Security for the High-Risk User: Separate and Unequal[1]

John Scott-Railton | Citizen Lab, Munk School of Global Affairs, University of Toronto

Abstract

The constant cyberattacks against governments and industry are widely known. Further from the public eye are the targeted attacks against civil society groups. These attacks aren't opportunistic online targeting—or cybercrime— but politically driven campaigns intended to disrupt, degrade, or steal civil society groups' private information. Many occur via the popular online platforms that under-resourced organizations use in place of the more expensive managed IT environments preferred by other sectors. Although attacks in any sector are costly, attacks against civil society often have much greater ramifications, including threats to life and liberty.

Digital threats against civil society deserve your attention—first, because of what they reveal about the default-insecure options in popular online platforms and, second, because addressing the most glaring cases will confer security for broader user populations.

Attacks against civil society

As civil society groups increasingly move their activities online, the Internet's historically ascribed promise to empower, connect, and enable is being demonstrated. Indeed, civil society is a poster child for the creative use and repurposing of inexpensive and free platforms for public advocacy. Out of the limelight, the same transformation has occurred operationally: technology has enabled geographically diverse movements to coalesce, exchange information, and marshal resources.

---

[1] Note, this document is slightly different from the accepted version. The IEEE paper can be found and cited at **Citation**: Scott-Railton J Security for the High Risk User: Separate and Unequal. I*EEE Security & Privacy*, Issue No.02 – Mar.-Apr. (2016 vol.14).

Yet the capacity to connect has vastly outpaced the ability to secure. Civil society groups tend to use commodity tools and popular online platforms. However, these platforms are built for mass adoption and don't include default settings that would naturally protect the data of higher-risk government or corporate users. For example, whereas autosharing makes online platforms user-friendly, the lack of default two-factor security (typically complementing a password with a second factor such as a short message service [SMS], or a one-time code) makes these same platforms precarious for higher-risk users. Moreover, the way that many civil society groups use technology extends beyond the managed environments, networks, and endpoints of private sector or government institutions. In addition, civil society groups often rely on volunteer or nonspecialist information technology talent to manage their resources.

> *"Digital threats against civil society deserve your attention...because of what they reveal about the default-insecure options in popular online platforms and...because addressing the most glaring cases will confer security for broader user populations."*

Civil society groups have been targeted to de-anonymize and jail or murder dissidents in the Middle East and Asia, threaten victims of crimes against humanity on several continents, and disrupt the moral stand of courageous citizens worldwide. Even organizations engaged in seemingly apolitical activities such as disaster relief might be targeted for the sensitive personal information they collect about disaster victims.

The Citizen Lab

Based in the Munk School of Global Affairs at the University of Toronto and directed by Professor Ronald J. Deibert, the Citizen Lab (https://citizenlab.org) is an interdisciplinary research laboratory that applies rigorous technical methods and policy analysis to understanding threats to a free and secure Internet.

Meanwhile, many adversaries— having noted the dropping costs to target, disrupt, and spy on civil society organizations—are turning to targeted digital attacks as a cheap way to monitor or disrupt a particular group's activities. Even when repressive regimes control networks, their targets might be outside their borders or might use encryption to frustrate passive interception, making targeted attacks very a ractive.

