UNDERGROUND ECONOMIES

GRAD SEC

OCT 17 2017

TODAY'S PAPERS

Click Trajectories: End-to-End Analysis of the Spam Value Chain

Kirill Levchenko* Andreas Pitsillidis* Neha Chachra* Brandon Enright* Márk Félegyházi‡ Chris Grier†
Tristan Halvorson* Chris Kanich* Christian Kreibich†◇ He Liu* Damon McCoy*
Nicholas Weaver†◇ Vern Paxson†◇ Geoffrey M. Voelker* Stefan Savage*

*Department of Computer Science and Engineering
University of California, San Diego

†Computer Science Division
University of California, Berkeley

◇International Computer Science Institute
Berkeley, CA

‡Laboratory of Cryptography and System Security (CrySyS)
Budapest University of Technology and Economics

***Abstract*—Spam-based advertising is a business. While it has engendered both widespread antipathy and a multi-billion dollar anti-spam industry, it continues to exist because it fuels a profitable enterprise. We lack, however, a solid understanding of this enterprise's full structure, and thus most anti-spam interventions focus on only one facet of the overall spam value chain (e.g., spam filtering, URL blacklisting, site takedown). In this paper we present a holistic analysis that quantifies the full set of resources employed to monetize spam email—including naming, hosting, payment and fulfillment—using extensive measurements of three months of diverse spam data, broad crawling of naming and hosting infrastructures, and over 100 purchases from spam-advertised sites. We relate these resources to the organizations who administer them and then use this data to characterize the relative prospects for defensive interventions at each link in the spam value chain. In particular, we provide the first strong evidence of payment bottlenecks in the spam value chain; 95% of spam-advertised pharmaceutical, replica and software products are monetized using merchant services from just a handful of banks.**

I. Introduction

We may think of email spam as a scourge—jamming our collective inboxes with tens of billions of unwanted messages each day—but to its perpetrators it is a potent marketing channel that taps latent demand for a variety of products and services. While most attention focuses on the problem of spam *delivery*, the email vector itself comprises only the *visible* portion of a large, multi-faceted business enterprise. Each click on a spam-advertised link is in fact just the start of a long and complex trajectory, spanning a range of both technical and business components that together provide the necessary infrastructure needed to monetize a customer's visit. Botnet services must be secured, domains registered, name servers provisioned, and hosting or proxy services acquired. All of these, in addition to payment processing, merchant bank accounts, customer service, and fulfillment, reflect necessary elements in the spam value chain.

While elements of this chain have received study in isolation (e.g., dynamics of botnets [20], DNS fast-flux networks [17], [42], Web site hosting [1], [22]), the relationship between them is far less well understood. Yet it is these very relationships that capture the structural dependencies—and hence the potential *weaknesses*—within the spam ecosystem's business processes. Indeed, each distinct path through this chain—registrar, name server, hosting, affiliate program, payment processing, fulfillment—directly reflects an "entrepreneurial activity" by which the perpetrators muster capital investments and business relationships to create value. Today we lack insight into even the most basic characteristics of this activity. How many organizations are complicit in the spam ecosystem? Which points in their value chains do they share and which operate independently? How "wide" is the bottleneck at each stage of the value chain—do miscreants find alternatives plentiful and cheap, or scarce, requiring careful husbanding?

The desire to address these kinds of questions empirically—and thus guide decisions about the most effective mechanisms for addressing the spam problem—forms the core motivation of our work. In this paper we develop a methodology for characterizing the end-to-end resource dependencies ("trajectories") behind individual spam campaigns and then analyze the relationships among them. We use three months of real-time source data, including captive botnets, raw spam feeds, and feeds of spam-advertised URLs to drive active probing of spam infrastructure elements (name servers, redirectors, hosting proxies). From these, we in turn identify those sites advertising three popular classes of goods—pharmaceuticals, replica luxury goods and counterfeit software—as well as their membership in specific affiliate programs around which the overall business is structured. Finally, for a subset of these sites we perform on-line purchases, providing additional data about merchant bank affiliation, customer service, and fulfillment. Using this data we characterize the resource footprint at each step in the spam value chain, the extent of sharing between spam organizations and, most importantly, the relative prospects for interrupting spam monetization at different stages of the process.

The remainder of this paper is organized as follows. Section II provides a qualitative overview of the spam ecosystem coupled with a review of related research.

Show Me the Money: Characterizing Spam-advertised Revenue

Chris Kanich* Nicholas Weaver† Damon McCoy* Tristan Halvorson*
Christian Kreibich† Kirill Levchenko*
Vern Paxson†‡ Geoffrey M. Voelker* Stefan Savage*

*Department of Computer Science and Engineering
University of California, San Diego

†International Computer Science Institute
Berkeley, CA

‡Computer Science Division
University of California, Berkeley

Abstract

Modern spam is ultimately driven by product sales: goods purchased by customers online. However, while this model is easy to state in the abstract, our understanding of the concrete business environment—how many orders, of what kind, from which customers, for how much—is poor at best. This situation is unsurprising since such sellers typically operate under questionable legal footing, with "ground truth" data rarely available to the public. However, absent quantifiable empirical data, "guesstimates" operate unchecked and can distort both policy making and our choice of appropriate interventions. In this paper, we describe two inference techniques for peering inside the business operations of spam-advertised enterprises: purchase pair and basket inference. Using these, we provide informed estimates on order volumes, product sales distribution, customer makeup and total revenues for a range of spam-advertised programs.

1 Introduction

A large number of Internet scams are "advertising-based"; that is, their goal is to convince potential customers to purchase a product or service, typically via some broad-based advertising medium.[1] In turn, this activity mobilizes and helps fund a broad array of technical capabilities, including botnet-based distribution, fast flux name service, and bulletproof hosting. However, while these same technical aspects enjoy a great deal of attention from the security community, there is considerably less information quantifying the underlying economic engine that drives this ecosystem. Absent grounded empirical data, it is challenging to reconcile revenue "estimates" that can range from $2M/day for one spam botnet [1], to analyses suggesting that spammers make little money at all [6]. This situation has the potential to distort policy and investment decisions that are otherwise driven by intuition rather than evidence.

In this paper we make two contributions to improving this state of affairs using measurement-based methods to estimate:

- *Order volume.* We describe a general technique—purchase pair—for estimating the number of orders received (and hence revenue) via on line store order numbering. We use this approach to establish rough, but well-founded, monthly order volume estimates for many of the leading "affiliate programs" selling counterfeit pharmaceuticals and software.
- *Purchasing behavior.* We show how we can use third-party image hosting data to infer the contents of customer "baskets" and hence characterize purchasing behavior. We apply this technique to a leading spamvertized pharmaceutical program and identify both the nature of these purchases and their relation to the geographic distribution of the customer base.

In each case, our real contribution is less in the particular techniques—which an adversary could easily defeat should they seek to do so—but rather in the data that we used them to gather. In particular, we document that seven leading counterfeit pharmacies together have a total monthly order volume in excess of 82,000, while three counterfeit software stores process over 37,000 orders in the same time.

On the demand side, as expected, we find that most pharmaceuticals selected for purchase are in the "male-enhancement" category (primarily Viagra and other ED medications comprising 60 distinct items). However, such drugs constitute only 62% of the total, and we document that this demand distribution has quite a long tail: user shopping carts contain 289 distinct products, including surprising categories such as anti-cancer medications

[1] Unauthorized Internet advertising includes email spam, black hat search-engine optimization [25], blog spam [21], Twitter spam [4], forum spam, and comment spam. Hereafter we refer to these myriad advertising vectors simply as spam.

UNDERGROUND ECONOMIES

- Economics drives both the attacks and the defenses

- What is for sale? Who sells it? How?
    - Defenders: Antivirus vendors, firewall vendors, etc.
    - *What about the attackers?*

- The idea is that we may be able to stem attacks if we can understand
    - the incentives
    - the choke points (might there be one bank we could shut down to cease spam?)

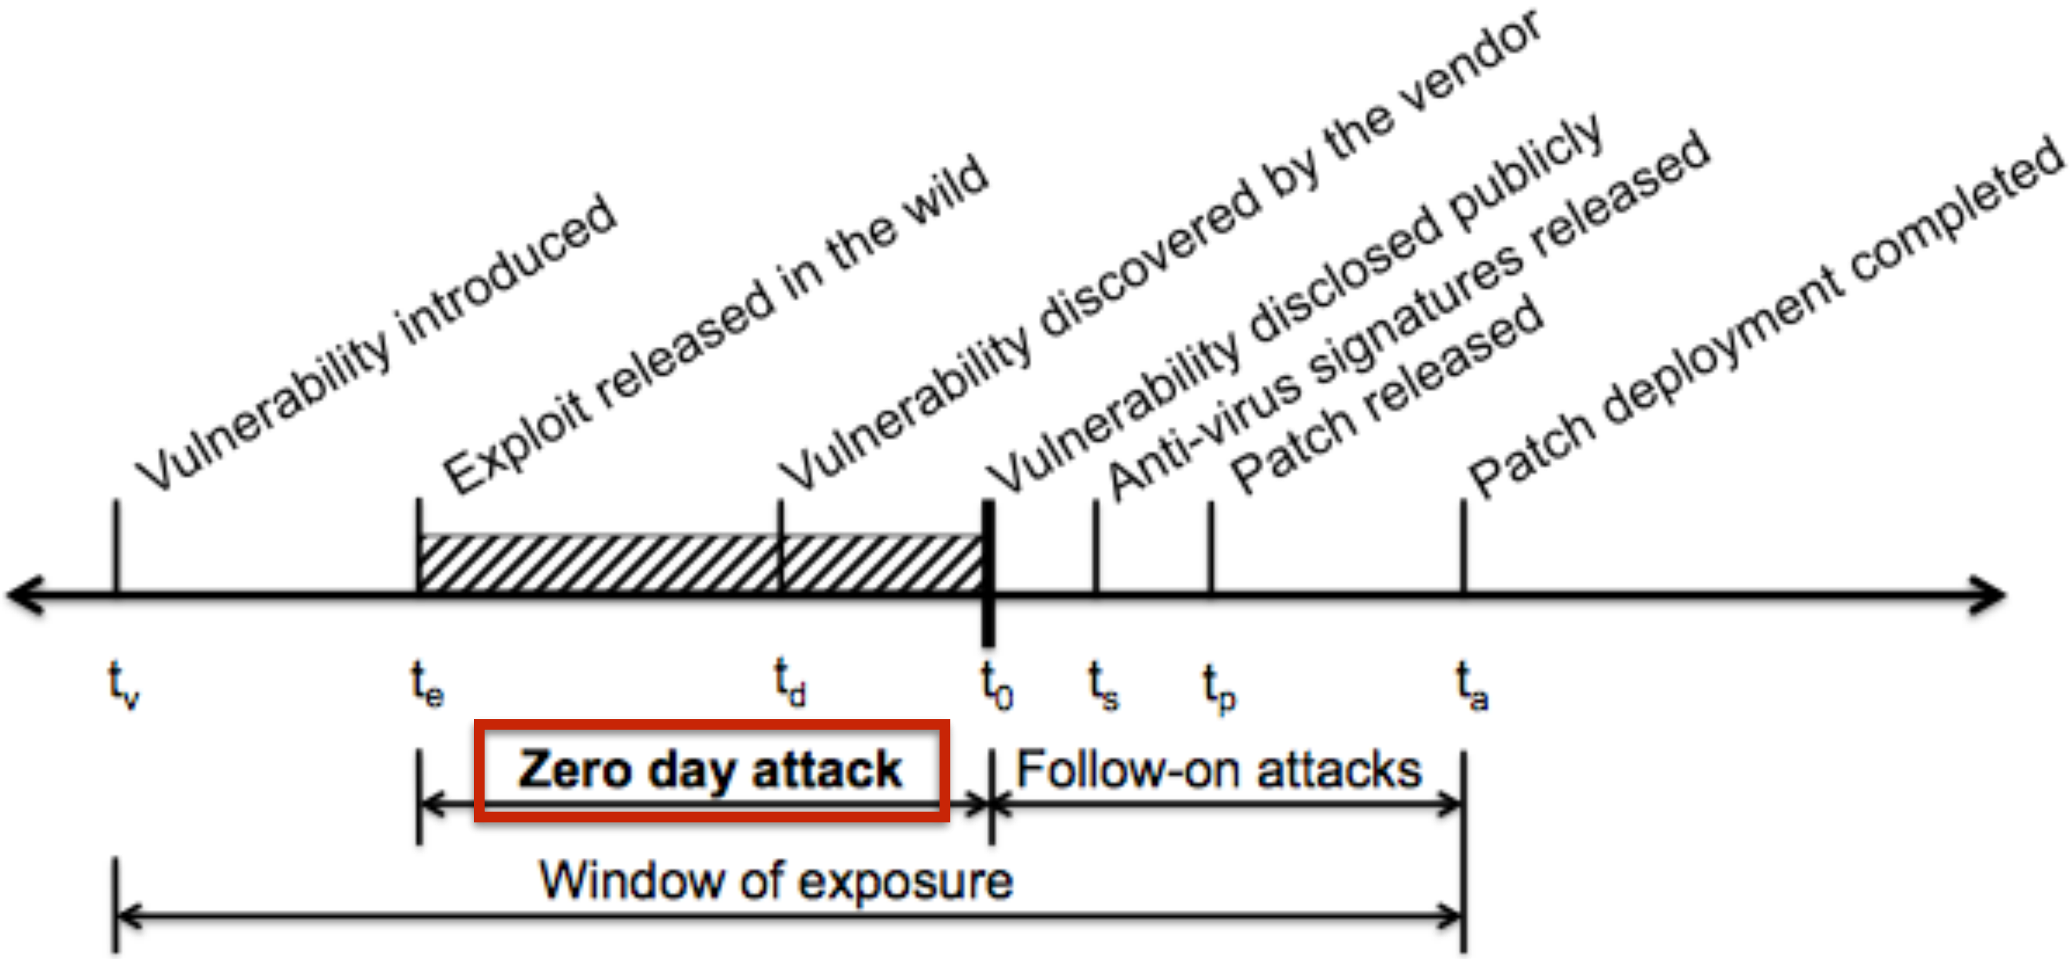
Vulnerability introduced
Exploit released in the wild
Vulnerability discovered by the vendor
Vulnerability disclosed publicly
Anti-virus signatures released
Patch released
Patch deployment completed
t_v
t_e
t_d
t_0
t_s
t_p
t_a
Zero day attack
Follow-on attacks
Window of exposure

- **Who buys**: Attackers, spies (and the companies who wrote the software) want to know about them

- **Through whom**: anonymous middlemen (e.g. Grusq) who match vulnerability finders up with buyers. Take commission (15% typical).

- **Payment**: Made in installments (cease payment when zero-day over)

Google offers a max of $3133.70 for information about flaws in their tech

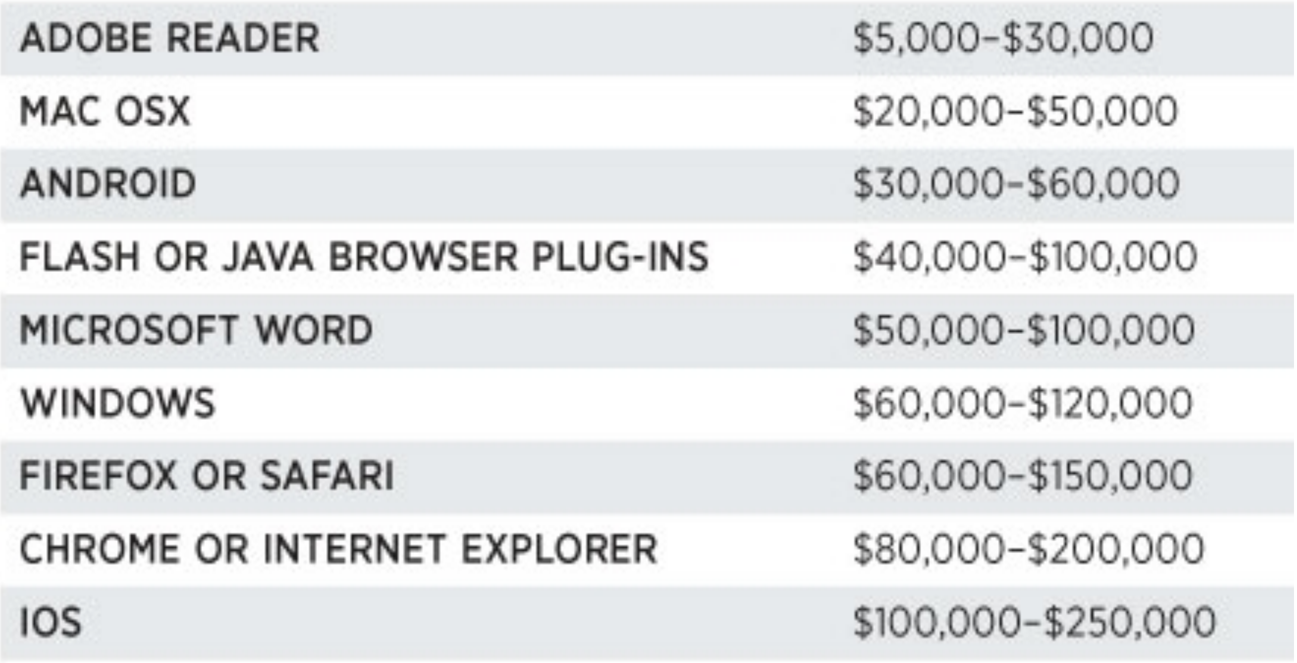

ADOBE READER	$5,000-$30,000
MAC OSX	$20,000-$50,000
ANDROID	$30,000-$60,000
FLASH OR JAVA BROWSER PLUG-INS	$40,000-$100,000
MICROSOFT WORD	$50,000-$100,000
WINDOWS	$60,000-$120,000
FIREFOX OR SAFARI	$60,000-$150,000
CHROME OR INTERNET EXPLORER	$80,000-$200,000
IOS	$100,000-$250,000

