ericsson.com

Are relay services the key to online privacy?

Mirja Kuehlewind Mirja Kuehlewind Master Researcher, Networks

10–12 minutes

In the past few years, several actions to strengthen user security and privacy on the Internet have been taken. Focusing on techniques that have been developed and deployed to prevent unintended third parties from accessing and manipulating communication. The wide-scale deployment of HTTPS to provide encrypted communication for web content access is one example. Efforts to deploy new protocols that encrypt associated data like domain name system (DNS) requests is another. These techniques prevent passive observers from knowing the exact data that is exchanged, but they can still deduce which parties communicate and sometimes even which services have been requested.

More recently, focus has centered on various mechanisms that enhance user privacy by concealing even more data and metadata from Internet communications. These mechanisms often rely on a relay service that intermediates the communication between two hosts. While at first glance it seems counterintuitive to involve yet another party in the communication process to protect user privacy, the logic behind using relays is that in the end, each party can only access a limited set of information, and therefore each party knows less than before.

These relay services specifically aim to separate two important pieces of information: knowledge of the identity of the person accessing a service is separated from knowledge about the service being accessed. This requires two levels of encryption.

Relay services in practice

Simple VPN services only add one level of encryption to the link between the client and the VPN server: the VPN server can still see which services are being accessed, and by whom. In the Internet Engineering Task Force (IETF), the leading standardization body for Internet technologies, most of the activities related to these goals of separating information are indicated by using the term Oblivious, but there's also MASQUE (Multiplexed Application Substrate over QUIC Encryption) and a new to-be-chartered group called PPM (Privacy Preserving Measurements) that apply this communication pattern to different use cases.

MASQUE is a new IETF working group that extends HTTP CONNECT to initiate and manage the use of QUIC-based relays. Catch up on the background of QUIC and MASQUE in our earlier blog posts. MASQUE is a tunnel-based approach similar to VPN services. However, it sets up an encrypted connection to a relay, or so-called MASQUE server, using QUIC as a tunnel transport, then forwards traffic through that tunnel to a target server or another relay (see Figure 1).

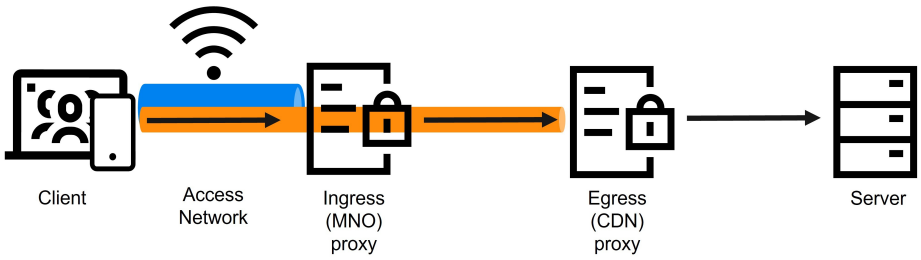

Figure 1: Setup with two MASQUE proxies, e.g. hosted by the Mobile Network Operator (MNO) and the Content Distribution Network (CDN).

At present, the most recent example of such services being deployed is Apple's new Private Relay service, which is in beta testing for iCloud+ users. When activated, Private Relay uses both the MASQUE and Oblivious DNS protocols for web traffic emitted by Safari, and for all DNS traffic.

In both cases, user traffic either for a web server or a DNS resolver is encrypted then first sent to a relay service that knows the user's identity and IP but doesn't see the user request itself. This relay service then forwards the encrypted traffic to another relay, which

can determine where to forward the end-to-end encrypted service request but doesn't know the user's identity or IP address.

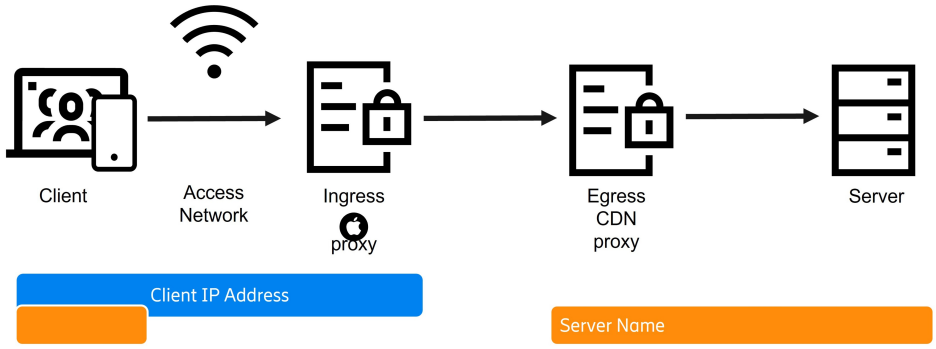

Figure 2: Apple's Private Relay setup - Only the ingress proxy can see the client's IP address and only the egress proxy knows the target server name.

