# DNS Resolver Selection

Some questions about privacy

# The question we’re asking:

Can you reduce the concentration of information about client activity by distributing DNS queries across different resolver services?

# Common DNS Resolver selection patterns

There is one configured resolver, and all queries go to that resolver.

There is a resolver used for all local network interfaces but a separate resolver configured for VPN interfaces.

Per interface DNS resolver lists

# Resulting DNS Resolver query patterns

There is one configured resolver, and all queries go to that resolver.

*So this resolver knows everything about the client’s query traffic.*

There is a resolver used for all local network interfaces but a separate resolver configured for VPN interfaces.

*The common resolver gets all query traffic except that for the VPN; the VPN gets all traffic passing through the tunnel.*

Per interface DNS resolver lists

*What it says on the tin:  each resolver gets queries only for the traffic on that interface.*

# Privacy implications of common approaches

Sending all the clients' queries to all available resolvers results in all of them knowing about all the resolution events.

Send all queries to one resolver results in that resolver seeing everything.

Having a separate resolver for a VPN interface avoids queries via that network leaking , but queries  for political-party.org, recreational-habit.com, and travel-destination.info all likely go to the same place; my-workplace.example is the only one likely to get separated out.

The privacy implication of interface lists depends on how you switch traffic among the interfaces.

# Can different resolver selections improve privacy?

# Can different resolver selections improve privacy?

Maybe?

# Can different resolver selections improve privacy?

Okay, it’s clear that carefully selecting the resolvers to include in the candidate set for their privacy properties will help.

Beyond that, though,  there are tradeoffs in client-based and name-based approaches.

Before going through those in detail, we want to pose these questions to the working group:

Is the WG willing to consider approaches that result in multiple DNS resolvers being discovered and used?

Is the WG willing to consider optimizations for reducing the concentration of information about client activity?

Is the WG willing to work on describing the operational impact of those optimizations?