The GNUnet Architecture

We Fix the Net!

Christian Grothoff

Team DÉCENTRALISÉ
Inria Rennes - Bretagne Atlantique

2.10.2014

"Never doubt your ability to change the world." –Glenn Greenwald

Status Quo

- Spy agencies do mass surveillance:
  - Cables, satellites, routers, phones, banking, physical mail, ...
  - Internet service providers (PRISM), cloud storage, ...

Status Quo

- Spy agencies do mass surveillance:
  - Cables, satellites, routers, phones, banking, physical mail, ...
  - Internet service providers (PRISM), cloud storage, ...
- Spy agencies do hacking:
  - Hardware: vendor cooperation, interdiction, saboteurs, ...
  - Software: 0-days (BND buys), ...
  - Networks: man-on-the-side (QUANTUM), ...
  - Standards: Dual-EC, IPSec, SSL, NIST ECC, ...

Status Quo

- Spy agencies do mass surveillance:
  - Cables, satellites, routers, phones, banking, physical mail, ...
  - Internet service providers (PRISM), cloud storage, ...
- Spy agencies do hacking:
  - Hardware: vendor cooperation, interdiction, saboteurs, ...
  - Software: 0-days (BND buys), ...
  - Networks: man-on-the-side (QUANTUM), ...
  - Standards: Dual-EC, IPSec, SSL, NIST ECC, ...
- Spy agencies do take control:
  - Influence trade negotations (hack EU, NGOs, etc.)
  - Sabotage UN climate conference negotations
  - “We kill people based on meta data.”

How can we secure networks to avoid totalitarianism?

The Internet is Fundamentally Broken

- Network generally learns too much: **no cleartext**
- Insecure defaults and system complexity
- Key, centralised Internet infrastructure is easily controlled:
  - Number resources (IANA)
  - Domain Name System (Root zone)
  - X.509 CAs (HTTPS certificates)
  - Dominant network service providers (Faceboogle)
- Encryption does not help if PKI is compromised, or plaintext is in the Cloud!

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`
- NS of **tum.de** is `dns1.lrz.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`
- NS of **tum.de** is `dns1.lrz.de`
- NS of `lrz.de` is `dns1.lrz.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`
- NS of **tum.de** is `dns1.lrz.de`
- NS of `lrz.de` is `dns1.lrz.de`
- NS of **in.tum.de** is `tuminfo1.informatik.tu-muenchen.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`
- NS of **tum.de** is `dns1.lrz.de`
- NS of `lrz.de` is `dns1.lrz.de`
- NS of **in.tum.de** is `tuminfo1.informatik.tu-muenchen.de`
- NS of `tu-muenchen.de` is `ws-han1.wip-ip.dfn.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`
- NS of **tum.de** is `dns1.lrz.de`
- NS of `lrz.de` is `dns1.lrz.de`
- NS of **in.tum.de** is `tuminfo1.informatik.tu-muenchen.de`
- NS of `tu-muenchen.de` is `ws-han1.wip-ip.dfn.de`
- NS of `dfn.de` is `ws-han1.wip-ip.dfn.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`
- NS of **tum.de** is `dns1.lrz.de`
- NS of `lrz.de` is `dns1.lrz.de`
- NS of **in.tum.de** is `tuminfo1.informatik.tu-muenchen.de`
- NS of `tu-muenchen.de` is `ws-han1.wip-ip.dfn.de`
- NS of `dfn.de` is `ws-han1.wip-ip.dfn.de`
- NS of **net.in.tum.de** is `dns1.lrz.de`

How broken is the Internet? A DNS case study.

What would a simple DNS lookup do? Say for `taler.net`?

- NS of **net** is `a.gtld-servers.net`
- NS of **taler.net** is `dns1.name-services.com`
- NS of `com` is `a.gtld-servers.net`
- CNAME of **taler.net** is **pixel.net.in.tum.de**
- NS of **de** is `n.de.net`
- NS of `de.net` is `ns1.denic.de`
- NS of `denic.de` is `ns1.denic.de`
- NS of **tum.de** is `dns1.lrz.de`
- NS of `lrz.de` is `dns1.lrz.de`
- NS of **in.tum.de** is `tuminfo1.informatik.tu-muenchen.de`
- NS of `tu-muenchen.de` is `ws-han1.wip-ip.dfn.de`
- NS of `dfn.de` is `ws-han1.wip-ip.dfn.de`
- NS of **net.in.tum.de** is `dns1.lrz.de`
- A of **pixel.net.in.tum.de** is 131.159.20.32

How broken is the Internet? A DNS case study.