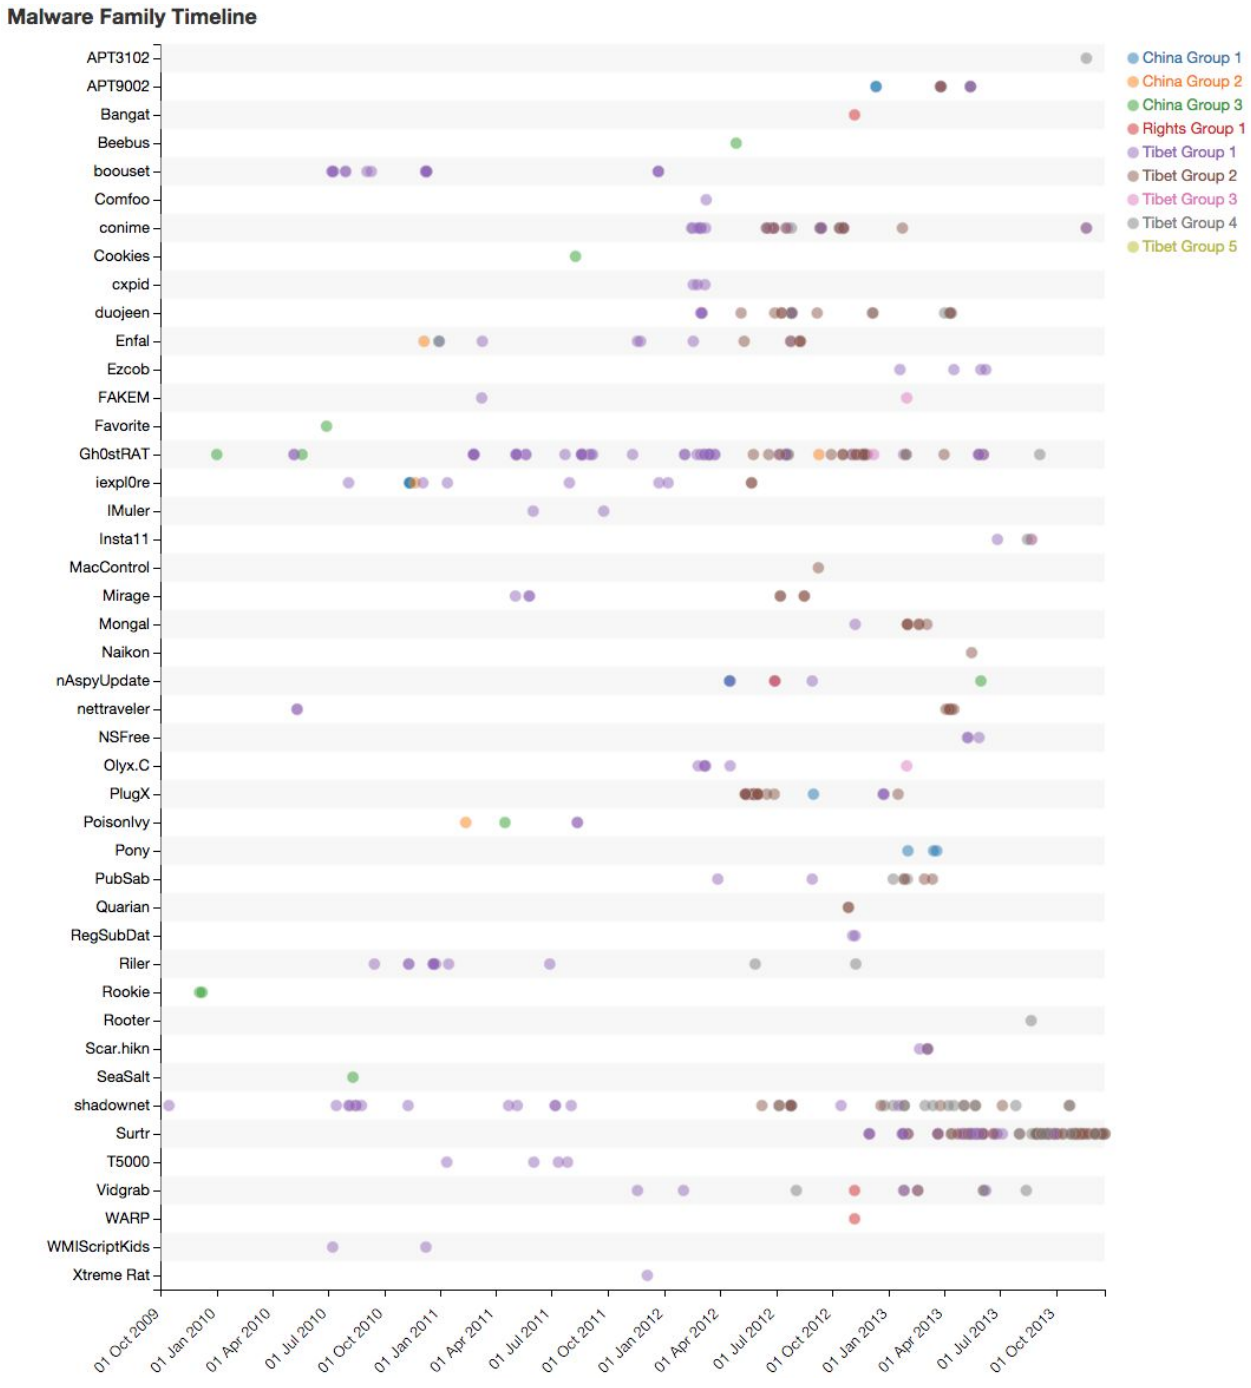

Fig 1. Malware Families used to target 9 groups enrolled in our Targeted Threats project (Tibet Groups are diaspora based, and focus on Tibetan issues, while China Groups focus on other China-related topics)

One need not look far to see evidence of this phenomenon; consider, for example, the recent conflict in Syria. The Citizen Lab (https://citizenlab.org) and other researchers have

documented successful attacks against aid, charity, media, human rights, and other organizations engaged in the Syrian context.[2]

In case after case, my Citizen Lab colleagues and I have observed cyberattacks leading not only to breaches of sensitive information but also to fear, distrust, and a sense of hopelessness among the targeted groups. In extreme cases, we've documented how attacks have contributed directly to arrests, detentions, and targeted killings.

A violent conflict isn't a prerequisite for an aggressive campaign. For example, we've documented more than 10 years of malware attacks against aid, pro-democracy, and development organizations concerning China and Tibet.

Moving away from anecdotes, our recent four-year study systematically examined and analyzed targeted attacks against 10 civil society groups.[3] During this period, we obtained more than 2,800 malware samples representing 44 distinct malware families, all used in targeted attacks (see ***Figure 1***).

What Do the Attacks Look Like?

