

# The GNU Name System

Christian Grothoff

Technische Universität München

30C3

# Thanks

- ▶ Matthias Wachs, Martin Schanzenbach
- ▶ Werner Koch, Florian Dold, Bart Polot, Sree Harsha Totakura, Simon Dieterle, Andreas Fuchs, Christian Fuchs, Stephan Posselt, Nils Durner, LRN, Ralph Holz, Gabor Toth
- ▶ Kenneth Almquist, Jacob Appelbaum, Daniel Bernstein, Ludovic Courtès, Krista Grothoff, Tanja Lange, Luke Leighton, Simon Josefsson, Nikos Mavrogiannopoulos, Ondrej Mikle, Stefan Monnier, Niels Möller, Chris Palmer, Martin Pool, Richard Stallman, Neal Walfield, Zooko Wilcox-O'Hearn
- ▶ This work was funded by the Deutsche Forschungsgemeinschaft (DFG) under ENP GR 3688/1-1.

# Where We Are



Source: esmont



Source: gawand.org



# Where We Are



## الموقع محظور

أسف إن الموقع الذي أردت تصفحه قد أُحجب وذلك بسبب إحتوائه على نشاط مخالف للقيم الاجتماعية أو السياسية أو الثقافية أو الدينية لخدمة الإمارات العربية المتحدة.

في حالة أردت فتح موقع قد أُحجب، الرجاء قم بتسمية استشارة الملاحظات الموجودة على موقعنا.

We apologize the site you are attempting to visit has been blocked due to its content being inconsistent with the religious, cultural, political and moral values of the United Arab Emirates.

If you think this site should not be blocked, please visit the [Feedback Form](#) available on our website.



**SITE BLOCKED** Source: wikileaks.org



# Adversary Model

- ▶ Any role
- ▶ Multiple Identities
- ▶ Computational Power
- ▶ Legal Power

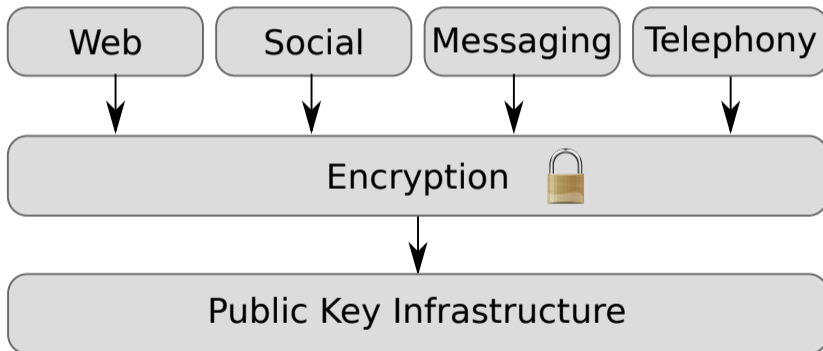


But cannot:

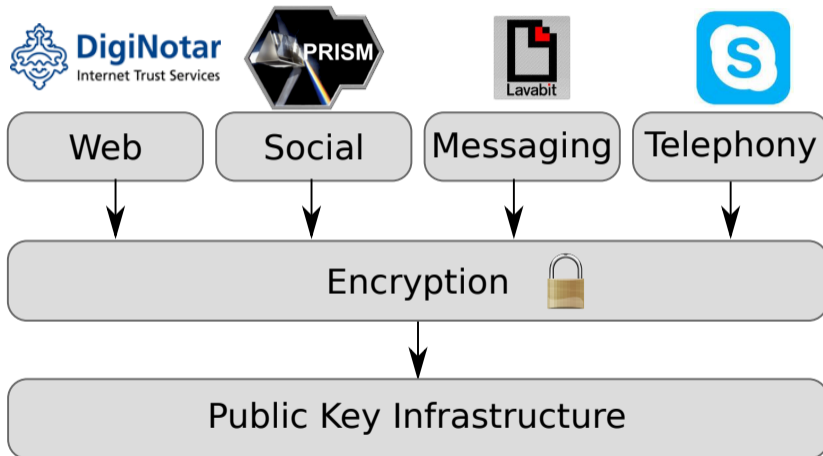
- ▶ Break or prevent crypto
- ▶ Compromise end-user system
- ▶ Prevent network communication



