UNIVERSITY OF TORONTO

MUNK SCHOOL OF GLOBAL AFFAIRS

Join the Global Conversation

The Citizen Lab

Research Brief
Number 33 – February 2014

Mapping Hacking Team's "Untraceable" Spyware

Bill Marczak, Claudio Guarnieri, Morgan Marquis-Boire, and John Scott-Railton

This post is the second in a series of posts that focus on the global proliferation and use of Hacking Team's RCS spyware, which is sold exclusively to governments.

Read the report's coverage in Wired Italy [in Italian] as well as coverage in SC Magazine, and Der Spiegel.

Read the first report in the series, "Hacking Team and the Targeting of Ethiopian Journalists."

Read the third report in the series, "Hacking Team's US Nexus."

SUMMARY

- Remote Control System (RCS) is sophisticated computer spyware marketed and sold exclusively to governments by Milan-based Hacking Team.[1] Hacking Team was first thrust into the public spotlight in 2012 when RCS was used against award-winning Moroccan media outlet Mamfakinch,[2] and United Arab Emirates (UAE) human rights activist Ahmed Mansoor.[3] Most recently, Citizen Lab research found that RCS was used to target Ethiopian journalists in the Washington DC area.[4]

- In this post, we map out covert networks of "proxy servers" used to launder data that RCS exfiltrates from infected computers, through third countries, to an "endpoint," which we believe represents the spyware's government operator. This process is designed to obscure the identity of the government conducting the spying. For example, data destined for an endpoint in Mexico appears to be routed through four different proxies, each in a different country. This so-called "collection infrastructure" appears to be provided by one or more commercial vendors—perhaps including Hacking Team itself.

- Hacking Team advertises that their RCS spyware is "untraceable" to a specific government operator. However, we claim to identify a number of current or former government users of the spyware by pinpointing endpoints, and studying instances of RCS that we have observed. We suspect that agencies of these twenty-one governments are current or former users of RCS: Azerbaijan, Colombia, Egypt, Ethiopia, Hungary, Italy, Kazakhstan, Korea, Malaysia, Mexico, Morocco, Nigeria, Oman, Panama,

Poland, Saudi Arabia, Sudan, Thailand, Turkey, UAE, and Uzbekistan. Nine of these countries receive the lowest ranking, "authoritarian," in *The Economist's* 2012 Democracy Index.[5] Additionally, two current users (Egypt and Turkey) have brutally repressed recent protest movements.

- We also study how governments infect a target with the RCS spyware. We find that this is often through the use of "exploits"—code that takes advantage of bugs in popular software. Exploits help to minimize user interaction and awareness when implanting RCS on a target device. We show evidence that a single commercial vendor may have supplied Hacking Team customers with exploits for at least the past two years, and consider this vendor's relationship with French exploit provider VUPEN.

INTRODUCTION

Background: Hacking Team and Remote Control System (RCS)

Hacking Team, also known as HT S.r.l., is a Milan-based company that describes itself as the "first to propose an offensive solution for cyber investigations."[6] Their flagship Remote Control System (RCS)[7] product, billed "the hacking suite for governmental interception,"[8] is a suite of remote monitoring implants (i.e., spyware) sold exclusively to government agencies worldwide.

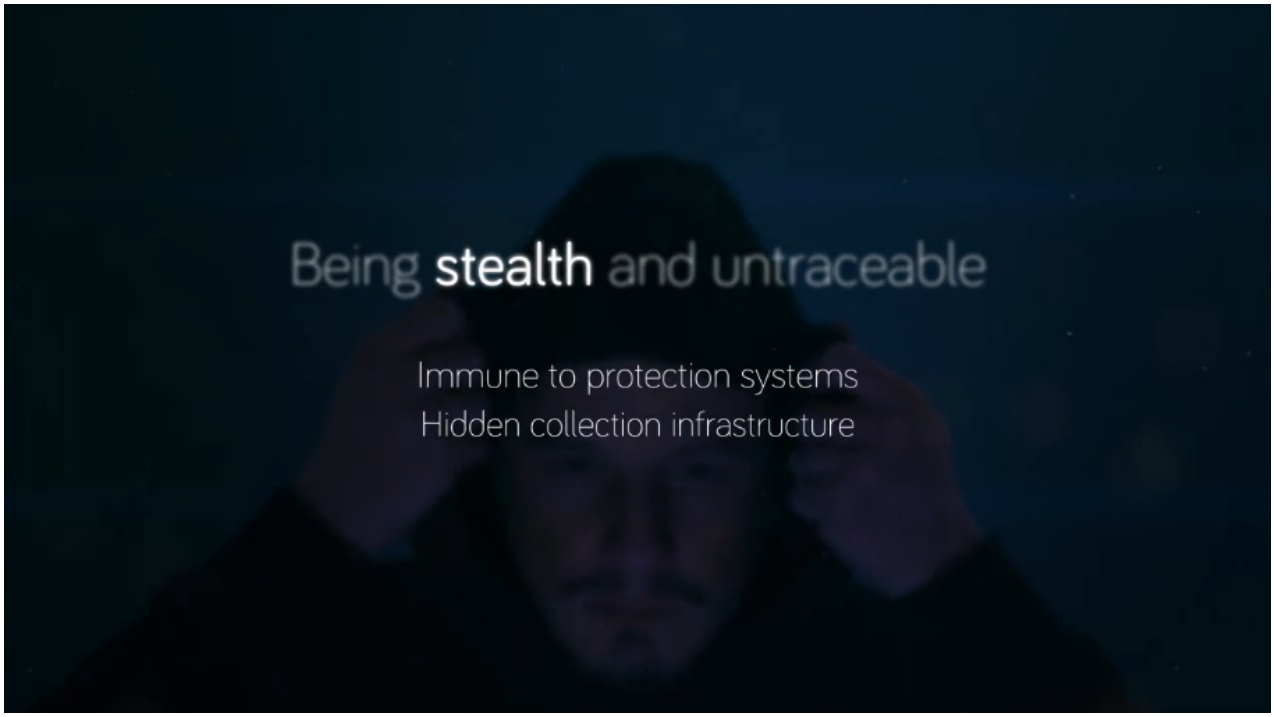

Screen capture from Hacking Team promotional video.

Hacking Team distinguishes RCS from traditional surveillance solutions (e.g., wiretapping) by emphasizing that RCS can capture data that is stored on a target's computer, even if the target never sends the information over the Internet.[9] RCS also enables government surveillance of a target's encrypted internet communications, even when the target is connected to a network that the government cannot wiretap. RCS's capabilities include the ability to copy files from a computer's hard disk, record skype calls, e-mails, instant messages, and passwords typed into a web browser.[10] Furthermore, RCS can turn on a device's webcam and microphone to spy on the target.[11]

While Hacking Team claims to potential clients that RCS can be used for mass surveillance of "hundreds of thousands of targets,"[12] public statements by Hacking Team emphasize RCS's potential use as a targeted tool for fighting crime and terrorism.[13]

Hidden Collection Infrastructure and Target Exploitation

Conclusively linking spyware to a government user is difficult. Clearly, a government must consume the information it gathers from the spyware, but direct communication between an infected computer and a government server would be easily linkable to the government and thus undesirable. Hacking Team advertises that the RCS "collection infrastructure"—the mechanism by which data gathered by the spyware is transmitted to the government—renders the spyware "untraceable" to a specific government.

