# Account Hijacking and the Underground Market

CS114/214 Selected Readings of CS Research

Zakir Durumeric

Stanford University

# Next Week's Lecture

***SocialHEISTing: Understanding Stolen Facebook Accounts***
Jeremiah Onaolapo et al.
USENIX Security Symposium, August 2021

The paper considers what happens when social media accounts are hijacked… especially…

What effects do account demographics have on how miscreants abuse accounts?

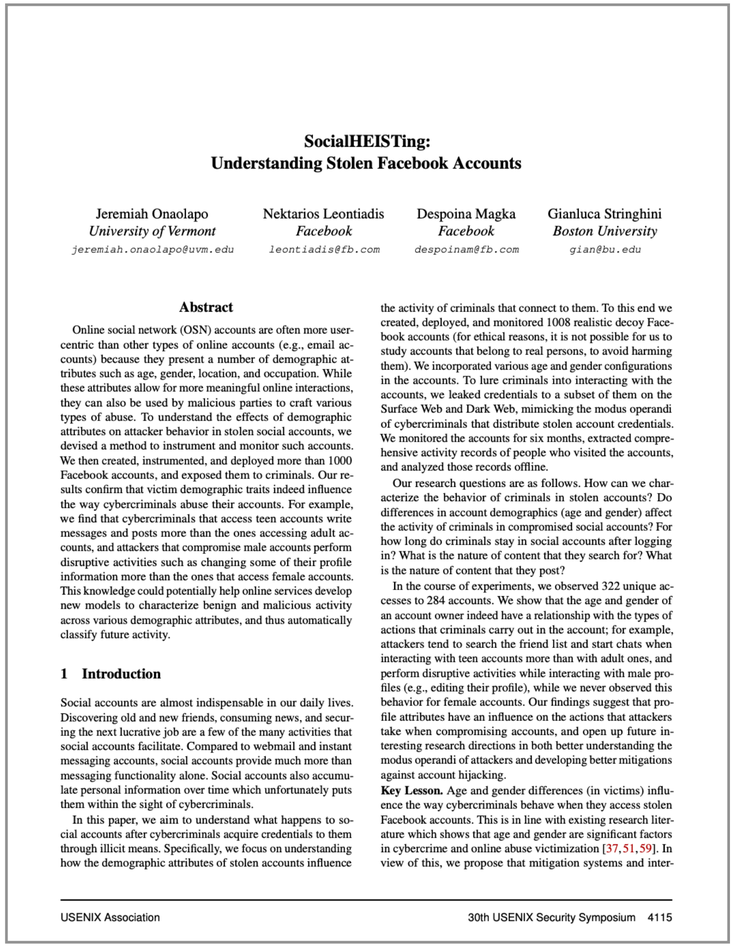

**SocialHEISTing:**
**Understanding Stolen Facebook Accounts**


Jeremiah Onaolapo
*University of Vermont*
jeremiah.onaolapo@uvm.edu

Nektarios Leontiadis
*Facebook*
leontiadis@fb.com

Despoina Magka
*Facebook*
despoinam@fb.com

Gianluca Stringhini
*Boston University*
gian@bu.edu



**Abstract**

Online social network (OSN) accounts are often more user-centric than other types of online accounts (e.g., email accounts) because they present a number of demographic attributes such as age, gender, location, and occupation. While these attributes allow for more meaningful online interactions, they can also be used by malicious parties to craft various types of abuse. To understand the effects of demographic attributes on attacker behavior in stolen social accounts, we devised a method to instrument and monitor such accounts. We then created, instrumented, and deployed more than 1000 Facebook accounts, and exposed them to criminals. Our results confirm that victim demographic traits indeed influence the way cybercriminals abuse their accounts. For example, we find that cybercriminals that access teen accounts write messages and posts more than the ones accessing adult accounts, and attackers that compromise male accounts perform disruptive activities such as changing some of their profile information more than the ones that access female accounts. This knowledge could potentially help online services develop new models to characterize benign and malicious activity across various demographic attributes, and thus automatically classify future activity.


**1 Introduction**

Social accounts are almost indispensable in our daily lives. Discovering old and new friends, consuming news, and securing the next lucrative job are a few of the many activities that social accounts facilitate. Compared to webmail and instant messaging accounts, social accounts provide much more than messaging functionality alone. Social accounts also accumulate personal information over time which unfortunately puts them within the sight of cybercriminals.

In this paper, we aim to understand what happens to social accounts after cybercriminals acquire credentials to them through illicit means. Specifically, we focus on understanding how the demographic attributes of stolen accounts influence the activity of criminals that connect to them. To this end we created, deployed, and monitored 1008 realistic decoy Facebook accounts (for ethical reasons, it is not possible for us to study accounts that belong to real persons, to avoid harming them). We incorporated various age and gender configurations in the accounts. To lure criminals into interacting with the accounts, we leaked credentials to a subset of them on the Surface Web and Dark Web, mimicking the modus operandi of cybercriminals that distribute stolen account credentials. We monitored the accounts for six months, extracted comprehensive activity records of people who visited the accounts, and analyzed those records offline.

Our research questions are as follows. How can we characterize the behavior of criminals in stolen accounts? Do differences in account demographics (age and gender) affect the activity of criminals in compromised social accounts? For how long do criminals stay in social accounts after logging in? What is the nature of content that they search for? What is the nature of content that they post?

In the course of experiments, we observed 322 unique accesses to 284 accounts. We show that the age and gender of an account owner indeed have a relationship with the types of actions that criminals carry out in the account; for example, attackers tend to search the friend list and start chats when interacting with teen accounts more than with adult ones, and perform disruptive activities while interacting with male profiles (e.g., editing their profile), while we never observed this behavior for female accounts. Our findings suggest that profile attributes have an influence on the actions that attackers take when compromising accounts, and open up future interesting research directions in both better understanding the modus operandi of attackers and developing better mitigations against account hijacking.

**Key Lesson.** Age and gender differences (in victims) influence the way cybercriminals behave when they access stolen Facebook accounts. This is in line with existing research literature which shows that age and gender are significant factors in cybercrime and online abuse victimization [37, 51, 59]. In view of this, we propose that mitigation systems and inter-