The malware attacks against civil society groups that we examined have shown widely varying levels of sophistication in social engineering. ***Figure 2*** displays a relatively unsophisticated attack, one that provides little incentive beyond curiosity and subject relevance for a victim to open it.

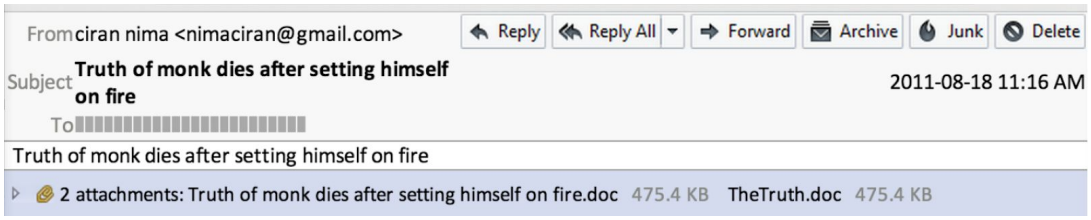

Fig. 2 Example of relatively unsophisticated social engineering. It provides little incentive beyond curiosity and subject relevance for a victim to open it

[2] Marczak B, Scott-Railton J, Marquis-Boire M & Paxson V 2014 When Governments Hack Opponents: A Look at Actors and Technology, 23rd USENIX Security Symposium, August 20-22, San Diego, CA.

[3] Crete-Nishihata M, Dalek J, Deibert R, Hardy S, Kleemola K, McKune S, Poetranto I, Scott-Railton J, Senft A, Sonne B, & Wiseman G 2014 Communities @ Risk: Targeted Digital Threats Against Civil Society. Citizen Lab, University of Toronto.

Other attacks are much more sophisticated. We often find evidence that the attackers had previous access to the sensitive internal communications of groups or informal networks of individuals.

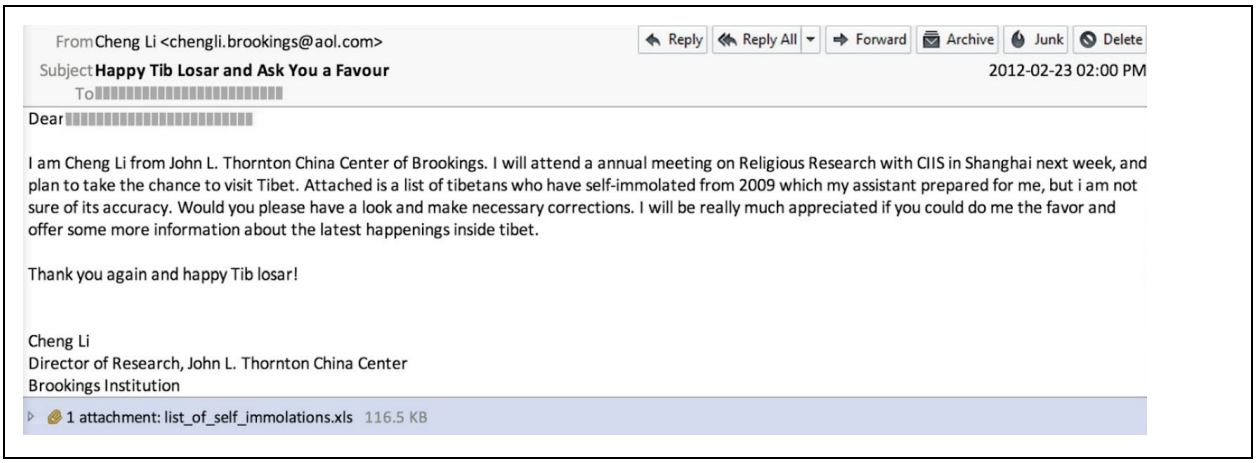

Fig 3. Example of a more sophisticated, personalized malware attack. The attackers likely used previously stolen materials to make this message both convincing and enticing to the recipient

For example, ***Figure 3*** shows an example of an attack that was personalized for the Tibetan activist who received it, appeared to come from a well-known scholar of contemporary China, and referenced a very important topic for the Tibetan human rights community. The attackers had done their homework and spent time crafting this message to be both convincing and enticing to the recipient. Although this attack isn't the most sophisticated social engineering we've observed, it's among the most advanced examples that we can share without revealing highly personal information about the targets.

Common Theme: "Just Enough" Technical Sophistication

Civil society is chronically under-resourced, often relying on unmanaged networks and endpoints, combined with extensive use of popular online platforms. Security, if any, is usually conferred only through behavioral precautions and the use of commodity antivirus products. This exceptionally low baseline provides a target-rich environment for attackers.

