Understanding Security Notifications At Scale

Frank Li - University of California Berkeley
Zakir Durumeric - University of Michigan
Michael Bailey - University of Illinois Urbana-Champaign
Vern Paxson - University of California Berkeley

Why study security notifications?

Lots of work in academia and industry on identifying security issues

Why study security notifications?

Lots of work in academia and industry on identifying security issues

However, those who find security issues are often not the same party as those who need the information

Why study security notifications?

Lots of work in academia and industry on identifying security issues

However, those who find security issues are often not the same party as those who need the information

Security notifications serve as a bridge

Why study security notifications?

Lots of work in academia and industry on identifying security issues

However, those who find security issues are often not the same party as those who need the information.

Security notifications serve as a bridge

There has been little academic study of security notifications

Our Research Agenda

Better understand the nature of these notifications and the most effective approach to conducting them

Our Research Agenda

Better understand the nature of these notifications and the most effective approach to conducting them

Today:

- Share our experiences and analysis from conducting several notification efforts
- Hear from you about your experiences and lessons learned

Experiences

We have measured and analyzed notification sent for:

- Heartbleed bug
- Security misconfigurations and vulnerabilities
- Compromised websites

The Heartbleed Bug

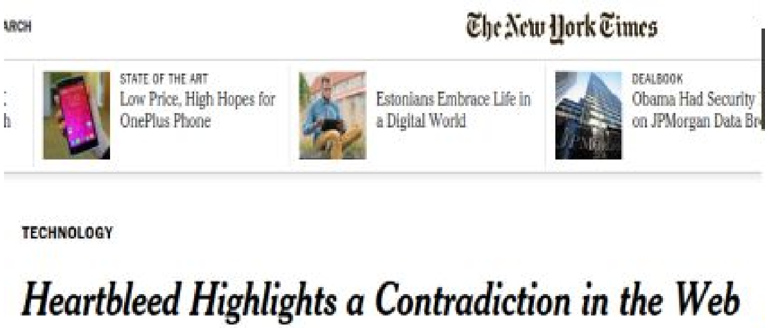

ARCH

The New York Times

STATE OF THE ART
Low Price, High Hopes for OnePlus Phone

Estonians Embrace Life in a Digital World

DEALBOOK
Obama Had Security on JPMorgan Data Br

TECHNOLOGY

Heartbleed Highlights a Contradiction in the Web

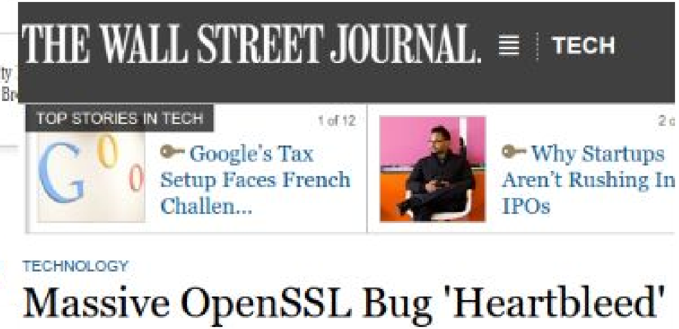

THE WALL STREET JOURNAL. ≣ TECH

TOP STORIES IN TECH 1 of 12

Google's Tax Setup Faces French Challen...

2 o

Why Startups Aren't Rushing Into IPOs

TECHNOLOGY

Massive OpenSSL Bug 'Heartbleed'

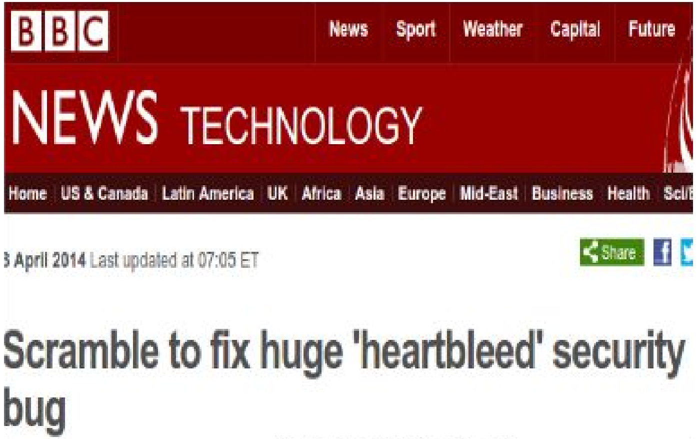

BBC News Sport Weather Capital Future

NEWS TECHNOLOGY

Home | US & Canada | Latin America | UK | Africa | Asia | Europe | Mid-East | Business | Health | Sci/E

8 April 2014 Last updated at 07:05 ET

Share

Scramble to fix huge 'heartbleed' security bug

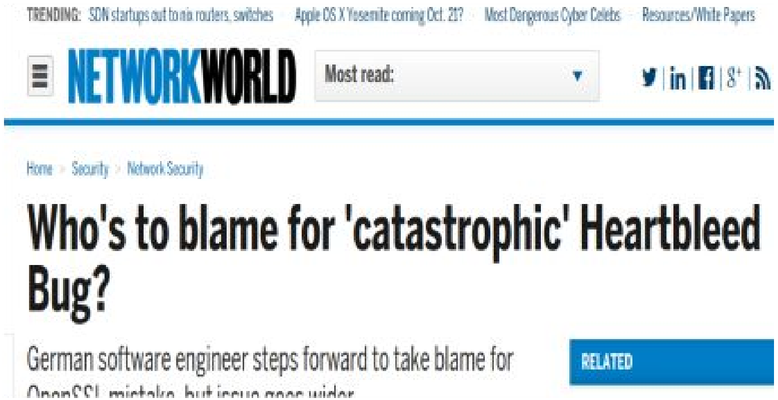

TRENDING: SDN startups out to nix routers, switches · Apple OS X Yosemite coming Oct. 21? · Most Dangerous Cyber Celebs · Resources/White Papers

NETWORKWORLD Most read: ▼

Home > Security > Network Security

Who's to blame for 'catastrophic' Heartbleed Bug?

