

# Automated Traffic Affects IETF Services

Robert Sparks

(data extracted from Cloudflare by Antoine Cordelle)

# Context and Caveats

- We've been reacting, affecting what we can measure.
- Our service delivery has been transforming.
  - About a year ago, most services were directly installed on a single server.
  - Today most are containerized and orchestrated. A few run as VMs.
  - Complex pages, like proceedings, have moved to a periodic pre-computation model
- Changes to the shape and capacity of our service delivery make comparing recent and past data challenging.
- We currently try to make our content as available as possible to well-behaved bots.

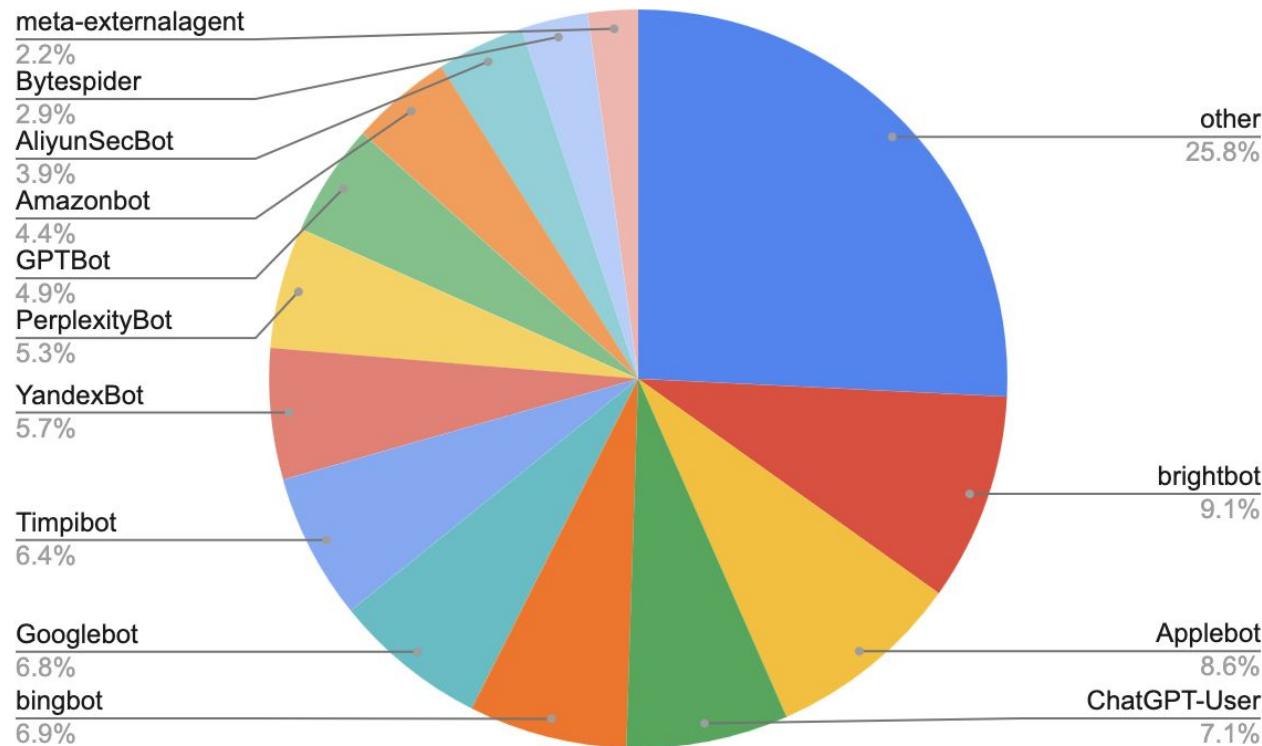
# Current Monthly Traffic for all ietf.org Web Services

14 June to 14 July

- 3.5B requests total - 3.23B junk immediately denied
- 277M remaining
  - 155M served from edge before cache
    - 62.1M Manage Challenged - 0.38% (227K) answered successfully
    - 1M served directly with 200
  - 82M requests served from cache.
  - 50M were the infrastructure talking to itself (logging/metrics)
  - 29.8M requests were served by origin