Citizen Lab research has consistently found that although the overall technical sophistication of attacks is typically low, their social engineering sophistication is much higher. This finding is borne out even when the same attack group has used more sophisticated attacks against other targets. For example, in four years of tracking attacks, we found only one zero-day exploit,

while the vulnerabilities used were often legacies that remained active (for example, the widely used CVE-2012-0158, which was active through 2015).[4]

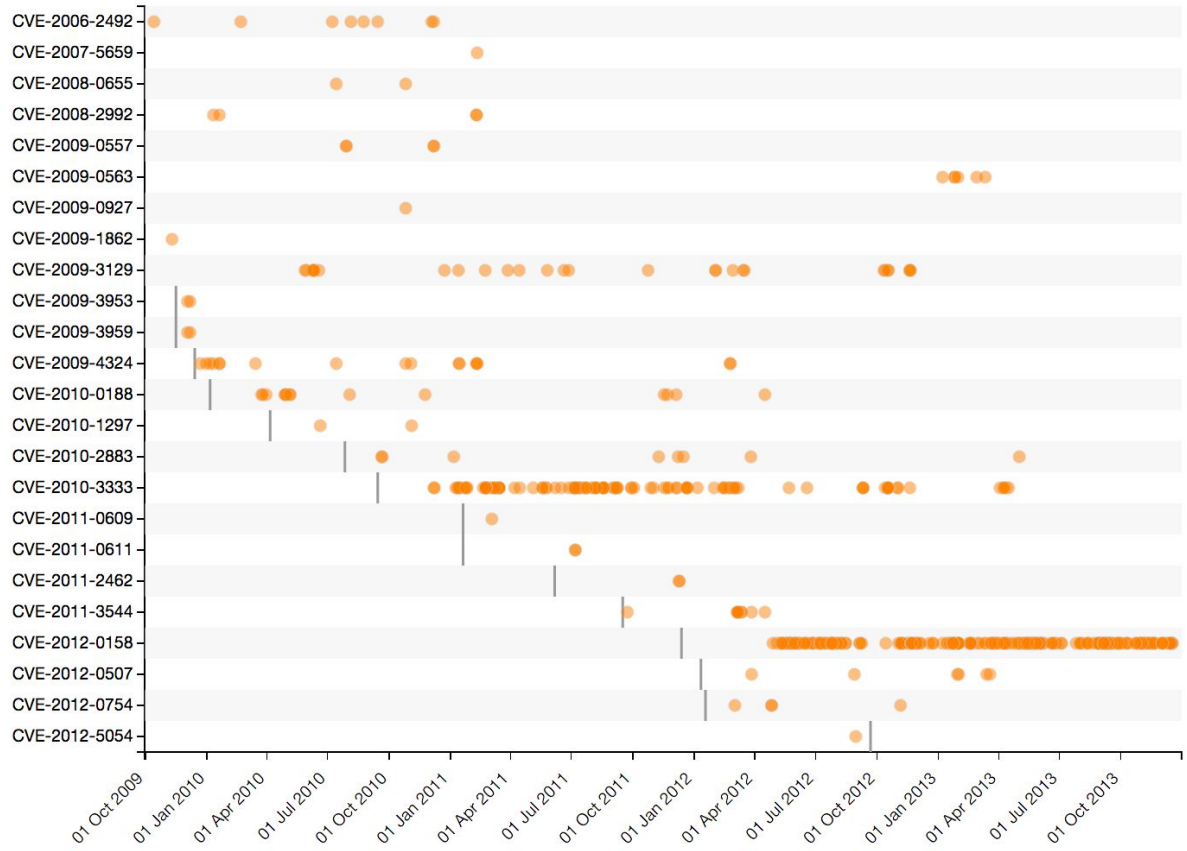

Figure 4. The long 'shelf life' of vulnerabilities used against civil society groups. Lines indicate date of Common Vulnerabilities and Exposures (CVE) announcement. Dots indicate use during campaigns.

Although this finding might indicate resource constraints on the part of attackers, it also suggests that the attacks continue to work for long periods (see ***Figure 4***). Economizing on sophisticated technology is rational: why "burn" a zero-day to gain access to a targeted organization if a simple remote access tool (RAT) will do?

Not Powerless, Just Outmatched

Civil society isn't powerless in the face of sustained cyberattacks. In Tibet, for example, a multiyear campaign in the diaspora community encouraged Tibetans to "detach from a achments," a humorous play on a Buddhist concept that served as a vehicle to encourage an

[4]Katie Kleemola, Masashi Crete-Nishihata, and John Scott-Railton. Targeted Attacks against Tibetan and Hong Kong Groups Exploiting CVE-2014-4114. Citizen Lab, June 15 2015.

important security behavior. A common exhortation was to not send files as attachments, but rather share them via Google Docs. This tactic conferred a degree of protection, because files opened in Google Docs can be previewed in the cloud and viewed without risk to users. The Citizen Lab has observed the recent growth of Google Docs use among civil society groups. attackers apparently have as well. In the past year, we've documented an increasing number of attacks being sent via Google Docs, apparently in direct response to this behavioral change by users.[5]

A second, exceptionally important security practice is also gaining ground among those who provide security advice to civil society groups: use two-factor authentication.[6] In some communities where hacking against civil society is especially widespread, behaviors appear to be changing, and adoption of two-factor authentication is increasing.

