

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: August 18, 2014

J. Pouwelse, Ed.
Delft University of Technology
February 14, 2014

**The Shadow Internet: liberation from Surveillance, Censorship and
Servers
draft-pouwelse-perpass-shadow-internet-00**

Abstract

This document describes some scenarios and requirements for Internet hardening by creating what we term a shadow Internet, defined as an infrastructure in which the ability of governments to conduct indiscriminate eavesdropping or censor media dissemination is reduced.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 18, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as

described in the Simplified BSD License.

Table of Contents

1.	Requirements Language	3
2.	Introduction	3
3.	Real world scenario: Arab Spring context	4
4.	Microblogging	5
5.	Onion routing and bandwidth accounting	7
6.	Driving scenarios	7
6.1.	20sec scenario	8
6.1.1.	Adversary model: A simplistic attacker	8
6.1.2.	Scenario details and architectural requirements	8
6.2.	Kill-switch scenario	10
6.2.1.	Adversary model: An advanced attacker	10
6.2.2.	Scenario details and architectural requirements	10
6.3.	Friend-to-friend scenario	11
6.3.1.	Adversary model: A powerful attacker	11
6.3.2.	Scenario details and architectural requirements	12
6.4.	Transmorph ability	13
6.5.	A single global conversation	13
6.6.	Spammers and hoaxes	14
7.	Design principles: simplicity and prior success	14
8.	Background rant: lack of coordination and fragmentation	14
9.	Current running code and related work	15
10.	Open issues	16
10.1.	Use cases and threat model	17
10.2.	System components, definitions and system architecture	17
10.3.	Current technology and gap	17
10.4.	Detailed system design and protocol specification	17
11.	Security Considerations	17
12.	IANA Considerations	17
13.	References	17
13.1.	Normative References	17
13.2.	Informative References	17
13.3.	URL References	17

1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

2. Introduction

The shadow Internet is an alternative communication infrastructure specifically crafted to be resilient to sniffing, blocking, filtering and shutdown. We discuss usage scenarios, architectural requirements and how it could be used to combat global surveillance and censorship.

Pervasive monitoring is a widespread attack on privacy [[ATTACK](#)]. It is defined as widespread (and often covert) surveillance through intrusive gathering of protocol artefacts, including application content, or protocol meta-data such as headers. Bits moving across the Internet are under surveillance at an unprecedented scale. Today, both Internet providers and governments possess the ability to monitor the moves of their digital citizens from central infrastructure points. The leaks by Edward Snowden show the capabilities of the NSA through their ANT product catalog [[CATALOG](#)]. Internet routers from various vendors are easily compromised in their entirety and "taps" may be inserted. Furthermore, some packets may be re-routed to modify them, called man-in-the-middle Internet route hijacking. The challenge is to make global pervasive surveillance expensive again. We need to figure out how to re-engineer the internet to prevent this kind of wholesale spying, see [[SCHNEIER](#)].

Censorship is another key threat to The Internet. We focus on media censorship and anonymity. The largest source of Internet traffic is video traffic. Video is a powerful medium for reaching and mobilizing mass audiences, with unique potential to bring about social change. Anonymous video distribution is needed to prevent repressive governments from stifling the free exchange of information that enables citizens to organize effective political opposition. As such, it constitutes a key tool for allowing populations to establish self-governance, and for safeguarding human rights around the world. Yet current methods of video distribution are controlled by a small number of gatekeepers, and are difficult to access privately or anonymously. Anonymous video distribution depends on technical structures to support the building of trust, community, and cooperative relationships between users. Until now, the difficulty of confronting this challenge has prevented any system from offering an effective solution.

Servers also cause problems. Email provider Lavabit is in court

appealing against a government order to hand over its encryption keys. Owners of servers can be forced to cooperate to install a "pen register", which tracks communication patterns. In the case of Lavabit, authorities demanded he hand over the encryption keys for its entire service and expose all customers. Essentially, a company director can be forced to choose between committing sabotage on their own security or go to jail. Therefore, servers are a key risk factor in the shadow Internet architecture. Servers often keep extensive logs, another risk factor. No privacy exists for the one billion regular users of online video. Rather, their watching patterns are meticulously recorded, including what they watched, when they watched it, and with whom they shared it.