German software engineer steps forward to take blame for OpenSSL mistake, but issue goes wider

RELATED

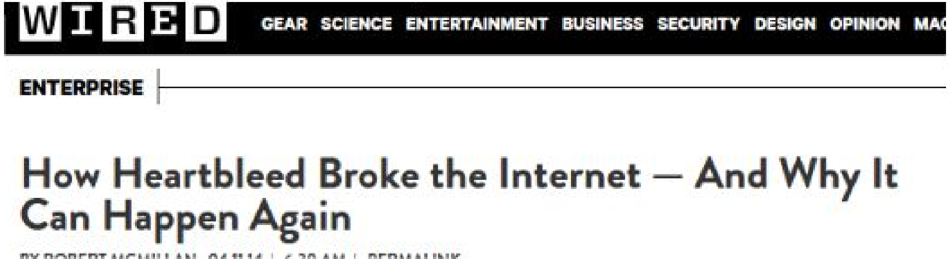

WIRED GEAR SCIENCE ENTERTAINMENT BUSINESS SECURITY DESIGN OPINION MAG

ENTERPRISE

How Heartbleed Broke the Internet — And Why It Can Happen Again

BY ROBERT MCMILLAN 04.11.14 | 6:30 AM | PERMALINK

What is Heartbleed?

- Allows access to sensitive data in memory, such as passwords, secret keys, etc., on OpenSSL servers

- Fix: Update to patched version, or disable TLS “Heartbeats”

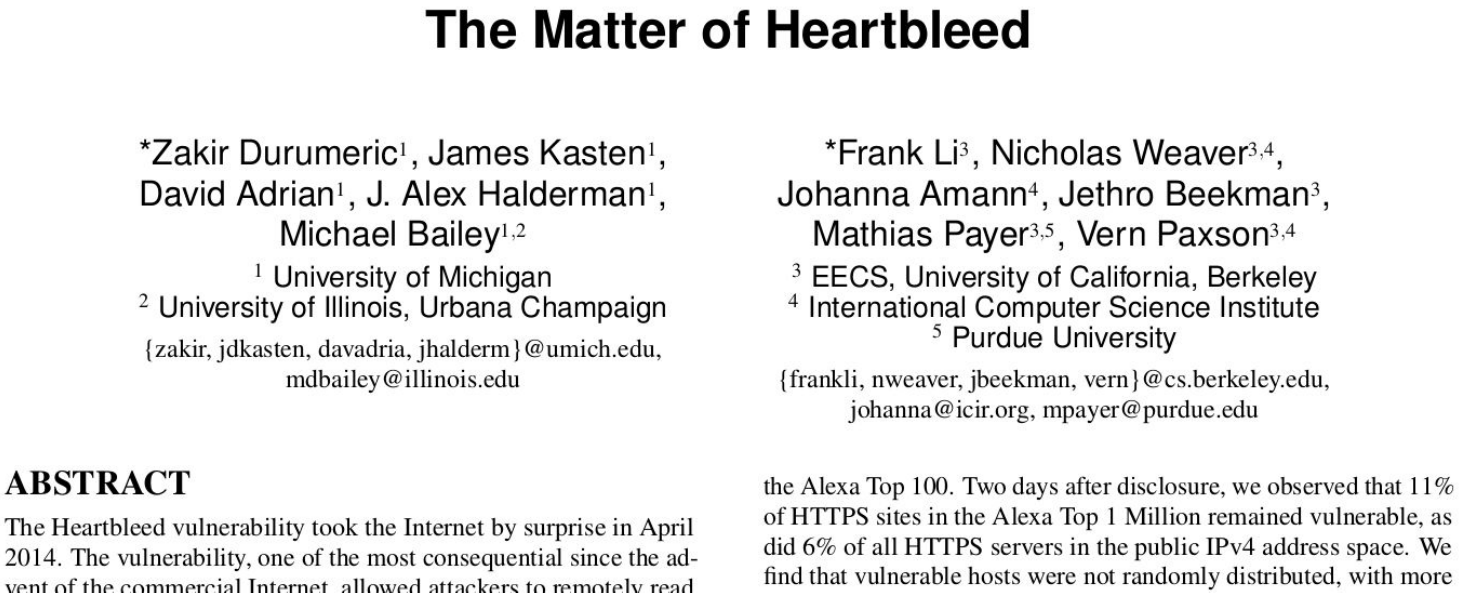

The Matter of Heartbleed

*Zakir Durumeric[1], James Kasten[1],
David Adrian[1], J. Alex Halderman[1],
Michael Bailey[1,2]
[1] University of Michigan
[2] University of Illinois, Urbana Champaign
{zakir, jdkasten, davadria, jhalderm}@umich.edu,
mdbailey@illinois.edu

*Frank Li[3], Nicholas Weaver[3,4],
Johanna Amann[4], Jethro Beekman[3],
Mathias Payer[3,5], Vern Paxson[3,4]
[3] EECS, University of California, Berkeley
[4] International Computer Science Institute
[5] Purdue University
{frankli, nweaver, jbeekman, vern}@cs.berkeley.edu,
johanna@icir.org, mpayer@purdue.edu