- Glue records and caching logic were not shown
- As deployed, DNSSEC fails on end-to-end authenticity and confidentiality
- DNS remains major source of traffic amplification attacks
- Some US court considered confiscating ccTLDs
- Censorship of non-TLD domain names is already common

Example #2: The IPv4 header (Sept. 1981)

<table>
<tr><td>Version</td><td>HDL</td><td>ToS</td><td colspan="2">Length</td></tr>
<tr><td colspan="3">Identification</td><td>Flags</td><td>Fragment offset</td></tr>
<tr><td colspan="2">TTL</td><td>T. Protocol</td><td colspan="2">Checksum</td></tr>
<tr><td colspan="5">Source IP address</td></tr>
<tr><td colspan="5">Destination IP address</td></tr>
<tr><td colspan="5">Options (optional)</td></tr>
<tr><td colspan="5">Data (Length–HDL bytes)</td></tr>
</table>

How broken is the Internet? Thoughts about IP

Some known issues with IP:

- Cannot prove IP address ownership (BGP hijacking, IP spoofing)
- Routers learn source address (meta data leakage)
- Routers learn payload (information leakage)
- Packet size typically too small for modern networks (inefficient)
- Packet size leaks information
- No congestion control $\Rightarrow$ DOS
- Much legacy baggage (fragmentation, ToS, options)
- IP? Really: IPv4, IPv6, NAT, 4in6, 6in4, 6over4, 6to4, NAT64, NAT66, Teredo, DS-Lite, NAT-PT, NAPT-PT, 4rd, 6rd, ...

How broken is the Internet? Thoughts about IP

Some known issues with IP:

- Cannot prove IP address ownership (BGP hijacking, IP spoofing)
- Routers learn source address (meta data leakage)
- Routers learn payload (information leakage)
- Packet size typically too small for modern networks (inefficient)
- Packet size leaks information
- No congestion control $\Rightarrow$ DOS
- Much legacy baggage (fragmentation, ToS, options)
- IP? Really: IPv4, IPv6, NAT, 4in6, 6in4, 6over4, 6to4, NAT64, NAT66, Teredo, DS-Lite, NAT-PT, NAPT-PT, 4rd, 6rd, ...

If IP was well-designed, network neutrality would not be debated.

Ideal packet (long-term vision)

32 byte destination $D = dG$ (ECC Point)
32 byte ephemeral key $S = sG$ (ECC Point)
$2^{16} - 128$ byte encrypted payload ($K = ECDHE(d, S)$)
64 byte HMAC

Once packets look like this, routers have no choice but to be neutral.

Migration strategy

- Physical infrastructure (routers, switches) will migrate last
- Need to rethink not just TCP/IP, but also client-server (PRISM!)
- Each user must be in control of his computation and data
- Interaction and cooperation must not use "trusted" third-party facilitators
- Need to build *decentralised* applications

Migration strategy

- Physical infrastructure (routers, switches) will migrate last
- Need to rethink not just TCP/IP, but also client-server (PRISM!)
- Each user must be in control of his computation and data
- Interaction and cooperation must not use "trusted" third-party facilitators
- Need to build *decentralised* applications

$\Rightarrow$ Rearchitect higher layers and applications first!

$\Rightarrow$ Deploy as *overlay* network

TCP/IP *below* is baggage we need to support "merely" for transition.

The GNUnet Vision (Simplified)

Internet

Faceboogle
DNS/X.509
TCP/UDP
IP/BGP
Ethernet
Phys. Layer

The GNUnet Vision (Simplified)

Internet	
Faceboogle	
DNS/X.509	
TCP/UDP	
IP/BGP	
Ethernet	
Phys. Layer	HTTPS/TCP/WLAN/...

The GNUnet Vision (Simplified)

Internet	
Faceboogle	
DNS/X.509	
TCP/UDP	
IP/BGP	
Ethernet	CORE (OTR)
Phys. Layer	HTTPS/TCP/WLAN/...

