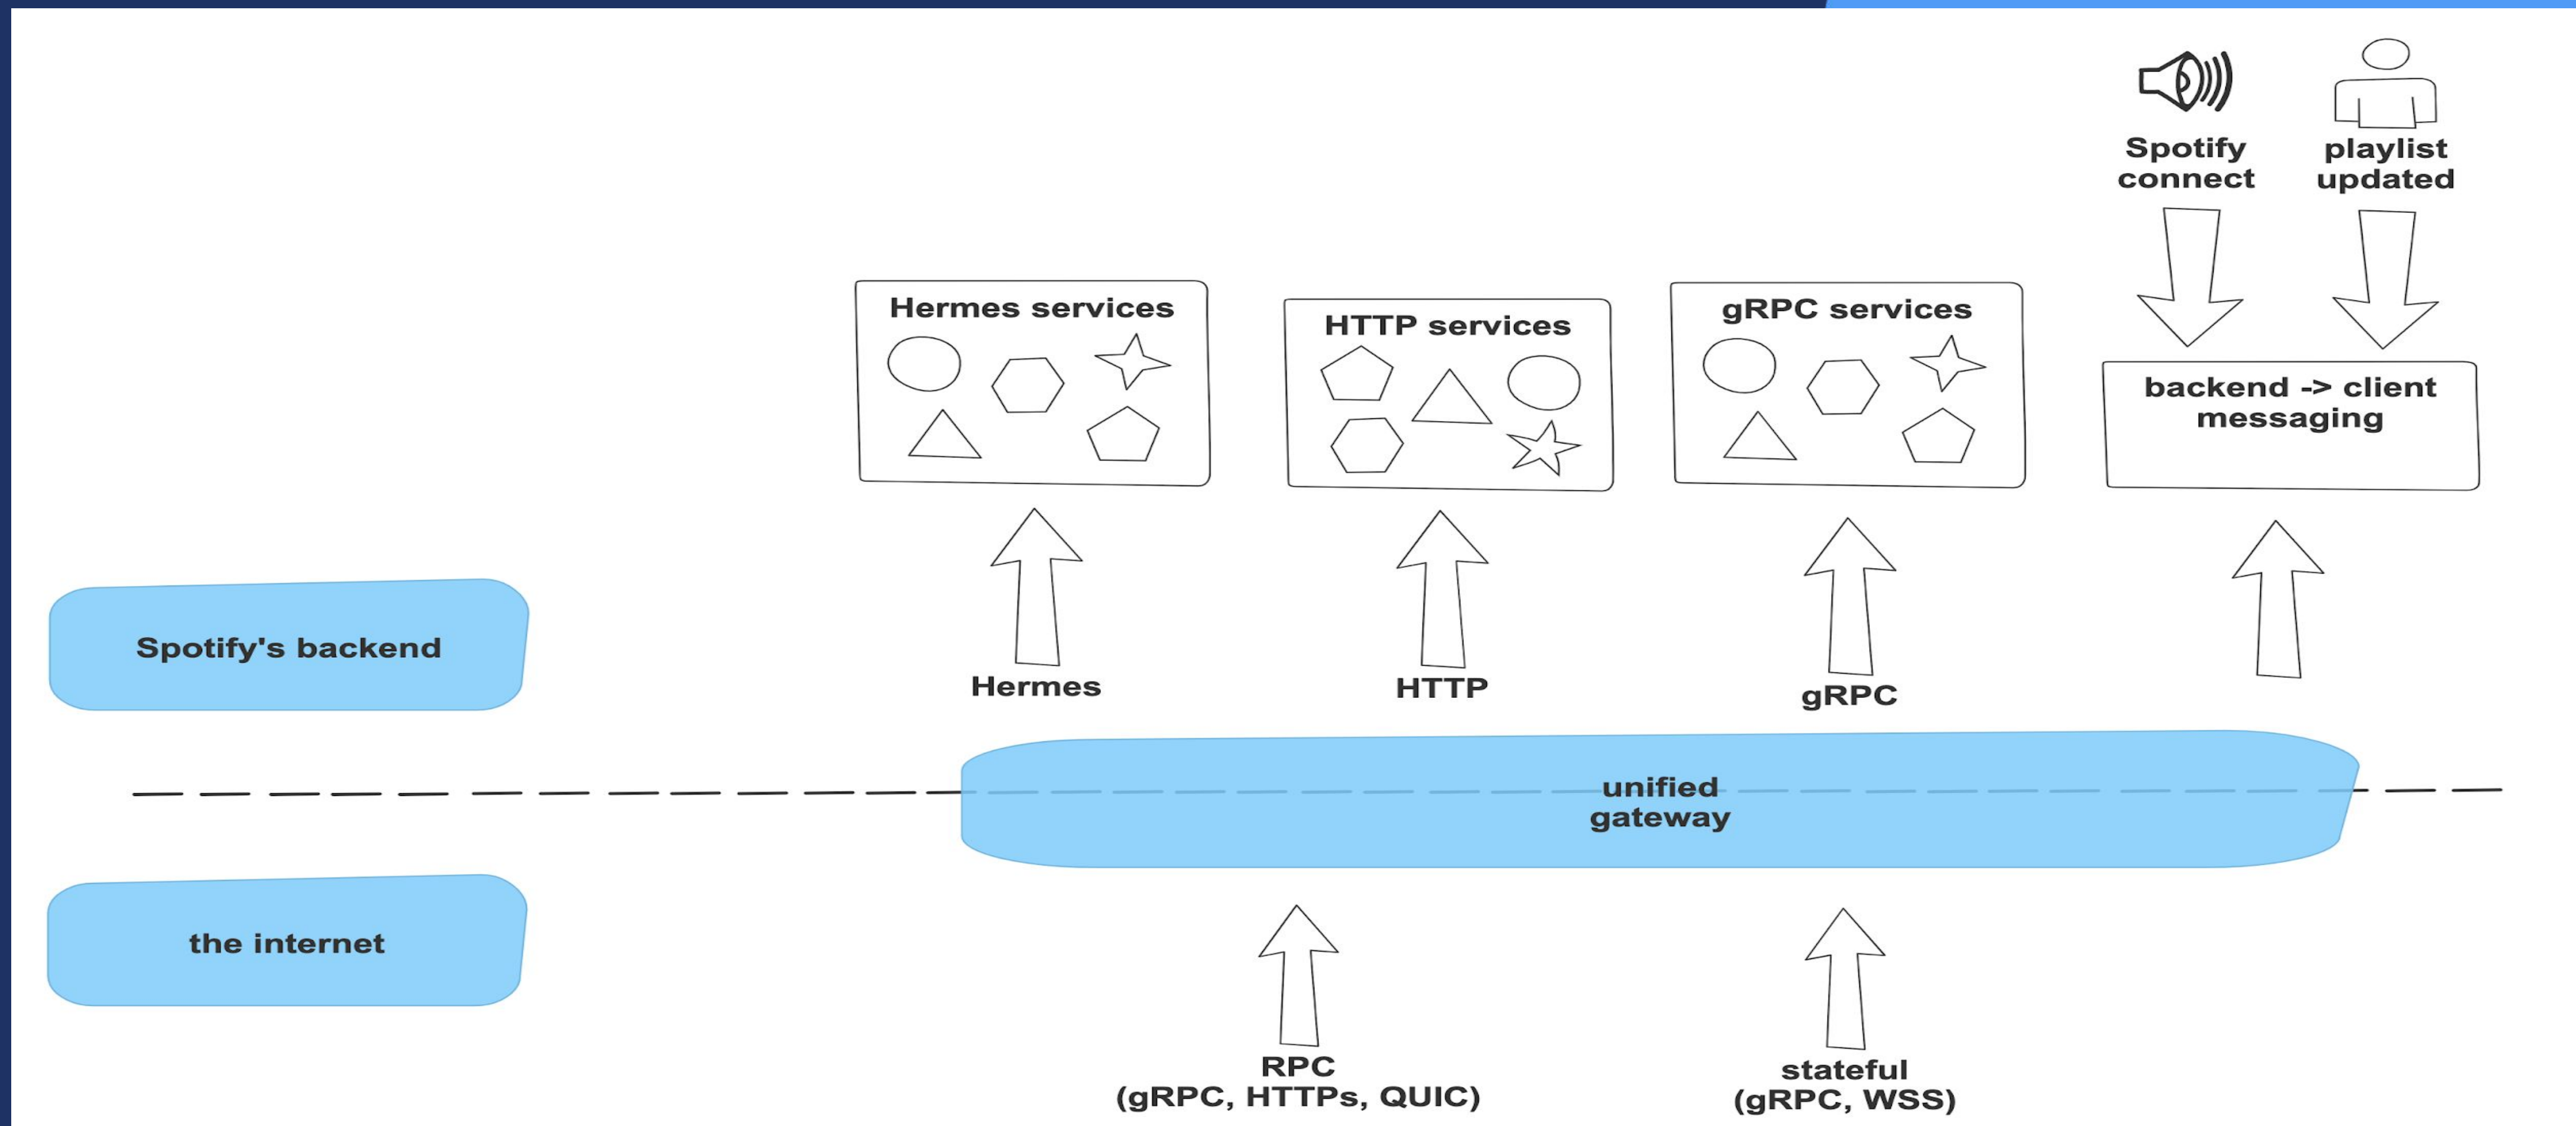


Old web gateway - limitations

- ❑ underinvested