# Today:

1. How are user accounts hijacked?
2. What does the underground marketplace for stolen credentials look like?

# Primarily Based On…

***Data Breaches, Phishing, or Malware?***
***Understanding the Risks of Stolen Credentials***

Kurt Thomas et al.
ACM Computer and Comm. Security, Nov. 2017

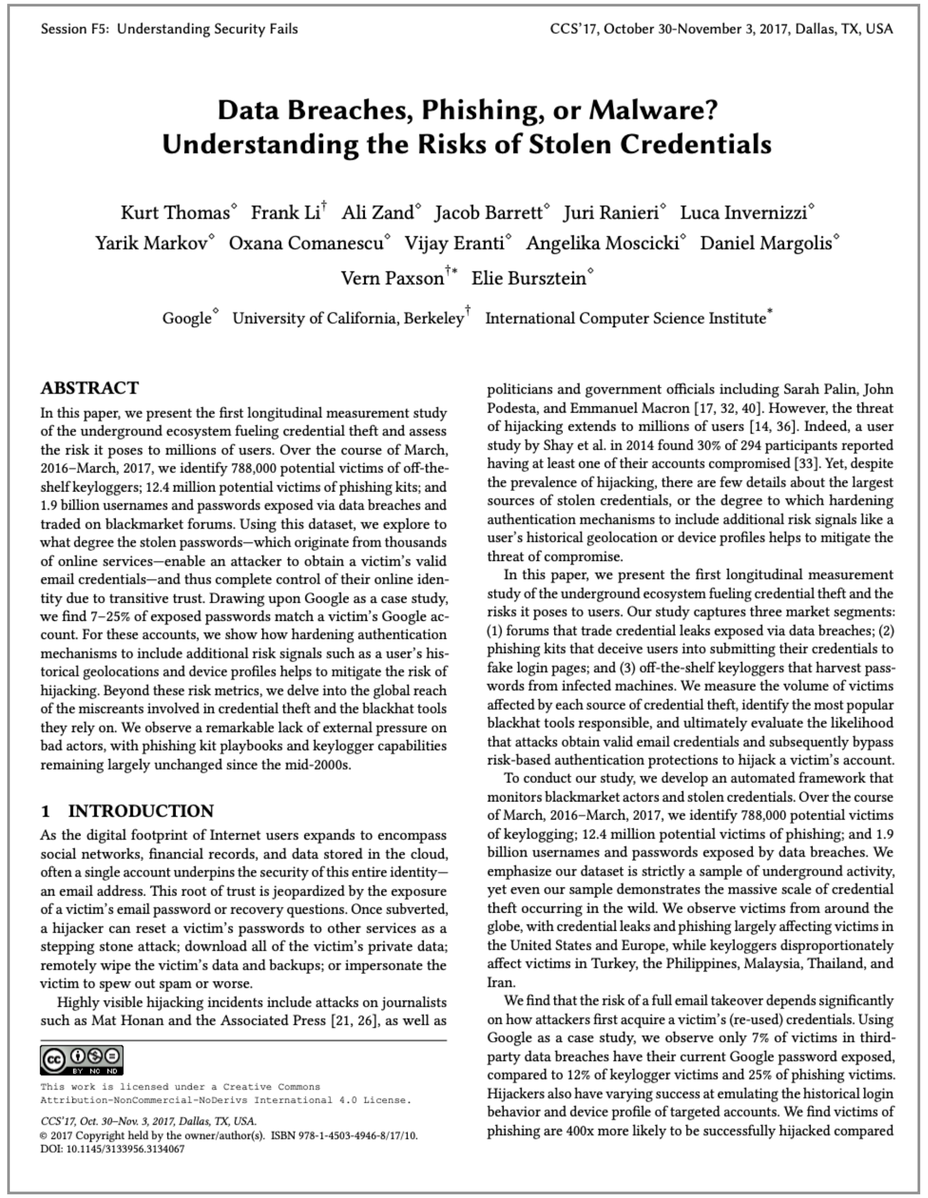



## Data Breaches, Phishing, or Malware? Understanding the Risks of Stolen Credentials

Kurt Thomas[◇] Frank Li[†] Ali Zand[◇] Jacob Barrett[◇] Juri Ranieri[◇] Luca Invernizzi[◇]
Yarik Markov[◇] Oxana Comanescu[◇] Vijay Eranti[◇] Angelika Moscicki[◇] Daniel Margolis[◇]
Vern Paxson[†*] Elie Bursztein[◇]

Google[◇] University of California, Berkeley[†] International Computer Science Institute[*]

**ABSTRACT**

In this paper, we present the first longitudinal measurement study of the underground ecosystem fueling credential theft and assess the risk it poses to millions of users. Over the course of March, 2016–March, 2017, we identify 788,000 potential victims of off-the-shelf keyloggers; 12.4 million potential victims of phishing kits; and 1.9 billion usernames and passwords exposed via data breaches and traded on blackmarket forums. Using this dataset, we explore to what degree the stolen passwords—which originate from thousands of online services—enable an attacker to obtain a victim's valid email credentials—and thus complete control of their online identity due to transitive trust. Drawing upon Google as a case study, we find 7–25% of exposed passwords match a victim's Google account. For these accounts, we show how hardening authentication mechanisms to include additional risk signals such as a user's historical geolocations and device profiles helps to mitigate the risk of hijacking. Beyond these risk metrics, we delve into the global reach of the miscreants involved in credential theft and the blackhat tools they rely on. We observe a remarkable lack of external pressure on bad actors, with phishing kit playbooks and keylogger capabilities remaining largely unchanged since the mid-2000s.

### 1 INTRODUCTION

As the digital footprint of Internet users expands to encompass social networks, financial records, and data stored in the cloud, often a single account underpins the security of this entire identity—an email address. This root of trust is jeopardized by the exposure of a victim's email password or recovery questions. Once subverted, a hijacker can reset a victim's passwords to other services as a stepping stone attack; download all of the victim's private data; remotely wipe the victim's data and backups; or impersonate the victim to spew out spam or worse.

Highly visible hijacking incidents include attacks on journalists such as Mat Honan and the Associated Press [21, 26], as well as politicians and government officials including Sarah Palin, John Podesta, and Emmanuel Macron [17, 32, 40]. However, the threat of hijacking extends to millions of users [14, 36]. Indeed, a user study by Shay et al. in 2014 found 30% of 294 participants reported having at least one of their accounts compromised [33]. Yet, despite the prevalence of hijacking, there are few details about the largest sources of stolen credentials, or the degree to which hardening authentication mechanisms to include additional risk signals like a user's historical geolocation or device profiles helps to mitigate the threat of compromise.

In this paper, we present the first longitudinal measurement study of the underground ecosystem fueling credential theft and the risks it poses to users. Our study captures three market segments: (1) forums that trade credential leaks exposed via data breaches; (2) phishing kits that deceive users into submitting their credentials to fake login pages; and (3) off-the-shelf keyloggers that harvest passwords from infected machines. We measure the volume of victims affected by each source of credential theft, identify the most popular blackhat tools responsible, and ultimately evaluate the likelihood that attacks obtain valid email credentials and subsequently bypass risk-based authentication protections to hijack a victim's account.

To conduct our study, we develop an automated framework that monitors blackmarket actors and stolen credentials. Over the course of March, 2016–March, 2017, we identify 788,000 potential victims of keylogging; 12.4 million potential victims of phishing; and 1.9 billion usernames and passwords exposed by data breaches. We emphasize our dataset is strictly a sample of underground activity, yet even our sample demonstrates the massive scale of credential theft occurring in the wild. We observe victims from around the globe, with credential leaks and phishing largely affecting victims in the United States and Europe, while keyloggers disproportionately affect victims in Turkey, the Philippines, Malaysia, Thailand, and Iran.