ABSTRACT

The Heartbleed vulnerability took the Internet by surprise in April 2014. The vulnerability, one of the most consequential since the advent of the commercial Internet, allowed attackers to remotely read

the Alexa Top 100. Two days after disclosure, we observed that 11% of HTTPS sites in the Alexa Top 1 Million remained vulnerable, as did 6% of all HTTPS servers in the public IPv4 address space. We find that vulnerable hosts were not randomly distributed, with more

ACM Internet Measurement Conference 2014

Detecting Vulnerable Hosts

Used the ZMap scanner to scan HTTPS servers

Ethical consideration: probe packet *does not* exploit Heartbleed or read any data from memory

Patch Rates

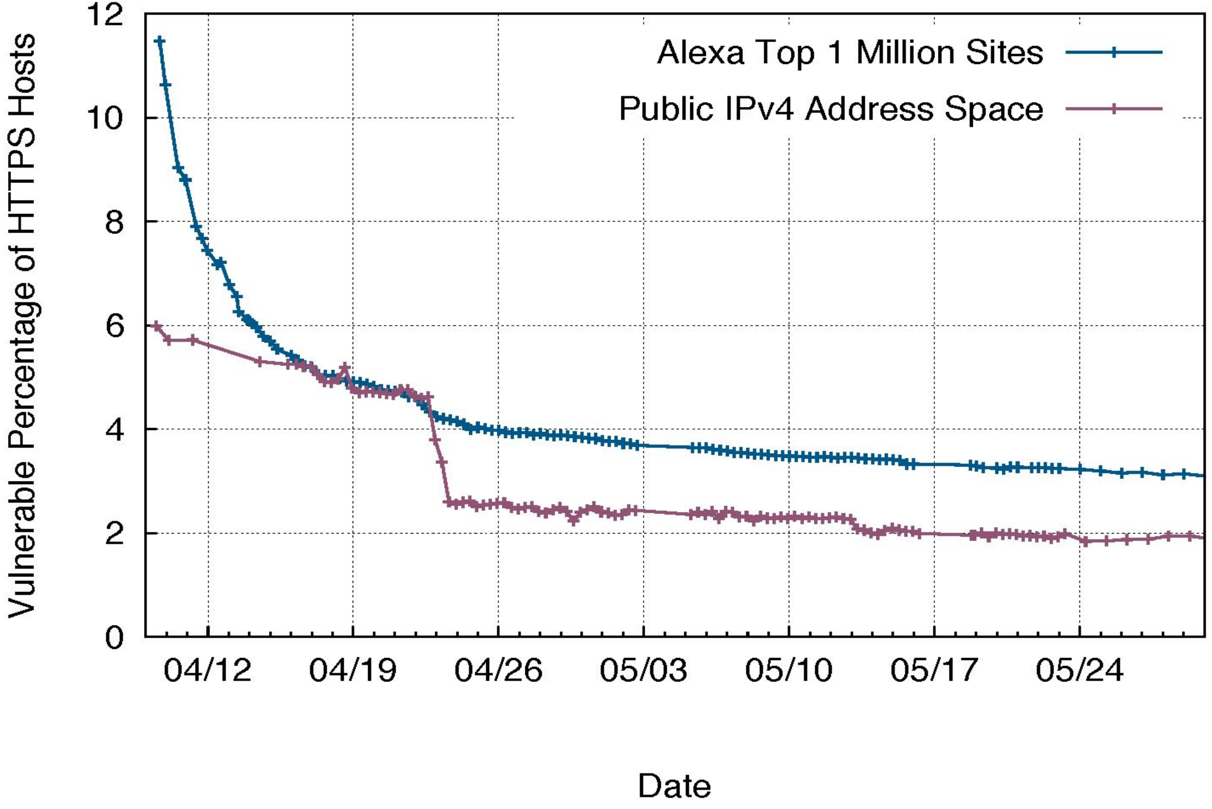
Vulnerable Percentage of HTTPS Hosts
0
2
4
6
8
10
12
04/12
04/19
04/26
05/03
05/10
05/17
05/24
Date
Alexa Top 1 Million Sites
Public IPv4 Address Space

Notification Effort

- April 24: Grabbed 4646 unique contact emails from WHOIS lookups for ~250k still-vulnerable IPs

- Randomly selected half to notify via email on April 28th, the other half notified on May 7th

