

MUNK
SCHOOL
OF
GLOBAL
AFFAIRS




COMMUNITIES @ RISK

Targeted Digital Threats Against Civil Society

November 11, 2014

<https://targetedthreats.net/>



EXTENDED ANALYSIS:

2.1

Summary, Methodology, and Data Overview

Summary

Communities @ Risk: Targeted Digital Threats Against Civil Society reports on an intensive study that analyzes targeted digital threats against 10 civil society organizations (CSOs) over a period of four years.

The report combines two major sections:

1. The **Executive Summary** provides detail on how the study was organized and why we feel it is important to read, a high-level overview of the key findings of the research, and considerations about next steps for several stakeholder communities in responding to targeted digital threats.
2. The **Extended Analysis** explains our methodology, and examines the detailed data we gathered during the study period. It is the evidentiary basis for the claims we make in the Executive Summary, and will likely be of interest to a more specialized audience (although we hope everyone will read it).

KEY FINDINGS

In the Executive Summary, we outline five high-level findings. We summarize them again below, while adding more granular details that are given extended treatment in the analysis that follows.

In the digital realm, CSOs face the same threats as the private sector and government, while equipped with far fewer resources to secure themselves.

Through cluster analysis we identify 10 distinct targeted malware campaigns. We find that five of these campaigns have connections to threat actors, previously reported to have targeted government and private industries. CSOs have limited resources and technical capacity, which makes responding to threats a challenge. We generally find that, due to resource constraints, CSOs focus their digital security strategies on user education and behavioural change rather than expensive technical solutions.

Counterintuitively, technical sophistication of malware used in these attacks is low, but the level of social engineering employed is high.

We develop the Targeted Threat Index, a metric for quantifying and characterizing the sophistication of targeted malware attacks. Using this metric, we find that the technical sophistication of targeted malware delivered to CSOs in our study is relatively low (e.g., relative to commercial “lawful intrusion” surveillance kits and conventional financially motivated malware), with much more effort given to socially engineering messages to mislead users.

Digital attacks against CSOs are persistent, adapting to targets in order to maintain access over time and across platforms.

Our analysis of attacks against CSOs over four years allows us to track how attackers change tactics. For numerous malware samples, we observe several versions of the malware appearing over the course of our study. These multiple versions show evidence of technical improvements to complement increasingly refined social engineering techniques. In some cases, we observe threat actors quickly changing tactics to adapt to shifting platform adoption and user behaviour.

Targeted digital threats undermine CSOs’ core communications and missions in a significant way, sometimes as a nuisance or resource drain, more seriously as a major risk to individual safety.

The impact of targeted digital attacks against technical systems is apparent and receives ample attention from researchers. However, we find evidence of wider impacts that are not always as obvious, including psychosocial strain and possible connections to physical harms (e.g., arrest and detention). Tracing connections between compromises and harm is challenging, because the relationship between digital compromises and the use of the compromised information by threat actors is indirect. Unlike the consequences of physical threats, which are often readily observable, the most serious impacts of digital threats are typically at least one step removed from the technology that has been exploited.

Targeted digital threats extend the “reach” of the state (or other threat actors) beyond borders and into “safe havens.”

The ways CSOs develop their perceptions of risk and threat stemming from targeted attacks depend in part on the physical proximity of their threat actor. Groups operating within the jurisdiction of a repressive regime have greater concerns over physical security and other direct interference from authorities. Conversely, groups situated

outside of a physical jurisdiction controlled by an adversary may prioritize digital threats over physical threats. For groups in diaspora and exile communities, targeted digital threats can be seen as a means for a powerful threat actor, such as a state, to extend their reach beyond borders and into “safe areas.”

EXTENDED ANALYSIS STRUCTURE

The Extended Analysis is structured into the following three sections. Each of these sections can be downloaded individually or read as a whole.

Summary, Methodology, and Data Overview outlines our mixed methods approach which incorporates analysis of technical and contextual data using methodologies from the field of information security and the social sciences, and presents a high level overview of our dataset.

Cluster Analysis provides detailed technical analysis of 10 distinct targeted malware campaigns.

Civil Society Perspectives and Responses reports on results from interview data and is a window into how groups under threat think about and respond to digital threats.

We also are publishing data that provide indicators of compromise (including YARA signatures of malware families, MD5 hashes of samples, and command-and-control servers), which are available on our github account and accessible through our project website.

Methodology

This section describes our methodology for data collection and analysis. Since our study involves the collection of potentially sensitive information from civil society organizations, and requires us to deal with personally identifiable information (PII), we consulted with the University of Toronto's Research Ethics Review Board during the design of our study. The methods described below have been submitted to and approved by this board.

STUDY PARTICIPANTS

We recruited participants via three channels: (1) an open call on our website, (2) outreach to organizations with which we had prior relationships, and (3) referrals from participating groups. As part of the study, these groups agreed to share technical data (e.g., emails with suspicious attachments) and participate in interviews. Their identity and any PII shared with us were kept strictly confidential.

Organizations with a mission concerning the promotion or protection of human rights were eligible to participate.¹ We also considered, on a case-by-case basis, organizations with a mission that does not directly address human rights, but which may engage in work related to human rights issues (e.g., media outlets that regularly report on human rights violations).

In total, 10 organizations participated in the study. The majority of these groups work on China-related rights issues, and five of these organizations focus specifically on Tibetan rights. The exceptions to the China- / Tibet-focused groups in our study are two large organizations that work on multiple human rights-related issues in various countries.

¹ For purposes of this study, "human rights" means any or all of the rights enumerated under the [Universal Declaration of Human Rights](#); the [International Covenant on Civil and Political Rights](#); and the [International Covenant on Economic, Social and Cultural Rights](#).

TABLE 1: Study organizations

ORGANIZATION CODE	DESCRIPTION	ORGANIZATION SIZE
Rights Group 1	Human rights organization focused on multiple issues and countries	Large (over 100 employees)
Rights Group 2	Human rights organization focused on multiple issues and countries	Large (over 100 employees)
China Group 1	Human rights organization focused on rights and social justice issues related to China	Small (1-20 employees)
China Group 2	Independent news organization reporting on China	Small (1-20 employees)
China Group 3	Human rights organization focused on rights and social justice issues related to China	Small (1-20 employees)
Tibet Group 1	Human rights organization focused on Tibet	Small (1-20 employees)
Tibet Group 2	Human rights organization focused on Tibet	Small (1-20 employees)
Tibet Group 3	Independent news organization reporting on Tibet	Small (1-20 employees)
Tibet Group 4	Human rights organization focused on Tibet	Small (1-20 employees)
Tibet Group 5	Human rights organization focused on Tibet	Small (1-20 employees)

Tibet Groups

Dharamsala is a small city in northern India set on the foothills of the Himalayas. His Holiness the Dalai Lama (HHDL) has lived in Dharamsala since 1959 following his escape from Tibet. Dharamsala is the base of the Central Tibetan Administration, which administers programs and schools for Tibetan refugees living in India and advocates for the rights of Tibetans in Tibet. It is also home to many Tibetan NGOs and independent media groups, and thousands of Tibetan refugees. This high concentration of prominent Tibetan institutions makes Dharamsala a prime target for malware campaigns. It has been called [one of the most hacked places in the world](#). For exiled Tibetans, this heightened level of digital risk compounds the many challenges of living as refugees in a developing country.

Three of the Tibet Groups in our study are headquartered in Dharamsala, and two

maintain regional offices there. Across these groups participants expressed challenges related to awareness of threats, low resources, and limited technical capacities.

Tibet Groups reported varying levels of awareness of digital risks in the community. While many participants noted that security awareness was generally increasing among Tibetans, others cautioned that some groups still do not have policies or response plans around targeted digital attacks and “continue to back burn things like security.”²

A major challenge identified by the Tibet Groups is a lack of technical capacity and resources in the community. Most Tibetan NGOs do not have dedicated system administrators. In some groups, staff members responsible for web development also take on double duty as system administrators. In addition to local staff, there are transient volunteers who come into the community to help with technical projects. As one of these volunteers noted, however, when volunteers leave the community projects sometimes end up unmaintained or completely abandoned.

While the unique circumstances of the Tibetan exile community are challenging, some groups are also taking proactive measures to increase digital security awareness. For example, one of our participating organizations prioritizes digital security in the community within its mission, focusing on raising awareness and user education. These grassroots initiatives demonstrate a growing commitment to addressing security challenges, despite ever-present resource limitations.

China Groups

The three China Groups all work on issues related to human rights and politics in China, but from outside of mainland China. China Groups 1 and 3 each have a central office and one regional branch. China Group 2 operates an independent news website from an office with limited staff. China Group 1 has a program manager that oversees technical projects, but does not have a dedicated system administrator on staff. Instead the group outsources management of its information technology infrastructure to a private company. China Group 3 has had a dedicated system administrator since its founding.

The work of these groups is politically sensitive and has attracted attention from Chinese authorities. China Groups 1 and 2 especially have come under pressure for

2 Tibet Group 1, Program Director, 2011

human rights advocacy and the dissemination of sensitive news, respectively. As China Group 1 explained, “Chinese authorities ... have very clearly in public designated us as an anti-China organization.”

These groups are all highly aware of targeted digital threats, and have experienced numerous prior incidents. All of the groups had received targeted malware in the past and their websites are consistently blocked in China. The website of China Group 2 has been repeatedly hit by distributed denial-of-service attacks.

Rights Groups

Rights Groups 1 and 2 are much larger organizations relative to the others in our study. Both have over 100 employees, multiple offices, enterprise level computing infrastructures, and dedicated IT teams and support desks.

These groups act as hub organizations. Rights Group 1, for example, supports multiple regional offices and CSO partners around the world. Rights Group 2, similarly, operates regional branches and is responsible for a large group of staff operating in numerous field locations.

Both groups contend with securing their head offices and maintaining awareness of threats faced by field offices. These challenges show that while the Rights Groups have greater resources they must grapple with a potentially wider spectrum of threats in multiple contexts and countries.

DATA SOURCES

Email Submissions: The majority of data collected consisted of emails identified by participants as suspicious, which were forwarded to a dedicated mail server administered by our research team. When available, these submissions included full headers, file attachments, and / or links.

Relying on forwarded emails presents a collection bias as the recipients must be able to identify that the emails are suspicious and remember to forward the samples to our research team. This collection method also limits the threats studied to those that are sent over email. Additionally, collection of forwarded email samples does not allow us to verify if a targeted organization was successfully compromised by an attack, or the scope of the attack. Recognizing this limitation, we added two more data collection methods to complement the collection of emails.

Network Intrusion Detection System: As an optional study component, we offered to install a network intrusion detection system (NIDS) inside the networks of the participants. In total, seven groups opted into the NIDS project. We used a combination of [community](#) and [commercial](#) rulesets, as well as a set of custom rules based on threats we analyzed from the email submissions. By placing a NIDS inside an organization's network, we were able to record incoming threats using vectors other than email, as well as detect and observe systems that had already been compromised.

Website Monitoring: We conducted external scans of the study organizations' websites to monitor for potential compromises such as [watering hole attacks](#). These scans were done with publicly available tools including [Cyberspark](#) and [URL Query](#).

Interviews and Fieldwork: To gain insights into the experiences of our groups, we conducted a series of semi-structured interviews over a four-year period and made site visits to their offices and locales. While there have been previous technical studies on targeted threats affecting CSOs, it is rare that the context surrounding these attacks and the experiences of the people facing them are properly explored. Interviews and site visits help provide insight into these vital elements.

When possible we conducted interviews with a senior staff member responsible for organizational programming (e.g., executive director, program manager), and a staff member responsible for technical support (e.g., system administrator, webmaster). The interviews explored the organizations' uses of and policies around technology, perceptions of digital security and threats, responses to threats, and the impact of threats. These interviews, coupled with site visits and participant observations, helped us understand the working conditions, routines, infrastructure, and local social and political context that form the day-to-day environment of our participants.

Interviews were held opportunistically and did not follow a set schedule. The total number of interviews per group is outlined in Table 2. The majority of interviews were audio recorded and transcribed. In some cases, conditions did not allow for audio recording and field notes were made instead. Interview transcripts were analyzed using line-by-line open coding of transcripts to identify emergent themes.³

3 Methods are described in Creswell, J.W. (2013). *Qualitative Inquiry and Research Design: Choosing Among Five Approaches*. London, UK: Sage.

TABLE 2: List of interviews conducted with participating groups*

GROUP	SUBJECTS	DATE
China Group 1	Executive Director, Program Manager (technical projects)	2010
China Group 3	System Administrator	2011
Rights Group 1	Chief Technical Officer, Program Manager	2012, 2014
Rights Group 2	Technical Officer	2011, 2014
Tibet Group 1	Executive Director, Program Director, Program Director (technical projects), Program Officer, Security Trainer	2011, 2012, 2013
Tibet Group 2	Executive Director	2013
Tibet Group 3	Editor-in-Chief	2014
Tibet Group 4	Technical Volunteer	2013
Tibet Group 5	Program Officer	2014

*NOTE: We were unable to conduct a site visit and interview with China Group 2, because they did not maintain participation in the project.

DATA ANALYSIS

Malware Analysis: We examined malware samples using static and dynamic analysis tools (e.g., IDA and OllyDbg), as well as manual analysis to extract information on exploits, malware functionality, malware family, command-and-control (C2) infrastructure, and other properties of the malware code (e.g., mutex and exported function names).

Email Content Analysis: We reviewed the subject line, body, and attachments for each submitted email and grouped the content into specific themes and categories. The header of each email was analyzed to determine if the sending email address was spoofed or the email address was otherwise designed to appear to come from a real person and / or organization. Indicators drawn from this analysis were used to assess the relative sophistication of the social engineering tactics found in the messages (we incorporate these indicators into our Targeted Threat Index described below). We conducted regular inter-rater reliability checks that flagged any potential edge cases and inconsistencies for discussion and re-evaluation.

Targeted Threat Index: We developed the Targeted Threat Index (TTI), which is a metric that characterizes and quantifies the sophistication of targeted attacks, to provide a consistent ranking of how advanced any given targeted malware attack is. The TTI score is calculated by taking a base value determined by the sophistication of the targeting method, which is then multiplied by a value for the technical sophistication of the malware. The base score can be used independently to compare emails, and the combined score gives an indication of the level of effort an attacker has put into individual threats.

Cluster Analysis: Through identification of patterns in malware families, development cycles, shared infrastructure, and social engineering tactics, we identified relationships between attacks and, when possible, linked them to known malware campaigns and threat actors.

DATA OVERVIEW

A high level overview of our datasets

EMAIL SUBMISSIONS

The malicious emails analyzed in this report span more than four years, from October 10, 2009 to December 31, 2013. During this period we collected 817 emails from the 10 groups participating in our study.

FIGURE 1: Cumulative number of email submissions per month during the study

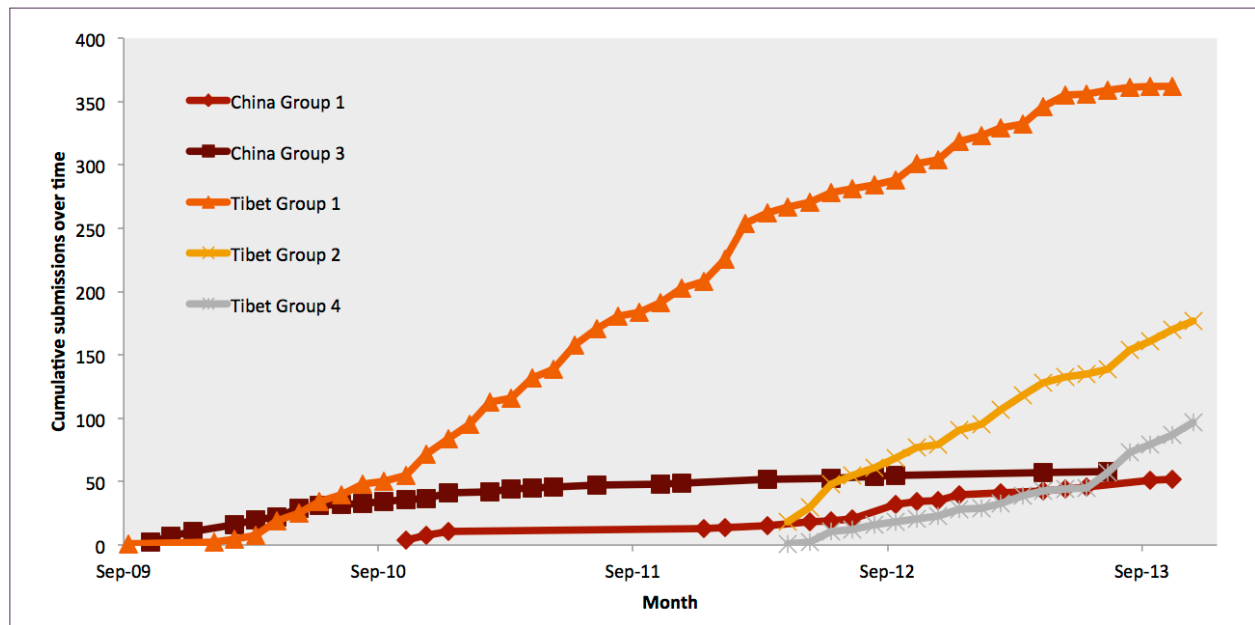
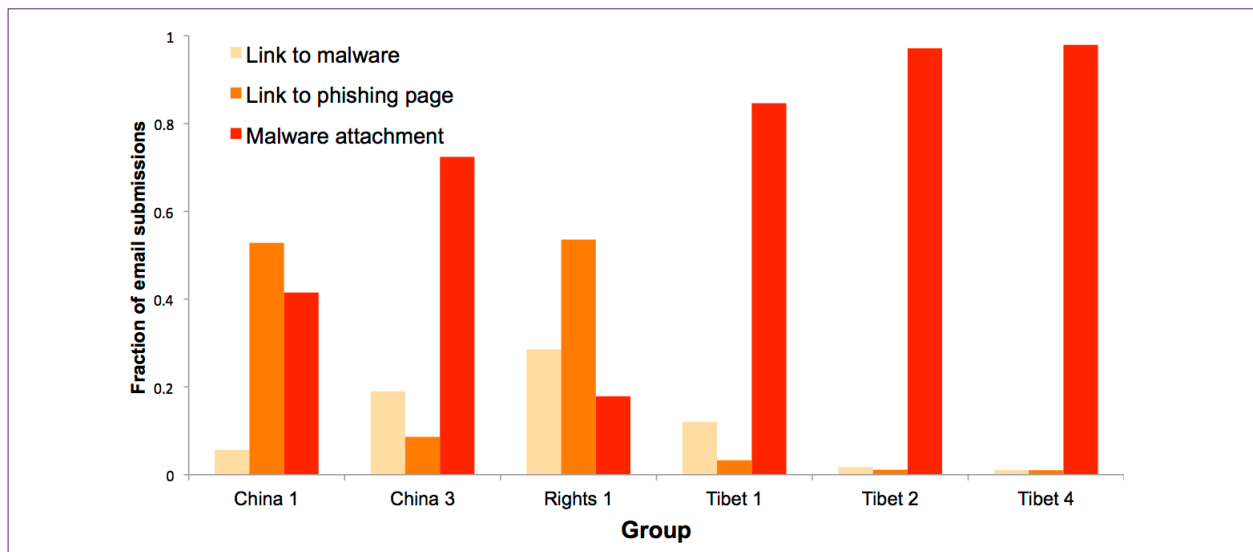


Figure 1 shows the cumulative number of email submissions per month over the course of the study. Although the first formal submission was received on November 28, 2011, some groups had existing archives of malicious messages received by their members, and they provided us with these older emails. Tibet Group 1 ac-

counts for the highest number of submissions relative to the other groups as it was one of the first groups in the study and is persistently targeted. Tibet Groups 2 and 4, which joined the study at a later date (April 2012), show a similar submission rate to Tibet Group 1, suggesting these groups are targeted at a comparable level.

FIGURE 2: Malicious emails by type for groups submitting 25 or more emails



We classify emails as malicious if they include attached malware, a direct link to malware or a drive-by download site, or a link to a phishing page. Figure 2 shows the number of emails of each type for the groups that submitted at least 25 emails to our system. The most common technique employed in these emails was a malicious attachment to the message. However, we observe a higher rate of phishing attacks on the China Groups and the Rights Groups. In particular, 46% of the emails submitted by China Group 1, and 50% of the emails submitted by Rights Group 1, direct the user to a phishing website.

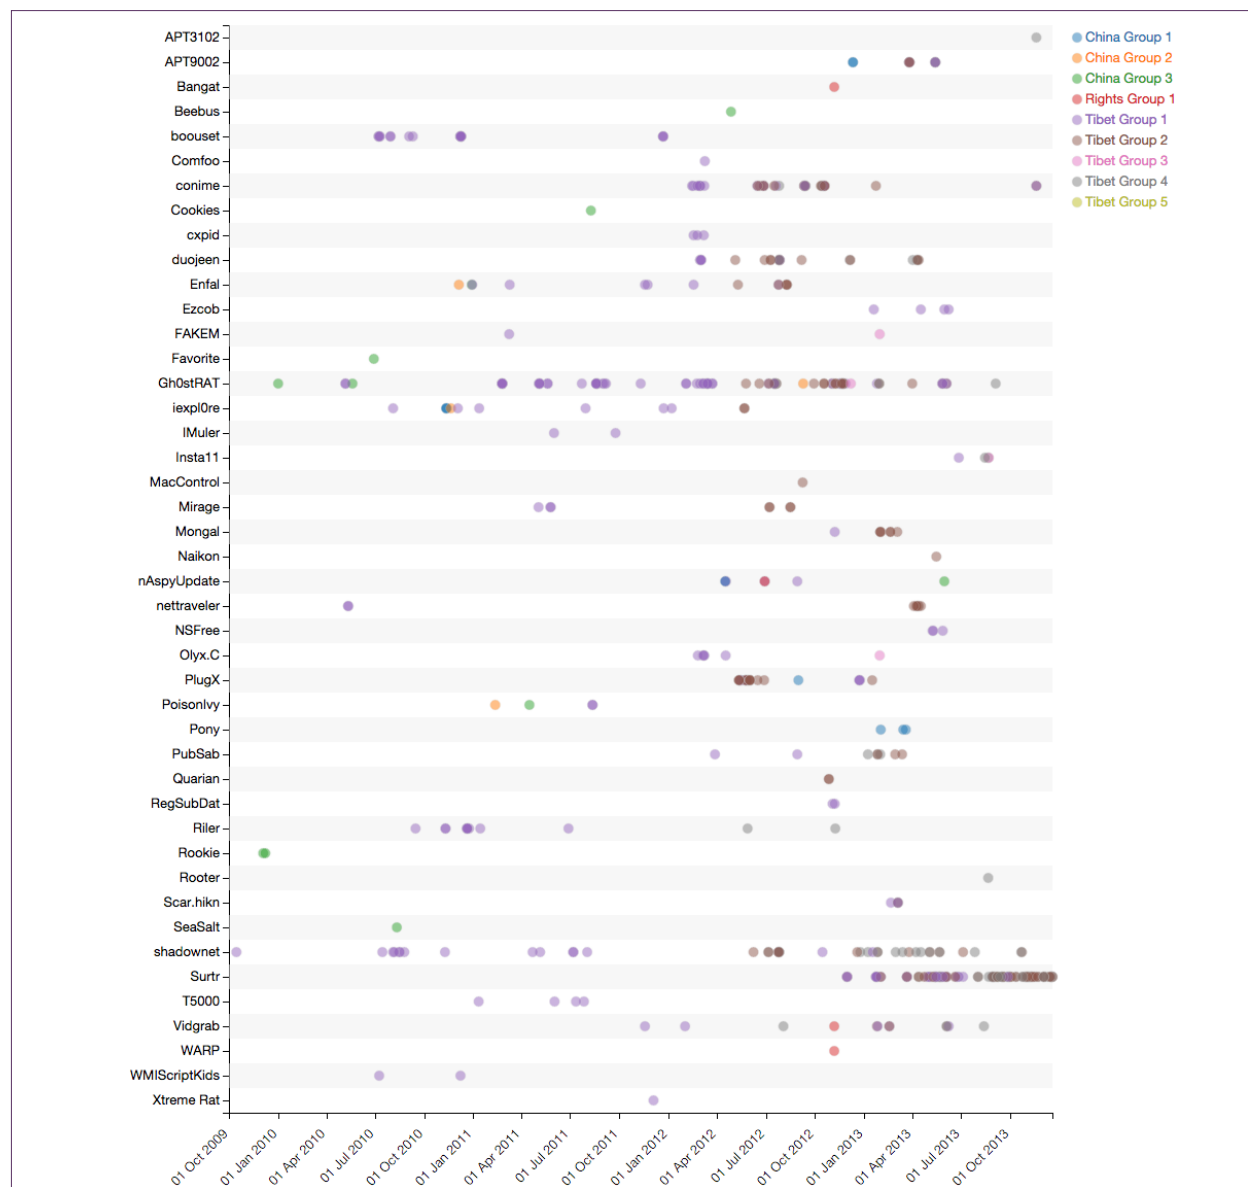
The rate of submissions to our project meant that it was feasible to manually analyze email attachments for malware as they were submitted. This analysis gives us higher confidence in our results than if we had automated the process. Antivirus (AV) signatures frequently fail to detect new or modified threats, and can overlook the kind of malicious payloads that can be identified with manual inspection (e.g., shellcode in an RTF exploit). In total, we analyzed 3,617 payload files and found 2,814 (78%) to be malicious.

MALWARE FAMILIES

We identified malware families through patterns in [network traffic](#) and characteristics in the code, such as strings seen in the binaries or names and locations of dropped files. In total, we identified 44 separate malware families (not including variants). The most frequently occurring families are Gh0st RAT, Surtr, Shadownet, Conime, Duojeen, and PlugX.

FIGURE 3: Malware family timeline

(The coloured dots represent attacks using a particular malware family against one of our study groups.)




CVEs

[Common Vulnerabilities and Exposures](#) (CVEs) is a dictionary of common names for publicly known security vulnerabilities. CVEs are each assigned a unique identifier code, with the form CVE-YYYY-NNNN, where YYYY indicates the year they were identified and NNNN are arbitrary digits. We identified 24 distinct CVEs used in 483 of the email attacks as displayed in Figure 4.

FIGURE 4: CVEs identified in samples during the study period.

(Vertical gray bars represent the date the CVE was created and orange dots represent targeted attacks using that CVE.)



The most common CVEs we observed were [CVE-2010-3333](#) (used in 112 attacks) and [CVE-2012-0158](#) (used in 294 attacks), which are both vulnerabilities in the way Microsoft Word handles RTF documents. Figure 4 clearly illustrates the shift in use from CVE-2010-3333 to CVE-2012-0158 in March and April of 2012. The popularity of these vulnerabilities is not limited to our dataset. They have been [widely used](#) in other attacks against a variety of targets.

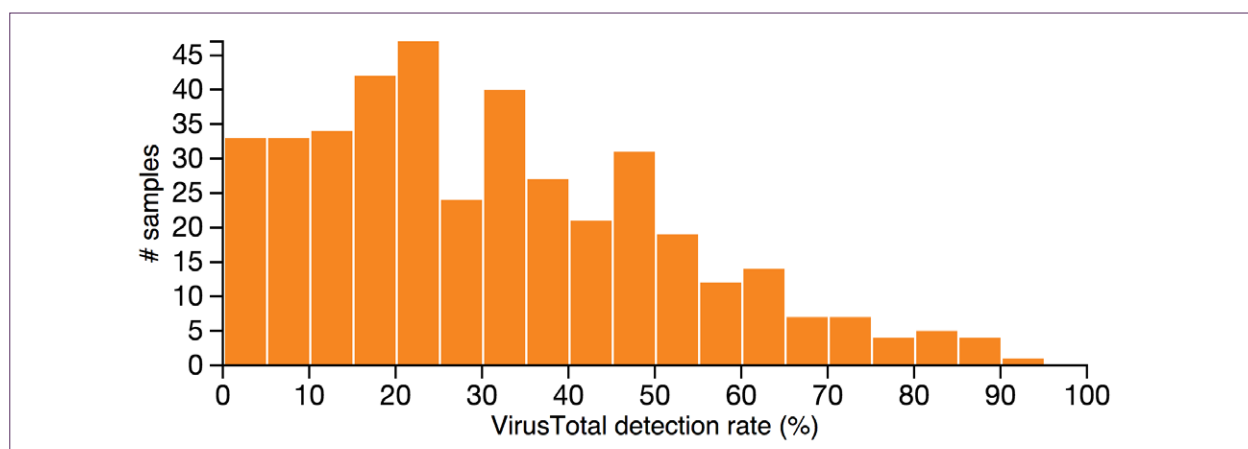
During four years of tracking attacks against our groups, we observed only one zero-day exploit. This attack used the Flash vulnerability [CVE-2012-5054](#), and was sent 22 days before the CVE entry was created.

These results show that vulnerabilities exploited by targeted malware attacks against CSOs are typically not technically advanced (compared with financially-motivated malware and commercial lawful intercept kits), and often use old (patched) vulnerabilities. For example, CVE-2012-0158 has been [patched](#) since April 10, 2012, but has remained the most common vulnerability used in attacks against the Tibet Groups for well over a year after the fix was issued. The repeated use of this vulnerability suggests the attackers are achieving successful compromises because target systems did not have the latest security updates. A possible explanation is that licensed software is cost-prohibitive for many organizations in the developing world, while pirated copies are easily available, leading many to use pirated operating systems and software.

ANTIVIRUS DETECTION

VirusTotal is a service that scans files through 53 different AV engines and provides a summary of malware detection results. We find that 369 of the 659 samples we received (56%) had been submitted to VirusTotal at the time of writing, with a median AV detection rate of 24% and mean detection rate of 25%. Detection rates were generally low, as 86% of these samples had a detection rate below 50%, meaning that less than half of the AV packages tested were able to identify them as malicious. These results suggest that simply running AV software, although potentially helpful, is not a very effective defence against these attacks.

FIGURE 5: Histogram of antivirus detection rates provided by VirusTotal



This low detection rate we observed is due in part to the extensive presence of CVE-2012-0158, which uses a number of techniques to hide the vulnerability from AV scanners.

