Extended Abstract: Nobody's Fault but Mine: Using Unauthenticated Unidirectional Pushes for Client Update

Paul Vines
Two Six Technologies
Arlington, Virginia, USA
paul.vines@twosixtech.com

Abstract

We identify a specific anticensorship task, configuration and software update, and define a new concept, Unauthenticated Push (UP) channels, which facilitate this task while minimizing user burdens and risks. We describe an example implementation using steganographic videos posted on public streaming services and outline how this and other existing systems align with the UP definition.

Keywords

censorship, steganography, anonymity

1 Introduction

Circumventing internet censorship is an ongoing contest between users and censors: currently, rapid updates of anticensorship tools and configurations enables users to circumvent censorship via efficient "direct" connections to proxies. These connections are successfully facilitated by a number of high-bandwidth / low-latency anticensorship channels like Obfsproxy[1], Snowflake[3], and others. However, two significant challenges remain: updating proxy information and updating circumvention software itself. The former is the "bridge distribution problem" or "rendezvous problem" where censors continually enumerate active proxies and block them by IP, necessitating proxy IPs to be rotated and conveyed back to users to re-enable access. For instance, when the Great Firewall deployed entropy-based blocking in 2022, updates to fully-encrypted circumvention protocols had to be deployed before users could regain access to the open internet[13]. Recent work on *signaling channels* (alt: *rendezvous channels*) have identified the importance of the rendezvous problem in modern circumvention research and developed a number of novel lower-bandwidth and/or higher-latency channels that use various third-party services to avoid IP-based censoring methods[8, 10–12, 14]. These channels and signaling protocols they are used for typically involve a 1-3 round-trip handshake of relatively small (kilobyte) messages to facilitate on-demand bootstrapping of new proxy connections.

We identify a related-but-distinct concept, *unauthenticated push (UP) channels*. These overlap with many signaling channels in that they use third-party services as intermediaries to deliver content to censored users without being blocked by the censor. However, they differ from (typical) signaling channels in three ways:

(1) They are specifically unidirectional
(2) They require no authentication or account association from the client-side
(3) They are higher-bandwidth to support operations like software updates rather than minimal rendezvous operations

The first two properties are aimed at enabling a model of operations emphasizing *minimal risk and burden on users*, approaching *receiver anonymity* (depending on threat model assumptions) and minimizing user intervention. This is deliberately traded off against increasing the risk and burden on the sending side, by requiring them to manage operational security concerns to maintain resilience (e.g. cycling accounts and avoiding bot-like behavior).

2 Example Model: Public Video Steganography

A well-motivated example implementation of the UP channel archetype is posting of steganographically encoded videos on a publicly accessible video hosting service, like flickr.com. The underlying transport mechanism used is a keyed steganographic encoding that requires a, relatively, expensive decryption operation to detect the encoded nature of a video. The users are required to possess a copy of the steganographic decoding algorithm and key, as part of a bootstrapping process. *However*, the design assumes both this software and key material are *public* and known to the adversary. This aligns with prior definitions of signaling channels as starting from *publicly addressable* connections that require no shared secrets[11].

The main points of interest for this system surround the *addressing* of the encoded videos: how do users automatically find them and how are censors prevented from blocking them. Users find videos using a time-based pseudorandom tag generator. For each time "epoch," a tag is generated from a publicly known list of possible tags (e.g. lists of trending topics). This tag is searched for, and the top n videos are queried and attempted to be decoded. Video steganography varies in efficiency, but real video durations also vary significantly. This variability means video steganography can suit data transmissions from a few kilobytes for small configuration updates (i.e. a unidirectional "signaling" use) to megabytes (e.g. anticensorship software updates).