Attackers have also noticed two factor authentication's growth and are creatively working to circumvent it. In a 2015 report, we documented an elaborate a empt to socially engineer the two-factor authentication used by Iranian freedom of expression and prodemocracy groups (see the "Faking Two-Factor Authentication").[7]

Faking Two-Factor Authentication

On a summer morning, an Iranian media organization worker awoke to a text message alert—purporting to be from Google—warning him of an unexpected sign-in attempt to his Gmail account. Minutes later, an email, also purporting to come from Google, warned him that someone in Iran was attempting to access his account.

The email contained a link to a "change password" dialog box that was prepopulated with his username, name, and user avatar. He was instructed to enter his old password as well as a new one. Entering a password on this page yielded a request for his two-factor code (see **Figure A**). You might have already guessed how the attack was supposed to work. The attacker was presumably monitoring the page in real time. As soon as the victim entered his password, the attacker would use the credential to log in to the genuine Gmail login. This would generate a two-factor code sent to the victim's phone, further cementing the fiction. Entering the code in the website would provide the attacker with the second factor—and access to the account.

---

[5] Ibid.

[6] Iulia Ion, Rob Reeder & Sunny Consolvo No One Can Hack My Mind. USENIX Eleventh Symposium on Usable Privacy and Security, July 22-24, 2015 Ottawa, Canada.

[7] Scott-Railton, J & Kleemola K. London Calling: Two-Factor Authentication Phishing From Iran, Citizen Lab, August 27 2015.

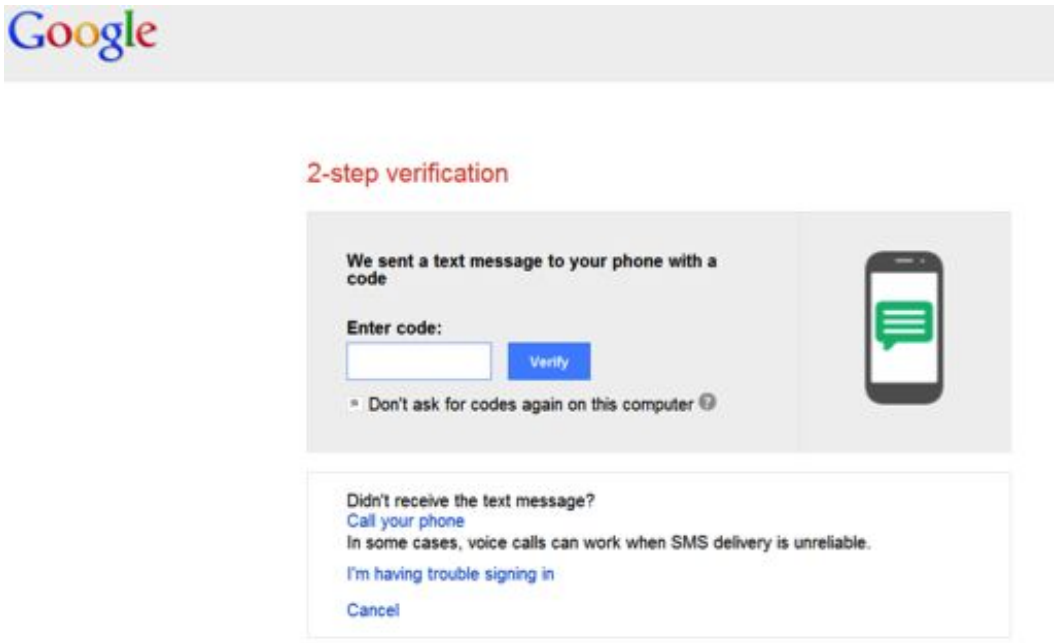

Figure A. Fake two-factor prompt. The phishing victim has just received a two-factor code from Gmail, triggered by the hacker's login. The hacker now waits for the victim to enter the twofactor code—and for access to the account.

Who Are the Attackers?

To better understand and mitigate risks for civil society, we must better understand the threat actors. The Citizen Lab's research regularly identifies three styles of progovernment electronic actors (PGEAs), each with a distinct set of skills and resources (see **Table 1**). Although each category's composition, resources, and skill sets vary widely, actors from each have proven to be formidable threats to the activities of civil society groups. These categories emphasize how techniques are obtained; however, the categories are sometimes more fluid in practice, with groups using tools from different categories depending on their target.

Table 1. Three styles of Pro Government Electronic Actors (PGEAs)

PGEA	Skills and resources
Nation State / Advanced Persistent Threat (APT)	Well-resourced efforts with access to advanced in-house capabilities and a wide range of attack avenues.

Purchasers of "lawful intercept" malware	Governments and agencies that use malware and exploits purchased from commercial providers, including FinFisher and Hacking Team
Cyber Militias	Use and re-purposing of Commercial Off The Shelf (COTS) Remote Access Tools (RATs), phishing, and other simple techniques. Varying ability to use exploits, rarely in possession of zero-days.