The GNUnet Vision (Simplified)

Internet

Faceboogle
DNS/X.509
TCP/UDP
IP/BGP
Ethernet
Phys. Layer

R^5N DHT (KBR)
CORE (OTR)
HTTPS/TCP/WLAN/...

The GNUnet Vision (Simplified)

Internet

Faceboogle
DNS/X.509
TCP/UDP
IP/BGP
Ethernet
Phys. Layer

CADET (SCTP+Axolotl)
R^5N DHT (KBR)
CORE (OTR)
HTTPS/TCP/WLAN/...

The GNUnet Vision (Simplified)

Internet	
Faceboogle	
DNS/X.509	GNU Name System
TCP/UDP	CADET (SCTP+Axolotl)
IP/BGP	R^5N DHT (KBR)
Ethernet	CORE (OTR)
Phys. Layer	HTTPS/TCP/WLAN/...

The GNUnet Vision (Simplified)

Internet

Faceboogle
DNS/X.509
TCP/UDP
IP/BGP
Ethernet
Phys. Layer

Applications
GNU Name System
CADET (SCTP+Axolotl)
R^5N DHT (KBR)
CORE (OTR)
HTTPS/TCP/WLAN/...

The GNUnet Vision (Simplified)

Internet

Faceboogle
DNS/X.509
TCP/UDP
IP/BGP
Ethernet
Phys. Layer

GNUnet

Applications
GNU Name System
CADET (SCTP+Axolotl)
R^5N DHT (KBR)
CORE (OTR)
HTTPS/TCP/WLAN/...

Fixing the Net: Building Blocks

- CORE: encrypted, off-the-record messaging between adjacent peers
- R^5N DHT: decentralised, censorship-resistant key-value store, also enables key-based routing (KBR) and route discovery
- GNU Name System: decentralised PKI, identity management and name system
- CADET: Confidential Ad-hoc Decentralised End-to-End Transport

Fixing the Net: Building Blocks

- CORE: encrypted, off-the-record messaging between adjacent peers
- R^5N DHT: decentralised, censorship-resistant key-value store, also enables key-based routing (KBR) and route discovery
- GNU Name System: decentralised PKI, identity management and name system
- CADET: Confidential Ad-hoc Decentralised End-to-End Transport
- Secure decentralised network size estimation
- Secure decentralised key revocation
- Efficient pair-wise set union (Eppstein) and set intersection (Bloom)
- Advanced cryptography:
  - Secure multiparty scalar product
  - Byzantine fault-tolerant consensus (set union)
  - Fouque's distributed key generation and cooperative encryption
  - Cramer-style electronic voting