Our research reveals that the RCS collection infrastructure uses a proxy-chaining technique which is roughly analogous to that used by general-purpose anonymity solutions like Tor in that multiple hops are used to anonymize the destination of information.[14] Despite this technique, we are still able to map out many of these chains and their endpoints using a specialized analysis.

Before a government can receive data, it must first infect one or more target devices with the RCS spyware. Frequently, this takes the form of phishing attacks that convince a user to open a cleverly disguised executable file, or authorize installation of an application. However, the use of exploits, which take advantage of bugs in computer software, can be a more effective technique. Exploits typically require less user interaction before a successful infection (e.g., opening a Microsoft Word document or simply viewing a webpage is enough). Since 2012, we have been tracking exploits that we have seen used to install commercial backdoors. Our research examines connections between these exploits and discuss their origin.

Roadmap

In this post, we begin by outlining our findings, first regarding the governments using Hacking Team's RCS spyware, and then regarding the exploits used to install RCS. We then present the methodology for our findings, including our technique for mapping proxy chains and identifying endpoints. Finally, we conclude by putting our findings into the context of the global surveillance marketplace.

SUSPECTED GOVERNMENT USERS OF RCS

*After extensive analysis, we believe we have identified a set of governments that are current or former end users of Hacking Team's RCS spyware. Our analysis is described in the section entitled "**Identifying RCS Proxy Chains**."*

Introduction

In response to a number of high-profile cases of repressive regimes apparently abusing Hacking Team's RCS spyware, Hacking Team Senior Counsel Eric Rabe stated that the company does not provide its products to "repressive" regimes:

"On the issue of repressive regimes, Hacking Team goes to great lengths to assure that our software is not sold to governments that are blacklisted by the EU, the US, NATO, and similar international organizations or any "repressive" regime."[15]

Hacking Team has also stated that RCS is not sold through "independent agents,"[16] and that all sales are reviewed by a board that includes outside engineers and lawyers. This board has veto power over any sale.[17] Before authorizing a sale, the company states that it considers "credible government or non-government reports reflecting that a potential customer could use surveillance technologies to facilitate human rights abuses," as well as "due process requirements" for surveillance.[18]

21 Suspected Government Users of RCS

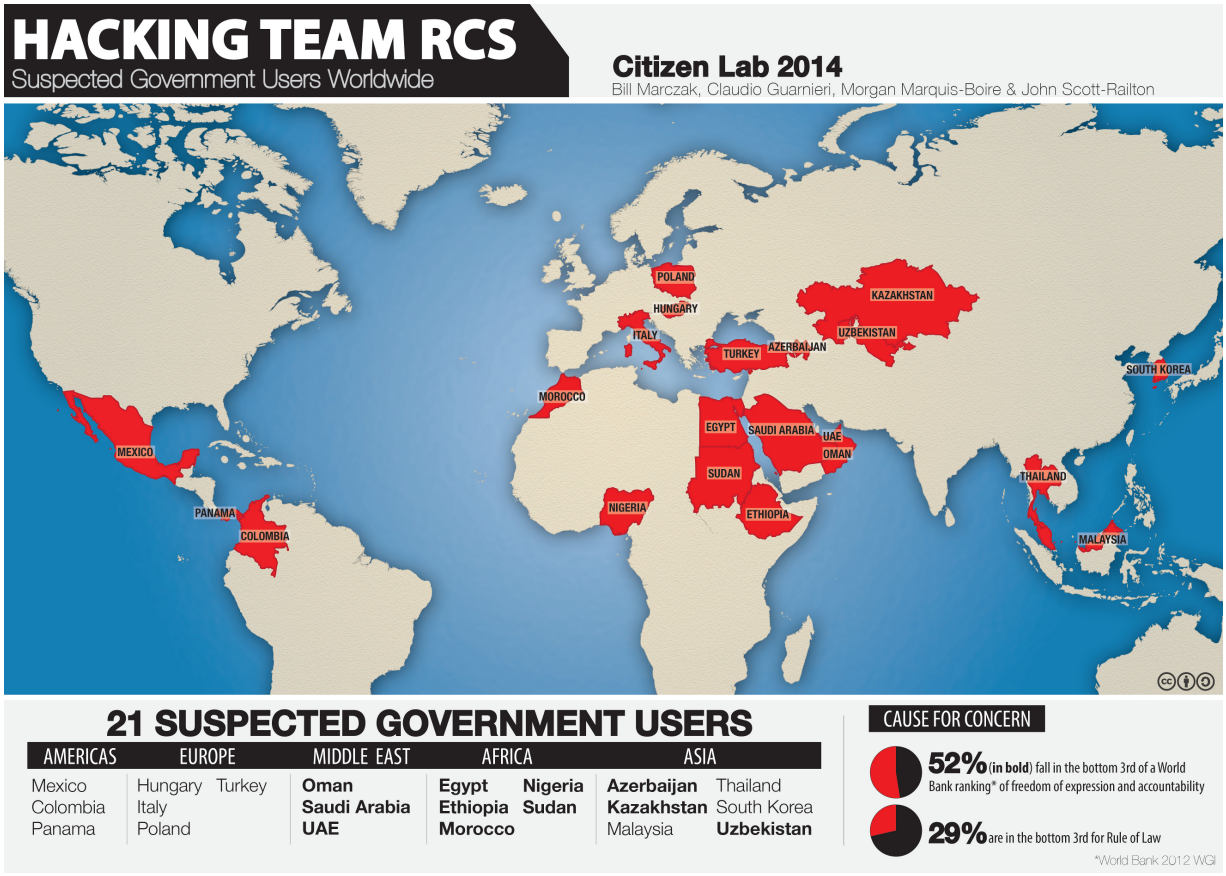

We suspect that twenty-one governments are using Hacking Team's RCS spyware. Except as otherwise noted, we identified these countries based on tracing endpoints of Hacking Team proxy chains: Azerbaijan, Colombia, Egypt, Ethiopia,[19] Hungary, Italy, Kazakhstan, Korea, Malaysia, Mexico, Morocco, Nigeria, Oman,[20] Panama,[21] Poland, Saudi Arabia, Sudan, Thailand, Turkey, United Arab Emirates,[22] and Uzbekistan.