Nation-state/advanced persistent threat. The most well-known and perhaps highest-profile PGEAs are the well-resourced and highly talented nation-state groups. These groups maintain a substantial in-house capability and typically operate from countries with well-developed science, technology, engineering, and mathematics (STEM) fields. Nation-state groups in, for example, China, Russia, and North Atlantic Treaty Organization countries, fall into this category. Targeted malware developed by these groups is often carefully crafted and obfuscated, and is designed to frustrate analysis and a ribution. The Citizen Lab has documented several cases of these threat actors, especially Chinese advanced persistent threat groups, targeting civil society groups as far back as 2009.[8]

Purchasers of "lawful intercept" malware. Not all countries or agencies can develop in-house malware. The Citizen Lab and others have reported on the troubling global proliferation of so-called "lawful interception" remote intrusion capabilities to governments and agencies. These technologies include tools like FinFisher and Hacking Team, which enable inexperienced users to quickly generate sophisticated implants. Such malware not only includes extensive remote access and administration capabilities but is also often built with obfuscation and anti-debugging features. Citizen Lab research, as well as a series of breaches, and investigative reporting, have contributed to identifying a growing range of countries, including Ethiopia, Bahrain, and Morocco, that are using these tools against pro-democracy activists, journalists, and dissidents.[9]

Cybermilitias. Conducting a targeted malware campaign doesn't require millions of dollars or decades of STEM investment. While studying conflicts ranging from Libya[10] and Syria,[11] to

[8] Infowar Monitor Tracking GhostNet: Investigating a cyber espionage network. Information Warfare Monitor. University of Toronto,Toronto, 2009.

[9] Bill Marczak, John Scott-Railton, and Sarah McKune, Hacking Team Reloaded? US-Based Ethiopian Journalists Again Targeted with Spyware. CitizenLab, March 9 2015.

[10] Scott-Railton, J 2012 Revolutionary Risks: Cyber Technology and Threats in the 2011 Libyan Revolution. United States Naval War College.

[11] Regalado D, Villeneuve N & Scott-Railton J Behind the Syrian Conflict's Digital Front Lines. FireEye, 2015.

Estonia[12] and beyond, researchers at the Citizen Lab and elsewhere have observed the evolution of PGEAs who use as their core toolkit commercial off-the-shelf RATs and other basic techniques borrowed and adapted from cybercrime forums. Increasingly, these groups also use packers and code obfuscation to circumvent antivirus detection. Although these groups' technical sophistication might be low, they are often keen and highly adaptive social engineers. The Citizen Lab has tracked cybermilitias stemming from the Arab Spring, and similar groups are beginning to appear elsewhere. A unique feature of cybermilitias is that they can use technically inexperienced but crafty individuals to seed malware assembled by a much smaller group. This lets them scale their ranks of nontechnical attackers quickly, perhaps more so than the other PGEA categories.

Lessons from Civil Society Groups

The resource and security problems that plague civil society are structural, political, and unlikely to be fixed by security engineers and researchers alone. Nevertheless, civil society's predicament highlights two lessons relevant to the problem of security for the general population.

Human Behavior Has Forever-Day Vulnerabilities

Users are still vulnerable to attacks with low technical sophistication but well-crafted social engineering.

It's human nature to want to help, to be curious, and to respond to a sense of fear and urgency. This natural urge presents an endless opportunity for attacks that rely on deception and trust exploitation. The PGEAs targeting civil society groups share the emphasis on targeting human behavior as the primary entry point for their campaigns. This "just enough" principle holds for many more sophisticated a ack groups, even with harder targets, as the head of the US National Security Agency's Tailored Access Operations recently pointed out at a security conference (www.youtube.com/watch?v=bDJb8WOJYdA). Social engineering has always been a game of probability, and most organizations contain members who are more likely to be taken in than others.

A growing cottage industry is "security training" that focuses on increasing civil society's awareness of surveillance and malware and on shifting security behavior. This development is

[12] The work of Rain Ottis is pioneering on this topic, see for example: Ottis, R. (2010) From Pitch Forks to Laptops: Volunteers in Cyber Conflicts. In Czosseck, C. and Podins, K. (Eds.) Conference on Cyber Conflict. Proceedings 2010. Tallinn: CCD COE Publications, pp 97-109.

promising but deserves a stronger evidence-based footing and numbers-driven repeat testing, such as the use of regular phishing simulations and penetration testing.

Even when large user populations become vigilant and change their behaviors, attackers have shown a remarkable ability to adapt their techniques in response. Still, a range of good security technologies attempt to account for some of these risks; the challenge is ensuring that they're systematically enabled.

Optional Security Features Don't Get Enabled

When a security feature is optional, harried users are less likely to enable it.

Most civil society groups depend on widely used online platforms and OSs. Despite facing substantial threats, they don't operate in managed environments or compute on managed endpoints. Security decisions are, to a great extent, left to individuals. Even when encouraged for a group, the norms and incentives often remain too weak to shape everyone's behavior. For example, the initial time and behavioral costs required to enable optional security features on accounts is often just enough to keep many from doing so. We've observed countless high-risk individuals abandoning security practices "just this once" or altogether, often over simple user experience or workflow frustrations.