"Shopping for zero-days" Forbes 2012

BUG BOUNTY PROGRAMS

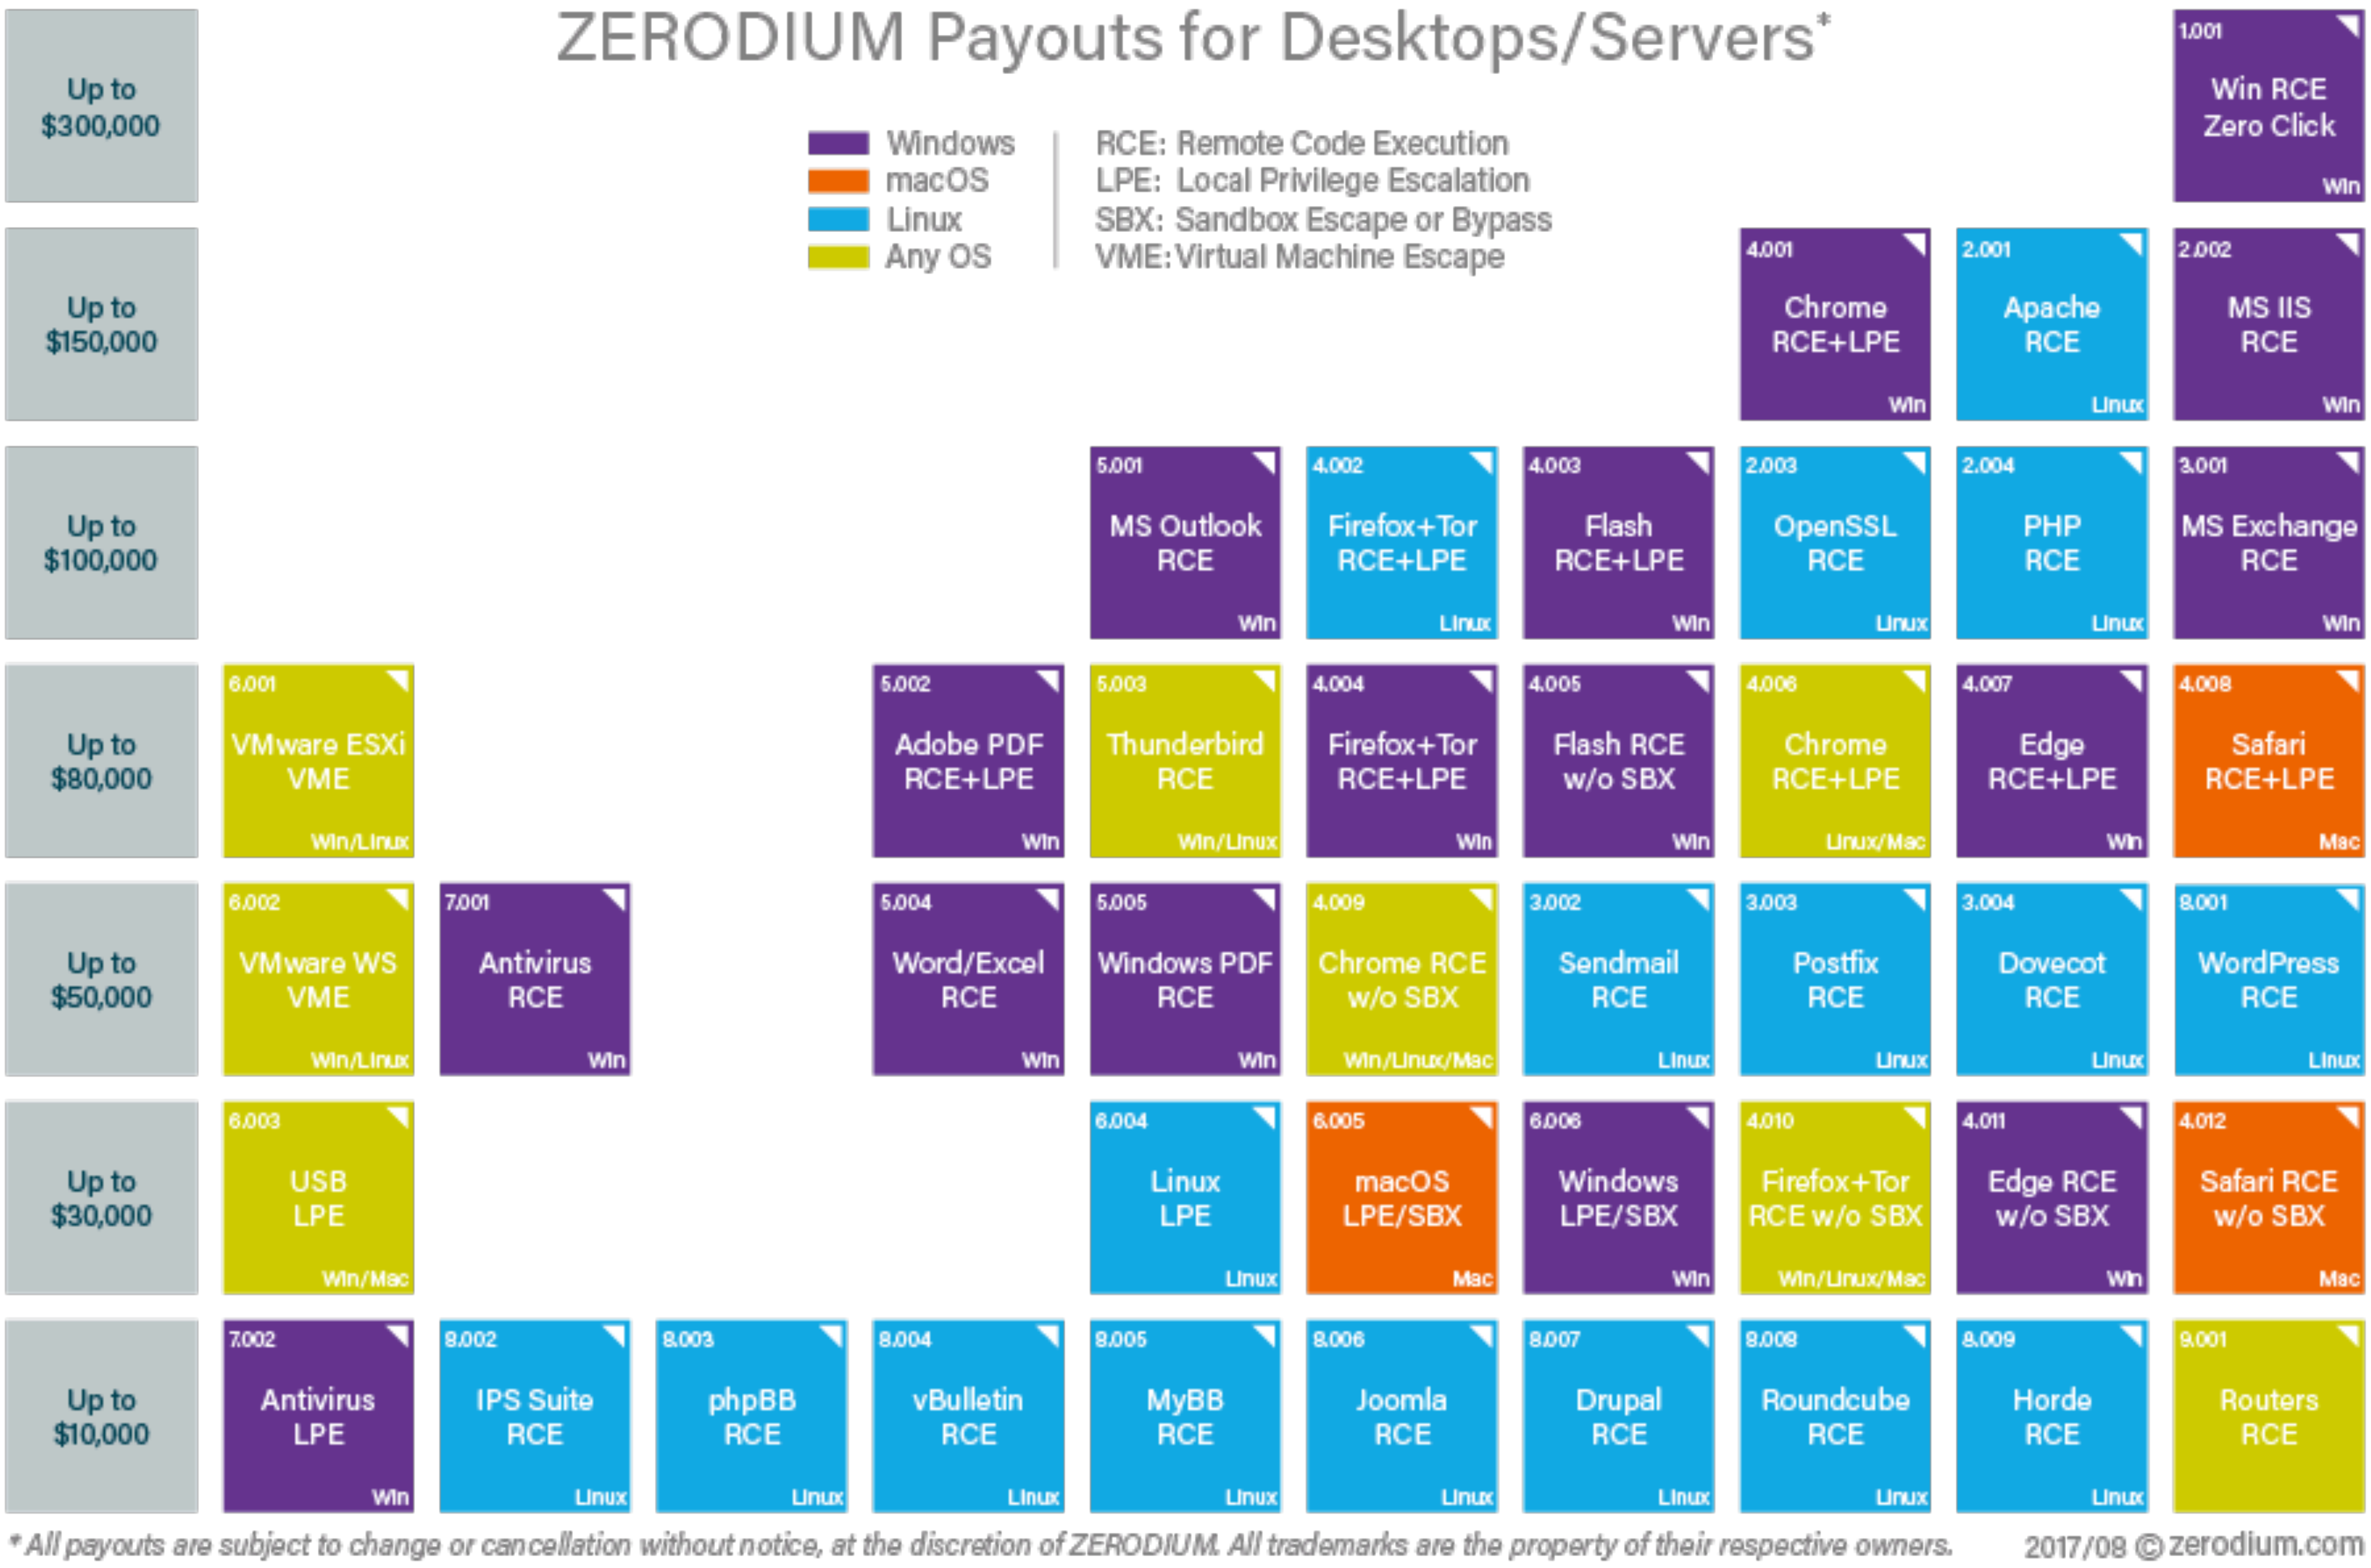

ZERODIUM Payouts for Desktops/Servers*

🟪 Windows
🟧 macOS
🟦 Linux
🟨 Any OS

RCE: Remote Code Execution
LPE: Local Privilege Escalation
SBX: Sandbox Escape or Bypass
VME: Virtual Machine Escape

Up to $300,000										1.001 Win RCE Zero Click Win
Up to $150,000								4.001 Chrome RCE+LPE Win	2.001 Apache RCE Linux	2.002 MS IIS RCE Win
Up to $100,000					5.001 MS Outlook RCE Win	4.002 Firefox+Tor RCE+LPE Linux	4.003 Flash RCE+LPE Win	2.003 OpenSSL RCE Linux	2.004 PHP RCE Linux	3.001 MS Exchange RCE Win
Up to $80,000	6.001 VMware ESXi VME Win/Linux			5.002 Adobe PDF RCE+LPE Win	5.003 Thunderbird RCE Win/Linux	4.004 Firefox+Tor RCE+LPE Win	4.005 Flash RCE w/o SBX Win	4.006 Chrome RCE+LPE Linux/Mac	4.007 Edge RCE+LPE Win	4.008 Safari RCE+LPE Mac
Up to $50,000	6.002 VMware WS VME Win/Linux	7.001 Antivirus RCE Win		5.004 Word/Excel RCE Win	5.005 Windows PDF RCE Win	4.009 Chrome RCE w/o SBX Win/Linux/Mac	3.002 Sendmail RCE Linux	3.003 Postfix RCE Linux	3.004 Dovecot RCE Linux	8.001 WordPress RCE Linux
Up to $30,000	6.003 USB LPE Win/Mac				6.004 Linux LPE Linux	6.005 macOS LPE/SBX Mac	6.006 Windows LPE/SBX Win	4.010 Firefox+Tor RCE w/o SBX Win/Linux/Mac	4.011 Edge RCE w/o SBX Win	4.012 Safari RCE w/o SBX Mac
Up to $10,000	7.002 Antivirus LPE Win	8.002 IPS Suite RCE Linux	8.003 phpBB RCE Linux	8.004 vBulletin RCE Linux	8.005 MyBB RCE Linux	8.006 Joomla RCE Linux	8.007 Drupal RCE Linux	8.008 Roundcube RCE Linux	8.009 Horde RCE Linux	9.001 Routers RCE

* *All payouts are subject to change or cancellation without notice, at the discretion of ZERODIUM. All trademarks are the property of their respective owners.*

2017/08 © zerodium.com

BUG BOUNTY PROGRAMS

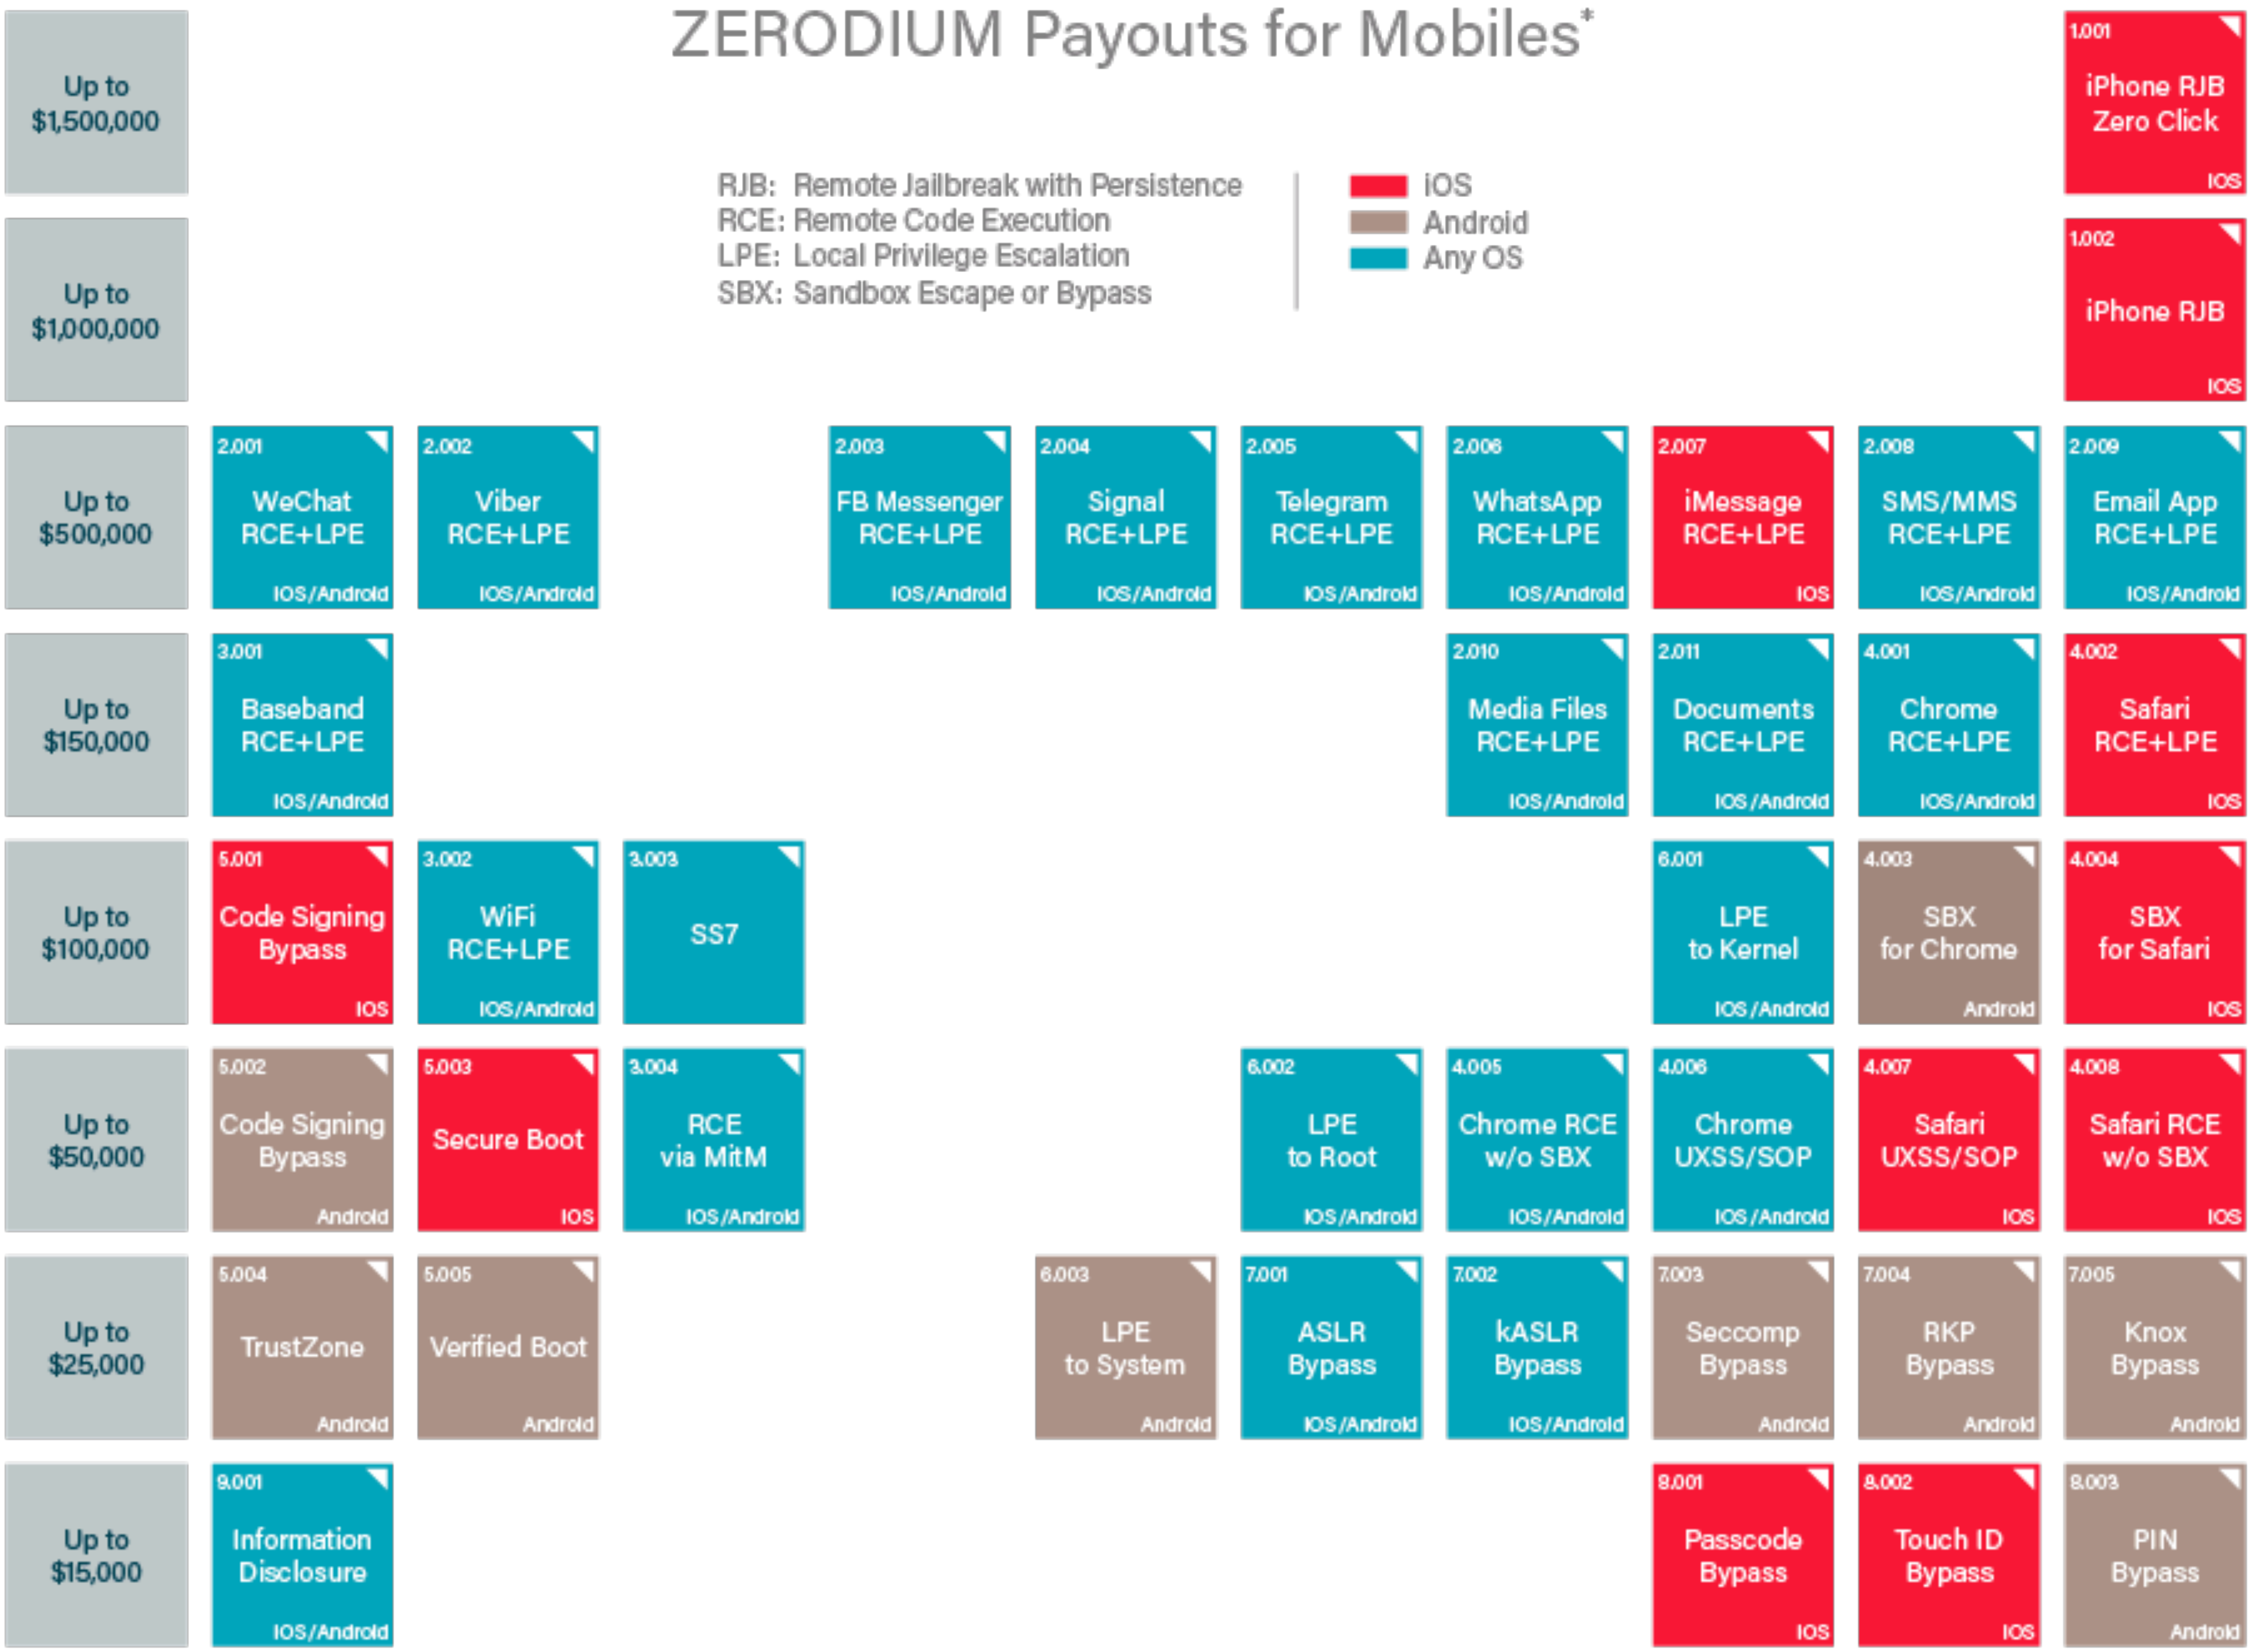

ZERODIUM Payouts for Mobiles*

RJB: Remote Jailbreak with Persistence
RCE: Remote Code Execution
LPE: Local Privilege Escalation
SBX: Sandbox Escape or Bypass

🟥 iOS
🟫 Android
🟦 Any OS

Up to $1,500,000									1.001 iPhone RJB Zero Click iOS
Up to $1,000,000									1.002 iPhone RJB iOS
Up to $500,000	2.001 WeChat RCE+LPE iOS/Android	2.002 Viber RCE+LPE iOS/Android		2.003 FB Messenger RCE+LPE iOS/Android	2.004 Signal RCE+LPE iOS/Android	2.005 Telegram RCE+LPE iOS/Android	2.006 WhatsApp RCE+LPE iOS/Android	2.007 iMessage RCE+LPE iOS	2.008 SMS/MMS RCE+LPE iOS/Android	2.009 Email App RCE+LPE iOS/Android
Up to $150,000	3.001 Baseband RCE+LPE iOS/Android						2.010 Media Files RCE+LPE iOS/Android	2.011 Documents RCE+LPE iOS/Android	4.001 Chrome RCE+LPE iOS/Android	4.002 Safari RCE+LPE iOS
Up to $100,000	5.001 Code Signing Bypass iOS	3.002 WiFi RCE+LPE iOS/Android	3.003 SS7					6.001 LPE to Kernel iOS/Android	4.003 SBX for Chrome Android	4.004 SBX for Safari iOS
Up to $50,000	5.002 Code Signing Bypass Android	5.003 Secure Boot iOS	3.004 RCE via MitM iOS/Android			6.002 LPE to Root iOS/Android	4.005 Chrome RCE w/o SBX iOS/Android	4.006 Chrome UXSS/SOP iOS/Android	4.007 Safari UXSS/SOP iOS	4.008 Safari RCE w/o SBX iOS
Up to $25,000	5.004 TrustZone Android	5.005 Verified Boot Android			6.003 LPE to System Android	7.001 ASLR Bypass iOS/Android	7.002 kASLR Bypass iOS/Android	7.003 Seccomp Bypass Android	7.004 RKP Bypass Android	7.005 Knox Bypass Android
Up to $15,000	9.001 Information Disclosure iOS/Android							8.001 Passcode Bypass iOS	8.002 Touch ID Bypass iOS	8.003 PIN Bypass Android

BUG BOUNTY PROGRAMS

Apple's bug bounty program hindered by low payouts, report says

By Mikey Campbell
Thursday, July 06, 2017, 04:13 pm PT (07:13 pm ET)

Apple's invite-only bug bounty program is off to a slow start as security researchers in search of high payouts are saving discovered exploits for high-price sales on the gray market.