Countries of Concern

While many of these countries are known for their lack of freedom of expression, and politicization of the justice system, several routinely violate basic due process rights, and commit ongoing serious human rights violations. Viewing these countries through the lens of Hacking Team's claimed sales practices, we highlight a few areas of concern:

- **Azerbaijan**: We identified an RCS endpoint in Azerbaijan (Azertelekom: 109.235.193.83) that was active between June and November 2013. Azerbaijan hit international headlines in 2013 when the results of the October presidential elections were accidentally released before voting began.[23] In the run up to the elections, state media released a compromising video of influential Azerbaijani

investigative reporter Khadija Ismayilova recorded from inside her home.[24] While not used in Ismayilova's case, RCS can enable the same type of monitoring by covertly activating a computer's webcam, as well as stealing private videos and pictures from a computer.

- **Kazakhstan**: We identified an RCS endpoint in Kazakhstan(JSC Kazakhtelecom Slyzhebnyi: 89.218.88.xxx). Human Rights Watch (HRW) reported numerous cases of Kazakh government critics facing arrest and detention in 2013. HRW reports that despite the government's adoption of an anti-torture statute, "torture remains common in places of detention," and "perpetrators of torture often go unpunished."[25]

- **Uzbekistan**: We found three RCS endpoints in Uzbekistan (Sarkor Telecom: 81.95.226.134, **81.95.224.10**, and Sharq Telekom: **217.29.123.184**). HRW's 2013 report on Uzbekistan made mention of the government's systematic use of torture with impunity, including beatings and rapes in detention. The report also mentioned lack of respect for due process, including the dissolution of the independent bar association, and forced disbarment of lawyers who take on controversial cases.[26]

- **Saudi Arabia**: We identified two RCS endpoints in Saudi Arabia (Etihad Etisalat: 37.242.13.10 and Al-Khomasia Shipping & Maintenance Co Ltd:[27] 62.149.88.20). HRW has reported "systematic violations of due process and fair trial rights" against detainees in Saudi Arabia, including the use of torture, denial of access to lawyers, inability for defendants to introduce evidence or confront their accusers at trial, and the lack of a penal code or precedent-driven jurisprudence.[28]

- **Sudan**: We identified one RCS endpoint in Sudan (VisionValley: 41.78.109.91), in a range of eight addresses called "Mesbar" (an Arabic word meaning a device used to "probe"). HRW reported that Sudanese authorities used excessive violence against protesters resulting in numerous deaths. HRW also raised concern over a growing number of politically motivated arrests and detentions and noted authorities' suspension of newspapers, and harassment of anti-government journalists.[29]

TARGET EXPLOITATION

How do governments install Hacking Team's RCS spyware on a target's computer? This section briefly outlines one method: the use of exploits, which takes advantage of bugs in computer software.

Introduction

There are many actors that sell exploits to governments. A 2012 Citizen Lab report examined[30] the possible connection between Hacking Team and VUPEN, a Montpellier-based company best known[31] for the sale of "offensive IT intrusion solutions and government grade exploits."[32] While VUPEN offered oblique denials in response to Citizen Lab's report, we have observed that very similar exploits continue to be used in the delivery of Hacking Team's RCS spyware. Recently, it was was reported that VUPEN had sold exploits to the US National Security Agency (NSA).[33]

Given the usefulness of exploits in the deployment of spyware, it is unsurprising that Hacking Team competitor FinFisher GmbH also supplies exploits as part of their "Government IT Intrusion and Remote Monitoring" solution. In the company's own words, "The FinFly Exploit Portal offers access to a large library of 0-Day and 1-Day Exploits for popular software like Microsoft Office, Internet Explorer, Adobe Acrobat Reader, and many more."[34] Hacking Team, FinFisher, and VUPEN promote their products in the same sessions at trade shows.[35]

Exploits

Here we provide a brief analysis of seven related exploits, which we have selected as representative of a larger corpus of exploits. One can discover related exploits by creating signatures based on these seven exploits.

The exploits we examine here all present themselves as malicious documents. It is possible that the vulnerabilities, pre-weaponization, were discovered and/or sold by different actors. However the exploit documents bear enough similarity to suggest that they are produced using the same procedure or program. Six of these documents appear to facilitate the installation of Hacking Team's RCS, while the other installs a remote access toolkit known as SpyNet. This suggests that there may be a common actor supplying the exploits independent of Hacking Team.

Exploit 1

Filename: scandale.doc
Hash: dab3e4423525c798d7937441a3b356d7633a30f229911b9bedd38eeff74717d
Exploit: Adobe Flash in Word document
Target: Moroccan citizen journalist group, Mamfakinch.[36] The message "Svp ne mentionnez pas mon nom ni rien du tout je ne veux pas d embrouilles…" was submitted to the Mamfakinch online news portal, along with a link to http://freeme.eu5.org/scandale%20(2).doc.
Payload: The final payload was Hacking Team RCS.[37] A more in-depth analysis of this attack can be found here.

Exploit 2

Filename: veryimportant.doc
Hash: cd1fe50dbde70fb2f20d90b27a4cfe5676fa0e566a4ac14dc8dfd5c232b93933
Exploit: RTF file with DOC extension; CVE-2010-3333.
Target: United Arab Emirates (UAE) human rights activist Ahmed Mansoor.
Payload: Downloads a second stage[38] from http://ar-24.com/0000000031/veryimportant.doc2. The second stage downloads a Hacking Team RCS payload[39] from http://ar-24.com/0000000031/veryimportant.doc3.

Exploit 3

Filename: الكوفحي رسالة.doc[40]
Hash: c166aff46cadce2db642047cdca65234c32c6634d9ed822eeeb2a911178d6cc3
Exploit: Adobe Flash in Word document; Adobe Flash "Matrix3D" Integer Overflow.
Target: Unknown, but the spyware payload used a command-and-control server (hamas.sytes.net) linked to attacks believed to be conducted by the UAE government.
Payload: Downloads a second stage from https://www.maile-s.com/yarab/stagedocJord. The second stage downloads a SpyNet payload from https://www.maile-s.com/yarab/Win32.scr, and downloads bait content from https://www.maile-s.com/yarab/kofahi.doc.
Analysis: A public mailing list post[41] credits Nicolas Joly of VUPEN for discovering this vulnerability. While VUPEN takes public credit for the discovery of this bug, it is also possible that the exploit used here was not written by VUPEN, but was instead independently discovered and/or weaponized by another party. Examination of the second stage of the payload shows it to be almost identical to veryimportant.doc2, simply using different URLs.

Exploit 4

Filename: مظلوم إماراتي.doc
Hash: 8dbaa77c4db80da6b110e770851932c65c322153fd9edc1df23bd2312584bd94
Exploit: Adobe Flash in Word document; the exploit does not have a CVE number but appears to have been silently fixed in Adobe Flash 11.4.
Target: A human rights activist in the UAE, and a journalist in the UAE.
Payload: Downloads a second stage from http://www.faddeha.com/stagedocuae1. Neither the second stage nor the payload were available for inspection.
Analysis: The metadata for this sample and for Exploit 3 are identical, suggesting that they were generated by the same actors.