- Scanned every 8 hours to track behavior

Notification Groups Patching Rates

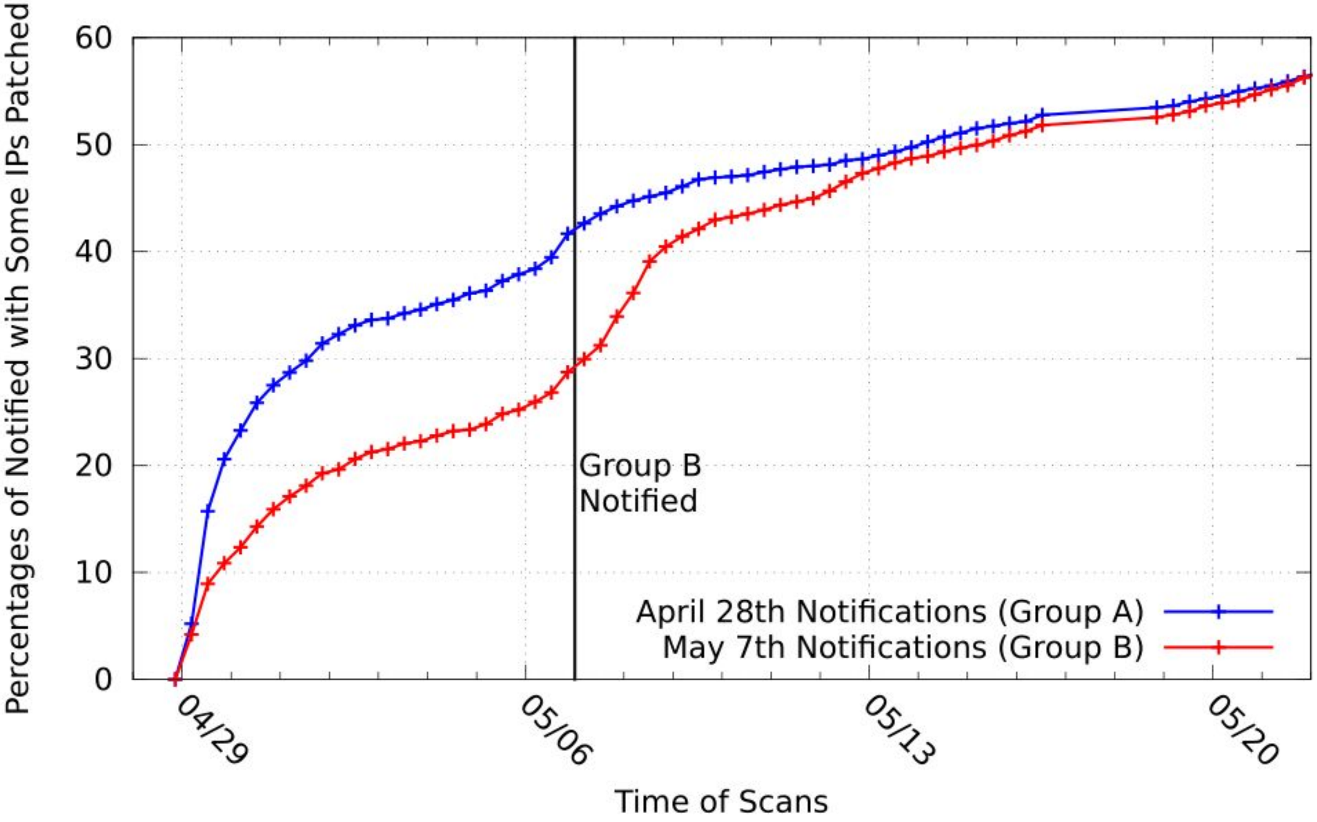

Notification Groups Patching Rates

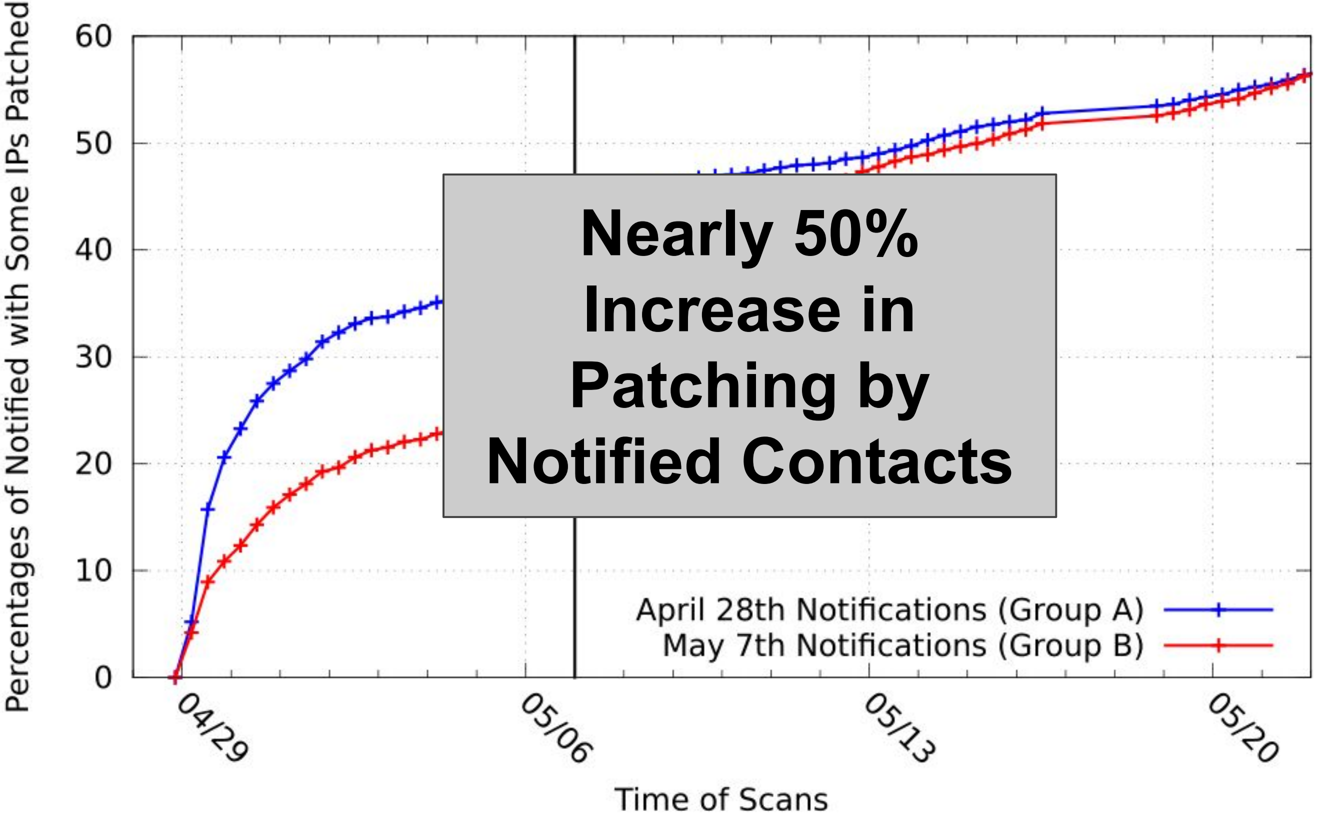

Nearly 50% Increase in Patching by Notified Contacts

First Round Responses

- Received 530 email responses
- 11.1% human responses, 40.2% automated, and 48.7% delivery failures

First Round Responses

- Received 530 email responses
- 11.1% human responses, 40.2% automated, and 48.7% delivery failures
- Of human contacts:
  - 92% positive
  - 5% neutral
  - 3% negative

First Round Responses

- Received 530 email responses
- 11.1% human responses, 40.2% automated, and 48.7% delivery failures
- Automated messages
  - Confirmations
  - Tickets
  - Trackers (many incorrectly configured)

Lessons Learned

- Notifications *can* be effective at promoting patching.

- Mass notifications are doable and can be well-received.

New Questions...

- How effective are notifications in other scenarios?
- How do we find reliable contacts for more hosts?
- What are best practices for effective notifications?

Security Misconfiguration Notifications

Security Misconfiguration Notifications

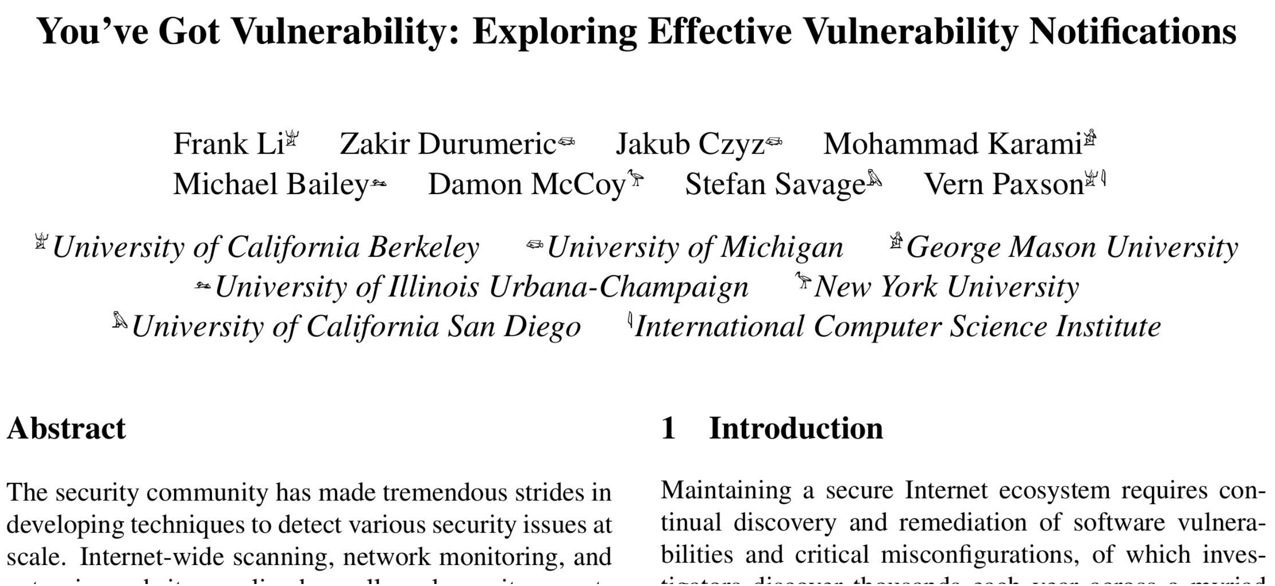

You've Got Vulnerability: Exploring Effective Vulnerability Notifications

Frank Li[𓀠] Zakir Durumeric[𓂀] Jakub Czyz[𓂀] Mohammad Karami[𓁐]
Michael Bailey[𓆓] Damon McCoy[𓅞] Stefan Savage[𓅐] Vern Paxson[𓀠𓇋]

[𓀠]University of California Berkeley [𓂀]University of Michigan [𓁐]George Mason University
[𓆓]University of Illinois Urbana-Champaign [𓅞]New York University