We find that the risk of a full email takeover depends significantly on how attackers first acquire a victim's (re-used) credentials. Using Google as a case study, we observe only 7% of victims in third-party data breaches have their current Google password exposed, compared to 12% of keylogger victims and 25% of phishing victims. Hijackers also have varying success at emulating the historical login behavior and device profile of targeted accounts. We find victims of phishing are 400x more likely to be successfully hijacked compared



*CCS'17, Oct. 30–Nov. 3, 2017, Dallas, TX, USA.*
 ISBN 978-1-4503-4946-8/17/10.
DOI: 10.1145/3133956.3134067

# How are accounts hijacked in practice?

# How are accounts hijacked in practice?

Typically accounts are hijacked opportunistically en masse

1. Data Breaches (Credential Dumps)
2. Phishing Attacks
3. Key Loggers (installed via malware)
4. Brute Force Guessing
5. Network Eavesdropping (Monster-in-the-Middle Attacks)

# How are accounts hijacked in practice?

1. Data Breaches (Credential Dumps)
2. Phishing Attacks
3. Key Loggers (installed via malware)
4. ~~Brute Force Guessing~~
5. ~~Network Eavesdropping (Monster-in-the-Middle Attacks)~~

Common | Hack | Tech | Code | Game | Groups | Web | GFX | Market | Money

## Hacks, Exploits, and Various Discussions

| Forum | Threads/Posts | Last Post |
|---|---|---|
| **Beginner Hacking**<br>This is for the entry level hacker wishing to learn more about the art of h4(k5.<br>Moderated By: Mentors<br>E-Whoring<br>Worms, Malware, and Viruses<br>Private Investigation Methods and Anonymity | 209,249<br>1,594,985 | **How to upload pictures in...**<br>Today 10:53 AM<br>by Mr.DoomHead |
| **Advanced Hacking**<br>If you feel you're past the beginner stages and want to delve deeper into computer security, analysis, and internet exploits you should participate here.<br>Botnets, IRC Bots, and Zombies<br>Cryptography, Encryption, and Decryption<br>Pentesting and Forensics<br>Decompiling, Reverse Engineering, Disassembly, and Debugging | 106,136<br>880,047 | **Can I get help with Delux...**<br>Today 10:43 AM<br>by DeluxeSoftware |
| **Hacking Tools and Programs**<br>Since every hacker needs tools and programs please post your favorites here.<br>Keyloggers<br>Remote Administration Tools | 146,889<br>1,631,757 | **Hacking the hacker back.**<br>Today 10:48 AM<br>by QU⚡CKSILVER |
| **Website and Forum Hacking**<br>We get a lot of discussions here about how to hack a website or forum so this is the area for those threads.<br>Moderated By: Mentors<br>SQL Injection Attacks<br>Requests for Hacking | 131,578<br>852,132 | **$1000 - for hack Twitter ...**<br>Today 10:36 AM<br>by Garymckinnonhacks |
| **Hacking Tutorials**<br>If you have a hacking tutorial please post it here for consideration to be in the Premium Hacking Tutorials.<br>Moderated By: Mentors<br>Free Ebook Hacking Tutorials<br>Video Tutorials | 28,440<br>516,094 | **[TUT] How to scan and exp...**<br>Today 10:21 AM<br>by kamrulahsan06 |
| **Wifi WPA WEP Bluetooth 4G LTE Wireless Hacking**<br>For hacking wireless networks, wep/wpa encryption, sniffers, backtrack, setup, connection problems, aircrack and other wireless related discussions please join this forum.<br>Moderated By: Mentors | 16,008<br>142,790 | **Wifi Password Obtaining**<br>Today 10:26 AM<br>by bob79 |
| **Skype, IRC, ICQ, AIM, XMPP, and IM Hacks**<br>If you are into exploiting instant messaging systems then we have a perfect forum for you. Post here for instant messenger hacks and exploits. | 10,238<br>94,961 | **Kaiten IRC BOTNET**<br>Today 07:16 AM<br>by GoldEagle |
| **Embedded Systems, Electronics, Gadgets, and DIY**<br>For all your hacking of electronics like radios, circuitry, alarms, radar detectors, cameras, remote controls or robots you can participate in this forum. Great area for DIY projects. Topics include Raspberry Pi and Arduino. | 3,901<br>38,847 | **Need help in a HW/SW Proj...**<br>Today 04:09 AM<br>by MoNo |

**Table 2: Breakdown of where we source credential leaks.**

| Source | Candidate documents | Confirmed leaks | Credentials extracted |
|---|---:|---:|---:|
| Paste sites | 3,317 | 1,666 | 4,855,780 |
| Search index | 26,208 | 1,304 | 10,856,227 |
| Public forums | 1,921 | 557 | 107,343,690 |
| Private forums | – | 258 | 1,799,553,568 |

# How are accounts hijacked in practice?

Google looked at credentials that were leaked online in 2017

1. **Data Breaches:** 1.9 billion credentials
2. **Phishing Attacks:** 12.4 million credentials
3. **Malware + Key Loggers:** 788,000 credentials

Every week, miscreants collect more than 235K credentials from phishing attacks and 14.9K from key loggers

Data breaches, phishing, or malware? Understanding the risks of stolen credentials. Thomas et al.

# How do leaked credentials affect Google accounts?

- 7% of third-party data breach victims
- 12% of malware + keylogger victims
- 25% of phishing victims

have their active Google password exposed

# Credentials don't always equal account compromise

How likely are accounts to be compromised in practice?

- Phishing victims are 400x more likely to be hijacked
- Keylogger victims are 40x more likely
- Databreach victims are 10x more likely

But… **26% of American adults have been notified of a data breach and 30% of adults have had an account compromised**

# Data Breaches and Credential Leaks

# Credential Leaks

Credential Leaks occur when major companies are hacked and their user databases are stolen

Compromised many companies, including Yahoo, MySpace, LinkedIn, Adobe, and Dropbox

Sometimes raw passwords are stolen. Sometimes just password hashes

Initially sold on the black market, but often publicly surface

# (Incorrect) Password Hashing

F = one way function (known as a "cryptographic hash)

Company stores **F(password).** To check if a password is correct, they check if **F(candidate)** == stored hash.

Better than storing raw password, but still bad! Why?

# Salted Password Hashing

F = one way function (known as a "cryptographic hash")

**Store:**

```
n = rand()
F(password || n), n
```

**To check if password is correct, calculate:**

```
F(candidate || n) == stored_hash
```

**Also!** Chose a very slow F() — use something that an attacker can't rapidly compute offline (e.g., Bcrypt or PBKDF2)

# Salted Password Hashing

F = one way function (known as a "cryptographic hash function")

Company stores **F(password).** To check if a password is correct, they check if **F(candidate)** == stored hash.