Exploit 5

Filename: Biglietto Visita.doc
Hash: c026ebfa3a191d4f27ee72f34fa0d97656113be368369f605e7845a30bc19f6a
Exploit: Adobe Flash in Word document; CVE-2013-5331. Was first seen used in November and December of 2013 and was a 0-day at that time.
Target: Unknown.
Payload: Downloads a second stage from http://176.58.111.219/ipaddrs/shell. Neither the second stage nor the payload were available for inspection. Was uploaded to VirusTotal along with OSX and Windows samples of Hacking Team RCS.
Analysis: The metadata of the document is as follows:

CDF V2 Document, Little Endian, Os: Windows, Version 5.1, Code page: 1252, Title: , Author: unknown, Template: Normal.dot, Last Saved By: unknown, Revision Number: 2, Name of Creating Application: Microsoft Office Word, Create Time/Date: Mon May 14 10:39:00 2012, Last Saved Time/Date: Mon May 14 10:39:00 2012, Number of Pages: 1, Number of Words: 7, Number of Characters: 41, Security: 0

While the previous exploit documents used embedded GZIP compressed flash, this one uses LZMA compression. This compression option was added to the Flash file format recently and is a manual option in Flash Professional CS6. At the time of the attack, it appeared that few anti-virus products supported automated analysis of LZMA compressed flash, making this an effective technique for avoiding detection.

Exploit 6

Filename: 1.doc
Hash: 519939a48ba9dbcc05abfeb4e7fbf9f9dda8c13b567ee6858decaadd09730770
Exploit: Adobe Flash in Word document; CVE-2013-0633.
Target: Unknown
Payload: Downloads a second stage from http://62.109.31.96/0000000025/1.doc2. The second stage downloads a Hacking Team RCS payload from http://62.109.31.96/0000000025/0000000025.exe.
Analysis: The exploit uses LZMA compression. The metadata is almost identical to that of Exploit 5.

Exploit 7

Filename: 1.doc
Hash: 1f9db646053a7bc6be1c8e8ef669079c9e9010306fa537ed555de2387952aa23
Exploit: Adobe Flash in Word document; CVE-2012-5054.
Target: Unknown
Payload: Downloads a second stage from http://62.109.31.96/0000000025/1.doc2. The second stage

downloads a Hacking Team RCS payload from http://62.109.31.96/0000000025/0000000025.exe.
Analysis: The exploit uses LZMA compression. The metadata is almost identical to that of Exploit 5.

Summary

Examination of the exploit document's metadata reveals that exploits 1, 3, and 4 share identical creation and last modification times (2012-05-15T10:39:00Z). Meanwhile, exploits 5, 6, and 7 also share a common time (Mon May 14 10:39:00 2012).[42] Creation time for all six of these exploit documents is 10:39:00, which suggests that all documents were created in the same manner. Additionally, all of the exploit documents described here download a second stage containing shellcode that then downloads and installs a third stage implant. We considered that the similarities in exploit payload and metadata were due to packaging by Hacking Team. However, only six of the seven exploits were used to deliver Hacking Team's RCS. It appears that the exploit document builder is backdoor-independent, allowing customers to select their preferred spyware for post-exploitation.

The exploits discussed above are representative of exploits used to deliver commercial backdoors that we have observed from 2012 until now. They appear to have all been created by the same actor. It appears that Hacking Team partners with a professional exploit vendor that has been providing their customers with exploits for the past two years.

IDENTIFYING RCS PROXY CHAINS

This section outlines the methodology we applied to identify RCS servers, and trace proxy chains to their endpoints. We first describe how we fingerprinted and detected the servers, and then how we traced the proxy chains.

Fingerprinting RCS Servers

We began this research by devising six fingerprints[43] for RCS servers by observing distinctive current and previous behavior (via historical scanning results) of servers listed in files detected as "DaVinci" or "FSBSpy" by at least one anti-virus engine on VirusTotal.[44] Using our fingerprints, we searched a range of public, historical scanning results:[45] the Internet Census,[46] Critical.IO,[47] Project Sonar,[48] Shodan,[49] and the ZMap SSL Scans.[50] See Appendix A for our scan results.

Two of our fingerprints, A1 and A2, are based on the response of RCS servers when they are issued an HTTP GET request. Fingerprint A2 looks for a specific type of webpage redirection, and fingerprint A1 looks for impersonation of the popular Apache Web server.

```
/HTTP\/1.1 (404 NotFound)?(200 OK)?\\r\\n(Connection: close\\r\\n)?Content-Type: text\/html\\r\\nContent-[lL]ength: [0-9]+\\r\\n(Connection: close\\r\\n)?(Server: Apache.*\\r\\n)?\\r\\n/ =~ banner and /Connection: close\\r\\n/ =~ banner and /<meta http-equiv=\\\"refresh\\\" content=\\\"0;url=http:\/\/[^\\]+\\\">/ =~ banner
```

The Ruby Boolean expression for Fingerprint A2 as formatted for the Critical.IO data[51]

Also, when issued an invalid HTTP request, many of the servers return a response with a distinctive typo[52] matching an open source Ruby-based webserver written by a Hacking Team employee. Since the project is open source, we do not treat this condition as sufficient to identify an RCS server.

The four fingerprints, B1, B2, B3, and B4, match SSL certificates returned by RCS servers, which have several distinctive formats. Certificates matching B1 identify the server as an RCS server.

