


# The Internet Is Broken

VERSIONE ITALIANA

This page collects evidence on the many layers the current Internet is vulnerable, insecure or plain broken and explains how those problems do not exist in our [architecture](#). The [W3C STRINT](#) paper on the subject is also a worthwhile read to get an overview.



## Ethernet, DHCP

When DHCP assigns IP numbers over the Ethernet, [Rogue DHCP](#) servers can perform [man-in-the-middle attacks](#) on devices being added to a local network.

In GNUet, no IP numbers need to be assigned and any other node in the network can safely be used for routing if it is willing and able to route. [Related problems](#) like denial-of-service attacks using DHCP do not exist in the GNUet set-up. Tricks like [PoisonTap](#)'s simulation of a local Internet are impossible thanks to CADET's protection mechanisms against sybil attacks.

## Internet Protocol (IP)

According to [Washington Post's "Net of Insecurity" series](#) the inventors of TCP/IP originally wanted to build basic end-to-end cryptography directly into the protocols, thus guaranteeing at least the authenticity of transmissions if not the content, within the possibilities of the late '70s. By impeding any public use of cryptography, the National Security Agency fundamentally broke the Internet early on. Since then we not only have an Internet which is unencrypted by default, it is also insecure as the provenience of any IP packet can be spoofed at will.

GNUet and secushare communications are encrypted bottom-up. By utilizing the cryptographic identities of entities also as their routing address, redirection of traffic becomes non-sensical. Any node in the network automatically becomes authoritative for itself, simply by generating its cryptographic identity that, thanks to the miracles of mathematics, is extremely unlikely that any other computer could intentionally recreate.

## Border Gateway Protocol (BGP)

[BGP](#), the protocol that was planned on three napkins and implemented with the intention that it would be a temporary hack to solve worldwide routing has been in use for decades now. It suffers from scalability and stability problems and offers plenty of possibilities for hijacking of Internet traffic since any participating ISP is technically able to remap any IP address range to any corner of the planet. This has happened several times already and can be considered a serious weapon in "cyber warfare". Essentially, the broken Internet is built on trust in hundreds of institutions - even sociology teaches us that this cannot work safely.

- <http://www.washingtonpost.com/sf/business/2015/05/31/net-of-insecurity-part-2/>
- <http://www.washingtonpost.com/sf/business/2015/06/22/net-of-insecurity-part-3/>
- <http://arstechnica.com/security/2013/11/repeated-attacks-hijack-huge-chunks-of-internet-traffic-researchers-warn/>

As described in CADET documentation and [video](#) presentations, GNUet finds its path from one end to the other end of a new Internet by itself - evading any attempts to misguide it. You heard the claim before: the Internet routes around [censorship](#). With GNUet this would actually be true, for the first time.

## Missing generic distribution and scalability layer

Up into the 90's several applications such as NNTP and IRC had experimented with multicast distribution logic. In 1992 the IETF attempted to solve the [scalability](#) issue once and for all, by ratifying the 'IP Multicast' protocol.

Why that failed and how secushare intends to address the issue with its own multicast layer in GNUet is described on the [pubsub](#) page. The consequence of a missing generic scalability facility is that in practice only commercial content delivery networks and cloud systems can deliver this crucial property for popular adoption of any Internet service.

## TCP, UDP

[TCP](#) has its own long history of vulnerabilities, mostly because of the aforementioned lack of cryptography.

There is no equivalent unsafe communications protocol over GNUet. All virtual circuits provided by CADET are cryptographically secured and enhanced with ratcheted forward secrecy. Hijacking connections is impossible, for example.

## Tor, the Onion Router

Tor is a fine end-to-end authentication and encryption layer when used with personal hidden services, superior to TLS and X.509 discussed below. Beyond the generally known issues however, the mere "socket" application programming interface (API) used to send and receive data over TCP opens up a potential for de-anonymization by shaping of traffic.

This is discussed on the [anonymity](#) page as it affects most attempts at anonymizing TCP streams. That page also explains in-depth how metadata protection can be improved by GNUet and secushare.

## Domain Name System (DNS)

Very vulnerable.