# Broken Pillars



# Broken Pillars



# We MUST Decentralize!

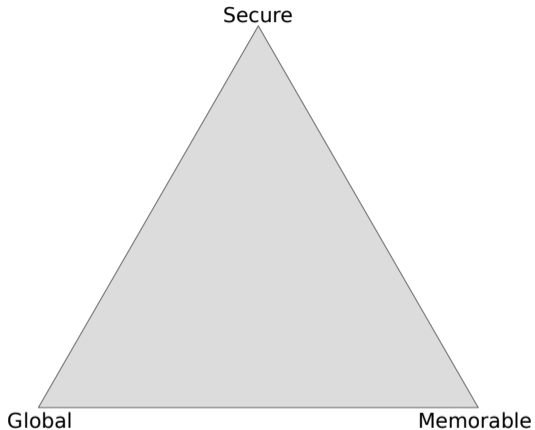
Centralized Internet infrastructure is easily controlled:

- ▶ Number resources (IANA)
- ▶ Domain Name System (Root zone)
- ▶ DNSSEC root certificate
- ▶ X.509 CAs (HTTPS certificates)
- ▶ Major browser vendors (CA root stores!)

Encryption does not help if PKI is compromised!




# Zooko's Triangle




A name system can only fulfill **two!**

# Zooko's Triangle



DNS, “.onion” IDs and `/etc/hosts/` are representative designs.

# Zooko's Triangle



DNSSEC security is broken by design (adversary model!)

# Namecoin

# Namecoin

- ▶ Memorable:

# Namecoin

- ▶ Memorable: Check
- ▶ Global:

# Namecoin

- ▶ Memorable: Check
- ▶ Global: Check
- ▶ Secure:

# Namecoin

- ▶ Memorable: Check
- ▶ Global: Check
- ▶ Secure: different adversary model!



# Namecoin

- ▶ Memorable: Check
  - ▶ Global: Check
  - ▶ Secure: different adversary model!
- ⇒ Availability of names (registration rate) is restricted

# Namecoin

- ▶ Memorable: Check
  - ▶ Global: Check
  - ▶ Secure: different adversary model!
- ⇒ Availability of names (registration rate) is restricted
- ⇒ Adversary must not have 51% compute power

# The GNU Name System

## Properties of GNS

- ▶ Decentralized name system with secure memorable names
- ▶ Delegation used to achieve transitivity
- ▶ Also supports globally unique, secure identifiers
- ▶ Achieves query and response privacy
- ▶ Provides alternative public key infrastructure
- ▶ Interoperable with DNS


## Uses for GNS in GNUnet

- ▶ Identify IP services hosted in the P2P network
- ▶ Identities in social networking applications

# Zone Management: like in DNS



gnunet-setup

General Network Transports File Sharing Namestore **GNS**

**Editing zone API5QDP7A126P06VV60535PDT50B9L12NK6QP64IE8KNC6E807G0** 

Preferred zone name (PSEU):


Master Zone  Private Zone  Shorten Zone


Name	Type	Value	Expiration	Public
<new name>				
+ +	<new record>			
	MX	5,mail.+	end of time	<input checked="" type="checkbox"/>
priv	<new record>			
	PKEY	3IQT1G601GUBVOS5C0JO87OEFB8N3DBJQ4L9SBI8PFLR8UKCVGHG	end of time	<input type="checkbox"/>
heise	<new record>			
	LEHO	heise.de	end of time	<input checked="" type="checkbox"/>
	AAAA	2a02:2e0:3fe:100::8	end of time	<input checked="" type="checkbox"/>
	A	193.99.144.80	end of time	<input checked="" type="checkbox"/>
home	<new record>			
大学	<new record>			
short	<new record>			
mail	<new record>			
homepage	<new record>			
fdfs	<new record>			
www	<new record>			

[Welcome to gnunet-setup.](#)


# Name resolution in GNS




Bob



Bob's webserver





**Bob Builder, Ph.D.**


**Address: Country, Street Name 23**

**Phone: 555-12345**


**Mobile: 666-54321**

**Mail: bob@H2R84L4JIL3G5C.zkey**

# Delegation






Alice




- ▶ Alice learns Bob's public key
- ▶ Alice creates delegation to zone **bob**
- ▶ Alice can reach Bob's webserver via **www.bob.gnu**