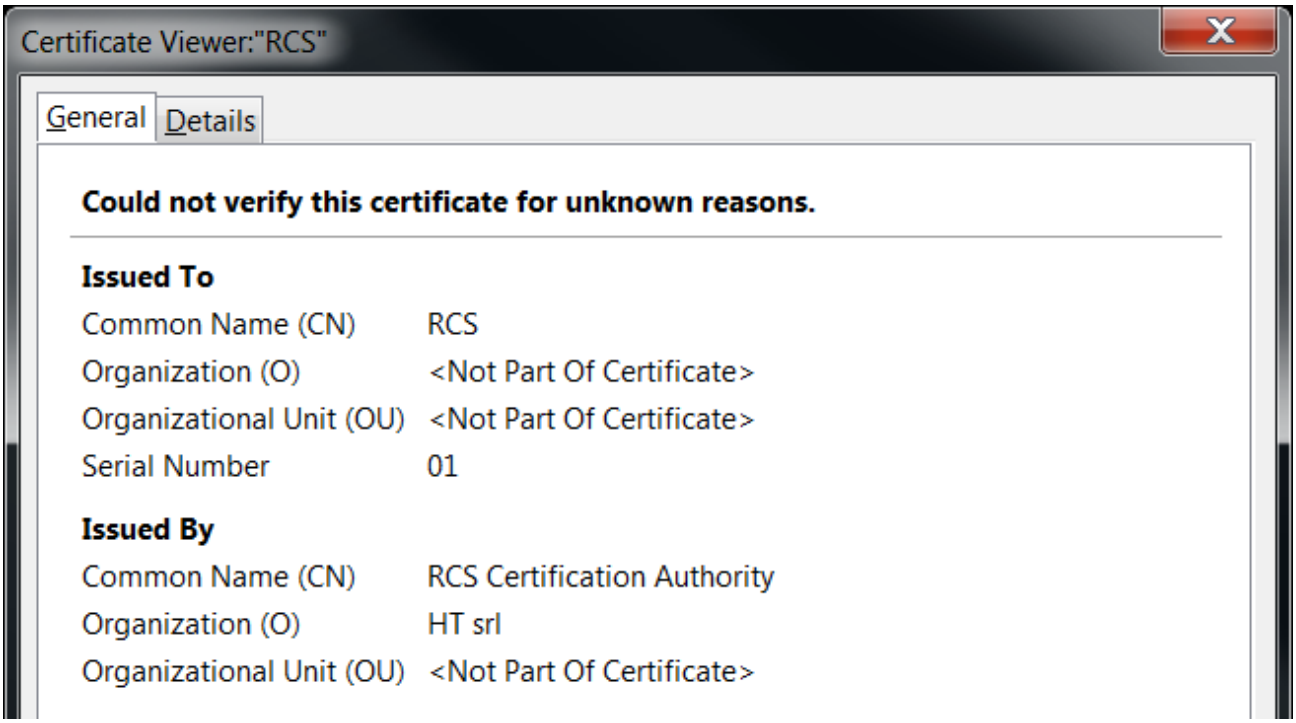

A certificate matching fingerprint B1

While examining SSL certificate results, we noticed that servers in several different countries were returning identical SSL certificates. We hypothesized that these servers were related to each other. As we describe in the next section, we found that servers with an identical SSL certificate were likely associated with a single government operator, and represented proxy chains that terminated in endpoints in a single country.

Identifying Endpoints and Mapping Proxy Chains

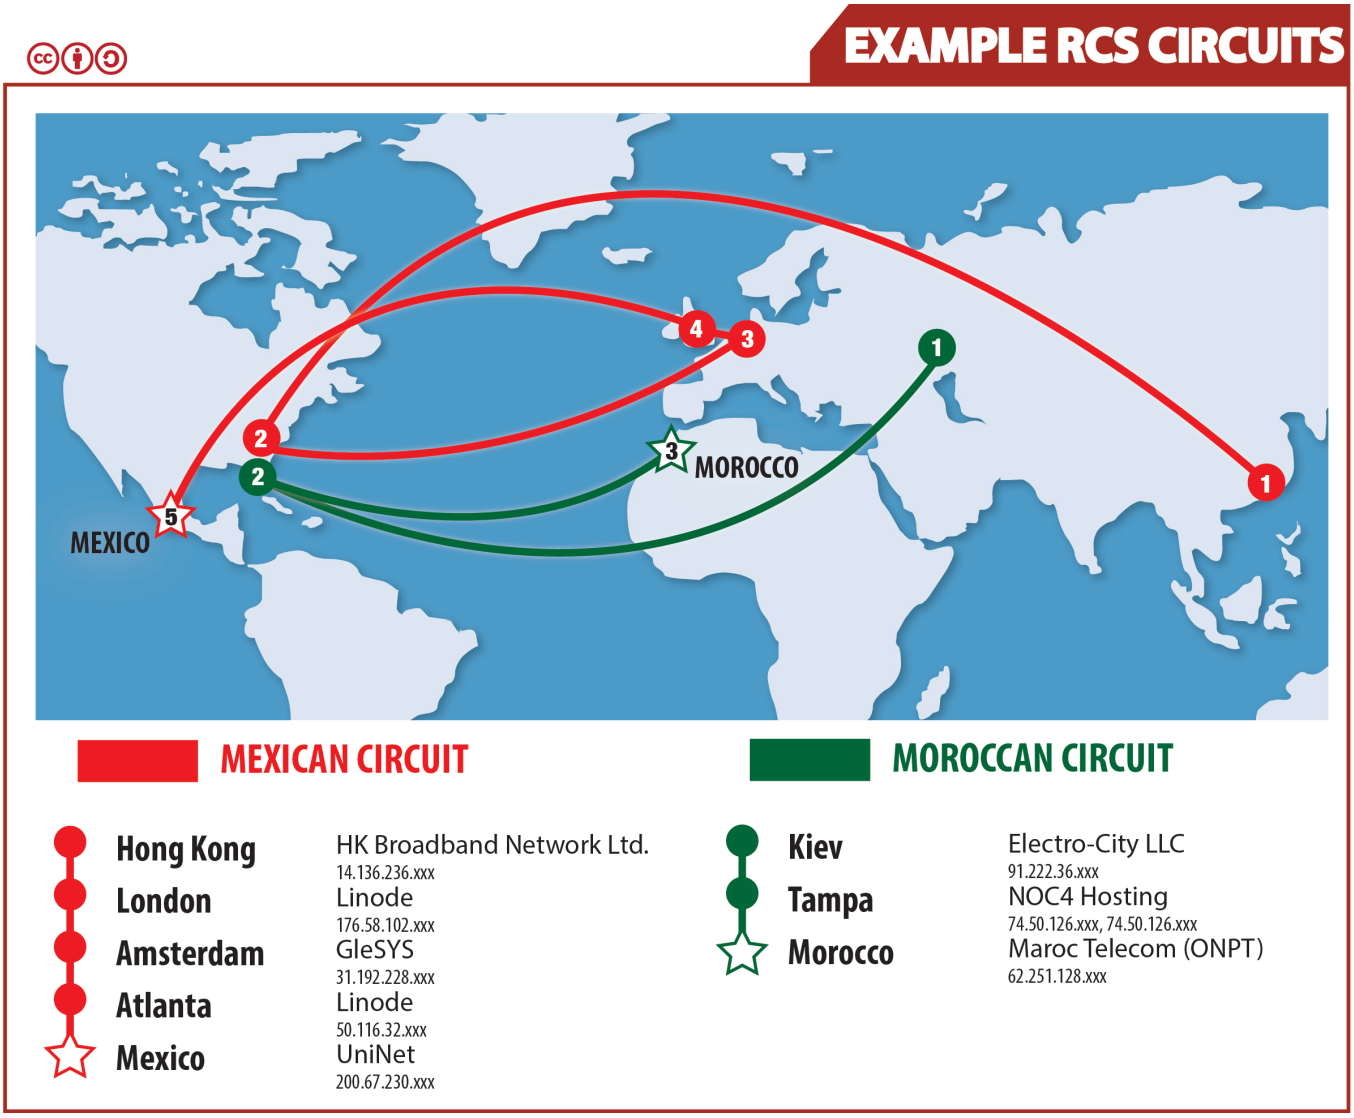

Visual representation of a Mexico and Morocco circuit.

We first provide a list of endpoints.