We outline scenarios for the shadow Internet and several requirements. A key principle of the shadow Internet is that it offers usable encryption without infrastructure. The shadow Internet is required to protect the privacy of its users, allowing them to remain anonymous if they wish, or to create digital personas for sharing and commenting on videos. Privacy will be achieved by having participants work together, each relaying the encrypted traffic of others. For sustainability, participants will need to relay several hours of traffic for each hour of video that they watch, so encouraging cooperation is a key part of our scenario. The presented scenarios strongly rely on a sense of community and mutual trust. The potential user-base of our system is large: over 6 billion hours of video are watched each month on YouTube alone. Our use-case and scenarios are focused on large-scale penetration and uptake.

3. Real world scenario: Arab Spring context

Video footage of human rights violations can be a powerful catalyst for change. Graphic depiction of atrocities for a world-wide television audience can influence outcomes. However, global news networks need to be able to obtain video footage from a hot zone.

Governments have demonstrated their ability to disable communications networks in times of crisis. During the 2011 Arab Spring, Egyptian authorities demanded that telecommunication companies sever their broadband connections and mobile networks--both local and European operators were forced to comply, and, as a result, digital Egypt vanished. Despite the country's decentralized infrastructure, an Internet blackout was relatively easy to carry out. The roles and consequences of social media (e.g., Facebook and Twitter) during that same period further illustrate the capacity governments have for Internet censorship and the challenges activists face in combating it. The April 6 Youth Movement from Egypt committed digital dissent in full public view. According to The New York Times[YOUTH], the movement "provided a structure for a new generation of Egyptians, who

aren't part of the nation's small coterie of activists and opinion makers, to assemble virtually and communicate freely about their grievances." But moving protest organizations to social media accessible to the public-at-large can hold surprising risks. On the ground, the movement's organization of labor strikes and protests in Facebook groups, many with thousands of followers, triggered arrest and imprisonment. Protesters in other countries quickly took note of Egypt's lesson and disabled their public Facebook profiles. In response, one government initiated social media searches on incoming, young, plane travelers by forcing them to login to Facebook upon arrival, thereby revealing online activities and any anti-government sympathies. [[FORCEDLOGIN](#)]

A glimmer of hope exists. The Arab Spring shows that a new generation is claiming their right to express themselves. Microblogging, social media in general and traditional satellite news broadcast networks are perceived as critical catalysts for political change. Generic computational fabric is soon getting in the hands of two billion people with the growth of smartphones and increasingly affordable communication. These smartphones are increasingly used to record and spread disruptive audiovisual material, even in regions without media freedom.

Democratic countries also face a dilemma. Restrictions on the free information flow is the topic of several proposed laws by elected representatives. The strength of copyright law impacts digital information flow. Politicians must decide between weak copyright law, as championed by civil rights activists versus strong copyright enforcement, as promoted by numerous players in the creative industries. Recent furor around SOPA, PIPA, etc. in the US plus the European Parliament vote on ACTA is highly relevant in this context.

The uniqueness of The Internet lies in the IETF standards. Moving certain bits to certain locations or offering a service requires no prior official approval. However, Internet-deployed mechanisms now exist which filter news and media in general for both surveillance and censorship. The Internet has ceased to provide reliable transport service for all users. The IETF can repeat its historical inter-networking role again by setting the standard for reliable flow of media packets.

4. Microblogging

Microblogging is an increasingly popular technology for lightweight interaction over the Internet. It differs from traditional blogging in that [[OPENMICRO](#)]:

- o Posts are short (typically less than 140 characters, which is the limit in SMS).
- o Posts are in plain text, but may contain links to photos or videos, often taken and uploaded with a mobile device.
- o People can reply to your posts, but not directly comment on them.
- o People learn about your posts only if they have permission to view them.
- o Your microblogging feed is discovered based on your identity at a domain or with a service.

The goal is creating a microblogging standard and facilitating a reference implementation for portable devices which is capable of operating in a hostile environment. This standard should be resilient against all known forms of censorship. This proposed draft standard SHALL provide: "information dissemination from a single smartphone to an audience of millions in the form of microblogging, enriched with pictures or streaming video which is guarded against all known forms of censorship such as: cyberspace sabotage, digital eavesdropping, infiltration, fraud, Internet kill switches, physical checkpoints and lawyer-based attacks with the best known protective methods".