- ❑ service discovery headaches
- ❑ high operational cost

Why migrate?

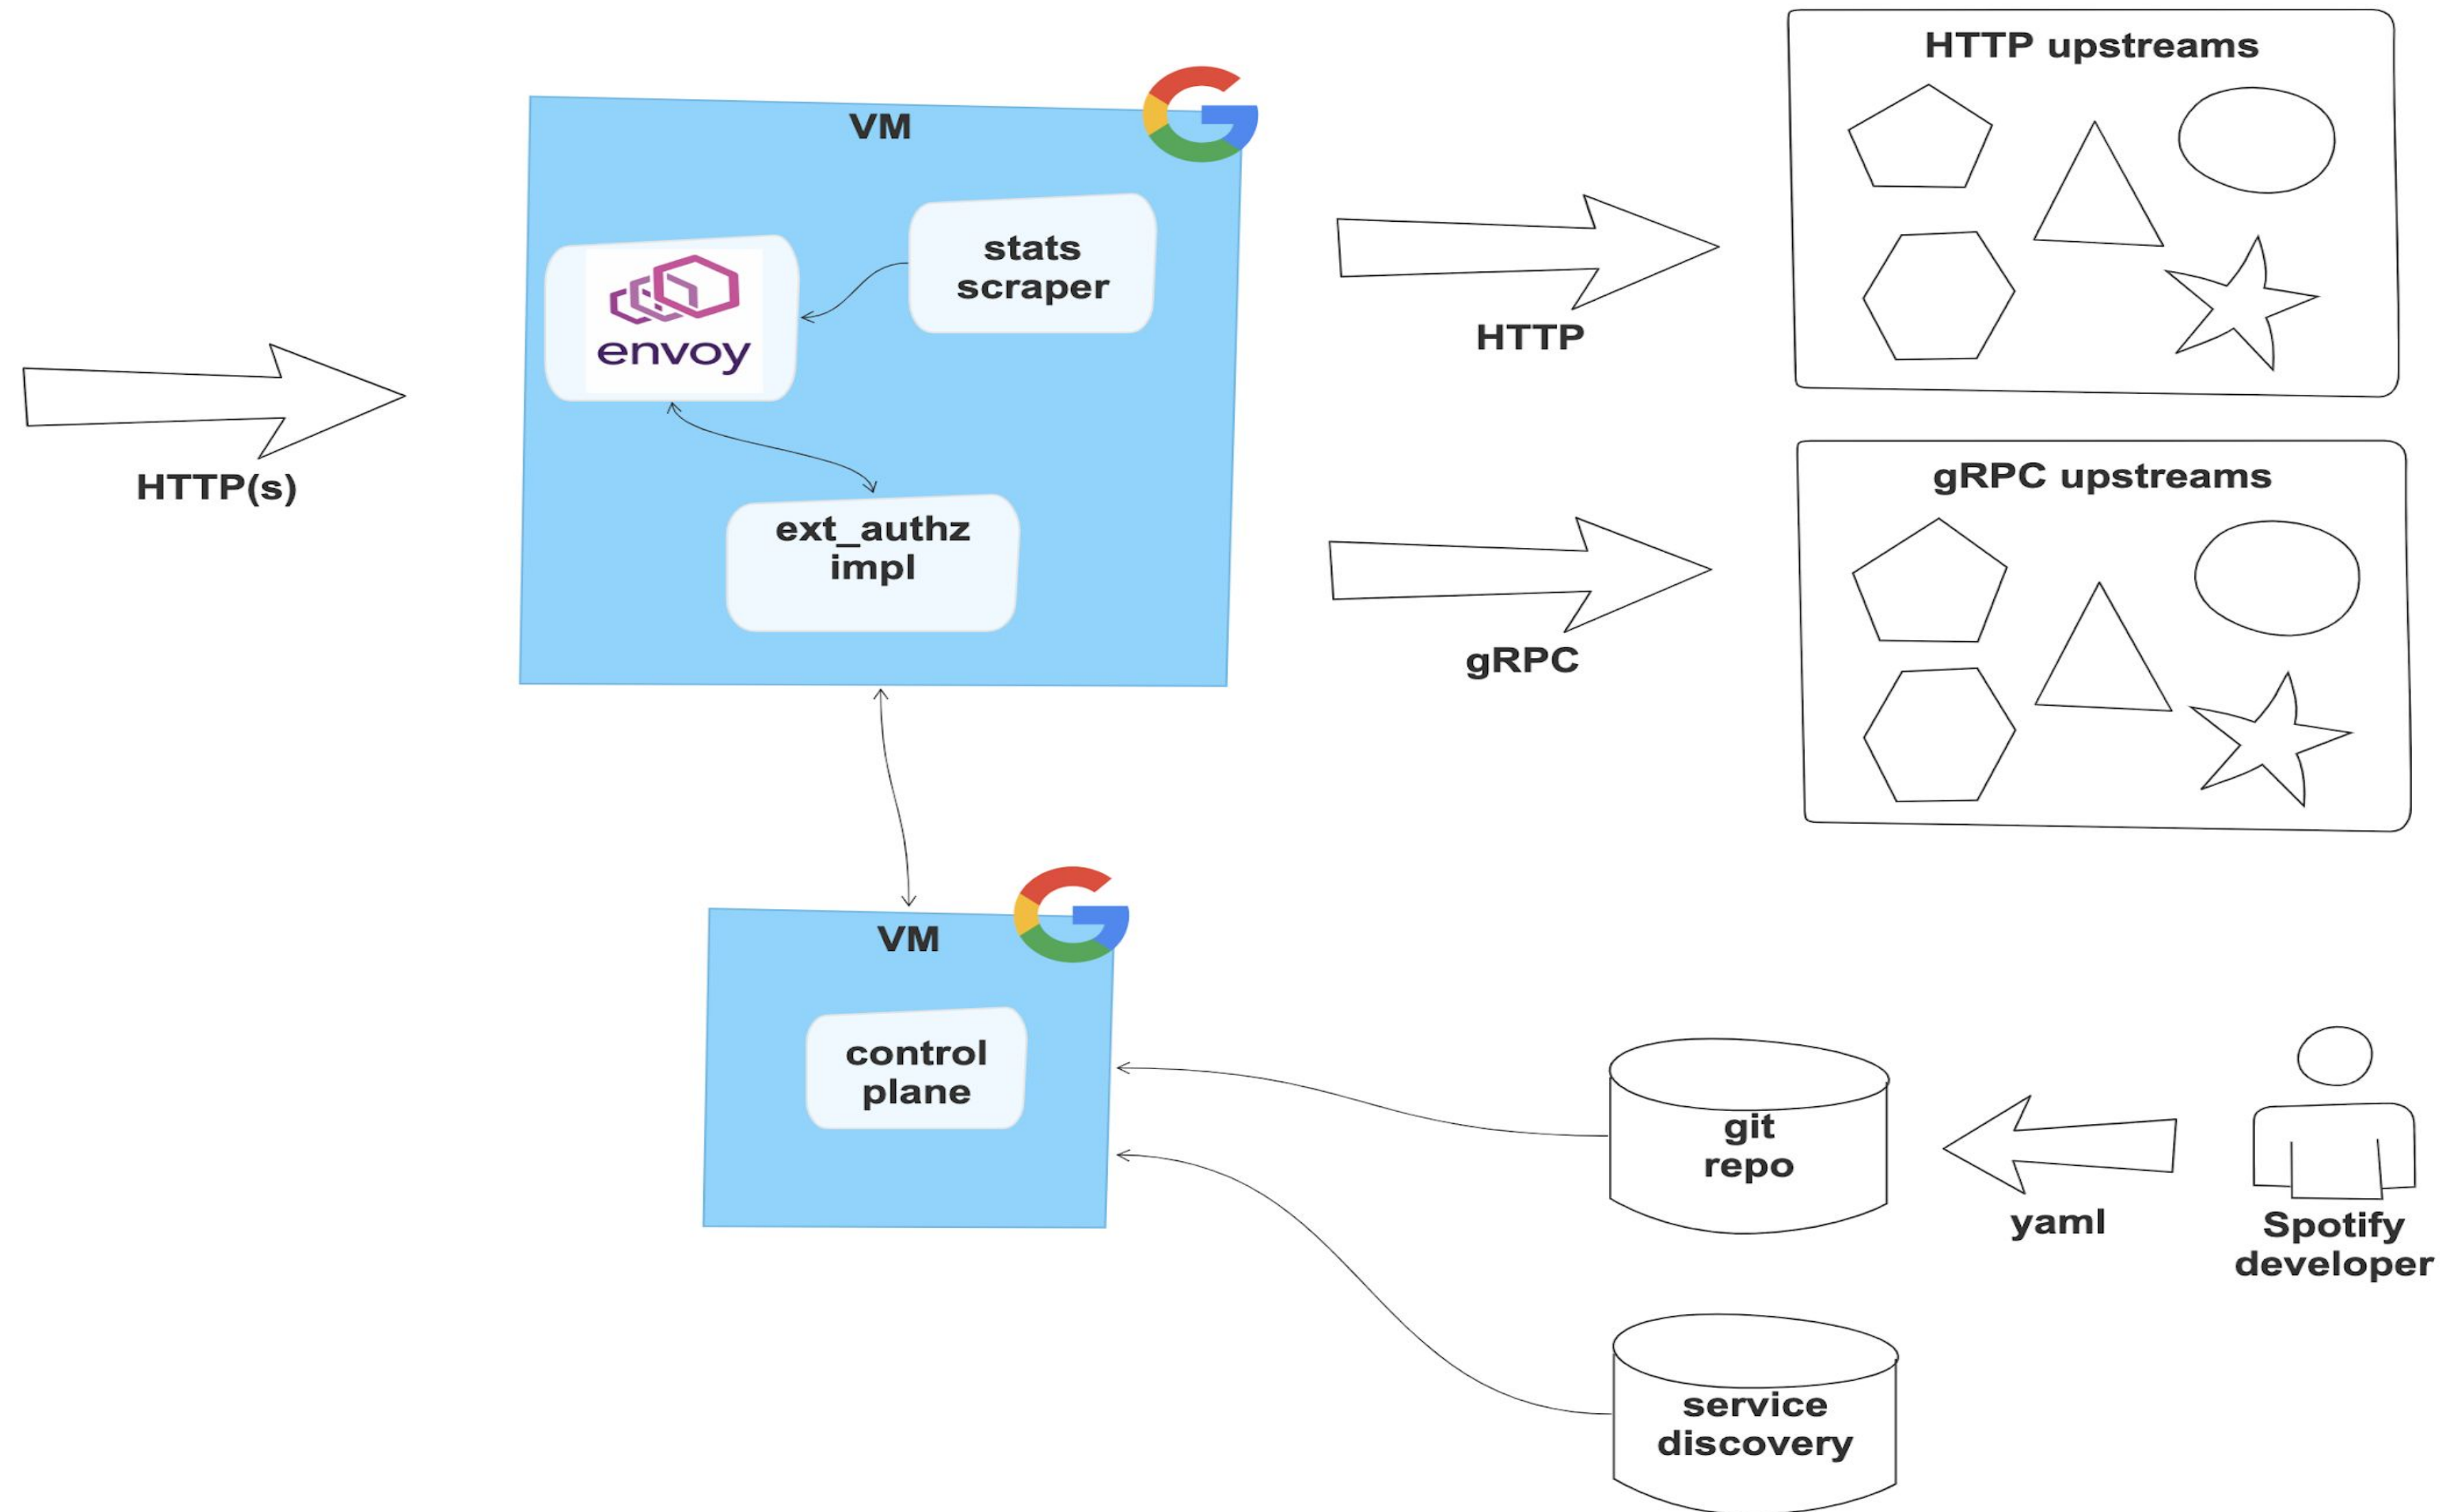
- ❑ unified perimeter
- ❑ avoid fragmentation or duplication of features