One of the simplest of these detection-reducing techniques is modifying the RTF header, since Microsoft Word will still be able to open the file, but fewer AV scanners will detect it as malicious. Another basic technique is encrypting malicious document and providing a password to open the file in the associated email. Simply adding a password to malicious files can help prevent AV detection.

Since there are four ActiveX controllers—ListView, ListView2, TreeView, and TreeView2—affected by this vulnerability and there are no strict syntax restrictions, there can be a large variance in the document templates into which malicious payloads are inserted. These can cause newer templates to initially have lower detection rates.

A notable technique observed was the creation of a MIME HTML (MHTML) file that uses the vulnerable ActiveX controllers. By default, MHTML files are opened by a browser; however, they can also be opened by Microsoft Word, which will trigger the exploit. Since Microsoft Word may not be the default application to open the file, automated sandbox programs may fail to detect the file as malicious.

The older CVE-2010-3333 vulnerability had similar issues with AV detection, because of the wide number of ways to encode the vulnerability. A small change in the way the vulnerability was written could evade signature detection while remaining functionally the same.

Although AV definitions are updated to account for evasion tricks, the lag between the use of evasion techniques in the wild and definition updates results in temporarily low detection rates, and hence the likelihood of successful compromises.

EMAIL CONTENT ANALYSIS

Subject line, body, and attachments: The content of the subject line, body, and attachments for each submitted email were content coded into 134 categories grouped under eight themes:

- Country / Region (referring to a specific geographical country or region)
- Ethnic Groups (referring to a specific ethnic group)
- Event (referring to a specific event)
- Organizations (referring to specific organizations)
- People (referring to specific people)
- Political (reference to specific political issues)
- Technology (reference to technical support)
- Miscellaneous (content without clear context or categories that did not fall into one of the other themes)

Email headers: The header of each email was analyzed to determine if the sending email address was spoofed, or the email address was otherwise designed to appear to come from a real person and / or organization (for example, by registering an email account that resembles a legitimate sender's name from a free email provider). We divide the results based on whether they attempted to spoof an organization or a specific person.

Results of this analysis confirm that message content and fraudulent senders are tailored to the interests of the target organizations.

Of the 520 total emails received by the Tibet Groups, 97% referenced content related to Tibetan issues. Email lures leveraged specific events of interest and respected persons in the Tibetan community. Emails referenced Tibet-related events, including holidays (Tibetan New Year), anniversaries (His Holiness the Dalai Lama's birthday), and protests (see Table 3). The most frequently referenced events were Tibetan self-immolations (31% of the emails leveraging event-related content).



Some of the attachments actually cannot be detected as a virus...We're not even sure if it...will cause any harm at all. It's just that the antivirus [is] saying that 'there's no threat,' but obviously there's something wrong with it."

—China Group 1

TABLE 3: Breakdown of top five categories in the Event theme for Tibet Groups

CATEGORY	NUMBER OF EMAIL RECORDS
Self-Immolation *	56
Tibetan National Uprising Day	24
HHDL Birthday	19
Flame of Truth Rally	13
Kalon Tripa Election	9

* Self-immolations are a controversial form of protest that Tibetans have used as a statement of opposition to Chinese government practices concerning Tibet. These protests have escalated in recent years. At the time of writing, it is estimated that since 2009, [approximately 132](#) Tibetans have self-immolated.

Of the 520 emails received by Tibet Groups, 272 (52%) were designed to appear to come from real organizations. In total 58 organizations were spoofed, of which 53 (91%) were Tibet-related groups (see Table 4). The most frequently spoofed organization was the [Central Tibetan Administration](#). The identities of four of the Tibet Groups in our study (Tibet Groups 1, 2, 3, and 5) were frequently spoofed internally and to external contacts. The frequency of emails with fraudulent contacts from Tibetan organizations shows an effort to have the message appear to come from within the Tibetan community and leverage existing trust relationships.



The emotions of the immolations [are] being used against people to have them click on [attachments].”

—Tibet Group 1

TABLE 4: Breakdown of top five categories in the Spoofed Organizations theme for Tibet Groups

CATEGORY	NUMBER OF EMAIL RECORDS
Central Tibetan Administration	58
Tibet Group 1	26
Tibet Group 2	13
Tibet Group 5	13
Tibet Group 3	11

We see a similar pattern for the China Groups. Of the 48 emails received by the China Groups, 46 (95%) referenced China. Content included references to Chinese political events such as the Communist Party of China (CPC) 18th Party Congress; the June 4, 1989 Tiananmen Square crackdown; and Chinese dissidents and prominent members of the CPC (see Table 5). Of the 48 emails, 13 (27%) spoofed real organizations (see Table 6). Two of our China Groups were spoofed (China Group 1, China Group 3). Rights Group 1 was also spoofed in one message to China Group 1. The remaining spoofed organizations were prominent human rights groups and intergovernmental organizations (e.g., the UN Office of the High Commissioner for Human Rights).

TABLE 5: Breakdown of top five categories in the Event theme for China Groups

CATEGORY	NO. OF EMAIL RECORDS
Jasmine Revolution	8
June 4, 1989, Tiananmen Square Crackdown	4
CPC 18th Party Congress	2
July 2009 Urumqi Riots	1
Chinese New Year	1

TABLE 6: Breakdown of top five spoofed organizations for China groups

CATEGORY	NO. OF EMAIL RECORDS
China Group 1	4
China Group 3	3
Office of the High Commissioner for Human Rights	3
Open Society Institute	2
Chinese Human Rights Defenders	2

The volume of email submissions from Rights Group 1 and Rights Group 2 was much lower than that from the Tibet and China Groups. However, we also observed content and email senders tailored to these organizations. Rights Group 1 received messages related to human rights issues in Africa and Russia. Of the 12 emails submitted, 92% were made to appear to come from Rights Group 1 email addresses (no other organizations were spoofed). The majority of these messages were phishing attempts with lures related to IT support, designed to gain access to Rights Group 1 email credentials. Rights Group 2 submitted two email samples, both of which were related to human rights issues in the Middle East. One message was made to appear to come from a Rights Group 2 email address.

While the content analysis results clearly show targeted attacks tailored to the interests of targeted groups, content coding alone does not provide a measure of the sophistication of social engineering used in the attacks. In the following section, we describe a metric to determine relative sophistication of attacks.

TARGETED THREAT INDEX

Our dataset includes a wide range of targeted malware threats that have varying levels of complexity. This range presents a challenge in ranking the relative sophistication of the malware and targeting tactics used by attackers.

While metrics such as the [Common Vulnerability Scoring System](#) exist for the purpose of communicating the level of severity and danger of a vulnerability, there is no standardized system for ranking the sophistication of targeted email attacks. This gap is likely because evaluating the sophistication of targeting is non-technical, and cannot be automated due to the requirement of a strong familiarity with the underlying subject material.

To address this gap, we developed the Targeted Threat Index (TTI) to assign a ranking score to the targeted malicious emails in our dataset. The TTI score is intended for use in prioritizing deeper analysis of incoming threats, as well as for getting an overall idea of how severely an organization is threatened.⁴

The TTI Score is calculated in two parts: (Social Engineering Sophistication Base Value) × (Technical Sophistication Multiplier) = TTI Score

TTI scores range from zero to 10, where 10 is the most sophisticated attack. Scores of zero are reserved for threats that are not targeted, even if they are malicious. For example, an email from a widely-spread spam campaign using an attached PDF or XLS file to bypass anti-spam filters would score zero. Sophisticated financially-motivated malware would also score zero if it was not part of a *targeted* attack.

Social Engineering Sophistication

To measure the targeting sophistication base value we assign a score that ranges from zero to five, which rates the social engineering techniques used to persuade a victim to open a malicious link or attachment. This score considers the content, presentation, and claimed sender identity of the email. This determination also includes the content of any associated files, as malware is often implanted into legitimate relevant documents to evade suspicion from users when the malicious documents are opened. The features for each score are detailed in Table 7 (for examples of emails with each of these scores see Appendix A).

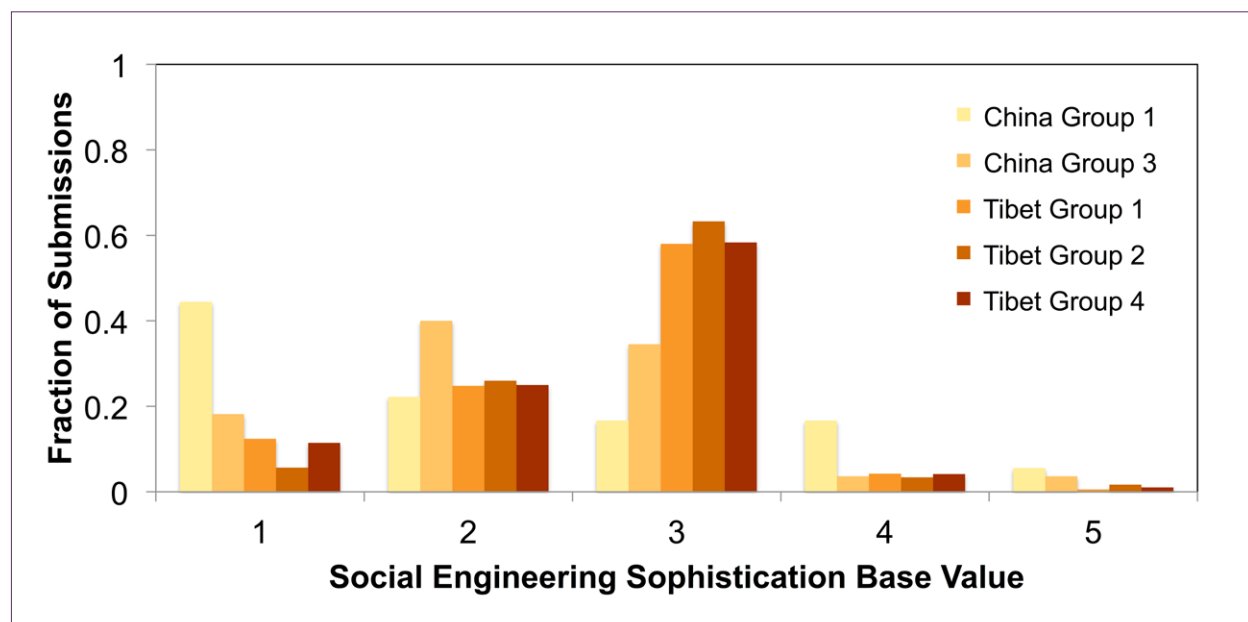
4 For further details on the TTI including detailed discussion of its design, limitations, and plans for future work see: Hardy, S., Crete-Nishihata, M., Kleemola, K., Senft, A., Sonne, S., Wiseman, G., Gill, P., Deibert, R. "Targeted Threat Index: Characterizing and Quantifying Politically-Motivated Targeted Malware." USENIX Security 2014. <https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-hardy.pdf>

TABLE 7: TTI base value score

VALUE	DESCRIPTION
0	Not targeted <ul style="list-style-type: none"> . Recipient does not appear to be a specific target. . Content is not relevant to recipient. . The email is likely spam or a non-targeted phishing attempt.
1	Targeted, Not customized <ul style="list-style-type: none"> . Recipient is a specific target. . Content is not relevant to recipient or contains information that is obviously false with little to no validation required by the recipient. . The email header and / or signature do not reference a real person or organization.
2	Targeted, Poorly customized <ul style="list-style-type: none"> . Recipient is a specific target. . Content is generally relevant to the target but has attributes that make it appear questionable (e.g., incomplete text, poor spelling and grammar, incorrect addressing). . The email header and / or signature may reference a real person or organization.
3	Targeted, Customized <ul style="list-style-type: none"> . Recipient is a specific target. . Content is relevant to the target and may repurpose legitimate information (such as a news article, press release, or a conference or event website) and can be externally verified (e.g., message references information that can be found online). Or, the email text appears to repurpose legitimate email messages that may have been collected from public mailing lists or from compromised accounts. . The email header and / or signature references a real person or organization.
4	Targeted, Personalized <ul style="list-style-type: none"> . Recipient is a specific target. . Email message is personalized for the recipient or target organization (e.g., specifically addressed or referring to individual and / or organization by name). . Content is relevant to the target and may repurpose legitimate information that can be externally verified or appears to repurpose legitimate messages. . The email header and / or signature references a real person or organization.
5	Targeted, Highly personalized <ul style="list-style-type: none"> . Recipient is a specific target. . Email is individually personalized and customized for the recipient and references confidential, sensitive information that is directly relevant to the target (e.g., internal meeting minutes, compromised communications from the organization). . The email header and / or signature references a real person or organization.

Figure 6 shows the targeting score for organizations in our study that submitted at least 50 emails. We can see that attackers targeting these groups put significant effort into the message lures. In particular more than half of the messages targeting the Tibet Groups in Figure 6 have a targeting score of 3 or higher. This result means threat actors are taking care to make the email appear to come from a legitimate individual or organization, and include relevant information (e.g., news reports or exchanges from public mailing lists). Higher targeting scores, which result from actions such as personalizing lures to an individual in the group, or including information that requires prior reconnaissance, were rare, but we nevertheless observed cases. For example, in the case of China Group 3, we observed an email that claimed to be from one of the organization's funders and referenced a specific meeting they had planned that was not public knowledge (social engineering score: 5).

FIGURE 6 : Social engineering base value of emails submitted per group (minimum 50 submissions)



Technical Sophistication

The technical sophistication multiplier ranks the relative technical sophistication of malware. This score is determined by measuring how well the payload of the malware conceals its presence on a compromised machine. We use a multiplier because advanced malware requires significantly more resources to customize for a particular target.

We focus on the level of obfuscation used to hide program functionality and avoid detection for the following reasons:

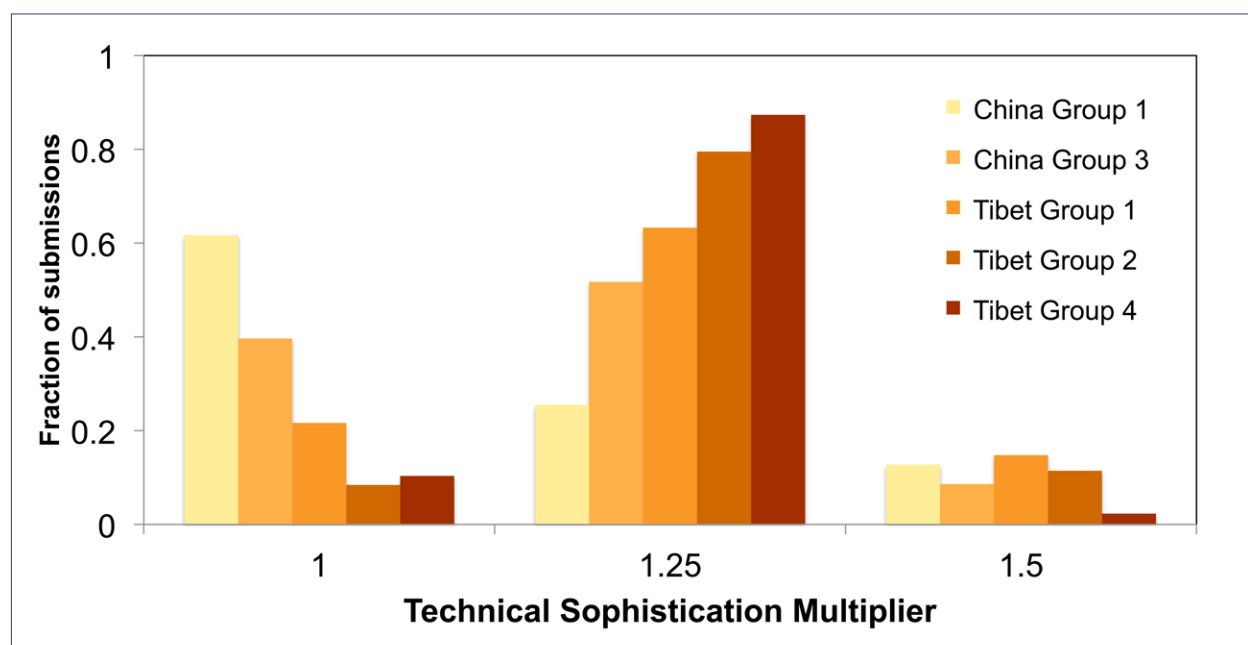
- It allows the compromised system to remain infected for a longer period;
- It hinders analysts from dissecting a sample, developing instructions to detect the malware, and disinfecting a compromised system; and
- Since most commonly used remote access trojans (RATs) have the same core functionality (e.g., key-logging, running commands, exfiltrating data, controlling microphones and webcams, etc.) the level of obfuscation used to conceal what the malware is doing can be used to distinguish one RAT from another.

TABLE 8: TTI technical sophistication multiplier

VALUE	DESCRIPTION
1	Not protected The sample contains no code protection, like packing, obfuscation (e.g., simple rotation of interesting or identifying strings), or anti-reversing tricks.
1.25	Minor protection The sample contains a simple method of protection, including: code protection using publicly available tools where the reversing method is available (e.g., UPX packing); simple anti-reversing techniques like not using import tables, or a call to <code>IsDebuggerPresent()</code> ; self-disabling in the presence of antivirus software.
1.5	Multiple minor protection techniques The sample contains multiple distinct minor code protection techniques (anti-reversing tricks, packing, virtual machine / reversing tools detection) that require some low-level knowledge. This level includes malware where code that contains the core functionality of the program is decrypted only in memory.
1.75	Advanced protection The sample contains minor code protection techniques along with at least one advanced protection method such as rootkit functionality or a custom virtualized packer.
2	Multiple advanced protection techniques The sample contains multiple distinct advanced protection techniques (e.g., rootkit capability, virtualized packer, multiple anti-reversing techniques), and is clearly designed by a professional software engineering team.

Figure 7 shows the technical sophistication multiplier values for emails submitted by the different organizations in our study. Our results show that malware used to target the groups in our study was relatively simple. The highest multiplier value we observed is 1.5 and even that value is seen infrequently. The majority of malware observed is rated either 1 or 1.25 according to our technical scoring criteria, with Tibet Groups observing a higher fraction of malware rated 1.25 and China Groups observing a higher fraction rated 1.

FIGURE 7: Technical sophistication multiplier of emails submitted per group (minimum 50 submissions)



Targeted Threat Index Results Overview

The TTI metric can help us better characterize the relative threat posed by targeted malware in several ways. Table 9 shows the technical sophistication multiplier and maximum / minimum TTI scores for malware families observed in our dataset. Since we primarily find simple malware, with a technical sophistication multiplier of 1 or 1.25, this value does a poor job of differentiating the threat posed by the different malware families to the CSOs. However, by incorporating both the technical sophistication and targeting base value into the TTI metric, we can gain more insights into how effective these threats are in practice.

If we consider the malware families with the highest technical sophistication, we can

see that their TTI values are relatively low, with scores mostly ranging from 1.5 to 4.5 (and one notable exception of 7.5). These tend to be malware families that are regularly used in targeted malware campaigns known to researchers. In particular, PlugX and PoisonIvy [have been found](#) used together in targeted attacks, and PlugX is still in active use and under continuous improvement. Despite their technical sophistication, these threats are not well executed and pose less of a risk to CSOs in which users may be able to identify and avoid these threats.

In contrast, the top five malware families in terms of TTI have lower technical sophistication multipliers (1.25) but much higher levels of social engineering. A notable exception is one highly targeted attack (social engineering score 5.0) that used PlugX (technical sophistication score 1.5) resulting in a TTI value of 7.5 (the highest score in the dataset). While this attack has a higher technical sophistication score than the top five malware families, the high TTI score is due to the level of targeting.

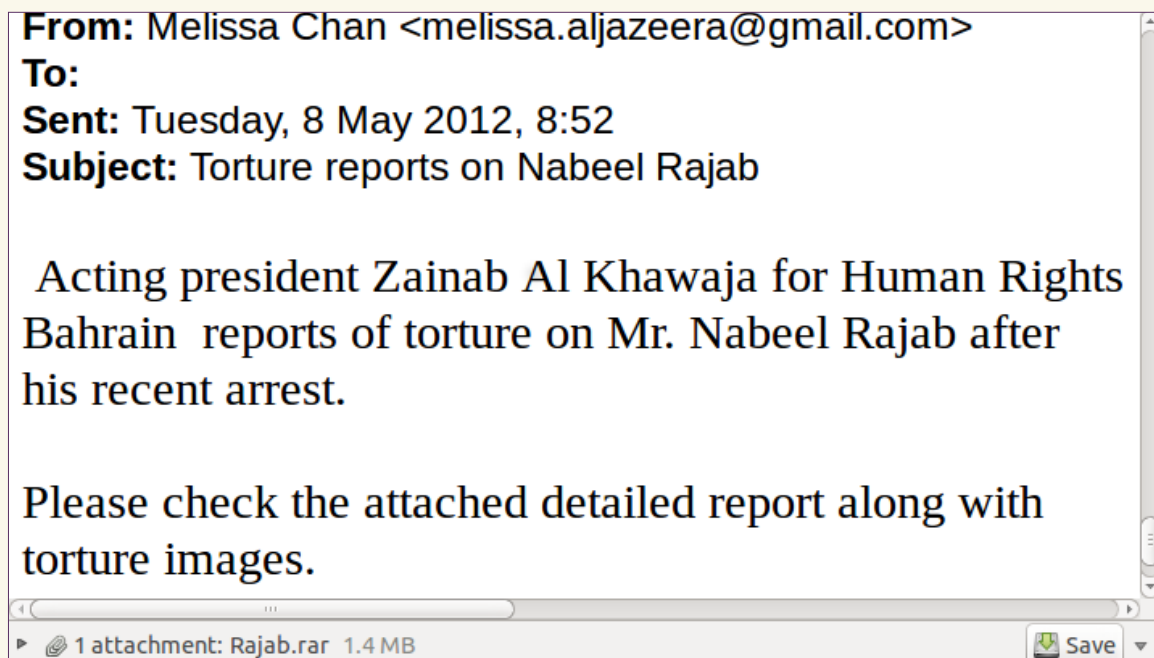
TABLE 9: Top malware families in our dataset by technical sophistication multiplier and final TTI score

TECHNICAL SOPHISTICATION		
Family	Max TTI	Technical Sophistication
PlugX	7.5	1.5
Gh0st RAT (LURKO), ShadowNet	6.25	1.25
Conime, Duojeen, IEXPLORE, GLASSES, cypid, Enfal, Surtr, Vidgrab	5	1.25
Cookies	5	1.0
TTI		
Family	Max TTI	Technical Sophistication
3102	3	1.5
nAspyUpdate	1.5	1.5
PlugX	7.5	1.5
PoisionIvy	3	1.5
WMIScriptKids	3	1.5

ANALYZING COMMERCIAL SPYWARE WITH THE TTI

Attacks using advanced commercial spyware such as FinFisher and DaVinci RCS do not necessarily rank higher on the TTI.

We analyzed a sample of FinFisher used against [Bahraini activists](#) and evaluated it with the TTI. The malware sample is technically advanced, scoring a 2.0, as a result of multiple advanced protection techniques, including a custom-written virtualized packer, MBR modification, and rootkit functionality. However, the email used in the attack is poorly customized and has several attributes that made it look suspicious to the intended target. The email attempts to reference an NGO called Bahrain Center for Human Rights, but mistakenly refers to it as “Human Rights Bahrain.” The message also lists the wrong name for the acting president of the group. It appears to come from a real journalist, Melissa Chan of Al Jazeera, but provides a suspicious gmail address (melissa.aljazeera@gmail.com). These attributes give the email a social engineering base value of 2. As a result, the attack scores an overall TTI score of 4.0, which is relatively low compared to many other attacks seen in our study. This result shows the importance of social engineering tactics: FinFisher is only effective if it is surreptitiously installed on a user's computer, which in some cases requires opening a malicious file (however, both FinFisher and Hacking Team offer optional [network injection products](#) that permit remote attackers to infect a device without user interaction).



ANALYZING COMMERCIAL SPYWARE WITH THE TTI (CONT'D)

Similar results can be observed with respect to attacks using DaVinci RCS, developed by Italy-based company Hacking Team, which has been used against activists and independent media groups. RCS also scores a 2.0 on our technical sophistication scale. We analyzed a targeted attack using RCS against a [dissident](#) in the United Arab Emirates. The email appears to come from “Arabic Wikileaks” (arabic.wikileaks@gmail.com) and asks the recipient to read a “very important message.” Again, while the malware used in these attacks is technically sophisticated, the social engineering lure is poorly customized (social engineering base value 2), resulting in an overall TTI score of 4.0.

From: **ARABIC WIKILEAKS** <arabic.wikileaks@gmail.com>
Date: 2012/7/23
Subject: هام جداً للإطلاع
To:

هام جداً للإطلاع وإبداء الرأي
ولكم الشكر

These results suggest that different threat actors possess varying levels and types of resources, and as a result use different attack methods. The majority of malware submitted in our study appears to be from actors that have in-house malware development capabilities, and the capacity to organize targeted campaigns. However, as this report shows, in many cases they spend significant effort on social engineering, but generally do not use technically advanced malware. Conversely, operators of FinFisher and DaVinci RCS have purchased advanced malware products, but in some cases paired them with relatively unsophisticated social engineering.

EXTENDED ANALYSIS:

2.2

Cluster Analysis

Targeted malware attacks are typically not discrete events. Rather, they are a part of systematic campaigns that use common malware, C2 infrastructure and social engineering tactics to target groups repeatedly over long periods of time. Threat actors using common tools and techniques may target multiple groups within a community.

To cluster attacks into campaigns, we analyze technical and contextual indicators to identify patterns. Where possible we draw connections between these clusters and previously reported campaigns and threat actors.

Malware attacks are clustered into campaigns by commonalities and patterns across the following indicators:

- Email headers: Originating-IP or common email addresses seen in Reply-To, Sender or Envelope-From email headers.
- Shared C2 infrastructure: Domain names and IP addresses to which malware beacons and/or from which it downloads additional modules.
- Static analysis: Commonalities between unusual strings or data structures seen in the malware samples or the files they drop.
- Malware development: Observable changes made to specific malware families over time.
- Social engineering tactics: Contextual patterns in targeted organizations, spoofed senders, and content of messages.

Based on the attributes described above, we identify 10 distinct campaigns, which we present in detail in the following sections.

IEXPLORE Campaigns

First Seen	August 3, 2010
Last Seen	May 21, 2012
Exploits	Windows: CVE-2010-0188; CVE-2010-3333
Malware Families	Windows: IEXPLORE RAT (aka C0d0s0)
Infrastructure	C2 domains: sixday.wikaba.com msupdate02.selfip.com msupdate02.selfip.info xinxin20080628.gicp.net humanbeing2009.gicp.net saveworld.gicp.net xinxin20080628.gicp.net 204.134.116.229 60.167.78.229 116.226.49.148 123.147.81.121 204.134.116.229
Targeted Groups	Tibet Group 1, Tibet Group 2, China Group 1, China Group 2
TTI range	2.5 - 5.0

BACKGROUND

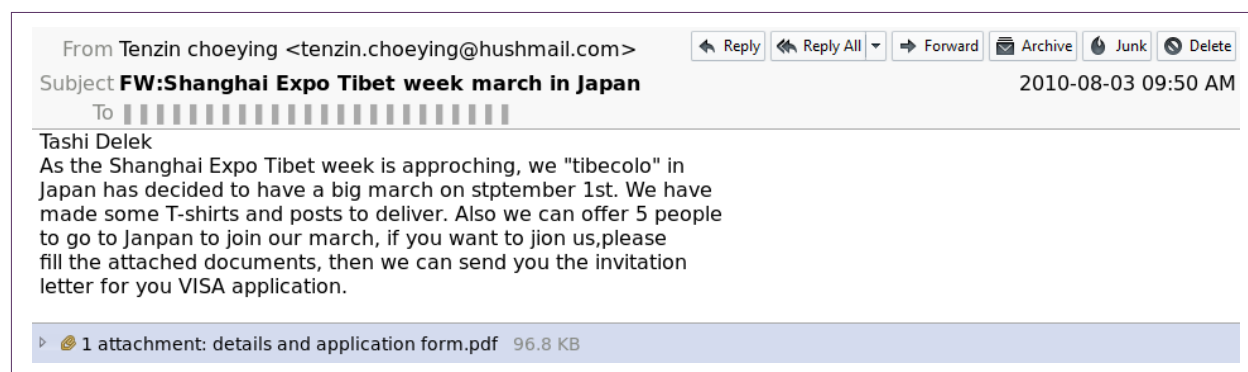
The IEXPLORE campaigns involved custom-developed Windows malware targeting four of the study groups with a unique email and delivery method used for each attempt. Each email was tailored specifically for the target in terms of subject, content, and the way the malware was attached and hidden. In addition, there was evidence that the malware was under active development during the campaign. The IEXPLORE campaigns serves as a typical example of “APT”- style operations.

CAMPAIGN TIMELINE

Attacks in this campaign are linked by the use of IEXPLORE RAT, which provides standard RAT functionality, including keylogging, file extraction, and control of microphone and webcam peripherals.⁵

We identified the IEXPLORE campaign through analysis of three separate attacks using this malware that were sent to China Group 1, China Group 2, and Tibet Group 1. This first series of attacks clearly shows how the attackers carefully customized social engineering tactics to the interests of the three different groups.

Evidence of this campaign first emerged in an August 3, 2010 email to Tibet Group 1 that referenced a protest against the Shanghai Expo in Japan. The malicious attachment was a PDF using CVE-2010-0188 to deploy IEXPLORE RAT.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	6b04821ad588b0f918318064a07dd5d6
C2	msupdate02.selfip.com

⁵ For a detailed technical analysis of the IEXPLORE RAT including a full enumeration of its commands and C2 communication protocol see Hardy, S. "IEXPLORE RAT," The Citizen Lab, August 2012, https://citizenlab.org/wp-content/uploads/2012/09/IEXPLORE_RAT.pdf

On November 11, 2010, China Group 1 received multiple emails addressed to the organization's director claiming to be from personal friends. The emails included an executable attachment in a password-protected archive, with the password provided in the body of the email. Packaging attachments in a RAR file makes them less likely to be discovered by an AV scanner. Password protecting the archive reduces the chances of AV detection even further. When executed, the malware connected to softwareupdate.8866.org (119.75.218.45). The level of personalization used in the message gives it a social engineering score of 4 and a total TTI of 5.0.