[𓅐]University of California San Diego [𓇋]International Computer Science Institute

Abstract

The security community has made tremendous strides in developing techniques to detect various security issues at scale. Internet-wide scanning, network monitoring, and

1 Introduction

Maintaining a secure Internet ecosystem requires continual discovery and remediation of software vulnerabilities and critical misconfigurations, of which inves-

USENIX Security 2016

Security Misconfiguration Notifications

Notifications for 3 classes of misconfigurations:

- Publicly Accessible Industrial Control Systems (ICS)
- DDoS Amplifiers
- Misconfigured IPv6 Firewall Policies

Security Misconfiguration Notifications

Publicly Accessible Industrial Control Systems (ICS):

- Remotely control physical infrastructure, but lacks important security features
- *Detection/tracking*: Protocol-specific fingerprints with ZMap
- *Fix*: Firewall or remove from public Internet

Security Misconfiguration Notifications

DDoS Amplifiers

- Protocols abused for DDoS attacks
- *Detection*: Monitoring DDoS attacks against a network
- *Tracking*: Custom protocol specific probing
- *Fix*: Firewall or disable protocols or abused functions

Security Misconfiguration Notifications

Misconfigured IPv6 Firewall Policies

- v6-only services may indicate firewall misconfiguration
- *Detection/tracking*: Large-scale probing with CAIDA’s Scamper tool
- *Fix*: Correct firewall policies, or disabling applications

Experiment Variables

- Who to contact?