Why Envoy?

- ❑ cloud native perimeter
- ❑ xDS
- ❑ vibrant contributor community

Enter Edge



The migration - requirements

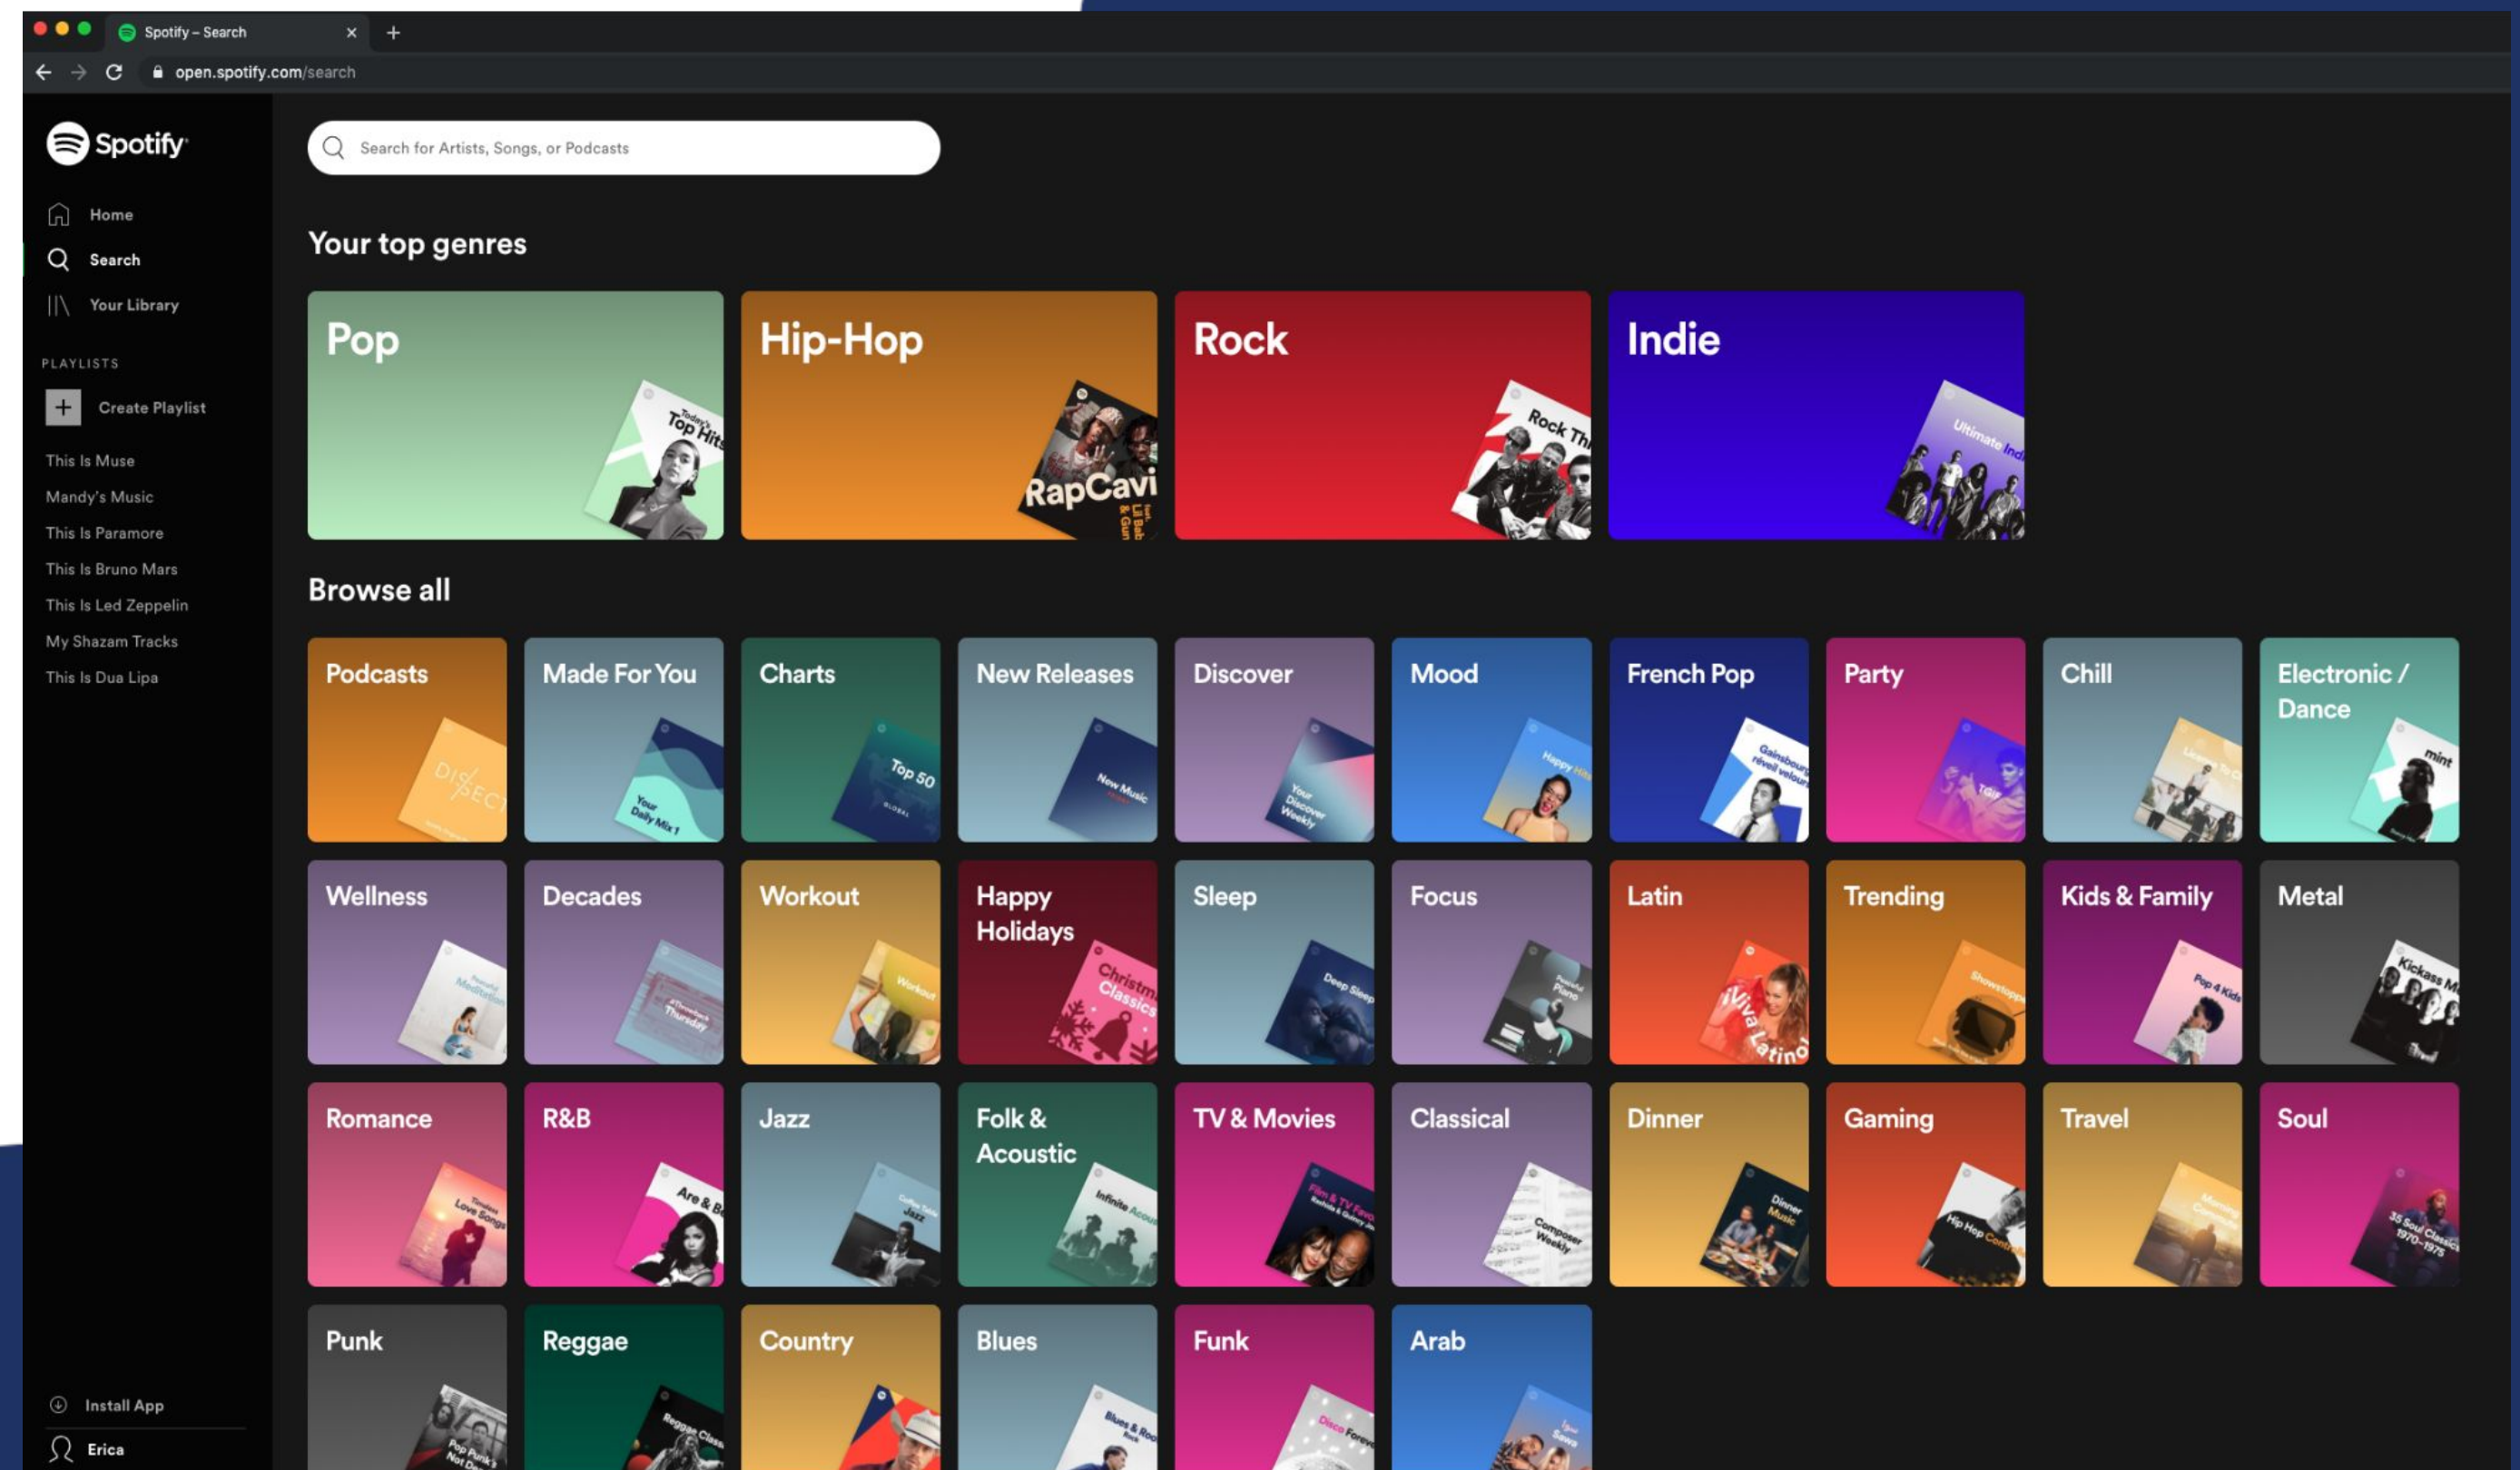
1. transparent to owners of upstreams
2. transparent to end users of upstreams
3. gradual
4. quick fallback

