

The implementation of an extensible socket API for modern networks.

SRE. Rayhaan Jaufeerally, Prof. Dr. Adrian Perrig, SRE. Brian Trammell

March 2019



The BSD socket API

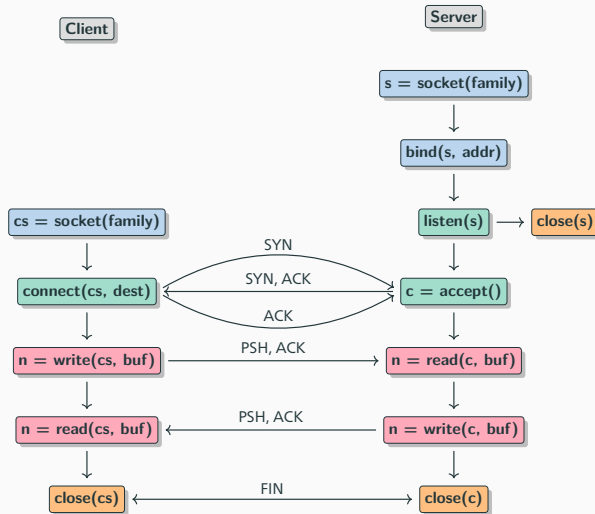


Figure 1: Illustration of an example usage of the BSD sockets API, annotated with TCP protocol semantics.

Problems with the BSD socket API

Simple design simple usage,

Expressing complex ideas is hard:

`setsockopt`s or `ioctl`s are required for complex options,

Such as `TCP_NODELAY` or `SO_REUSEADDR`,

Only one address pair for a connection,

No path selection, network property selection,

Only one uplink.

TAPS was already specified by TAPSWG.

This thesis:

- TAPS-like implementation in the Go programming language,

- Supporting infrastructure (e.g. beaconing service),

- SPAIR6,

- Demo application.

What does the Go TAPS implementation contain?

Support for UDP, TCP, TLS, SCION UDP,

Interface selection using information about local interfaces,

End to end path selection in SCION during dial,

Racing not fully implemented (only helper function for clients to “race” dials).

SPAIR6

Idea for path selection

End user AS's already have more than one uplink, but paths are automagically chosen,

It should not be too hard to expose the path data from the BGP router to the end host,

The end host could signal back to the network which path it wants to use.

The SPAIR6 architecture

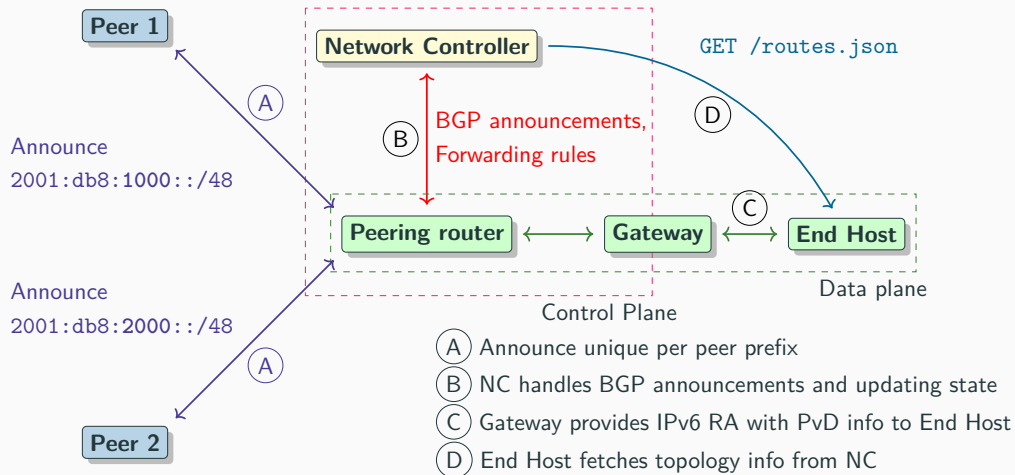


Figure 2: An illustration of the Multiprefixing approach.

A TAPS Implementation

Dialer specifies what properties are required and what the destination is,
Resolution of paths and candidate destinations is handled seamlessly,
Interface exposes a message oriented transport.

Dialer interface

```
1  // Dialer is used to establish a connection
2  type Dialer struct {
3      PropSpec      *PropertySpecification
4      CapProf       CapacityProfile
5
6      RequireDNSSEC bool
7      RequireDoH    bool
8
9      Local  net.Addr
10     Remote net.Addr
11     Creds  *Credentials
12
13     FastOpenReq []byte
14     FastOpenResp []byte
15
16     SCIONPathChooser func([]*sd.PathReplyEntry) *sd.PathReplyEntry
17 }
```

Connection interface

```
// Conn defines an interface all TAPS transports must provide.
type Conn interface {
    net.Conn
    // TransportProperties returns the TAPS transport properties supported.
    TransportProperties() []TransportProperty
    // Dial initializes the connection using the provided dialer.
    Dial(d *Dialer) error
    // SetFramer sets a TAPS framer to be used in Send and Recv.
    SetFramer(f Framer)
    // TAPS specific Send.
    Send(message interface{}, opts []MessageProperty, done chan MsgRef, err chan
        ↪ MsgFail) MsgRef
    // TAPS specific Recv.
    Recv(message chan interface{}, err chan error)
}
```

In this thesis lifetime expiry is checked in the TAPS library before sending.