> *"Attackers readily exploit this lack of security 'group immunity.'"*

Attackers readily exploit this lack of security 'group immunity.' For example, it's rare to find a civil society group that has comprehensively adopted two-factor security. Although attackers have a wide range of potential approaches, threat actors targeting civil society often rely heavily on credential theft, even when using malware.

A Specific Plea to Industry: Harden Authentication for Common Users

Currently, only some private sector companies and governments can afford serious security. As a result, common users must depend on companies to "have their backs."

Security-versus-usability tradeoffs are real, but it's arguable that today's free services place too much responsibility on amateur and often harried users. One area stands out in particular: authentication. User credentials are a constant target of attacks, and hardening the authentication process would benefit both common users and higher-risk civil society groups.

> *"The question is: Should two-factor authentication be a default-on option for popular platforms?"*

Two-factor authentication is already available as an optional feature for many popular online platforms. Despite an extensive history of attackers attempting to find workarounds, and growing evidence of attacks against SMS-based two-factor authentication, it remains a very effective tool that quickly raises the attackers' cost per compromise. The question is: Should two-factor authentication be a default-on option for popular platforms?

Enabling key security choices by default can have powerful effects. When companies enroll employees in retirement by default, participation increases dramatically.[13]

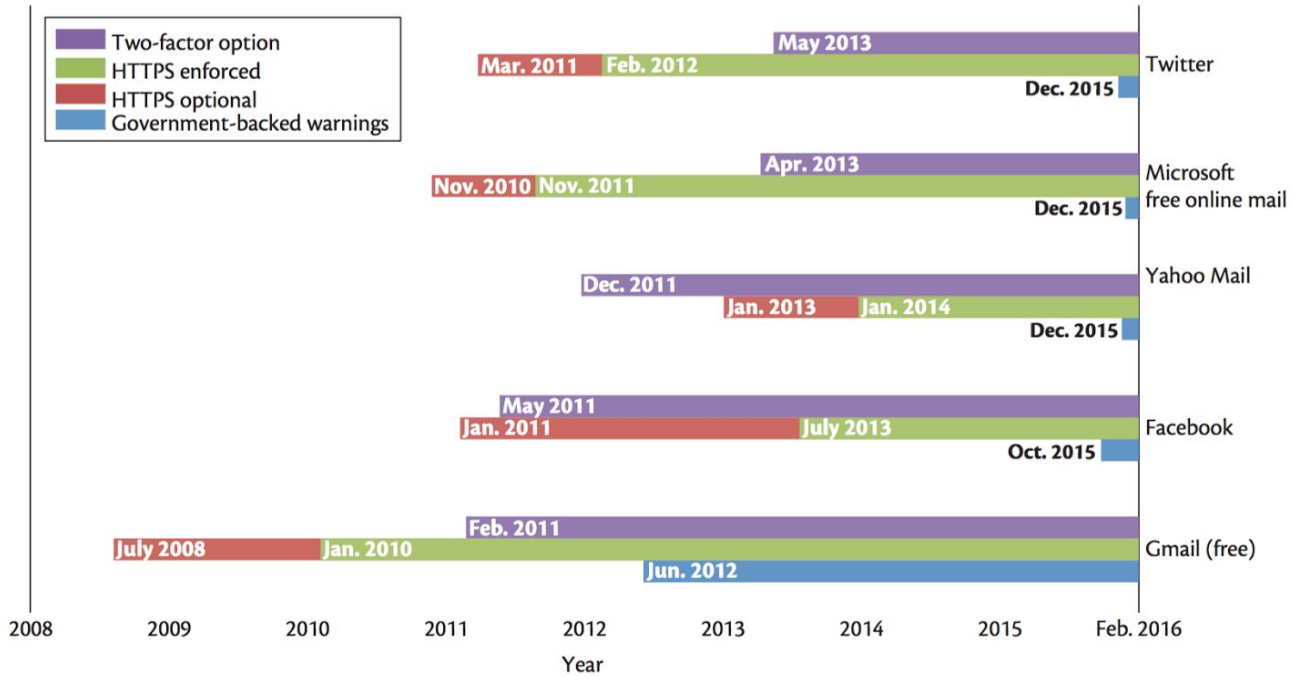

[13] The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States, John Beshears, James J. Choi, David Laibson, Brigitte C. Madrian. in Social Security Policy in a Changing Environment, Brown, Liebman, and Wise. 2009

Figure 5. Approximate date, based on public announcements, when key security features became available on popular online platforms' free services.[14]

There's also an analogy to the history of automobile safety.[15] Seat belts were originally a nonstandard option on vehicles sold in the US. It wasn't until seat belts became standard in cars *and* required by law that people started to use them ubiquitously.[16] This eventually made seatbelt wearing unobtrusive and no longer viewed as inconvenient.