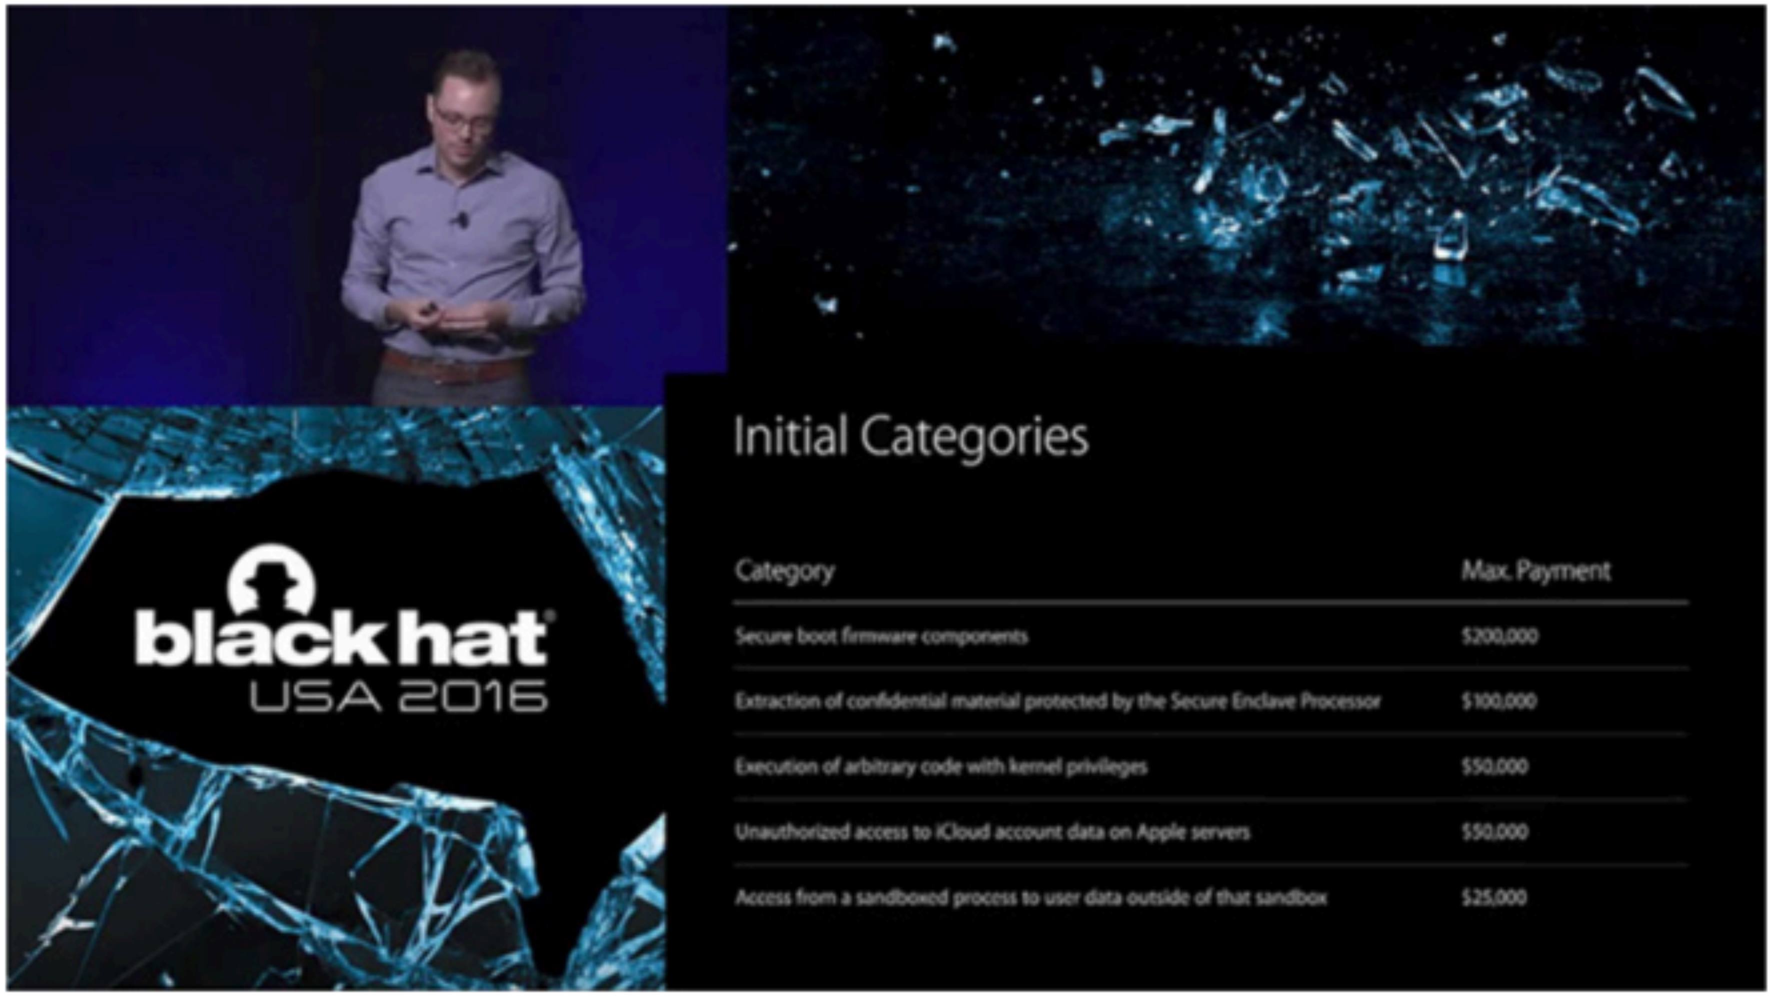

Initial Categories

Category	Max. Payment
Secure boot firmware components	$200,000
Extraction of confidential material protected by the Secure Enclave Processor	$100,000
Execution of arbitrary code with kernel privileges	$50,000
Unauthorized access to iCloud account data on Apple servers	$50,000
Access from a sandboxed process to user data outside of that sandbox	$25,000

Apple's Ivan Krstic announces the bug bounty program at Black Hat USA 2016.

$200k < $1.5M

iOS bugs are too valuable to report

BUG BOUNTY PROGRAMS

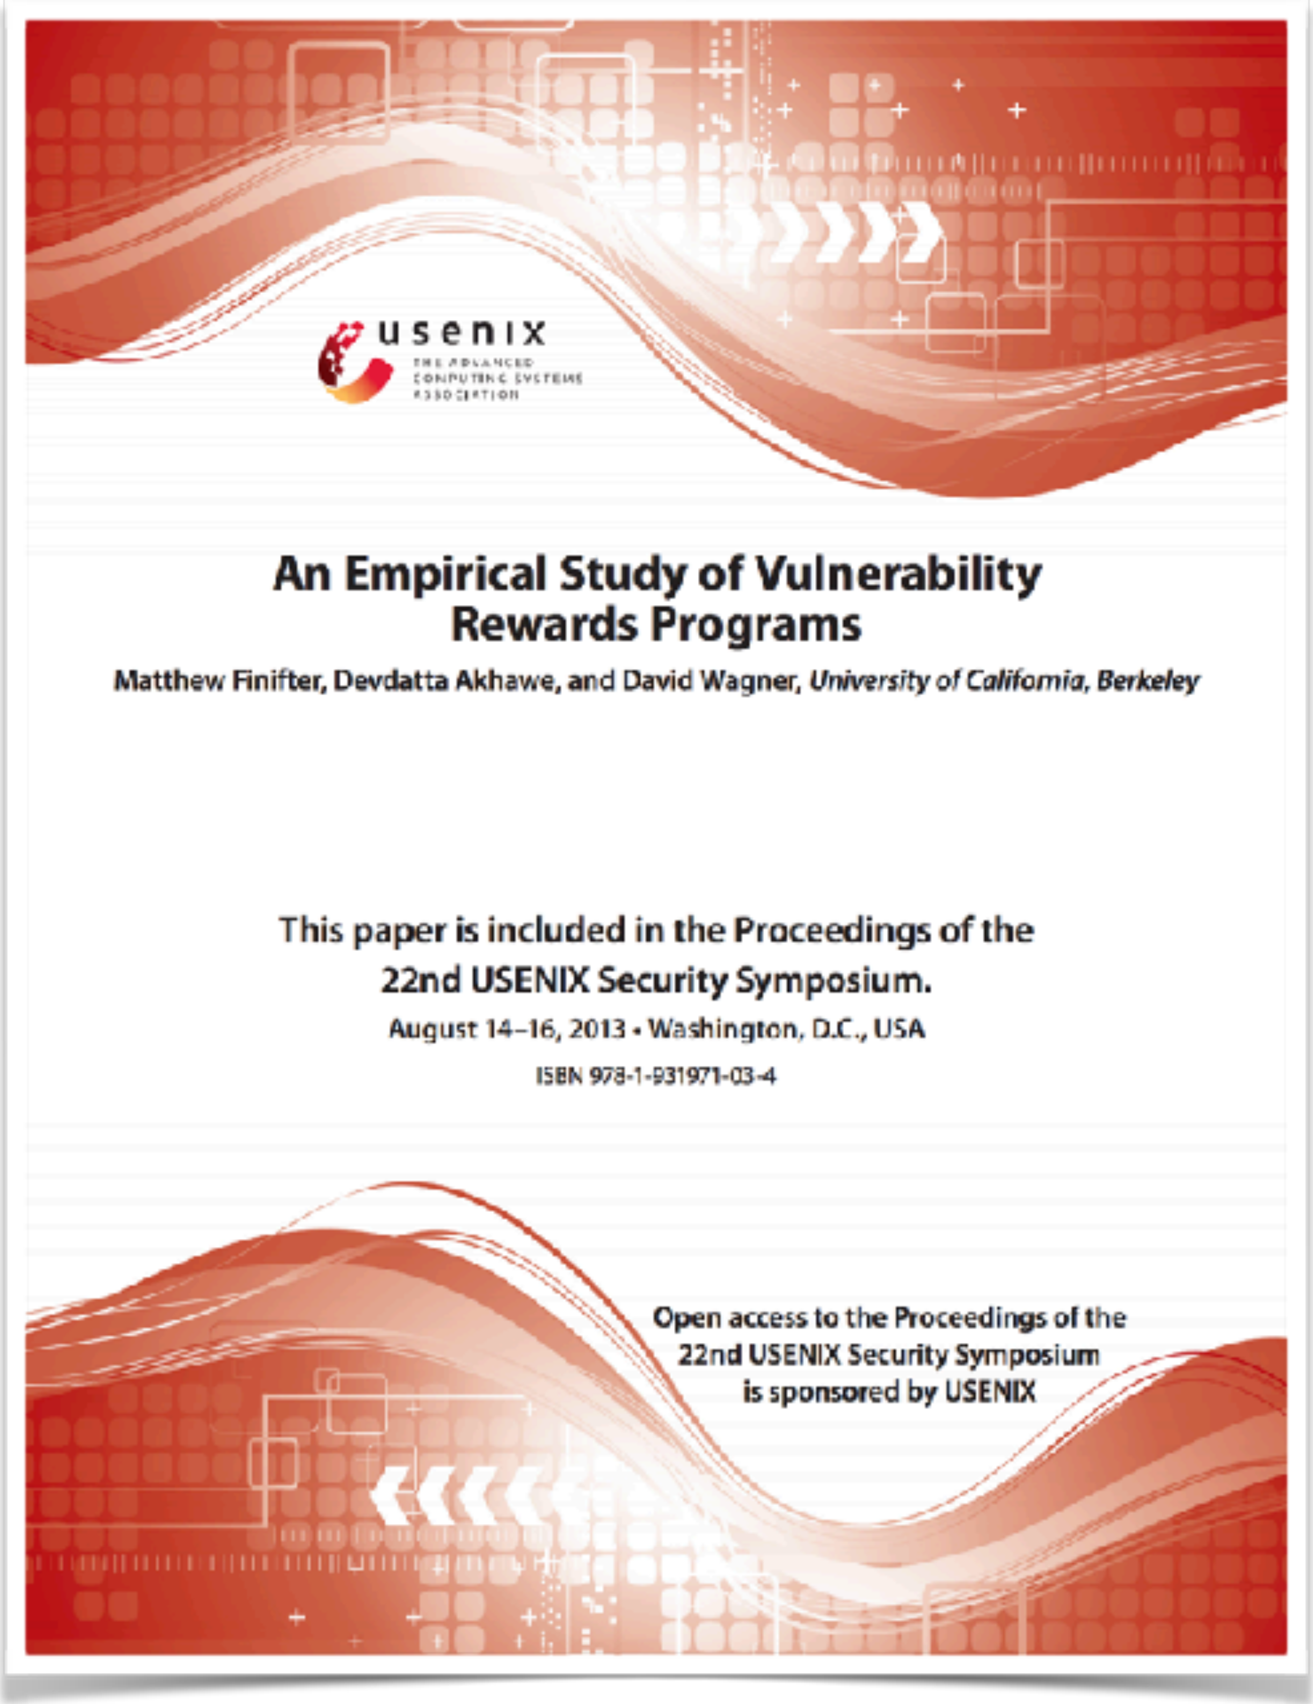

usenix
THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

An Empirical Study of Vulnerability Rewards Programs

Matthew Finifter, Devdatta Akhawe, and David Wagner, *University of California, Berkeley*

This paper is included in the Proceedings of the
22nd USENIX Security Symposium.
August 14–16, 2013 • Washington, D.C., USA
ISBN 978-1-931971-03-4

Open access to the Proceedings of the
22nd USENIX Security Symposium
is sponsored by USENIX

Studied Chrome & Firefox VRPs

VRPs yield patched vulnerabilities

28% of Chrome's patches

24% of Firefox's patches

VRPs are a good deal (for vendors)

Nowhere near full-time salary

What about today's bug bounty programs? What about 3rd parties?

SPAM

- Unsolicited, annoying email (or posts on blogs, social networks, etc.) that seeks to
  - Sell products
  - Get users to install malicious software

- Typical defenses
  - Look for key words in the messages
  - Block certain senders (**SpamHaus** blacklist of IP addrs)

- But what is the economics behind it all?
  - How do they send out so much email?
  - Are they selling real things? How?

SENDING SPAM

- Tons of email to send, and easy to block a single IP address from sending

- Need lots of IP addresses
  - But since SMTP (email) uses TCP, we need to actually be able to operate those IP addresses

- Buy lots of computers? (expensive)

Compromise lots of computers!

BOTNETS

- Collection of compromised machines (bots) under unified control of an attacker (botmaster)
- Method of compromise decoupled from method of control
  - Launch a worm/virus, etc.: remember, payload is orthogonal!
- Upon infection, a new bot "phones home" to *rendezvous* with botnet **"command-and-control" (C&C)**
- Botmaster uses C&C to push out commands and updates

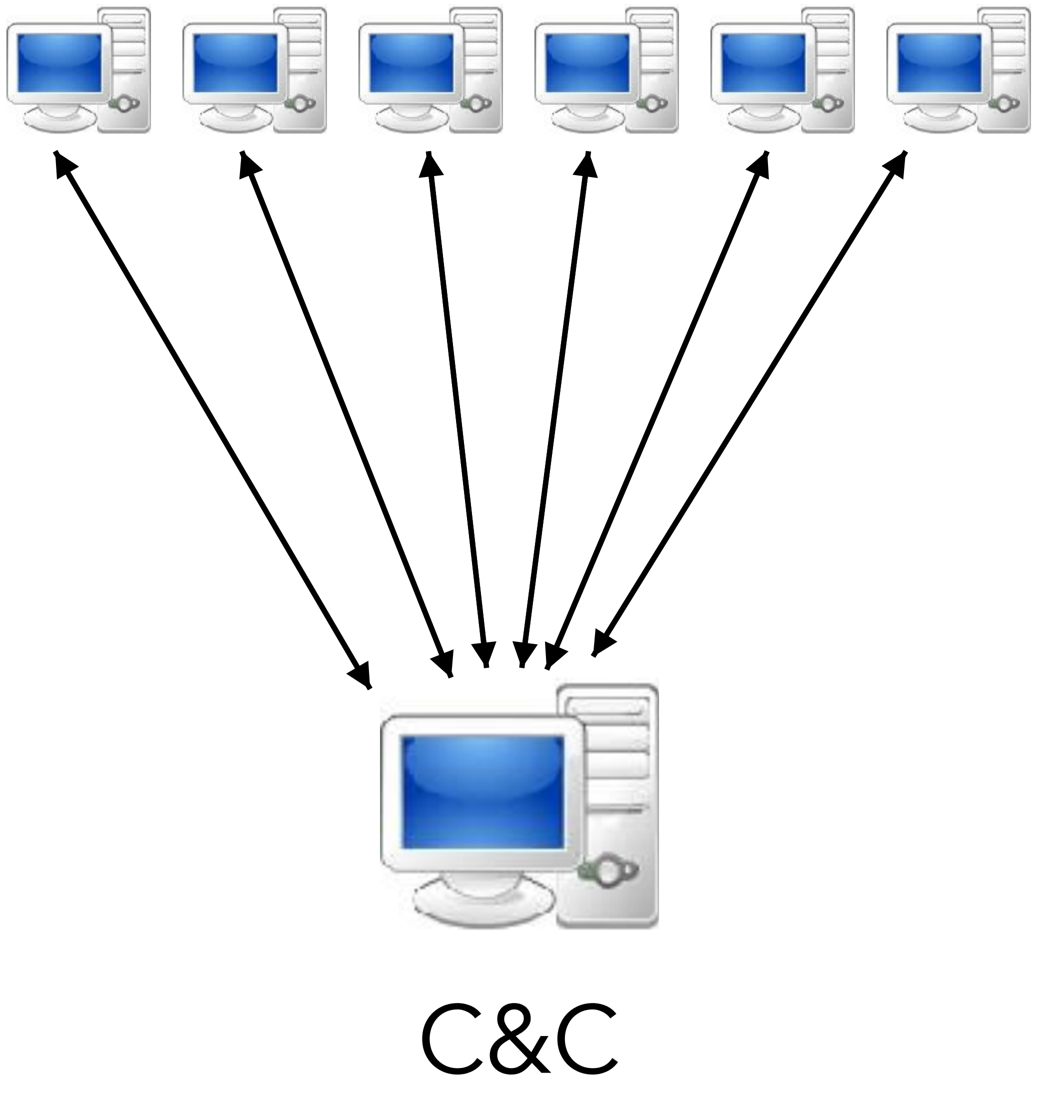

Topology can be star (like this), hierarchical, peer-to-peer…

SUPPORTING CLICKS

- Ideally a user will click on an embedded URL
- Result is more complex than just going to a web server
  - Defensive measures: URL and domain **blacklisting** & **takedown notices** by ISPs
- Confuse defenses (esp. blacklisting) with moving targets:
  - **Redirection sites** (legit-looking URL, like a URL shortener, or just manage DNS yourself and create throwaway domains that redirect to a more permanent domain)
  - **Bulk domains**: purchased from a reseller or as part of an affiliate program (more later)
- But web servers are static, so how do we keep them from being shut down due to blacklisting and takedown notices?