Endpoint IP	Country	First Seen	Last Seen
109.235.193.83	Azerbaijan	6/2/2013	11/26/2013
190.242.96.49	Colombia	10/21/2013	1/7/2014
41.33.151.150	Egypt	3/10/2013	10/29/2013
216.118.232.xxx	Ethiopia	11/18/2013	2/3/2014
81.183.229.xxx	Hungary	6/16/2012	Active
2.228.65.226	Italy	10/26/2012	Active

82.104.200.51	Italy	9/17/2012	12/2/2013
88.33.54.xxx	Italy	6/4/2012	Active
95.228.202.xxx	Italy	9/18/2012	Active
95.228.202.xxx	Italy	9/17/2012	Active
95.228.202.xxx	Italy	9/18/2012	Active
95.228.202.xxx	Italy	9/18/2012	Active
95.228.202.xxx	Italy	9/17/2012	Active
95.228.202.xxx	Italy	9/15/2012	Active
89.218.88.xxx	Kazakhstan	8/21/2013	Active
211.51.14.129	Korea	8/26/2012	1/7/2014
203.217.178.xxx	Malaysia	5/28/2012	Active
189.177.47.xxx	Mexico	1/30/2014	Active
189.177.65.13	Mexico	11/13/2013	12/10/2013
189.177.74.147	Mexico	11/1/2013	11/1/2013
201.157.43.60	Mexico	10/13/2013	1/7/2014
200.67.230.2	Mexico	5/25/2012	Active
41.248.248.xxx	Morocco	6/3/2012	Active
41.248.248.xxx	Morocco	7/25/2012	Active
41.248.248.xxx	Morocco	6/12/2012	Active
41.248.248.xxx	Morocco	5/27/2012	Active
81.192.5.xxx	Morocco	7/25/2012	Active
62.251.188.xxx	Morocco	5/31/2012	Active
197.210.255.178	Nigeria	9/15/2013	10/21/2013

95.49.xxx.xxx[53]	Poland	8/10/2012	Active
37.242.13.10	Saudi Arabia	1/7/2014	1/7/2014
62.149.88.20	Saudi Arabia	6/5/2012	7/2/2013
41.78.109.91	Sudan	12/14/2012	1/12/2014
203.149.47.xxx	Thailand	10/4/2013	Active
95.9.71.180	Turkey	11/13/2013	11/19/2013
81.95.226.134	Uzbekistan	8/7/2013	9/2/2013
81.95.224.10	Uzbekistan	1/22/2013	1/26/2013
217.29.123.184	Uzbekistan	7/21/2013	9/16/2013

We illustrate the process of linking servers together, identifying endpoints, and mapping circuits, with the example of a circuit that forwards traffic to Mexico.

The Mexico Circuit

First, we noted that the following group of servers all returned exactly identical SSL certificates:

IP	Provider	Country	First Seen
14.136.236.xxx	38Cloud	HK	2013-04-13
31.192.228.xxx	GleSYS	NL	2013-04-11
176.58.102.xxx	Linode	UK	2013-02-09
50.116.32.xxx	Linode	US	2013-08-04

In general, we also noted that a number of Hacking Team servers did not return SSL certificates, but instead used a global sequential IPID (i.e., these servers had properties akin to “zombie hosts” in the terminology of idle scanning).[54]

The IPID is a field in every IP packet that is used to identify fragments of a packet. As a packet flows across the internet, it may be fragmented into smaller packets if it is too big to traverse a link between its source and destination all at once. Each packet fragment has the same IPID as the original packet, allowing the fragments to be reassembled into the original packet at the destination. The sender generally does not know which packets will be fragmented *a priori*, so must assign an IPID to almost every packet. Many older operating

systems assign IPIDs to every packet using a global sequential method: the first packet sent to any destination has IPID 0, followed by 1, 2, and so on. Modern OS versions assign IPIDs randomly.

For our purposes, if a server has a global IPID, then we can use it as a counter for the number of packets that the server has sent to anyone. Furthermore, anyone can probe the server for this value by sending a request (e.g., TCP SYN[55]) to the server, and looking at the IPID value in the response (e.g., SYN/ACK). By probing the IPID value twice, once at time t1 and once at t2, one can see if the server sent any packets between t1 and t2.

The RCS servers we found with global IPIDs appeared to be located in countries that we saw less frequently in our scan results, whereas servers with SSL certificates appeared to be in countries that we saw more frequently. Based on this observation, we hypothesized that the servers returning SSL certificates were forwarding traffic to the servers with global IPIDs.

In order to establish whether a server, X, forwards traffic to another server with a global IPID, Y, we first waited until Y was idle (i.e., not sending any packets, as measured by repeated IPID probes). We then sent an IPID probe to Y, followed by a request (HTTP HEAD[56]) to X, followed by another IPID probe to Y. We looked for a gap (a difference greater than one) between the first and second IPID values; a gap would imply that Y responded to, and thus processed and received, our request to X. If this is the case, we say that X is a proxy for Y. We repeated the experiment more than ten times to be sure, checking in each case that Y was idle. We discovered that the four servers mentioned above were proxies for the following server:

IP	Provider	Country	First Seen
200.67.230.xxx	UniNet	MX	2012-05-25

The next natural question was: is this server the endpoint? Or is it a proxy for another server?

We noticed that when we sent an HTTP HEAD or HTTP GET request to the MX server, the IPIDs of all IP packets returned by the MX server were consecutive. In other words, the MX server could not have been communicating with any other server by creating new packets, in responding to our request.[57] However, MX could have been forwarding all *existing* packets of our request by rewriting destination addresses (like Network Address Translation).[58] This would not induce the creation of any new packets, so it would be consistent with observed consecutive IPIDs. However, this type of forwarding would still be measurable in latency (round trip time) differences between the server in question and neighbouring servers not related to the spyware. In order to determine whether this was the case, we compared the latency of the MX server (measured using hping in both TCP and ICMP modes) with neighbouring servers in the IP space. If the latency of the MX server was higher than neighbouring servers, it could indicate that the MX server was a proxy as opposed to an endpoint. We found the latency to be identical to neighbouring servers. Thus, based on these two tests, we concluded that MX is an endpoint. To summarize, we used the following test to determine whether a server was an endpoint:

1. Determine whether the server has a global IPID. If not, the test will not work.
2. Wait for the server to be idle. If the server is not idle, the test will not work.
3. Issue an HTTP GET request to the server, and look at the IPIDs of all packets returned by the server in response (we used tcpdump). If the IPIDs returned are not consecutive, then the server may not be an endpoint.

4. Compare the latency of the server (ICMP and TCP RTTs) with surrounding servers. If there is a difference, then the server may not be an endpoint.

Recall that so far, we have established that all four servers are proxies for MX, and MX is an endpoint. The next natural question was about the topology of the forwarding. Is each server directly forwarding traffic to MX, or are the servers arranged in a proxy chain where one server forwards to the next server, etc., and eventually one of the servers forwards to MX? In order to determine this, we computed the circuit latency, an observer-independent measure of the latency required for a given server to relay a request to MX and receive a response.[59]

IP	Country	Circuit Latency
14.136.236.147	HK	1584.51ms
50.116.32.138	US	1115.06ms
31.192.228.60	NL	865.25ms
176.58.102.218	UK	834.25ms
200.67.230.2	MX	65.05ms

While the latencies support the hypothesis that the servers are arranged in a proxy chain, it is also possible that we failed to detect other servers that are part of this group. In this case, the topology may be very different than what we have hypothesized.