On November 19, 2010, China Group 2 received an email containing a story about a high-profile, high-rise apartment building fire in Shanghai. The message was written in Chinese and repurposed text from a news article on the event.

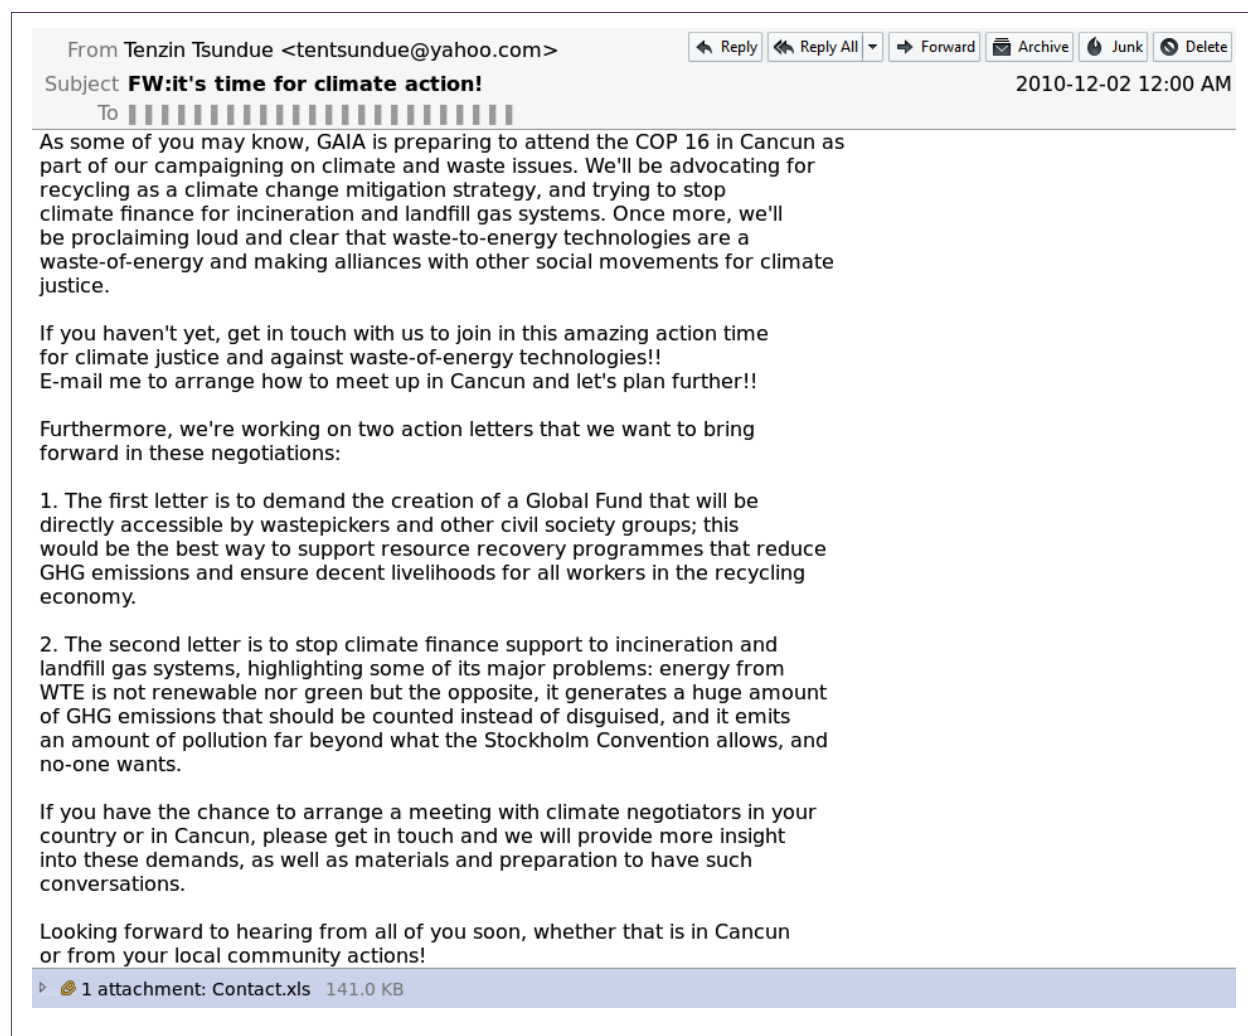
Attached to the email were four images and two executable files (.scr extensions) designed to look like images using the Unicode right-to-left override character. When each executable file is run, it will install and launch the malware, drop an image, open the image, and delete itself. The malware connects to xinxin20080628.gicp.net (114.60.106.156). The attack has a social engineering score of 3 and a total TTI of 3.75.

FIGURE 8: Image of a high-rise fire used to trick recipients into running the malware



The remaining attacks we analyzed targeted Tibetan groups exclusively.

On December 2, 2010, Tibet Group 1 received an email that included an Excel spreadsheet attached to an email that appeared to be from organizers of a conference on climate change.



EXTENDED ANALYSIS: 2.2 Cluster Analysis

Social engineering	3
Technical	1.25
TTI	3.75
MD5	8d4e42982060d884e2b7bd257727fd7c
C2	60.167.78.229

On January 11, 2011, Tibet Group 1 received an email about an annual review of Tibetan human rights issues that contained an executable file designed to appear to be a video of a speech by HHDL.

From Guzel Makan <guzelmakan70@ymail.com>

Reply
Reply All
Forward
Archive
Junk
Delete

Subject **Fw: 2010--the Tibet Human Rights Year in Review**
2011-01-11 10:07 PM

To [REDACTED]

----- Forwarded Message -----

From: Dalai Lama <info@dalailama.org>
To: Watan <watanqim@yahoo.com>
Sent: Mon, January 9, 2011 3:00:09 AM
Subject: Someone has sent you a message from Dalai Lama

Message from sender:
Press statement of Dalai Lama on 2010--the Tibet Human Rights Year in Review
By Henryk Szadziewski
Manager, Tibet Human Rights Project

1 attachment: Speech video of Dalai Lama.rar 68.0 KB

Social engineering	2
Technical	1.25
TTI	2.5
MD5	21a1ee58e4b543d7f2fa3b4022506029
C2	61.132.74.68

In July 2011, IEXPL0RE was sent to Tibet Group 1 again. This time it used a .rar archive file containing a malicious .hlp file.

Tibet Group 1 received two more emails with IEXPL0RE in late December 2011 and early January 2012. On December 22, an email referencing Uyghur refugees

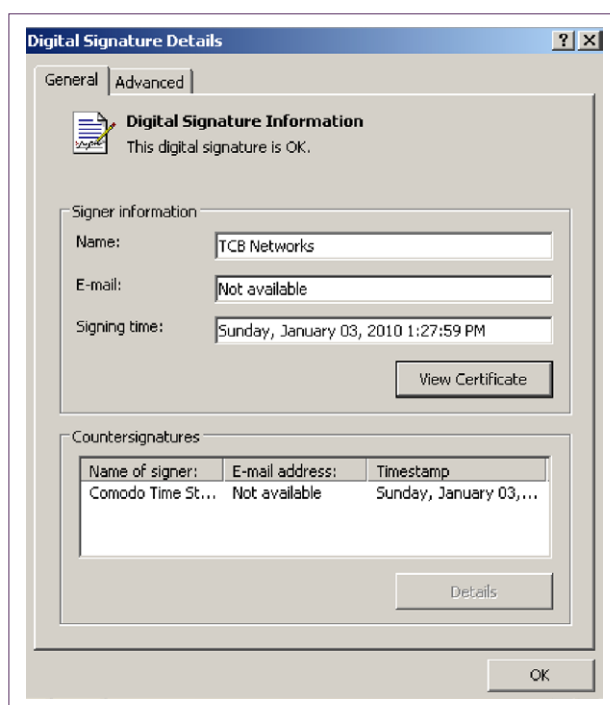
was received; on January 6, an email in Chinese about Taiwan issues copied from a blog post was received. Both had .rar attachments with the same file, which used CVE-2010-3333. Each of these emails score a social engineering base value of 2.0, and a technical score of 1.25 for an overall TTI of 2.5.

On May 21, 2012, a newer version of the RAT payload was distributed via email in multiple RTF documents to Tibet Group 2. This attack was coupled with a higher degree of social engineering. Separate emails with the same payload and content were sent to both the Executive Director and Program Coordinator, addressing them by name. The email contained an invitation to a legitimate event and included the email signature of a real person, with an attached file purporting to contain information about the event. The sender notes that the recipients were identified as contacts for HHDL, and asks them for help contacting His Holiness in order to invite him to the event. This level of personalization gives the attack a social engineering score of 4 (total TTI 5.0)

The attached RTF file drops a DLL alongside a legitimate program vulnerable to DLL hijacking, allowing the malware to run without a warning to the victim that it is not digitally signed. [Strokelt](#), a program for using mouse gestures, uses a file named config.dll without verifying the authenticity of the file. By replacing config.dll with the RAT downloader, the malicious code is run

while appearing legitimate to the operating system (see Figure 9).

FIGURE 9: Valid digital signature for Strokelt program, which is used to launch malicious config.dll file



OBSERVATIONS

This series of attacks represents a typical ‘APT’-style campaign. Multiple groups were targeted, with each attack custom developed for each group, including tailored social engineering. The evolution of the RAT payload, as seen in the series of samples targeting Tibet Groups, suggests that the malware was under active development. The social engineering tactics and development cycles observed in this campaign demonstrates the organized and persistent nature of the attackers.

Mobile Malware

BACKGROUND

The use of malware targeting mobile platforms in espionage campaigns is relatively rare, but is likely a vector that will become more common due to the increasing ubiquity of mobile computing.

During investigations of C2 servers associated with the [Luckycat campaign](#), [Trend Micro](#) found two malicious Android APKs in early stages of development that could collect device information, as well as download and upload files by remote command. Based on the available information, it was unclear how the attackers intended to deliver the mobile malware to targets.

In 2013, researchers at [Kaspersky](#) reported the compromise of an email account of a high-profile Tibetan activist that was then used by attackers to send targeted malware to the activist's contacts. The emails referenced the World Uyghur Congress and included a malicious APK file that appeared to be an application with information on the event. The malware allowed attackers to collect data from infected devices including contacts, call logs, SMS messages, geolocation, and phone data (phone number, OS version, phone model, and SDK version).

Researchers in our group have also found evidence of commercial surveillance products that target multiple mobile platforms (e.g., Android, IOS, BlackBerry, Symbian) developed by [Hacking Team](#) and [FinFisher](#).

In other recent work, researchers found that participants in the Occupy Central protests in Hong Kong [received links through WhatsApp](#) to an Android application that appeared to be associated with the protest organizers, but was actually malware that could send a variety of information back to attackers.

In our study, we identified the use of compromised Android applications sent as part of a targeted attack against a prominent figure in the Tibetan community. This attack lever-

aged a genuine email that was likely exfiltrated by attackers, and attached compromised versions of the chat application KakaoTalk and mobile radio application TuneIn.⁶

VECTOR OF ATTACK

On December 4, 2012, an information security expert who works within the Tibetan community sent a private email to a member of the Tibetan Parliament in Exile, based in Dharamsala, India. That email attached genuine versions of Kakao Talk⁷ and TuneIn⁸ as APK files.

On January 16, 2013, an email purporting to be from this same information security expert was sent to a high profile political figure in the Tibetan community. The email contained the same text as the message from December 4, but attached compromised versions of the Kakao Talk and TuneIn Android APKs.

In order for the malware to be installed, the user must permit applications to be installed from sources other than the Google Play store. This permission is not enabled by default in Android. However, as many members of the Tibetan community (particularly those living in Tibetan areas in China) have restricted access to the Google Play service, they are required to permit applications to be installed from outside sources. It is common for APKs to be circulated outside of Google Play. In addition to permitting the “allow from unknown sources” option, the user must also approve the additional permissions requested by the application. Users may be duped into accepting these permissions by assuming they are required for the regular functionality of the application or by not reviewing them carefully before approving. Once these permissions are approved, they are used to authorize the additional data-gathering capabilities of the malware, which is configured to autostart on device boot.

We later confirmed that the original recipient of the legitimate email had his email account compromised. Therefore, it appears likely that the attackers harvested the

6 We previously reported this attack in a blog post, Citizen Lab, “Permission to Spy: An Analysis of Android Malware Targeting Tibetans,” April 18, 2013, <https://citizenlab.org/2013/04/permission-to-spy-an-analysis-of-android-malware-targeting-tibetans>. See a Tibetan translation of this post here: <https://citizenlab.org/2013/04/android-malware-in-tibetan>.

7 KakaoTalk is a chat app that is developed by a South Korean company (Kakao Corporation). Members of the Tibetan community have used KakaoTalk and other applications as alternatives to WeChat (another chat app popular in Asia) after concerns were raised regarding that application’s general security and the potential for Tencent (the China-based developer of WeChat) to monitor users at the behest of the Chinese government.

8 TuneIn is a media player application for listening to Internet Radio. TuneIn is used by Tibetans to listen to streams such as Voice of America’s Tibetan service, to engage with their culture, and to stay on top of world news.

original email from the compromised account, and over the course of a few weeks developed malicious versions of the attached APKs. The use of private information in this attack gives it a social engineering score of 5. The technical score of the malware is 1.25 (see the section below for details on the malware's functionality). The total TTI is 6.25.

MALWARE ANALYSIS

The functionality and certificates used for the malicious versions of the KakaoTalk and TuneIn APKs are identical. Both applications were repackaged into modified APKs and signed with an illegitimate certificate (KakaoTalk malware MD5 cbc474e34f26b4afd-02932d8cae9e401 Tunein Malware MD5 ba760392f171e2f05d0352cc1e00190c). Below, we reproduce the original and fake certificates used for KakaoTalk. Notice that fields in the illegitimate certificate have been populated with what appears to be an assortment of nonsensical characters from a QWERTY keyboard:

Original legitimate certificate:

```
Owner: OU=kakaoteam, O=kakao, C=ko  
Issuer: OU=kakaoteam, O=kakao, C=ko  
Serial number: 4c707197
```

Illegitimate certificate:

```
Owner: CN=qwe, OU=asd, O=zxc, L=rty, ST=fgh, C=vbn  
Issuer: CN=qwe, OU=asd, O=zxc, L=rty, ST=fgh, C=vbn  
Serial number: a3e5475
```

The following permissions are added by the malware that do not exist in the legitimate version:

```
android.permission.GET_ACCOUNTS
android.permission.ACCESS_NETWORK_STATE
android.permission.READ_SMS
android.permission.INTERNET
android.permission.ACCESS_FINE_LOCATION
android.permission.WRITE_SETTINGS
android.permission.WRITE_SECURE_SETTINGS
android.permission.WRITE_APN_SETTINGS
android.permission.MOUNT_UNMOUNT_FILESYSTEMS
android.permission.PROCESS_OUTGOING_CALLS
android.permission.DEVICE_POWER
adnroid.permission.ACCESS_CHECKIN_PROPERTIES
android.permission.INTERNET
adnroid.permission.CHANGE_WIFI_STATE
android.permission.MODIFY_PHONE_STATE
android.permission.BLUETOOTH_ADMIN
android.permission.BLUETOOTH
android.permission.BIND_DEVICE_ADMIN
android.permission.USES_POLICY_FORCE_LOCK
android.permission.CHANGE_CONFIGURATION
```

Note that two of the additional permissions requested by the malware are misspelled, rendering these permissions unusable:

```
adnroid.permission.ACCESS_CHECKIN_PROPERTIES
adnroid.permission.CHANGE_WIFI_STATE
```

The malicious versions of both applications have the same functionality enumerated below:

- On a periodic basis the user's contacts, call history, SMS messages, and cellular network configuration are written to an encrypted file called info.txt.
- The malware periodically contacts the C2 server "android.uyghur.dnsd.me" to retrieve updated configuration information, such as URLs and login credentials. This configuration information directs the malware to an upload location for the info.txt file. The site hosting the C2 appears to emulate the appearance of the Baidu website (a Chinese search engine), but includes encrypted configuration data hidden in the comments. By masking the C2 as a seemingly innocuous website, requests would appear to be legitimate on casual inspection. The configuration data contained in the comments directs the malware to upload captured data from the device to an FTP server and contain a pointer to a new C2 that would allow the attackers to change the C2 should that need arise.
- The malware intercepts SMS messages and searches for a special code sent by the attacker, which, if detected, responds to the sender with the base station ID, tower ID, mobile network code and mobile area code of the infected phone in question. This message is not displayed to the user, and they are never made aware of it.

OBSERVATIONS

The compromised Android applications that we detected as part of our study, as well as mobile malware described by other security researchers, show that mobile devices are indeed targets for espionage attackers. These attacks serve as early examples of a trend that seems likely to grow alongside the rapid spread of mobile computing.

As described above, there are particular security risks for users residing in locations where access to standard secure channels for installing mobile applications is restricted. As users are required to distribute and install APKs of unknown provenance, they are at increased risk of malicious applications, particularly if those applications use fake certificates (as was the case in this attack).

OS X Campaigns

First Seen	May 2011
Last Seen	Early 2013
Exploits	CVE-2009-0563; CVE-2011-3544; CVE-2012-0507 ; CVE-2009-3129
Malware Families	Revir/IMuler, Olyx / Lamadai / PubSab, MacControl
Infrastructure	C2 domains: freeavg.sytes.com (Olyx.C), mail.hiserviceusa.com (Olyx.C), yahoo.xxuz.com (Olyx.C), coremail.info (SabPab.A), rtx556.onedumb.com (SabPab.A), www.teklimakan.org (iMuler), IPs: 112.213.126.118 (Olyx.C), 100.42.217.91 (Olyx.C), 198.74.124.3 (SabPab.A), 199.192.152.100 (SabPab.A), 61.178.77.158 (MacControl)
Targeted Groups	Tibet Group 1, Tibet Group 2, Tibet Group 3, Tibet Group 4
TTI Range	2.0 - 3.75

BACKGROUND

While Windows was the most commonly targeted operating system in our study, it was not the only platform targeted. Malware targeting OS X is increasingly paired with Windows malware, giving attackers a better chance of compromising the machine, whatever the operating system. This approach can take the form of code that determines the target's operating system, such as a web page that uses JavaScript to detect the operating system and then download a cross-platform exploit with appropriate payload.

Four of the Tibet Groups in the study⁹ received targeted malware specifically designed for OS X. Tibet Groups 1, 2, and 4 received emails with malware in an attachment or link. Malware was detected on the network of Tibet Group 3 by a NIDS on their office network.

The OS X malware seen in our study ranges in sophistication from simple programs

⁹ Tibet Groups 1, 2, 3, and 4

that rely entirely on social engineering, paired with targeted emails that are not customized for the target (TTI: 2.0), to moderately customized emails with malware that has minor code protection (TTI: 3.75). While the technical sophistication of the malware does not vary widely, all of the malware families observed show active and consistent development over the course of the study.

MALWARE ANALYSIS

The malicious emails used a combination of social engineering, and exploits against a variety of vulnerabilities, to install malware on the victim's computer.

The vectors we observed include:

- An attached .zip file containing an executable
- An attached Word document using CVE-2009-0563
- A link to a Java .jar file using CVE-2011-3544
- A link to a Java .jar file using CVE-2012-0507

The subject and body text of all of the emails targeting the Tibet Groups contained information relating to Tibetan news and activities (e.g., current world events, upcoming rallies, and self-immolations).

We see Word document vectors first being sent in early 2013. Interestingly, these attacks use a vulnerability made public back in 2009. The use of this vulnerability may be due to the Java vulnerabilities having a higher chance of being patched by the Tibetan community, after they received substantial media attention. However, as the Word documents were all part of one campaign, it is likely just coincidence, as an email carrying the later Java vulnerability was received while the Word campaign was still underway.

We observed three malware families targeting OS X, all of which are simple RATs with low technical sophistication scores:

- Revir/IMuler (technical score: 1.0)
- Olyx/Lamadai/PubSab (technical score: 1.0)
- MacControl (technical score: 1.25)

Revir / IMuler

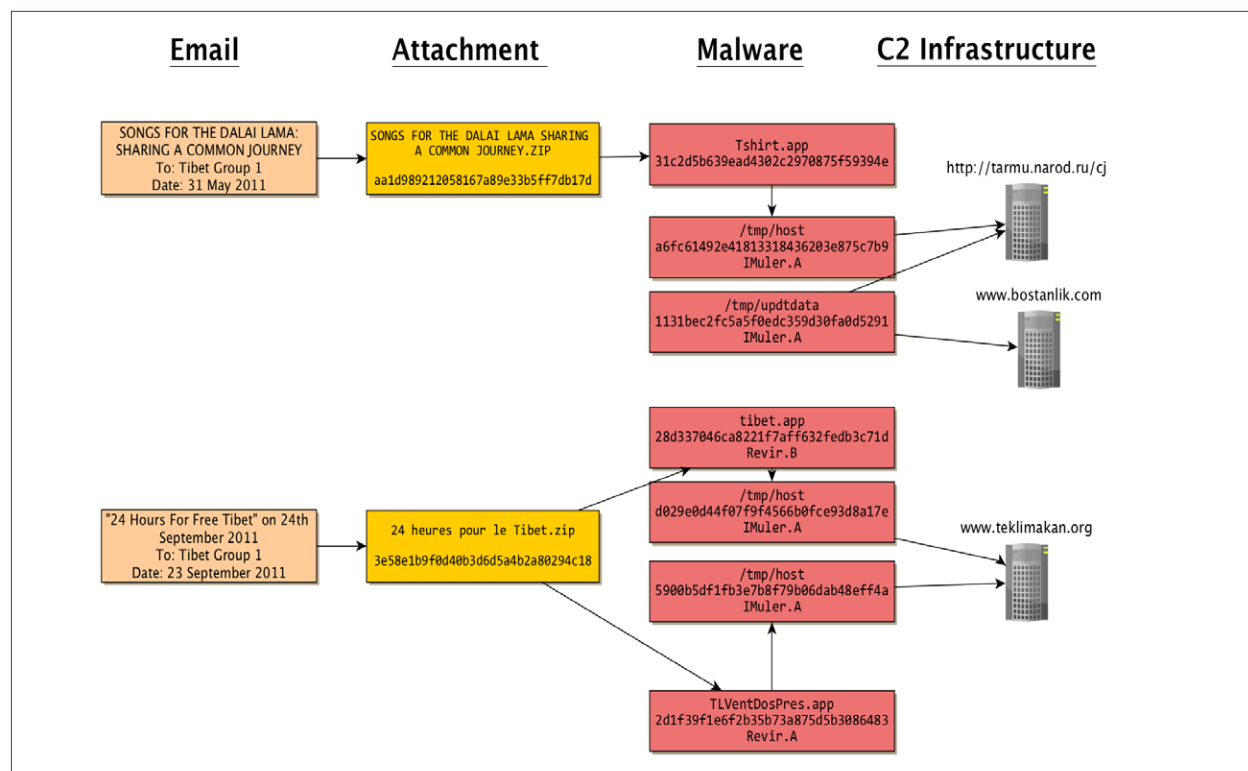
Revir and IMuler are names for individual components of one particular targeted attack for OS X, and are often used interchangeably. Revir is the name for the “dropper” or shell program that deploys the embedded malicious payload (in this case, IMuler), as well as a clean payload that is then opened to distract the user.

The clean payload that is used can differentiate the early variants of Revir that we have seen. Using the [F-Secure naming scheme](#) Revir.A carries a PDF and Revir.B carries a JPG. Later variants Revir.C and Revir.D allow for any type of clean decoy file, and also encrypt the payload.

IMuler acts as a simple remote access trojan, providing the attacker with the ability to upload and delete the victim’s files, download and run additional malware, and take screenshots. The

variants we observed have no reverse engineering protection in the code, although later versions starting with [IMuler.B](#) will look for the Wireshark network analyzer and stop running to evade analysis if it is found.

We observed two attacks against Tibet Group 1 using the Revir/IMuler combination. The first, an email sent in May 2011, was a combination of Revir.B and IMuler.A and was the earliest Mac malware attack seen in the study. This email’s contents were about a legitimate event featuring HHDL. The second email, sent in September 2011, stepped up the attack by containing both Revir.A/IMuler.A and Revir.B/IMuler.A combinations. This email purported to be from a legitimate Tibetan rights organization and referred to an upcoming event.

FIGURE 10: Revir/IMuler attacks

The attack in September 2011 is particularly interesting because the Revir/IMuler components show very clear development progress compared to the version sent in May. The May version is a two-stage program. The initial program dropped as /tmp/host downloads a second program as /tmp/updtdata, which is then used for communication with the C2 server. The September version integrates the second program into the first, merging functionality. This change means that the download of a second executable is not required, eliminating a more suspicious component of the infection process.

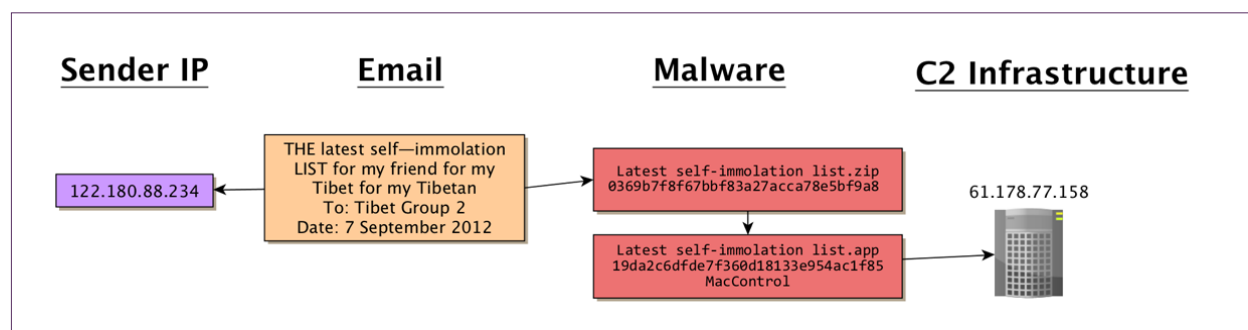
The malware version sent in September 2011 was also used in [another campaign reported by Eset in March 2012](#), using photos of a topless model as the lure to run the attachment.

MacControl

On September 7, 2012, we identified an attack targeting Tibet Group 2 using another malware family, MacControl. The samples seen from this family have a technical score of 1.25

This email repurposed content from Radio Free Asia and claimed to contain a list of self-immolations giving it a social engineering score of 3 (TTI 3.75). The attached executable connects to a C2 server at 61.178.77.158:80 and functions as a standard RAT.

FIGURE 11: MacControl attacks



Another email received by the same organization three weeks later contained a malicious Excel file that installed Gh0st RAT with the variant-identifying text LURK0. This RAT shared the same C2 as the MacControl, connecting back to 61.178.77.158 on port 8080.

This pairing of MacControl with Gh0st RAT has been used in attacks against Uyghur users, as reported by [Kaspersky](#) and [AlienVault](#).

Outside of our study participants, we have also seen MacControl campaigns similar to those reported by Kaspersky and AlienVault, targeting Tibetan and Uyghur communities. These differ slightly than those described above in that they use different flag text in the Gh0st RAT component, and connect to a nearby IP (61.178.77.169). They are also delivered using a Word vulnerability, while the email sent to the in-study recipient contained an executable inside a .zip file.

Olyx / Lamadai / PubSab

Olyx, Lamadai, and PubSab (or SabPub) are variants of the same malware that are differentiated by the C2 server used and the location where the malware hides on a compromised system. These names are often used interchangeably by different antivirus or security companies. Further complicating matters, there is often overlap between names: for example, Olyx.C is the same as Lamadai.B.

Olyx.A

- » Threat location: /Library/Application Support/google/startp
- » Launcher: ~/Library/LaunchAgents/www.google.com.tstart.plist

Olyx.B (Lamadai.A)

- » Threat location: /Library/Audio/Plug-Ins/AudioServer
- » Launcher: ~/Library/LaunchAgents/com.apple.DockActions.plist

Olyx.C (Lamadai.B)

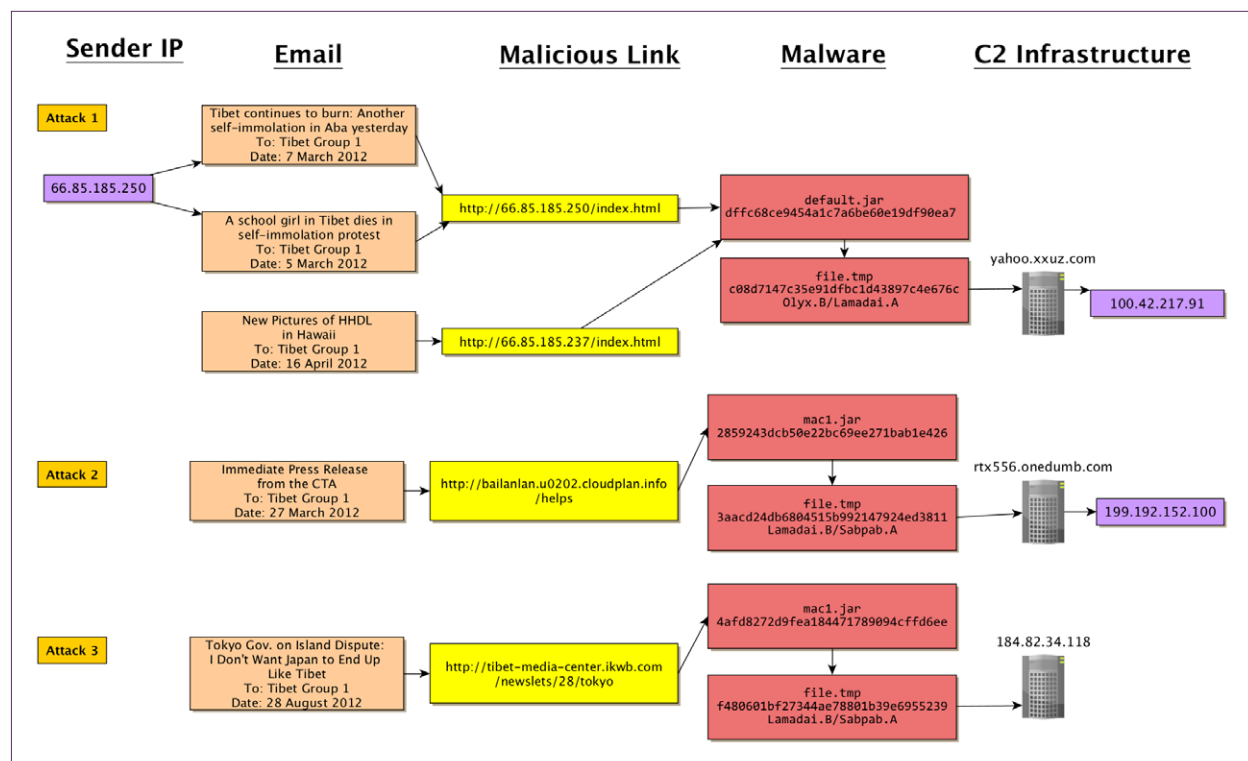
- » Threat location: Applications/Automator.app/Contents/MacOS/DockLight
- » Launcher: ~/Library/LaunchAgents/com.apple.DockActions.plist

PubSab.A

- » Threat location: ~/Library/Preferences/com.apple.PubSabAgent.pfile
- » Launcher: ~/Library/LaunchAgents/com.apple.PubSabAgent.plist

Olyx.C was observed in emails sent to Tibet Group 1, and via a NIDS on the network of Tibet Group 3.

