

Lost in the Edge: Finding Your Way with Signposts

Charalampos Rotsos, Heidi Howard, David Sheets,
Richard Mortier,[†] **Anil Madhavapeddy**, Amir Chaudhry,
Jon Crowcroft

<http://anil.recoil.org/papers/2013-foci-slides.pdf>

University of Cambridge, UK
[†] University of Nottingham, UK
anil@recoil.org

13th August, 2013

Contents

- 1 Introduction
 - Challenge & Constraints
 - Building on DNS
- 2 Signposts
 - Architecture
 - Components
- 3 Conclusions
 - Implications
 - Questions

The Challenge

Centralised cloud-hosted services are convenient but create risks:

- Loss of data and services due to service shutdown (whether for commercial or political reasons)

The Challenge

Centralised cloud-hosted services are convenient but create risks:

- Loss of data and services due to service shutdown (whether for commercial or political reasons)
- Global passive observers recording ~~all~~ 1.6% traffic

The Challenge

Centralised cloud-hosted services are convenient but create risks:

- Loss of data and services due to service shutdown (whether for commercial or political reasons)
- Global passive observers recording ~~all~~ 1.6% traffic
- Inefficient and inconvenient synchronisation in mobile and offline environments

The Challenge

Centralised cloud-hosted services are convenient but create risks:

- Loss of data and services due to service shutdown (whether for commercial or political reasons)
- Global passive observers recording ~~all~~ 1.6% traffic
- Inefficient and inconvenient synchronisation in mobile and offline environments

Our Approach

Use DNS to enable *personal clouds*, making it easy to deploy apps that function securely and efficiently across our own device network, across the Internet edge.

Constraints

Compatibility. Can't require users to change all their apps.

Security. Need to control access to our personal devices:
requires *authentication* and *confidentiality*.

Connectivity. Need to be able to interconnect devices whatever
network is available.

Constraints

Compatibility. Can't require users to change all their apps.

Security. Need to control access to our personal devices:
requires *authentication* and *confidentiality*.


Connectivity. Need to be able to interconnect devices whatever
network is available.

Data vs Orchestration

What's the *minimal* network infrastructure that we can deploy to represent individual users on the core Internet?

Regaining Connectivity


- Network Address Translation (NAT) killed end-to-end IP addressing



- Packet filtering makes tunnel setup dynamic (Full-cone NAT? Is UDP blocked? IPSec?)

Regaining Connectivity


- Network Address Translation (NAT) killed end-to-end IP addressing



- Packet filtering makes tunnel setup dynamic (Full-cone NAT? Is UDP blocked? IPSec?)
- Redirection and proxies (e.g., Wifi hotspots) require traversal

Regaining Connectivity

- Network Address Translation (NAT) killed end-to-end IP addressing



- Packet filtering makes tunnel setup dynamic (Full-cone NAT? Is UDP blocked? IPSec?)
- Redirection and proxies (e.g., Wifi hotspots) require traversal
- Multipath is increasingly available (e.g., 3G + Wifi)

Contents

- 1 Introduction
 - Challenge & Constraints
 - Building on DNS
- 2 Signposts
 - Architecture
 - Components
- 3 Conclusions
 - Implications
 - Questions

DNS

DNS is **THE** Internet naming standard:

- Supported in almost every embedded device.
- Naturally hierarchical and cacheable.
- Flexible and "extensible".
- Resolver infrastructure exists almost everywhere (including censorship).

DNS Today

```
# host recoil.org
recoil.org has address 89.16.177.154
recoil.org mail is handled by 10 dark.recoil.org.
recoil.org mail is handled by 20 mx-caprica.easydns.com.
```

DNS Today

```
# host recoil.org
recoil.org has address 89.16.177.154
recoil.org mail is handled by 10 dark.recoil.org.
recoil.org mail is handled by 20 mx-caprica.easydns.com.
```

Why can't we have stronger DNS bindings between edge devices?

```
# host ipad.home.anil.recoil.org
ipad.home.anil.recoil.org has address 192.168.1.19
```

DNS Manipulation

DNS is **already** manipulated: content networks differentiate results by the query source so the nearest CDN node can serve data

Indeed,

“DNS servers can play games. As long as they appear to deliver a syntactically correct response to every query, they can fiddle the semantics.” — RFC3234

DNS Manipulation

DNS is **already** manipulated: content networks differentiate results by the query source so the nearest CDN node can serve data

Indeed,

“DNS servers can play games. As long as they appear to deliver a syntactically correct response to every query, they can fiddle the semantics.” — RFC3234

Names for The Average Joe

But there's nowhere for **individuals** to easily host their own little name services online. Change this, and everything improves.