# Name Resolution










# Name Resolution




# Name Resolution




# Name Resolution




# Name Resolution



# Name Resolution



# Name Resolution




# GNS as PKI (via DANE/TLSA)

The screenshot shows a web browser window with the address bar displaying <https://freedom.gnu>. A security warning dialog box is open, titled "freedom.gnu" and "Identity verified". The dialog has two tabs: "Permissions" and "Connection". The "Connection" tab is active, showing a green padlock icon and the text: "The identity of this website has been verified by GNS CA." Below this, there is a link for "Certificate Information". Another green padlock icon indicates: "Your connection to freedom.gnu is encrypted with 256-bit encryption." Below this, it states: "The connection uses TLS 1.2." and "The connection is encrypted using AES\_256\_CBC, with SHA1 for message authentication and ECDHE\_RSA as the key exchange mechanism." There is also a "Site information" section with an information icon and the text: "You have never visited this site before today." At the bottom of the dialog is a link: "What do these mean?". The background of the browser shows the GNU Operating System website. The top navigation bar is red and contains the text "Operating System". Below this, there are links for "español [es]", "فارسی [fa]", "français [fr]", "hrvatski [hr]", and "italiano [it]". A secondary navigation bar is white with a red background and contains links for "Why", "Licenses", "Education", "Software", "Documentation", and "Help". The main content area has a blue header with the text "What is GNU?". Below this, there is a paragraph of text: "operating system that is [free software](#)—it respects your freedom. [of GNU](#) (more precisely, GNU/Linux systems) which are [What we provide](#)." Below the text is a screenshot of a presentation slide titled "What is free software?" and "The Free Software Definition". The slide features a cartoon drawing of a ram's head and a pie chart.

The [GNU Project](#) was launched in 1984 to develop the GNU system. The name "GNU" is a recursive acronym for "GNU's Not Unix!". "[GNU](#)" is pronounced *g'noo*, as one syllable, like saying "grew" but replacing the *r* with *n*.

A Unix-like operating system is a [software collection](#) of applications, libraries, and

# Security Issue: DHT





## Query Privacy: Terminology

$G$  generator in ECC curve, a point

$n$  size of ECC group,  $n := |G|$ ,  $n$  prime

$x$  private ECC key of zone ( $\in \mathbb{Z}_n$ )

$P$  public key of zone, a point  $P := xG$

$l$  label for record in a zone ( $\in \mathbb{Z}_n$ )

$R_{P,l}$  set of records for label  $l$  in zone  $P$

$q_{P,l}$  query hash (hash code for DHT lookup)

$B_{P,l}$  block with information for label  $l$  in zone  $P$  published in the DHT under  $q_{P,l}$

# Query Privacy: Cryptography

Publishing  $B$  under  $q_{P,I} := H(dG)$

$$h := H(I, P) \tag{1}$$

$$d := h \cdot x \pmod n \tag{2}$$

$$B_{P,I} := S_d(E_{HKDF(I,P)}(R_{P,I})), dG \tag{3}$$

# Query Privacy: Cryptography

Publishing  $B$  under  $q_{P,I} := H(dG)$

$$h := H(I, P) \tag{1}$$

$$d := h \cdot x \pmod n \tag{2}$$

$$B_{P,I} := S_d(E_{HKDF(I,P)}(R_{P,I})), dG \tag{3}$$

Searching for  $I$  in zone  $P$

$$h = H(I, P) \tag{4}$$

$$q_{P,I} = H(dG) = H(hxG) = H(hP) \Rightarrow \text{obtain } B_{P,I} \tag{5}$$

$$R_{P,I} = D_{HKDF(I,P)}(B_{P,I}) \tag{6}$$

# Revocation

## Revocation Basics

- ▶ Revocation certificate (RC): message signed with private key
  - ▶ Peer receives new valid RC, floods to all neighbours
  - ▶ All peers store all valid RCs forever
- ⇒ Expensive operation ⇒ proof-of-work

# Revocation

## Revocation Basics

- ▶ Revocation certificate (RC): message signed with private key
  - ▶ Peer receives new valid RC, floods to all neighbours
  - ▶ All peers store all valid RCs forever
- ⇒ Expensive operation ⇒ proof-of-work

## Revocation Magic

- ▶ Peers maybe offline during initial flood
  - ▶ Network might be temporarily partitioned
- ⇒ Need to reconcile revocation sets on connect

Whenever two peers establish a P2P connection, they must compute the set union of their RC sets!