The campaign against Tibet Group 1 consisted of five emails that contained links to malicious .jar files that exploited Java vulnerabilities (CVE-2011-2544 or CVE-2012-0507). All of these emails appeared to come from real people or organizations, and referenced Tibetan themes giving them a social engineering score of 3. The malware is basic with a technical score of 1. The total TTI is 3.

FIGURE 12: Olyx/Lamadai/PubSab attacks

On January 29, 2013, the NIDS on the network of Tibet Group 3 detected evidence of a Java vulnerability being used to serve multi-platform malware pretending to be Adobe Flash Player. This sample was distributed through a web page that does web browser user agent detection. Tibet Group 3 did not submit any emails containing this link, so the specific attack vector used is unclear.

The website hosting the malware was hxxp://services.addons.mozilla.publicvm.com, and had an .xpi file for Firefox and a .jar containing Olyx.C for Mac. The way the malware was served was different than other similar attacks in that it checked both browser and OS, not just OS, to determine which malware program would be used.

On February 5, 2013, we received additional alerts that showed similar malicious pages were visited by Tibet Group 3, again without indication of the original attack vector. A web page was flagged by the NIDS due to a suspicious encoded string that decoded to a tinyurl.com redirector. This link led to a page on hxxp://adobeupdate.publicvm.com, which had attacks for IE, Firefox, Java (Windows, but not OS X), and

a standard Windows binary. There may also have been an OS X attack, but we were unable to locate it from the information recorded by the NIDS.

The original page that served up the attack also had a distinctive comment in the script that identified it as a legitimate script from the website of the US State Department (www.state.gov). At the time of our analysis the server was hosting 100 other domains. The www.state.gov script we found on the page suggests the IP was hosting a fake US State Department website that included the malicious link.

OBSERVATIONS

Mac OS X was once commonly seen as a more secure alternative to Windows. Targeted groups in the Tibetan community shared this assumption. For example, in a 2008 [Washington Post](#) article on targeted attacks against Tibetan groups, a volunteer providing technical assistance to a Tibetan NGO noted that the group had attempted to mitigate attacks by using “more secure platforms such as Apple computers.” While the number of malware vectors targeting OS X is small compared to the many vulnerabilities used against Windows targets, it is clear that OS X malware is becoming an important tool for attackers targeting human rights organizations. All of the malware families described here are under active development and we will likely see more attacks targeting OS X at greater levels of technical sophistication.

DNF Campaigns

First Seen	November 26, 2010
Last Seen	March 4, 2013
Attack Vectors	Targeted malicious emails
Exploits	Windows: CVE-2009-3129, CVE-2011-3544, CVE-2012-0158 Mac: CVE-2011-3544
Malware Families	cxgid, FAKEM (HTML variant), Olyx, Scar.hikn
Infrastructure	C2 domains: www.usciro.com, server.universityexp.com, mail.miyazaki-housou.com, forum.livetldownload.com, forum.mercifulland.com, www.snowhataj.com, www.holyplateau.com, mail.hiserviceusa.com, www.hiserviceusa.com, mail.loveargon.com
Targeted Groups	China Group 1, Tibet Group 1, Tibet Group 2, Tibet Group 4
TTI Range	2.5 - 5.0

BACKGROUND

We identified the Domain Name Family (DNF) campaign by clustering attacks together on the basis of a set of malware families that communicate with domains registered under a series of suspicious names.

Analysis of attacks using Olyx Scar.hikn, cxgid, and FakeM malware families revealed that these samples connect to a set of domains that are registered to series of names: Mily Luna, Philip Fischer (adonis.fischer@yahoo.com), William Bottle (john.felder@hotmail.com), and XieZhong Customer. Searching through domain registration information revealed a large number of domains registered under these names in a short time frame that were related to Tibet, Japan, education and business. All of these domains used a common hosting company called XinNet. Most of these domains have since expired, but historical registration data can be retrieved using services such as [DomainTools](#).

The registration information provided by “Mily Luna” includes conflicting fields (i.e., address in Nepal, but city and province as Hong Kong SAR), which further demonstrates that this information does not reflect a real user:

Domain Name : miyazakihousou.com
PunnyCode : miyazakihousou.com
Creation Date : 2009-02-01 10:53:24
Updated Date : 2012-02-11 10:51:20
Expiration Date : 2013-02-01 10:47:33

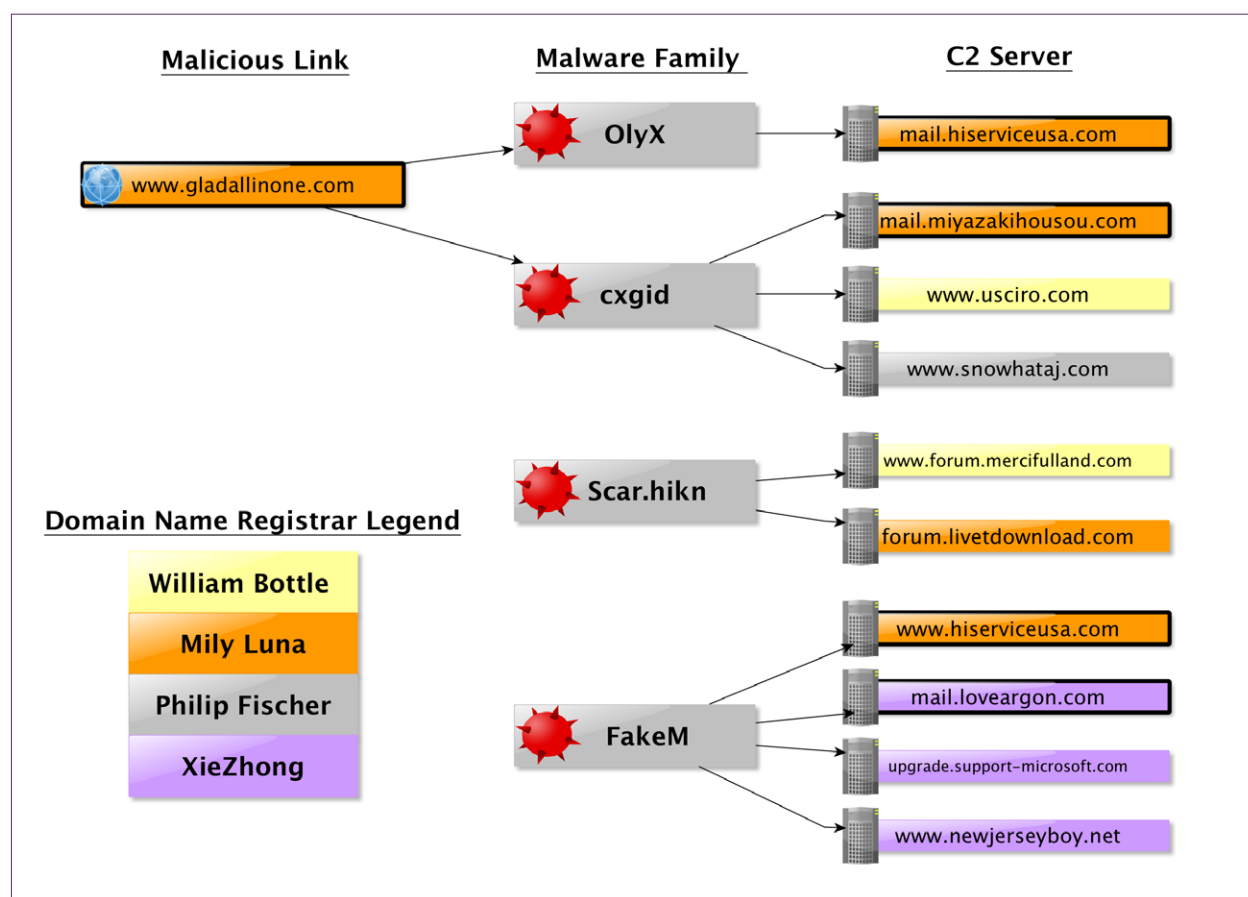
Registrant:

Organization : mily luna
Name : mily luna
Address : No.358,dapho road,Kathmandu, Nepal
City : xianggangtebiexingzhengqu
Province/State : xianggangtebiexingzhengqu
Country : china
Postal Code : 2000000

Some of these C2 domains are registered using email addresses at xiezhong.com. Through domain registration records, we were able to find more than 25 additional domains connected to this cluster, including many registered to “John Smith” (world-freeusa@gmail.com). While we did not see any malicious activity related to these domains, some of the domains are suspiciously named (kaspersky-ru.org, kaspersky-us.org, thetibetpost.net). In the case of kaspersky-us.org in particular, [VirusTotal](#) shows that only 1/51 antivirus products detect the site as malicious, but the one that does is made by Kaspersky. VirusTotal also includes a URL query report showing thetibetpost.net as malicious.

Figure 13 illustrates the connections between malware families, C2 domains, and the domain registrants in the DNF campaign. One FakeM sample used one of the Mily Luna and Xie Zhong domains as C2s. FakeM has been observed being used in conjunction with cxgid by [other researchers](#), but we have not seen other reports identifying the infrastructure found in this cluster.

FIGURE 13: Relationship between domain hosting malicious link, malware family, C2, and name used to register domains



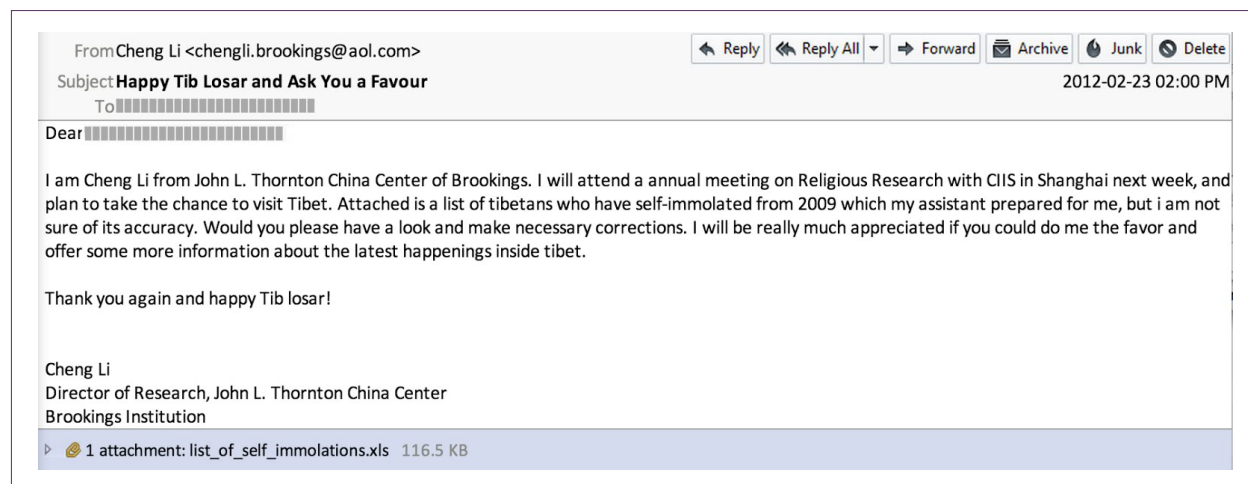
ADAPTIVE ATTACKERS

We observed DNF Campaign attacks between November 2010 and March 4, 2013 that targeted Tibet Group 1, Tibet Group 2, and Tibet Group 4. The social engineering scores of these emails were between 2-3 and the technical scores of the malware were 1.25 (TTI 2.5 – 3.75). In June 2011, we were sent an automated AV detection notice from China Group 1. The alert identified a sample that was also sent in emails to Tibet Group 1, and which connected to the DNF-related domain upgrade.support-microsoft.com. This link suggests a staff member of China Group 1 likely received a malicious email from the DNF campaign and opened the payload, triggering the AV detection.

Initially, attacks in this campaign exclusively used Windows malware. However, the attackers demonstrated their ability to quickly adapt tactics to meet new requirements.

On February 23, 2012, an email was sent to the director of Tibet Group 1. It addressed the director personally and appeared to come from Mr. Cheng Li, a prominent China scholar based at the Brookings Institution. The email requests the assistance of Tibet Group 1 in verifying information on Tibetan self-immolations. The name and title provided in the email all match real details for Mr. Cheng Li provided on his Brookings Institute staff page.

The director of Tibet Group 1 noted to us that at first it was flattering to be asked to consult a well-known China expert on Tibetan issues. However, the director quickly noticed that the email was sent from a suspicious AOL account (chengli.brookings@aol.com). This account appears to have been registered by the attackers for this specific attack. Attached to this email was an Excel spreadsheet that used CVE-2009-3129 to install cxgid malware. The malware connects to mail.miyazakihousou.com (112.213.126.18), which is a domain registered under the name Mily Luna.

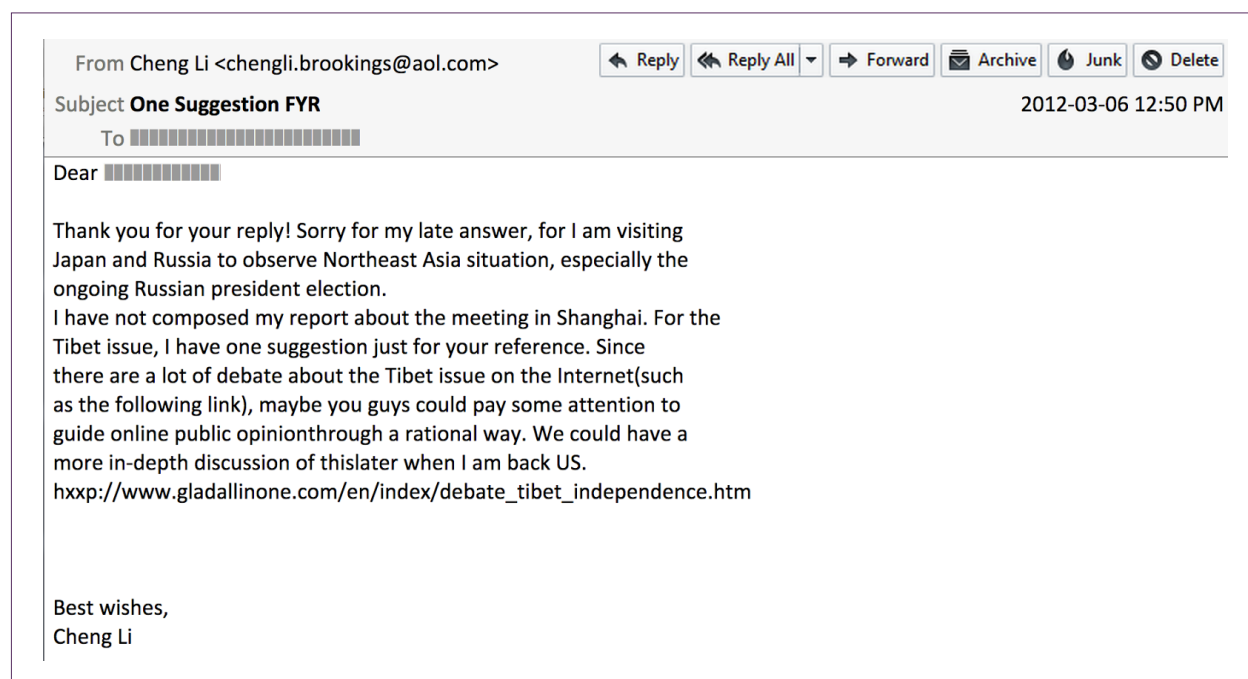


Social engineering	4
Technical	1.25
TTI	5.0
MD5	64e2d3b91977bb0c293cac3e97669f03
C2	mail.miyazakihousou.com (112.213.126.18)

The director of Tibet Group 1 flagged this message to us as one of the most targeted messages they had ever received. After consultations with Tibet Group 1, we decided to run an experiment on the attackers to test how responsive they would be to interaction from a target. Working with the director, we crafted a reply to “Cheng Li” and sent it on March 2:

“Thank you for your inquiry. I’d be happy to help out—I’m having trouble opening the document on my mac though, I think there may be an issue with the Chinese character font? I think if you sent me a Word version that might be easiest, as it would also allow me to make comments in the document.”

On March 6 “Cheng Li” replied, apologizing for his late response due to work-related travel. He encouraged the director to review information on Tibet issues on a website. The link provided pointed to a website containing a Java vulnerability that had payloads for both Windows and OSX systems. The payload for Windows was the same cxgid sample sent in the original email. The payload for OSX was Olyx and connected to mail.hiserviceusa.com (112.213.126.118). Both the malicious website and C2 were domains registered under the name Mily Luna.



Social engineering	4
Technical	1.25 - cxgid sample 1.0 - Olyx sample
TTI	5.0 - cxgid sample 4.0 - Olyx sample
MD5	f9beda8a6eef73f60d3911e890fb11fe - cxgid sample 7f016da6d8fafd03b5bb536ce4106f53 -Olyx sample
C2	mail.hiserviceusa.com (112.213.126.118)

OBSERVATIONS

While other malware campaigns we identified typically use free subdomains, this cluster primarily relied on registered domains. The use of registered domains provided a useful variable around which to cluster attacks. Registered domains can also be blacklisted more easily than free services providing subdomains.

The DNF Campaign also demonstrates the adaptability of the attackers. Upon receiving a message from Tibet Group 1 indicating the targeted user was using a Mac, the attackers quickly responded with malware targeting OS X.

APT 1 Campaigns

First Seen	April 2010
Last Seen	August 16, 2012
Exploits	N/A
Malware Families	Bangat, GLASSES, WARP, WEBC2-QBP
Infrastructure	C2 domains: ash22ld.compress.to, ewplus.com (compromised site), johnbell.longmusic.com, 66.228.132.8 (hard coded ip)
Targeted Groups	Tibet Group 1, Rights Group 1
TTI Range	5.0

BACKGROUND

On February 19, 2013, [Mandiant](#) released a [report](#) that shed light on a prolific cyber espionage group they call APT1 (also referred to as “[Comment Crew](#)” or “[Byzantine Candor](#)”), which had targeted a large number of organizations in a wide range of industries, stealing terabytes of data. Mandiant’s report traced APT1 operations to China and claims that the group may in fact be the Second Bureau of the People’s Liberation Army General Staff Department’s Third Department, also known as Unit 61398.

APT1 has been active since at least 2006. Mandiant has observed the group breaching 141 organizations from 20 major industry sectors. Within Mandiant’s report there is no mention of CSOs as targets among these compromised organizations. However, in previous reports and released datasets, there are indications that civil society is targeted by APT1. Both [Mandiant](#) and [Shadowserver](#) have included a Tibetan-themed domain (tibethome.org) in their APT1-related domain lists, which suggests that Tibet-related organizations may have been targeted, but no further details on Tibet-related operations were included. In 2012, a [Bloomberg article](#) listed the nonprofit organization International Republican Institute among target organizations compromised by APT1 in June 2011, but no technical details of the attack were released.

In the course of our study, we found evidence that APT1 targeted Tibet Group 1,¹⁰ and successfully compromised the networks of Rights Group 1.

TARGETING TIBET GROUP 1

On April 28, 2010, the director of Tibet Group 1 was sent an email from a Yahoo! webmail address. The sender makes a personal plea to Tibet Group 1 to help find his Tibetan wife who he claims went missing since after visiting Tibet.



Social engineering	4
Technical	1.25
TTI	5.0
MD5	6fb3ecc3db624a4912ddbd2d565c4995
C2	ewplus.com (204.14.88.45)

Some details of the email immediately flag it as suspicious: the name in the email address is “Nate Herman” although the email body is signed “Martin Lee.” The forwarded email included full headers, so we were able to obtain more information about its origin (Yahoo! includes the sender’s source IP in the headers when an email is sent over the webmail interface). In this case, the originating IP was 69.95.255.26, which

¹⁰ We originally published analysis of the APT1 related attack against Tibet Group 1 in a blog post, Hardy, S. APT1’s GLASSES – Watching a Human Rights Organization, February 25, 2013, <https://citizenlab.org/2013/02/apt1s-glasses-watching-a-human-rights-organization>

is registered to One Communications, Inc. / EarthLink Business, and is very close to IPs used in a [similar attack](#)—demonstrating that this attack is not isolated, and the IPs are likely being reused for other malware campaigns.

This email contains a link to a ZIP file located at hxxp://tcw.homier.com/attchments/details.zip (MD5: 6fb3ecc3db624a4912d-dbd2d565c4995). The “homier.com” domain belongs to Homier Distributing Company, Inc. and appears to have been compromised. A search for this subdomain shows other instances of malware hosted there, including a case detailed in a [ThreatExpert report](#) describing a malicious file stored in `/images/update.bin`, and [another malicious program](#) getting the file `/attachments/SalaryAdjustment.zip`.

Analysis of the files revealed malware that shares a number of similarities to malware described in Mandiant’s APT1 report that they call “GOGGLES” —a simple downloader that is controlled via encoded markers in files accessed over HTTP. The malware sent to Tibet Group 1 shares both a large percentage of code and the same C2 infrastructure as the program described in the APT1 report, which suggests the two pieces of malware are both used by APT1.

We call this malware GLASSES because it is related to GOGGLES, and used a compromised eyeglasses storefront website as its C2 server. The GOGGLES code is more

sophisticated than the GLASSES code. In addition to a more effective method of hiding the command data, it also has more countermeasures to protect against reverse engineering and hide itself on the infected system. For this reason, it is very likely that GOGGLES is a later version of GLASSES.

COMPROMISING RIGHTS GROUP 1

In August 2012, Rights Group 1 was made aware of a serious compromise of their network infrastructure. Following incident response from a third party, Rights Group 1 shared workstation hard drives with the Citizen Lab that were suspected to have been compromised as part of the intrusion. The attackers had access to the network infrastructure of Rights Group 1 from January 2011 to August 2012. During this time the attackers were able to move laterally through the network, install RATs, extract sensitive data and passwords, and impersonate staff identities. The incident affected computers beyond the ones to which we had access, but these hard drives provide enough data to reveal malware and C2 infrastructure that is linked to APT1.

We conducted forensic analysis of six Windows workstation hard drives used by Rights Group 1 staff members. This analysis found that three of the drives were compromised by multiple versions of malware that matched a tool used by APT1 called Bangat, which is used to

establish footholds in a network and maintain persistence. Bangat has standard backdoor functionality, including features to start keyloggers, gather system information, and take screenshots.

Comparing the samples retrieved from the compromised hard drives to Bangat samples available from the [Contagio APT1 malware collection](#) reveals close similarities. Rights Group 1 samples included the same functionality and important strings as the APT1 contagio samples. These included temporary file names (~MC_[#]~, where # are numbers) and DES key (!b=z&7?cc,MQ>) used for encryption. Binary comparison between the two samples reveals an approximate 97% match.

One of the compromised hard drives included a variant of an HTTP backdoor used by APT1 that Mandiant calls WARP. This malware has no RAT functionality and is primarily used to gather system information and download stage two malware. Therefore, we believe that WARP was used as a dropper to install Bangat onto the compromised system.

Binary comparison between the WARP sample from Rights Group 1 and a WARP sample from the Contagio APT1 malware collection (md5 C0134285A276AB933E-2A2B9B33B103CD) revealed a 90% similarity. The main differences between samples is that the Rights Group 1 sample does not have functions from wininet.dll in the import table, and uses LoadLibrary and GetProcAddress to import them.

All three of the compromised hard drives included samples that communicated to 66.228.132.8 as a C2. This IP address also had an HTML comment on its default webpage that indicated it also served as a C2 for WEBC2-QBP, another malware family described by Mandiant in their [APT1 report](#). The same C2 (66.228.132.8) was also used by two Bangat samples in the Contagio APT1 collection (MD5s BD8B082B7711BC980252F988BB0CA936, E1B6940985A23E5639450F8391820655).

TABLE 10: Overview of malware retrieved from compromised hard drives

HARD DRIVE	MALWARE	FILE CREATION DATE	MD5	C2
HD1	Warp - Dropper for Bangat	n/a	2b941110e046a03894d-41f90272c3012	n/a
HD1	Bangat (irmon32.dll)	May 15, 2012	21afca59b9aaa26676adbf72ccff7b9	hurrissonstone.dnset.com, dynosessfich.myMom.info
HD1	Bangat (Nwsa-pagent.dll)	July 5, 2012	429de63ec18eda4f0699b-2145bab5480	66.228.132.8
HD2	Bangat (rasauto32.dll)	June 11, 2012	45dc7e-b8e76143846f242940ff-369cb4	66.228.132.8
HD2	Bangat (Nwsa-pagent.dll)	June 19, 2012	429de63ec18eda4f0699b-2145bab5480	johnbell.longmusic.com
HD3	Bangat (rasauto32.dll)	June 11, 2012	5dc7e-b8e76143846f242940ff-369cb4	66.228.132.8

OBSERVATIONS

The APT1 campaigns illustrate one of the broader findings of this report. *While large, resourceful threat actors like the APT1 group are frequently documented targeting government and industry, the same actors use similar tools, techniques, and procedures to target CSOs as well.* While government and industry have the resources and expertise to respond to such threats, in many cases CSOs do not. Even large CSOs are vulnerable to this problem. While Rights Group 1 is a large and well-resourced organization relative to others in our study it was compromised for over a year-and-a-half before the threat was identified.

NetTraveler Campaigns

First Seen	April 30, 2010
Last Seen	September 12, 2012
Attack Vectors	Targeted malicious emails, Watering hole attacks
Exploits	Windows: CVE-2009-3129; CVE-2010-0188; CVE-2012-3333; CVE-2012-0158 Mac: CVE-2012-0507
Malware Families	Windows: Conime, Gh0st RAT, RegSubDat, Netpass Mac: Dockster
Infrastructure	Email Sender IPs: 209.11.241.144 C2 domains: 209.11.241.144, akashok.w63.1860host.com:80 (69.43.161.162), ww2.akashok.w63.1860host.com:80 (204.13.161.108), gen2012.eicp.net:1080 (61.178.77.98), 61.178.77.98:8080, 61.178.77.98:1080, 61.178.77.98:80, itsec.eicp.net:443 (1.203.31.195), www.eaglesey.com (120.50.35.46), itsec.eicp.net:8088 (209.11.241.144)
Targeted Groups	Tibet Group 1, Tibet Group 2, Tibet Group 3, Tibet Group 4, Tibet Group 5, China Group 3
TTI Range	2.5 - 4.0

BACKGROUND

In June 2013, Kaspersky released a [report](#) detailing the operations of a threat actor that compromised over 350 victims in 40 different countries. Kaspersky called the main malware used in these campaigns “NetTraveler” after an internal string found in the tool, “NetTraveler is Running!” Targets identified in this report included Tibetan and Uyghur groups, the energy industry, military contractors, scientific research centres, universities, government institutions, and embassies.

The Kaspersky report identifies the IP 209.11.241.144 as a “mothership” server

used as a VPN and a C2 in the campaigns. We see 209.11.241.144 as a sender IP for 19 emails in our study. Searching for other emails that share the same malware we find a total of 34 emails, which we can split into seven campaigns based on how the common C2 infrastructure is used. Additionally, there was one email from this sender that we were not able to cluster, as the attachment was a password-protected ZIP file and the password was not evident. Attacks using infrastructure related to NetTraveler targeted all five of the Tibet Groups in our study as well as China Group 3.

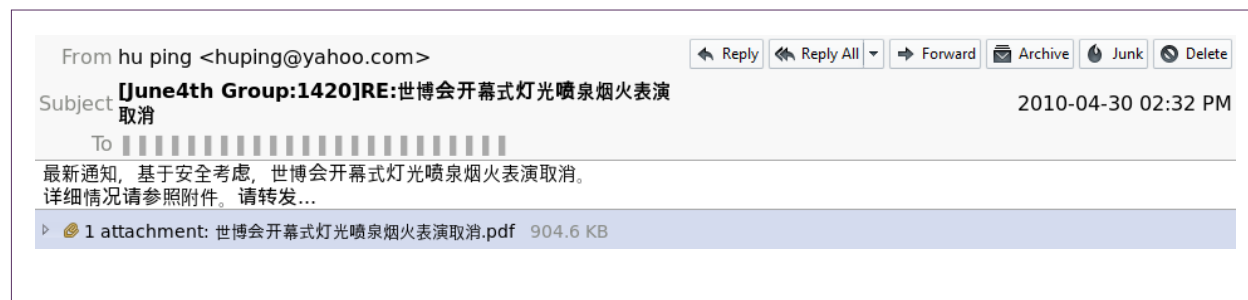
NetTraveler operators are also known to use watering hole attacks against Tibetan websites. In December 2012, F-Secure [reported](#) on malware that relied on an entirely different method of attack and compromise but used the same infrastructure. A website related to HHDL at www.gyalwarinpoche.com was compromised and served the CVE-2012-0507 Java exploit (the same as used in the [Flashback](#) malware) to compromise computers running OS X. This malware, which F-Secure calls Dockster, connects back to the same IP that sent many of the malicious emails we observed, itsec.eicp.net:8088 (209.11.241.144). [Kaspersky](#) has documented similar watering hole attacks against Uyghur-related websites.