# Why do leaks affect more than vulnerable site?

*Das et al.*: 43% of passwords re-used

*Wash et al.*: 31% of users reuse passwords

Google: 36% of leaked passwords inverted

**Table 4: Top 10 passwords across all plaintext leaks.**

| Rank | Top Passwords | Number of Credentials | Percent of Credentials |
|---|---|---|---|
| 1 | 123456 | 6,387,184 | 0.35% |
| 2 | password | 2,759,747 | 0.15% |
| 3 | 123456789 | 2,249,344 | 0.12% |
| 4 | abc123 | 985,709 | 0.10% |
| 5 | password1 | 888,836 | 0.05% |
| 6* | homelesspa | 855,477 | 0.05% |
| 7 | 111111 | 855,257 | 0.05% |
| 8 | qwerty | 829,835 | 0.05% |
| 9 | 12345678 | 828,848 | 0.05% |
| 10 | 1234567 | 740,464 | 0.04% |

# Credential Leaks

| Rank | Source | Number of credentials | Plaintext after inversion |
|---|---|---|---|
| 1 | Unknown $^{p}$ | 558,862,722 | 100.0% |
| 2 | MySpace $^{p}$ | 322,014,681 | 100.0% |
| 3 | Badoo | 125,322,081 | 33.0% |
| 4 | Adobe$^{\diamond}$ | 123,947,902 | 0.0% |
| 5 | LinkedIn | 112,322,695 | 85.6% |
| 6 | VK $^{p}$ | 76,865,954 | 99.6% |
| 7 | Tumblr$^{*}$ | 73,355,694 | 0.0% |
| 8 | Dropbox$^{\dagger}$ | 68,669,208 | 0.0% |
| 9 | Zoosk | 57,085,529 | 68.2% |
| 10 | IMesh$^{\ddagger}$ | 51,283,424 | 0.0% |
| 11 | LastFM | 41,631,844 | 85.4% |
| 12 | Fling $^{p}$ | 40,724,332 | 100.0% |
| 13 | Neopets $^{p}$ | 35,822,980 | 100.0% |
| 14 | Mate1 $^{p}$ | 27,383,966 | 100.0% |
| 15 | Unknown $^{p}$ | 26,351,372 | 99.8% |
| 16 | 000webhost $^{p}$ | 15,249,241 | 100.0% |
| 17 | Taobao $^{p}$ | 15,051,549 | 100.0% |
| 18 | NexusMods $^{p}$ | 6,759,631 | 100.0% |
| 19 | Unknown $^{p}$ | 5,728,163 | 99.7% |
| 20 | Unknown $^{p}$ | 4,901,088 | 100.0% |