The focus on microblogging is driven by feasibility. Creating a standard for overcoming censorship for social networks, search engines or web browsing in general is extremely challenging. Mitigating the threats posed by Internet kill switches requires focus on the most feasible viable standard. The related work listed in this document shows existing operational systems. Existing systems cover all functionality we desire, however none of them cover all aspects and little interoperability exists.

As early as 2006, long before Arab spring events, it was reported that individuals in wide swathes of the Arab world were using Bluetooth technology to bypass police restrictions. According to news reports[[DATING](#)], communication between men and women in this region had been made possible by cellphone technology. When Bluetooth-capable phones are in close proximity, they can engage directly in digital and social chatter--no other infrastructure is needed. Moreover, when sharing photos or bandwidth-hungry videos with friends it also pays to be close. Government provided cellphone networks might not be filtering you, but can still be dreadfully slow. It therefore pays to use cell phones' Bluetooth-based, direct file-transfer features--and it comes as no surprise that wireless-transfer apps have seen millions of installs. A query of Google

Trends for the phrase "Bluetooth transfer" reveals a geographical spread of this interesting social phenomenon[TREND]. It seems millions of mobile phone owners are already employing the social practice of wireless data exchange. Viability is increased by building upon this practice.

5. Onion routing and bandwidth accounting

Tor pioneered the technique of relaying traffic to improve privacy and security. We extend their valuable work in our scenarios. However, our attacker model differs significantly and our goal is robust video streaming, instead of browsing. Another key difference is our focus on communities and utilizing trust between participants.

Onion routing consumes a lot of bandwidth and processing capacity. Within Tor there is no direct rational incentive to operate a relay node or exit node. In our outlined scenario and mechanisms we reward cooperating (bandwidth donating) individuals with priority over freeriders. This is an old idea, yet never realized. Large-scale uptake is a key challenge.

We believe bandwidth accounting is essential for anonymous streaming, as it creates an incentive mechanism motivating people to participate in a collective system and thus contributing to its sustenance. This, in turn, makes pervasive surveillance harder. Within the scenarios we strive to leave out design and implementation details. However, incentive compatibility is a MUST HAVE requirement and explicitly included.

In our scenario we assume that it is possible to create and distribute proof-of-work certificates. Such certificates are the technological basis for the incentives. Helping others with becoming anonymous through onion routing yields a cryptographically signed proof-of-work certificate. Such proof-of-work certificates somehow provide cooperating individuals with priority or faster service.

6. Driving scenarios

Recent Arab spring events have shown the power of ubiquitous camera-phones, new media and microblogging. This document proposes to use smartphones, wifi and USB sticks for multimedia transport and playback. The architecture, features and driving scenarios are specifically crafted to enable compliant implementations as a single smartphone app without any additional server infrastructure.

Each scenario is focused on certain threats in a hostile environment. The adversary becomes stronger in several of the following scenarios and we also focus on the social media context.

6.1. 20sec scenario

First scenario, called "20sec", defines an open microblogging standard. This first scenario duplicates existing microblogging practices with an open standard in a fully decentralized setting. The scenario requirements are performance equal to central-server based approach (e.g. the ability to reach 20 million people in 20 seconds).

6.1.1. Adversary model: A simplistic attacker

Eavesdropping is a common and easy passive attack in a hostile environment. In this scenario we assume the attacker has full access to the network between the user and any Internet server. Specifically, the adversary can observe, block, delay, replay and modify all traffic coming from any server. Furthermore, all servers such as DNS servers, web servers, swarm trackers, CDN cloud servers and access portals are assumed to be under direct or indirect control of the adversary.

The adversary cannot compromise traffic between smartphones or other participating devices. The adversary cannot compromise smartphones or other participating devices. The adversary cannot break standard cryptographic primitives, such as block ciphers and message-authentication codes.