RCS SAMPLES

Along with the scanning and fingerprinting of RCS servers previously described, we have also been able to identify likely government users by analyzing "bait content" associated with Windows and OS X RCS samples found in the wild.

Italy

Italy, the home country of Hacking Team, is one of the most prolific users of RCS. We identified several endpoints operated by Telecom Italia, but also Fastweb Network Administration Staff and C.S.H. & M.P.S. S.r.l. Interestingly, C.S.H. & M.P.S. S.r.l. was apparently involved in 2011 in the surveillance and arrest of Luigi Bisignani, an Italian ex-journalist allegedly deeply involved in secretly influencing Italian politics.[60] We also observed numerous RCS samples (see Appendix B) that can be attributed to Italy, for which the context and targets are currently unknown. One of the most recent cases involved a document called Biglietto Visita.doc (Italian for "business card") embedding a Flash 0-day exploit (see Exploit 5 in the Target Exploitation section above). The document was uploaded from Italy along with Windows and OSX versions of the RCS spyware.

Oman

We identified an RCS sample uploaded to VirusTotal from Oman[61] that contained a bait document about Omani poetry, purportedly authored by Dr. Mohammed Mahrooqi at the University of Nizwa in Oman.

Panama

We identified a server in Italy, 93.95.219.97, which was associated with two domain names, noticiaspty.com and blackberry-upgrade.com we believe were used for attacks in Panama. According to our scans, the IP was apparently last active around 9 September 2013.[62] A Google search for the noticiaspty.com domain name reveals links to ".jad" files (BlackBerry applications) sent publicly via Twitter.

Two links to BlackBerry applications from @teresa_var.[63,64]

We also searched Topsy and found a deleted tweet that also links to a .jad file on this website.[65] The Tweet contains a link to a BlackBerry application that is purportedly a leaked telephone recording of prominent lawyer Zulay Rodriguez.

The intended targets are as yet unclear, however almost all the Twitter users followed by and following @teresa_var are Panamanian politicians, many belonging to the *Democratic Revolutionary Party* (PRD).[66] The tweets were all sent in December 2011.

Other Samples

We found many other samples of RCS in the wild. Using the results of our mapping of RCS proxy chains, we managed to attribute samples to countries including: **Egypt, Hungary, Italy, Kazakhstan, South Korea, Morocco, Saudi Arabia, Turkey** and **Uzbekistan**. See **Appendix B** for sample hashes.

CONCLUSION

Hacking Team has made a number of statements that seem intended to reassure the public, as well as potential regulators, that they conduct effective due diligence and self-regulation regarding their clients, and the human rights impact of their products. They also market their RCS product as untraceable. Our research suggests that both of these claims ring hollow.

Hacking Team says:

We have established an outside panel of technical experts and legal advisors, unique in our industry, that reviews potential sales. This panel reports directly to the board of directors regarding proposed sales.

First, with respect to human rights, we have encountered a number of cases where bait content and other material are suggestive of targeting for political advantage, rather than legitimate law enforcement operations. Moreover, in an earlier post in this series, we identified the targeting of a US-based news organization. In

other cases, however, the material did appear to be indicative of possible criminal investigations.[67] Similarly, we have also found Hacking Team endpoints in regimes with both high and very low rankings in governance, rule of law, and freedom of expression.

We find it reasonable to conclude that some uses of this hacking suite are legally-sanctioned, properly overseen criminal investigations conducted under due process and rule-of-law. It is equally reasonable, however, to conclude that some uses are abusive, partisan, or unaccountable. **Our findings of the global proliferation of Hacking Team belies their claims of high-quality due diligence.** While they claim to rely on an outside panel for guidance on potential sales, little information is available about its members, processes, or the grounds under which a sale might be rejected.

With respect to its secrecy; despite marketing by Hacking Team that promises secrecy, as well as clear efforts at obfuscation, we have also shown once more that commercial trojan manufacturers may be guilty of hype when they make claims like "stealth" and "untraceable." We think it unlikely that the purchasing officers of client countries were warned that their use of these tools could be exposed by non-nation-state researchers.

Our report also highlights a professional alignment between exploit sellers and companies that sell surveillance trojans. While these actors are natural business partners, the conclusion we draw is that the marketplace for exploits and surveillance software, despite the opacity and competitiveness for government contracts, can also be cozy, with vendors regularly working together to sell products and solutions to clients. While this collaboration may offer a one-stop-shop experience for purchasers, it also helps tie vendors, campaigns, companies, and countries back together when investigated.

In conclusion, the combination of global proliferation, as well as dubious promises about "stealth" features points to the dangers—to many stakeholders—of an unregulated marketplace defined by lack of transparency and accountability.

ACKNOWLEDGEMENTS

Thanks to Bernhard Amann, Collin D Anderson, Zakir Durumeric, Drew Hintz, Ralph Holz, Shane Huntley, Andrew Lyons, Vern Paxson, Mark Schloesser, and Nicholas Weaver.

FOOTNOTES

[1] http://hackingteam.it/index.php/customer-policy
[2] http://slate.me/1eSTeUF
[3] https://citizenlab.org/2012/10/backdoors-are-forever-hacking-team-and-the-targeting-of-dissent/
[4] https://citizenlab.org/2014/02/hacking-team-targeting-ethiopian-journalists/
[5] The countries are Azerbaijan, Ethiopia, Kazakhstan, Nigeria, Oman, Saudi Arabia, Sudan, UAE, and Uzbekistan. See https://www.eiu.com/public/topical_report.aspx?campaignid=DemocracyIndex12
[6] http://hackingteam.it/index.php/about-us
[7] Also referred to as DaVinci or Galileo.
[8] http://hackingteam.it/images/stories/galileo.pdf
[9] http://hackingteam.it/components/com_gk3_photoslide/thumbs_big/59191101.jpg
[10] https://www.securelist.com/en/analysis/204792290/Spyware_HackingTeam
[11] http://www.theverge.com/2013/9/13/4723610/meet-hacking-team-the-company-that-helps-police-hack-into-computers
[12] ibid.