# Phishing

# Phishing Attack

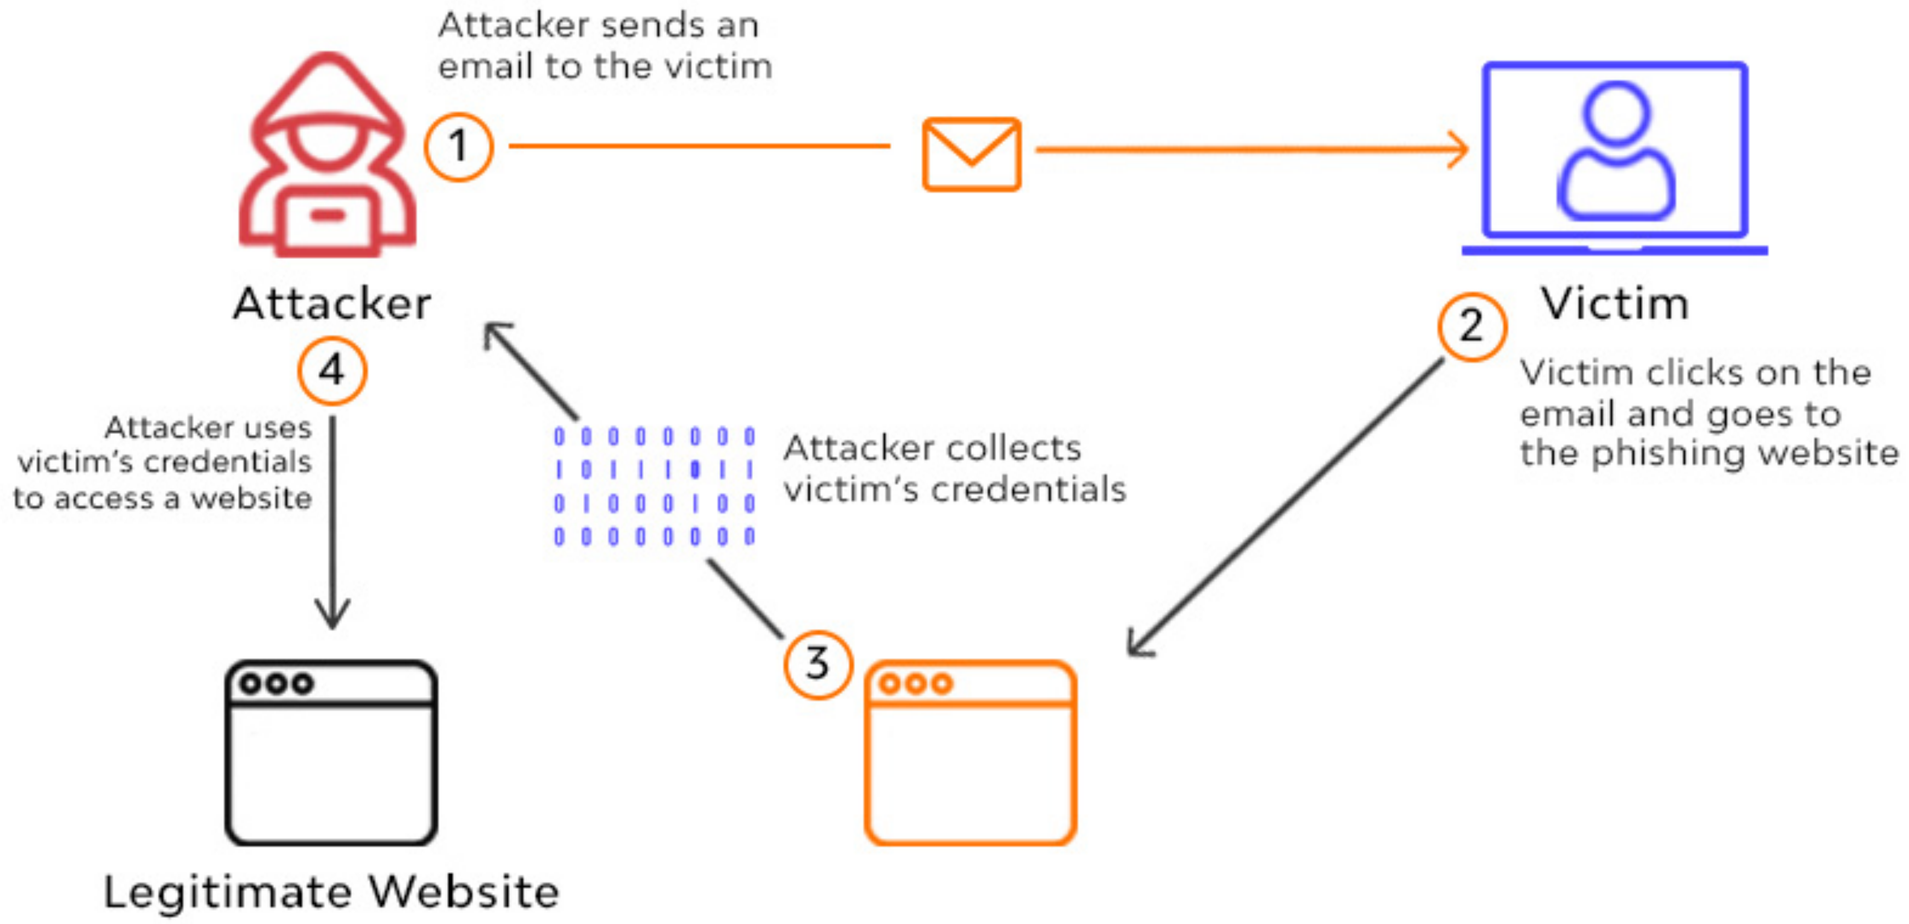

# Phishing Attack

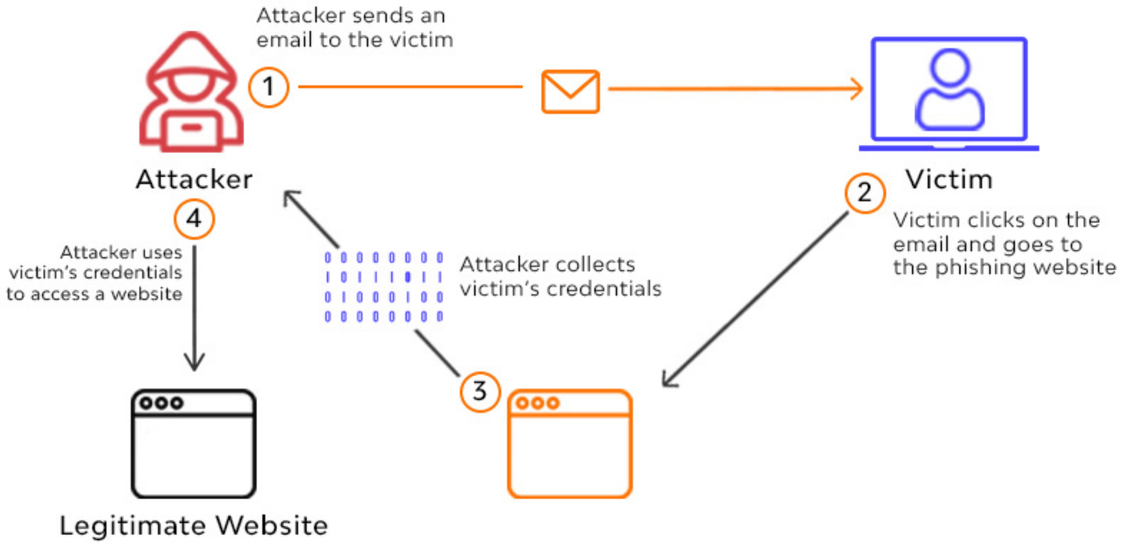


9% of victims who visit a phishing site will submit some amount of personal information

Different study: depending on email, up to 5-8% of recipients will click on link in a well crafted phishing message