# Efficient Set Union

(based on “What’s the difference? Efficient Set Reconciliation without Prior Context”, Eppstein et al., SIGCOMM’11)

- ▶ Alice and Bob have sets  $A$  and  $B$
- ▶ The sets are very large
- ▶ ...but their symmetric difference  $\delta = |(A - B) \cup (B - A)|$  is small
- ▶ Now Alice wants to know  $B - A$  (the elements she’s missing)
- ▶ ...and Bob  $A - B$  (the elements he’s missing)
- ▶ How can Alice and Bob do this efficiently?
  - ▶ w.r.t. communication and computation

## Bad Solution

- ▶ Naive approach: Alice sends  $A$  to Bob, Bob sends  $B - A$  back to Alice
- ▶ ... and vice versa.
  
- ▶ Communication cost:  $O(|A| + |B|)$  :(
- ▶ Ideally, we want to do it in  $O(\delta)$ .
- ▶ First improvement: Don't send elements of  $A$  and  $B$ , but send/request hashes. Still does not improve complexity :(
  
- ▶ We need some more fancy data structure!

# Bloom Filters

**Constant size** data structure that “summarizes” a set.

Operations:


$d = \text{NewBF}(\text{size})$  Create a new, empty bloom filter.

$\text{Insert}(d, e)$  Insert element  $e$  into the BF  $d$ .

$b = \text{Contains}(d, e)$  Check if BF  $d$  contains element  $e$ .  
 $b \in \{\text{“Definitely not in set”}, \text{“Probably in set”}\}$




## BF: Insert




$$H(\text{Element \#1}) = (2, 3, 7)$$

## BF: Insert



$$H(\text{Element \#1}) = (2, 3, 7)$$


## BF: Insert



$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H(\text{Element \#2}) = (1, 3, 5)$$


## BF: Insert



$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H(\text{Element \#2}) = (1, 3, 5)$$


## BF: Membership Test



$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H(\text{Element \#2}) = (1, 3, 5)$$

## BF: Membership Test (false positive)



$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H(\text{Element \#2}) = (1, 3, 5)$$

## Counting Bloom Filters

BF where buckets hold a **positive integer**.

Additional Operation:

*Remove*( $d, e$ ) Remove element from the CBF  $d$ .

⇒ False negatives when removing a non-existing element.

# Invertible Bloom Filters

Similar to CBF, but

- ▶ Allow **negative counts**
- ▶ Additionally store **(XOR-)sum of hashes** in buckets.


Additional Operations:

$(e, r) = \text{Extract}(d)$  Extract an element ( $e$ ) from the IBF  $d$ , with result code  $r \in \{\text{left}, \text{right}, \text{done}, \text{fail}\}$

$d' = \text{SymDiff}(d_1, d_2)$  Create an IBF that represents the symmetric difference of  $d_1$  and  $d_2$ .




## IBF: Insert



$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H'(\text{Element \#1}) = 4242$$


## IBF: Insert



$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H'(\text{Element \#1}) = 4242$$

## IBF: Insert




$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H'(\text{Element \#1}) = 4242$$

$$H(\text{Element \#2}) = (1, 3, 5)$$

$$H'(\text{Element \#2}) = 0101$$

## IBF: Insert



$$H(\text{Element \#1}) = (2, 3, 7)$$

$$H'(\text{Element \#1}) = 4242$$

$$H(\text{Element \#2}) = (1, 3, 5)$$

$$H'(\text{Element \#2}) = 0101$$

## IBF: Extract

1	0101	pure bucket
1	4242	
2	4343	
0	0000	
1	0101	
0	0000	
1	4242	

- ▶ Pure bucket  $\Rightarrow$  extractable element hash
- ▶ Extraction  $\Rightarrow$  more pure buckets (hopefully/probably)
- ▶ Less elements  $\Rightarrow$  more chance for pure buckets

# Symmetric Difference

We can directly compute the symmetric difference without extraction.

- ▶ Subtract counts
- ▶ XOR hashes

# The Set Union Protocol


1. Create IBFs
  2. Compute SymDiff
  3. Extract element hashes
- 
- ▶ Amount of communication and computation only depends on  $\delta$ , not  $|A| + |B|$   
:)
  - ▶ How do we choose the initial size of the IBF?
  - ▶  $\Rightarrow$  Do difference estimation first!

