Alternative Services: Problems and Principles

Context: HTTP serving is complex

A deployment of any significant size is comprised of
- Multiple servers
- Maybe multiple server configurations
- Potentially multiple server operators

The service needs to maintain ~~near~~ constant availability
- Only as a whole
- Individual instances do not need to

At any given time, there is a right server to be talking to

Often that is the current one

Sometimes it isn’t … probably

Alt-Svc can be wrong

In a multi-operator configuration, Alt-Svc might not be good for identifying servers run by a different operator (#1673)

Problems arise most often *when values are cached*

RFC 7838 has ~~rules~~ heuristics for cache invalidation

- These are not good enough
- These are only tied to changes the client observes
- What we want is invalidation on
  - Network path changes
  - Server configuration changes
- …but we can't consistently detect when those happen

Information provided over an authenticated TLS connection is the final authority

Mistakes cannot produce security bugs

…they might cause performance bugs though

Information sources

DNS can only contain hints

- #include “usual_dnssec_debate.h”

Cached data ~~can~~ will be wrong

Information sent over HTTP could be authoritative

- But should we treat it as such?
- Is the current server instance a good source of information about other server instances?
- Or is it only good for a signal that it believes that it is no longer the best place to talk to? (for shutdown, load shedding, etc…)

Performance bugs are bugs

Any hints should therefore be complete and accurate

…as much as possible, at least

Many performance bugs

- Make before break can hide problems
  - Alternatives can fail without any visible effect
- Alternative information might be usable, but wrong
  - Wrong server instance
  - Improvement over current, but not the best
- Caching hints might improve performance … or not
  - Cached data can become outdated at any time
  - Lack of cached data can mean getting
  - Restarting without cached data still means relying on cached data
  - … from DNS, which could be stale too

Reusing connections is better than making new ones

Except when you reuse the wrong connection

Connection reuse issues

Can a connection to an alternative be reused? (#1696)

How does this interact with ORIGIN? (#1691)