Software architecture: overview

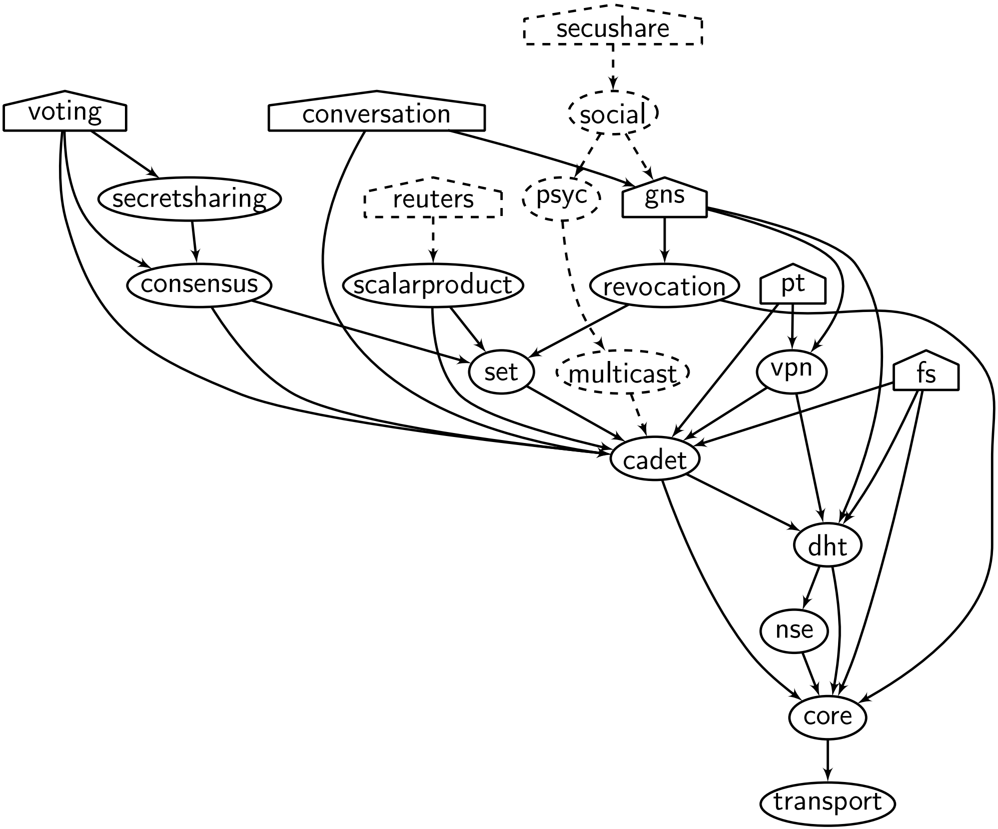

Fixing the Net: Applications

- Anonymous file-sharing
- IP-over-GNUnet
- Voice-over-GNUnet
- Decentralised social networking (future)
- Decentralised cooperative news distribution (future)
- Privacy-preserving constraint negotiation (future)
- Taler: Taxable Anonymous Libre Electronic Reserves (future)