6.1.2. Scenario details and architectural requirements

Smartphone owner Alice with wifi-based Internet access records an eye-witness video. She attaches this video to a microblog entry and shares the story and the video content automatically with friends Bob and Charlie, who are subscribers of her news feed. Alice does not need to trust any central server with her credentials, nor has to prove her identity to a central (web) server. Bob and Charlie are both behind a NAT middlebox compliant to the BEHAVE recommendations [[RFC4787](#)]. No assistance of a coordinating server (e.g. STUN or TURN) is required to traverse this NAT box using UDP messages. This scenario assumes direct or NAT-based Internet access (the next scenario deals with packet forwarding).

Performance should be equal to a central-server based approach, providing the ability to reach 20 million people in 20 seconds. This first scenario duplicates existing microblogging practices with an open standard in a fully decentralized setting. The 20sec scenario requires that solutions provide seamless backwards compatibility with existing leading solutions (e.g. Twitter, Sina Weibo, chyrp, heello) by using content import tools. Proposed open solutions **MUST** permit easy bulk transcoding and ingest of existing news feeds into this

open standard.

An essential feature of the 20sec scenario is all central gatekeepers or communication to them is possibly compromised. Ownership of data is fundamental to autonomy. To meet the anti-censorship goal, 20sec assumes an infrastructure which is not dependent and completely decoupled from potentially hostile servers such as DNS servers and web servers. 20sec MUST be based on full self-organization. The infrastructure consists purely of devices running compliant implementations. No central server requires installation or maintenance, making this infrastructure independent on any type of funding or business model. 20sec requires an overlay which is highly resilient. Smartphones, tablets and PCs are able to utilize this P2P overlay for microblogging. Existing solutions such as [\[OPENMICRO\]](#) require a central webserver and OAuth-like authentication primitives. This prior work is not suitable for our 20sec scenario, as we aim to remove all server, ultrapeer or superpeer reliance and equality of all participants in the overlay.

When Alice downloads the smartphone app and runs it for the first time, the application performs a bootstrap phase. On this initial startup, the microblogging software looks for at least one other peer in the overlay. The simplest method of bootstrapping is to use a list of peers currently online, together with their port number. See the example below.

```
# file: Central-Bootstrap-Servers.txt
# default bootstrap peers
server1.always-online.org 6420
host1.never-offline.ro 6420
sealand.routed.org 6420
168.0.0.13 6420
```

A file sharing program needs a fresh list of peers to bootstrap. Thus a pre-defined list of peers is included in the software installer. As peers can go offline it is important that at least one peer out of possibly thousands on the list is still online. This pre-existing address list of possibly working peers must therefore remain valid for as long as possible. Bootstrapping is done by contacting peers in the list, possibly in parallel. If a single peers replies, the smartphone app of Alice is connected. Once connected, a fresh list of working peer Internet addresses COULD be requested. Several ideas have been proposed on bootstrapping systems without an "online bootstrap server" list. For instance, simply by smart brute force pinging, as described by the University of Denver [\[BOOTSTRAP\]](#).

It is RECOMMENDED compliant implementations explore and implement

efficient alternatives for decentralized initial bootstrapping.

6.2. Kill-switch scenario

This scenario describes a situation without any Internet access. We assume the government has essentially "killed" the Internet, in an Arab spring like scenario. It is focused on ad-hoc packet forwarding between smartphones.

6.2.1. Adversary model: An advanced attacker

The adversary has disabled all Internet-based communication.

We assume the adversary cannot eavesdrop, jam, delay, replay, modify or spoof wireless communication between smartphones. The adversary cannot compromise smartphones or other participating devices. The adversary cannot break standard cryptographic primitives, such as block ciphers and message-authentication codes.

6.2.2. Scenario details and architectural requirements

Smartphone owner Alice has no Internet access. She records a video, attaches this video to a microblog entry in her phone app. Friends Bob and Charlie are subscribed to her news feed. Bob and Charlie are at some point within range of the wifi, bluetooth or other wireless capability of Alice. This fresh microblog entry plus video is shared automatically. Bob obtained the message from Alice using a smartphone app which is periodically scanning if other devices are around and if they possibly have fresh news. This periodic synchronization SHOULD be energy-efficient. Bob sees no noticeable decrease in battery lifetime after he obtained unconstrained news access. Charlie later goes to a square where numerous people have gathered, most of which are highly interested in the latest videos. The fresh messages automatically spreads in this crowd.

