

Happy Families in Tor

Alexander Hansen Færøy

2025-06-11 All Hands

Context and problem: Old family system is causing us trouble:

- Poor UX for Relay Operators: Each time they add a new relay, they need to update all their other relays with information about the new relay.

This issue is fortunately largely mitigated by our relay operators using various shared tooling for configuration management.

- Our relay operators are really good at what they are doing. They seem to have a lot of resources, and a lot of computation power. This leads to them being hungry to run more and more relays as network bandwidth and computation power increases.
- Unfortunately, the C implementation of Tor doesn't scale very well with modern multi-core CPU's. This leads to people running multiple instances of C Tor per computer.

The Network Team believes Arti will solve this, and it's one of the many reasons we are currently working on Project 141 (Arti Relays).

- Memory and storage space size issues (particularly bad for Apple iOS users):

```
~/.tor % wc -c cached-microdescs
24724290 cached-microdescs
~/.tor % grep family cached-microdescs | wc -c
20225420
```

23.58 MB of cached micro descriptors consists of 19.29 MB (82%) family info.

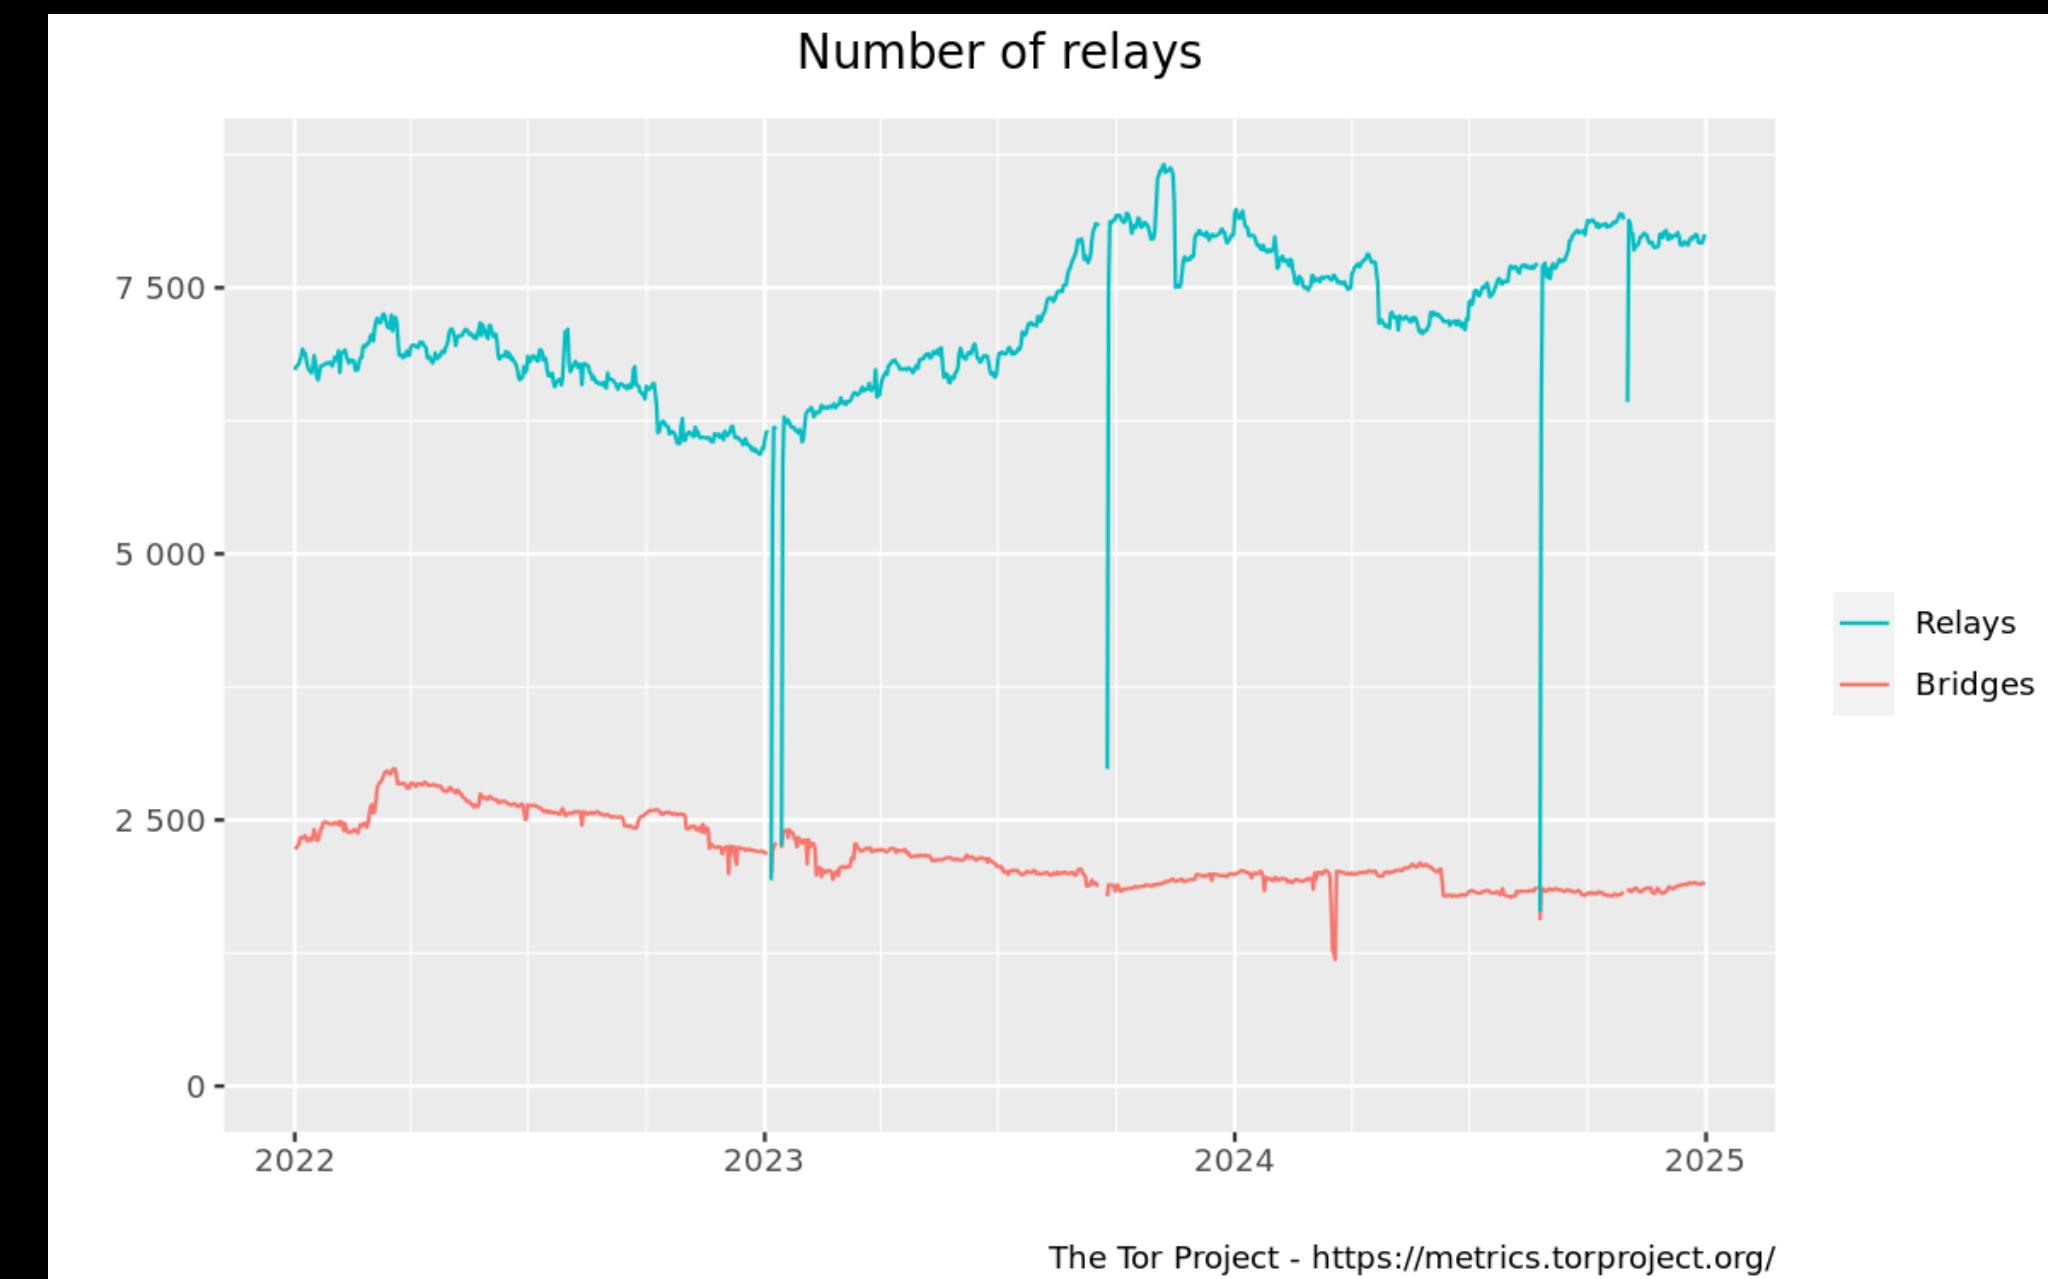
Context and problem: Old family system is causing us trouble:

- Because of many requests from our relay community, we decided to allow more relays per IP. This is generally good for the network, but it made our space issues more problematic due to $O(n^2)$ behaviour of our current family system.

Timeline:

2023-01-31 AuthDirMaxServersPerAddr was bumped from 2 to 4 relays per IP.

2023-06-28 AuthDirMaxServersPerAddr was bumped from 4 to 8 relays per IP.



Solution: Proposal 321: Better performance and usability for the MyFamily option (v2)

- Proposal written by Nick Mathewson during 2020 as part of the work on “Walking Onions”.
- Gives each relay operator a new identity key: the family identity key.

This key gives us some new opportunities: we can now both validate ownership of an individual relay, but also a set of relays (a family).

- Makes analysis work a lot easier: you already have a 256-bit public key that you can use for “group by”-like operations on relay data.

Before this, we would have to build a set of relays from the family value, add the relay ID itself, and hash the set together to create a “group by” identifier.

- But Network Team, when can we have this shiny new piece of technology?

“Thanks” to the situation in 2025-Q1, this got implemented in:

- Arti 1.4.1 (released 3rd of March, 2025);
- C Tor 0.4.9.2-alpha (released 2nd of April, 2025).

See also: <https://blog.torproject.org/happy-families/>

Questions?