Network Architecture: Egyptian Edition

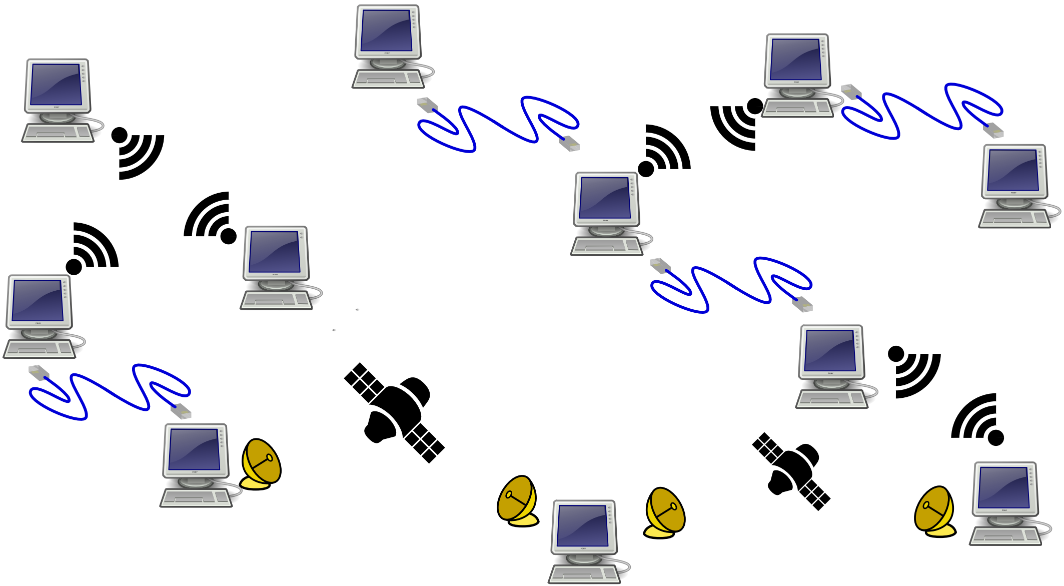

Network Architecture: With Infrastructure

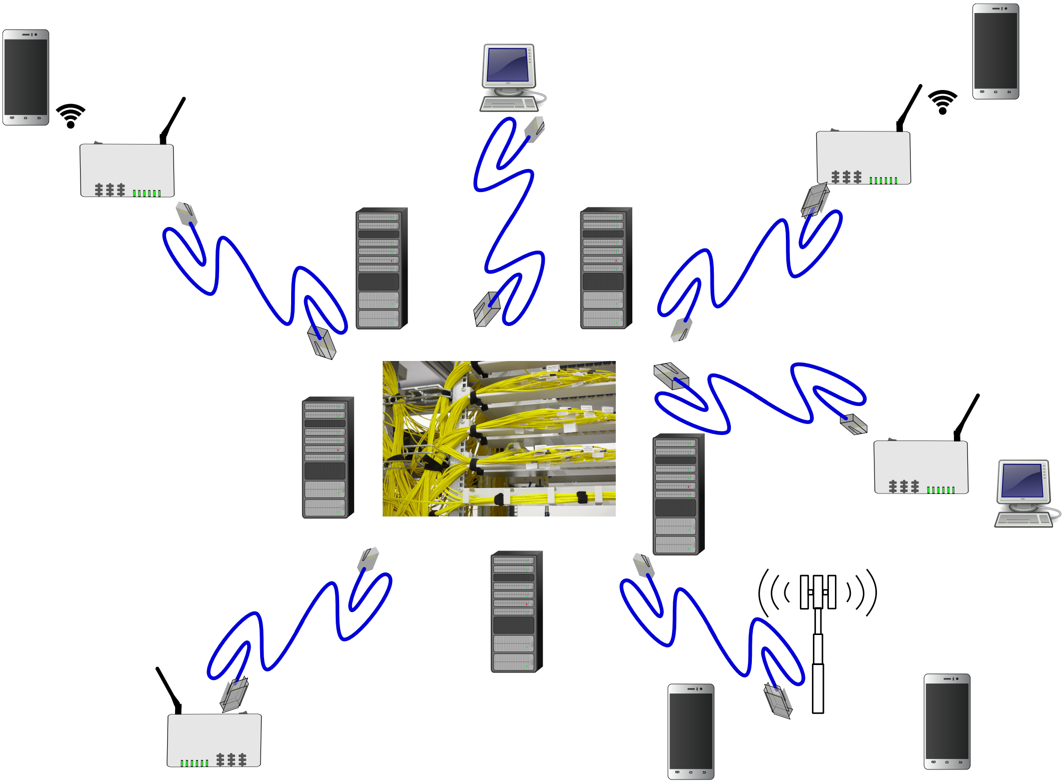

GNUnet and performance

- Cryptography and bandwidth overheads are for most applications irrelevant
- For IP-replacement, some investment in cryptographic hardware may be warranted
  $\Rightarrow$ opportunity for Europe to become technical leader
- Routing currently scales with $O(\sqrt{n} \log n)$
  $\Rightarrow$ more research warranted, but may suffice already
- Decentralised administration scales with $O(n)$ vs. $O(1)$ for centralised
  $\Rightarrow$ usability is critical, more development needed
- Education maybe even harder:
  How could users distinguish secure systems from insecure systems?

System cost

Short-term overlay:

- Software: 1–5 M€ and 2–5 years to achieve usability
- NAT: ratios of 1:2 users at $\approx$ 50 € COTS
- DHT: ratios of 1:1000 to 1:10000 users at $\approx$ 3,000 € COTS

Long-term full infrastructure migration:

- Router: tens of millions of € to develop:
  high-speed router at 10 GBit/s needs to do 20,000 DH public key operations/s;
  - Xeon E3 takes $\approx$ 150,000 cycles/op
  - Cortex-A9 takes $\approx$ 580,000 cycles/op
    $\Rightarrow$ router needs custom ASIC

  $\Rightarrow$ Final costs then likely comparable to modern routers
- But: networks include way more than high-speed routers (3G, Satellite, ...)

Overlay networks as "parallel universes"

- Can deploy many overlay network designs in parallel
- Co-exist with existing Internet using same hardware
- May be effected to some degree by security issues in underlay (availability, performance, DoS, connectivity, censorship, surveillance)
- Overlay networks typically operate globally, hard to constrain by region

Overlays do not change jurisdiction issues!

Thoughts on jurisdiction

- Few modern governments follow or enforce existing laws:
  - Prohibition of torture
  - Geneva Convention
  - Human rights (privacy, surveillance, asylum, food, shelter)
  - Due process
  - Anti-corruption, taxation, freedom of information

Thoughts on jurisdiction

- Few modern governments follow or enforce existing laws:
  - Prohibition of torture
  - Geneva Convention
  - Human rights (privacy, surveillance, asylum, food, shelter)
  - Due process
  - Anti-corruption, taxation, freedom of information
- "Die Zeit" online recently titled that German government proposes to break fundamental constitutional principles without good reason ("ineffective" measure)

Thoughts on jurisdiction

- Few modern governments follow or enforce existing laws:
  - Prohibition of torture
  - Geneva Convention
  - Human rights (privacy, surveillance, asylum, food, shelter)
  - Due process
  - Anti-corruption, taxation, freedom of information
- "Die Zeit" online recently titled that German government proposes to break fundamental constitutional principles without good reason ("ineffective" measure)

  $\Rightarrow$ Neither constitutions nor ordinary laws constrain the corpocracy.