> *It wasn't until seat belts became standard in cars and required by law that people started to use them ubiquitously*

Does this analogy also apply to the tech sector? For an industry historically wary of government regulation, continuing to raise the security bar for all users would signal that self-regulation is possible. There is recent evidence of this. For example, full-session HTTPS by default (and most recently, HTTP Strict Transport Security, STARTTLS, and Forward Security)[17] was a critical step and an indication that, in the long run, coordinated efforts are possible, despite early worries about full-session HTTPS's feasibility. Meanwhile, Google, Facebook, Yahoo, Twitter, and Microsoft (among others) also deserve credit for notifying users of suspected government-backed attacks, and for actively investigating threats on their platforms (see **Figure 5**).

Companies' collaborative efforts to share information about threats is another important move. So, too, is the growing sophistication of other methods of acting on signals of user identity, such as requiring authentication when unfamiliar devices, locations, and browsers are detected.

Decisions about making security features optional vs. default on are constrained by economic, usability, and marketing concerns. Implementing default-on two-factor authentication in its current form on popular platforms might add time to user signup processes, degrade user experiences compared to competitors, and increase the costs associated with account recovery. Most two-factor techniques on popular platforms send SMSs by default, which would

---

[14] For Gmail HTTPS optional, although HTTPS was available previously (by typing HTTPS://www.gmail.com into the menu bar), users couldn't enable it as "default on" before the date shown. The date for Microsoft HTTPS enforced is based on publicly cited information indicating a rollout and might be slightly inaccurate.

[15] Kosslyn, Justin Have a Very Malware-Free Christmas, Slate, December 23, 2015

[16] Alma Cohen & Liran Einav 2011 The Effects of Mandatory Seat Belt Laws on Driving Behavior and Traffic Fatalities. Discussion Paper # 341. Center for Law, Economics and Business, Harvard University.

[17] EFF, EFF's Encrypt The Web Report.

dramatically increase costs if two-factor was the default. Finally, users would (temporarily) risk losing access to accounts and the important personal information they contain.

Many issues with mandatory two-factor authentication are addressed in managed environments, or are accepted as the cost of doing business for institutions such as banks that are liable for account breaches and can use fees to defray the costs of sending SMSs.

With sufficient research and user testing, perhaps focused on experimenting with tactics to measurably increase two-factor use for high-risk groups, it might be possible to develop a roadmap for overcoming these concerns. Users who have received warnings of government-backed attacks are a natural place to start, as are highly-targeted diaspora communities and countries undergoing political unrest. Mitigating the cost of sending so many SMSs is more difficult. However, future innovation and improved support for other second factors, such as physical devices, apps, mobile-apps, and other tokens and standards (e.g. U2F), will likely help diminish the issue of costs. This also addresses the threat of malicious carriers and other parties interfering with SMS traffic to intercept and use 2 Factor codes.

Civil Society as Canary in the Coal Mine

The security of civil society groups has an image problem among security professionals and academia. Although targeted attacks, like malware campaigns, are often highly effective, inflicting severe harm, their techniques are not always at the technological cutting edge. Moreover, they rarely affect tens of thousands or millions of users directly. Thus, many security professionals view such attacks as "edge cases" and give them less attention than other attack categories such as large-scale campaigns, or the most technologically advanced attacks. Yet these edge cases might be harbingers of the techniques that will be levied against common users within a particular geographic location or community. Thus, more research and data sharing between industry and academia is desperately needed to understand them.

> *"...these edge cases might be harbingers of the techniques that will be levied against common users…"*

Security professionals might already be aware of the threats against civil society occurring on their platforms. Yet there are few avenues for action, and serious time and resource constraints make investigations difficult. Similarly, notifying potentially affected parties in a fraught political environment might seem like an institutional risk.

> *"...harried, high-risk users must maintain superhuman standards of behavioral security … or face potentially devastating compromises."*

Research groups, including the Citizen Lab, and civil society groups are studying the problem and building relationships with at-risk groups. This research is under-resourced compared to the private sector, yet the problems are no less important. Exploring new avenues for partnerships, data sharing, and engaging after attacks would be an excellent way to exercise social responsibility. Given the emphasis on social engineering and credential theft, research on high-risk user populations is urgently needed.

As the problem stands, harried, high-risk users must maintain superhuman standards of behavioral security and make individual choices to increase security—or face potentially devastating compromises.

Large platforms have a responsibility to protect high-risk users, even if these users weren't consciously invited. Many are moving in this direction, warning of government-backed attacks and investing in increasingly sophisticated risk analysis around alternate signals of user identity.

> *"...the many small individual choices ...should be replaced by big choices by companies deciding to take a leadership role…"*

In this article, I emphasized default-on two-factor authentication as a similarly important goal, despite its associated costs. Experimenting with aggressive techniques to increase adoption, initially focused on known high-risk user groups, would make a tangible dent in the success of threat actors targeting not only civil society but the rest of us too. Ideally, the many small individual choices to enable secure options should be replaced by big choices by companies deciding to take a leadership role in solving this problem.

Acknowledgements

I thank many colleagues including Iulia Ion, Masashi Nishihata, Bill Marczak, Justin Kosslyn, Ron Deibert, Eric Sears, Peter Railton, Ned Moran, Matt Braithwaite, Brandon Dixon, Jonathan Pevarnek, Gary Belvin, and Kristen Dennesen.