CAMPAIGN 1

The first appearance of an attack that used infrastructure related to NetTraveler was sent to China Group 3 on April 30, 2010. The email attached a PDF that used CVE-2010-0188, and connected to C2 servers at akashok.w63.1860host.com:80 (69.43.161.162) and ww2.akashok.w63.1860host.com:80 (204.13.161.108). The sender IP matches the mothership server identified by Kaspersky (209.11.241.144).

C2 traffic from this malware appears as follows:

```
GET /wl/netpass.asp?action=gettext HTTP/1.0
GET /wl/netpass.asp?hostid=...&hostname=...&hostip=...^filename=18155523-sys.log&filestart=0&filetext=begin:...
```



Social engineering	2
Technical	1.25
TTI	2.5
MD5	00403a09dbc87103e2efe060aec07566
C2	akashok.w63.1860host.com (69.43.161.162) ww2.akashok.w63.1860host.com (204.13.161.108)

CAMPAIGN 2

The first of three campaigns using the Conime malware family involved seven emails, five of which were distinct, mostly concerning March 10th Tibetan Uprising demonstrations. Coinme samples used in these attacks have a technical score of 1.25. These emails were sent between February 13 and March 7, 2012 and were all targeted at Tibet Group 1. The campaign used a combination of malicious XLS and RTF documents exploiting CVE-2010-3333. The majority of these attacks score a TTI of 3.75. One email only scores 2.0 on social engineering sophistication and a 2.5 overall TTI. We see the mothership server (209.11.241.1440) and 120.50.35.60 used as a mail sender. All of the attacks in this campaign used 61.178.77.98 (without an associated DNS name) as a C2.

CAMPAIGN 3

The second of three campaigns using the Conime malware family involved seven emails, three of which were distinct, sent to Tibet Groups 1, 2, and 4. These emails all scored a social engineering sophistication value of 3.0, for a combined TTI of 3.75. Two of the distinct emails had an attached XLS document; one was encrypted, the other was not. The third email used a malicious RTF document exploiting CVE-2010-3333. The encrypted XLS was sent on July 25, 2012, and the other emails were sent between September 10 and 12, 2012. We again see 209.11.241.1440 as an email sender IP. All of these exploits dropped the same variant of Conime, which connected to gen2012.eicp.net:1080 (61.178.77.98) as a C2.

CAMPAIGN 4

The third campaign using Conime was more varied than the other two, and was targeted at Tibet Groups 1, 2, and 4. Fifteen emails were received, eleven of which were distinct (although one showed only minor changes), ranging from 2.0 to 4.0 on the

social engineering sophistication score. These emails were sent over a longer timeframe than the other campaigns, extending between June 14, 2012 and September 12, 2012. Vulnerabilities used included both major RTF (CVE-2010-3333, CVE-2012-0158) and XLS (CVE-2009-3129) versions. One email, received by Tibet Group 2, received a social engineering sophistication score of 4.0. This email was highlighted to us by the recipient as highly targeted, and referenced an upcoming conference call about grant funding. Like the previous two NetTraveler campaigns, the malware connected directly to 61.178.77.98.

CAMPAIGN 5

This campaign used a variant of Gh0st RAT, with a flag text of “Snow.” Identical emails, concerning a visit of HHDL to Portland, were sent to Tibet Groups 2 and 4 on January 28, 2013. The emails have a social engineering score of 2, with an overall TTI score of 2.5. The attackers again used 209.11.241.144 as a mail sender and 61.178.77.98 as a C2.

CAMPAIGN 6

This campaign used a different malware family, RegSubDat, which was contained in an RTF using CVE-2012-0158, attached to an email sent to Tibet Group 1. Again we see mail sent from 209.11.241.144, but in this case the malware connected to a different C2

server: itsec.eicp.net:443 (1.203.31.195). This attack scored 3.0 on the social engineering sophistication value for an overall TTI of 3.75.

CAMPAIGN 7

The final sample from the NetTraveler group was observed in an email message sent to Tibet Group 1 on March 15, 2012. This malware was sent from the same mother-ship server (209.11.241.144) described above, but rather than attaching the malicious file, as had been done for all prior attacks, this email contained a link to an infected XLS file. The file was hosted at www.eaglessey.com (120.50.35.46), but was no longer present when we attempted to access it.

OBSERVATIONS

The NetTraveler campaign serves as another example of a campaign that targets CSOs alongside industry and government targets. These campaigns are conducted by a prolific threat actor that has targeted a variety of different sectors. Our findings confirm the targeting of Tibetan groups identified by Kaspersky, as all five of our Tibet Groups were targeted. This campaign demonstrates an adaptive attacker that uses a variety of vulnerabilities for different applications, including targeting of both Mac and Windows platforms.

PlugX Campaigns

First Seen	February 10, 2011
Last Seen	January 15, 2013
Attack Vectors	Targeted malicious emails
Exploits	CVE-2012-0158 (RTF, DOC, and XLS), CVE-2012-1889 (Internet Explorer), CVE-2012-5054 (Flash), CVE-2009-4324 (PDF), CVE-2007-5659 (PDF)
Malware Families	PlugX, Poison Ivy
Infrastructure	C2 domains: sociapub.flower-show.org:8080 (14.102.252.142), 114.142.147.51:8080, system.windowsdeupdate.com:8080 (174.139.12.84, 98.126.14.13), web.windowsdeupdate.com:7070 UDP (74.139.12.84, 98.126.14.13), , new.edamobile.com:443 (58.64.200.114), jinyuan2011.zapto.org:443 (123.129.19.145)
Targeted Groups	Tibet Group 1, Tibet Group 2, China Group 1, China Group 2
TTI Range	1.5 - 7.5

BACKGROUND

PlugX is a well-known family of malware that researchers have observed being used in targeted attacks against Tibetan organizations, NGOs, government institutions, and private companies.

Trend Micro has [published a report](#) on PlugX, describing a long-standing campaign that previously used Poison Ivy, another malware family. Jaime Blasco at Alien Vault [claims to have tracked down the author](#) of PlugX, who is allegedly based at a Chinese security company.

The PlugX samples seen in our study can be clustered into four campaigns, based on email sender IP and C2 infrastructure. Examining email topics, vulnerabilities used, and compile paths (as described in the Alien Vault blog post) suggests that the four

campaigns are from the same source. We have also seen a Poison Ivy sample used in this campaign.

The attack vectors and vulnerabilities used in PlugX are more varied than other attacks in our dataset. The vulnerabilities used include instances of CVE-2012-0158 in three separate file formats, an Internet Explorer vulnerability (CVE-2012-1889) that will install PlugX as a drive-by download, and a Flash vulnerability (CVE-2012-5054) hosted on an external website. The earlier Poison Ivy attack used two older PDF vulnerabilities. The Flash vulnerability was particularly notable; it was a zero-day at the time of the attack, leaving any user who clicked the malicious link it was hosted on vulnerable to compromise.

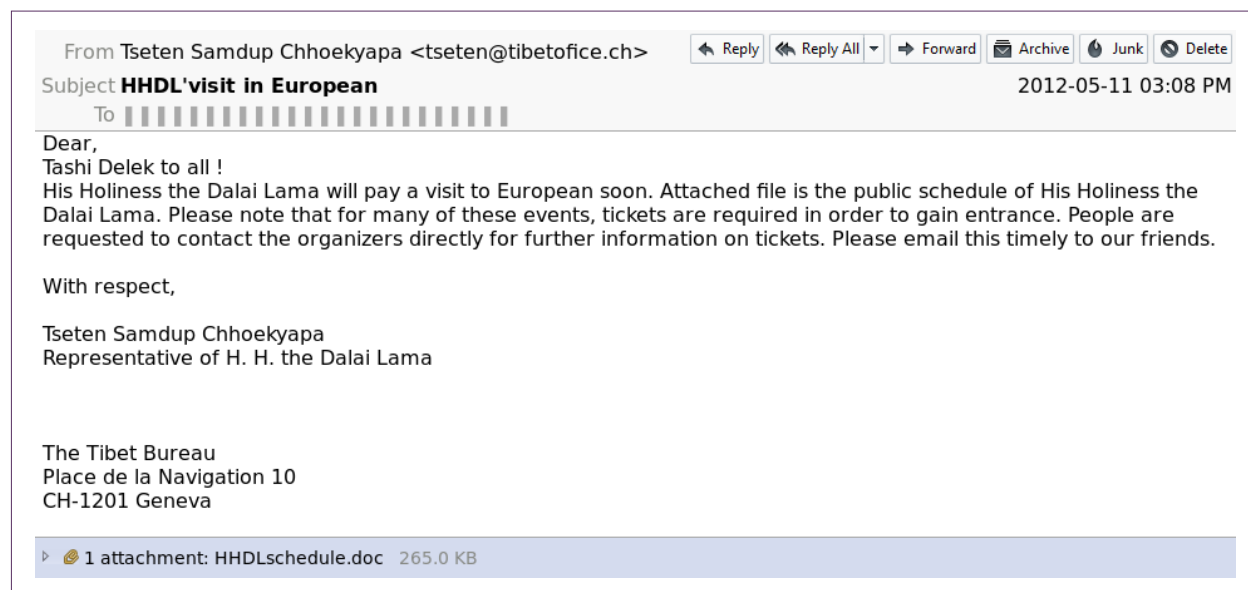
CAMPAIGN 1

The first set of attacks consists of fifteen emails, five of which were unique, sent from May 11, 2012 to June 1, 2012. Tibet Group 1 and Tibet Group 2 were both targeted with at least four out of the five emails. These emails show many signs of coming from the same source, including a common return address of `kandid77@rambler.ru`, a sender IP of 98.126.14.13, and common C2 infrastructure.

Two different C2 domain names were used: `system.windowsdeupdate.com` (TCP port 8080) and `web.windowsdeupdate.com` (UDP port 7070). These DNS names both pointed to the same IPs, which include 174.139.12.84 and 98.126.14.13.

All of these emails have a social engineering score of 3.0 and an overall TTI of 4.5. One example, sent to Tibet Group 2, spoofed a [legitimate Tibetan official](#) and contains a Word document that outlines the schedule of an [actual European tour](#) taken by the Dalai Lama.

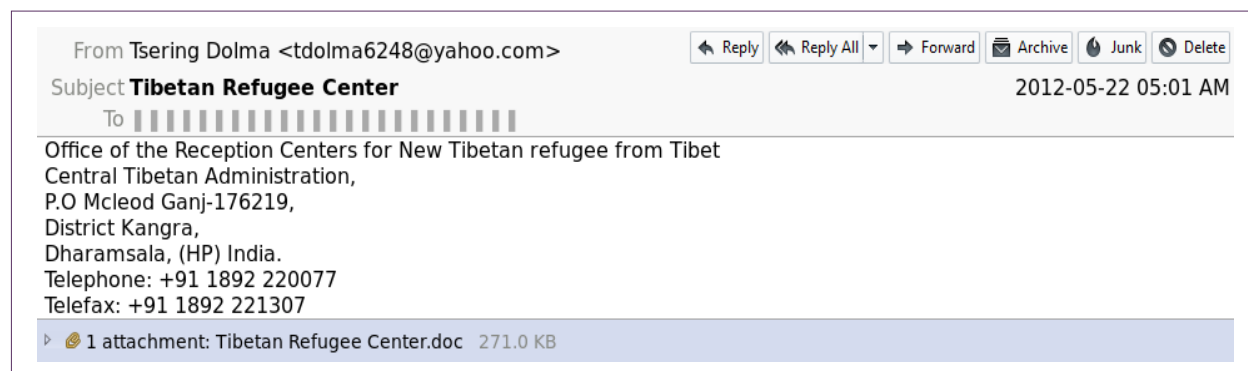
EXTENDED ANALYSIS: 2.2 Cluster Analysis



Social engineering	3
Technical	1.5
TTI	4.5
MD5	1aa5dde570575d0b001a48e62a412f14
C2	system.windowsdeupdate.com (174.139.12.84)

CAMPAIGN 2

On May 22, 2012, an email was sent from the IP address 69.46.75.74 to Tibet Group 2, which claimed to be from an individual named Tsering Dolma, with an email signature belonging to the Central Tibetan Administration, and with the return address of 'tdolma6248@yahoo.com.' This email contained an attached RTF with CVE-2012-0158 that was used to install PlugX.



Social engineering	2
Technical	1.5
TTI	3.0
MD5	74de1701d68d7e0a9f86bb6297246ebd
C2	new.edamobile.com (58.64.200.114)

CAMPAIGN 3

Three emails were sent to Tibet Group 2 and China Group 1 between June 15 and August 30, 2012. Each email had unique content, attack vectors, sender email address and IP, and vulnerability used. The vulnerabilities included a Word variant of CVE-2012-0158, the Flash vulnerability CVE-2012-5054, and Internet Explorer vulnerability CVE-2012-1889.



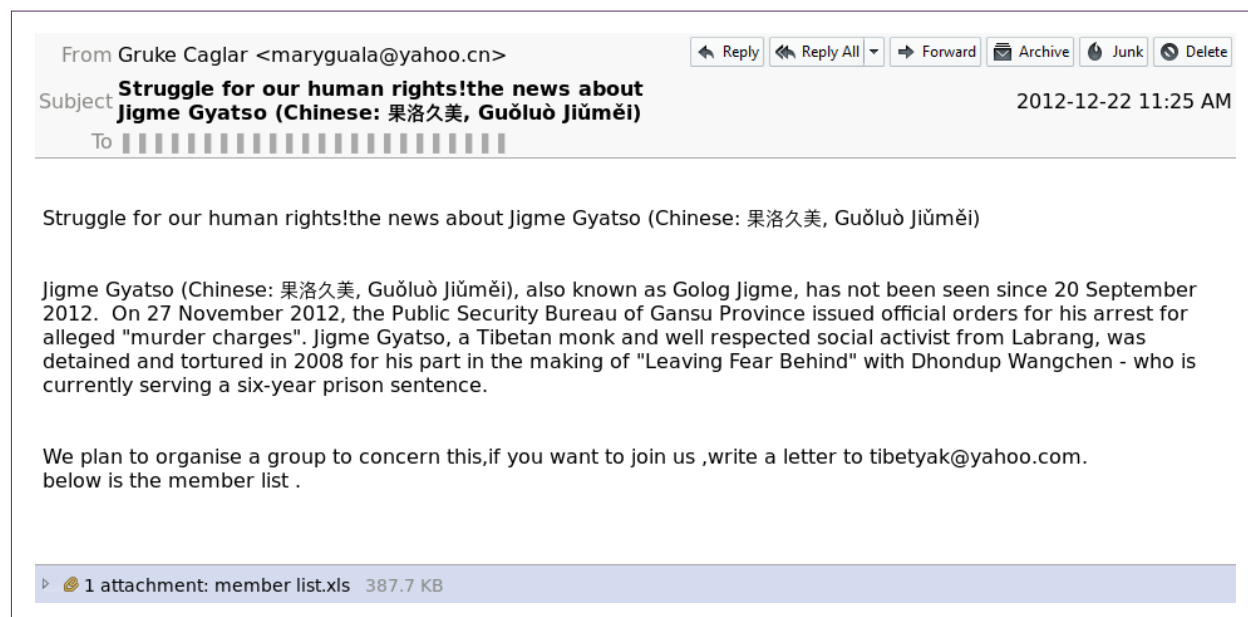
Social engineering	3
Technical	1.5
TTI	4.5
MD5	81f3a6e7a73a9845c6eb9a3d46597223
C2	114.142.147.51:8080

The Flash vulnerability CVE-2012-5054 was a zero-day at the time it was used in an attack against China Group 1. The attack was delivered in an email that was highly customized for the recipient and used a malicious link in the message as the vector. It referred to a group of individuals who had recently been involved in internal private meetings and appeared to be a forwarded message from the director of the organization. The highly targeted nature of this attack, combined with the technical sophistication of the PlugX malware family, resulted in a TTI score of 7.5, the highest seen in the study.

CAMPAIGN 4

The last campaign consisted of four unique emails sent to Tibet Groups 1 and 2 between December 22, 2012 and January 15, 2013. These emails all included attachments that used CVE-2012-0158. The C2 domain used was jinyuan2011.zapto.org:443, which resolved to 123.129.19.145 at the time of the attack. Three of these four emails scored 2.0 on the social engineering sophistication score (and 3.0 TTI overall), and one scored 3.0 on social engineering for an overall TTI of 4.5.

One of these four emails repurposed legitimate text about a Tibetan monk who had been detained:



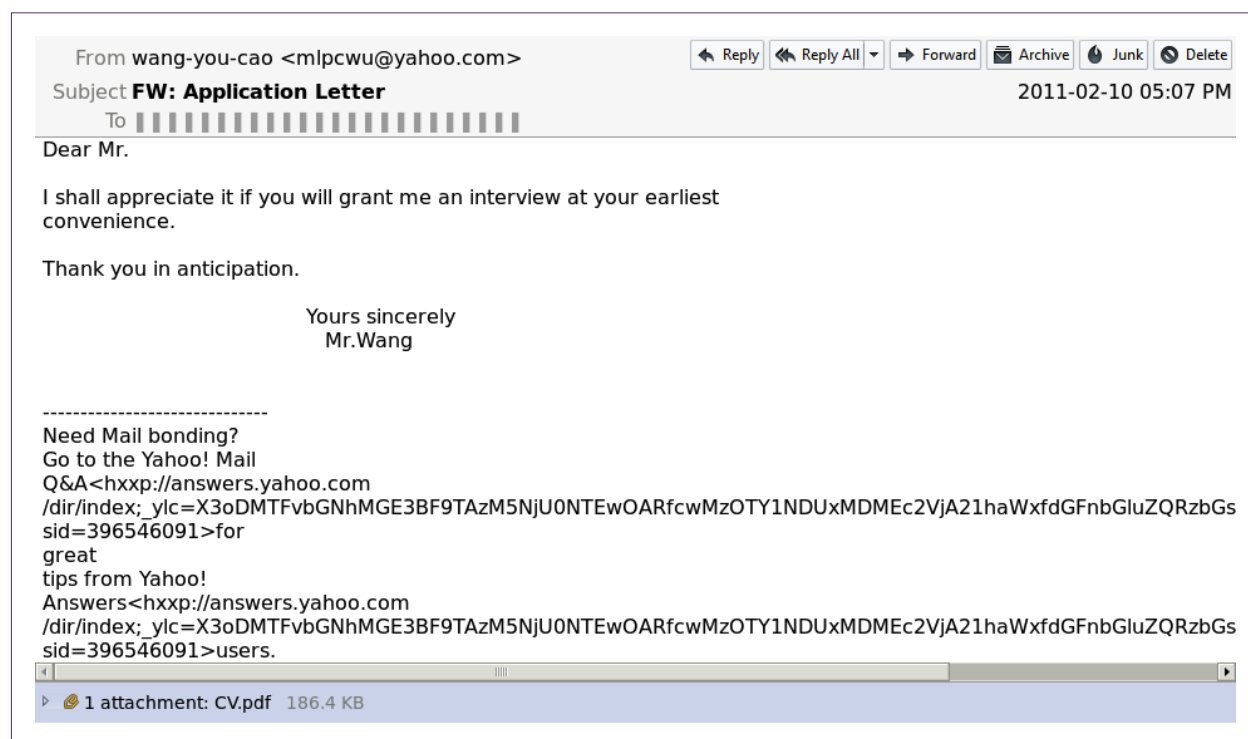
Social engineering	2
Technical	1.5
TTI	3.0
MD5	61af83d594c0e69b201d1ad20d3fd28d
C2	jinyuan2011.zapto.org (123.129.19.145)

This campaign was interesting in that it did not use the encrypted BOOT.LDR files, instead using NvSmartMax.dll.url, and logging keyboard data to NvSmart.hlp. This functionally corresponds to [observations](#) made by Kaspersky researchers that PlugX was becoming more mature. In particular, we observed that the identifying strings and logging data were removed in this campaign. It is particularly interesting that while the malware itself is being improved, potentially in response to published reports from threat researchers, the quality of the targeting in this campaign has gone down.

POISON IVY

In September 2012, Trend Micro [described](#) the use of PlugX in a campaign that had previously used the Poison Ivy RAT and targeted government and private companies in Japan. We also saw evidence in our study of Poison Ivy being used in conjunction with PlugX in an attack sent to China Group 2 on February 10, 2011, over a year

before our first observed PlugX attack. This email included a PDF with two vulnerabilities, CVE-2009-4324 and CVE-2007-5659.



Social engineering	1
Technical	1.5
TTI	1.5
MD5	49c9cf000fa1d789f3df8d739f997eb8
C2	sociapub.flower-show.org (14.102.252.142)

The Poison Ivy RAT connects to a C2 at sociapub.flower-show.org:8080 (14.102.252.142), the same Poison Ivy C2 domain [observed by Trend Micro](#) on July 11, 2012. This attack has also been seen elsewhere in the wild, as noted in a [Threat Expert report](#) describing the same malware seen with a different file size and MD5 hash (9ADFC6DD86D5FF36F2CAB781663E1075).

OBSERVATIONS

The PlugX campaign provides yet another example of a campaign that targets civil society organizations alongside government and industry groups, using the same infrastructure and malware to compromise targets. Aside from these similarities, the campaign otherwise had a number of unique characteristics that separated it from others in our research. Most notably, it was the only instance of a zero-day vulnerability seen in our study. Given that zero-days are highly effective, as software developers have yet to patch the vulnerability, they are highly lucrative and sought after. It is notable that the malicious attacks would use this zero-day to target a CSO, as once such an exploit is exposed it runs the risk of being identified and having the vulnerability fixed. The PlugX campaign also included a broader variety of attack vectors than what was seen in most campaigns. The attached files included the zero-day Flash vulnerability, an exploit for Internet Explorer, as well as the standard Microsoft Office exploits seen elsewhere.

TseringKanyaq Campaigns

First Seen	May 4, 2012
Last Seen	July 26, 2013
Attack Vectors	Targeted malicious emails
Exploits	Windows: CVE-2012-0158, Mac: CVE-2009-0563, CVE-2012-0507, CVE-2013-1331
Malware Families	Windows: Shadownet, Duojeen Mac: PubSab
Infrastructure	Email sender: 163.com, myopera.com, gmx.com C2 domains: newwolfs29.zxq.net, newwolfs21.blog.163.com, dplcoopsociety.us.dwyu.com, laraider2.he1.ifreeurl.com, pomehra.typepad.com, tbt society.info, nedfortibt.info, duojeen.info, appleboy1111.blogspot.com, coremail.info
Targeted Groups	Tibet Group 1, Tibet Group 2, Tibet Group 4
TTI Range	3.0 - 3.75

BACKGROUND

Unlike the previous campaigns that were grouped by shared infrastructure or connections to previously reported threat actors, the “TseringKanyaq” cluster was first identified through contextual analysis.

This cluster consists of a series of attacks targeting Tibet Groups, which had either ‘tseringKanyaq@yahoo.com’ or ‘d.kanam@yahoo.com’ in the reply-to address field of the malicious emails. Following the identification of this pattern, further attacks that shared common infrastructure were linked.

The addresses ‘d.kanam@yahoo.com’ and ‘tseringKanyaq@yahoo.com’ do not match known email addresses or names of persons in the Tibetan community. However, “tseringKanyaq” may be a misspelling of [Kanyag Tsering](#), a Tibetan monk from the Kirti Monastery in the Ngaba region of Tibet. This region has been the scene

of a number of Tibetan self-immolations, and Kanyag Tsering has provided reports of the incidents to international media. He is a well known and respected member of the Tibetan community who, due to his work in getting information from inside Tibet to journalists, has developed a significant media presence. We met with Kanyag Tsering and showed him our analysis of this cluster. He confirmed that the address 'tseringKanyaq@yahoo.com' does not belong to him. Despite the possible intentional similarity of 'tseringKanyaq@yahoo.com' to the name of a notable Tibetan monk, the purpose behind the consistent use of these addresses in the reply-to field is unknown.

We identified three malware families in this cluster, which were used to target Windows (ShadowNet and Duojeen) and OS X (PubSab). The ShadowNet malware family is associated with the [ShadowNet espionage](#) group, which was discovered by the Information Warfare Monitor and the ShadowServer Foundation in 2009 and was revealed to be targeting Tibetan organizations and Indian military and government institutions. The malware we found did not connect to infrastructure related to the previous ShadowNet campaign.

All three malware families used in the TseringKanyaq campaign were also used by the [LuckyCat campaign](#), which was discovered by Trend Micro in 2012. LuckyCat is notable for targeting companies based in India and Japan working in aerospace, energy, engineering, shipping, and military research in addition to Tibetan activists. We find infrastructure connections between the TseringKanyaq and LuckyCat campaigns, which suggests some level of coordination.

EMAIL PROVIDER INFRASTRUCTURE

This campaign is marked by a period of gaps in which attacks stop and later re-emerge with similar identifying features (most notably the consistent use of the Reply-To addresses), some level of improvements to the malware C2 infrastructure, and/or changes to the social engineering tactics.

We divide these gaps into three distinct periods in which the attackers utilize different email providers to send attacks.

- 163.com: May 4, 2012 - July 9, 2012
- myopera.com: July 24, 2012 - September 5, 2012
- gmx.com: October 14, 2012 - July 26, 2013

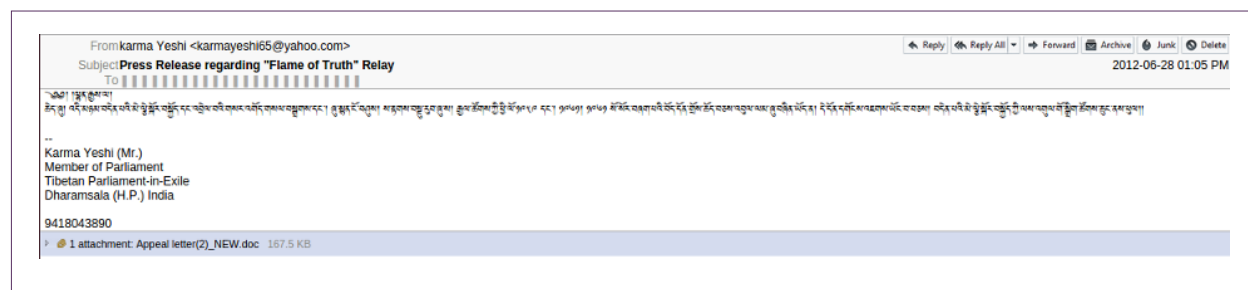
Between these three periods there is no overlap in the use of these mail providers. We clearly observe the attackers moving from one provider to the next. The use of these different providers could be due to a number of possibilities: the domain provider may have shut down the accounts due to notifications or detection of abuse, changes to provider infrastructure may have created difficulties in maintaining the accounts, or the attackers may have moved on to the next domain infrastructure once the advantages of an alternate provider became apparent.

163.COM CAMPAIGNS

The first wave of attacks occurred between May 4 and July 9, 2012. Four email attacks were sent during this period targeting Tibet Groups 2 and 4. All of these attacks spoof prominent organizations in the Tibetan community and repurpose legitimate content, which gives them a social engineering sophistication base value of 3. All of the malware samples in this cluster have a technical score of 1.25, for a total TTI of 3.75.

The first attack on May 4, 2012, was sent to Tibet Group 2 with an email that repurposed content concerning a petition campaign. The actual email sender was psjiangzuo@163.com (174.139.21.26). The attached file dropped Duojeen malware that connects to www.xiuxiu.in (173.231.22.201).

The second attack on June 28, 2012, sent to Tibet Group 2, purported to come from Karma Yeshe, a member of the Tibetan Parliament in Exile (TPiE), and provided information (in Tibetan) on the Flame of Truth Rally, a campaign launched by the TPiE to express solidarity with Tibetans who have self-immolated. The actual sender of the mail is sysutyubu@163.com (222.212.213.197). The attachment also drops Duojeen malware.

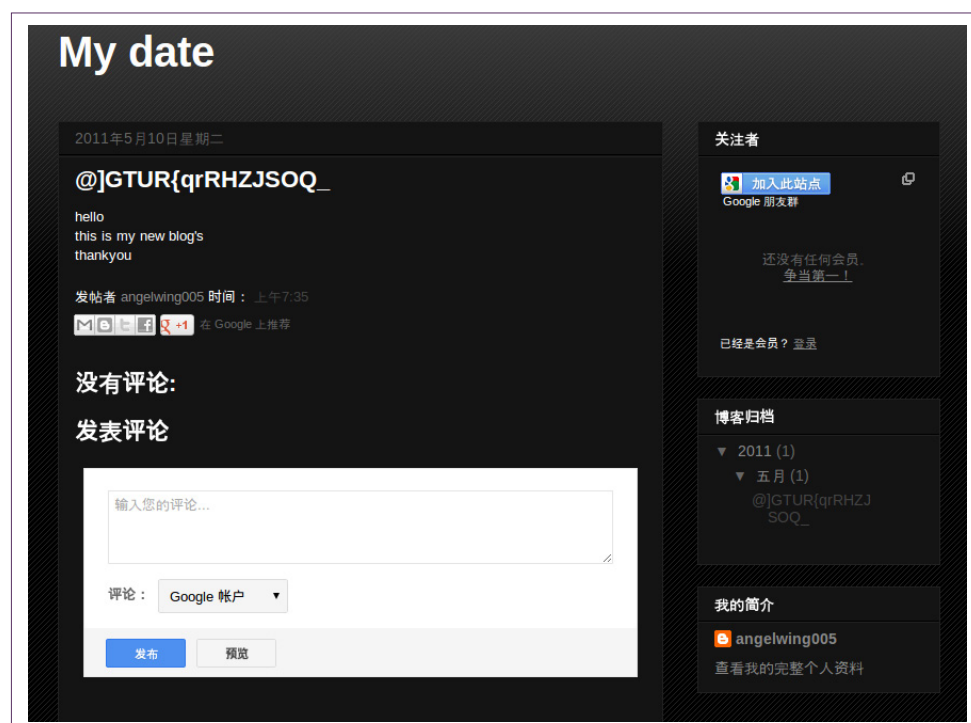


Social engineering	3
Technical	1.25
TTI	3.75
MD5	e9b9c09002197882ed1140054d20623a
C2	dplcoopociety.us.dwyu.com (184.82.238.34)

On July 5, 2012, Tibet Groups 2 and 4 both received identical emails with content related to a recent self-immolation. The malware used was ShadowNet, which leverages [Windows Management Instrumentation](#) (WMI), a system tool meant for administrators. Its intended usage as a tool for collecting system information and automation makes it an ideal mechanism for gathering and exfiltrating data. Use of legitimate Windows features can make it more difficult for administrators to identify activity as malicious.