[13] http://www.corpwatch.org/article.php?id=15868
[14] Except that the infrastructure for RCS does not appear to be shared; each government user has its own, fixed proxy chain(s), with no infrastructure shared between different governments.
[15] http://news.cnet.com/8301-13578_3-57573707-38/meet-the-corporate-enemies-of-the-internet-for-2013/
[16] http://www.eluniverso.com/noticias/2013/12/11/nota/1901271/firma-hacking-team-fue-contactada-estado-ecuatoriano
[17] http://www.ibtimes.co.uk/hacking-team-murky-world-state-sponsored-spying-445507
[18] http://hackingteam.it/index.php/customer-policy
[19] We identified Ethiopia as a likely government user based on the targeting of Ethiopian journalists in Washington DC. See https://citizenlab.org/2014/02/hacking-team-targeting-ethiopian-journalists/.
[20] We identified Oman as a likely government user based on the apparent targeting of one or more Omani writers, as described later.
[21] In addition to finding an endpoint in Panama, we also found indications that opposition Panamanian politicians were targeted, as described later.
[22] We identified UAE as a likely government user based on the targeting of UAE activist Ahmed Mansoor; the attack used a server registered to Royal Group UAE, which is chaired by the son of the founder of the country.
[23] http://www.washingtonpost.com/blogs/worldviews/wp/2013/10/09/oops-azerbaijan-released-election-results-before-voting-had-even-started/
[24] http://www.washingtonpost.com/blogs/worldviews/wp/2013/08/07/intimate-video-emerges-again-of-reporter-investigating-azerbaijan-presidents-family/
[25] http://www.hrw.org/world-report/2014/country-chapters/kazakhstan?page=1
[26] http://www.hrw.org/world-report/2014/country-chapters/uzbekistan?page=2
[27] Al-Khomasia (http://www.alkhomasia.biz/site/) is a Saudi maritime shipping company, part of Abdulkader Al-Bakri's eponymous group (http://www.albakri.biz/site/). A domain name registered to the Al-Bakri Group, emailsrvc.com, pointed to the same IP and was consistent with one of our fingerprints for an RCS server as early as 12/22/2010.
[28] http://www.hrw.org/world-report/2014/country-chapters/saudi-arabia
[29] http://www.hrw.org/world-report/2014/country-chapters/121334
[30] https://citizenlab.org/2012/10/backdoors-are-forever-hacking-team-and-the-targeting-of-dissent/
[31] http://www.lemonde.fr/technologies/article/2013/02/19/hackers-d-etat_1834943_651865.html
[32] http://www.vupen.com/english/company.php
[33] http://www.zdnet.com/nsa-purchased-zero-day-exploits-from-french-security-firm-vupen-7000020825/
[34] https://www.documentcloud.org/documents/810501-769-gamma-group-product-list-finfisher.html
[35] http://www.issworldtraining.com/iss_wash/track4.html
[36] http://slate.me/1eSTeUF
[37] 53cd1d6a1cc64d4e8275a22216492b76db186cfb38cec6e7b3cfb7a87ccb3524
[38] b5462a2be69d268a7d581fe9ee36e8f31d5e1362d01626e275e8f58029e15683
[39] 277cae7c249cb22ae43a605fbe901a0dc03f11e006b02d53426a6d11ad241a74
[40] The translation is "Kofahi Letter;" the letter was purportedly from Nabil al-Kofahi, a reformist Muslim Brotherhood official in Jordan.
[41] http://seclists.org/bugtraq/2012/Sep/46
[42] Exploit 2 is an RTF and does not contain Word document metadata.
[43] We keep relevant fingerprints secret, out of an abundance of caution to prevent disruption to ongoing legitimate criminal investigations that may involve the use of Hacking Team spyware.
[44] FSBSpy is a *dropper*, a piece of malware with limited functionality that is able to install more malware. It has long been suspected that FSBSpy is a first stage dropper used to verify an infection before, possibly manually, issuing the deployment of the final and more full-featured RCS, with which it shares multiple portions of code. Our scan results confirm a strong link between FSBSpy and RCS — we observed no differences between these servers, and even identified server groups used by both RCS samples and FSBSpy

samples. Thus, in this post, we treat FSBSpy as an RCS product.
[45] We also contacted a team at TU Munich (https://pki.net.in.tum.de/), who searched their private scanning results for us; an interested reader could reproduce all of our server findings without access to this data.
[46] We searched the following service probes: 80-TCP_crossdomainxml, 80-TCP_FourOhFourRequest, 80-TCP_GetRequest, 80-TCP-HTTPOptions, 80-TCP-RTSPRequest, 443-TCP-SSL23SessionReq, 443-TCP-SSLSessionReq. Data is available at http://internetcensus2012.bitbucket.org/.
[47] https://scans.io/study/sonar.cio
[48] We searched the IPv4 SSL certificate scans from 9/10/2013 to 2/10/2014 (data available at https://scans.io/study/sonar.http) and the HTTP GET Port 80 scans from 10/29/2013 to 01/30/2014 (data available at https://scans.io/study/sonar.ssl.
[49] http://www.shodanhq.com/
[50] https://scans.io/study/umich-https
[51] In words, A2 checks for a response that contains a distinctive header, and also redirects to another website. The most common redirect we observed was http://www.google.com/, but we also identified other redirects, including two redirects to http://www.blackberry.com/btsc/kb18327, a page describing vulnerabilities in the BlackBerry PDF distiller (CVE-2009-2643), which could allow an attacker to execute arbitrary code via a malicious PDF file.
[52] The typo was "HTTP1/1 400 Bad request." The typo was corrected in version 0.1.7.
[53] We observed too many IPs in this range to list here. It appears that the endpoint in this case has a dynamic IP address. All IP addresses that we identified in Poland appear to be endpoints of the same proxy chain.
[54] http://en.wikipedia.org/wiki/Idle_scan
[55] We sent TCP SYNs using the *hping* utility.
[56] We sent HTTP requests using the *httping* utility. We used HTTP requests to induce traffic, because Hacking Team's RCS uses HTTP as its communication protocol.
[57] Also, the consecutive IPIDs told us that all packets were communicating with a single IP protocol stack (i.e., it was not the case that some packets were being handled by one server, and some were being forwarded to another server). In particular, it was not the case that one server accepted the TCP connection, and then forwarded packets on the established TCP connection to another server.
[58] http://en.wikipedia.org/wiki/Network_address_translation
[59] We sent HTTP HEAD requests to each of the five servers (including MX) using the *httping* utility. In split mode (*-S*), *httping* returns a value for the latency representing the round trip time to the server, plus the time taken by the server to forward to MX, and receive a response from MX. We obtain the circuit latency by subtracting from this value the round trip time to the server. We measured circuit latencies from three separate datacenters (London, Atlanta, Dallas). We measured the minimum time, as well as the average round trip time, and subtracted the latter from the former. We took 20 measurements of each, from each datacenter. We picked the minimum difference from across the 3 datacenters for each of the five IPs.
[60] http://www.tomshw.it/cont/news/bisignani-usava-un-trojan-spia-per-intercettare/32498/1.html
[61] 209a986d8e17d361424dc11ffc69511b
[62] The Internet Archive recorded the contents of noticiaspty.com on 14 June 2013, which is consistent with one of our scanning fingerprints: https://web.archive.org/web/20130614133421/http://noticiaspty.com/.
[63] Available on Twitter and WebCitation.
[64] Available on Twitter and WebCitation.
[65] http://topsy.com/trackback?url=http%3A%2F%2Fwww.noticiaspty.com%2Fzulay_rodriguez.jad
[66] http://en.wikipedia.org/wiki/Democratic_Revolutionary_Party
[67] For example, we found two samples that communicated with a Saudi Arabian endpoint that had financially-themed bait content, suggesting possible anti-money laundering operations.