The migration - stage 1

- ❑ non-critical service that we own
- ❑ optimize for speed over safety
- ❑ global DNS change (#yolo)

The migration - stage 2

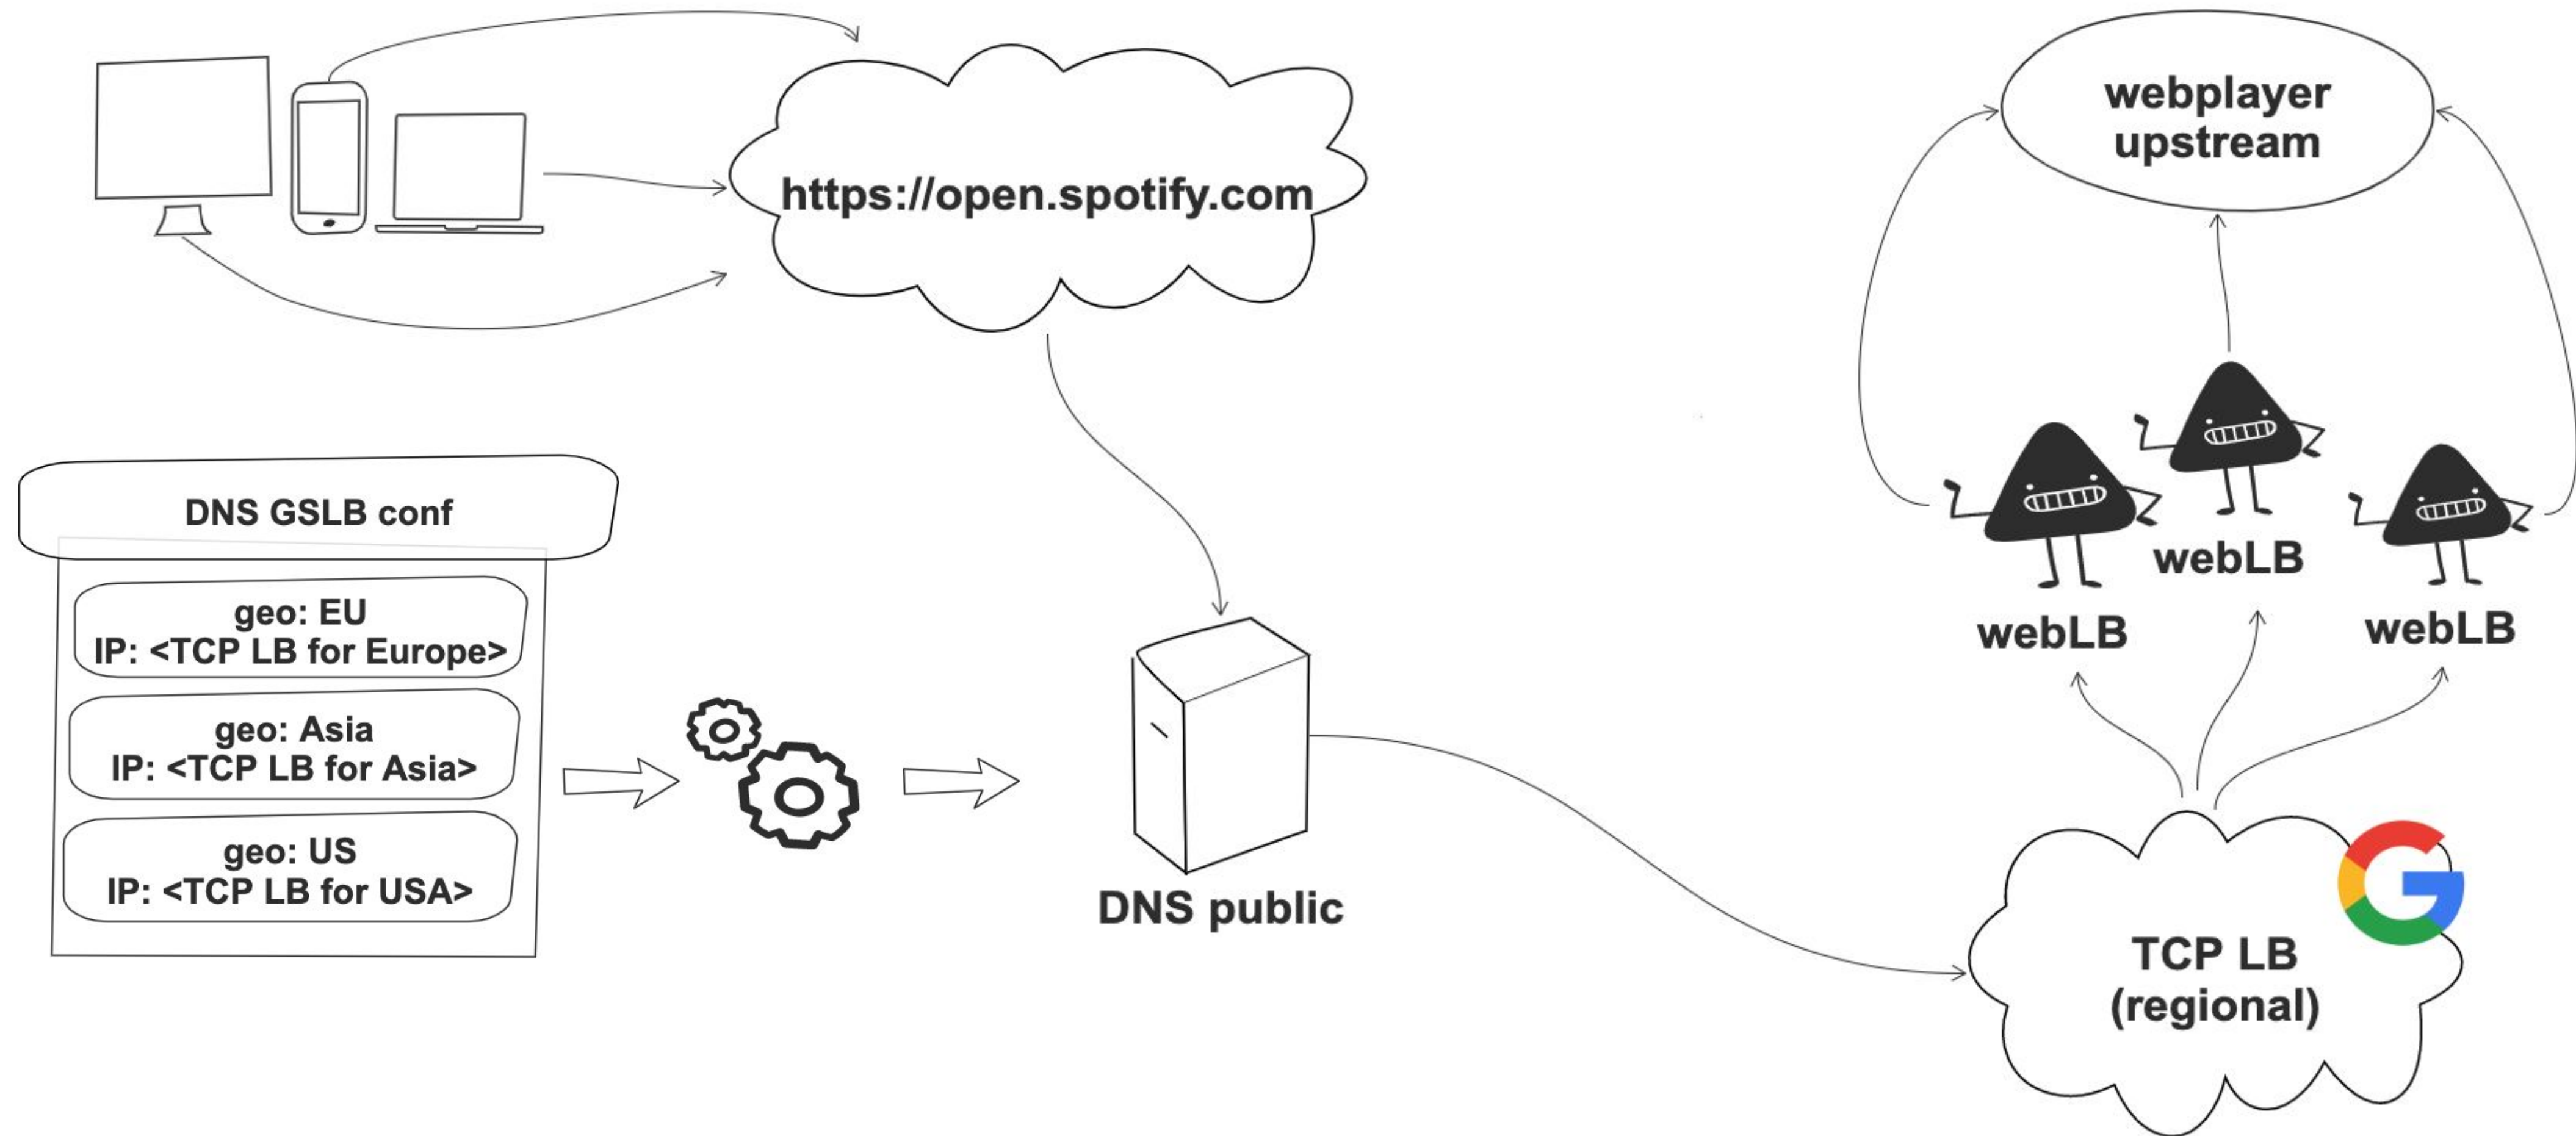
- critical service, owned by another team
- optimize for safety