Note that this scenario differs from Delay-Tolerant Networking (DTN), as being investigated by a Working Group within the Internet Research Task Force [[RFC4838](#)] and scientists[BUBBLE]. The DTN focus is on finding routes to an explicitly given destination, usually by maintaining routing tables. Their system model and terminology cannot be applied in our context, for instance, "Endpoint Identifiers" which identify the original sender and final destination. In our Internet-Free scenario sender Alice does NOT explicitly send a message with destination Bob.

A wealth of related work exists in this area. General solutions are found in mobile ad hoc networks (MANET), which provide self-organized IP routing among wireless devices, and delay-tolerant networks (DTN),

which use a simple store-and-forward primitive to communicate over heterogeneous links. Mobile ad hoc networks have been studied within the Internet Research Task Force (IRTF) since 1997, leading to several standards published by the IETF's MANET Working Group, while delay-tolerant networks are currently the focus of the IRTF's DTN Research Group. We hope that much of that knowledge can be reused, despite our scenario differing slightly from DTN (as being investigated by the IRTF [[RFC4838](#)])

6.3. Friend-to-friend scenario

This third scenario uses friend-to-friend networking to remove the requirement for active networking and wifi sensing. The smartphones of Alice and Bob need to be synced manually. This scenario SHOULD deliver a privacy-by-design type of microblogging service.

6.3.1. Adversary model: A powerful attacker

We must assume from the Arab Spring scenario the existence of a powerful adversary. For instance, the adversary has disabled all Internet-based communication. The adversary even actively monitors wireless communication. Protocol designers have identified the following threats [[BRIAR](#)] for similar circumstances:

- o The adversary can observe, block, delay, replay, and modify traffic on the underlying network. Thus, the microblogging service must ensure end-to-end security without relying on the security of the underlying network.
- o Wireless communication is regularly monitored. Responding to any wireless requests from a stranger is a direct threat to the user and extremely harmful.
- o Possession of encrypted electronic messages or encryption technology in general is extremely harmful to the smartphone owner.
- o The adversary has a limited ability to compromise smartphones or other participating devices. If a device is compromised, the adversary can access any information held in the device's volatile memory or persistent storage.
- o The adversary can choose the data written to the microblogging layer by higher protocol layers.
- o The adversary cannot break standard cryptographic primitives, such as block ciphers and message-authentication codes.

Encryption is not a sufficient requirement of the friend-to-friend scenario, everything **MUST** be hidden. Possession of smartphones apps with encryption is already dangerous for the owner.

6.3.2. Scenario details and architectural requirements

Reports from repressive regions indicate that USB sticks are commonly used to transport sensitive information. See for instance this extensive report on North-Korea [[NKOREA](#)]. In the friend-to-friend scenario a network of friends is trusted to transport news manually, by simply carrying it around. Smartphones with NFC capability or manual USB transfer are used to duplicate and move messages. Thus Alice delivers her fresh news message to Bob, which is later given manually to Charlie.

As direct social connections are sparse and proximity of friends is not continuous, this scenario **SHOULD** facilitate usage of friends-of-friends or further removed social ties to relay news messages. This requires the development of a decentralized social network, for instance, with digital signatures of friendship certificates. In effect this would create a "decentralized social network", completely autonomous and owned by all participants. We assume Alice only has Bob in her friendlist and Bob only has Charlie in his friendlist. An **OPTIONAL** feature is that the smartphone apps running on the smartphone Alice and Charlie detect that they have friendship path through Bob. Fresh news is thus exchanged.

The interception of a single smartphone **MUST NOT** expose the app itself, any friend list or worse: the entire social network. We assume Alice is placing herself in danger with electronic tools for "subversive activities against the democratic republic". Information hiding techniques are essential or even life-critical. Possibly based on Zero-Knowledge Proof (ZKP) protocols [[ZEROKNOW](#)]. The smartphone app **MUST** pose as a harmless entertainment feature of a smartphone or use another mechanism to become a "stealth app". The requirement of such a stealth app is that a somewhat knowledgeable person will not detect the presence of the app and will not discover any video content, hence making the app checkpoint-proof. The app itself should be hidden, i.e., it should not be visible in the app list of the phone but, for example, be activated by dialing a secret telephone number. In addition, the app should be able to virally spread and be able to bypass any governmental restrictions on the official app store.