# Difference Estimation


- ▶ Needed: Estimator accurate for *small* distances
- ▶ Turns out we can re-use IBFs for difference estimation
- ▶ *Sample* the set by looking at hashes, create multiple IBFs




# Strata Estimator




# Strata Estimator




# Strata Estimator



# Strata Estimator



# Estimation




Estimate as  $(3 + 7) \cdot 2^2$ .

(Number of extracted hashes scaled by expected number of elements in the remaining IBFs)

## The “.zkey” pTLD

- ▶ “LABELS.PKEY.zkey” format
  - ▶ PKEY is the public key of the zone
  - ▶ Works a bit like “.onion”
- ⇒ Globally unique identifiers!



A business card for Bob Builder, Ph.D. The card features the TUM logo in the top left corner and a circular logo with a vertical line in the top right corner. A QR code is located in the bottom left corner. The contact information is listed in the bottom right corner.

**TUM**

**Bob Builder, Ph.D.**

**Address: Country, Street Name 23**

**Phone: 555-12345**

**Mobile: 666-54321**

**Mail: bob@H2R84L4JIL3G5C.zkey**

## NICKnames

- ▶ “alice.bob.carol.dave.gnu” is a bit long for Eve (“.gnu”)
- ▶ Also, we need to trust Bob, Carol and Dave (for each lookup)
- ▶ Finally, Alice would have liked to be called Krista (just Bob calls her Alice)

## NICKnames

- ▶ “alice.bob.carol.dave.gnu” is a bit long for Eve (“.gnu”)
- ▶ Also, we need to trust Bob, Carol and Dave (for each lookup)
- ▶ Finally, Alice would have liked to be called Krista (just Bob calls her Alice)
- ▶ “NICK” records allow Krista to specify her preferred NICKname
- ▶ GNS adds a “NICK” record to each record set automatically
- ▶ Eve learns the “NICK”, and GNS creates “krista.short.gnu”



## NICKnames

- ▶ “alice.bob.carol.dave.gnu” is a bit long for Eve (“.gnu”)
- ▶ Also, we need to trust Bob, Carol and Dave (for each lookup)
- ▶ Finally, Alice would have liked to be called Krista (just Bob calls her Alice)
- ▶ “NICK” records allow Krista to specify her preferred NICKname
- ▶ GNS adds a “NICK” record to each record set automatically
- ▶ Eve learns the “NICK”, and GNS creates “krista.short.gnu”
- ▶ Memorable, short trust path in the future! TOFU!
- ▶ Krista better pick a reasonably unique NICK.

## Shadow Records

- ▶ Records change
- ▶ Expiration time controls validity, like in DNS
- ▶ DHT propagation has higher delays, compared to DNS

## Shadow Records

- ▶ Records change
- ▶ Expiration time controls validity, like in DNS
- ▶ DHT propagation has higher delays, compared to DNS
- ▶ SHADOW is a flag in a record
- ▶ Shadow records are only valid if no other, non-expired record of the same type exists

## Practical Concerns

- ▶ Name registration
- ▶ Support for browsing
- ▶ New record types
- ▶ Integration with applications
- ▶ State of the implementation

## Registering a name in GNS

- ▶ Bob gives his PKEY to his **friends** via QR code
- ▶ or registers it at the **GNUnet fcfs** authority *pin.gnu* as "bob"
- ▶ → Bob's friends can resolve his records via *\*.petname.gnu*
- ▶ → or *\*.bob.pin.gnu*

# From DNS to GNS

Names are not globally unique, but ...

... we need support for Virtual Hosting!

... we need support for SSL!

# From DNS to GNS

Names are not globally unique, but ...


... we need support for Virtual Hosting!

... we need support for SSL!

Solution: Client Side SOCKS Proxy

# Legacy Hostname (LEHO) Records



LEHO records give a hint about the DNS name the server expects.





# Legacy Hostname (LEHO) Records

LEHO records give a hint about the DNS name the server expects.



## Long-Term Vision

- ▶ Integration with browser and HTTP server
- ▶ HTTP server receives “GNS-Zone: PKEY” instead of “Hostname”
- ▶ HTTP client uses “TLSA” record of GNS, instead of “LEHO”

## Relative Names

- ▶ GNS records can contain “.+”
- ▶ CNAME: “server1.+”
- ▶ MX: “mail.+”
- ▶ “.+” stands for “relative to current zone”

Supporting this for links in browsers would be nice, too.