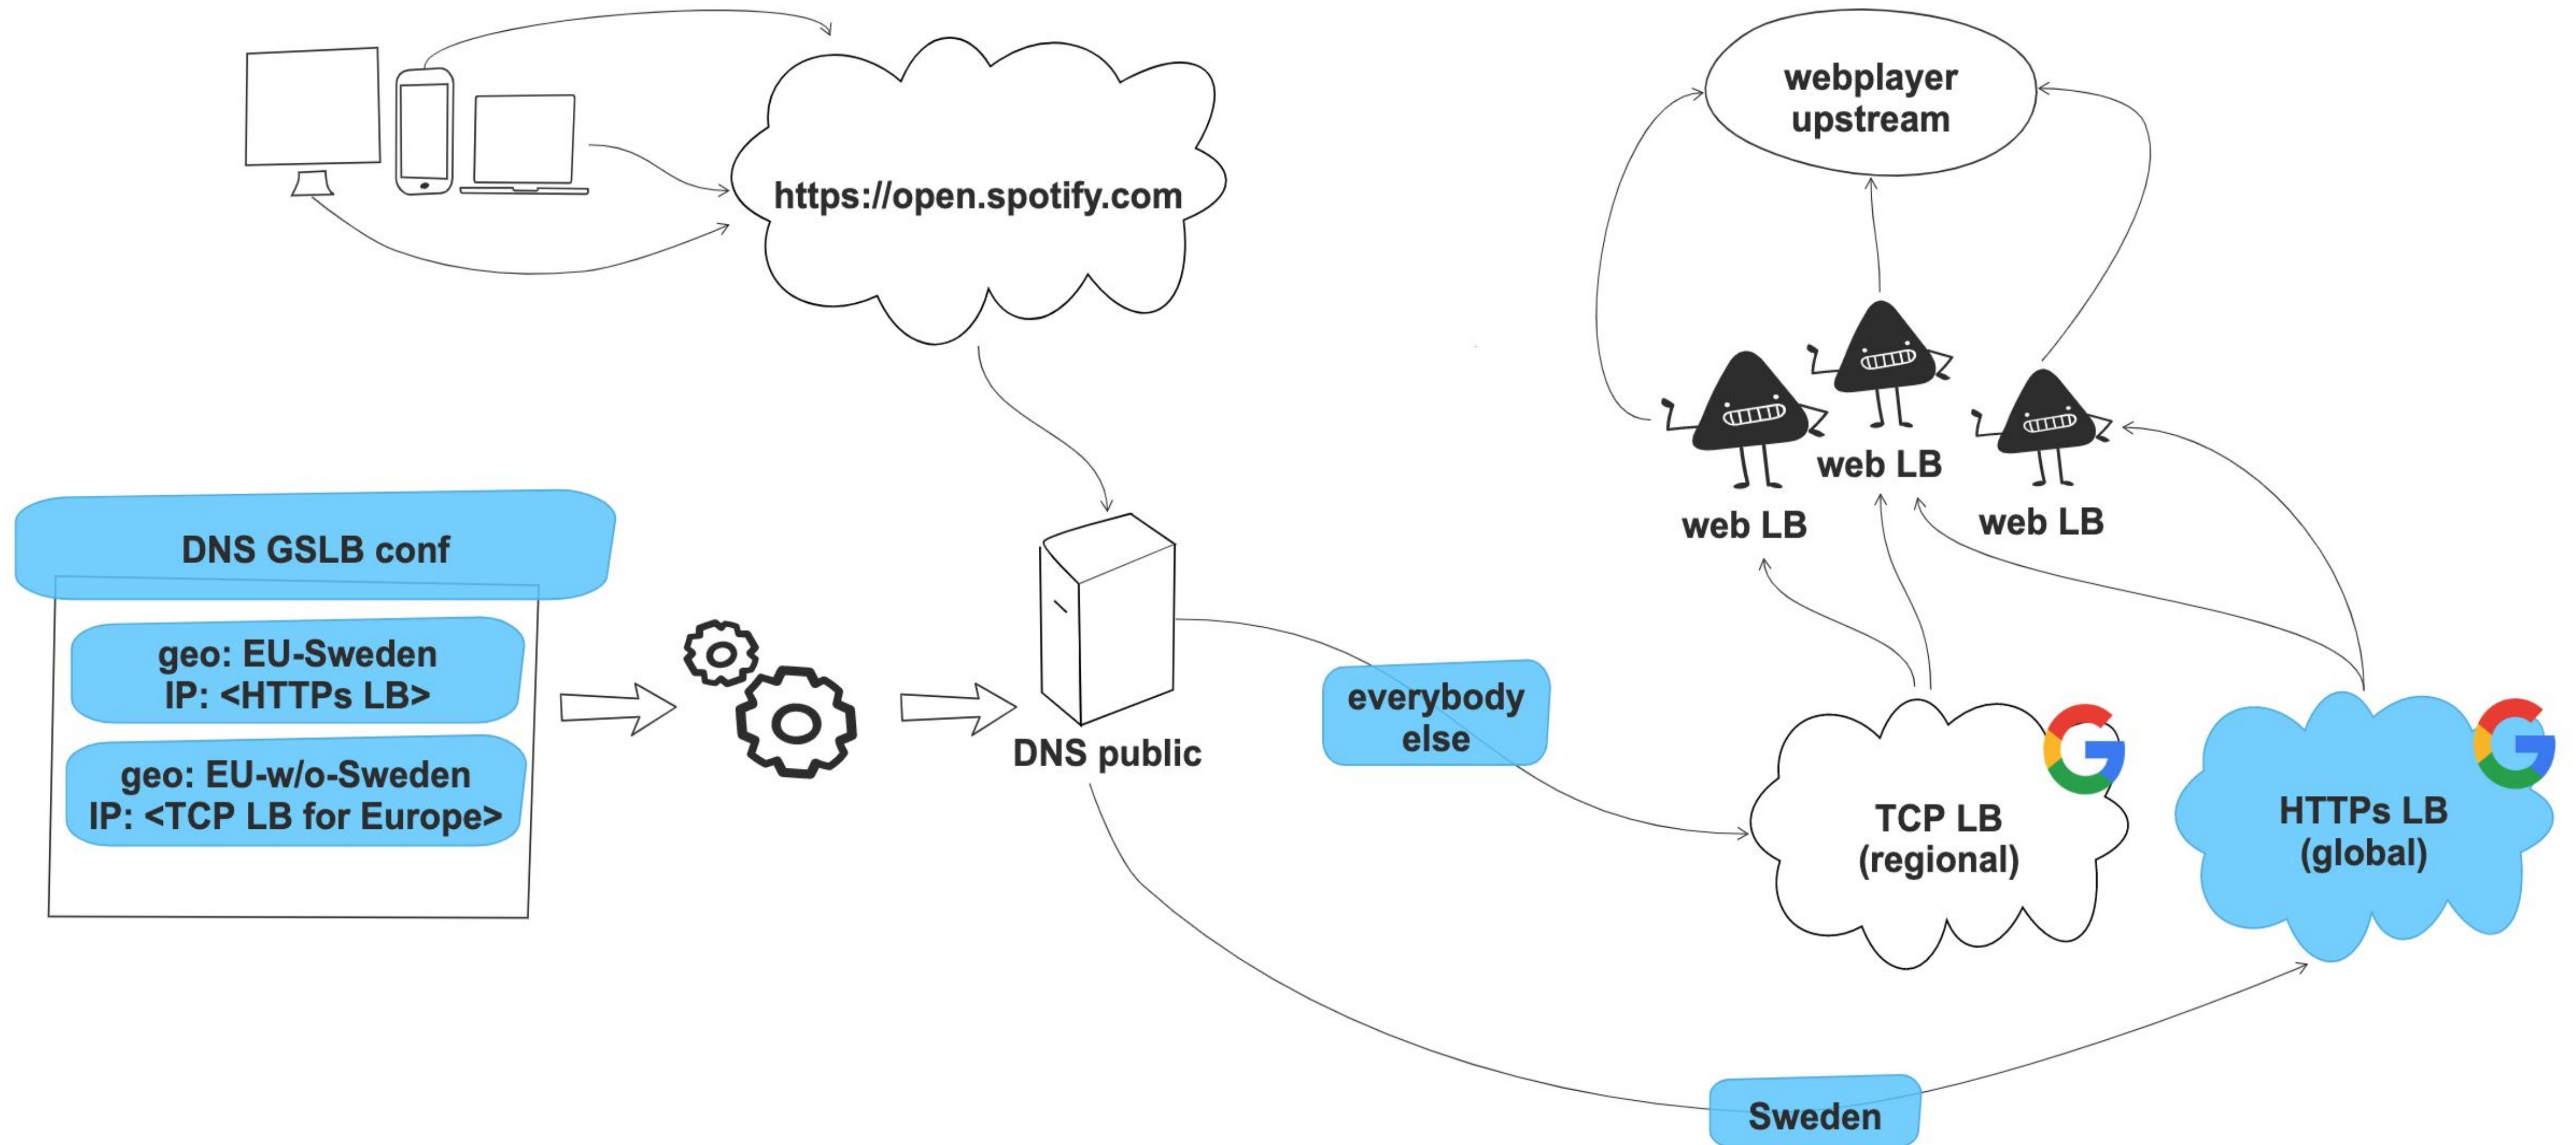
Stage 2 - strategy

1. shift some traffic to Edge
2. monitor, react
3. repeat, until all traffic is shifted to Edge