  WHOIS contact, our local US-CERT, host’s local CERT

Experiment Variables

- Who to contact?

  WHOIS contact, our local US-CERT, host’s local CERT

- What to say to server admins (WHOIS contacts)?

  Terse message

  Terse message with a link to detailed info site

  Verbose message with details

Notification Methodology

- Found abuse contacts via WHOIS
- Grouped hosts by their abuse contacts
- Randomly assigned contacts to control vs CERT vs WHOIS groups

Experiment Groups

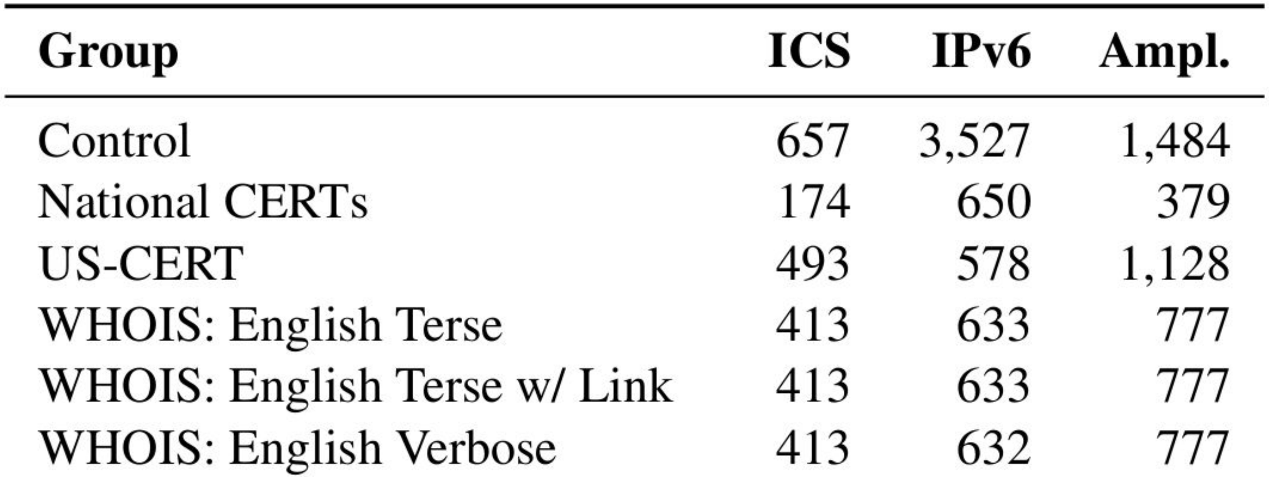

Group	ICS	IPv6	Ampl.
Control	657	3,527	1,484
National CERTs	174	650	379
US-CERT	493	578	1,128
WHOIS: English Terse	413	633	777
WHOIS: English Terse w/ Link	413	633	777
WHOIS: English Verbose	413	632	777