SPAMBOT

Botnet used for sending spam

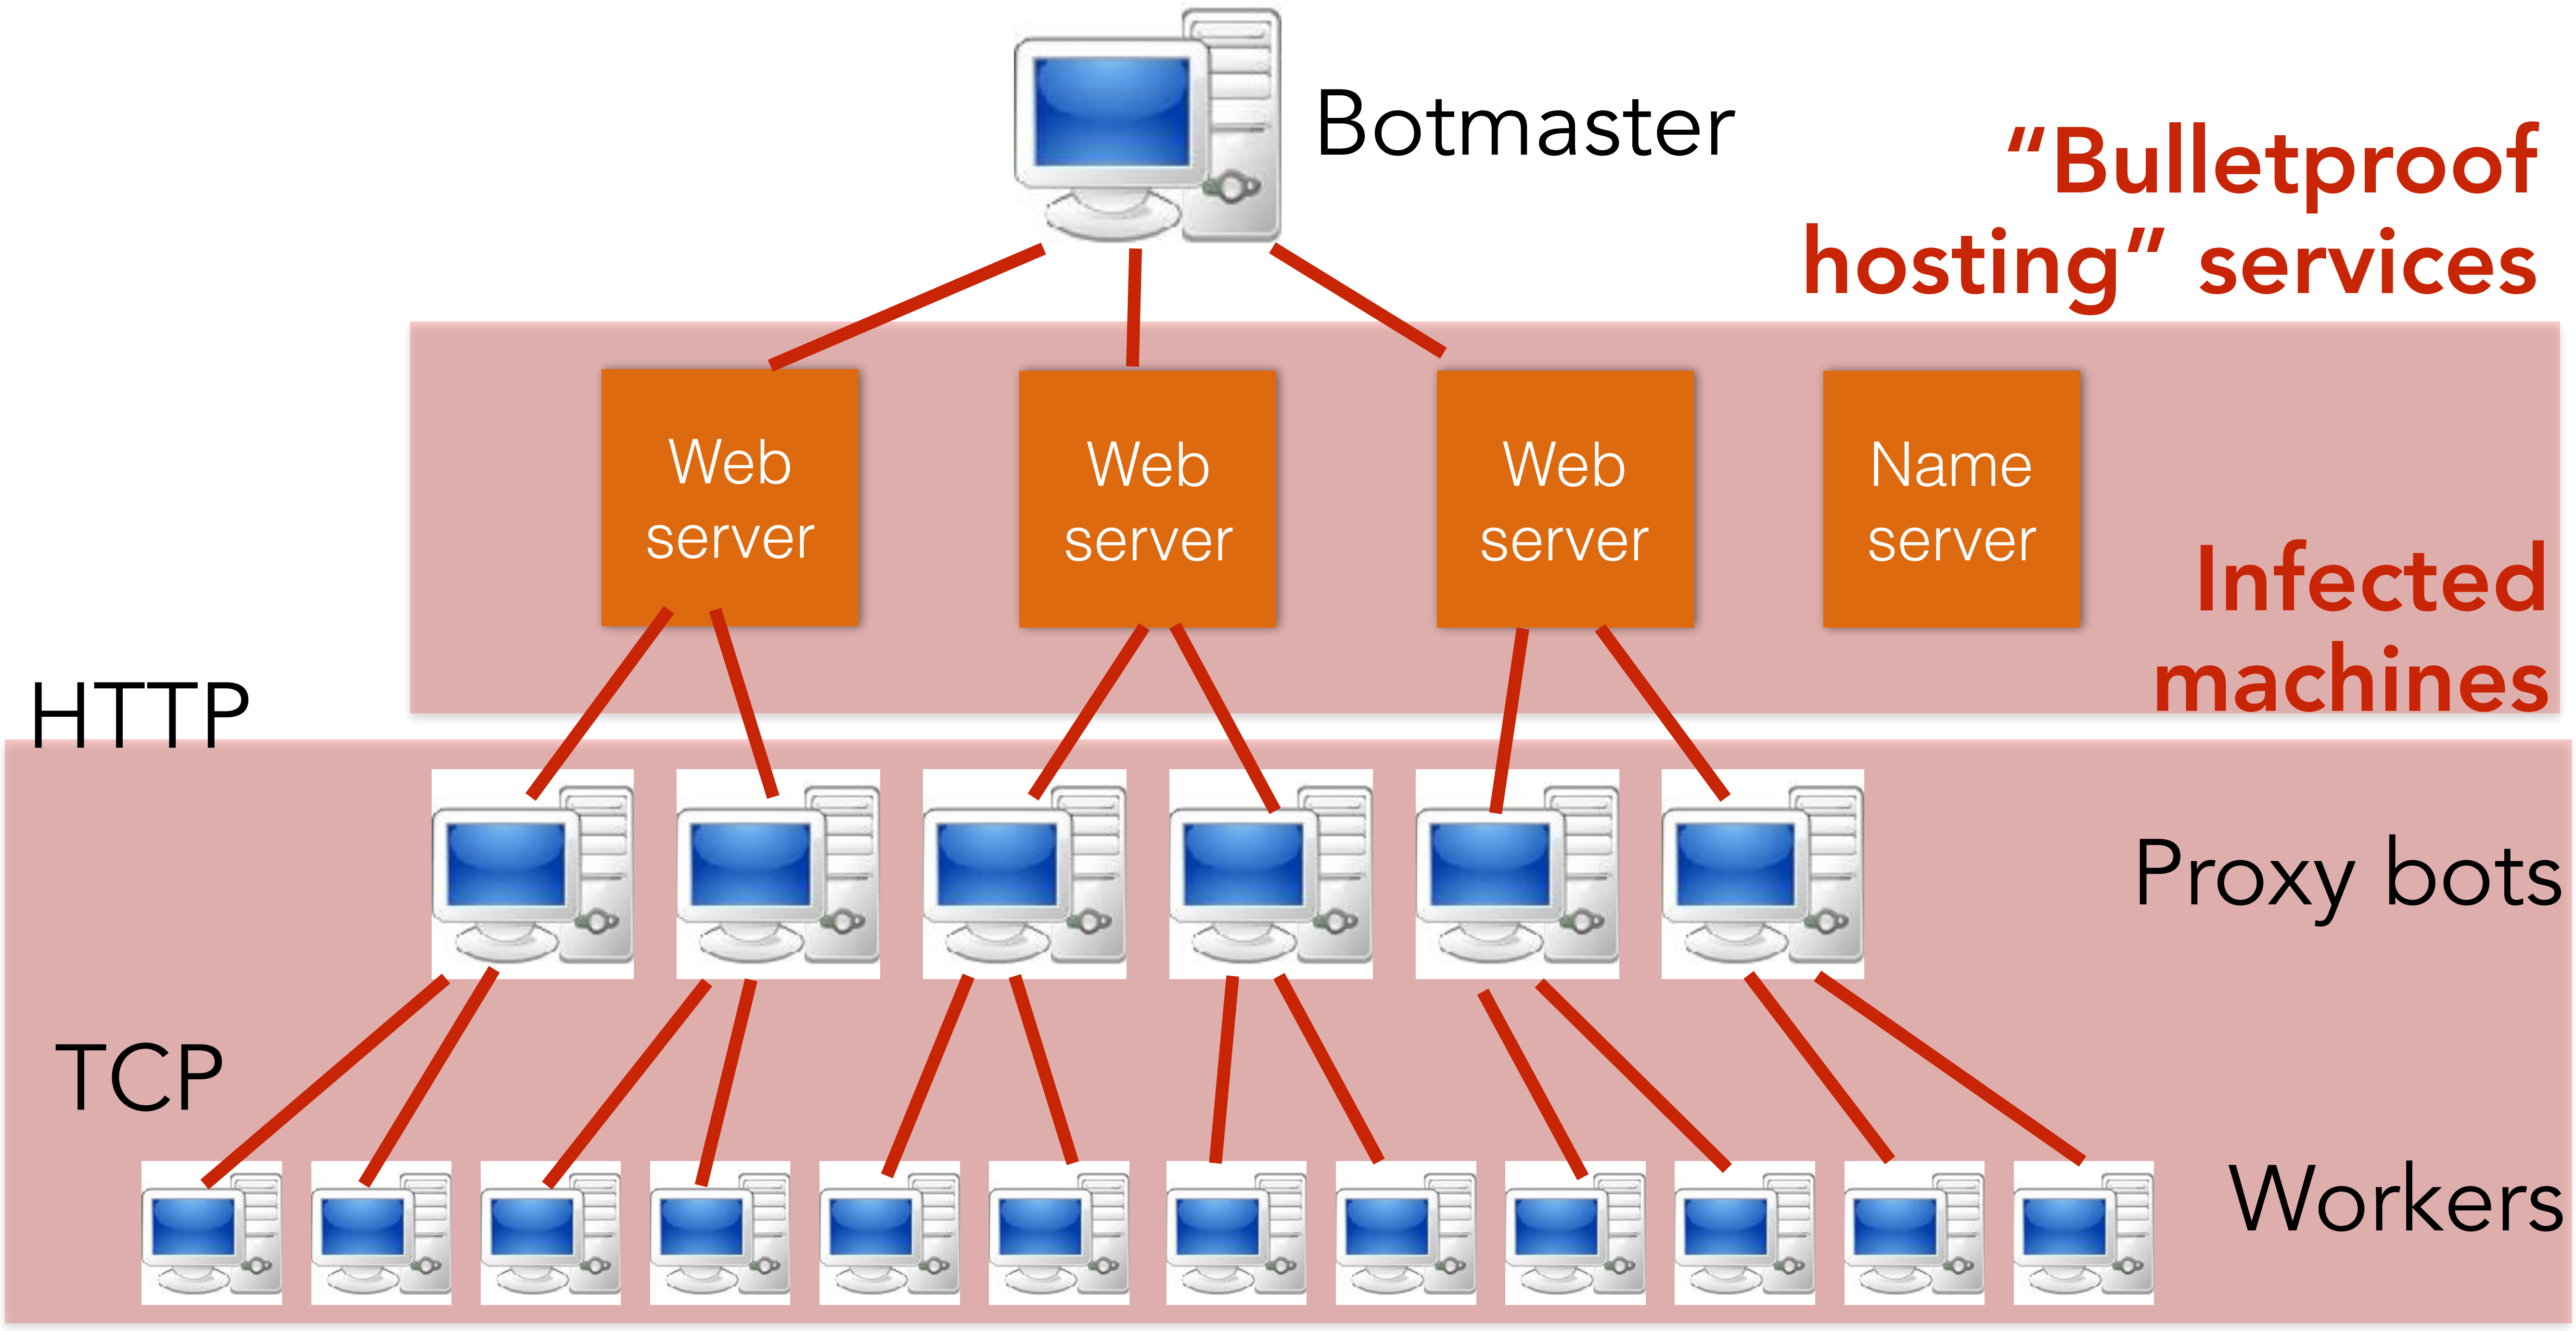

BULLETPROOF HOSTING SERVICES

- Services / specific hosts are often blocked by appealing to their ISPs (“please block this user..”)
- Bulletproof hosting services will refuse to block you (for a price)
- Many have been taken down
  - Often linked to criminal organizations
- Storm botnet: Controller likely run by Russian Business Network
  - Used Atrivo as their bulletproof hosting service

WHY SO MANY LEVELS OF INDIRECTION?

- Many workers send email
- User clicks: gets sent to a proxy bot, who redirects to a web server
- Why proxies?
    - To subvert defenses that block IP addresses
    - Keep the IP address for a given host (buydrugs.ru) moving
- **"Fast flux"** network
    - Short-lived TTLs in DNS responses (hostname to IP address mapping changes quickly)
    - Web proxies to a set of fixed web servers

AN ASIDE ABOUT BOTNETS

MONETIZING BOTNETS

- General malware monetization approaches apply:
  - Keyloggers (steal financial, email, social network, etc. accounts)
  - Ransomware
  - Transaction generators
    - Watch user's surfing
    - Wait to log into banking site and inject extra money, then alter web server replies to mask change in user balance
    - Or wait until the user clicks and inject your own, too.

MONETIZING BOTNETS

- Additionally, botnets give you massive scale
    - DDoS
    - Click fraud
    - Scam infrastructure
        - Hosting web pages (e.g., for phishing)
        - Redirection to evade blacklisting/takedown notices
    - Spam

None of these cause serious pain for the infected user!

Users have little incentive to prevent these

ADVERTISING YOUR BOTNET

How do you advertise the capabilities of your amazing botnet?

Some DNS root servers advertise query volume "see how much attack traffic we can fend off!"

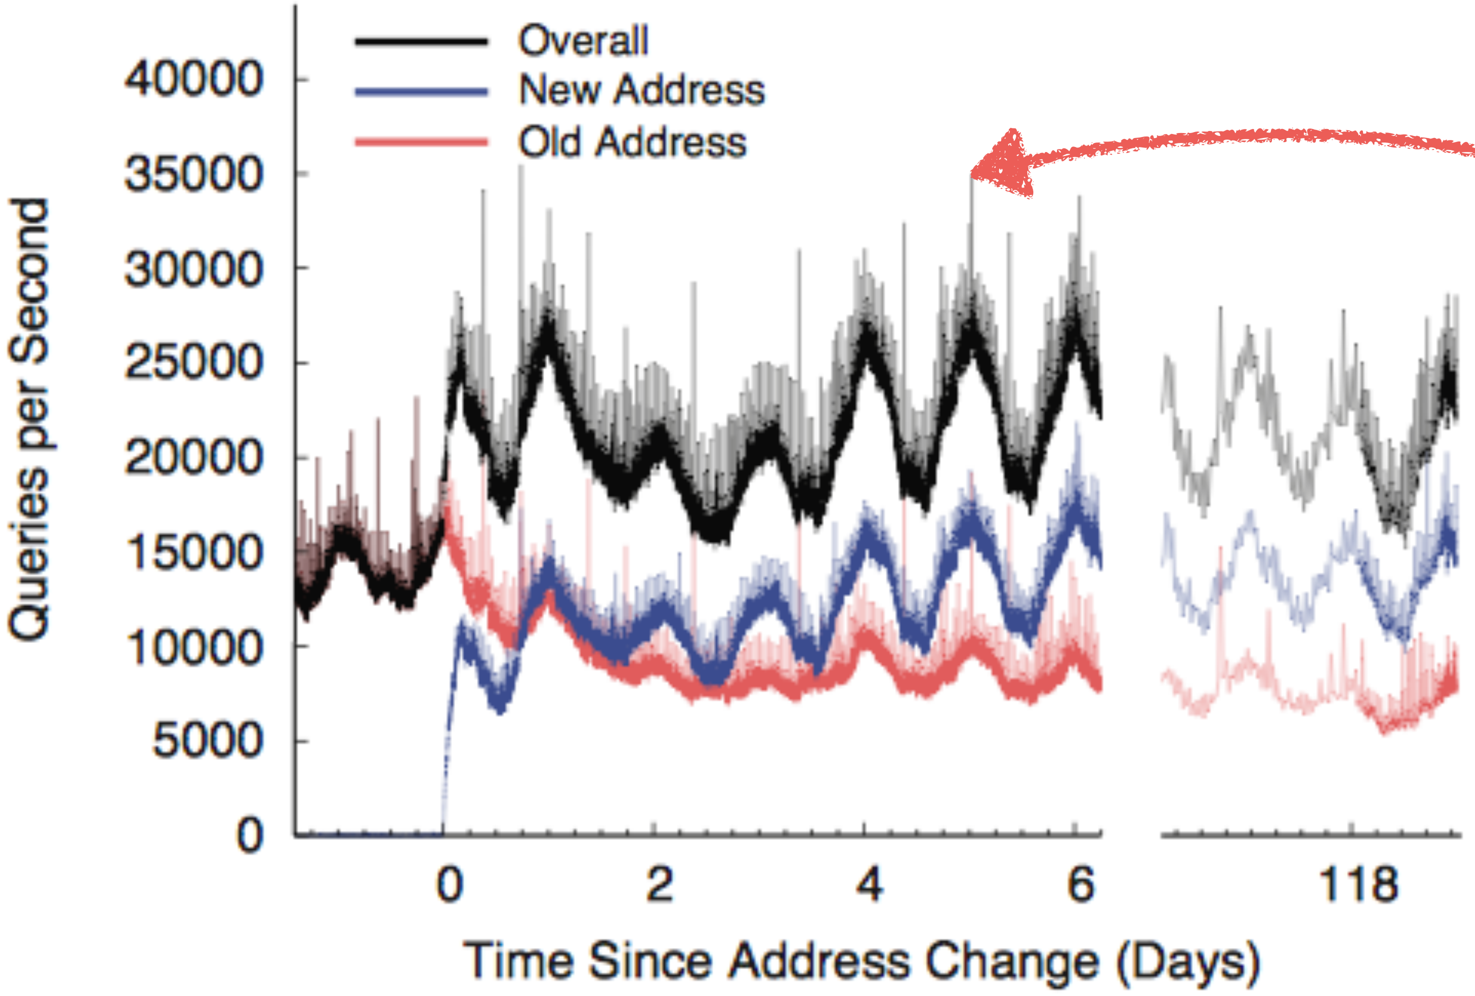

"Look for the surge 4 days from now"

THE IMPORTANCE OF BOTNETS

- Botnets represent the “great modern threat” of the Internet

- Why not worms?
  - Greater control over botnets
  - Less emergent
  - Quieter
  - Flexible

TAKING DOWN BOTNETS

- Approach #1: **prevent** the initial bot infection
  - Infection is decoupled from bot's participation in the botnet, so this is equivalent to preventing malware infections in general - *hard*

- Approach #2: **Take down** the C&C master server

- Botmaster counter-measures?

  - Move the C&C around: each day (e.g.) bots generate a large list of possible domain names.
  - Try a random subset looking for C&C server.
  - Server signs its replies

Counter-counter measure?

BACK TO SPAM

AFFILIATE PROGRAMS

Markets drive efficiency and specialization:
some specialize in botnets, others in spam

- You can join an affiliate program!
  - You send out emails and get a commission (30–50%)

- Affiliate program provides:
  - Storefront templates, shopping cart management
  - Analytics support
  - Advertising materials
  - Central web service interface for affiliates to track conversions and to register for payouts
  - Domains bought in bulk
  - …

GETTING PAID

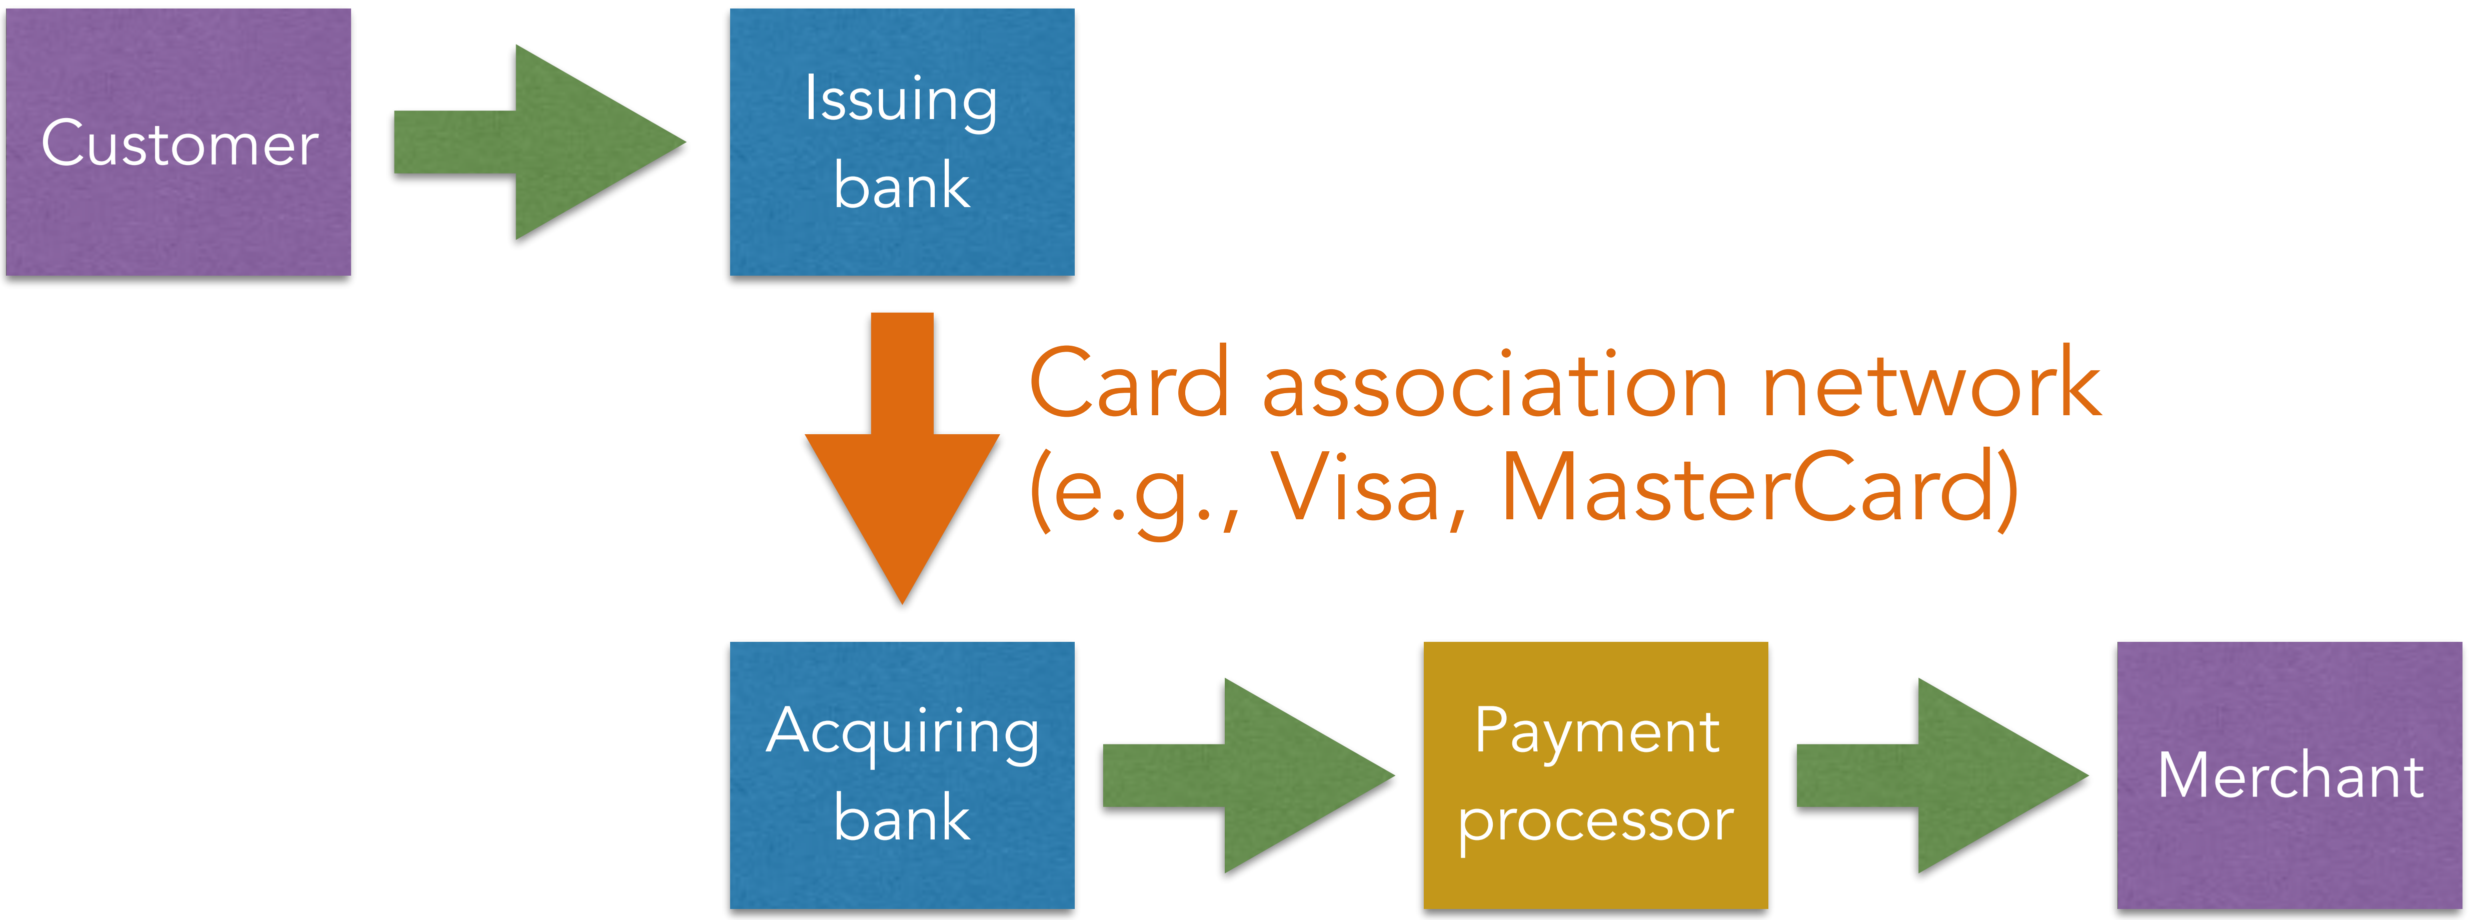

Customer

Issuing bank

Card association network (e.g., Visa, MasterCard)