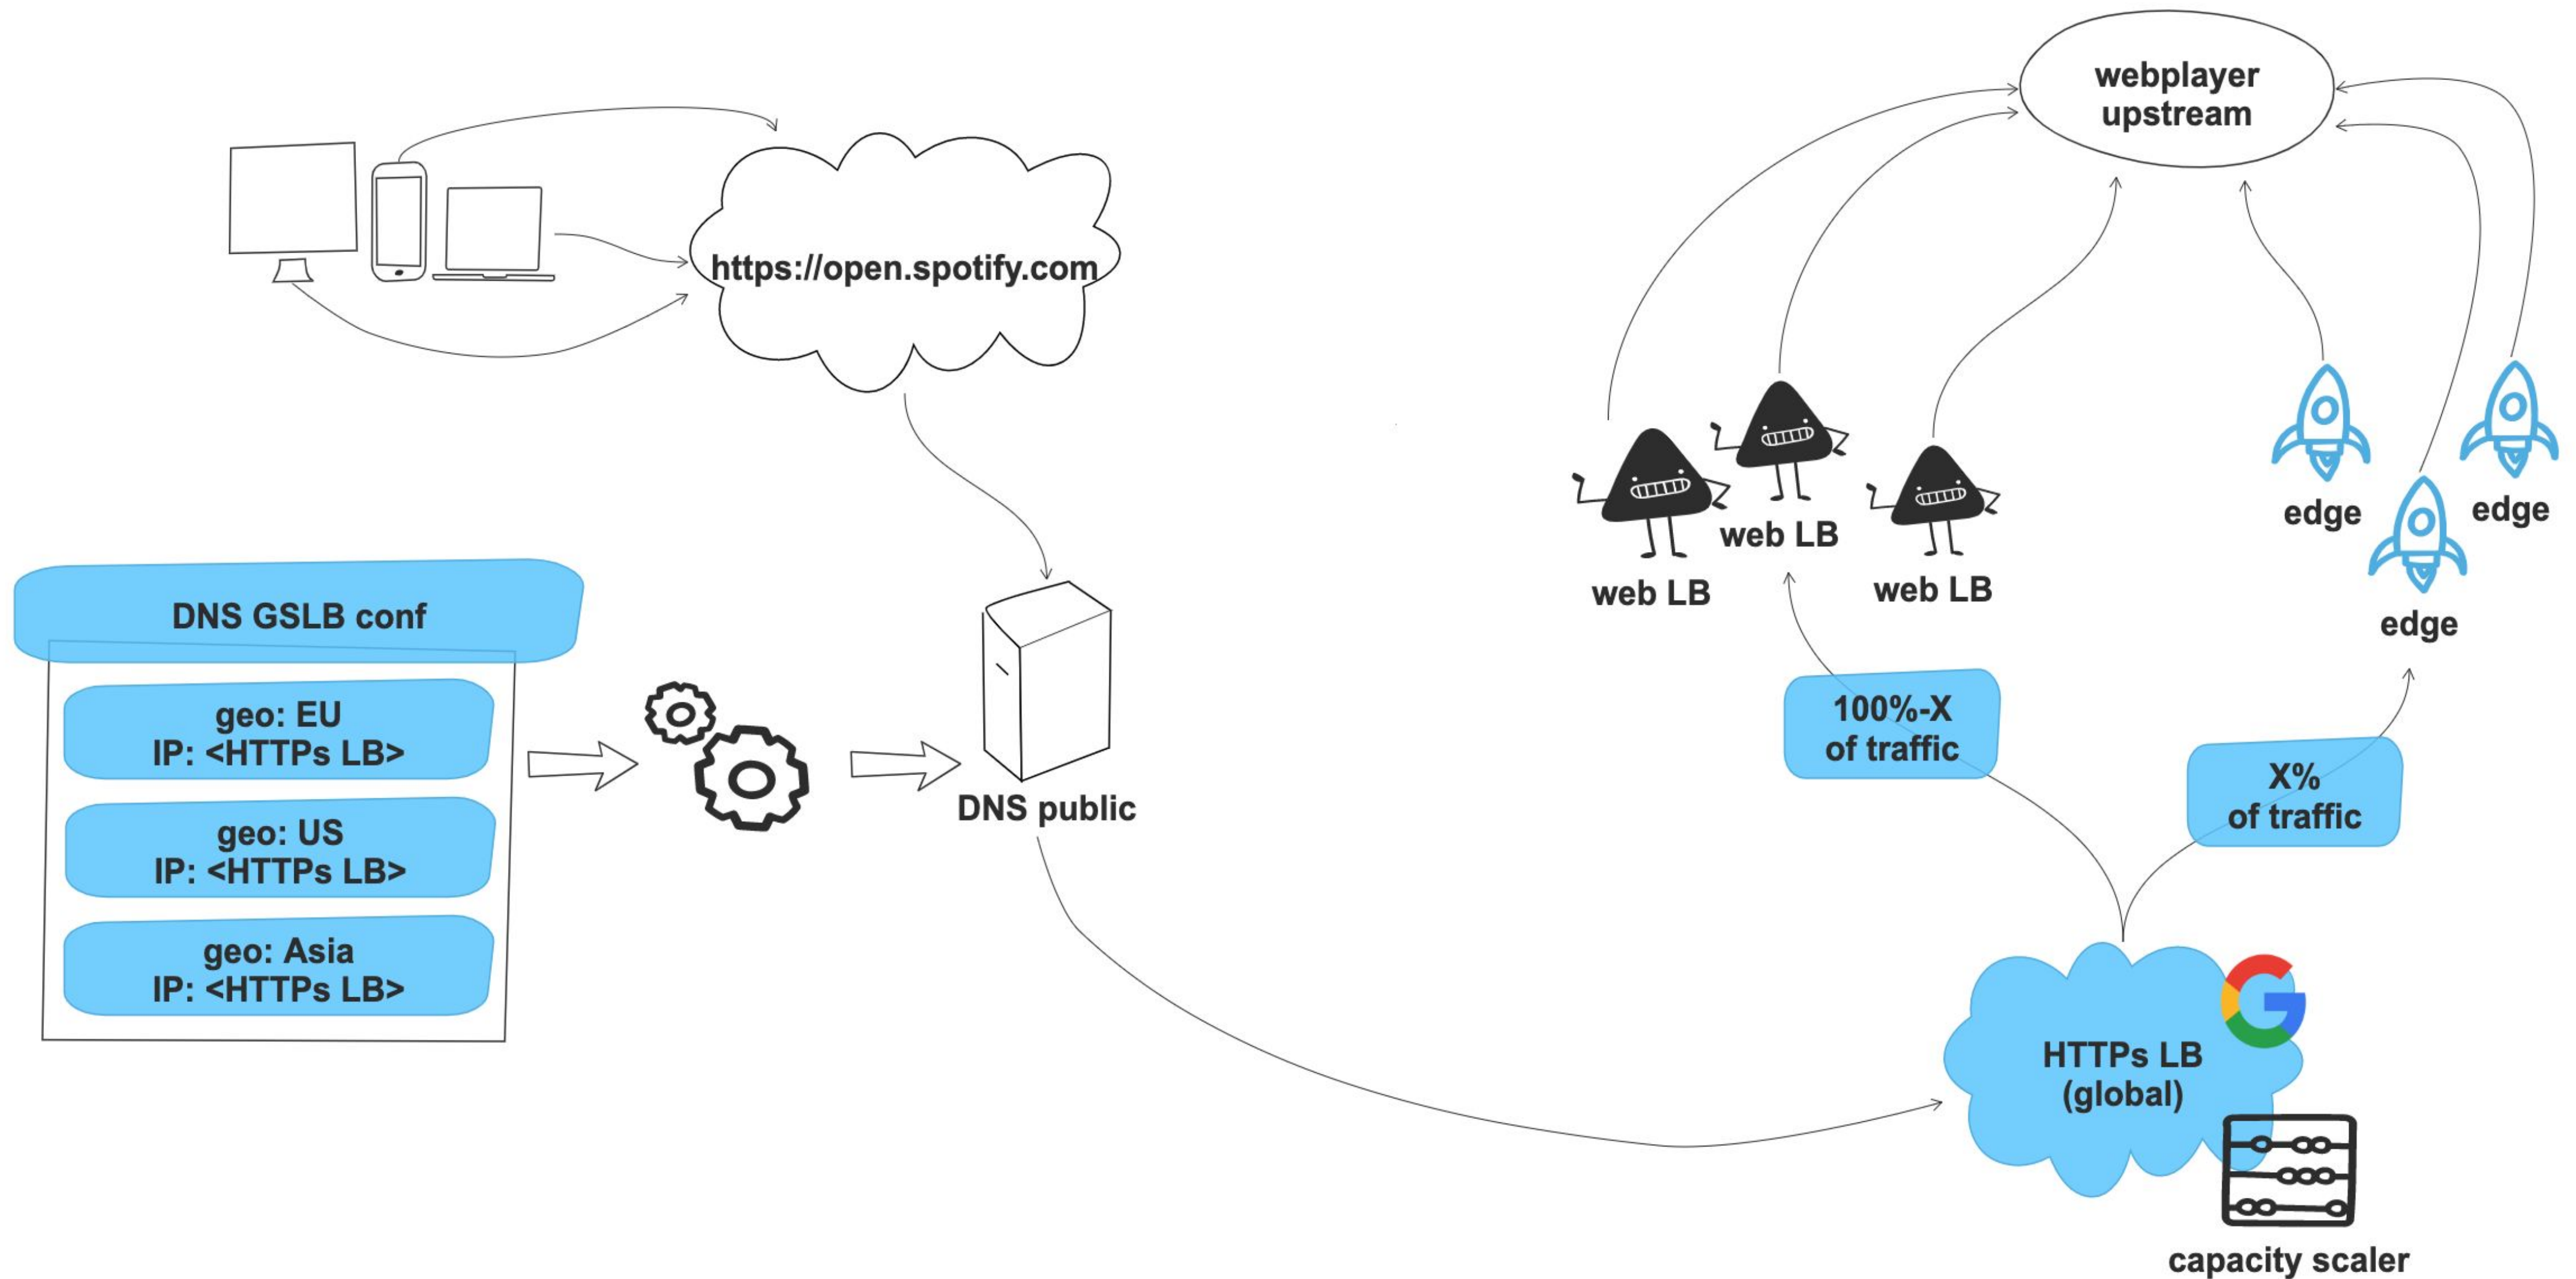
webplayer - before Edge



webplayer on Edge - swap LBs



webplayer on Edge - ramp up traffic



The migration - stage 3

- ❑ 130+ remaining HTTP upstreams
- ❑ reduce migration surface
- ❑ add missing config params to Edge
- ❑ script away!

What went well

- ❑ “panic button” indispensable
- ❑ coordination and communication, trust was built
- ❑ in time for Spotify hack week!
- ❑ contributing back PRs! #1 #2

Learnings and failures

- ❑ zero downtime: not quite
- ❑ grasping the failure modes
- ❑ software entropy
- ❑ state drift

Future plans - short term

- ❑ support for more auth schemes
- ❑ rate limiting
- ❑ automated mitigation against bad deploys

Future plans - long term

- ❑ Move all traffic behind edge
- ❑ remove header decoration from ext_authz service
- ❑ declarative infrastructure
- ❑ leverage K8s

Key takeaways

1. design for failure
2. try, fail, adjust, repeat
3. if possible, reduce migration surface
4. you can do it!



Thanks for listening!

atc@spotify.com

Join the Band!

ATC is hiring (wink-wink)

spotifyjobs.com