A third: 7% of visitors submit info that visit phishing site

Phishing sites are capable of capturing more than just raw username and password

# Phishing Infrastructure

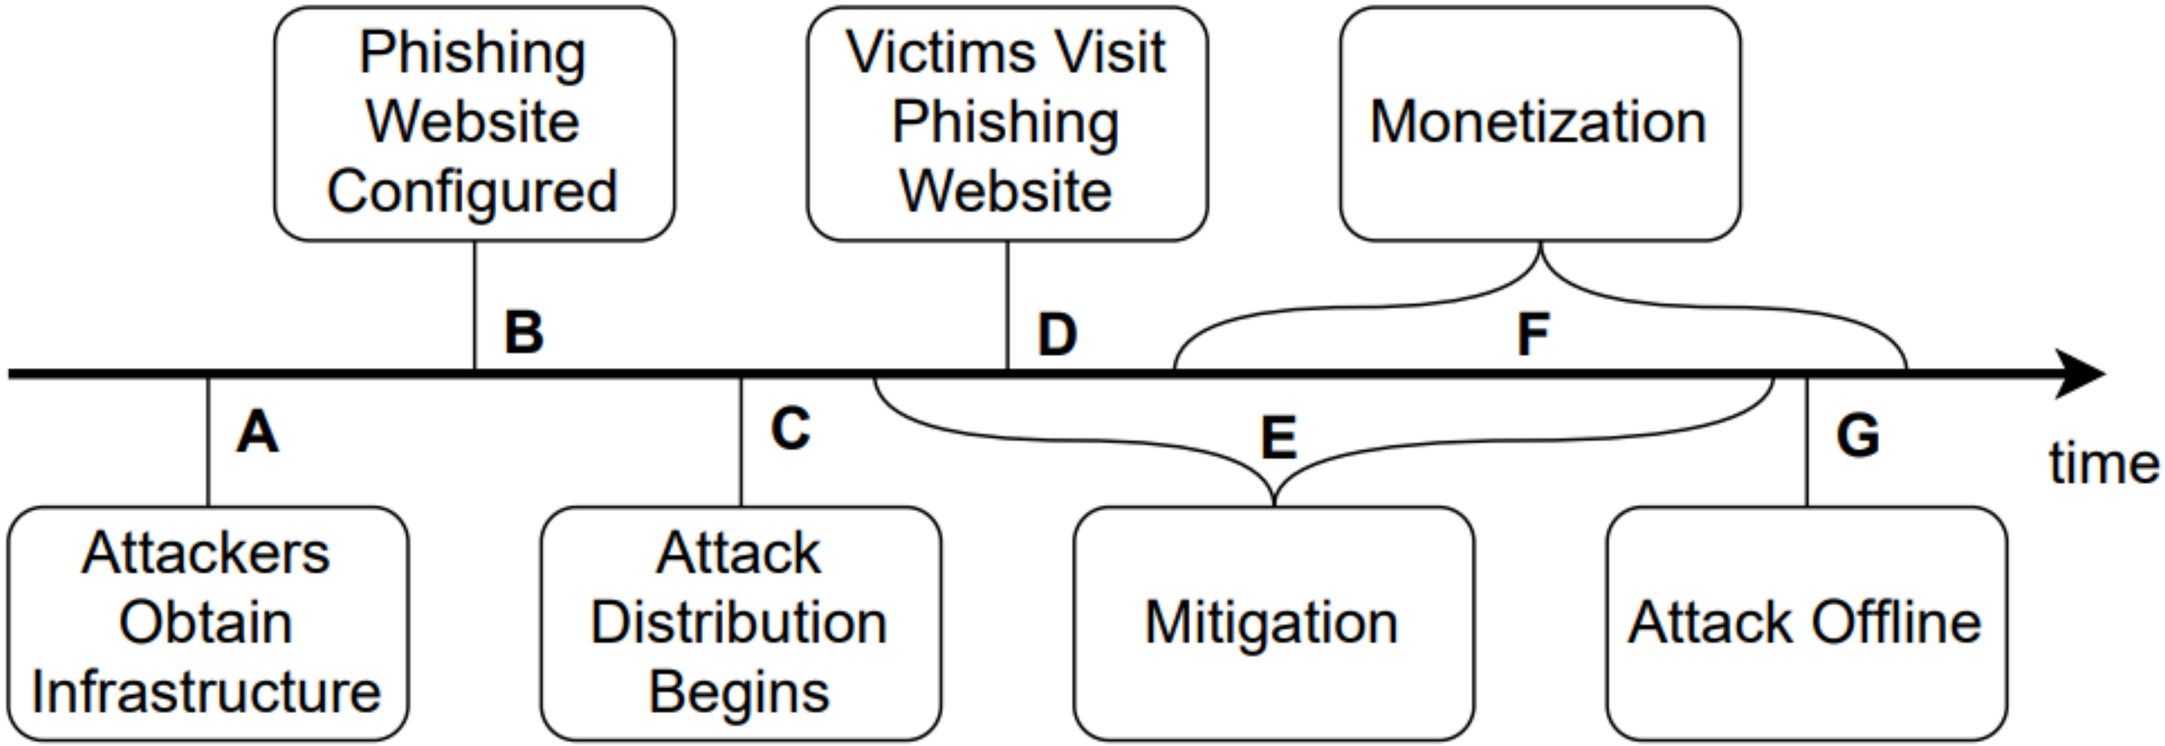

# Attacks are Frequent and Short Lived

Average phishing attack spans 21 hours between the first and last victim visit

Phishing sites are detected and start be blocked after 9 hours

40% of victims will visit sites after they've been detected

Long tail distribution — Top 10% largest attacks in our accounted for 89.13% of targeted victims

# Phishing Kits

Attackers sell “ready-to-deploy” packages for creating and configuring phishing content

Not just email — social media, direct message, SMS, WhatsApp

Forward stolen credentials to the operator in one of three ways: through SMTP to an email address controlled by the operator, via FTP, or by connecting to a remote database.

# What Do Phishing Kits Collect?

| Data type | Phishing kits | Keyloggers |
|---|---|---|
| Email | 81.4% | 97.8% |
| Password | 83.0% | 100% |
| Geolocation | 82.9% | 73.6% |
| Phone number | 18.1% | 0.1% |
| Device information | 16.2% | 67.9% |
| Secret questions | 7.4% | 0.1% |
| Full name | 45.8% | 85.3% |
| Credit card | 39.9% | 2.1% |
| SSN | 8.8% | 0.1% |

# Detection at Google

Commonly used phishing kits let Google uncover when stolen credentials are sent to a Gmail address — identifying 12M credentials sent to 19K drop locations over a one year period (likely a huge underestimate of total)

Phishing largely affecting victims in the U.S. and Europe — number is increasing as time goes on

# Malware and Keyloggers

# Popular Keyloggers — HawkEye and Predator Pain

Similar to Phishing Kits, attackers have created easy-to-deploy key logging software

Around 40% of variants collect credit card information and 8.8% of variants collect social security numbers

Google saw around 14,879 credentials per week from key loggers — but likely a huge underestimate

# Exploit as a Service

Today, host compromise is decoupled from host monetization

World has evolved into a *pay-per-install* model of malware distribution, where miscreants pay for compromised hosts via the underground economy

Typically machines are compromised via driveby downloads that target browser and plugin vulnerabilities (e.g., PDF viewers, Flash, and Java)