This scenario requires modification and enhancement based on real-world experience from human rights activist [[EGYPTSTUDY](#)].

6.4. Transmorph ability

Prior scenarios expanded the threat model. This and the following scenarios are focused on the social media context. News is created in a region without freedom and then needs to travel to the outside world. We refer to this simply as the freedom/non-freedom border. Different transport protocols, dynamics and different solutions are needed on the two sides of this border.

We now expand the friend-to-friend scenario with a transmorph ability, the ability of news to cross the freedom/non-freedom border.

Alice is a well known blogger in an region with extreme censorship. Her identity on Twitter has millions of followers. However, she has no direct ability to reach a Twitter.com server or Internet in general. We assume Alice only has Bob in her friendlist and Bob only has Charlie in his friendlist. Charlie is able to smuggle a collection of messages out of the country. The messages originating from Alice should be transmorphed into a series of Twitter post belonging to her.

The identities used in Twitter are highly identifiable labels, with a certain trust level. This hard identity with millions of followers is a stark contrasts with anonymity. Current anti-censorship technology lacks the ability to first have stealth encrypted transport of news, cross the freedom/non-freedom border and then transmorph this news into a public accessible form with a highly identifiable label.

6.5. A single global conversation

Existing technologies, such as [\[TOR\]](#) in combination with XMPP or the Orbot smartphone app facilitate protected point-to-point communication. However, a desired scenario is to facilitate more current the Twitter-like social media practices, best typified as a "global conversation".

Furthermore, current social media revolves around video-rich, real-time interaction with groups, hashtag-based discovery and social networking. All of these aspects are not offered or are incompatible with current-generation of privacy enhancing technology. More knowledge is needed about reputation models in news reporting and information flows. In the current microblogging age, can the number of real-person followers be seen as your reputation? Do several news sources of moderate reputation which report the same news story yield together an increased reputation score?

This work should combine privacy enhancement with microblogging.

6.6. Spammers and hoaxes

This final scenario is focused on spam. All technology addressing one of the above scenarios **MUST** also have the capability to deal with spam. Unfortunately, this ability to deal with spam is in conflict with simplicity.

Alice and Bob are exchanging the fresh messages from their social network (similar to Internet-free or Friends-only). Eve is actively trying to disrupt the system by injecting news channels with a mix of genuine news, obviously fake messages (consuming valuable system resources and user attention) and hoaxes. These falsehoods made to masquerade as truth result in erosion of overall trust in the system.

Systems **SHOULD** offer capabilities to report spam, mechanisms for fact validation and reputations of (pseudo) identities.

7. Design principles: simplicity and prior success

Designing and crafting software which is completely self-organizing has clear limits [[CAPLIMIT](#)] and requires a certain level of expertise [[LEVELS](#)]. In order to avoid repeating mistakes from the past, this document aims to base its design principles on existing new media successes. For microblogging this means following market leading solutions and enhance them with censorship resilience. We recognize the following success factors: Simplicity, Real-time responsiveness, Near-effortless news creation, News items are bundled in channels, combine public broadcasting and person-to-person private messaging, following a channel is single direction, more followers yields more visibility, keyword search with push of updates and ability to deal with spam.

8. Background rant: lack of coordination and fragmentation

Computers communicating on equal footing has been part of the IETF standards for many decades. Recently several loosely connected standard initiated around explicitly driven by the P2P paradigm for applications such as Internet telephony video streaming. An essential problem in this domain is the lack of coordination and standard setting for P2P technology. A large part of the innovation around P2P seems to happen in single-person Open Source projects and small groups which lack the engineering capacity to make generic, re-usable and documented components. Given their running code-driven nature, money and time is not available for attending standards-setting meetings, writing formal specifications and defining quality control testing suites. Profit-driven organizations should have the resources to overcome these resource shortage issues. However, due to the dynamic, disruptive and litigious nature of P2P few examples

exist of companies which are capable of supporting an IETF standard setting activity for several years.