Acquiring bank

Payment processor

Merchant

Facilitates payment

SHIPPING GOODS

- Business-to-business websites will make connections across many different goods
  - Alibaba, EC-Plaza, ECTrade, …
- Commonly offer “drop shipping”
  - The spambot operator does not need to purchase any warehouse/storage

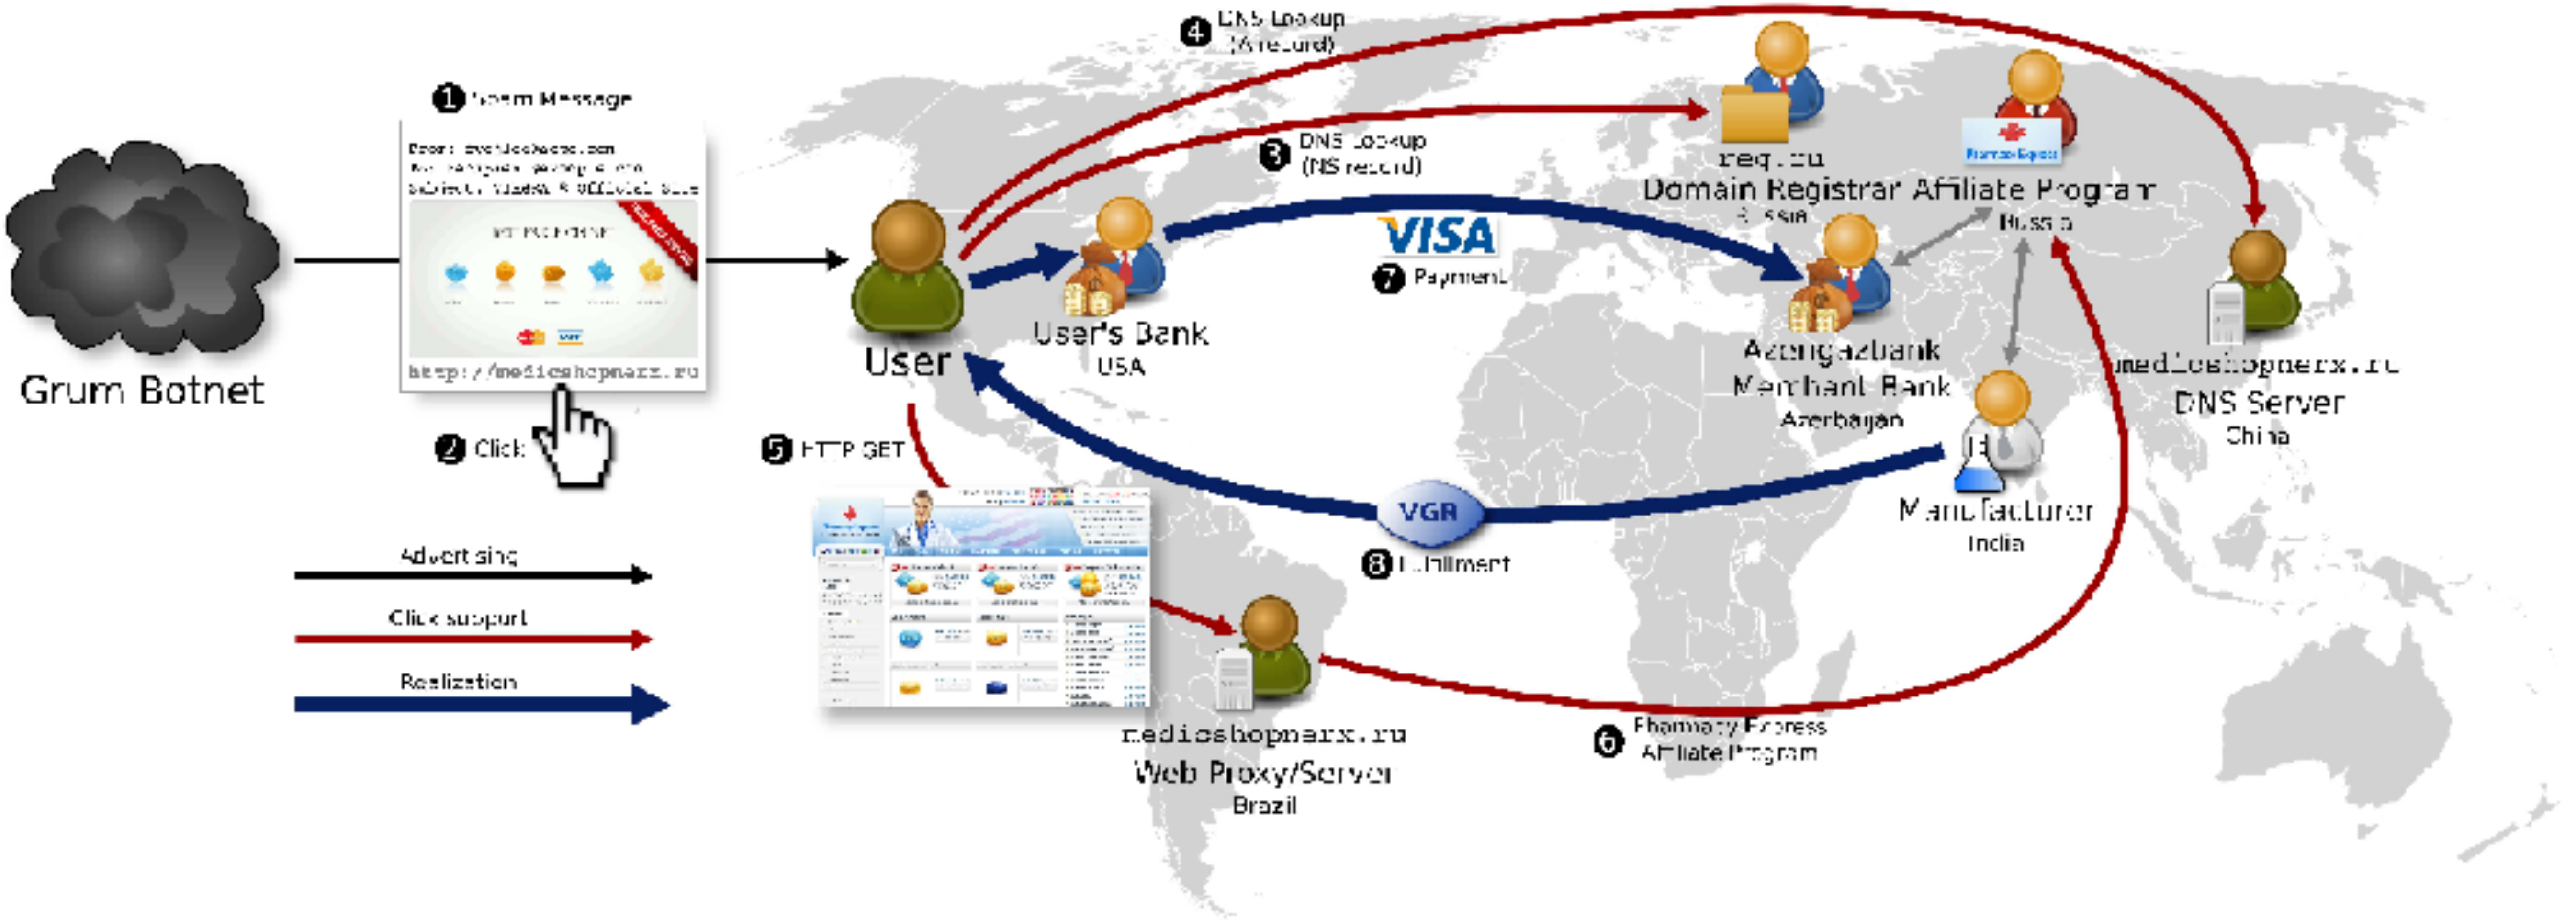

Figure 1: Infrastructure involved in a single URL's value chain, including advertisement, click support and realization steps.

1. Spam delivered
2. User clicks
3. Domain registered by reg.ru
4. Nameserver hosted in China
5. Renders storefront
6. Analytics updated at affiliate
7. User makes payment; acquiring bank in Azerbaijan
8. Supplier in Chennai, India delivers 10 days later

ANALYZING SPAM CLICK TRAJECTORIES

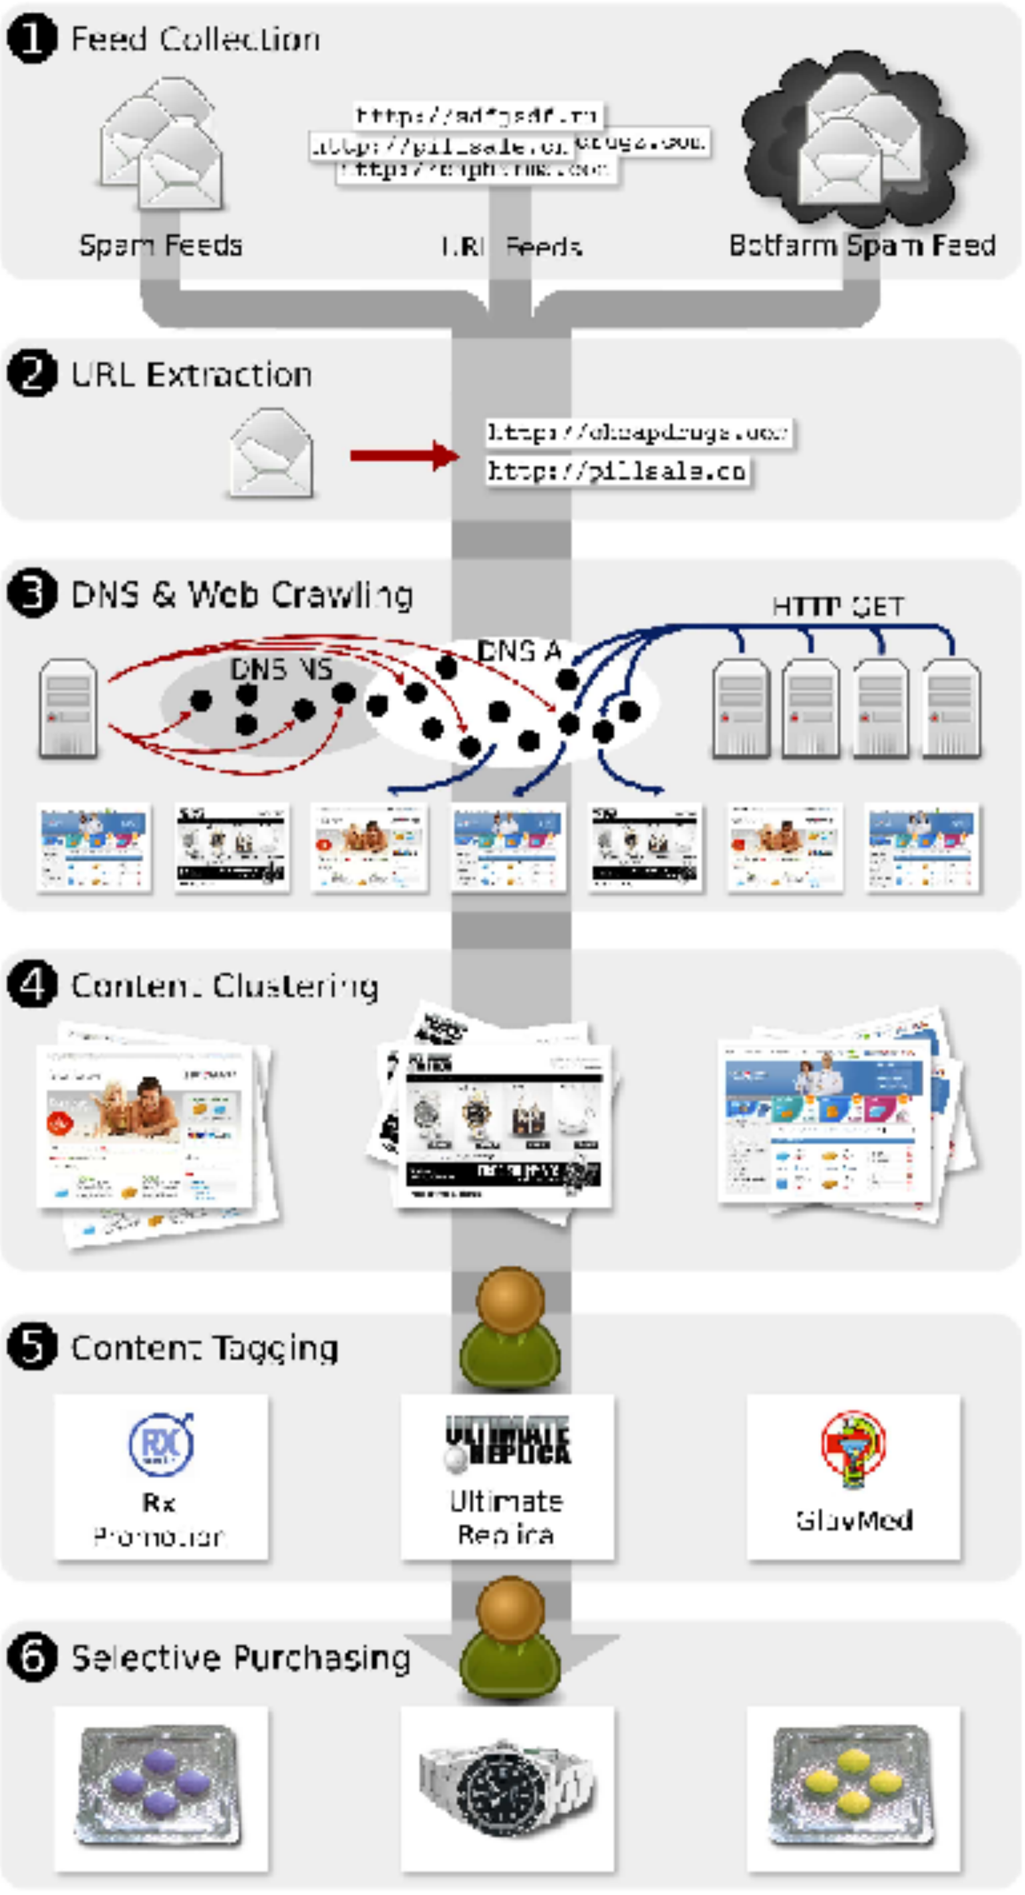

Figure 2: Our data collection and processing workflow.