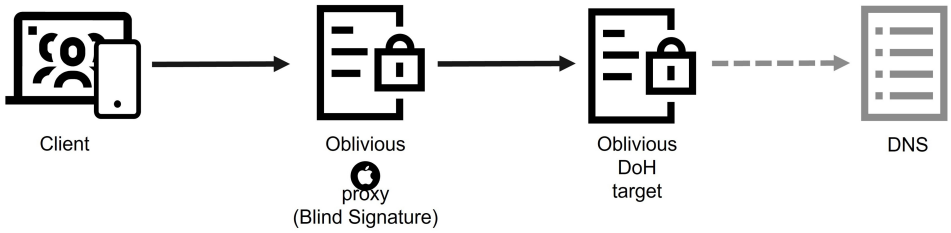

Figure 3: Use of ingress and egress proxies to provide the Oblivious DNS over HTTP (DoH) service

While this approach seems straightforward, it's a substantial change in communication patterns. Instead of sending traffic directly to a service, essentially all user traffic is routed to the same small set of intermediaries, and traffic received at the target server comes from a more limited set of entities as well. This significantly changes how traffic flows and is observed in the networks, as well as for the application service providers. Consequently, deploying these kinds of services on a large scale will have an impact on how we manage our networks.

Sharing the right data with the right entity

Without relay services, traffic effectively broadcasts cleartext information to potentially unknown third parties that may listen on-path, especially when transport information isn't otherwise encrypted. Tunnels between selected relays can empower users to

control the data and metadata that could reveal privacy-sensitive information. An example is usage patterns and details of services accessed that can reveal if a user is at home or not to anybody passively listing on the network path. The challenge is to ensure that the right data, and only the right data, is shared with the right entity. Ideally, users only share their identity with an entity they already have a trusted relationship with, like the Access Network provider or service provider. A setup like this is more complicated than what we have today and may make some of the network management techniques currently deployed more complex, but there are also opportunities for better collaboration between the network and, say, the application at the endpoint.

Explicit trust relations provide the basis for more targeted information exchanges with intermediates. Today many network functions – for example for performance optimization or zero rating – passively listen and try to derive useful data from what is revealed. However, the details of what information is available could change anytime due to the ongoing deployment of encryption techniques, general protocol evolution, and new services like Private Relay, or simply by change of application behavior. Requesting explicit information from a relay service instead, or a relay from the endpoint, provides better guarantees that the information is correct and useful and won't suddenly break if traffic changes due to the deployment of encryption, or new end-to-end services. The latter point might even be the more important one.

The network management techniques deployed at present often rely on information that is exposed by most traffic but without any guarantees that the information is accurate. For example, zero rating today often relies on the Service Name Identifier (SNI) in TLS. However, a simple technique like domain fronting can circumvent and thereby 'cheat' the respective network control functions. In the future, the SNI will likely be encrypted and therefore won't be usable anymore for a passive observer.

Using relays to avoid information ambiguity

Similar challenges exist for techniques such as TCP optimizers that

provide protocol specific in-network performance enhancements for unencrypted traffic. These techniques are used today as they can be deployed by the network without collaboration or coordination with other entities, and as such provide a relatively straight forward approach to improve performance under challenging network conditions. However, these techniques often rely on cleartext information from the protocol (for example, TCP header) or even unencrypted application layer data. This kind of traffic interception or traffic manipulation has caused protocol ossification in the past and therefore can hinder deployment of new protocol features. As the portion of encrypted traffic further increases (HTTPS and QUIC traffic) these techniques are less applicable. Instead, having an explicit collaboration between one or more endpoints and an in-network relay to exchange information explicitly avoids protocol ossification and ambiguity of the information provided.

A further example is parental control, which relies on DNS filtering today and therefore is also one of the functions that service providers indicated as being affected when Private Relay is used. The example perfectly illustrates the main problem of this technique: it breaks easily if, for example, another DNS server is used, or information simply becomes better protected. Using explicit relays instead of a DNS-based solution to provide this function not only makes this service less fault-prone. It also provides an opportunity for improvements by involving the content provider, or a relay that acts on behalf of the content provider, and therefore can provide much better information as input to the content control decision itself.

Turning a challenge into opportunity: empowering collaboration and user privacy

Finally, the cases described above show there's also an opportunity in this change. While adding relays seems technically more complex, business relations can become simpler. Whenever collaboration with a network service provider or content provider is needed, this can also be proxied by the relay hosting provider, and therefore a potentially much smaller set of entities to cooperate

with. Especially this is a chance for mobile network operators to establish new, well defined business relationships, and hopefully more easily.

The deployment of relay-based services will indeed change the communication pattern and traffic flows on the Internet. Instead of direct communication between two parties that is observable by all on-path elements, intermediates will be involved, and only limited information will be distributed to an explicitly selected set of trusted parties. However, given the importance of the Internet, more user control of privacy is long overdue and explicit collaboration between these parties could be the key for better in-network support of new and emerging services. Now is the right time to focus our work on these new communication patterns and make the best use of the emerging technologies for network collaboration!

Learn more

Read more about Ericsson's network security solutions

Find out more about our network services, fit for today's shifting needs

Learn about the benefits of network security automation

Listen: Why 5G is the most secure platform, with our Chief Product Security Officer