The ShadowNet attacks used a WMI Script that contained links to one of three blogs to which the malware attempts to connect. The blog then has a string with encoded C2 information as shown in Figure 14 below.

FIGURE 14: Sample of blog post used to transmit C2 information to infected machines

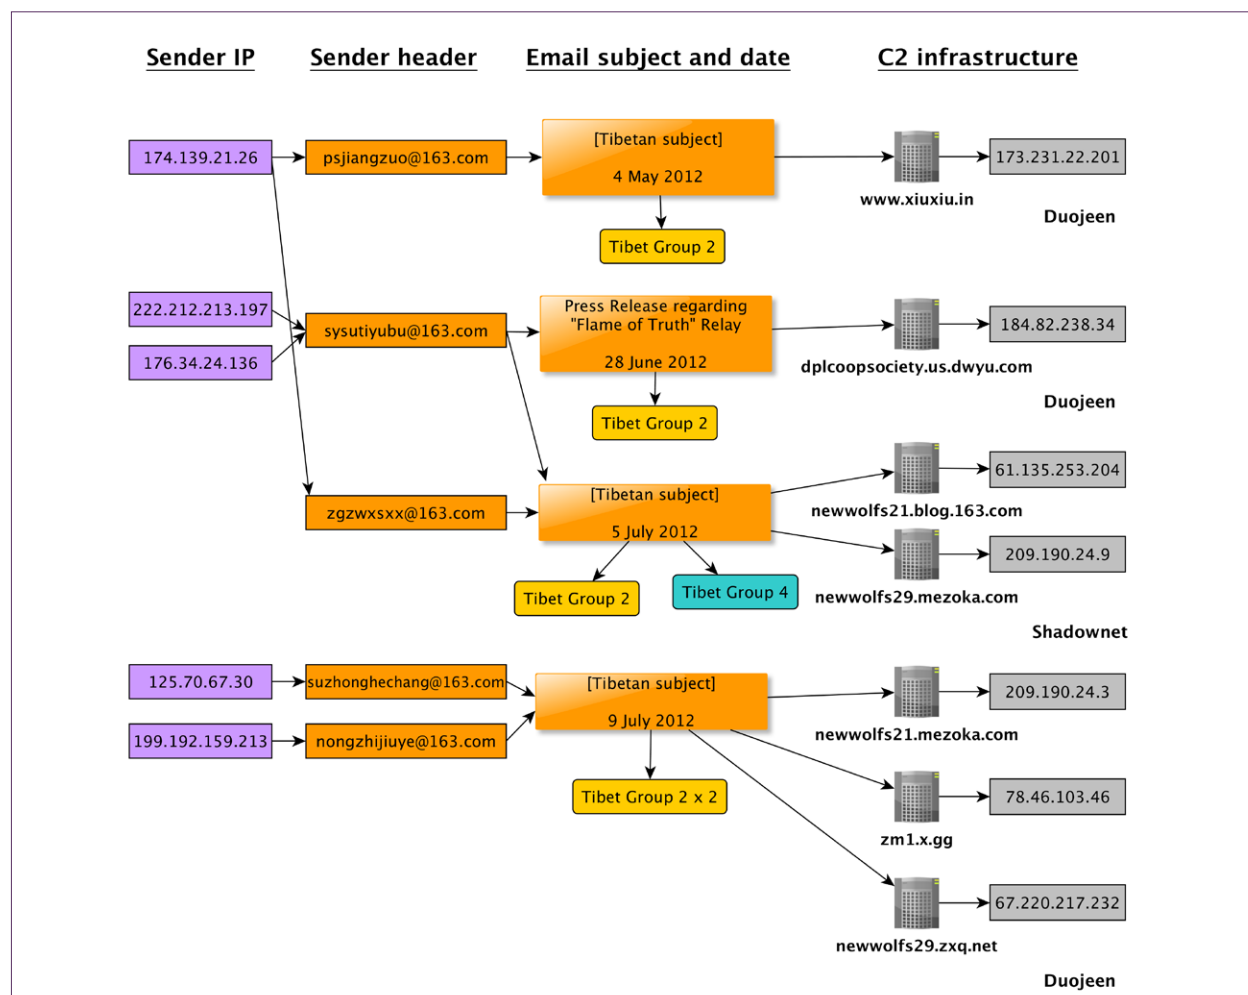


EXTENDED ANALYSIS: 2.2 Cluster Analysis

Once a connection to the C2 is made, system information and data can be sent to the attackers. In the case of the July 5 attacks, the malware first connects to newwolfs21.blog.163.com, after which it retrieves the C2 newwolfs29.mezoka.com.

On July 9, 2012, Tibet Group 2 received two identical emails sent to two separate organizational accounts. The emails contained content regarding self-immolations. The malware used, Duojeen, retrieves system information, establishes a connection with zml.x.gg, and sends the collected information. It then retrieves second stage malware from <http://newwolfs29.zxq.net/winxp.rar>. As with the previous attack, the two emails were sent from different accounts: suzhonghechang@163.com (125.70.67.30) and nongzhijiuye@163.com (199.192.159.213).

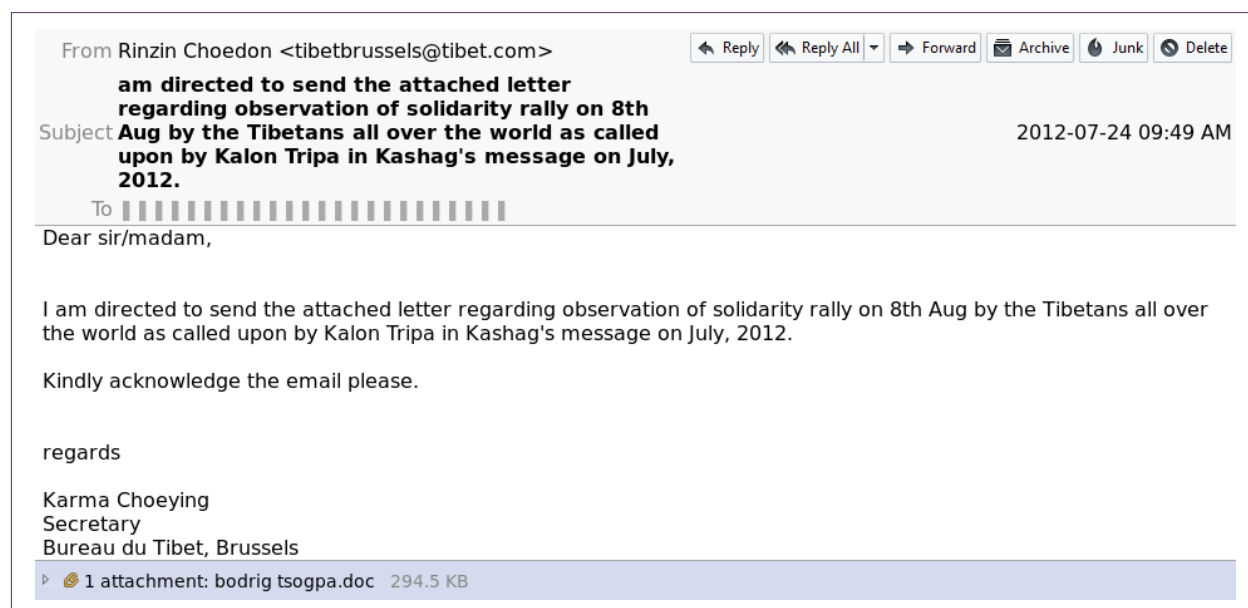
FIGURE 15: Email sender, IP, and C2 infrastructure for 163.com tseringKanyaq emails



MYOPERA.COM CAMPAIGNS

From July 24 to September 5, 2012, the attackers moved their mail provider to myopera.com. The attacks continued to use ShadowNet and Duojeen malware and the same common C2 infrastructure as the previous campaign. During this period, we observed three attacks sent to Tibet Groups 1, 2, and 4, in some cases to multiple accounts within the organizations. All of these attacks spoofed prominent groups in the Tibetan community and repurposed legitimate content, which gives them a social engineering score of 3. All of the malware samples in this cluster have a technical score of 1.25, for a total TTI of 3.75. Interestingly, in this campaign we observed some emails being sent from Tor exit nodes, which shows the attackers making a new effort to hide their location.

Between July 24 and 25, Tibet Groups 1, 2, and 4 received identical emails that appeared to be from the Tibet Office in Brussels, with content regarding an upcoming rally.



Social engineering	3
Technical	1.25
TTI	3.75
MD5	cc0b8b8e42fdd59cc4b32b3a06e57281
C2	newwolfs29.mezoka.com (209.190.24.9)

The attached malware was ShadowNet, and it connected to newwolfs21.blog.163.com to retrieve C2 instructions. In each instance, the email was sent from a different account (newwolfs41@myopera.com, tenzin600@myopera.com, tibettibetan3@myopera.com, tenzin600@myopera.com, and mytenzin@myopera.com); however, the sender IP was the same for each account (184.82.49.114). We later discovered that this sender IP was also used as a C2.

During an investigation of one of the C2s used in the campaign (newwolfs20.x.gg/mits/) we found open directories that included a sample from the Sparksrv malware family, which we presume was intended to act as stage two malware. This sample used 184.82.49.114 as a C2.

Between July 25 and 26 Tibet Groups 1, 2, and 4 received identical emails that spoofed Karma Yeshe, a member of the TPiE, with an update on the Flame of Truth rally campaign. The malware in this case was Duojeen.

[illegible]

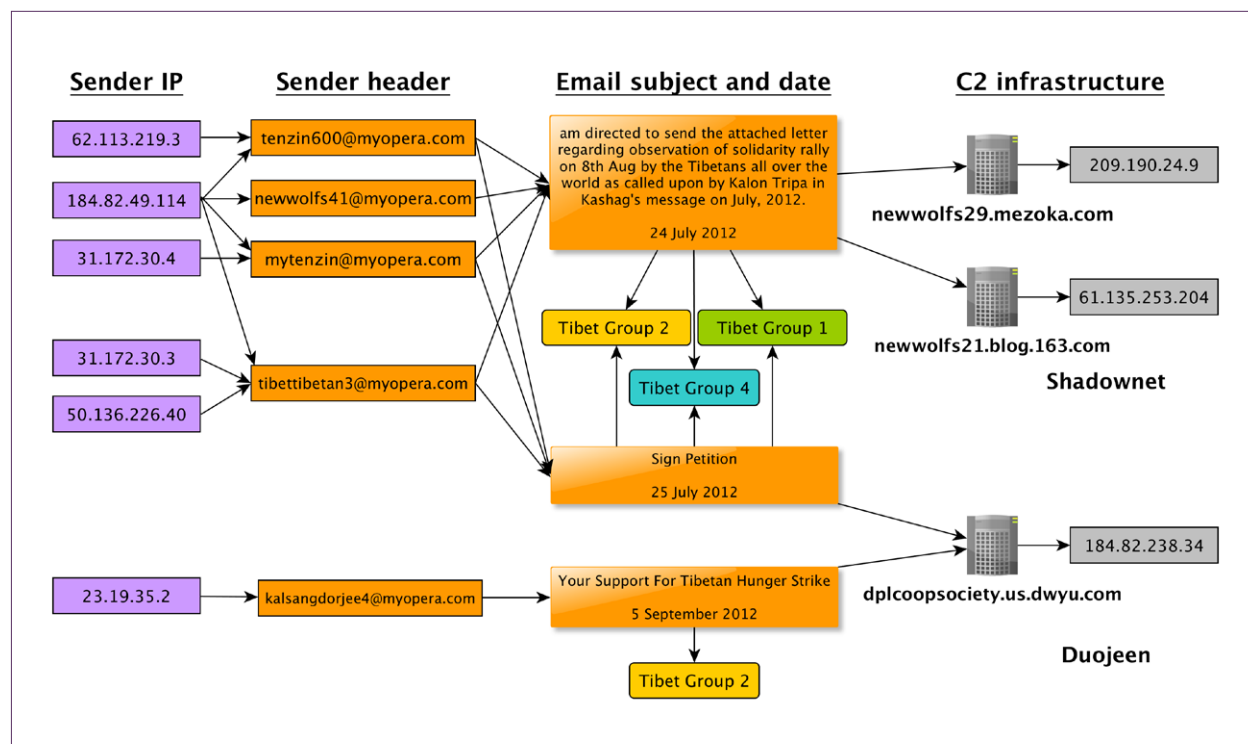
Social engineering	3
Technical	1.25
TTI	3.75
MD5	f208137dfb3271b5cd3c67492e2522dc
C2	dplcoopsociety.us.dwyu.com (184.82.238.34)

In three of the attacks, the sender IP traced back to a Tor exit node maintained by the Chaos Computer Club in Germany. The timing of this shift in tactics is interesting, as it comes after a series of attacks that used a C2 as the mail sender. Using Tor to mask the real location of the mail sender may therefore have been an effort to improve operational security. However, the attackers use of Tor is inconsistent and following this series of attacks was observed only one other time.

TABLE 11: IP address of email senders used during MyOpera campaign

SENDER	IP	LOCATION
mytenzin@myopera.com	31.172.30.4	Germany - Tor exit node
tenzin600@myopera.com	62.113.219.3	Germany - Tor exit node
tibettibetan3@myopera.com	31.172.30.3	Germany - Tor exit node
tibettibetan3@myopera.com	50.136.226.40	US - Comcast

The final attack in this campaign was sent to Tibet Group 2 on September 5, 2012. It contains information regarding a hunger strike initiative undertaken by the Tibetan Youth Congress. The attachment contained Duojeen malware that connected to dplcoopsociety.us.dwyu.com (184.82.238.34).

FIGURE 16: Email sender, IP, and C2 infrastructure for myopera.com tseringKanyaq emails

GMX.COM CAMPAIGNS

From October 14, 2012 to July 26, 2013, the attackers switched to gmx.com as their mail provider. The use of gmx.com is interesting, because the email headers from this provider include a unique user ID number, which can be used to track malicious accounts using this service. However, we see no overlap in user IDs from the gmx.com accounts used in the attacks, which suggests that the attackers used a script to generate new accounts, and as a result the user ID number is always different.

Within this period we observe the attackers using reply-to 'tseringKanyaq@yahoo.com' and also beginning to use reply-to 'd.kanam@yahoo.com' in messages. We cluster attacks by these two accounts in the subsections below.

gmx.com 'tseringKanyaq' campaigns

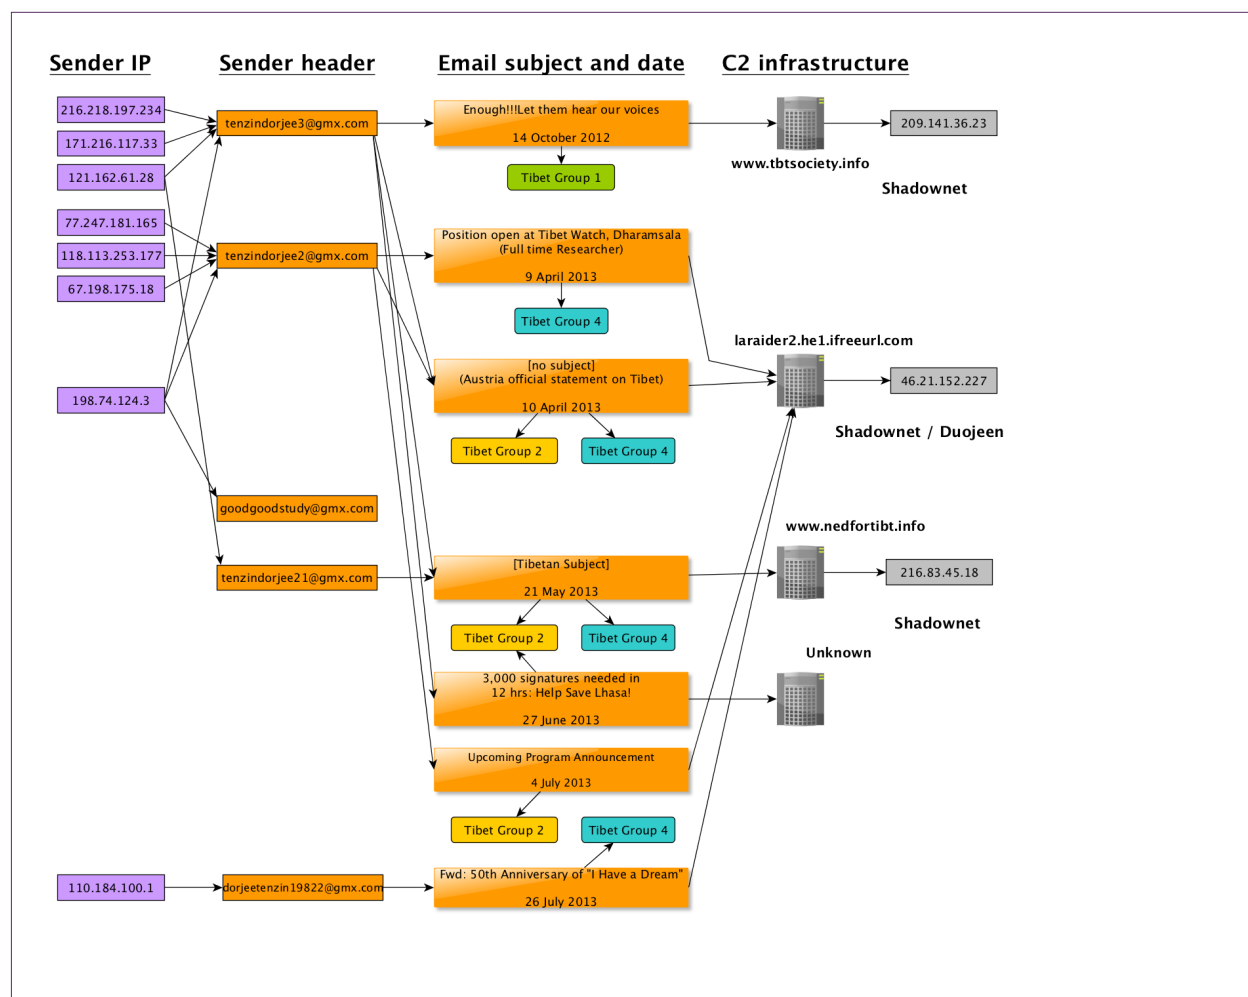
From October 14, 2012 to July 26, 2013, we observed 11 attacks targeting Tibet Groups 1, 2, and 4, using the gmx.com mail provider and a reply-to address of 'tseringKanyaq@yahoo.com.'

As with the previous campaigns, these messages all spoofed real people and organizations in the Tibetan community. Each email in this campaign has a social engineering sophistication score of 3 and a technical score of 1.25 for a total TTI of 3.75.

Five attacks in this campaign used ShadowNet connecting to www.tbtsociety.info (209.141.36.23), laraider2.he1.ifreeurl.com (46.21.152.227), or www.nedfortibt.info (216.83.45.18). Three attacks dropped Duojeen connecting to laraider2.he1.ifreeurl.com (46.21.152.227).

In one instance we see the connection of an email sender traced back to a Tor exit node (IP 77.247.181.165).

FIGURE 17: Email sender, IP, and C2 infrastructure for gmx.com tseringKanyaq emails



gmx.com 'd.kanam' campaigns

From December 18, 2012 to May 2, 2013, we observed 12 attacks targeting Tibet Groups 2 and 4 using the gmx mail provider and a reply-to address of 'd.kanam@yahoo.com.'

Message lures in this campaign all spoofed real people and / or organizations and repurposed legitimate content from Tibetan groups. Each email in this campaign has a social engineering sophistication score of 3.

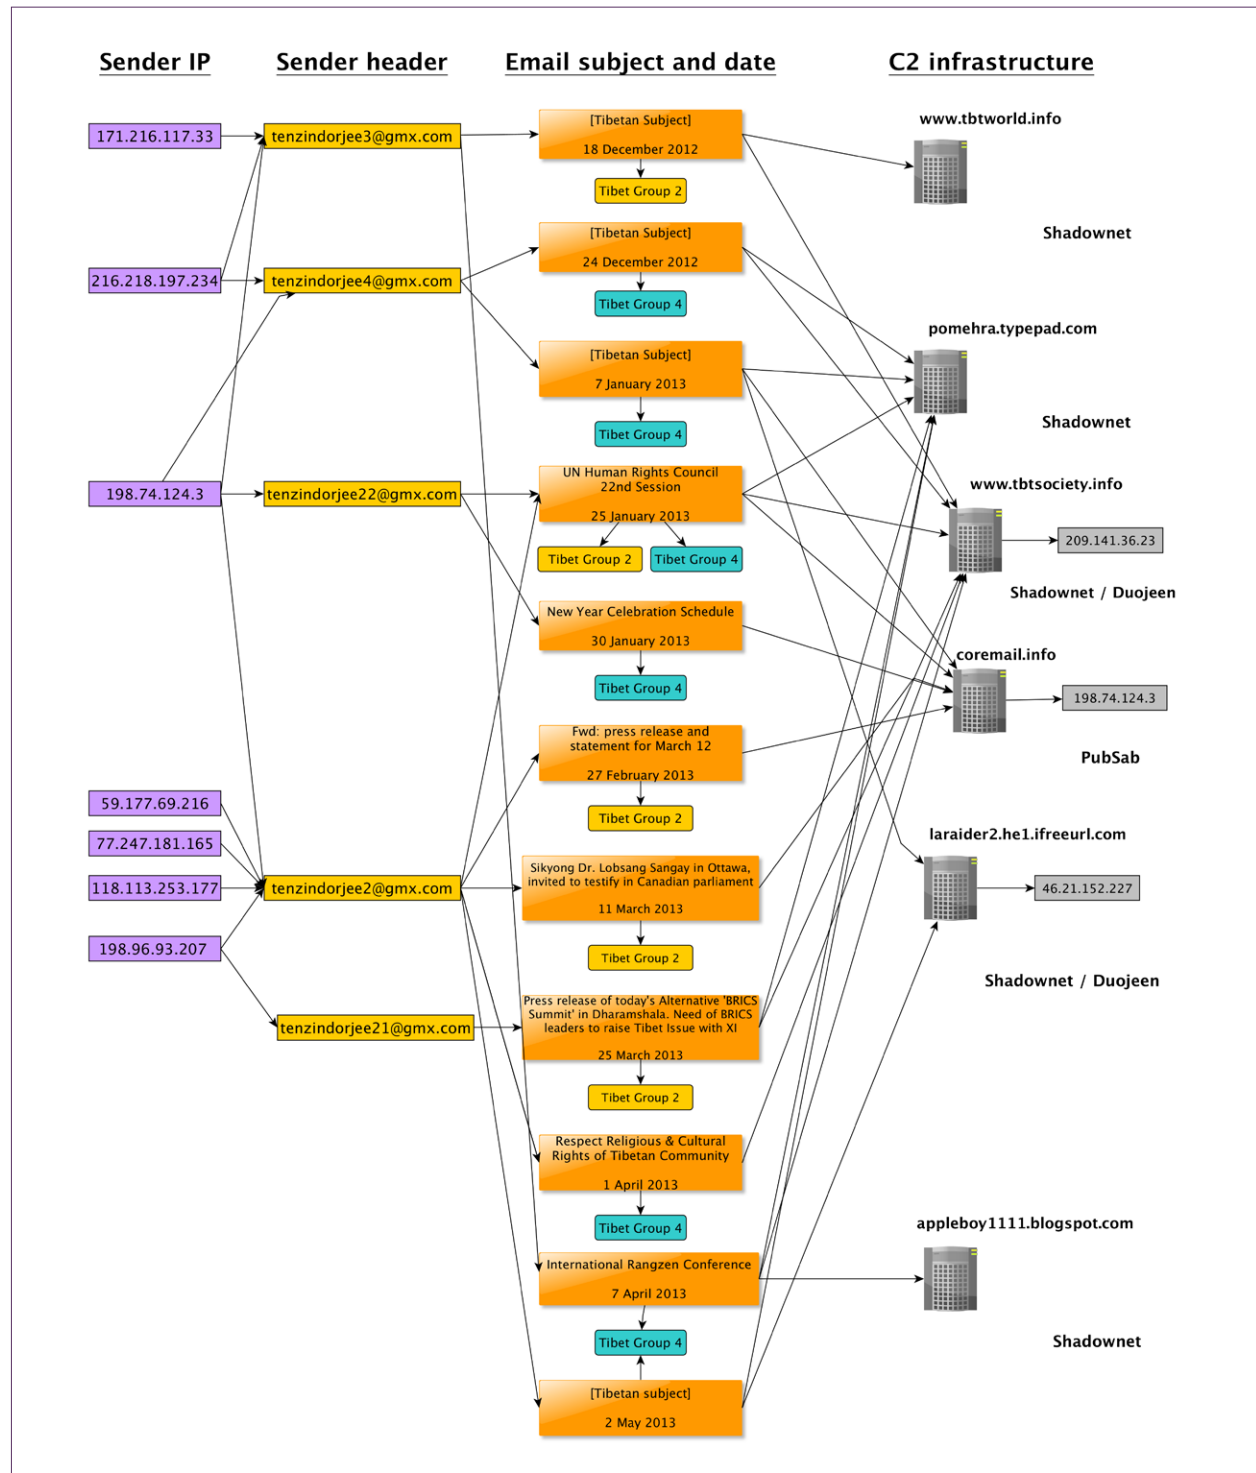
Seven attacks in this campaign used ShadowNet connecting to www.tbtsociety.info, laraid2.he1.ifreeurl.com

(46.21.152.227), and pomehra.typepad.com (204.9.177.195). Each of these attacks has a technical score of 1.25.

Two attacks used Duojeen connecting to www.tbtsociety.info (216.83.45.18). Each of these attacks has a technical score of 1.25.

Within this wave of attacks we also observed the use of PubSab implanted in Word documents that used the exploit CVE-2009-0563. In this instance the malware connected to coremail.info (198.74.124.3).

FIGURE 18: Email sender, IP, and C2 infrastructure for gmx.com d.kanam emails



CONNECTIONS TO LUCKYCAT CAMPAIGN

We identify a number of connections between the TseringKanyaq cluster and the [LuckyCat campaign](#), as well as Sparksv and Duojeen campaigns related to LuckyCat.

The LuckyCat campaign used a variety of malware including Duojeen, ShadowNet and PubSab, which we also see used in the TseringKanyaq campaigns. Beyond this common set of malware we also see connections to infrastructure linked to LuckyCat and related campaigns.

The LuckyCat campaign utilized a series of free hosting and VPS services for its C2 infrastructure. One of the VPS services is hosted on duojee.info. On July 6, 2011, Tibet Group 1 received a malicious email containing ShadowNet malware. While this sample does not include the reply-to address tseringkanyaq.yahoo.com or d.kanam@yahoo.com, the C2 infrastructure has connections to the campaign. The malware connects to and retrieves C2 information from appleboy1111.blogspot.com. We observed a previous version of the script with encoded C2 information that points to duojee.info. The script was then updated to point to www.tbtsociety.info (216.83.45.18), which is a C2 we see used repeatedly in the gmx.com campaigns. The transition to this C2 shows evidence of the attackers shifting infrastructure that was previously linked to LuckyCat to new infrastructure we see used in the TseringKanyaq campaign.

We also see connections to Sparksrv campaigns that have been linked to LuckyCat. Sparksrv is malware used by the LuckyCat campaign as a second stage tool to add additional functionality after the first stage dropper successfully infects a target. Our analysis of open directories on a TseringKanyaq-related C2 revealed Sparksrv on the server, which suggest it was also being used as a second stage in this campaign.

Trend Micro identifies rukiyeangel.dyndns.pro as a C2 used for Sparksrv campaigns related to LuckyCat. In two attacks sent on December 24, 2012 and April 10, 2013, we see the email sender IPs originating from 198.74.124.3 and 216.218.197.234, respectively. The IP 198.74.124.3 currently resolves to coremail.info, which was used as a C2 for PubSab attacks in the TseringKanyaq campaign. [Passive DNS records](#) show that 198.74.124.3 and 216.218.197.234 previously resolved to rukiyeangel.dyndns.pro.

OBSERVATIONS

This campaign has several interesting characteristics relative to others in our study. The ShadowNet malware is the only example of WMI malware we observed. While this quality makes the malware relatively easy to remove, it also makes it more difficult for the user to identify. This campaign also relies on a highly disposable C2 infrastructure.

This cluster is also the only campaign to be first identified through contextual clues rather than a strict reliance on shared code or C2 infrastructure. The frequent

use of ‘tseringKanyaq@yahoo.com’ and “d.kanam@yahoo.com” in the Reply-to field is in some cases the only indicator tying the attacks together. Despite using a variety of different domains to send the malicious emails, it remains unclear why the same email address was reused so often.

Notably, this campaign also has links to other malware tools and campaigns related to ShadowNet and LuckyCat that have targeted a range of communities and sectors including Tibetans.

DTL Campaigns

First Seen	December 21, 2011
Last Seen	July 4, 2013
Attack Vectors	Targeted malicious emails
Exploits	Windows: CVE 2010-3333, CVE-2012-0158
Malware Families	Windows: 9002, 3102, Mongal, Nsfree, Boouset, Gh0st RAT, LURK0 (Gh0st RAT variant), CCTV0 (Gh0st RAT variant), Surtr (Remote and GtRemote), T5000
Infrastructure	C2 domains: dtl.eatuo.com, dtl.dnsd.me, dtl6.mo00.com, tbwm.wlyf.org
Targeted Groups	Tibet Group 1, Tibet Group 2, Tibet Group 3, Tibet Group 4
TTI Range	2.5 - 6.75

BACKGROUND

We identified a distinct campaign of targeted malware attacks against Tibetan groups that used the shared infrastructure of four C2 domains (dtl.eatuo.com, dtl.dnsd.me, dtl6.mo00.com, tbwm.wlyf.org). Tracking the IP address resolution of these domains over time, we observed that at certain periods they resolve to the same IP and therefore belong to a shared C2 infrastructure. We call this cluster the DTL campaign, because of the use of “dtl” in most of the C2 server domain names. We see this infrastructure used for a series of campaigns that involve 9 malware families: T5000, 9002, Boouset, Mongal, Nsfree, Gh0st RAT, LURK0 (Gh0st RAT variant), CCTV0 (Gh0st RAT variant), and Surtr. We also identified one other malware family (3102) that is likely related due to code overlap.