# Torpig

In 2009, a team of researchers seized control of Torpig Botnet

Victims are infected through drive-by-download attacks

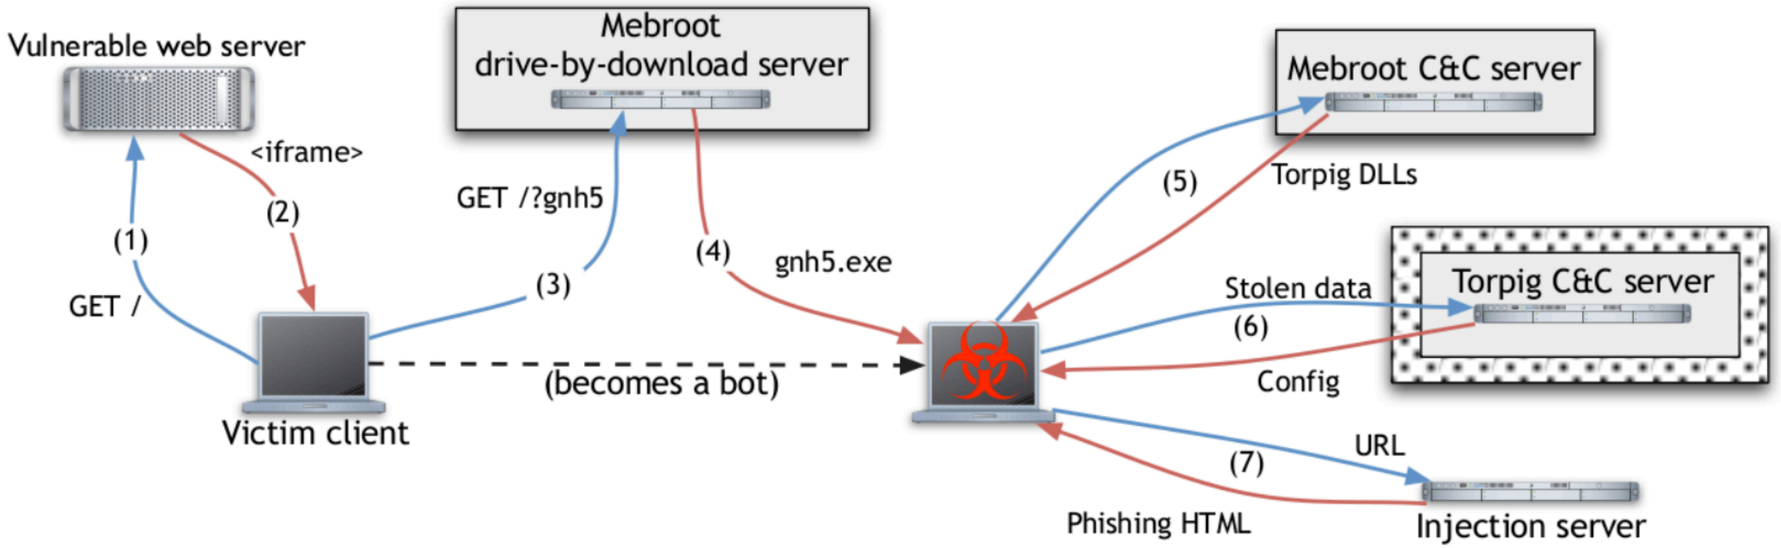

# Torpig

Injects malicious Javascript into visited websites to steal credentials

Sends credentials to a central server

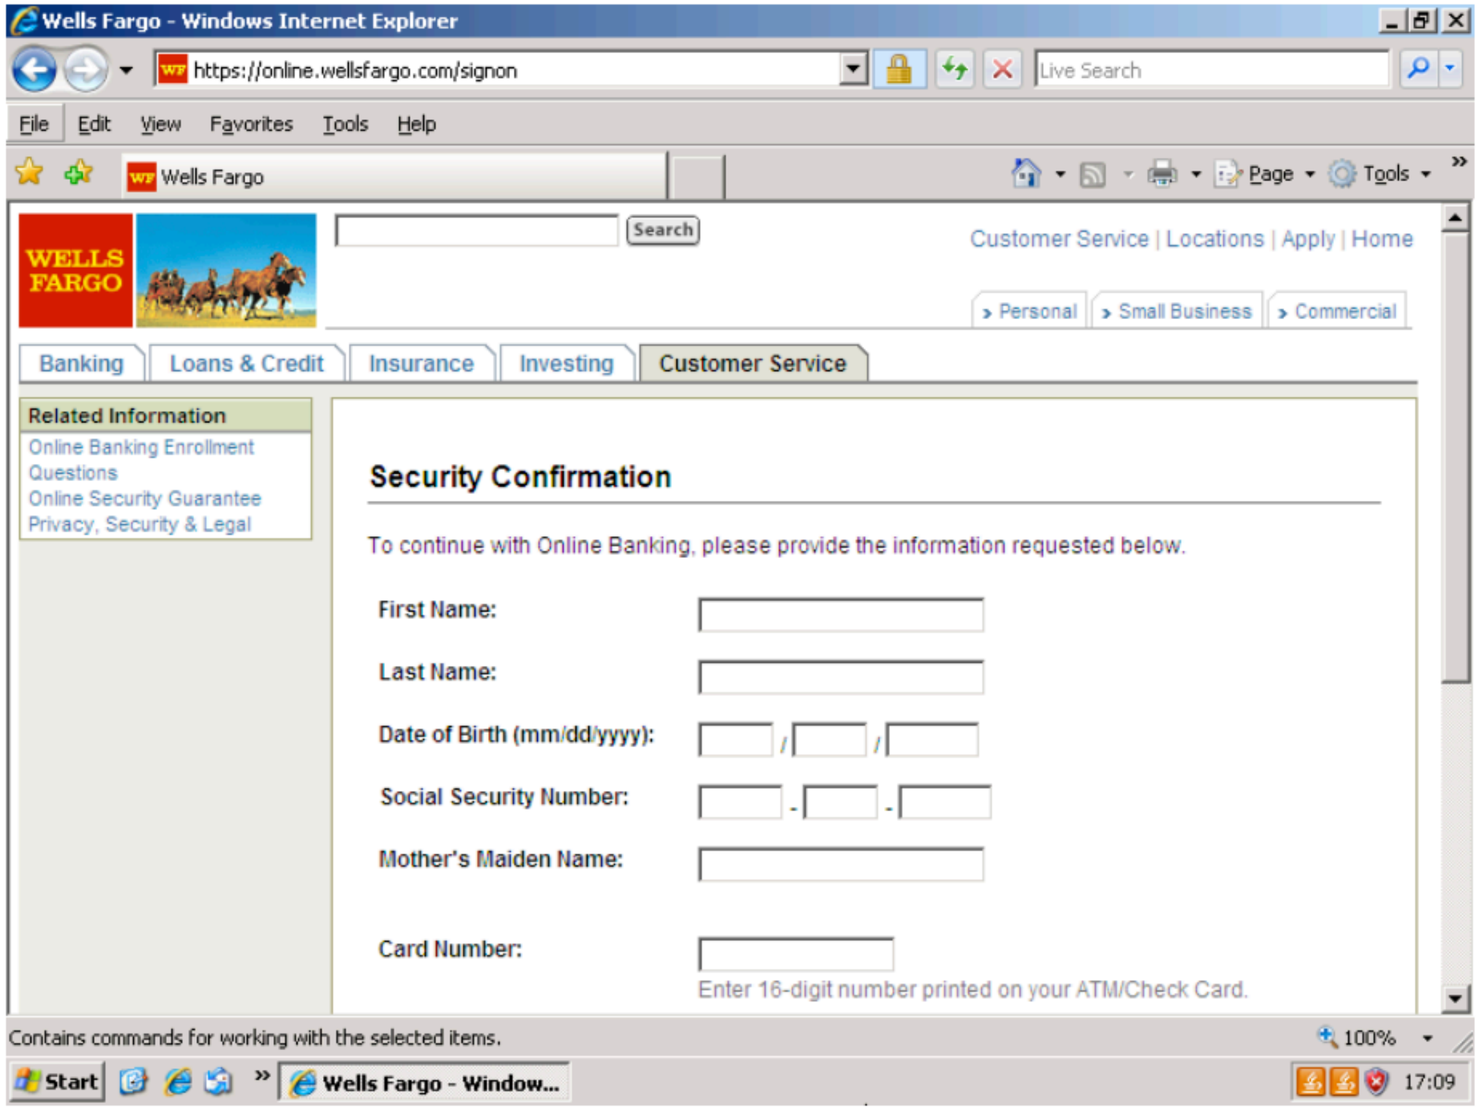

# Torpig

Infected an estimated ~185K machines

| Data Type | Data Items (#) |
|---|---|
| Mailbox account | 54,090 |
| Email | 1,258,862 |
| Form data | 11,966,532 |
| HTTP account | 411,039 |
| FTP account | 12,307 |
| POP account | 415,206 |
| SMTP account | 100,472 |
| Windows password | 1,235,122 |

```
[gnh5_229]
[MSO2002-MSO2003:pop.smith.com:John Smith:
    john@smith.com]
[pop3://john:smith@pop.smith.com:110]
[smtp://:@smtp.smith.com:25]
```

```
[gnh5_229]
POST /accounts/LoginAuth
Host:  www.google.com
POST_FORM:
Email=test@gmail.com
Passwd=test
```