Thoughts on jurisdiction

- Few modern governments follow or enforce existing laws:
  - Prohibition of torture
  - Geneva Convention
  - Human rights (privacy, surveillance, asylum, food, shelter)
  - Due process
  - Anti-corruption, taxation, freedom of information
- "Die Zeit" online recently titled that German government proposes to break fundamental constitutional principles without good reason ("ineffective" measure)

  $\Rightarrow$ Neither constitutions nor ordinary laws constrain the corpocracy.

But: physical laws do constrain corpocracy!

Code is law

- Client-server: master-slave

Code is law

- Client-server: master-slave
- TCP/IP: mass surveillance

Code is law

- Client-server: master-slave
- TCP/IP: mass surveillance
- Peer-to-peer: anarchy

Code is law

- Client-server: master-slave
- TCP/IP: mass surveillance
- Peer-to-peer: anarchy
- Tor: privacy as an option

Code is law

- Client-server: master-slave
- TCP/IP: mass surveillance
- Peer-to-peer: anarchy
- Tor: privacy as an option
- GNUnet: privacy by default

You will obey the code. Let's make it work for you (and that means GNU).

Is GNUnet a “darknet”?

- From the point-of-view of mass surveillance, hopefully yes

Is GNUnet a “darknet”?

- From the point-of-view of mass surveillance, hopefully yes
- For users and liberal society, it should be more like a shield

Is GNUnet a “darknet”?

- From the point-of-view of mass surveillance, hopefully yes
- For users and liberal society, it should be more like a shield
- For criminals, they should gain nothing (and cybercriminals should loose)

Is GNUnet a “darknet”?

- From the point-of-view of mass surveillance, hopefully yes
- For users and liberal society, it should be more like a shield
- For criminals, they should gain nothing (and cybercriminals should loose)
- For the totalitarian state, it enables liberal anarchist terrorism.

What about Legal Intercept?

- We must not compromise design or protocols
- We must not enable intercept in the network
- Traditional methods will continue to work:
  - Bug the environment (rooms, cars, etc.)
  - Take physical control of end-systems to install malware or compromise hardware
  - This will not scale, but neither would courts if they actually exercised oversight

We must not enable mass surveillance.
It must be *costly* and *dangerous* to intercept.

Conclusion

- Exist plenty of ideas for building more secure networks
- Need to **do** systems programming and software engineering to make them real
- Full migration will take **decades**
- Can validate and begin to deploy using overlay techniques

"A society that gets rid of all its troublemakers goes downhill." –Robert A. Heinlein

Do you have any questions?

References:

- Christian Grothoff, Bart Polot and Carlo von Loesch. *The Internet is Broken: Idealistic Ideas for Building a GNU Network*. **W3C/IAB Workshop on Strengthening the Internet Against Pervasive Monitoring (STRINT)**, 2014.
- Matthias Wachs, Martin Schanzenbach and Christian Grothoff. *A Censorship-Resistant, Privacy-Enhancing and Fully Decentralized Name System*. **13th International Conference on Cryptology and Network Security (CANS)**, 2014
- Bart Polot and Christian Grothoff. *CADET: Confidential Ad-hoc Decentralized End-to-End Transport*. **13th IEEE IFIP Annual Mediterranean Ad Hoc Networking Workshop (MedHocNet)**, 2014.
- Julian Kirsch, Christian Grothoff, Monika Ermert, Jacob Appelbaum, Laura Poitras and Henrik Moltke. *NSA/GCHQ: Das HACIENDA-Programm zur Kolonisierung des Internet*. In **Heise Online** 8'2014. Heise Zeitschriften Verlag, 2014.
- Judith Horchert, Christian Grothoff, Christian Stöcker. *NSA-System Treasuremap: "Jedes Gerät, überall, jederzeit"*. In **Spiegel Online Netzwelt** 9'2014. Spiegel-Verlag, 2014.

Let's BUILD A GNU ONE