2.1 Adversary Analysis

An adversarial analysis of this hypothetical system can observe several potential threat surfaces requiring varying levels of sophistication and access to coercion or control over the third-party service:

(1) Blocking the third-party flickr.com service as-a-whole
(2) Mass-scanning and blocking encoded videos
(3) Enumerating and blocking videos using the pseudorandom tags
(4) Banning accounts posting encoded videos

Free and Open Communications on the Internet 2025(2), 43–45

(5) Tracking anticensorship users receiving encoded content

Completely blocking the service is a blunt but effective approach that holds for most "indirect" anticensorship tools (including recently published signaling channels). The argument against it is traditionally based on adversaries attempting to avoid collateral damage of denying the service to legitimate users. In this case, loss of access to flickr.com specifically may not be a significant deterrent to many censoring countries (although we note it is not currently blocked in many censored regions). However, many other videosharing serviecs exist, and availability of streaming video content *is* an expectation of many countries. If major international videosharing platforms are blocked, the country is likely to develop its own (e.g. Asparat in Iran[2] which blocks access to mainstream streaming platforms).

Untargeted scanning and blocking videos is ultimately a steganographic algorithm concern, which we do not explore but note that the field continues to advance and mass scanning is generally cost-prohibitive[7]. Using the pseudorandom tags to reduce the set of content that needs to be scanned can make this feasible - effectively emulating the users of the system. However, the tags are selected to include those used by unassociated videos, which may change over time: the top-n videos for a given tag at time t are not necessarily the same videos at times *t-1* or *t+1*, thus the adversary cannot *pre-emptively* search for the videos, but also cannot be sure the relevant videos will be freshly uploaded at the relevant timewindow.

Attacks on the senders (video uploaders) and receivers (users) are really where we see the UP archetype properties represented by this model system. If videos were always uploaded by the same account, it would reduce the argument for resilience to a question of whether the adversary can exert sufficient influence over the third-party service to ban a single account. Similarly, users could feasibly be targeted if video viewing data was gathered and the videos were sufficiently unpopular that most viewers of the account were anticensorship users. In the past, other systems have handled these tradeoffs by moving towards authenticated viewing to reduce adversary observation at the cost of imposing burdens of signup and direct linkability to the user: e.g. using private messaging services like email or Telegram for bridge distributions[9, 12].

Instead, an UP channel design pushes additional burdens on the sender to improve security while avoiding costs for the users: we assume the sender is capable of more sophisticated behaviors than maintaining a single account: they can maintain a stable of accounts, organizing creating new ones as necessary and mixing non-steganographic videos. The effect is to impose an *uneven* burden resting on the sender, who has greater resources. Similarly, the model uses unauthenticated access requests from the users to both reduce traceability and burden by not requiring account signups from users.

3 Unauthenticated Push Channels Broadly

The UP channel concept represents a lens through which to understand anticensorship techniques: it emphasizes making burden tradeoffs in favor of the user, requiring no authentication or account signup actions by the users, and providing resilience to censorship even when *publicly addressed*. This overlaps with some current signaling channels but not others: Skyhook[10], PushRSS[14], SQS[8], AMPCache[3] and Meek[5] all satisfy these properties, and some (Meek and Skyhook) can clearly support transmission of larger content transfers beyond small signaling exchanges. Others, such as Raven[12] and CloudTransport[4] do not satisfy these properties because they require the users to maintain an account for authenticated access to the service infrastructure. Looking more broadly, the email- and telegram-based bridge request paths used by Tor also do not satisfy these properties, since they require user authentication via Telegram or a specific email provider to function.

4 No-Cost Scalability

Another property of the model channel described above, and similar approaches making use of *uncharged* third-party content hosting, is the ability to increase data delivery and sustain adversarial denial-of-service (DoS) attacks without incurring monetary costs. Steganographic videos can easily reach megabytes of data transmission, and once uploaded are effectively free for the sender to distribute to any number of users. This contrasts with other channels like meek, SQS, AMPCache, and Skyhook[3, 5, 8, 10] that have major concerns with adversary sybil attacks that use the system as-intended to drive up hosting costs (i.e. financial denial-of-service).

Extending to other forms of steganographic content can receive the same types of protection: aggregating and serving user created content for free is the business model of many popular services, including audio, pictures, and more niche content like machine learning models[6]. All of these can provide the "no-cost scalability" offered by the video example system so long as they support some level of queryable indexing and a suitable steganographic encoding exists.