In November 2013, FireEye published a [report](#) that also identified the DTL cluster linking seven malware samples to four C2 domains, three of which we also observed (dtl.eatuo.com, dtl.dnsd.me, dtl6.mo00.com). FireEye only saw DTL campaigns using

the malware family 9002, which we observed being used for three attacks, of which one sample had a matching MD5 to one provided in the FireEye report (9f5e9e6b0c-87cad988f4a486e20bbc99). Our visibility into DTL campaigns only revealed Tibetan targets. However, one attack sent to a Tibet Group had an email body and attachment written in Uyghur. Other researchers have identified attacks related to the DTL campaigns targeting [Uyghur groups](#). The Uyghur samples sent to the Tibet Group may therefore be the operators accidentally sending the wrong lure.


Interestingly, FireEye observed campaigns using DTL infrastructure targeting a range of government and industry entities, showing their scope goes beyond CSOs. Such targets included entities within the following sectors (using FireEye's categories): U.S. federal government, state and local government, services/consulting/VAR, financial services, telecommunications, aerospace/defense/airlines, energy/utilities/petroleum refining, healthcare/pharmaceuticals, entertainment/media/hospitality, insurance, chemicals/manufacturing/mining, high-tech, and higher education.

MALWARE DEVELOPMENT PATTERNS

Of the nine malware families seen in DTL campaigns, the most frequently used were LURK0 and CCTV0, which are both variations of the Gh0st RAT codebase. LURK0 and CCTV0 are named for the five-character header that appears in network traffic when the malware is run. Both pieces of malware have standard RAT functionality including keylogging, file listing, and data exfiltration. Our observations of the DTL campaign show active development of these RATs over the period of two years that are unique to this cluster.

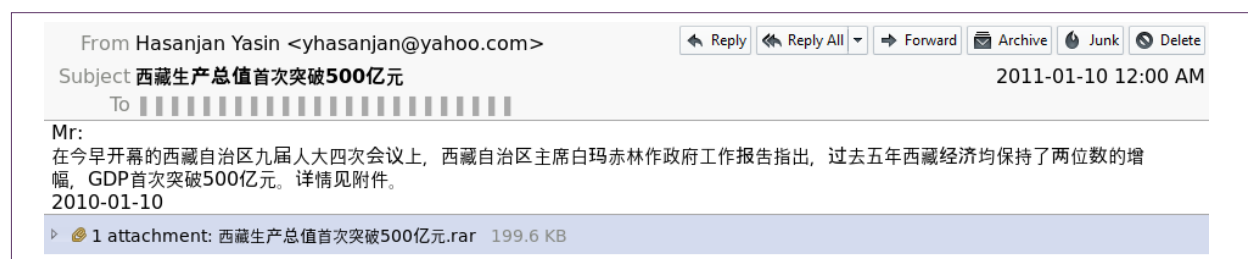
We found relations between malware samples using binary comparison tools to attempt to determine shared code bases, and comparing various identifiers in the samples. For example, LURK0 creates registry keys with names that are a variation on "DbxUpdate" and then uses a mutex to see if it is already running on the infected system. These names can be customized and used to attempt to distinguish between campaigns using the same malware family. Another useful feature for analysis is compilation times. Although these times can be easily modified, if related samples all have the same compilation date, they were likely created with the same builder. Through analysis of these features and tracking of shared C2 infrastructure we divide DTL-related attacks into a series of eight campaigns. We also discuss two related campaigns that, while not using the same infrastructure, use malware that shares code and identifying features and are likely developed by the same group.

FIGURE 19: Timeline of DTL-related malware families and the domains/IPs and ports to which they connect (IP addresses to which the domains resolved, where available, are represented by background colour)




CAMPAIGN 1: T5000

In 2011, four emails were sent to Tibet Group 1 with T5000 malware attached. On January 10, the group received an email with a .rar archive attachment containing an executable file. The email and attachment were in Chinese.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	2cf577eda241158e3c3b5431f30b9aeb

On July 11, Tibet Group 1 received an email in Tibetan, with a Microsoft Help (.hlp) file attached. The T5000 malware embedded in the file successfully connected to deepinlife.dyndns.info.



Social engineering	3
Technical	1.25
TTI	3.75
MD5	604d501e9e0ce7c175060b8512f706b7
C2	deepinlife.dyndns.info

On July 26, Tibet Group 1 received the fourth of the emails, with another help file attached. This sample connected directly to 180.178.53.37 without DNS resolution.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	6b7482e846643938b97e0078379763c5
C2	180.178.53.37

T5000 is the first instance we see of the “DTL” name, using Dtl.dat as the name of the network configuration file. Although this campaign does not use DTL domain names for C2 servers, we can identify it as part of this cluster due to the shared sender IPs. All of these emails used gmx email account and were sent from either 66.103.141.24, 69.73.160.142, 65.124.5.107, 64.124.5.107, or 209.234.204.31.

In November 2013, [Cylance reported](#) on attacks using T500 that targeted human rights groups and the automotive industry. The name they gave this threat actor was “Grand Theft Auto Panda,” as “they appear to be punching people in the face and stealing their cars.”

CAMPAIGN 2: LURK0 (SOFTMY.JKUB.COM)

In February 2011, Tibet Group 1 was sent two identical emails using LURK0 malware that used softmy.jkub.com as a C2. While this campaign does not use DTL-related C2s, the LURK0 samples have features that are otherwise unique to the DTL cluster. The samples originally created a registry key named “DbxUpdate” and a mutex named “111.” The components in these samples have compilation dates of “2010-09-26 04:31:01” and “2010-12-09 03:22:21.”

The emails sent repurposed text about Chinese authorities clamping down on activists following an online call for protests.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	37457f46709b793d13a25da0d4c896fa
C2	softmy.jkub.com

In May 2011, Tibet Group 1 was again targeted by attacks using LURK0 samples that connected to softmy.jkub.com. These samples created a folder named “DbxUpdateET” and a mutex named “ET” with compilation times in March 2011. The May and February attacks all utilized DLL hijacking of linkinfo.dll to maintain persistence on the system. The May emails, shown below, described a recent award ceremony

hosted by the Canadian Multicultural Council. We were unable to find references to this event online and therefore cannot confirm if it was real.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	f024c4febb69195750c7af06e15aa1f7
C2	softmy.jkub.com

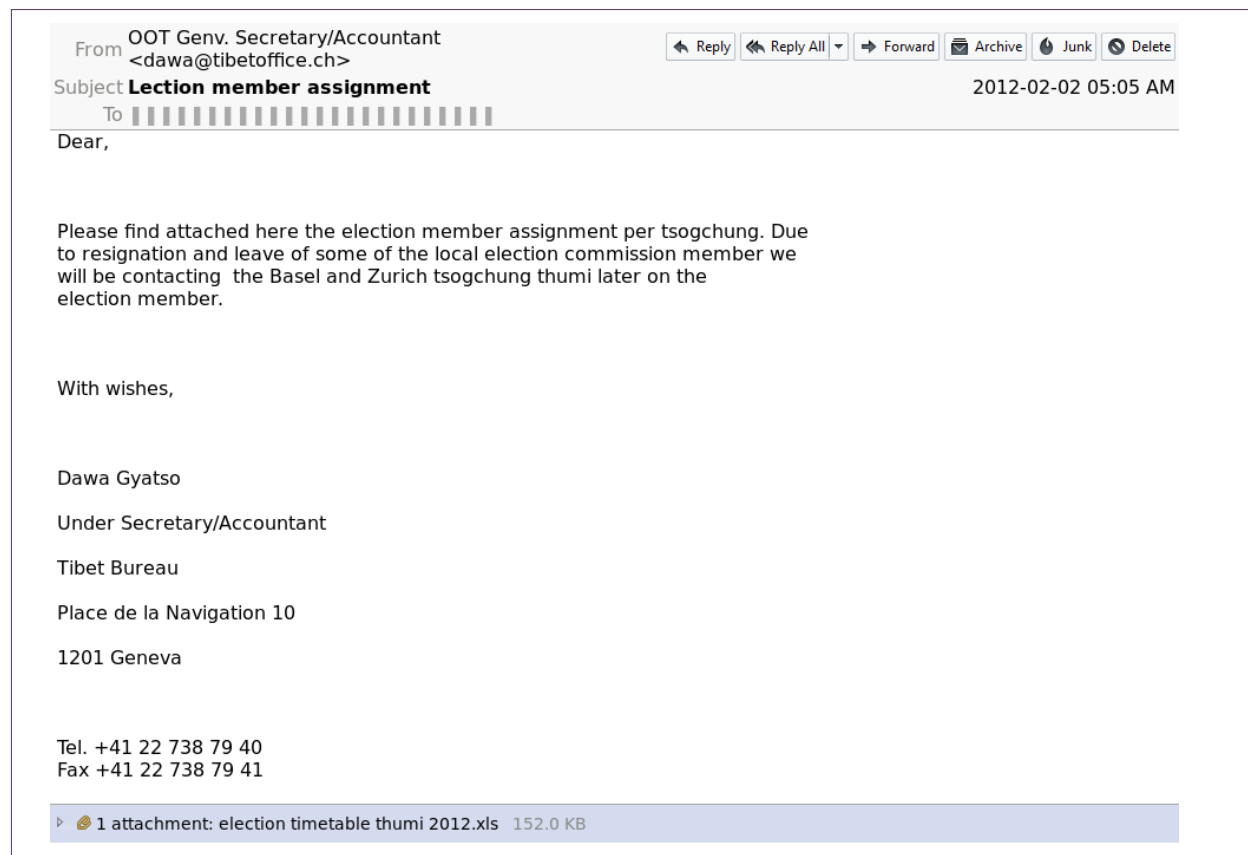
CAMPAIGN 3: BOOUSER (DTL)

In December 2011, multiple members of Tibet Group 1 were targeted by identical emails that appeared to be from a funder providing a report from a grantee. The attached word document dropped Boouser malware that connected to dtl.dnsd.me as a C2. This campaign is the first instance of attacks using DTL-related domains. Boouser is a simple piece of malware (technical score 1.0) with limited code obfuscation that sends unencrypted data back to the C2. It features standard RAT capabilities including a keylogger and the ability to execute remote commands. The social engineering score of these attacks is 3 (TTI 3).

CAMPAIGN 4: LURK0 - CRAZYTOWN EDITION (DTL)

From February 2 to March 14, 2012, a campaign of 10 LURK0 attacks targeted Tibet Groups 1 and 2 using dtl.dnsd.me and dtl.eatuo.com domains as C2s. These were the first LURK0 attacks to use DTL domains as a C2. These samples also performed DLL hijacking on linkinfo.dll and created a key named DbxUpdateET. The mutex name was changed to “ETUN.” These samples use the internal name “ButterFly.dll.” Nine of the attacks used a common tactic of attaching a rar file containing benign jpeg image files and shortcuts that actually link to a LURK0 dropper. These samples had August 15, 2011 as a compilation date.

The emails referenced a number of topics including writings from a Tibetan activist and a recap of a rally held the day before to commemorate the March 10th Tibetan Uprising Day. Another email spoofed the legitimate email address of an individual at the Tibet Bureau in Geneva, and attached a malicious document containing information on the organization of an undetermined election.



Social engineering	3
Technical	1.25
TTI	3.75
MD5	216ca9c711ba602be9ec626d1d44ff99
C2	dtl.dnsd.me (192.198.85.101)

CAMPAIGN 5: LURK0 UNDER DEVELOPMENT

In late March 2012, another LURK0 attack was observed that had considerable differences from the previous wave of attacks. Although this new attack also used compressed .rar files as a vector, unlike previous attacks it did not perform a DLL hijack, instead dropping a file called win32.exe. This file writes several new files to disk: IconConfigEt.DAT, containing a DLL with the core RAT functionality; iexplore.exe, which copies IconConfigEt.DAT to IconCacheEt.DLL, overwrites the DAT file and then runs the DLL file; and temp.exe which simply creates a shortcut to iexplore.exe. This functionality changes the persistence mechanism from DLL hijacking to the creation of an executable that launches on startup with the innocuous name of 'iexplore.exe.' Once launched, this executable runs the DLL. This sample also features a new mutex name: "ER." These emails spoofed recognized Tibetan NGOs and referenced content about self-immolations in Tibet, and as a result scored 3.0 on the social engineering sophistication base value. The technical score for these emails was 1.25 for a total TTI of 3.75.

An additional attack in May 2012 indicated further development. This attack utilized a trojaned Word document as the attack vector with password protection to hinder AV detection. The sample was similar to malware seen during the March attack but this time with a different mutex name ("ERXXXXXXX") and different names for the droppers. Additionally, instead of using a separate DLL, the sample dropped two files named iexplore.exe. One of these files simply ran the other which was signed with a digital certificate issued to Shenzhen OuMing Keji Co., Ltd.

Emails sent to Tibet Group 1 as part of this attack included repurposed text about self-immolations in Tibet, as well as an email on celebrations of the birthday of HHDL.

CAMPAIGN 6: LURK0 (DTL)

In July 2012, another LURK0 campaign of 11 attacks emerged that targeted Tibet Groups 1, 2, and 4 using the dtl.dnsd.me domain as a C2. These LURK0 samples have additional features compared to prior versions. The version of the zlib compression library used for encrypting communications between infected hosts and the C2 was upgraded from 1.1.4 to 1.2.3. These samples also created an executable called iexplore.exe (instead of performing DLL hijacking like in the earlier attacks). However, compared to the previous attacks they featured fewer layers of droppers and extra files. Configuration data like campaign codes and C2 information were changed to being stored in configuration files that could be easily modified. These samples all featured a compilation date of “2012-05-28 05:35:16” and a mutex name of “ERXXXXXXX” while maintaining the “DbxUpdateET” registry key name.

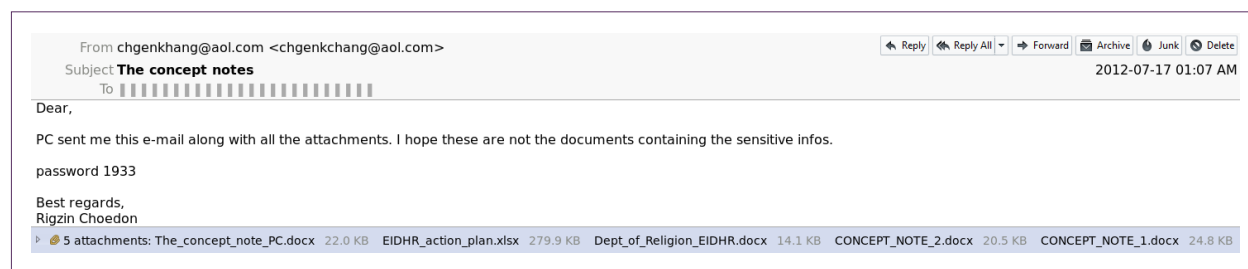
Two of these attacks had lures related to HHDL’s birthday and contained encrypted Word files with the password contained in the message body.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	f2a0787388dd6373336b3f23f204524a
C2	dtl.dnsd.me (184.105.64.183)

Five of these attacks used .doc implants with decoy documents containing what appears to be a Tibetan organization’s legitimate proposal to the European Instrument for Democracy and Human Rights (EIDHR). The timing of these attacks is noteworthy, as a [genuine EIDHR call for proposals](#)—including proposals for

“Actions Aimed at Fighting Cyber-Censorship and to Promote Internet Access and Secure Digital Communication”—was pending at the time, with a July 20 deadline for concept notes. The emails were received by the groups on July 16 and 17, just a few days before the deadline.



Social engineering	3
Technical	1.25
TTI	3.75
MD5	0fe550a5d1187d38984c505ef7741638
C2	dtl.dnsd.me (184.105.64.183)


Three attacks had identical email lures referencing a South African group’s visit to Dharamsala, India that appear to have been repurposed from a legitimate private communication. The email appears to be a request to the Tibetan organizations hosting the planned trip, with the malicious attachment containing an authentic travel itinerary as a decoy. This is a highly targeted attack based on private communications, and as a result receives the highest social engineering sophistication score (5, TTI 6.75).

CAMPAIGN 7: CCTV0

In November 2012 a campaign targeted Tibet Groups 1, 2, and 4 using dtl.dnsd.me and dtl.eatuo.com domains as C2s. The first wave of these samples had compilation dates of October 15, 2012 and later samples had compilation dates of November 11, 2012. These samples changed the five-character code visible in network traffic from ‘LURK0’ to ‘CCTV0’, which prevents strict IDS rules looking for “LURK0” in network traffic from detecting the malware. The samples featured an embedded DLL with the internal name “ETClientDLL.dll” instead of “ButterflyDll.dll” which was seen in earlier attacks. The initial samples in this wave would query a benign third-party website to determine

the user's IP. When this website's results page was modified by its creators, the samples would crash trying to parse the page, leading to the feature being removed in later samples. Another significant change to later samples was padding of the resource section of the executables with extra data, resulting in a larger file size to avoid AV heuristics. Executables were padded with random data to get a different hash every time, making it more difficult for malware researchers and AV companies to share indicators.

TTI scores for the 10 emails varied. Three identical emails were sent to Tibet Group 1 and multiple accounts at Tibet Group 2. The lures used in these attacks were relatively poorly customized.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	1c44d9cf686f53f1194cdee2aefb99c2
C2	dtl.dnsd.me (199.36.72.214)

EXTENDED ANALYSIS: 2.2 Cluster Analysis

Later attacks included lures with more detailed information.

From Tenzin Gyalpo--Prinvate Office <ohhdl@gmail.com>

Subject **Schedule**

To [REDACTED]

[Reply](#)
[Reply All](#)
[Forward](#)
[Archive](#)
[Junk](#)
[Delete](#)

2012-11-20 04:46 AM

Dear Supporters and Friends:

In the attachment of "Schedule.xls" is the public schedule of His Holiness the Dalai Lama both in India as well as abroad. Please note that for many of these events, tickets are required in order to gain entrance. People are requested to contact the organizers directly or visit the websites given below for further information on tickets. In general, most of the events in India are free where as the majority of events abroad require paid tickets. For your information, as a long-standing policy His Holiness the Dalai Lama does not accept any fees for his talks. Where tickets need to be purchased, organizers are requested by our office to charge the minimum entrance fee in order to cover their costs only.

Please note that the dates are subject to change.

Contact

Office: ohhdl@dalailama.com

Website Feedback: webmaster@dalailama.com

Mailing address:
 The Office of His Holiness the Dalai Lama
 Thekchen Choeling
 P.O. McLeod Ganj
 Dharamsala
 Himachal Pradesh (H.P.) 176219
 India

Telephone:
 91 1892 221343
 91 1892 221879

Fax:
 91 1892 221813

1 attachment: Schedule.xls 141.5 KB

Social engineering	3
Technical	1.25
TTI	3.75
MD5	16b82aa9f537811490fdf2e347ec106f
C2	mychangeip1.ddns.info (110.152.229.247)

The changes seen in samples used in this campaign provide insight into the development path of the LURK0 family. Although we did not find C2 infrastructure overlap with other campaigns using LURK0, we did find other similarities. Other campaigns

that used LURK0 continued to use an older version of zlib, retained the LURK0 network header, and used different internal names (such as continuing to use the Butterfly moniker). A possible explanation for this observation is that while threat actor groups may share tools, development and customization of malicious software is decentralized.

CAMPAIGN 8: SURTR

From November 2012 until September 2013, the primary malware used in the DTL campaign changed to a new family called Surtr.¹¹ These attacks continued to use the same C2s as the earlier families. This malware targeted Tibet Groups 1, 2, and 4. Unlike other families in the DTL cluster, Surtr downloads an additional component that contains its main functionality after infection. We have seen two versions of this used with the internal names Remote and GmRemote.

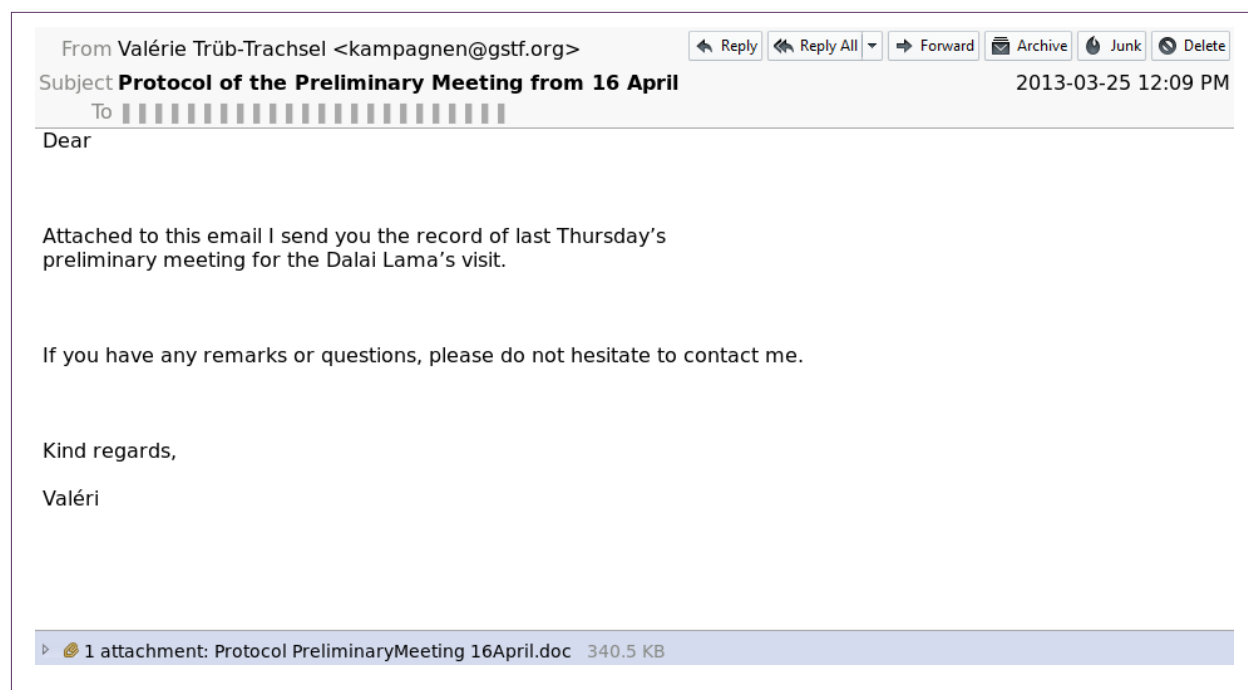
Although the Surtr and LURK0/CCTV0 malware families do not share a large amount of code, they exhibit similarities in behaviour. While some of these similarities, such as the use of zlib in both LURK0/CCTV0 and Surtr, are likely coincidental, others are much more specific. For example, similar registry key names used for configuration information and campaign codes, expanding of the resource section to avoid identical hashes, and similar formatting for sending system information are some of the similarities. Internal names used for the NSFree family and the several LURK0/CCTV0 variations follow a similar scheme, such as the filenames ‘NSFreeDll’ and ‘BTFreeDll’ and the creation of folders named MicET, MicBT, and MicNS.

In addition to these similarities, LURK0/CCTV0 and Surtr have also been used in conjunction with one another. For example, during our analysis we observed LURK0 being downloaded and installed as a stage two after initial infection with Surtr.

11 We first reported technical details on Surtr in Kleemola, K., Hardy, S. “Surtr: Malware Family Targeting the Tibetan Community” Citizen Lab, August 2 2013, <https://citizenlab.org/2013/08/surtr-malware-family-targeting-the-tibetan-community/>

CAMPAIGN 9: 9002

In the first half of 2013, we observed three emails sent out to Tibet Groups 1, 2, and 4 containing the 9002 malware. On March 25, an email was sent to Tibet Groups 1 and 2 with a CVE-2012-0158 attachment from a gmx account with sender IP 66.103.141.24. A different email sent from another gmx account, with sender IP 64.124.5.107 and another malicious attachment, was then seen the next day, March 26, sent to the same groups.



Social engineering	3
Technical	1.25
TTI	3.75
MD5	2c8ef540ae77f1184ddfd3e0a1f810b
C2	dtl.dnsd.me (74.121.190.38)

On May 13, 2013, we saw a similar pattern. A new email was sent from 209.234.204.31 using a gmx account to Tibet Groups 1, 2, and 3, again using CVE-2012-0158. This email was then seen again on May 14, sent to many other targets including more staff at Tibet Group 1, and a number of other Tibetan NGOs, and CTA offices.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	22640ef1d8663a45653d2a6e12604b09
C2	dtl.dnsd.me (74.121.190.38)

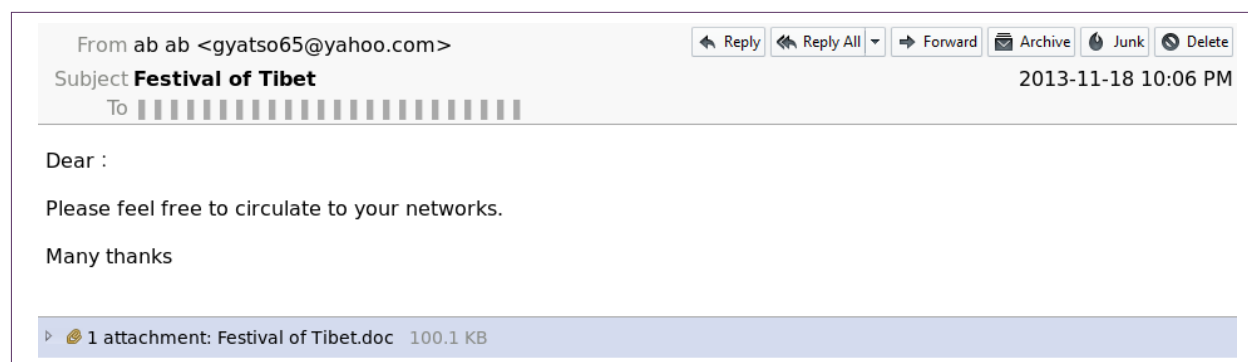
All of the 9002 samples connected to C2 servers at dtl.dnsd.me and dtl.eatuo.com on port 3123. This campaign uses campaign codes of the form “Tmdd,” where m is the month and dd is the day (e.g., T315 for emails sent on March 15).

An interesting feature of the 9002 malware is that it shares exported function names and embedded filenames with Surtr, making it very likely that it was developed alongside Surtr by the same group.

RELATED 3102 CAMPAIGN

3102 is a family of malware that appears similar to 9002, but with additional protection and anti-reversing features. We observed one campaign using 3102 with techniques similar to the original 9002 campaign.

On November 18, 2013, Tibet Groups 1 and 2 each received an email with a Tibetan theme from a Yahoo! address. While these emails contained the same subject, body, and attachment using CVE-2012-0158, they each had different recipient lists visible in the To: and Cc: headers. The inclusion of visible recipient lists is a method also used in the 9002 campaign.



Social engineering	2
Technical	1.25
TTI	2.5
MD5	6bd6b50af9361da2361ff34a8ca99274
C2	183.86.194.130

OBSERVATIONS

The DTL campaign is notable for the variety of malware families used, the active development cycles of the malware, and connections to targeting of government and private industry.

The malware development cycle used here was easily traceable, and during the course of our study we were able to identify a large number of changes made and their

effects. While most campaigns focused on the use of a small number of malware families, we identified upwards of 10 distinct families in this cluster. Although some of these families were variants of each other, the range of malware used demonstrates the adaptability of attackers and their persistence in developing new techniques to compromise their targets.

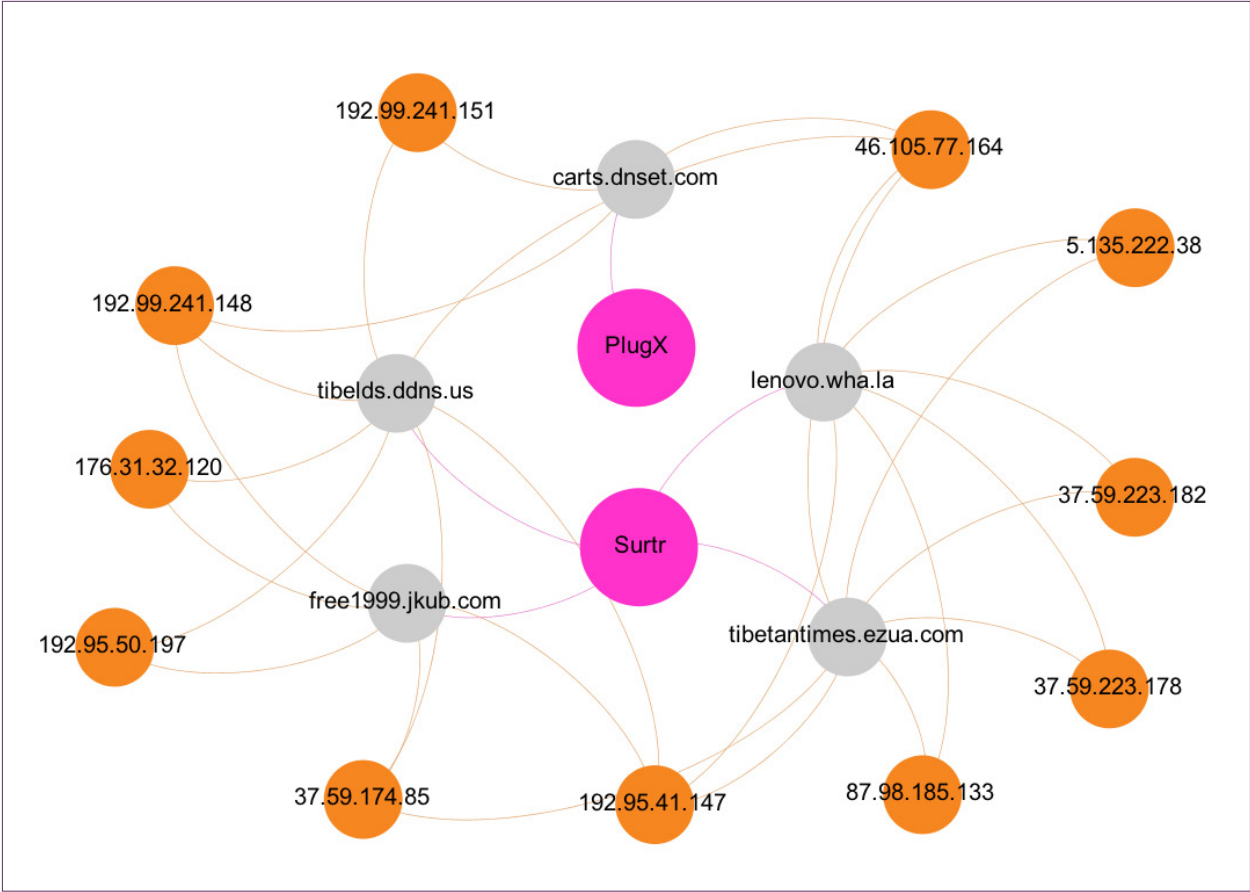
In this campaign we see DTL-related infrastructure that was also used in the attacks against industry and government targets that were reported by FireEye. However, we see little overlap in the malware families. The DTL attacks reported by FireEye exclusively used 9002, whereas we see 9002 and nine other families in our dataset. This lack of similarity suggests that DTL operators may differentiate the tools used in their operations based on the target type.

