



Design Requirements for Civil Internetworking (Position paper)

Christian Grothoff

► To cite this version:

Christian Grothoff. Design Requirements for Civil Internetworking (Position paper). Protecting online privacy by enhancing IT security and strengthening EU IT capabilities, LIBE committee and the STOA panel together with the Luxembourg Presidency, Dec 2015, Brussels, Belgium. hal-01244744

HAL Id: hal-01244744

<https://hal.inria.fr/hal-01244744>

Submitted on 16 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Design Requirements for Civil Internetworking

Position paper

Christian Grothoff
christian.grothoff@inria.fr

November 23, 2015

1 Early mistakes

Internet protocols need a general overhaul to make them suitable as the main communication infrastructure for a sustainable civil liberal society. When the US military created the Internet, the fundamental goal was to interconnect existing networks, which was a simple practical objective. What is more lasting were the second level goals *in order of importance* as elaborated by David Clark in [2]:

1. Internet communication must continue despite loss of networks or gateways
2. The Internet must support multiple types of communications service.
3. The Internet architecture must accommodate a variety of networks.
4. The Internet architecture must permit distributed management of its resources.
5. The Internet architecture must be cost effective.
6. The Internet architecture must permit host attachment with a low level of effort.
7. The resources used in the internet (sic) architecture must be accountable.

We note that the only security-related goals relate to availability and usability. While some Internet architects were aware of encryption due to their proximity to the NSA, encryption was classified and thus could not even legally be included in the emerging Internet design that was deliberately non-classified.

As a result of these 1970 design goals, the lack of data protection remains visible at all layers of the Internet architecture today. For example, the IP header from 1981, which is still in use today, unnecessarily leaks information about the sender address and the encapsulated transport protocol. Additionally, the checksums in TCP/IP are insufficient to provide authenticity, and it is difficult to prove ownership of a destination address.

With the emergence of e-commerce, a thoroughly broken cryptographic protocol called “secure sockets layer” (SSL) was introduced to provide a minimum of authenticity and confidentiality. The regulatory environment, rushed deployment and lack of understanding of cryptography by the designers resulted in a complex design that took the Internet community 20 years to sufficiently understand and thereby obsolete. The surviving complex cryptographic libraries remain a treasure trove for vulnerabilities.

2 Civil Needs for Private Communication

Today, the actual use and thus the social requirements for a global network differs widely from those goals of 1970. While the Internet remains suitable for military use, where the network equipment is operated by a command hierarchy and when necessary isolated from the rest of the world, the situation is less tenable for civil society.

Due to fundamental Internet design choices, Internet traffic can be misdirected, intercepted, censored and manipulated by hostile routers on the network. And indeed, the modern Internet has evolved exactly to the point where, as Matthew Green put it, “the network is hostile”. Even if we consistently deployed authenticated encryption on the Internet, the network would still learn way too much. Authenticated encryption does not obfuscate sender and receiver, and also discloses the times, frequency and volume of the communication, which enables reverse engineering of the set of pages visited via traffic analysis. Thus, it is not enough to encrypt, we need to also protect the metadata.

Metadata is by definition data about data, and thus provides a higher level abstraction of the underlying data. Therefore, leaking metadata is not a cosmetic issue. As Michael Hayden, the former head of both NSA and CIA explained, the US government “kills based on metadata”. The democide that has resulted from the NSA’s application of statistics and machine learning to metadata leaked by information networks [3] is thus partially the responsibility of computer scientists, as we have failed to build and deploy systems that adequately protect this critical information.

Authenticated encryption is also of limited utility as long as we stick to centralised service providers operating on the “Big Data” business model. The NSA’s PRISM program shows how the hunger for data by the spies is paralleled by the hunger for data by the advertising-based service providers, resulting in an unholy union. Privacy and data protection issues will continue as long as the dominant business model on the Internet is for service operators to collect and monetize client information.

3 Design Requirements for a GNU Network

Liberal society needs a network architecture that uses the anti-authoritarian decentralised peer-to-peer paradigm and privacy-preserving cryptographic protocols. Civil networks need to be anti-authoritarian, as modern networks are a critical common used for virtually all civil, political and economic activity. Just like a liberal society cannot allow anyone including the government to exercise totalitarian control over roads, food, water or electricity, we similarly cannot allow GCHQ/NSA to achieve their stated goal of information dominance.

We propose to rearchitect the Internet and create the GNUnet, an evolution of the Internet with the fundamental goal of lessening the threat of networks supporting totalitarian control over the population. This includes both control over their networking and computing, as well as ensuring that the network protects privacy. To clarify this, we propose the following ten objectives for the design of the GNUnet, in this order:

1. The GNUnet must be implemented as free software, to ensure that citizens enjoy the four essential freedoms of free software [4].
2. The GNUnet must only disclose the minimal amount of information necessary, to second and especially third parties. Each user decides what information can be shared and with whom.

3. The GNUnet must be decentralised and survive Byzantine failures in any position in the network. Increased malicious participation may cause security and privacy assurances to deteriorate commensurate with the resources expended by the adversary.
4. The GNUnet must make it explicit to the user which entities must be trustworthy when establishing secured communications.
5. The GNUnet must use compartmentalization to isolate sensitive information against Byzantine failures even among layers of the implementation on the same device.
6. The GNUnet must be open and permit new systems to join.
7. The GNUnet must be self-organizing and not depend on administrators.
8. The GNUnet must support a diverse range of applications and devices.
9. The GNUnet must scale and be cost effective.
10. The GNUnet must provide incentives for participants to contribute more resources than they consume.

We note that many of Clark's original Internet design goals correspond to design goals the bottom of the priority list for the GNUnet: enabling host attachment, diversity of use, cost effectiveness and accountability remain important, but need to fall behind the need for data protection and Byzantine fault tolerance.

The objectives formulated for the GNUnet also go beyond the simple slogan of "privacy-by-design" [1]. While objective (2) is enshrining data protection as a key goal, privacy-by-design only calls for respect for the user's privacy, while objectives (1) and (3) are more broadly protecting the user's autonomy. In contrast to privacy-by-design, for the GNUnet the use of proprietary software or centralized service providers is not acceptable, as they make citizens dependent on software vendors or service providers.

The last point hints at the importance of economic considerations for future network designs. In particular, it is crucial that we open opportunities for alternative business models that are not based on advertising and Big Data, for example by providing privacy-preserving micropayments. For this to happen, liberal financial regulation, that does not strangle emerging privacy-preserving payment systems, will be key. Only then we can hope to break up the unified front of the "Big Data" industry and "mass surveillance" spy agencies against citizens.

References

- [1] Ann Cavoukian. Privacy by design. Technical report, Information & Privacy Commissioner, Ontario, Canada, 2009.
- [2] D. Clark. The Design Philosophy of the DARPA Internet Protocols. *SIGCOMM Comput. Commun. Rev.*, 18(4):106–114, August 1988.
- [3] Yves Eudes and Christian Grothoff. Skynet, le programme ultra-secret de la NSA créé pour tuer. *Le Monde*, (20.10.2015):14, January 2015.
- [4] Richard M. Stallman. *Free Software, Free Society: Selected Essays of Richard M. Stallman*. GNU Press, Boston, Massachusetts, 2002.