5 Future Work

From this definition of UP channels and example model system, we plan to implement this scheme in several parallel steps: first, adapting existing steganography and service interaction libraries to provide a working UP channel based on video steganography on flickr.com. This can empirically demonstrate the limits of current steganographic encodings, refine the pseudorandom address generation scheme, and enable a holistic security analysis of the system relative to modern censor behaviors and capabilities. In parallel, we will refine the theoretical definition of UP channels in relation to existing services and identify additional candidate services and content types that are suitable for use.

6 Conclusion

In this extended abstract we have laid out an argument for defining a specific anticensorship channel concept: the Unauthenticated Push (UP) channel. This conceptual definition emphasizes minimizing burdens and risks on receivers (users) and supporting specific anticensorship tasks like configuration (e.g. proxy distribution) and software updates by using public broadcast channels and increased sender operational security measures. We explored a hypothetical design of a video steganographic channel and observed how it fulfilled the UP criteria while remaining secure and resilient.

References

[1] Yawning Angel. 2023. obfs4: The obfourscator. https://gitlab.com/yawning/obfs4 Access on 5/30/2023.

[2] Frud Bezhan. 2012. Iran's 'YouTube' Stumbles Out Of The Gate. https://www.rferl.org/a/irans-youtube-mehr-video-sharing-site/24796887.html

[3] Cecylia Bocovich, Arlo Breault, David Fifield, Xiaokang Wang, et al. 2024. Snowflake, a censorship circumvention system using temporary {WebRTC} proxies. In *33rd USENIX Security Symposium (USENIX Security 24)*. 2635–2652.

[4] Chad Brubaker, Amir Houmansadr, and Vitaly Shmatikov. 2014. CloudTransport: Using Cloud Storage for Censorship-Resistant Networking. In *Privacy Enhancing Technologies Symposium*. Springer. https://petsymposium.org/2014/papers/paper_68.pdf

[5] David Fifield, Chang Lan, Rod Hynes, Percy Wegmann, and Vern Paxson. 2015. Blocking-resistant communication through domain fronting. *Privacy Enhancing Technologies* 2015, 2 (2015). https://www.icir.org/vern/papers/meek-PETS-2015.pdf

[6] huggingface-pricing 2025. HuggingFace Pricing. https://huggingface.co/pricing Accessed on 4/21/2025.

[7] Jayakanth Kunhoth, Nandhini Subramanian, Somaya Al-Maadeed, and Ahmed Bouridane. 2023. Video steganography: recent advances and challenges. *Multimedia Tools and Applications* 82, 27 (2023), 41943–41985.

[8] Michael Pu, Andrew Wang, Anthony Chang, Kieran Quan, and Yi Wei Zhou. 2024. Exploring Amazon Simple Queue Service (SQS) for Censorship Circumvention. *Free and Open Communications on the Internet* (2024).

[9] Tor 2022. Request a bridge from Telegram. https://forum.torproject.org/t/request-a-bridge-from-telegram/3277/2 Accessed on 4/21/2025.

[10] Paul Vines. 2024. Ten Years Gone: Revisiting Cloud Storage Transports to Reduce Censored User Burdens. *Free and Open Communications on the Internet* (2024).

[11] Paul Vines, Samuel McKay, Jesse Jenter, and Suresh Krishnaswamy. 2024. Communication Breakdown: Modularizing Application Tunneling for Signaling Around Censorship. *Privacy Enhancing Technologies* 2024, 1 (2024). https://www.petsymposium.org/popets/2024/popets-2024-0027.pdf

[12] Ryan Wails, Andrew Stange, Eliana Troper, Aylin Caliskan, Roger Dingledine, Rob Jansen, and Micah Sherr. 2022. Learning to Behave: Improving Covert Channel Security with Behavior-Based Designs. *Proceedings on Privacy Enhancing Technologies* 3 (2022), 179–199.

[13] Mingshi Wu, Jackson Sippe, Danesh Sivakumar, Jack Burg, Peter Anderson, Xiaokang Wang, Kevin Bock, Amir Houmansadr, Dave Levin, and Eric Wustrow. 2023. How the Great Firewall of China detects and blocks fully encrypted traffic. In *32nd USENIX Security Symposium (USENIX Security 23)*. 2653–2670.

[14] Diwen Xue and Roya Ensafi. 2023. The Use of Push Notification in Censorship Circumvention. *Free and Open Communications on the Internet* (2023).