- [https://en.wikipedia.org/wiki/Domain\\_Name\\_System#Security\\_issues](https://en.wikipedia.org/wiki/Domain_Name_System#Security_issues)
- <http://www.networkworld.com/article/3134516/security/answers-to-is-the-internet-broken-and-other-dyn-ddos-questions.html>

At the 30c3 congress we had an exciting panel featuring authors of DNSSEC, CurveDNS, Namecoin and GNS. It became clear in the debate how the GNU Naming System is the most advanced in addressing all of the involved issues.

## TLS & X.509

With TLS/SSL typically wrapped around HTTP and X.509 as its method for authentication the Internet has encryption on the wrong network layer and relies on external institutions for authenticity which can be subverted by governments at will.

- [Big corps are tracking people by means of long-lived TLS sessions;](#)
- [HSTS Tracking;](#)
- [82.9% of webserver supporting forward secrecy were using weak Diffie-Hellman parameters;](#)
- [Wikipedia List of TLS Vulnerabilities;](#)
- See also the reasons why we developed the ['Certificate Patrol'](#) addon for Torbrowser/Firefox and the many articles we collected that proved us we weren't paranoid enough.

See above in the 'IP' box on how GNUet encrypts everything from the bottom without needing external authorities. See GNS and the secushare distributed social graph for better ways to trade authentication information rather than X.509 and DNSSEC/DANE.

## Authentication and Security in IT and the Internet of Things

While the business world is keen on selling internet-enabled things, the absence of reasonable methods of authentication are a continuous threat for these devices to be subverted and used against their owners. The [security](#) page describes just how the distributed social graph of secushare can reduce if not eliminate the biggest of threatening scenarios regarding unmaintained devices in IT security in general.

- <http://www.washingtonpost.com/sf/business/2015/11/05/net-of-insecurity-the-kernel-of-the-argument/>
- <http://www.washingtonpost.com/sf/business/2015/07/22/hacks-on-the-highway/>
- <http://www.washingtonpost.com/sf/business/2015/06/22/net-of-insecurity-part-3/>

If pervasive and cryptographic, [social](#) can truly be the keyword to solve technological problems that are in fact sociological. If all network interactions first require a permission from humans to even take place, the secushare communication protocols are the only pieces of software that really need to be fully bug-free. Any bugs that may exist in application software and even system kernels no longer come to play unless when abused by [people](#) that were granted access and therefore supposed to use them respectfully.

It is not a stretch to say that the GNU Internet would indeed be a safer place with most known problems of the broken Internet resolved.

## E-Mail

Not only has e-mail been designed without privacy in mind, which is why [PGP encryption on top of mail](#) is so hard and unnatural to use, it was also invented at a time when all participants on the network were friendly professionals who would not send unsolicited spam mails.

The problem with [spam](#) is how the mail system still has no awareness of who is who. It tries to blacklist the evil sources in a universe where the nasty sources are exploding in number instead of whitelisting the trustworthy senders since in many regular mail systems there is no feedback channel from the user interface, telling the mail servers which mail sources were actually welcome.

Among the 600 spam mails I received today, there were 4 legitimate mails that shouldn't have ended up in the spam folder. This makes e-mail technically an unreliable communication medium - you can never be sure it actually works.

This technological absurdity has created a multi-million market for corporate solutions. Since homegrown indie spam protection isn't sufficiently effective, anyone who'd like to have an acceptable e-mail experience needs to use a surveillance economy offering such as G-Mail or Microsoft Hotmail. Even then, that really important mail might have ended up in the spam folder, which makes corporate walled gardens the ultimate peaceful haven: if you want no trouble from spammers, use Facebook, Twitter or [Riot/Telegram/Signal/Whatsapp](#).

Why is their model working? Because they map the social graph information of who is a real person rather than a spambot, directly onto the delivery mechanism.

As described in [business](#) and [like](#), in a GNU Internet you cannot send an unsolicited message to a complete stranger. You either have a social network link that credibly confirms you are a real person, or, for example, you submit a plea on that person's public website to get back to you, automatically adorned by cryptographic authentication, so when that person responds to you they will be put through directly into your inbox. In other words, in a GNU Internet, the spam business and criminality model is terminated.