Dataset

156 orders over 2 months

PURCHASE PAIRS

- Most *affiliate programs* provide a confirmation page with an order number

- This order number usually just increments

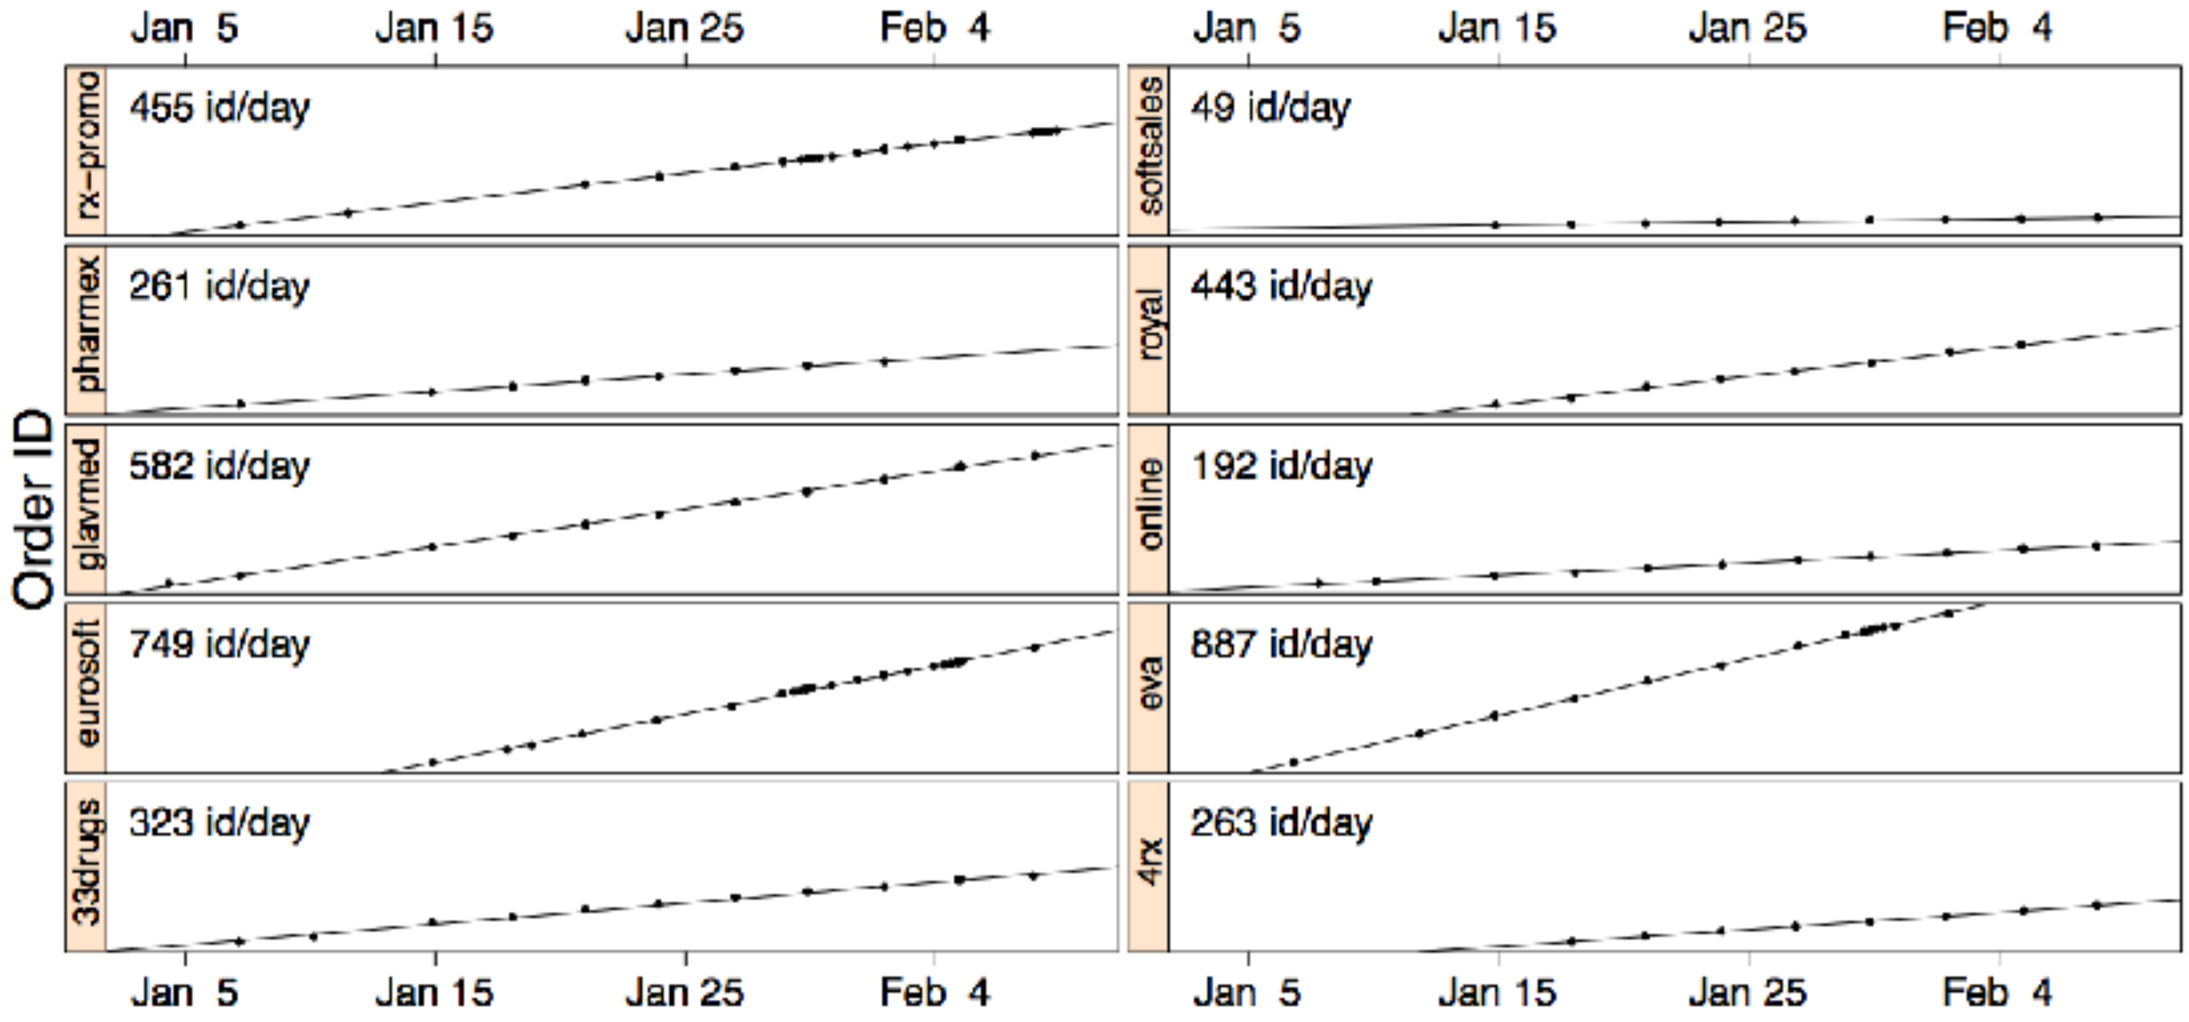

PURCHASE PAIRS

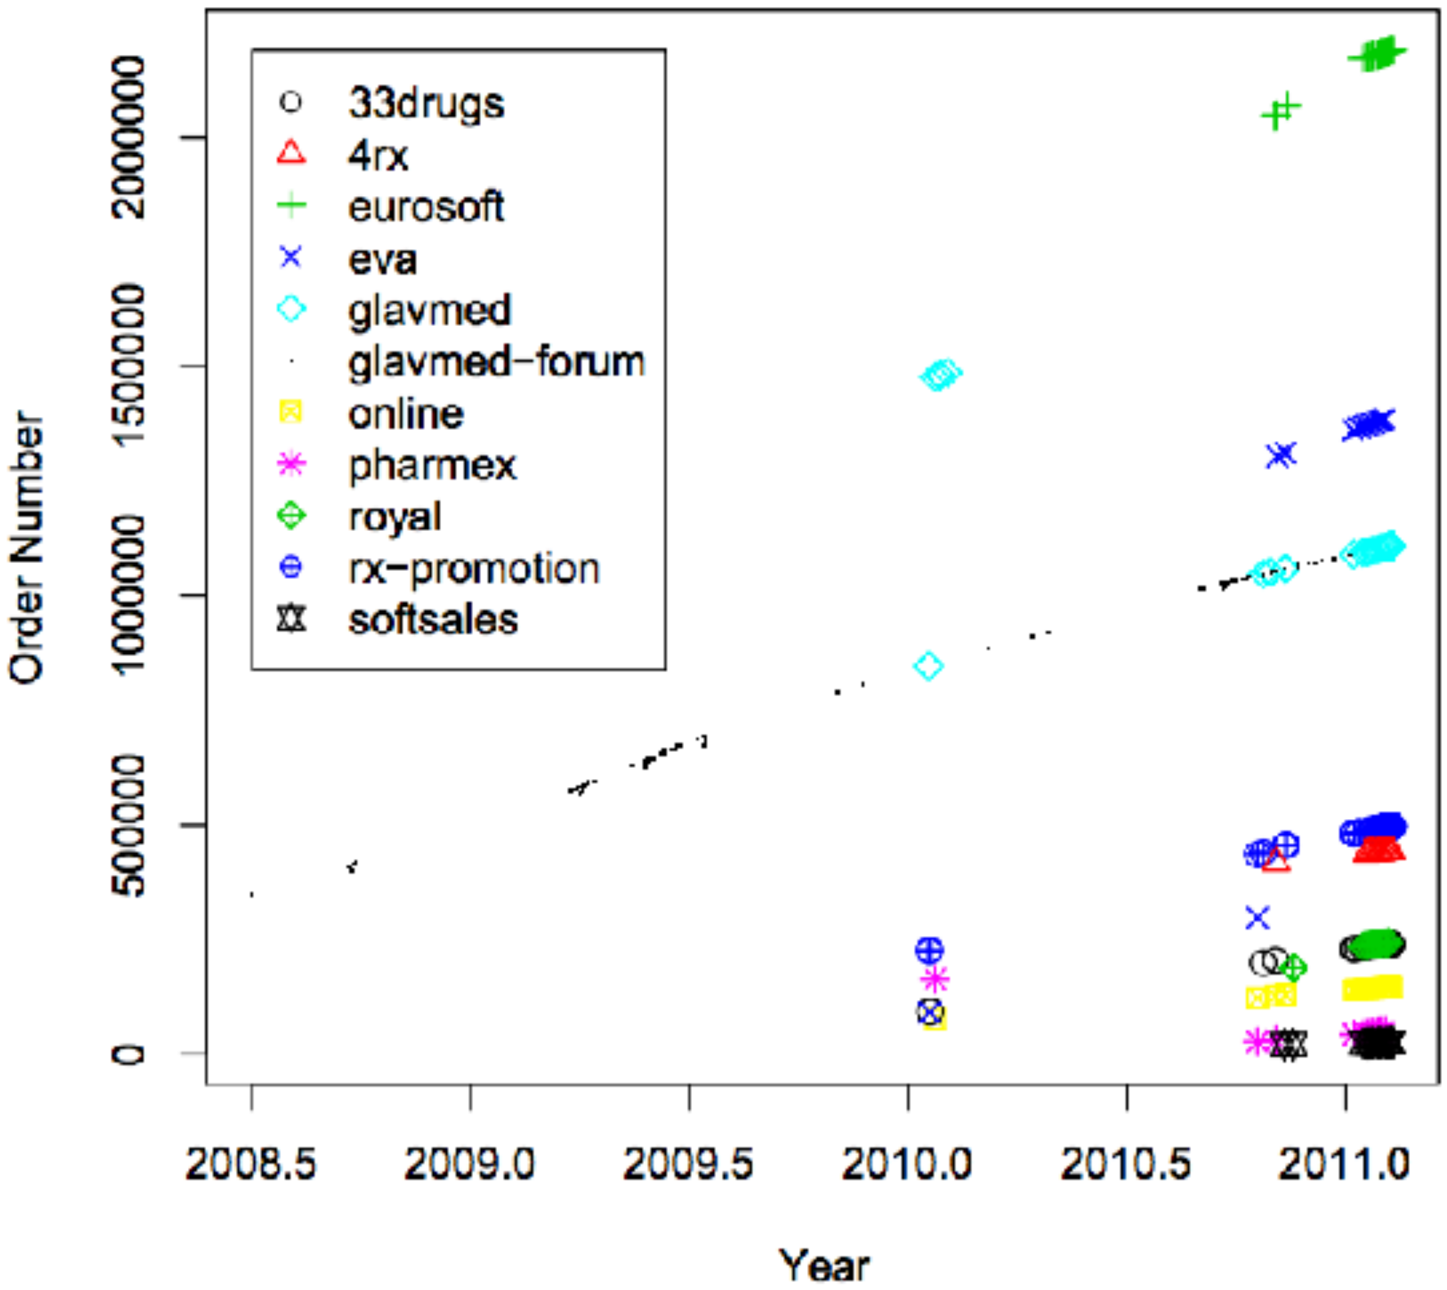

Figure 2: Order numbers (y-axis) associated with each affiliate program versus the time of attempted purchase (x-axis).

INFERRING WHAT PEOPLE BUY

- EvaPharmacy (a top 5 spam-advertised pharmacy affiliate program):
  - 2/3 of outsourced image hosting was to compromised 3rd party servers
- They contacted the owners of these servers and asked for logs
- Correlated image logs with purchases

METHODOLOGICAL SHORTCOMINGS

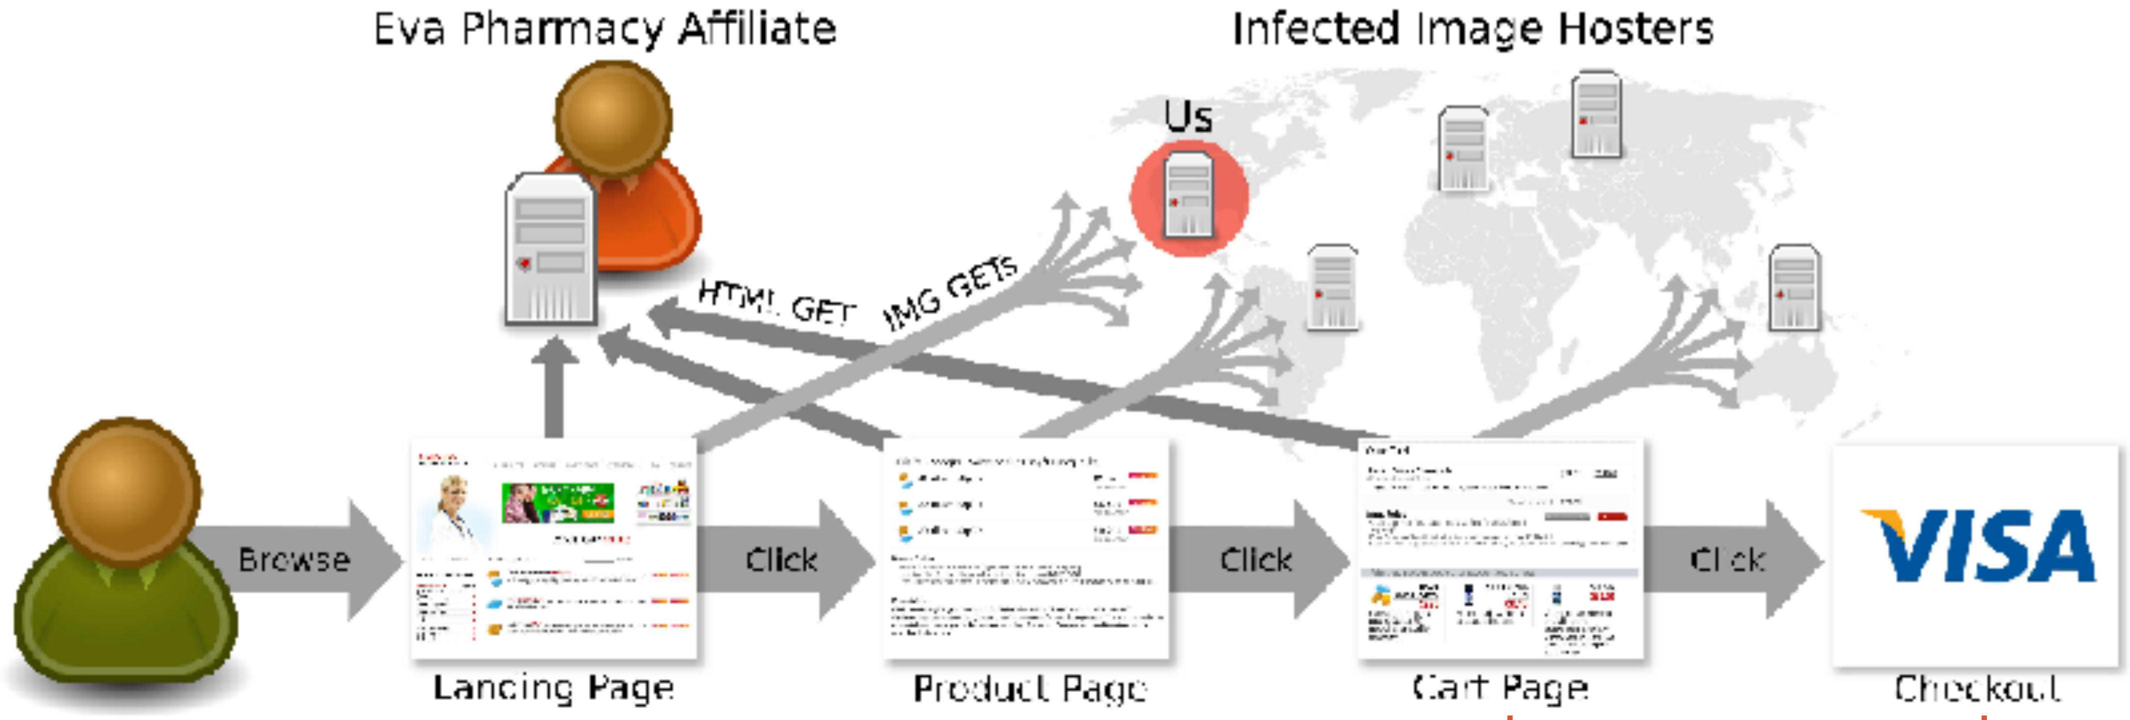

2. Images often independent of dosage/count

(cannot infer exact amount)

1. Checkout page does not include unique images

(can only infer it was in cart)

3. Not all affiliates sell the same formularies

(EvaPharmacy study limited)

4. Almost all visitors from spam email

(potential bias in behavior?)

WHO/WHAT GETS SOLD

- Three most common products sold:
  - Pharmaceuticals (vast majority)
  - Replica luxury goods
  - Counterfeit software

- Run by relatively few affiliate programs

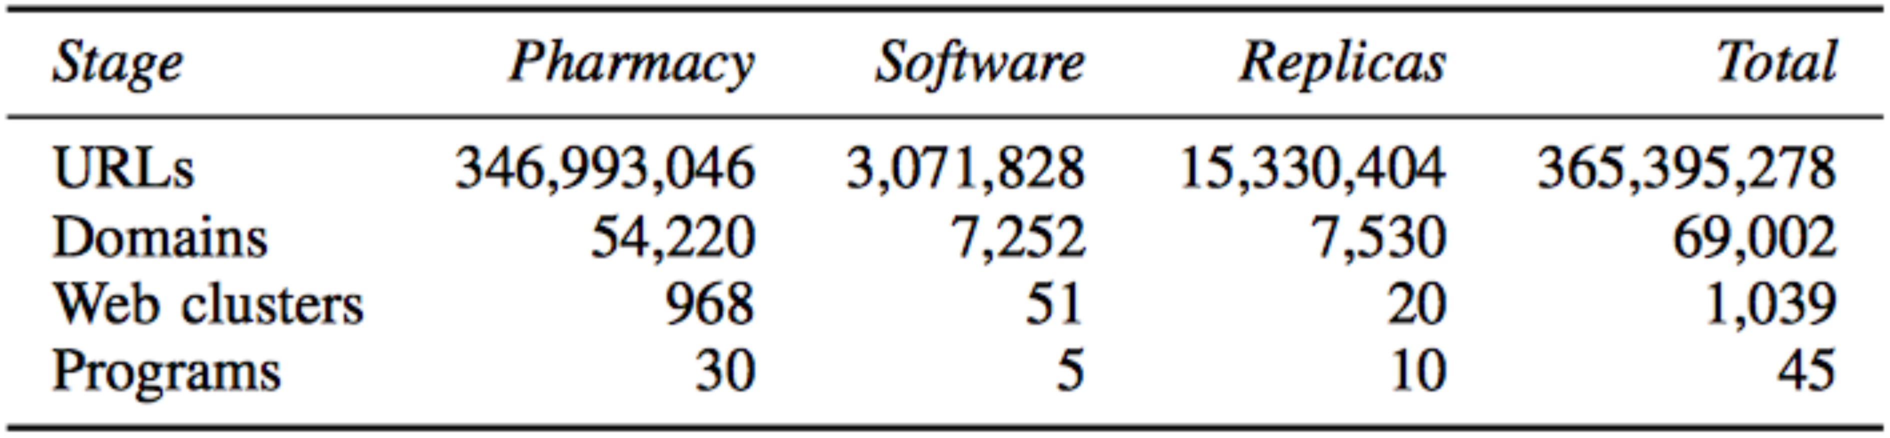

Stage	*Pharmacy*	*Software*	*Replicas*	*Total*
URLs	346,993,046	3,071,828	15,330,404	365,395,278
Domains	54,220	7,252	7,530	69,002
Web clusters	968	51	20	1,039
Programs	30	5	10	45

Table III: Breakdown of clustering and tagging results.

FEW AFFILIATE PROGRAMS CONSTITUTE THE MAJORITY

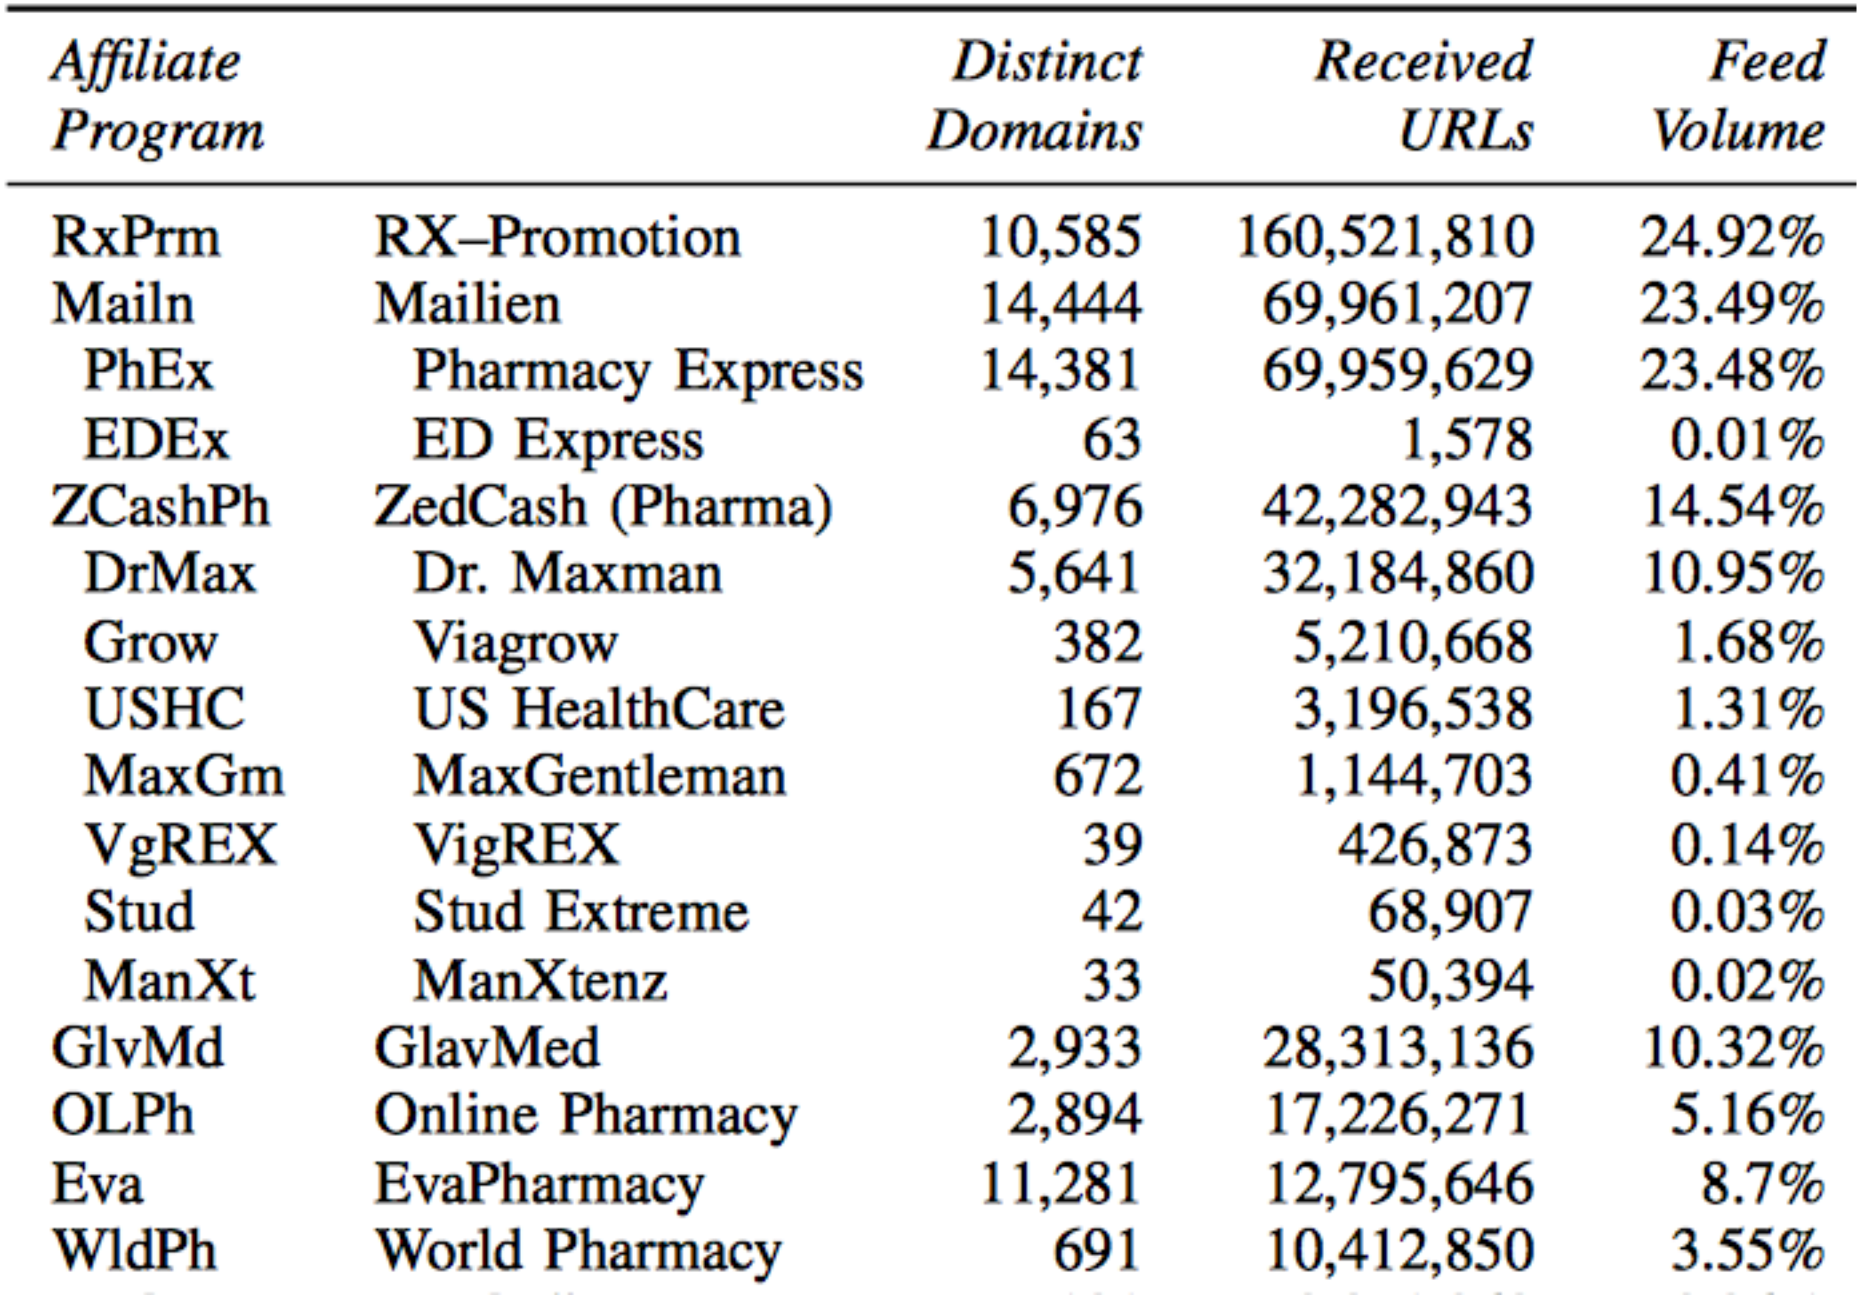

Affiliate Program		*Distinct Domains*	*Received URLs*	*Feed Volume*
RxPrm	RX–Promotion	10,585	160,521,810	24.92%
Mailn	Mailien	14,444	69,961,207	23.49%
PhEx	Pharmacy Express	14,381	69,959,629	23.48%
EDEx	ED Express	63	1,578	0.01%
ZCashPh	ZedCash (Pharma)	6,976	42,282,943	14.54%
DrMax	Dr. Maxman	5,641	32,184,860	10.95%
Grow	Viagrow	382	5,210,668	1.68%
USHC	US HealthCare	167	3,196,538	1.31%
MaxGm	MaxGentleman	672	1,144,703	0.41%
VgREX	VigREX	39	426,873	0.14%
Stud	Stud Extreme	42	68,907	0.03%
ManXt	ManXtenz	33	50,394	0.02%
GlvMd	GlavMed	2,933	28,313,136	10.32%
OLPh	Online Pharmacy	2,894	17,226,271	5.16%
Eva	EvaPharmacy	11,281	12,795,646	8.7%
WldPh	World Pharmacy	691	10,412,850	3.55%