➡ 20M of traffic from bots

# \*.ietf.org bot requests to the edge, 14 June to 14 July



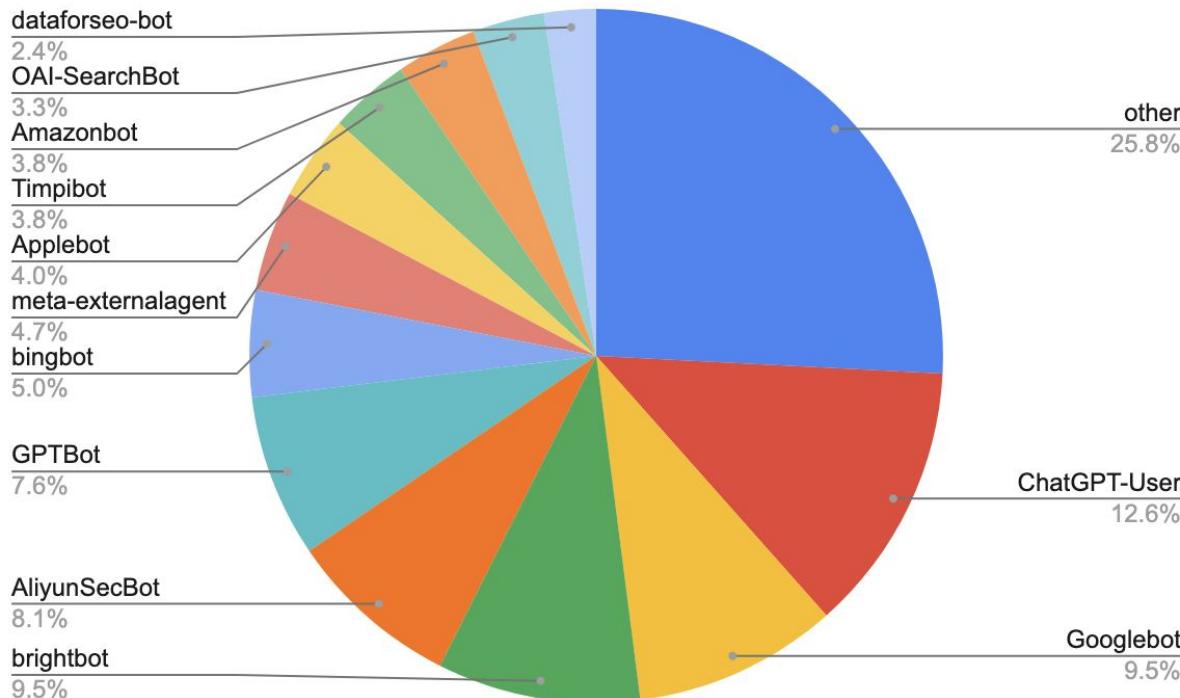
# Current Monthly Traffic for datatracker.ietf.org

14 June to 14 July

- 70.9M requests total
  - 30.4M handled by Cloudflare configuration at the edge, before cache
    - 22.2M were Managed Challenged
    - 75K served directly with 200
  - 21.7M requests served from cache
  - 18.7M requests served by the datatracker

➡ 7.4M of traffic from bots

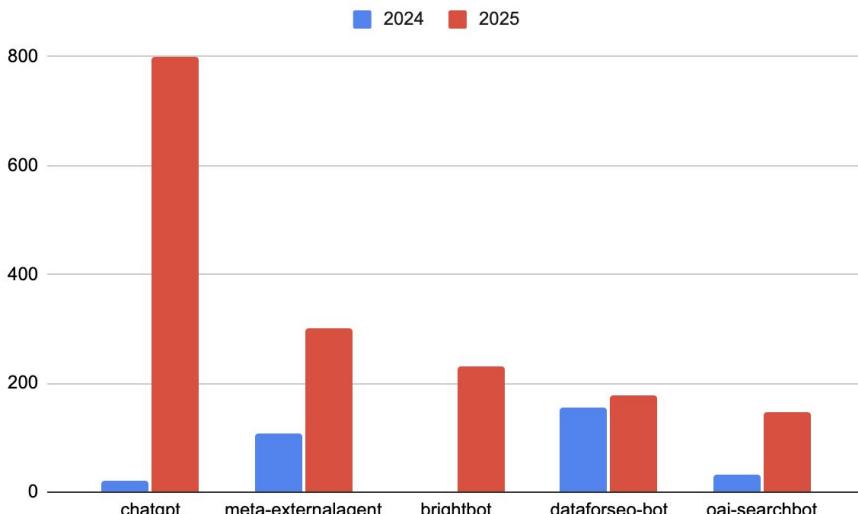
# datatracker bot requests to the edge: 14 June to 14 July



# Changes since last year

Requests served by the datatracker origin in one month

More requests



Fewer requests