DNS Security


Authentication. DNSSEC provides a standard, deployed security model where identity chains are established by trusting the registrars or other trust anchors

Confidentiality. DNSCurve adds confidentiality, repudiability, integrity, and authentication to name resolution through an Elliptic Curve Cryptographic tunnel; can trade compatibility against overhead, with 255-bit Curve25519 keys offering complexity equivalent to 3072-bit RSA

Contents


- 1 Introduction
 - Challenge & Constraints
 - Building on DNS
- 2 Signposts
 - Architecture
 - Components
- 3 Conclusions
 - Implications
 - Questions

Architecture



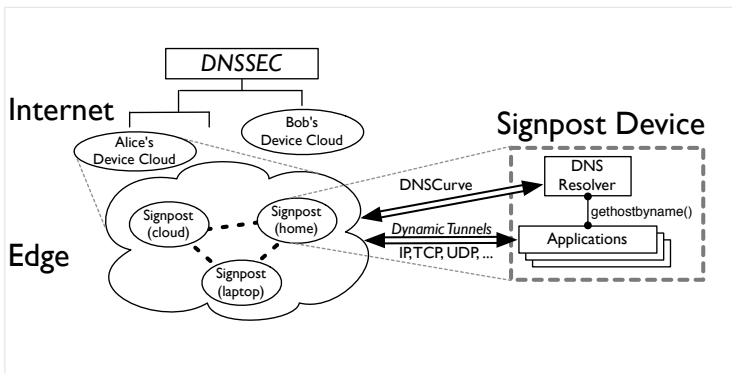
At the edge, devices interconnect using tunnels created in response to authenticated, confidential DNSCurve queries. Connections access-controlled via authenticated query source.

Architecture



At the edge, devices interconnect using tunnels created in response to authenticated, confidential DNSCurve queries. Connections access-controlled via authenticated query source.

Architecture



At the edge, devices interconnect using tunnels created in response to authenticated, confidential DNSCurve queries. Connections access-controlled via authenticated query source.

Contents

- 1 Introduction
 - Challenge & Constraints
 - Building on DNS
- 2 Signposts
 - Architecture
 - Components
- 3 Conclusions
 - Implications
 - Questions

Active Edge Resolution

- Incremental, parallel resolution via 0 TTL responses containing multiple results.

Active Edge Resolution

- Incremental, parallel resolution via 0 TTL responses containing multiple results.
- Bootstrap trusted public keys between devices via resurrecting duckling. No passwords during resolution.

Active Edge Resolution

- Incremental, parallel resolution via 0 TTL responses containing multiple results.
- Bootstrap trusted public keys between devices via resurrecting duckling. No passwords during resolution.
- Degrade gracefully from P2P to personal cloud service to shared provider.

Active Edge Resolution

- Incremental, parallel resolution via 0 TTL responses containing multiple results.
- Bootstrap trusted public keys between devices via resurrecting duckling. No passwords during resolution.
- Degrade gracefully from P2P to personal cloud service to shared provider.
- Resolution triggers tunnel establishment scripts; currently support (**L2**) Tuntap/SSH, OpenVPN, (**L3**) IPSec, (**L4+**) Privoxy/Tor via SOCKS

Active Edge Resolution

- Incremental, parallel resolution via 0 TTL responses containing multiple results.
- Bootstrap trusted public keys between devices via resurrecting duckling. No passwords during resolution.
- Degrade gracefully from P2P to personal cloud service to shared provider.
- Resolution triggers tunnel establishment scripts; currently support (**L2**) Tuntap/SSH, OpenVPN, (**L3**) IPSec, (**L4+**) Privoxy/Tor via SOCKS
- Seamless operation with extra host support (e.g., OpenFlow)

Identity Management

- Automatic, internal key management in a personal trust hierarchy simplifies hygiene.
- TSIG/SIG0 DNSSEC signatures used to demonstrate subnamespace authority.
- Manage keys for SSH, PGP, *Curve in parallel.
- Provides low-friction revocation, making rollover usable by mortals (?)

Programming Model

Currently: Sockets API decouples `getaddrinfo(3)` from `connect(2)`, so less powerful.

With Signposts:

- Applications bind names to flows in one call, separating connection establishment from data transfer,
- Signpost nodes select environmentally optimal routes via long-poll DNSCurve updates
- Signpost resolver proxies DNS on `localhost`, late-binding lookups only when traffic is sent (e.g., TCP SYN)

Work-in-Progress

Resolution. Looking to more efficient path establishment than “try everything at once”

Identity. Automating key derivation & management

Programming. Exploring details, e.g., need to patch OpenSSL, provide local OpenFlow switch; more in *The Case for Reconfigurable I/O Channels*, RESoLVE 2012 (<http://anil.recoil.org/papers/>)

Implementation. May be easier to support applications that use sockets via lightweight VMs (e.g., <http://openmirage.org> with *Message Switch*, <http://github.com/djs55/message-switch>)

Contents

- 1 Introduction
 - Challenge & Constraints
 - Building on DNS
- 2 Signposts
 - Architecture
 - Components
- 3 Conclusions
 - Implications
 - Questions

Alternatives & Possibilities

Signpost uses DNS as a device-facing interface for compatibility – but could support alternative mechanisms for upstream resolution:

- Perspectives (<http://perspectives-project.org/>) offers a P2P trust network
- Namecoin (<http://namecoin.info/>) provides decentralized naming but has economic issues.

When widely deployed, a set of Signposts could help with:

- Tor. Constructing a mix zone, perhaps using *Dustclouds* (<http://anil.recoil.org/papers/2010-iswp-dustclouds.pdf>)
- Dissent (<http://dedis.cs.yale.edu/2010/anon/>), simplifying its use by Average Joe.

Contents

- 1 Introduction
 - Challenge & Constraints
 - Building on DNS
- 2 Signposts
 - Architecture
 - Components
- 3 Conclusions
 - Implications
 - Questions

Thank you!

Questions?

<https://github.com/signposts>

<https://github.com/mirage>