# E-Crime

# Commoditization

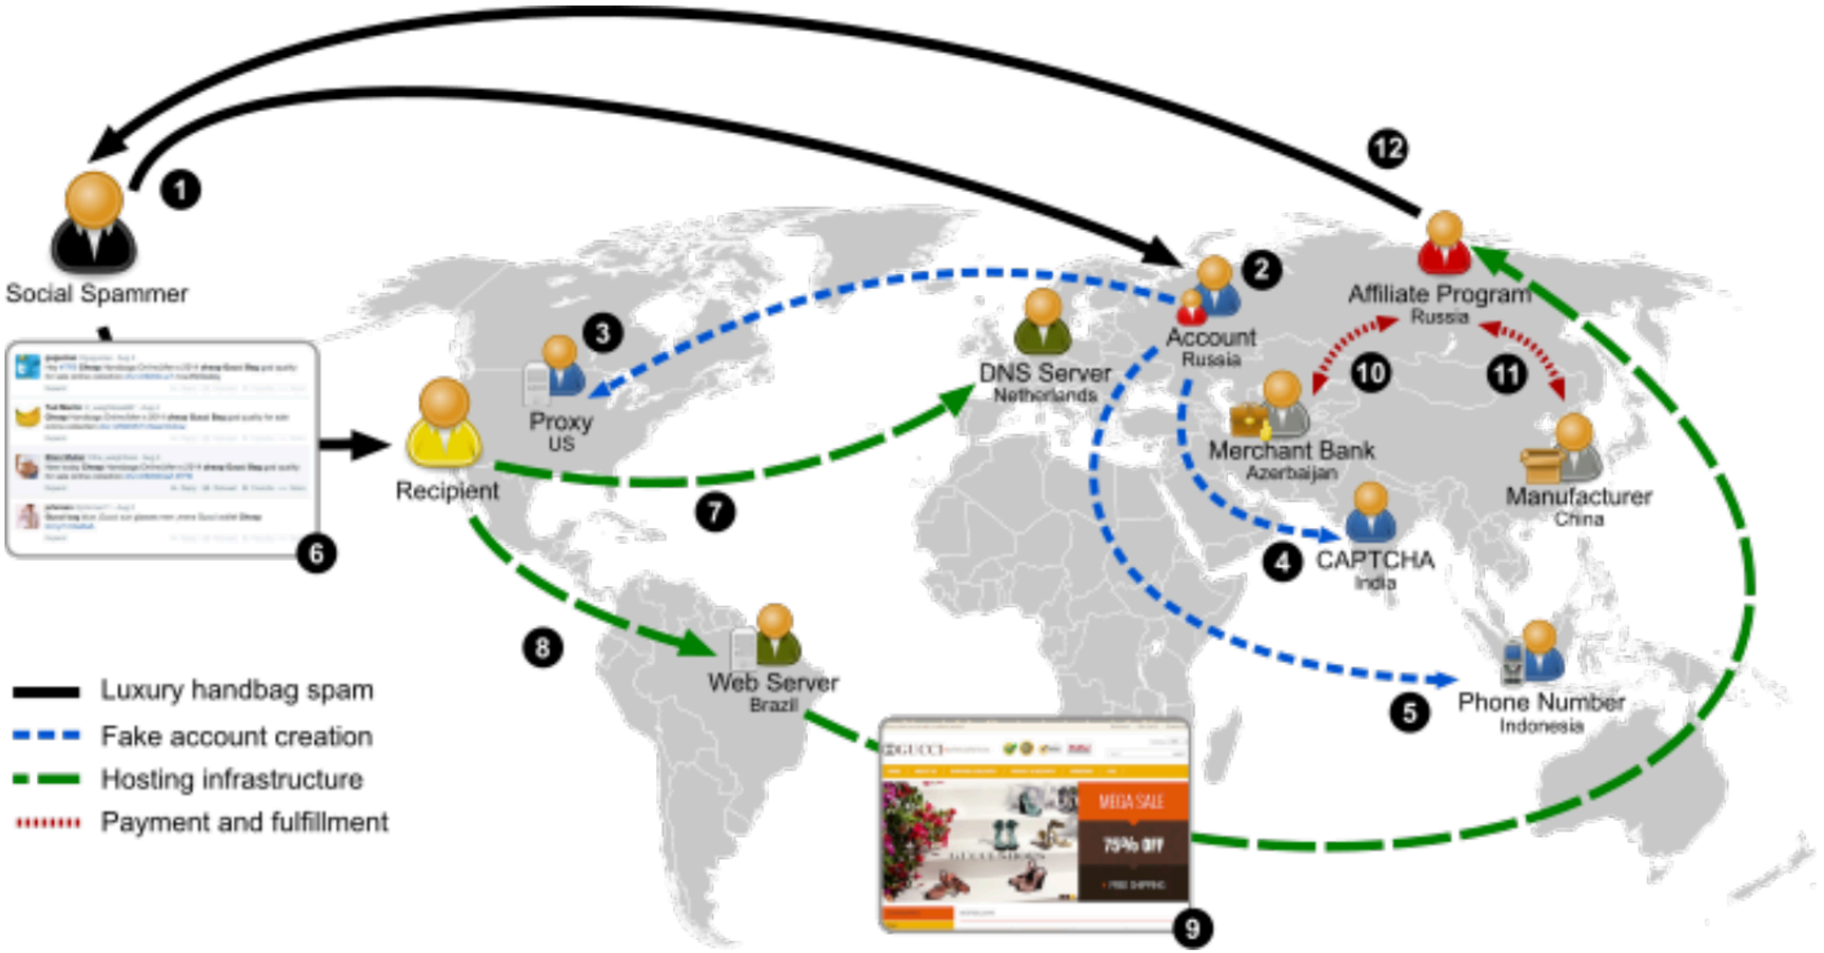


Framing Dependencies Introduced by Underground Commoditization. Thomas et al.

# How much do … cost?

| Support Center | Resource | Estimated Cost | Volume or Period |
|---|---|---|---|
| *Compromised Hosts* | Blackhole exploit kit [27] | $1,500 | 1 year |
| | Nuclear exploit kit [27] | $1,500 | 1 year |
| | Neutrino exploit kit [27] | $450 | 1 month |
| | Phoenix exploit kit [27] | $1,000–1,500 | 1 month |
| | Pay-per-install: US/UK [15] | $100–180 | 1,000 |
| | Pay-per-install: Europe [15] | $20–160 | 1,000 |
| | Pay-per-install: Other [15] | <$10 | 1,000 |
| *Human Services* | CAPTCHAs [106] | $1–2 | 1,000 |
| | SMS challenge [143] | $200 | 1,000 |
| | Mobile SIMs [143] | $140–420 | 1,000 |
| | English blog content [107] | $2–4 | 1 |
| | Chinese blog content [156] | $0.25 | 1 |
| *Networking & Hosting* | Proxy: 150 IPs | $25 | 1 month |
| | Proxy: 15,000–30,000 IPs | $250 | 1 month |
| | DDoS: 800 Mbps [70] | $10 | 1 month |
| | DDoS: 100 Gbps [30] | $200 | 1 day |
| *Accounts & Engagement* | Hotmail account [145] | $4–30 | 1,000 |
| | Yahoo account [145] | $6–15 | 1,000 |
| | Twitter account [145] | $1–20 | 1,000 |
| | Facebook PVA [145] | $80–400 | 1,000 |
| | Google PVA [145] | $80–500 | 1,000 |
| | Twitter followers [136] | $4–20 | 1,000 |
| | Twitter retweets [136] | $79–550 | 1,000 |
| | Facebook likes [36] | $15–70 | 1,000 |

Framing Dependencies Introduced by Underground Commoditization. Thomas et al.

# Estimated Profit

| Profit Center | Strategy | Estimated Revenue | Time Frame |
|---|---|---|---|
| *Spamvertised products* | Pharamcuticals [97] | $12–92 million | 2007–2010 |
| | Luxury knock-offs [152] | $68 million | 2013–2014 |
| *Scareware & Ransomware* | Fake anti-virus [133] | $130 million | 2008–2010 |
| | CryptoLocker [159]* | $3 million | 2013–2014 |
| *Clickfraud* | ZeroAccess [115] | $36 million | 2013 |
| | DNS Changer [149]* | $14 million | 2007–2011 |
| *Financial Scams* | Pump and dump [150]* | $120 million | 2008–2013 |
| | 419 scammers [8]* | $200 million | 2006 |
| *Credit Card Theft* | ATM withdrawl scam [118]* | $45 million | 1 day |
| | Zeus banking trojan [9]* | $70 million | 2009–2010 |
| | Re-selling stolen cards [35]* | $300 million | ?–2013 |

Framing Dependencies Introduced by Underground Commoditization. Thomas et al.