Ongoing Surtr Campaigns

First Seen	August 1, 2013
Last Seen	ongoing
Attack Vectors	Targeted malicious emails
Exploits	CVE-2012-0158
Malware Families	Surtr (GtRemote, Remote), PlugX
Infrastructure	C2 Domains: carts.dnset.com; free1999.jkub.com; kevin.zzux.com; lenovo.wha.la; patton.mrslove.com; tibelds.ddns.us; tibetantimes.ezua.com; zeeza.info
Targeted Groups	Tibet Group 1, Tibet Group 2, Tibet Group 4, Tibet Group 5
TTI Range	1.25 - 3.75

A new campaign using Surtr as the primary malware family emerged in August 2013, about one month after the DTL attacks stopped. This campaign's C2 infrastructure consists of free ChangeIP domains. This campaign uses throwaway AOL and Gmail accounts designed to impersonate real people and legitimate organizations to deliver malicious emails. There is no overlap with infrastructure seen in any other campaigns, but this could simply be the result of the use of dynamic DNS and free subdomains.

FIGURE 20: The five most commonly seen domains and IP addresses to which they have resolved, per malware family



Although this cluster uses dynamic DNS, the domains will often resolve to the same IPs at the same time, as shown in Table 12.

TABLE 12: IP resolution of Surtr-associated domains per date range

LENOVO.WHA.LA			TIBETANTIMES.EZUA.COM		
2014-06-04	2014-07-10	37.59.223.182	2014-06-04	2014-07-10	37.59.223.182
2014-07-11	2014-08-26	87.98.185.133	2014-07-10	2014-08-29	87.98.185.133
2014-08-27	2014-09-27	5.135.222.38	2014-08-29	2014-09-27	5.135.222.38
2014-09-28	2014-10-23	176.31.149.75	2014-09-27	2014-10-23	176.31.149.75

FREE1999.JKUB.COM			TIBELDS.DDNS.US			*CARTS.DNSET.COM		
2014-06-02	2014-09-10	176.31.32.120	2014-06-02	2014-07-10	176.31.32.120	x	x	x
2014-09-10	2014-09-10	37.59.223.183	2014-07-11	2014-08-11	46.105.77.164	x	x	x
x	x	x	2014-08-12	2014-09-10	192.99.241.151	2014-09-05	2014-09-05	192.99.241.151
2014-09-10	2014-10-23	192.99.241.148	2014-09-11	2014-10-23	192.99.241.148	2014-09-12	2014-10-23	192.99.241.148

This is the only campaign outside of DTL that we have seen use Surtr. Like the DTL campaign, the primary attack vector is malicious email attachments using CVE-2012-0158.

We have identified 67 attacks related to this campaign since we first discovered it in August 2013. It began with two unique emails sent to Tibet Groups 2 and 4 on August 1, 2013. The message to Tibet Group 2 appeared to be an internal mailing list used by its steering committee and staff. The message requests that the list administrator approve a mailing list posting. The message sent to Tibet Group 4 purported to be from “Tibeta Associatio” (sic) and referenced Tibetan autonomy in the subject line but had no message in the email body. Both attacks use .rtf files that drop Surtr (GtRemote) and connect to free1999.jkub.com.

For the first year of the campaign, Surtr was used exclusively. In July 2014, this campaign began using a variant of PlugX. This version removed the identifying strings found in previous versions of the malware. The variant still used the DLL [side loading](#) technique found in previous versions, albeit with a different legitimate executable. It also contained the same functionality. A notable difference was that the malware did not load a properly formed executable into memory, in what appears to be an attempt to hinder analysis.

At the time of writing, this campaign remains the main source of attacks targeting Tibetan Groups, and we are continuing to monitor developments.

End of section. Find the full report at targetedthreats.net

EXTENDED ANALYSIS:

2.3

Civil Society Perspectives and Responses

Over the course of the study we conducted interviews and site visits with nine of the 10 groups participating in our study.¹² The interviews were intended to provide greater context into the perceptions and implications of the attacks documented in this report. The interviews were transcribed and coded to identify emergent themes.¹³ This section reflects those themes as section headers, which are as follows:

- Information communication technologies (ICTs) as an enabler and threat to civil society groups
- How our participants perceive digital risks and threats
- The impact of targeted attacks
- Civil society responses to targeted attacks

Each section, in turn, draws on participants' responses, alongside our synthesis, to provide a window into how groups under threat think about and respond to digital threats. In general, we documented groups at different stages of addressing digital security. Some groups had taken on digital security as a core part of their mission before the study began. Others had only begun to notice issues about digital security a few years ago but by the time of interviewing spoke to us about digital threats as a structural problem for their operations. Still, many were in the process of trying to decide on what measures to take, and how to implement them systematically.

ICT: Enabler and Threat to Civil Society Groups

ICTs are central to the activities of the groups, and help them balance an historic asymmetry between them and powerful, well-resourced state interests.

[Technology] is the only way you can have any serious impact if you're a size like us and you are trying to go up against the Chinese government which has considerabl[y] more resources than us. So if you want to try to

12 We were unable to conduct interviews and a site visit with China Group 2.

13 Analysis followed line-by-line open coding methods and grounded theory approaches described in Creswell, J. W. (1998). *Qualitative inquiry and research design: Choosing among five traditions*. London, UK: Sage.

*have an impact, technology is the leverage.*¹⁴

The two large human rights organizations in our study both operate distributed programming across multiple countries with hundreds of staff and partner organizations. The size and complexity of these operations makes these groups highly reliant on technology.

*Now we do work in more than 80 countries at any given time, and we have staff based in many cities across the world.*¹⁵

The theme of diaspora communications was central to groups related to Tibet and China. The Tibet Groups, for example, find that technology enables them to reach into restricted environments from “safe” areas overseas.

*In comparison to other communities, there is almost nowhere that is so physically separated because of the restraints imposed by China....Even in Pyongyang, you got foreign journalists...Tibet is just closed off....Technology is the only way that that can be bridged.*¹⁶

ICTs are also the primary tool that organizations used to connect with the highly distributed and fractured community in the Tibetan diaspora and in Tibet.

*Those of us born and raised in exile, and certainly our parents and other generations, crave to go home to this land they are so attached to. Then technology comes along and it's like BOOM! You can have it all...on some level even though the Chinese are still there and physically we can't do it, but in this other space we can.*¹⁷

14 China Group 1, Director, 2010

15 Rights Group 2, Technical Officer, 2011

16 Tibet Group 1, Program Officer, 2011

17 Tibet Group 1, Director, 2011

WECHAT: CONNECTIVITY AND RISKS

WeChat is a mobile chat application developed by Chinese company Tencent, which has gained [huge user numbers](#) around the world with a high concentration in China. Tibet Groups cited WeChat, which is highly popular in their community, as an example of the tension between connectivity and security. This connectivity was seen as beneficial, but not without risks.

“New ground shows up like WeChat...and it threatens to both undermine [our efforts] and offers us some ideas for what it is people want and [what they’re] willing to compromise for the sake of connection.”—Director, Tibet Group 1

“Tibetans in India and Nepal, Tibetans in the West are all being connected by using the same app...and forging these new connections...in some ways, we’re seeing really good things come out of it in terms of all the news we’re getting from Tibet and seeing all the footage from the self-immolations and protests coming from WeChat and in that sense it’s becoming important, but people are not so attuned to the risks.”—Program Officer, Tibet Group 1

Tibet Groups voiced concerns over the increasing popularity of WeChat due to censorship and surveillance requirements on companies operating in China, and the close relationship between Tencent and the Communist Party of China. Adding to these concerns are a series of documented incidents of Tibetans in Tibet being [arrested](#) for content they shared on WeChat, like images of HHDL.

Civil society and its champions are not the only groups who felt that technology could enable movements to push back against the status quo. A Tibetan group noted that they thought the Chinese government, was also very concerned about its potential.

...self immolations and protests...the Tibetan cultural pride, the songs...this is the reason the Chinese are cracking down so hard and going after everyone the censorship and the surveillance is...[happening] because the technology has showed them what’s possible. There is a movement now where there wasn’t one or where it had almost disappeared before... the technology has enabled that.¹⁸

Yet the ability to connect is constantly eroded by efforts to monitor and interfere with groups’ activities. Participants also recognized that their reliance on technology intro-

18 Tibet Group 1 Director, 2011

duced new risks of monitoring, coercion, and electronic attack.

[Technology is] this funny thing where it's a life line, and then it's...maybe your ticket to jail.¹⁹

Groups need members of their communities to maximize the use of ICTs but also to do so securely.

We...need this technology, but we need everybody to know how to use it and be able to be secure and be safe.²⁰

Targeted digital threats are also changing how some CSOs see the promise of technology.

I think that civil society is feeling the heat around targeted attacks and surveillance and I think it's affecting the public sphere and our ability to feel comfortable communicating in what used to be understood as a free and open medium.²¹

In practice, CSOs are in a constant process of navigating through new communication environments, and tradeoffs between connectivity and security. Contrasting theories have popularly characterized ICTs either as “liberation technologies” that can empower political movements, or as levers of control for governments to suppress these very same movements.²² Our participants suggest that the reality is somewhere in between.

How Civil Society Groups Perceive Risk and Threats

We tried to elicit participants' informal “threat models” as a precursor to understanding how these models shaped their response to digital risks.²³ All of the groups in our study work on political issues that can potentially be seen as threatening to specific authorities. The context of this work makes many groups perceive the attacks against them as politically motivated.

19 Tibet Group 1, Director, 2011

20 Tibet Group 1, Program Officer, 2011

21 Rights Group 2, Technical Officer, 2014

22 See, e.g., Diamond, L. “Liberation Technology,” *Journal of Democracy* 21, no. 3 (2010): 69–82; Morozov, E. *The Net Delusion: The Dark Side of Internet Freedom* (New York: PublicAffairs, 2011).

23 A threat model assesses the risk and relative impact of threats against an entity that are specific to the context in which it is situated.

Although political concerns were a backdrop to the attacks, groups tended to focus on the tactical goals of the attackers rather than the greater political objectives. Groups described the probable goals of attacks as efforts to “hinder our operations,” “keep an eye on things,” or cause as much “inconvenience and chaos as possible [to] somehow affect our ability to do what we do.”²⁴

While groups shared similar views on the operational objectives of attackers, their sense of risk stemming from attacks depended on the physical proximity of their adversary. A program manager with extensive field experience working with CSOs in multiple countries explained that groups operating within the jurisdiction of an adversary have greater concerns over physical security and other direct interference from authorities.

*...in many places it's a very physical sort of thing. Our biggest challenge...[was when]...local authorities wandered in and took computers. So it's not like we expect the attacks to be all coming in over the wire. In most places where we operate it's probably not even the easiest place for them to get at although it is certainly a lot more subtle.*²⁵

If a group is situated outside of a physical jurisdiction controlled by an adversary then targeted digital threats may be a higher priority concern than physical threats.

*Take for example Ukraine where they aren't necessarily expecting a lot of challenges with their local government, but they might be a target for cross border action, that is one of the places where the digital threats become a particular vector or focus rather than just one of the many things that they are thinking about.*²⁶

DIASPORA AND EXILE COMMUNITY THREAT MODELS

Many of the participants in our study working on China and Tibet issues are embedded within geographically distributed diaspora networks. Their missions often include collecting information from closed areas “inside,” while transmitting other information back in. As the director of an organization working on China explained:

24 Rights Group 1, Chief Technical Officer, 2011; Tibet Group 2, Executive Director, 2014

25 Rights Group 1, Program Manager, 2014

26 Rights Group 1, Program Manager, 2014

Our main target and focus is on the mainland. And so the people who [we] are trying to promote, the people who we are trying to give platforms for their issues and their problems and their voices are inside.²⁷

A central concern of these groups was the security of the bidirectional flow of information to and from at-risk persons in China and Tibet. A group working on China perceived its adversary as primarily interested in this information exchange.

Anyone who doesn't want [our organization] to be able to safely and securely get information from inside...and send it back in. What kind of information? Individual case information, human rights information... Anyone who has an incentive not to have people know about it would have the incentive to compromise our operations and make it hard for us -- and that's a lot of people.²⁸

Due to the antagonistic response these activities elicit from the Chinese government, the Tibet and China Groups generally perceived the actors behind targeted digital attacks on their community to be directed by or related to agents of the PRC.

I think in most cases, [the staff] believe it's coming from China.²⁹

Tibetan groups in exile perceived Tibetans living “inside” as having the highest likelihood and impact of harm from digital attack. To demonstrate the potential consequences of targeted digital threats, a Tibetan security trainer explained how he presents the connection between actions outside affecting at-risk groups inside Tibet.

I kind of force them into thinking about like the risk associated with Tibetans inside Tibet, and also kind of like stress the fact ‘We are in a free country, we don't have to worry about ourselves...What can the Chinese government do to me? Nothing. But what happens when you have done something that's harming somebody else inside Tibet?’ ...that's when a lot of people think about it a little more.³⁰

The threat models of groups working on China and Tibet show a priority given to digital threats due to the proximity of these groups to their perceived adversary and the years of persistent attacks they have experienced.

27 China Group 1, Director, 2010

28 China Group 1, Director, 2010

29 China Group 1, Technical Project Manager, 2014

30 Tibet Group 1, Digital security trainer, 2013

...most communities just aren't...under as persistent, sophisticated and threatening attacks, so they just have not gotten to the same place in prioritizing digital security.³¹

Our study exclusively addresses digital threats and does not cover all potential threats a group may face. While digital threats are a primary security concern for the majority of groups in our study, they are just one piece of a holistic risk environment for CSOs.

Impact of Targeted Threats on Civil Society Groups

Most participants were clear that they saw potentially severe consequences from targeted digital attacks, with the greatest danger being to communications with those “inside,” and hence vulnerable to arrest or harassment. Nevertheless, a few participants surprised us by noting that their organization downplayed the possibility of negative consequences from attacks (or had done so in the past). Often this reaction took the form of citing the ‘openness’ of their organization’s work, and suggesting that there was ‘nothing sensitive’ in their exposed data.

This is an organization predicated on virtues of transparency...when I first came here there were a lot of conversations like, ‘Why should we encrypt our email, we are not hiding anything.’³²

As evidence of attacks against that human rights organization piled up, however, perceptions began to change.

...there has been a sea change in the four years since I have been here...they recognize a riskier threat environment and how it is dealt with.³³

Awareness of risks was, however, still a work in progress.

...in the last year or two any organization that I have been working with in a closed society or dealing with sensitive topics at least has some sort of hazy consideration that this stuff should be a concern and that’s maybe a change from

31 Tibet Group 1, Director, 2011

32 Rights Group 1, Program Manager, 2014

33 Rights Group 1, Program Manager, 2014

*when I first started. But...it's very piecemeal.*³⁴

These responses raise an important issue: the relationship between digital compromises and the use of the compromised information by adversaries is indirect. Unlike the consequences of physical threats, which are often directly observable to a non-expert, the most serious impacts of digital threats are typically at least one step removed from the technology that has been exploited. Making the link between technological threats and “real-world” harm can be challenging, in part due to limited evidence.

*... [there has] been little evidence of the direct impact on people's safety as a result of some of these threats.... [but] our understanding of how surveillance and repressive practices have operated even in pre-digital times provides us with sufficient evidence to understand that there may be a connection. I believe that there has been an increase on the reach of this harm by specialist state actors...*³⁵

CONNECTING SURVEILLANCE AND HARM

We often heard stories of arrest and detention from groups in the study (and through other Citizen Lab projects) that appeared to be linked to electronic surveillance.

Members of the Tibetan community shared with us accounts of Chinese authorities confronting Tibetans with call records and chat transcripts during interrogations. Meanwhile, research on [Ethiopia](#) has revealed that detainees have been presented with similar evidence during interrogations.

In [Syria](#) there are also reports of interrogators presenting detainees with records of communications, and cases where accounts of detainees are used to seed malware to contact lists.

In [Bahrain](#), meanwhile, activists were arrested after posting pseudonymous tweets critical of the government. The real identities of these persons may have been discovered by authorities through a technique in which an attacker sends the pseudonymous Twitter account a link to a webpage containing an embedded remote image. When the victim clicks on the link or opens the email, their IP address is revealed to the attacker. Authorities can then link the IP to the target's true identity through their ISP.

We strongly suspect that these cases are only the tip of the iceberg, and that the digital element in many cases of harm goes unrecognized due to lack of investigation, not lack of incidents.

34 Rights Group 1, Program Manager, 2014

35 Rights Group 2, Technical Officer, 2014

We also think, however, that by downplaying the consequences of targeted digital threats, some participants were showing us something interesting about the resilience and adaptability of their communication styles, which have co-existed with an adversary that has used extensive monitoring for many years.

...it can be a nuisance, it can be a distraction, it can waste time, but...in the grand scheme of things, it's not as though the movement on a whole operates in a way that is dependent upon secure conversations.³⁶

Nevertheless, the same participant was very clear that serious (even physical) harm could come to individuals and groups “inside” through targeted attacks against them or their contacts.

PSYCHOLOGICAL IMPACTS AND COPING STRATEGIES

While tracking the consequences of a targeted attack for networks of trust and reputation can be challenging and require investigation, some participants spoke in detail about the psychological impact of compromise.

It was quite upsetting. I think it sort of paralyzed us emotionally—the two of us that were affected—for a few days.³⁷

In this incident, the emotional harm was perceived as more impairing, and less easily mitigated, than the breach to the computer system.

...the act of cleaning our computers was something that was relatively straightforward...but it was the emotional impact that sort of threw us.³⁸

Further work is needed to document the connection between targeted digital threats and psychosocial strain to move towards a more complete understanding of how targeted individuals and organizations evolve and adapt their coping strategies. One interesting coping strategy prevalent among Tibet Groups was to explain attention from adversaries as a signal that their work was important, and was having an impact.

The reason you are a target is because you are doing something that is bothering somebody and to be proud actually of the work that you do that has drawn the attention of these people who clearly want to mess us up somehow.³⁹

36 Tibet Group 2, Director, 2014

37 Tibet Group 2, Director, 2014

38 Tibet Group 2, Director, 2014

39 Tibet Group 2, Director, 2014

As one participant put it, the challenge was to balance the frustration of being compromised with feelings of encouragement:

...we work so hard the entire day and then at the end of the day when you find out that your website has been attacked and people can't get access to it, you get frustrated, but at the same time you also get more encouraged to know that... feeling that your work is making an impact and the Chinese government has to go to the extent of spending time following us and attacking us and spending large amounts of money just for that.⁴⁰

BEING TARGETED: A TEACHABLE MOMENT

Beyond trainings and awareness campaigns, what brought threats home, unsurprisingly, was being targeted or compromised. Being attacked personalized the problem, and turned warnings into tangible concerns.

[It] is visible for users in places that they understand—your email, your Twitter account—even if they don't understand the implications, the connections...they now see it as something real and personal.⁴¹

Tibet Groups felt that the persistent targeting of their community has helped them raise awareness of digital security and highlight the need for vigilance.

It has made Tibetans more aware of the potential of the Chinese government. We always think about the Chinese government creating problems for us diplomatically, we don't think of the cyber world...and how they can maneuver their way into it.⁴²

[The attacks]... give us as a reminder to be more careful.⁴³

Groups made it clear that greater awareness is a work in progress, and that documenting the connection between attacks and specific harms to individuals and groups is a promising way forward.

The basic goals should be to get people to realize that these threats are real... the

40 Tibet Group 3, Editor-in-Chief, 2014

41 Rights Group 2, Technical Officer, 2014

42 Tibet Group 3, Editor in Chief, 2014

43 Tibet Group 5, Program Officer, 2014

*chain of events from clicking on something, to some person being in peril.*⁴⁴

Civil Society Responses to Targeted Digital Threats

As groups struggled with targeted threats, many placed an emphasis on awareness raising and user education as a primary method of responding. These kinds of strategies are important for civil society and applicable to a variety of problems. The fact that many of the digital threats they experience rely on social engineering makes this a promising direction. However, responses from the groups also suggest that resource constraints, and limits on available technical expertise, may have constrained other avenues of response.

RESOURCE AND CAPACITY CONSTRAINTS

The most commonly mentioned challenge to addressing digital threats was, unsurprisingly, resource issues in general, and technical resources and skills in particular:

*Every community with a struggle is under-resourced, and if this hasn't moved up the priority list, they don't have the capacity to do this type of stuff or implement it or there's not even enough awareness that is needed for them to be able to [get] people to pay attention.*⁴⁵

Organizations in the Global South

These challenges are especially acute for groups in the global South. All Tibet Groups in our study had their operations or a portion of their operations based in Dharamsala, India. In this context, the groups are operating within a refugee community grappling with persistent targeted attacks and conventional development challenges. Resources are sparse. These groups cannot afford enterprise computing infrastructures, or the expensive security solutions adopted by larger, well-resourced counterparts.

Complicating these challenges is the problem of “brain drain” of technically skilled people in the community. Participants told us many Tibetans with specialized technical training

44 Tibet Group 1, Program Director, 2011

45 Tibet Group 1, Director, 2011

leave Dharamsala in search of better job prospects elsewhere. Tibet Groups saw the need to create opportunities for Tibetans with technical skills to work in the community:

*[Tibetans living in exile in India] are essentially a refugee population and all these folks want to get jobs. So, if we can actually bring them in and give them jobs in supporting their own community, that's kind of a goal.*⁴⁶

Others, on the same topic, suggested that the problem was improving, slowly:

*I think now it's coming up slowly and slowly, but there was a time in the Tibetan community where we really lacked webmasters, where we really lacked people who are well equipped and who have good knowledge in terms of Internet security.*⁴⁷

These challenges are not unique to the Tibet Groups. Rights Group 1 explained that local partners supported by the group faced similar technical capacity difficulties:

*Security is hard and it's much harder than it needs to be....the challenge of trying to keep your stuff in some kind of secure state as is currently defined is just well beyond what any typical partner organization is able to deal with...For most of our partners they are lucky if they have a young guy who understands a bit about computers.*⁴⁸

A particularly common problem for groups in the global South is the use of pirated software (unpatched or pre-backdoored software is often a source of insecurity). The use of pirated software is widespread in the Tibetan community due to prohibitive licensing costs. Similarly a program officer in a large rights group explained it is difficult to convince a local partner to purchase a software license “when you can jump out to the local market, [and] for a dollar buy a disk.”⁴⁹

Larger Organizations

Two of our participating organizations, Rights Groups 1 and 2, had significantly higher technical capacities and financial resources, and they approached information security in a manner similar in some ways to a large company. For example, the groups have senior management in charge of security and technology, IT support teams / help desks, and occasionally hire companies to provide security consultations

46 Tibet Group 1, Program Director, 2011

47 Tibet Group 5, Program Officer, 2014

48 Rights Group 1, Program Manager, 2014

49 Rights Group 1, Program Manager, 2014

and incident response services. However, although the size and resources of these groups afford them certain advantages over smaller groups, they experience equally vexing security challenges.

The complexity of global operations and distributed staff and partners creates problems for introducing and sustaining security awareness, even as attacks seem to increase:

*We have seen a colossal uptake on attacks on the home office or we are just much more aware of them than we used to be. We anticipate that such things are also happening at the field office level and to our partners, but we have much less visibility into that.*⁵⁰

The lack of network visibility among in-country offices was cited as a particular challenge:

*We don't have a unified network with all our field offices... so we don't have the same enterprise level of security and capacity there...[the field offices and NGO partners] have to face a range of threats that are from the physical world as well.*⁵¹

Both groups face challenges adapting technology policies for regional offices and partner organizations. Rights Group 2, for example, contended with securing its head office and maintaining awareness of threats faced by field offices:

*There's not a lot of security awareness in the organization. There's ... small pockets of knowledge, but the rest of the organization will prove to be the weakest link....people don't understand, especially people that work in the field don't understand the sensitivity of the work the organization does, so they tend to be a bit more lax about... certain things.*⁵²

Bureaucratic processes were seen as hindering the adoption of new security policies, given the challenge of informing decision makers about emerging security issues:

*I think it has been very top-down, you know some [policy] comes from the top, they go to the bottom and there is no way to inform what's going on in the decision process.*⁵³

50 Rights Group 1, Program Manager, 2014

51 Rights Group 1, Program Manager, 2014

52 Rights Group 2, Technical Officer, 2014

53 Rights Group 2, Technical Officer, 2011

TRAINING AND USER EDUCATION

While the two large groups were able to invest in security appliances and dedicated technology support, the other groups in our study focused on user education and awareness as the primary security strategy.

The majority of groups focused on internal training programs and training with partner groups. Five were able to conduct these trainings themselves,⁵⁴ while four others⁵⁵ drew on third-party support. These trainings varied widely, ranging from short explanations of security policies to sustained user education programs.

Several key themes emerged from our interviews about digital security trainings: the value of understanding the local context, the need for training based on organization-specific threat models, and the value of focusing on behavior rather than simply teaching a wide range of tools.

Training informed by local context and threat models

While user education and training were a major part of groups' strategies, most highlighted the importance of situating trainings in local context and using accessible language and concepts.

*If we could break it down for people in a way that they understand, if we could give them metaphors and other ways to understand what exactly this means for us, and paint the bigger picture, it has an impact.*⁵⁶

Rights Group 1 explained that conducting formal risk assessments of its partner organizations was key in developing appropriate educational strategies.

Other groups engaged in training shared similar comments and noted the importance of ensuring trainings are in line with both technical and contextual realities. Several interviewees pointed out that keeping abreast of new technical developments and context-specific risks was challenging and time-consuming, highlighting the value of intermediary organizations that perform this role within a particular targeted civil society context. Indeed, Tibet Group 1 went so far as to structure its mission to focus on digital security awareness and education programs for the Tibetan community. The

⁵⁴ Rights Groups 1, 2; China Groups 1, 3; Tibet Group 1

⁵⁵ Tibet Groups 2, 3, 4, 5

⁵⁶ Tibet Group 1, Director, 2011

group provided training support to all the other Tibetan organizations who participated in the study.

Moving Beyond Tools

All of the groups identified a common set of user practices for preventing infection of malware: not opening unsolicited attachments, being careful with web links, keeping systems up-to-date, and generally remaining vigilant online. Explaining the safe and secure use of tools was an aspect of training, but many groups focused more on how to change behaviour and develop a security mindset rather than train specific tools.

*[We are] trying to equip people with a different mindset, so that they are changing their behaviors...so they...run through a mental filter before doing something.*⁵⁷

Tibet Group 1, which regularly provided trainings to its peers, was particularly adamant about the need to focus on user behaviour over specific tools:

*We would really like to see resources shift from trying to mitigate problems through tools, to mitigating problems through education and educating people about their practices.*⁵⁸

The Tibet Groups felt that user education had to be a community-wide effort and not something isolated to particular organizations or individuals. For example, some Tibetan groups have been promoting a “Detach from Attachments” campaign that encourages users to move away from sharing documents through email attachments and shift to alternative cloud-based platforms like Google Drive. The campaign uses a mix of humor and references to Tibetan culture and is a good example of user education that is connected to a specific threat model and local context.

Encouraging behavioural change and implementing new organizational policies can be challenging. Tibet Group 5 explained that while malicious attachments were a priority threat for the group, moving to alternative document platforms was difficult due in part to generational gaps in the group’s membership. Users from older generations, they explained, were resistant to changing familiar practices, like the use of attachments.

Understanding these organizational challenges and breaking down trainings into simple incremental steps that can be adapted to specific environments were identified as keys

⁵⁷ Tibet Group 1, Director, 2011

⁵⁸ Tibet Group 1, Program Director, 2011

to success. A Tibetan security trainer described the importance of showing people small victories from their point of view and demonstrating how they can learn to achieve ICT objectives in ways that are safe and secure, but also do not appear too difficult for them to use in their daily workflows.⁵⁹ Others agreed, saying that their goal was to find techniques that made it “as simple as [possible] to do the right thing...”⁶⁰

Indeed, there were some cases where, if implemented effectively, modest behavior modifications could have a considerable impact. While we observed organizational challenges in implementing practices like “Detach from Attachments,” based on what we have seen, the campaign could be effective against some of the current threats against the Tibetan community. More than 80% of malware submitted to us by Tibet Groups used a malicious email attachment. Furthermore, for two of the Tibet Groups in our study, simply not opening attachments would mitigate more than 95% of targeted malware threats that use email as a vector.⁶¹

However, this is just one mitigation strategy focused on a single vector of attack. Threat actors are highly motivated and will likely adapt their tactics as users change their behaviors. For example, it is possible that if every user in a particular community began to avoid opening attachments, attackers would move on to vectors such as watering hole attacks or attacks on cloud-based document platforms.

As the groups themselves noted, user education and awareness-raising activities need to be ongoing, and must be informed by local context, threat models, and the latest technical information.

⁵⁹ Tibet Group 1, Security Trainer, 2013

⁶⁰ Rights Group 1, Program Manager, 2014

⁶¹ This determination is based on two groups that had submitted at least 40 emails.