WHAT GETS SOLD

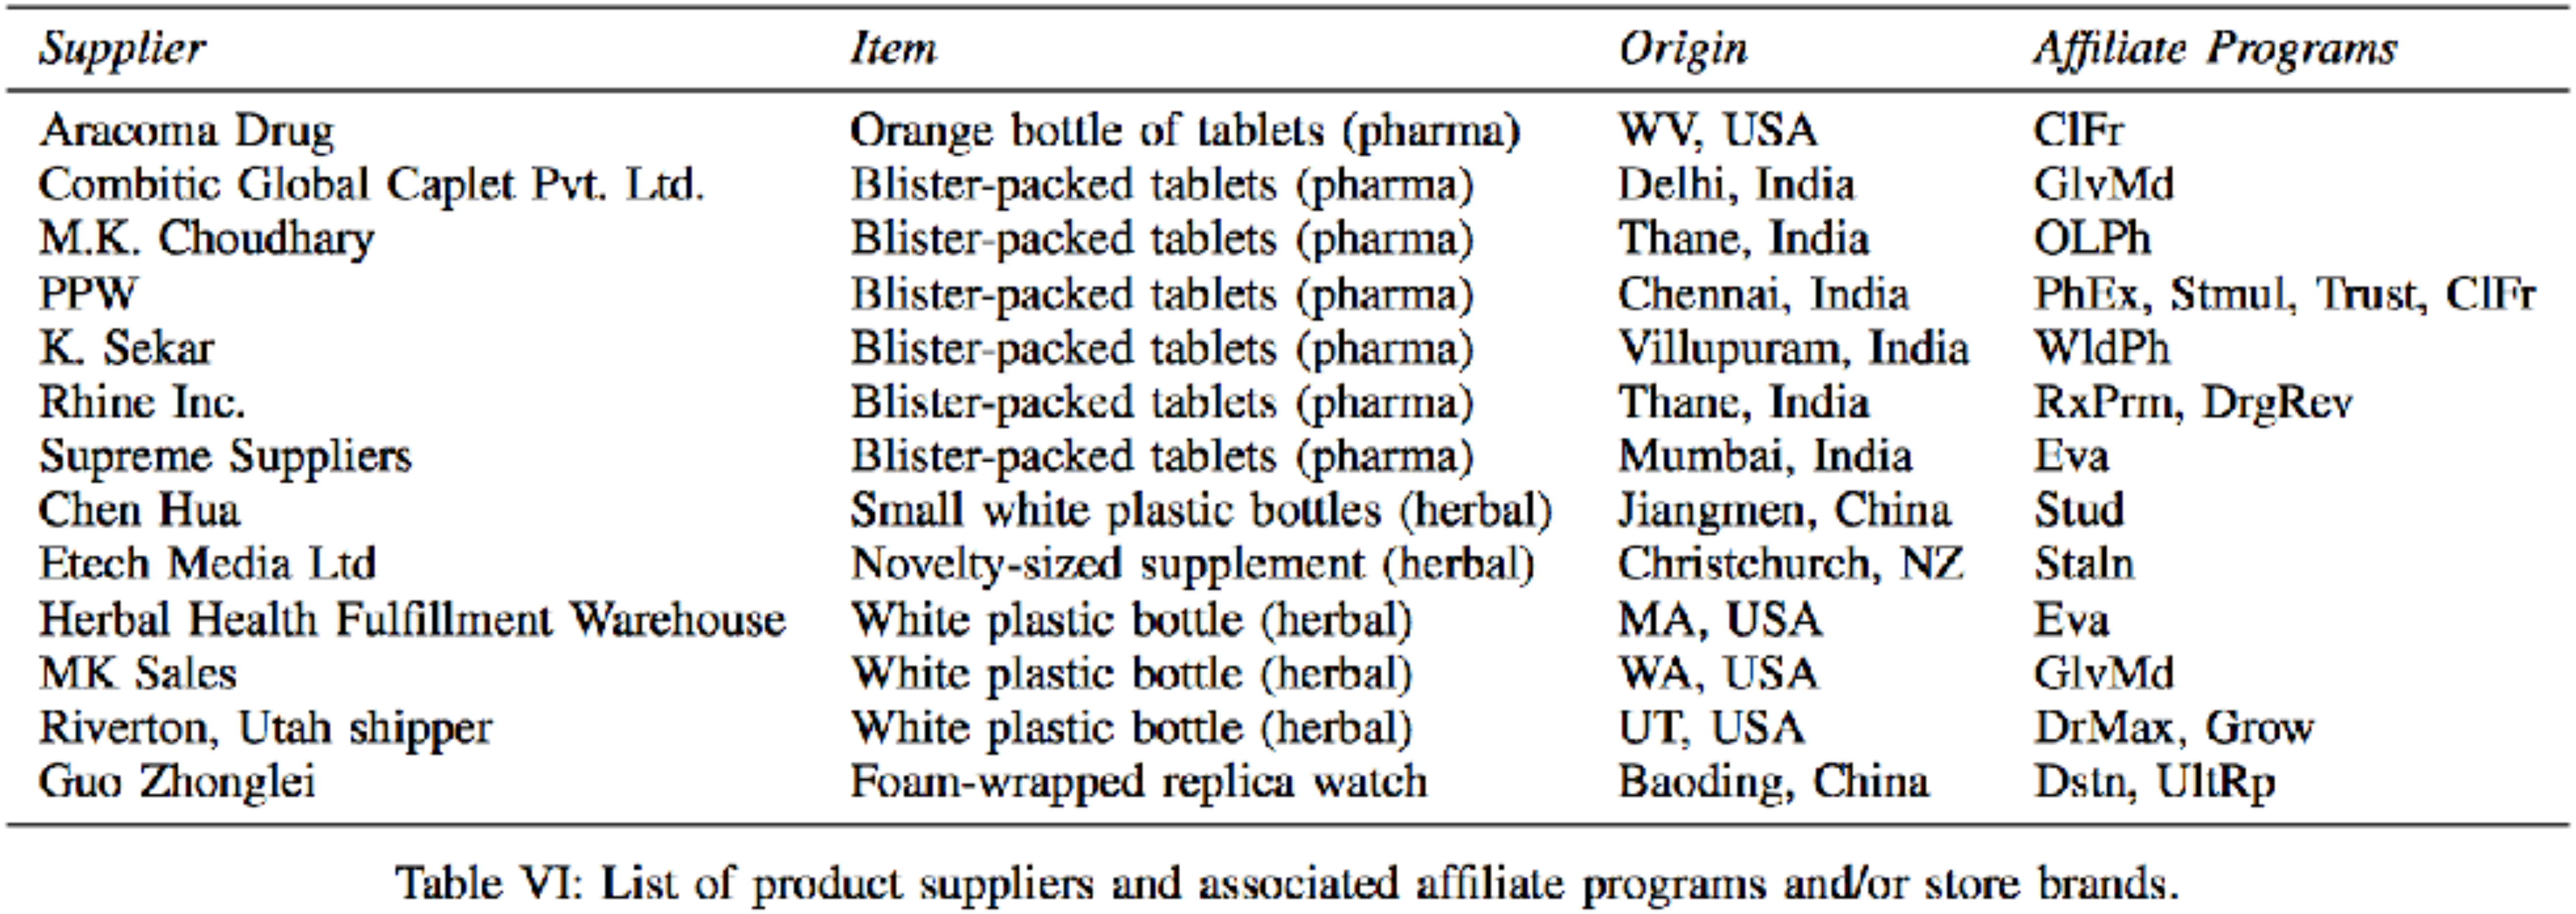

Supplier	*Item*	*Origin*	*Affiliate Programs*
Aracoma Drug	Orange bottle of tablets (pharma)	WV, USA	ClFr
Combitic Global Caplet Pvt. Ltd.	Blister-packed tablets (pharma)	Delhi, India	GlvMd
M.K. Choudhary	Blister-packed tablets (pharma)	Thane, India	OLPh
PPW	Blister-packed tablets (pharma)	Chennai, India	PhEx, Stmul, Trust, ClFr
K. Sekar	Blister-packed tablets (pharma)	Villupuram, India	WldPh
Rhine Inc.	Blister-packed tablets (pharma)	Thane, India	RxPrm, DrgRev
Supreme Suppliers	Blister-packed tablets (pharma)	Mumbai, India	Eva
Chen Hua	Small white plastic bottles (herbal)	Jiangmen, China	Stud
Etech Media Ltd	Novelty-sized supplement (herbal)	Christchurch, NZ	Staln
Herbal Health Fulfillment Warehouse	White plastic bottle (herbal)	MA, USA	Eva
MK Sales	White plastic bottle (herbal)	WA, USA	GlvMd
Riverton, Utah shipper	White plastic bottle (herbal)	UT, USA	DrMax, Grow
Guo Zhonglei	Foam-wrapped replica watch	Baoding, China	Dstn, UltRp

Table VI: List of product suppliers and associated affiliate programs and/or store brands.

ACQUIRING BANKS

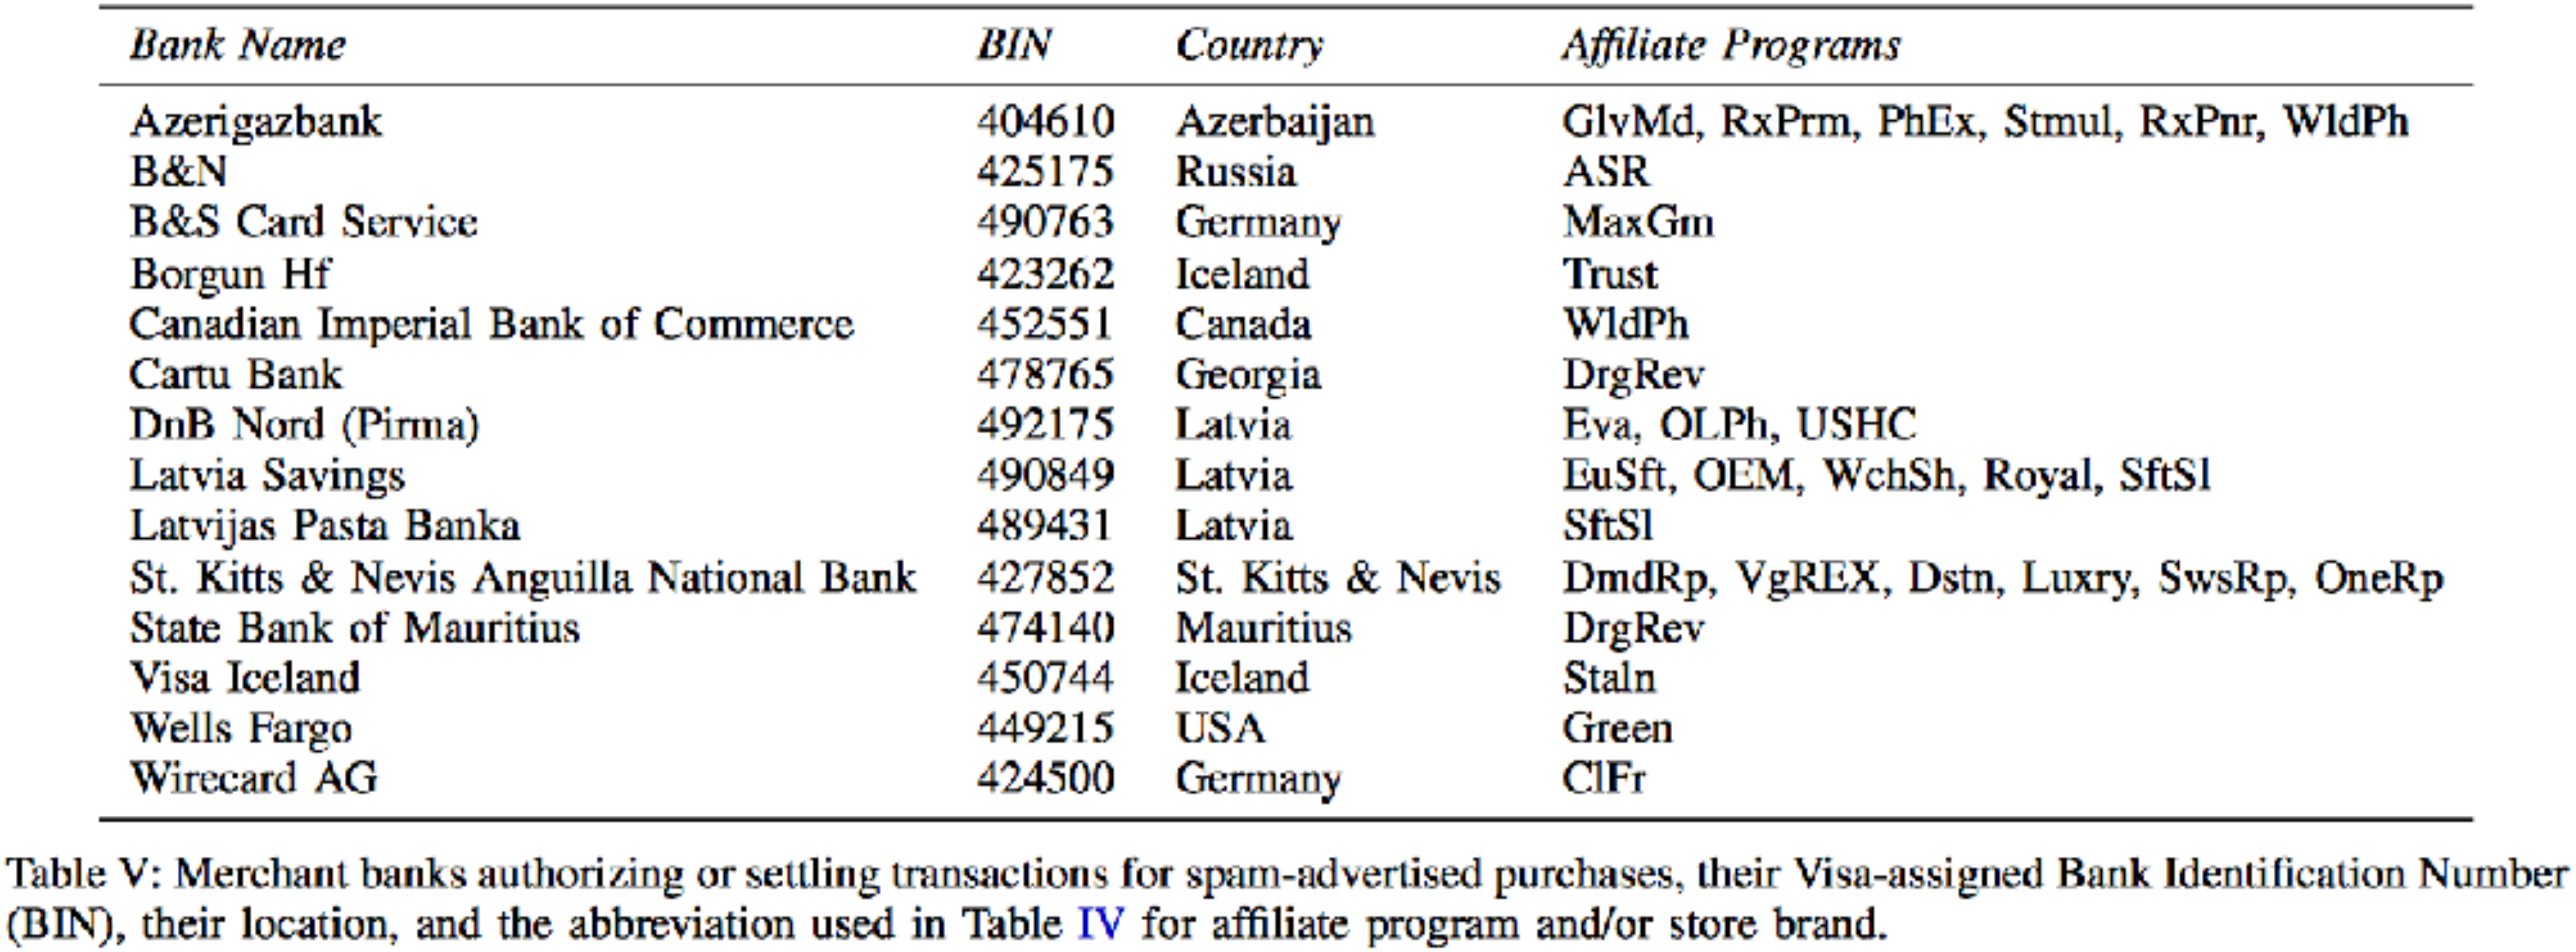

Bank Name	*BIN*	*Country*	*Affiliate Programs*
Azerigazbank	404610	Azerbaijan	GlvMd, RxPrm, PhEx, Stmul, RxPnr, WldPh
B&N	425175	Russia	ASR
B&S Card Service	490763	Germany	MaxGm
Borgun Hf	423262	Iceland	Trust
Canadian Imperial Bank of Commerce	452551	Canada	WldPh
Cartu Bank	478765	Georgia	DrgRev
DnB Nord (Pirma)	492175	Latvia	Eva, OLPh, USHC
Latvia Savings	490849	Latvia	EuSft, OEM, WchSh, Royal, SftSl
Latvijas Pasta Banka	489431	Latvia	SftSl
St. Kitts & Nevis Anguilla National Bank	427852	St. Kitts & Nevis	DmdRp, VgREX, Dstn, Luxry, SwsRp, OneRp
State Bank of Mauritius	474140	Mauritius	DrgRev
Visa Iceland	450744	Iceland	Staln
Wells Fargo	449215	USA	Green
Wirecard AG	424500	Germany	ClFr

Table V: Merchant banks authorizing or settling transactions for spam-advertised purchases, their Visa-assigned Bank Identification Number (BIN), their location, and the abbreviation used in Table IV for affiliate program and/or store brand.

SO HOW MUCH ARE SPAMBOTS MAKING?