## Wild West Web

The entire architecture of the [Web](#) is optimized for the data collection economy, containing surveillance taps in [HTTP](#) (ETag, Cookie, content negotiation...), HTTPS (persistent TLS session identifiers), Javascript (AJAX, WebRTC, Canvas API, logging of mouse movement, measurement of keyboard hesitation and typing skills, font selection frameworks etc), HTML directly (HTML video) and indirectly by allowing for inclusion of de-anonymizing scripts, fonts and images from surveillance websites such as Facebook and Google. Web browsers themselves provide built-in "features" such as "Safe Browsing", which identify the user with Google immediately at the start of the session. The entire architecture of the web introduces an artificial greed for real-time access to servers and even cloud systems to address the dependency on reliability.

- <https://ashkansoltani.org/2012/02/25/cookies-from-nowhere/>
- [https://en.wikipedia.org/wiki/HTTP\\_ETag#Tracking\\_using\\_ETags](https://en.wikipedia.org/wiki/HTTP_ETag#Tracking_using_ETags)
- <http://lucb1e.com/rp/cookielesscookies/>
- <https://security.stackexchange.com/questions/12679/how-can-i-prevent-tracking-by-etags>
- <https://iranthreats.github.io/resources/webrtc-deanonymization/>
- ...

The [OnionScan](#) project identifies insecure usage patterns in mostly web technologies that lead to subsequent de-anonymization of Tor-based services. It illustrates how the web needs a revamp in order to not undermine the newly acquired privacy when using a more private kind of Internet. The [business](#) page suggests an entirely different mode of operation for an ethical kind of worldwide web.

## Missing generic micropayment layer

For a brief moment in the mid-90's Internet commerce was undecided between going for advertising or introducing a micropayment system. Unfortunately the contenders in the payment market didn't publish their source code, greedily thinking they could establish a cash cow for themselves. In the meantime the advertising industry discovered the value in collecting personal data of customers. Now the Internet does not even know or consider an alternative to the surveillance economy which threatens constitutional pillars of democracy.

- <http://nymag.com/intelligencer/2018/04/an-apology-for-the-internet-from-the-people-who-built-it.html>

[Taler](#) is a GNU free software implementation of the sort of micropayment system the Internet should have had then and now. It works both over the broken as over a GNU Internet and provides an ethical alternative to both the surveillance and the blockchain economy (Bitcoin and remixes). The latter provides no method for human society to get its fair share in transaction fees, thus [acting like a digital tax haven](#), not to mention the ecological damage.

## Facebook vs. Federation

The absence of a generic solution for [scalability](#) required and empowered corporations to develop proprietary server-side distribution systems, starting in 1994 with the "Bundesdatenautobahn", followed up by Akamai, leading to today's cloud infrastructure. Combined with the data aggregation economy, which works best for those who have the most data, an oligarchic architecture of near monopoly has formed. The absence of low-level end-to-end cryptography has further enabled these players to monetize intimate data. The absence of an oblivious and natural payment system has made them gatekeepers of Internet-based trade.

Traditional Internet legends suggest that the alternative to cloud systems is the [federation](#) of private servers using open standards such as SMTP, XMPP, Diaspora, GNU Social and so on. Follow the link to see a description why that hasn't worked out for several decades.

The secushare and GNUet approach is to completely revamp the underlying infrastructure, providing for the missing architectural layers so that humans are no longer dependent on corporations to achieve digital freedom of interaction and intimacy.

## Unsafe encryption over open standards

See also our critique on [end-to-end cryptography over insecure federated protocols](#) such as SMTP and [XMPP](#).

## Does it have to be GNUet and secushare?

Of course not. Anyone could have started in 2001 thinking about these things and get to the point where we are today, unfortunately we don't know of any other technology that is organically looking at the entire problem stack, not just some parts of it, but that is not a hindrance to start. You may want to avoid repeating mistakes that others made before you (computer science has always been good at repeating mistakes, so this is a really hard challenge).

But in the end it doesn't matter where the solution comes from as long as it does indeed work, that's why we support the idea that a [legislation](#) could define the framework, the wishlist of features a future Internet should provide, then the industry is welcome to find solutions in due time, with or without secushare and GNUet. Even if secushare was a ready-to-use reality, it would still need a hand getting established. Legislation might be the only hand that can establish an ethical choice over commercial convenience.