Other guarantees can be gained from deeper integration with the transport stack, or wider network.

Implementation of Transmission profiles

We implemented transmission profiles by mapping to DSCP codepoints.

Future work would be to integrate with more advanced systems such as COLIBRI.

DNS over HTTPS

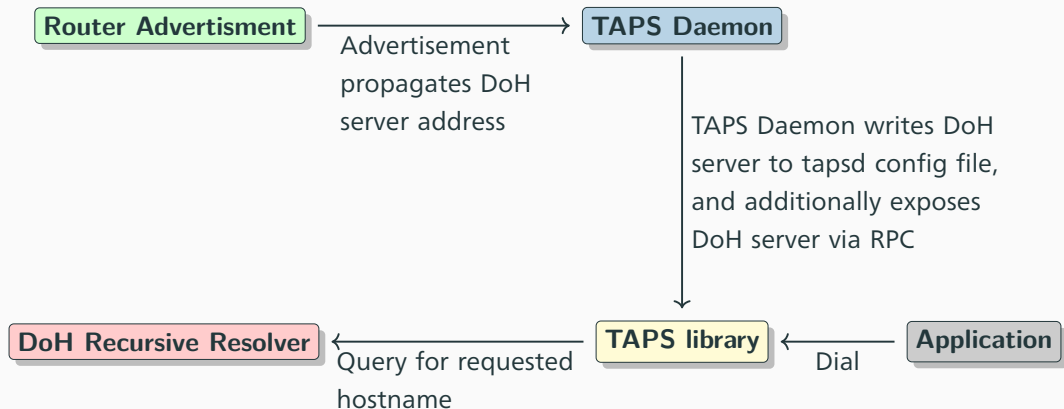


Figure 3: Diagram showing the propagation of DNS over HTTPS configuration information through the system.

Name resolution in TAPS

Name resolution handled completely independently of application

Can take progressive steps to deploy new security protocols without applications noticing,

Supports DNS over HTTPS if network advertises it,

Supports DNSSEC by default,

Supports RAINS lookups for SCION addresses.

Discovering network services

Global system overview

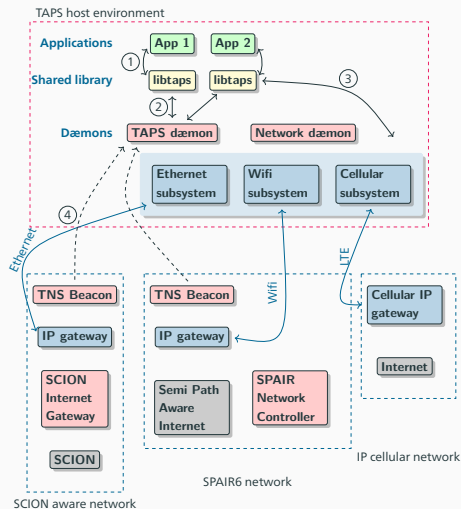


Figure 4: Global system overview

Listens for IPv6 Router Advertisement messages with the Provisioning Domain (PvD) option,

Fetches the URL contained in the PvD option and decodes the JSON document,

Populates a map of interfaces with the contents of the JSON document,

Performs lookups on behalf of applications as to which interface should be used,

Data elements contained in PvD

SCION_ONHOST — Use local SCION resources,

SCION_IPGWW — Use SCION IP Gateway,

SPAIR6_URL — Location of the SPAIR6 route server,

DoH_URL — DNS over HTTPS resolver to use.

TAPS daemon interface

```
1  // InterfaceQueryRequest is used to query which interface to use.
2  type InterfaceQueryRequest struct {
3      // RequireDNSSEC ensures any addresses returned are validated via DNSSEC.
4      RequireDNSSEC bool
5      // RequireDoH ensures queries are only made via DNS over HTTPS to ensure privacy
6      // ↪ and integrity to the resolver.
7      RequireDoH bool
8      Destination string
9      MTU          uint
10     ASPathExact  []uint32
11     ASPathSome   []uint32
12     ASPathExclude []uint32
13     // LatencyOptimized prefers links with low latency.
14     LatencyOptimized bool
15     // BandwidthOptimized prefers links with high bandwidth.
16     BandwidthOptimized bool
17 }
```

Questions?

Thank you!