- To understand, we would have to know:
  - Order volume (how much is sold as a result of an affiliate program over time?)
  - Purchasing behavior (what are people buying?)
- Prior understanding was vague at best

AFFILIATE PROFIT

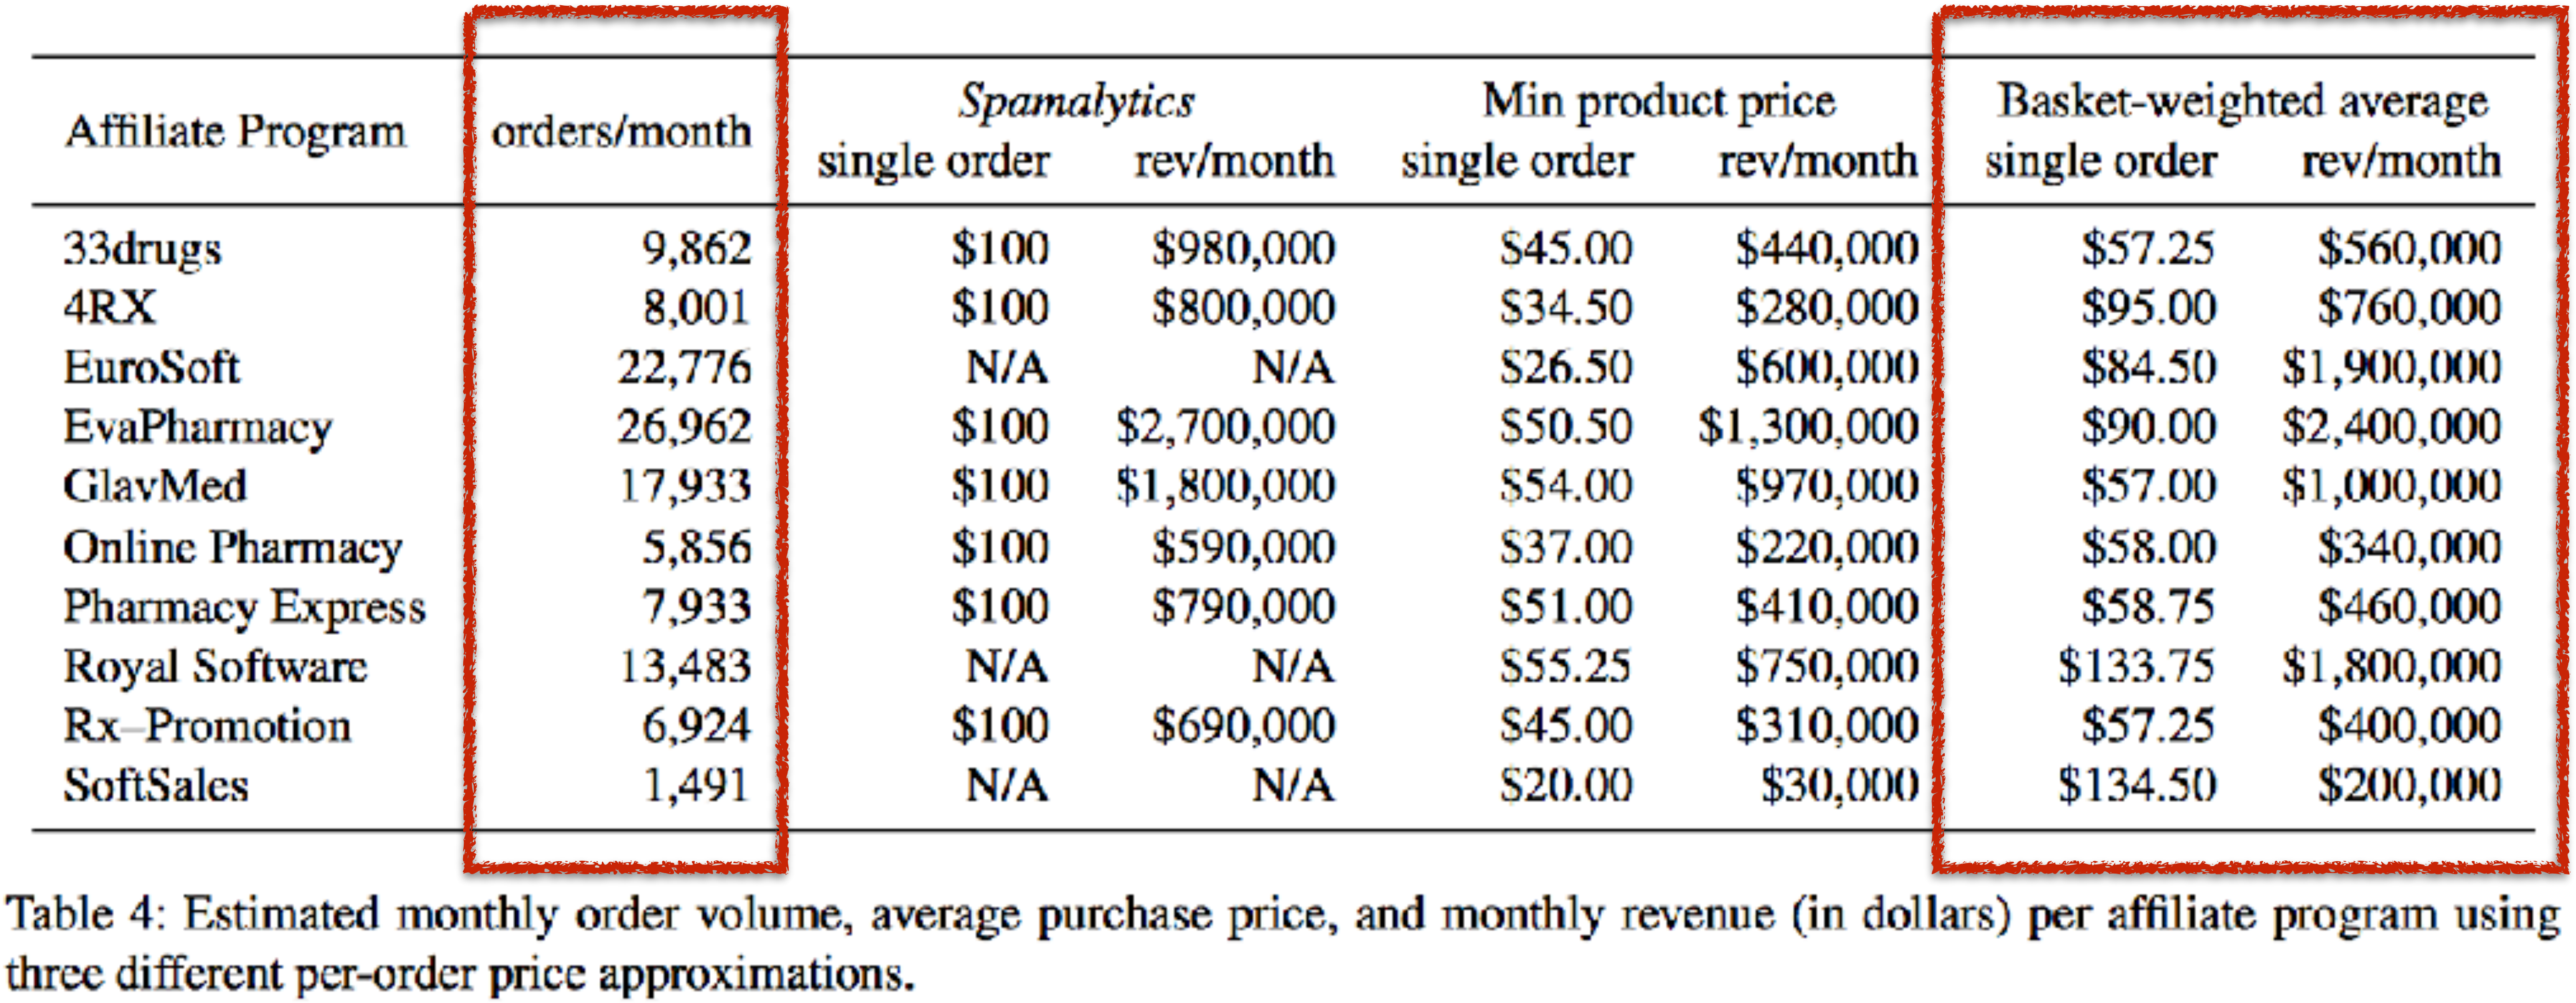

Affiliate Program	orders/month	*Spamalytics*		Min product price		Basket-weighted average	
		single order	rev/month	single order	rev/month	single order	rev/month
33drugs	9,862	$100	$980,000	$45.00	$440,000	$57.25	$560,000
4RX	8,001	$100	$800,000	$34.50	$280,000	$95.00	$760,000
EuroSoft	22,776	N/A	N/A	$26.50	$600,000	$84.50	$1,900,000
EvaPharmacy	26,962	$100	$2,700,000	$50.50	$1,300,000	$90.00	$2,400,000
GlavMed	17,933	$100	$1,800,000	$54.00	$970,000	$57.00	$1,000,000
Online Pharmacy	5,856	$100	$590,000	$37.00	$220,000	$58.00	$340,000
Pharmacy Express	7,933	$100	$790,000	$51.00	$410,000	$58.75	$460,000
Royal Software	13,483	N/A	N/A	$55.25	$750,000	$133.75	$1,800,000
Rx–Promotion	6,924	$100	$690,000	$45.00	$310,000	$57.25	$400,000
SoftSales	1,491	N/A	N/A	$20.00	$30,000	$134.50	$200,000

Table 4: Estimated monthly order volume, average purchase price, and monthly revenue (in dollars) per affiliate program using three different per-order price approximations.

Over 100k orders/month
in this dataset alone

Some have guessed that
"spammers make little
money at all"

So who's actually buying this junk?

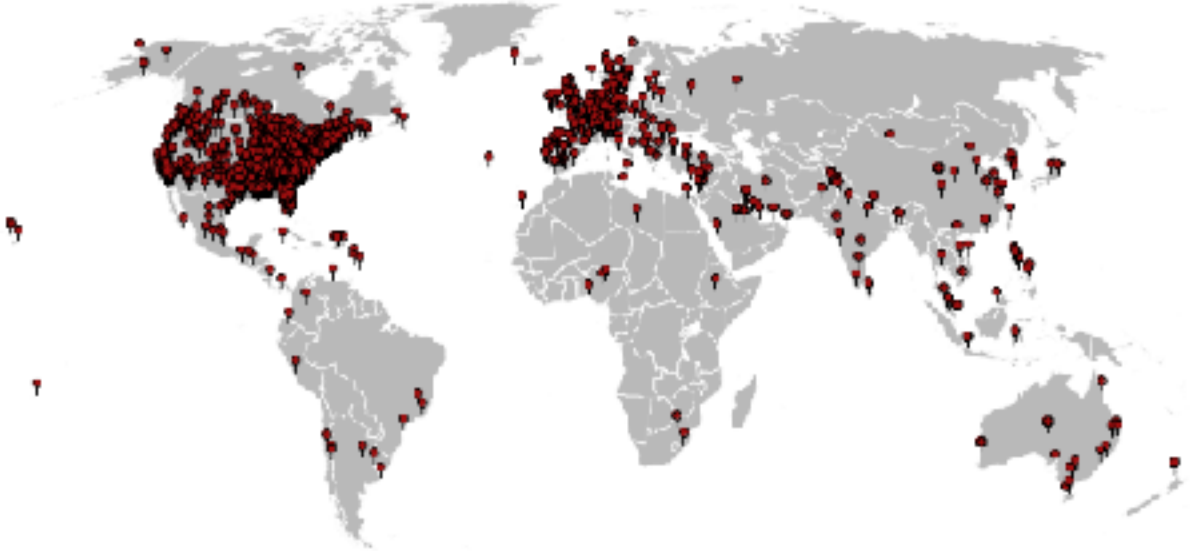

Figure 6: The geographic distribution of those who added an item to their shopping cart.

Stop buying this junk!

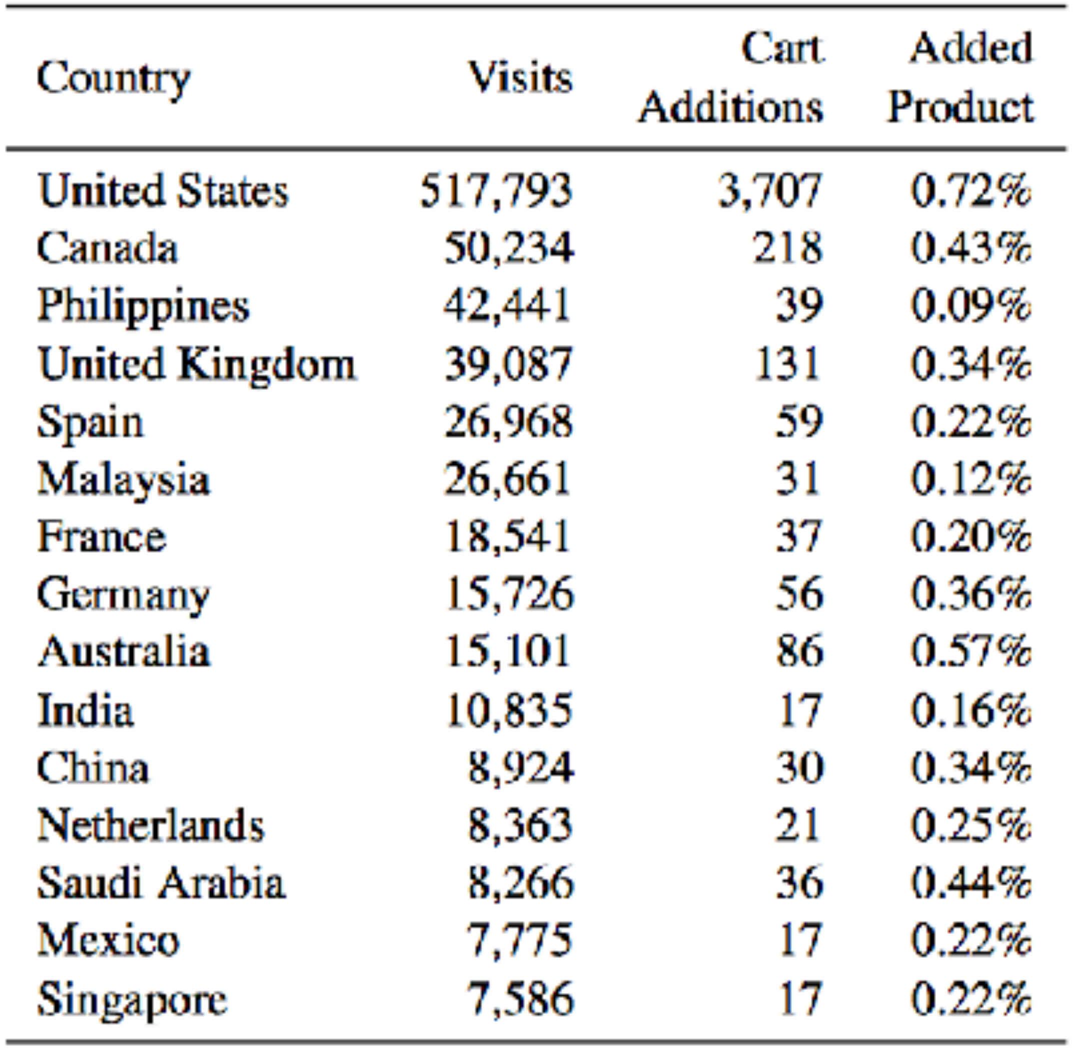

Country	Visits	Cart Additions	Added Product
United States	517,793	3,707	0.72%
Canada	50,234	218	0.43%
Philippines	42,441	39	0.09%
United Kingdom	39,087	131	0.34%
Spain	26,968	59	0.22%
Malaysia	26,661	31	0.12%
France	18,541	37	0.20%
Germany	15,726	56	0.36%
Australia	15,101	86	0.57%
India	10,835	17	0.16%
China	8,924	30	0.34%
Netherlands	8,363	21	0.25%
Saudi Arabia	8,266	36	0.44%
Mexico	7,775	17	0.22%
Singapore	7,586	17	0.22%

Table 2: The top 15 countries and the percentage of visitors who added an item to their shopping cart.

What are you buying?

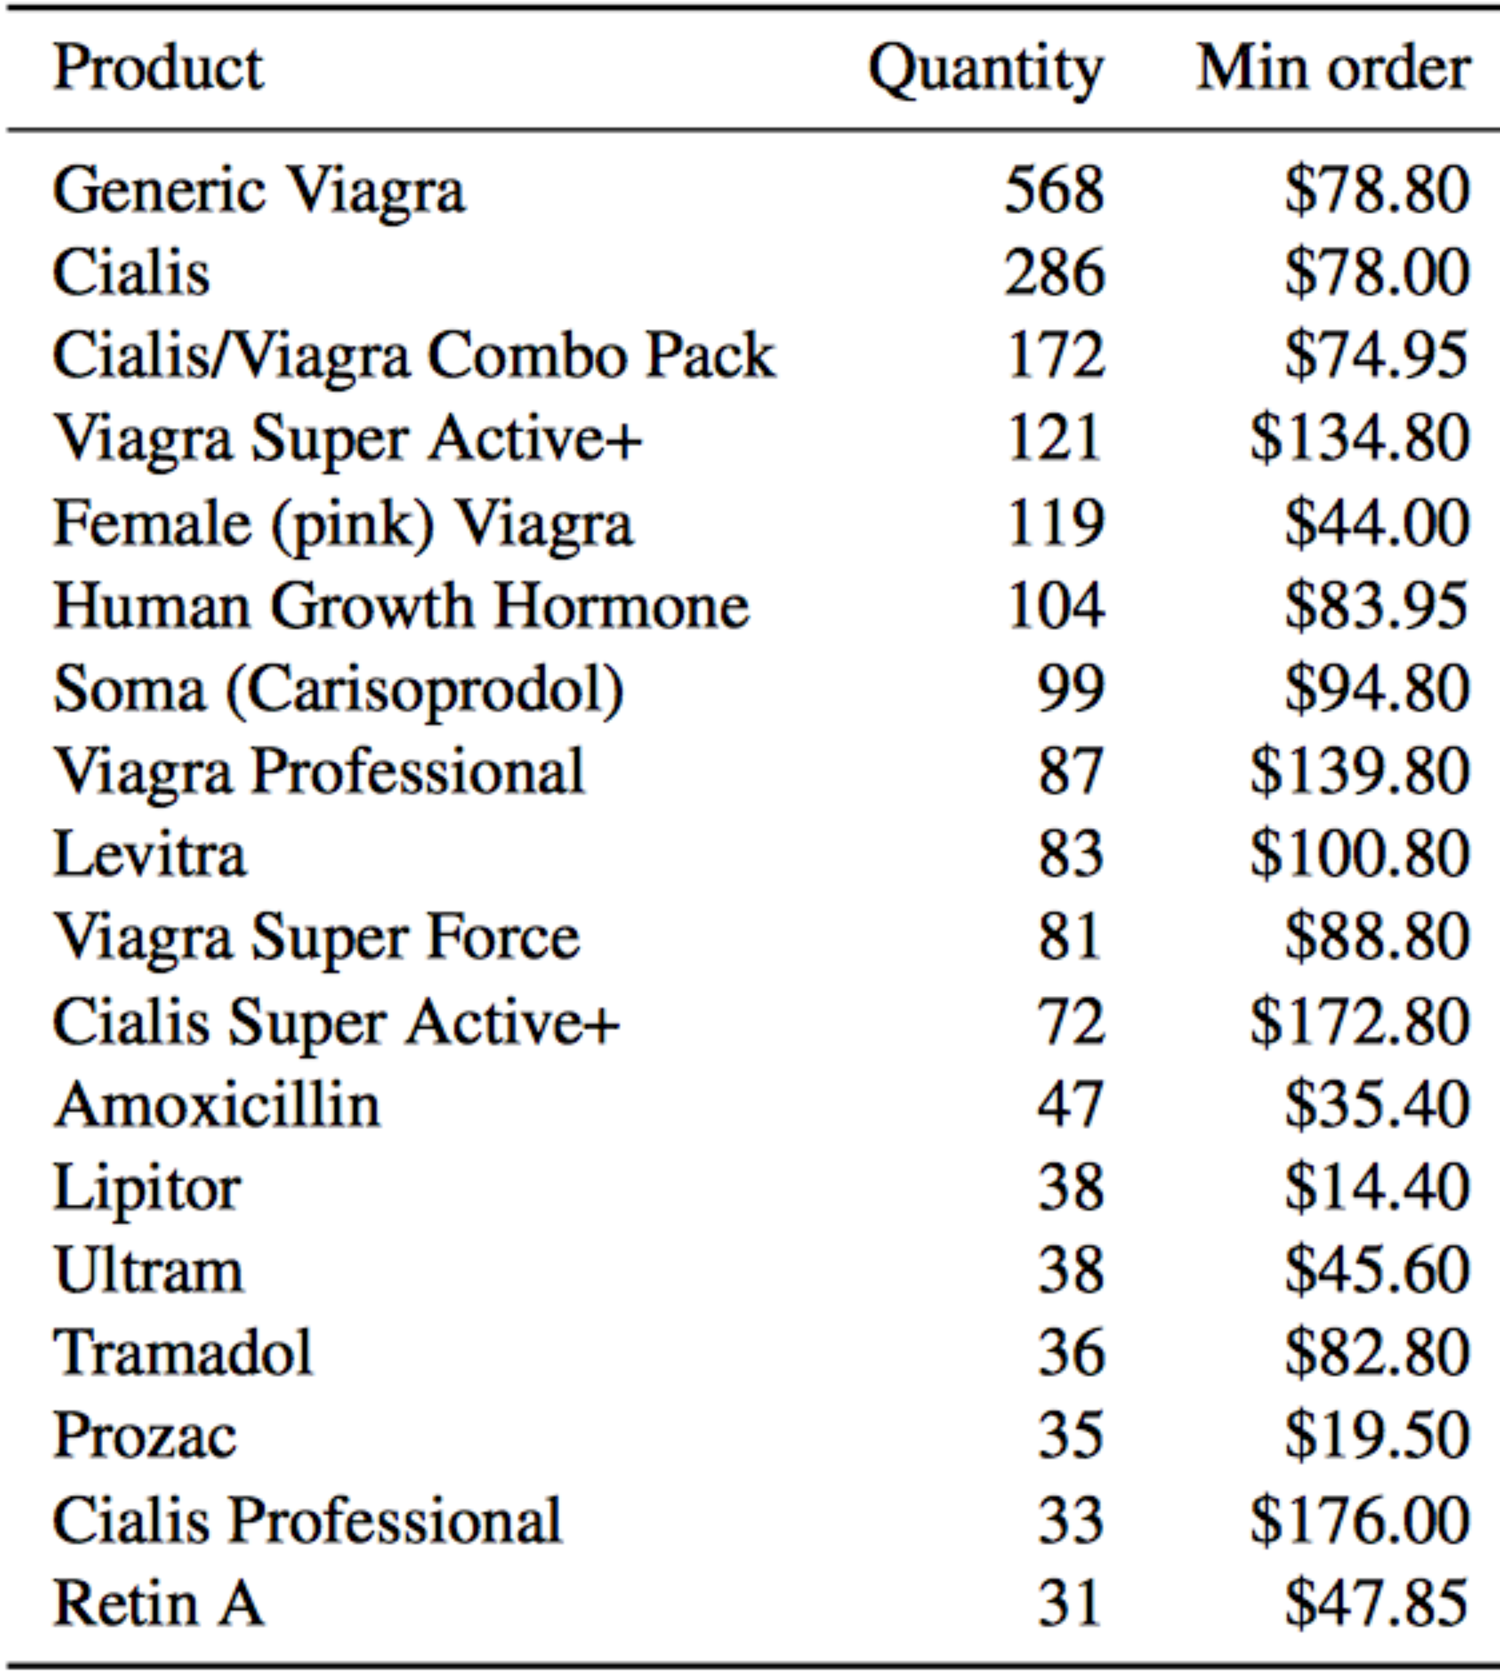

Product	Quantity	Min order
Generic Viagra	568	$78.80
Cialis	286	$78.00
Cialis/Viagra Combo Pack	172	$74.95
Viagra Super Active+	121	$134.80
Female (pink) Viagra	119	$44.00
Human Growth Hormone	104	$83.95
Soma (Carisoprodol)	99	$94.80
Viagra Professional	87	$139.80
Levitra	83	$100.80
Viagra Super Force	81	$88.80
Cialis Super Active+	72	$172.80
Amoxicillin	47	$35.40
Lipitor	38	$14.40
Ultram	38	$45.60
Tramadol	36	$82.80
Prozac	35	$19.50
Cialis Professional	33	$176.00
Retin A	31	$47.85

"Why do you rob banks?"

"Because that's where the money is"

Why does the emergence of the underground economy matter?

- Many of the centralized components of these networks get pursued and shut down
- Markets lead to efficiencies and specializations
  - Lowers barrier to entry: only need a single skill
  - Some underground market activities are legal
- Competition spurs innovation
  - Accelerates the arms race
  - Defenders must assume a more pessimistic threat model
- Facilitates non-$ Internet attacks
  - Provides actors (political, nation-state) with cheap attack components

WHY STUDYING IT MATTERS

And why *continuing* to study it matters

- Like any complex system, these markets can themselves be infiltrated
  - Some research on infiltrating affiliate programs & botnets, taking over C&C
- Can identify choke points
  - Many hosting services have been shut down
  - Draws attention to shady banks
  - Draws attention to shady doctors
    - Early spambot had one doctor writing 1500+ prescriptions per day