Results

Results

Our notifications had no effect on DDoS Amplifiers…

- Prior notification efforts
- Population bias

Remediation Rates

IPv6

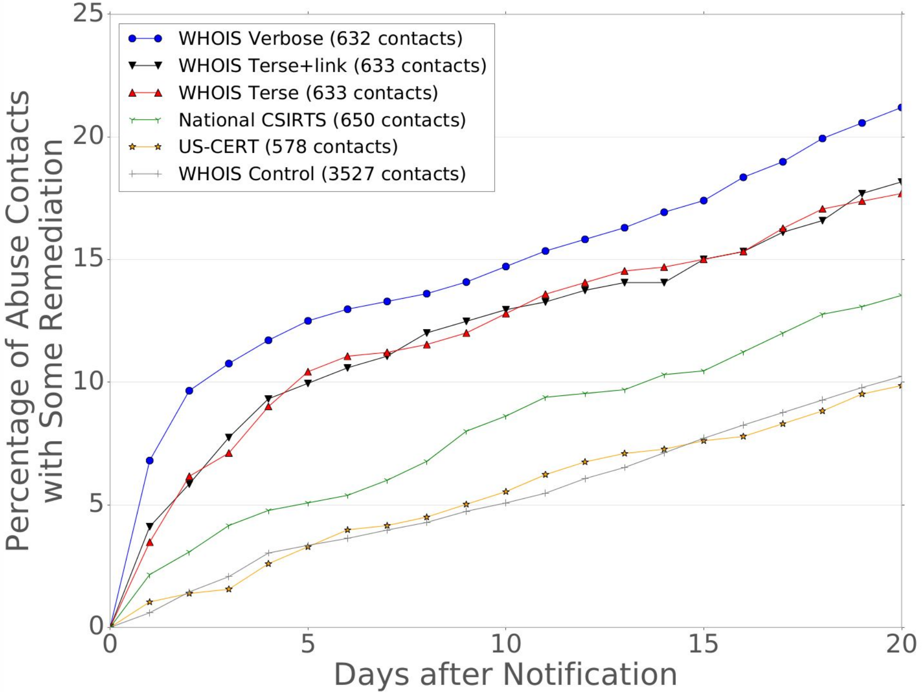

ICS

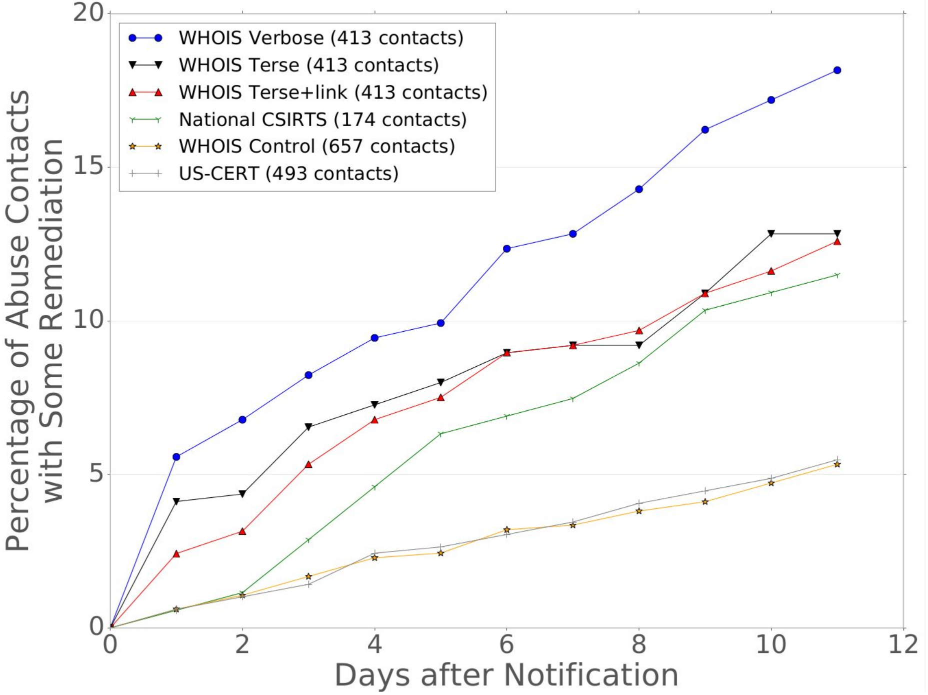

Remediation Rates

IPv6