As presented during IETF 81 area directorate, there is "not a clear long-term architecture yet for you to build actual classes of P2P applications using IETF technologies". Forming an overlay is hard and scalable privacy-preserving unstructured search solutions are only barely out of the scientific research community.

From the above we conclude that a key obstacle to the success of this proposal is implementation and uptake. A draft document, active community and reference implementation ideally evolve together over time. To overcome this issue a continuous incremental improvement approach is advised. The preferred way is incremental development of single a reference implementation, based on free software.

9. Current running code and related work

DISCLAIMER: this section needs significant expansion and listing of projects with running code and self-organization.

Several Open Source projects have running code and partially implemented the above four scenarios. We will briefly list them here.

[TOR] A free software implementation of second-generation onion routing, a mechanism enabling users to communicate anonymously over the Internet. This flagship project has boosted online anonymity for over a decade and is the key example for the cat and mouse dynamics of privacy/surveillance technologies. The Orbot project provides an Android implementation of Tor. Due to the usage of the client server/model, exit node principle plus lack of reputations this architecture is not compatible with our scenarios.

[DIASPORA] A free personal web platform implementing a distributed social networking service. This partially operational system is based on a client/server model and thus not compatible with our ad-hoc scenarios.

[BRIAR] Briar is a secure news and discussion system designed to be used by journalists, activists and civil society groups in authoritarian countries. Briar differs from existing circumvention tools and mesh networks in three significant ways: needs no external infrastructure, can operate over any mixture of available media and builds on social relationships. The aims of this project are similar to our scenarios, but this project lacks running code and has few active developers.

[BUBBLE] DTN researchers have simulated closely related scenarios. Dissemination in the Arab Spring scenario is likely to involve an explicit copy between people who trust each other, referred to as social-based forwarding in this study.

[TWIMIGHT] The Twimight project by ETH-Zurich university shows that decentralized microblogging already exists. Researchers developed an Android application that uses Twitter servers in normal conditions, but switches to a Bluetooth-based disaster mode when Internet connectivity is lost.

[MUSUBI] The Musubi smartphone app represents another key, censorship-free, technology advancement. Developed by Stanford University, it offers instant messaging service and media sharing capabilities similar to WhatsApp, Ping, and Blackberry Messenger. What makes it unique is that all data and processing resides on the smartphones, not in the cloud. This decentralization removes the need for central processing and provides significant decoupling from the underlying infrastructure. Exchange of cryptographic keys is integrated in the friending process--Musubi essentially builds a decentralized social graph. Unfortunately, Musubi is also limited--all data transfers go through central servers, as it lacks NAT-traversal capability.

[TRIBLER] DISCLAIMER2: this project is coordinated by the author. This project has created Open Source firmware for a Samsung Internet-connected television which gives it the ability to find, share and stream news videos within a fully self-organizing overlay; operated only by remote control [[REBELLIONTV](#)]. It is also available as generic zero-server file sharing software for the PC which has been installed by 1.2 million users. It uses the Dispersy elastic database for providing: keyword search, content discovery, content voting and spam prevention using crowd sourcing [[DISPERSY](#)]. For swarm-based streaming and generic message transport it uses the IETF protocol developed within the PPSP working group, called Libswift [[LIBSWIFT](#)]. All this code is created by a single team and specifically designed to facilitate evolution into the prior described scenarios. An Libswift demo streaming app is available on the Android market.

10. Open issues

Deliverables planned and issues which need to be addressed.

TODO: ADD REF Privacy definition:

<http://tools.ietf.org/html/draft-iab-privacy-terminology-01>

TODO: REF

<http://www.ietf.org/id/draft-iab-privacy-considerations-03.txt>

TOD0: <http://datatracker.ietf.org/doc/search/> P2P

TOD0: <http://datatracker.ietf.org/doc/rfc4981/> SEARCH survey

TOD0: <https://datatracker.ietf.org/doc/draft-ietf-p2psip-reload/>

[10.1.](#) Use cases and threat model

[10.2.](#) System components, definitions and system architecture

[10.3.](#) Current technology and gap

[10.4.](#) Detailed system design and protocol specification

[11.](#) Security Considerations

tbd.

[12.](#) IANA Considerations

tbd.