## New Record Types

- ▶ PKEY: delegate to another GNS zone
- ▶ NICK: preferred names for shortening
- ▶ LEHO: legacy hostname

## New Record Types

- ▶ PKEY: delegate to another GNS zone
- ▶ NICK: preferred names for shortening
- ▶ LEHO: legacy hostname
- ▶ GNS2DNS: delegate to DNS
- ▶ VPN: peers hosting TCP/IP services
- ▶ PHONE: call users using `gnunet-conversation`

## DNS Delegation

- ▶ Delegate to DNS using GNS2DNS records
- ▶ GNS2DNS record specifies:
  - ▶ Name of DNS resolver (i.e. “ns1.example.com” or “piratedns.”)
  - ▶ DNS domain to continue resolution in (i.e. “example.com” or “piratebay.org”)
- ▶ GNS will first resolve DNS resolver name to A/AAAA record
- ▶ GNS will then resolve “*left.of.gns2dns.example.com*” using DNS

## VPN Delegation

- ▶ Delegates to GUNet VPN
- ▶ VPN record specifies:
  - ▶ Identity of hosting peer (no anonymity!)
  - ▶ Service identifier (hash code)
- ▶ GNS can map VPN record to A/AAAA record of `gnunet-vpn` tunnel

## PHONE service

- ▶ PHONE record specifies:
  - ▶ Identity of hosting peer (no anonymity!)
  - ▶ Line number (to support multiple phones per peer)
- ▶ `gnunet-conversation` uses *reverse lookup* for caller ID



## PHONE service

- ▶ PHONE record specifies:
  - ▶ Identity of hosting peer (no anonymity!)
  - ▶ Line number (to support multiple phones per peer)
- ▶ `gnunet-conversation` uses *reverse lookup* for caller ID

## ybti assembly plug

- ▶ Bart Polot will present more about Conversation
- ▶ Florian Dold will present more fun GNUnet crypto
- ▶ Julian Kirsch will present Knock

# Application Integration

- ▶ SOCKS proxy (`gnunet-gns-proxy`)
- ▶ NSS plugin
- ▶ DNS packet interception (`gnunet-dns-service`)
- ▶ GNS (C) API
- ▶ GNS (IPC) protocol
- ▶ GNS command-line tool

# Application Integration

```
FILE *p;
char *cmd;
char line[128];
struct in_addr ip;

if (-1 == asprintf(&cmd, "%s %s\n", "gnunet-gns -r -u", name))
    return -1;

p = popen(cmd,"r");

if (p != NULL )
{
    if (fgets( line, sizeof(line), p ) != NULL)
    {
        if (line[strlen(line)-1] == '\n')
        {
            line[strlen(line)-1] = '\0';
            if (inet_pton(AF_INET, line, &ip))
            {
                //Do something
            }
        }
        else
        {
            fclose (p);
            free (cmd);
            return -1;
        }
    }
}
...

```

## Current State

- ▶ GNS part of GNUnet since 0.9.3
- ▶ Crypto changed to Curve25519 in 0.10.0
- ▶ Internationalized Domain Names are supported

## Current State

- ▶ GNS part of GNUet since 0.9.3
- ▶ Crypto changed to Curve25519 in 0.10.0
- ▶ Internationalized Domain Names are supported
- ▶ Installation is “non-trivial” (for your parents)
- ▶ SOCKS proxy is known to be problematic
- ▶ No GUI for TLSA/CERT records yet

## Current State

- ▶ GNS part of GNUet since 0.9.3
- ▶ Crypto changed to Curve25519 in 0.10.0
- ▶ Internationalized Domain Names are supported
- ▶ Installation is “non-trivial” (for your parents)
- ▶ SOCKS proxy is known to be problematic
- ▶ No GUI for TLSA/CERT records yet

## GNS Key Exchange Party Plug

- ▶ Matthias Wachs will describe process at Lightning Talks 2
  - ▶ Install GNUet today & create private key
  - ▶ Use `gnunet-bcd` to create business cards
- ⇒ Print business cards at Wau Holland tomorrow!

End

**Thank you!**

grothoff@in.tum.de

Get the code:

<https://gnunet.org/gns>