[13.](#) References

[13.1.](#) Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[13.2.](#) Informative References

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", [BCP 127](#), [RFC 4787](#), January 2007.

[RFC4838] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, "Delay-Tolerant Networking Architecture", [RFC 4838](#), April 2007.

[13.3.](#) URL References

[ATTACK] <https://datatracker.ietf.org/doc/draft-farrell-perpass-attack>, "Pervasive Monitoring is an Attack".

- [CATALOG] <http://www.spiegel.de/international/world/catalog-reveals-nsa-has-back-doors-for-numerous-devices-a-940994.html>, "Shopping for Spy Gear: Catalog Advertises NSA Toolbox".
- [SCHNEIER] <http://www.theguardian.com/commentisfree/2013/sep/05/government-betrayed-internet-nsa-spying>, "The US government has betrayed the internet. We need to take it back".
- [YOUTH] <http://www.nytimes.com/2009/01/25/magazine/25bloggers-t.html>, "Revolution, Facebook-Style".
- [FORCEDLOGIN] <http://online.wsj.com/article/SB125978649644673331.html>, "Iranian Crackdown Goes Global".
- [NKOREA] http://audiencescapes.org/sites/default/files/Report_Summary_Quiet_Opening_North%20Korea_InterMedia.pdf, "A QUIET OPENING: NORTH KOREANS IN A CHANGING MEDIA ENVIRONMENT".
- [EGYPTSTUDY] <http://conferences.sigcomm.org/imc/2011/docs/p1.pdf>, "Analysis of country-wide internet outages caused by censorship".
- [OPENMICRO] <http://xmpp.org/extensions/xep-0277.html>, "XEP-0277: Microblogging over XMPP".
- [DATING] <http://www.washingtonpost.com/wp-dyn/content/article/2006/08/05/AR2006080500930.html>, "Saudi Youth Use Cellphone Savvy To Outwit the Sentries of Romance".
- [TREND] <http://www.google.com/trends/?q=bluetooth+transfer>, "Google Trends query".
- [BOOTSTRAP] <http://grothoff.org/christian/dasp2p.pdf>, "Bootstrapping Peer-to-Peer Networks".
- [ZEROKNOW] http://www.cse.ust.hk/~liu/luli/PT_Trans_final.pdf, "Pseudo trust: Zero-knowledge based authentication in anonymous peer-to-peer protocols".
- [CAPLIMIT] <http://doi.ieeecomputersociety.org/10.1109/MC.2012.54>, "The CAP Theorem's Growing Impact".
- [LEVELS] <http://blog.incubaid.com/2012/03/28/the-game-of-distributed-systems-programming-which->

level-are-you/, "The Game of Distributed Systems Programming. Which Level Are You?".

- [TOR] <http://www.torproject.org>, "Tor Project: Anonymity Online".
- [DIASPORA] <http://diasporaproject.org/>, "Diaspora is a fun and creative community that puts you in control.".
- [BRIAR] <https://fulpool.org/btp.pdf>, "Secure communication over diverse transports".
- [BUBBLE] <http://dx.doi.org/10.1109/TMC.2010.246>, "BUBBLE Rap: Social-Based Forwarding in Delay-Tolerant Networks".
- [TWIMIGHT] <http://dl.acm.org/citation.cfm?id=2159576.2159601>, "Twitter in disaster mode: smart probing for opportunistic peers".
- [MUSUBI] <http://dl.acm.org/citation.cfm?id=2187866>, "Musubi: disintermediated interactive social feeds for mobile devices".
- [TRIBLER] <http://dl.acm.org/citation.cfm?id=2206767>, "Tribler: P2P search, share and stream".
- [REBELLIONTV] <http://www.tribler.org/trac/wiki/SwiftTV>, "RebellionTV a.k.a. Libswift on a television project".
- [DISPERSY] www.frayja.com/pub/dispersypaper2012.pdf:donotdistribute, "Dispersy elastic database".
- [LIBSWIFT] <http://www.libswift.org>, "IETF PPSP streaming protocol implementation".

Author's Address

Johan Pouwelse (editor)
Delft University of Technology
Mekelweg 4
Delft
The Netherlands

Phone: +31 15 278 2539
EMail: J.A.pouwelse@tudelft.nl