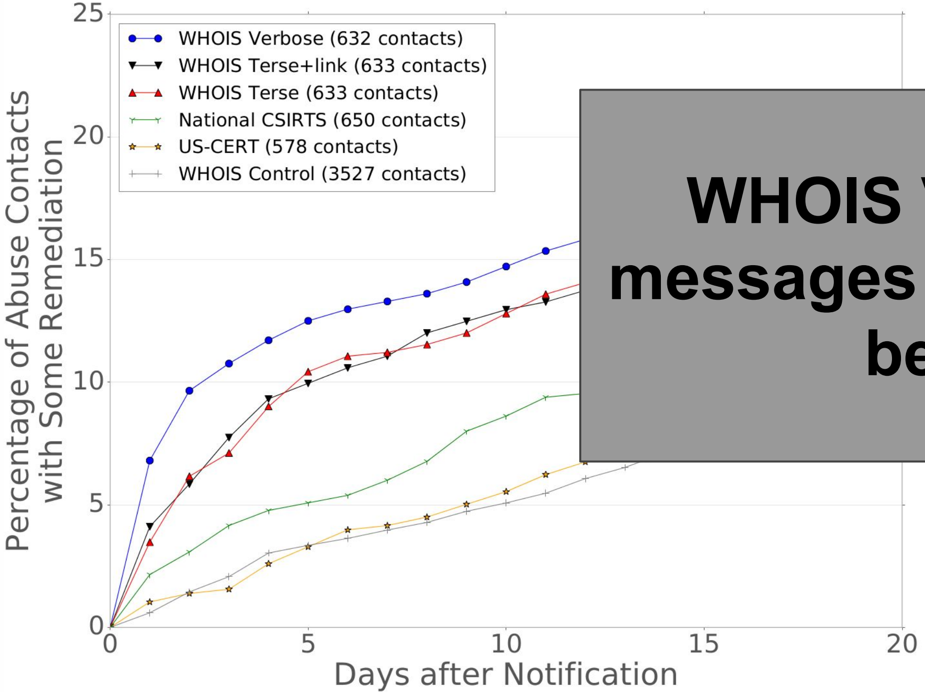

ICS

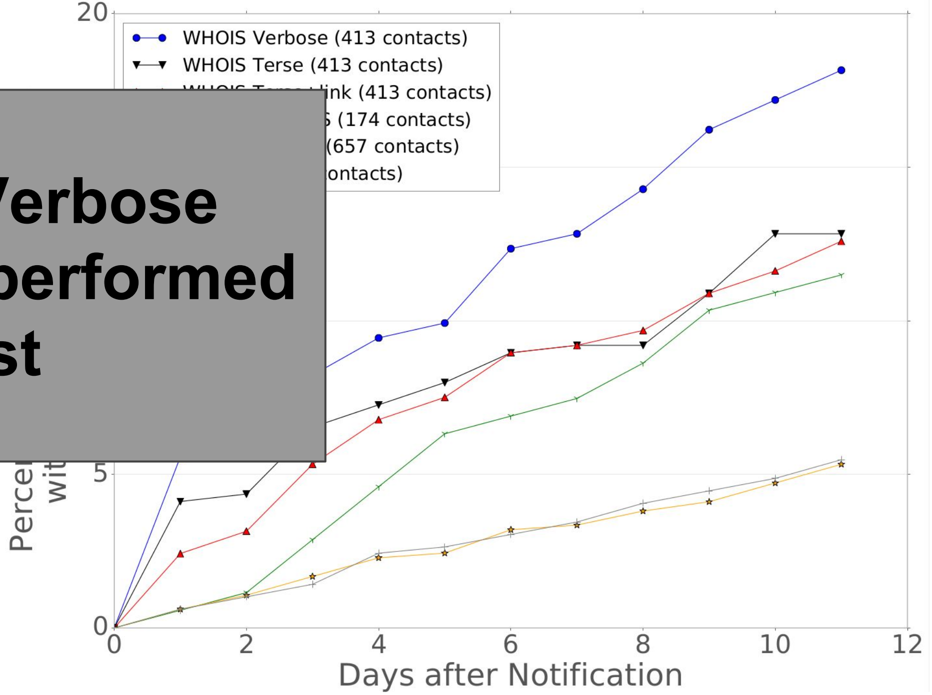

WHOIS Verbose messages performed best

Remediation Rates

IPv6

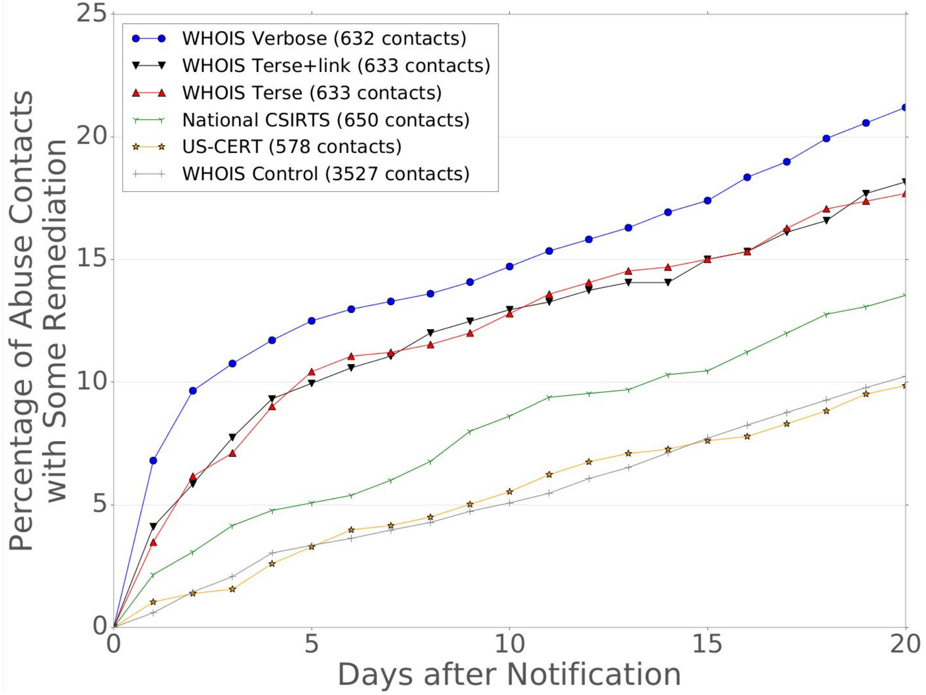

ICS

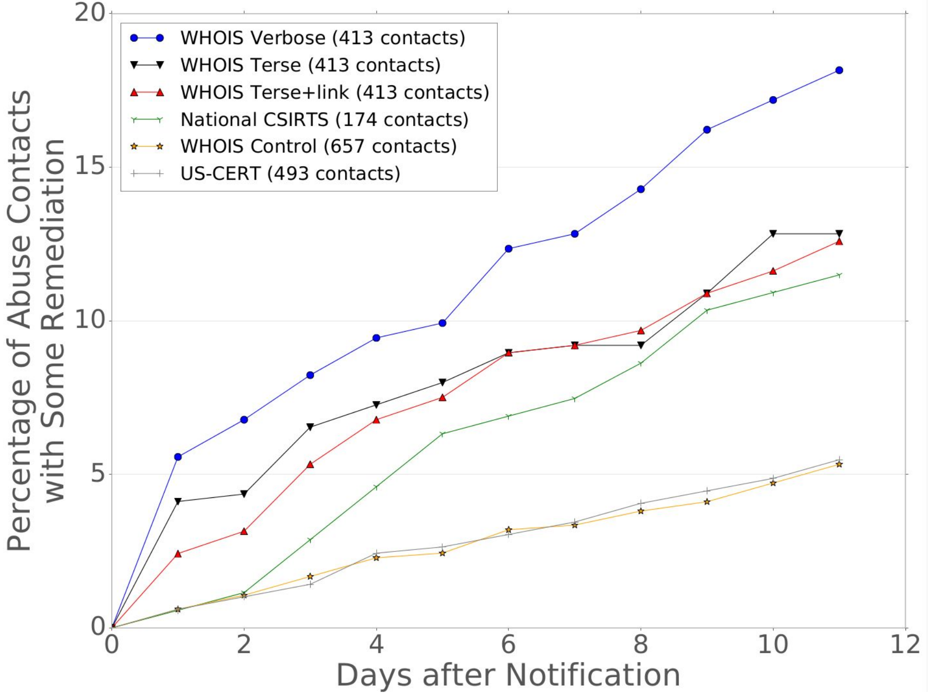

Remediation Rates

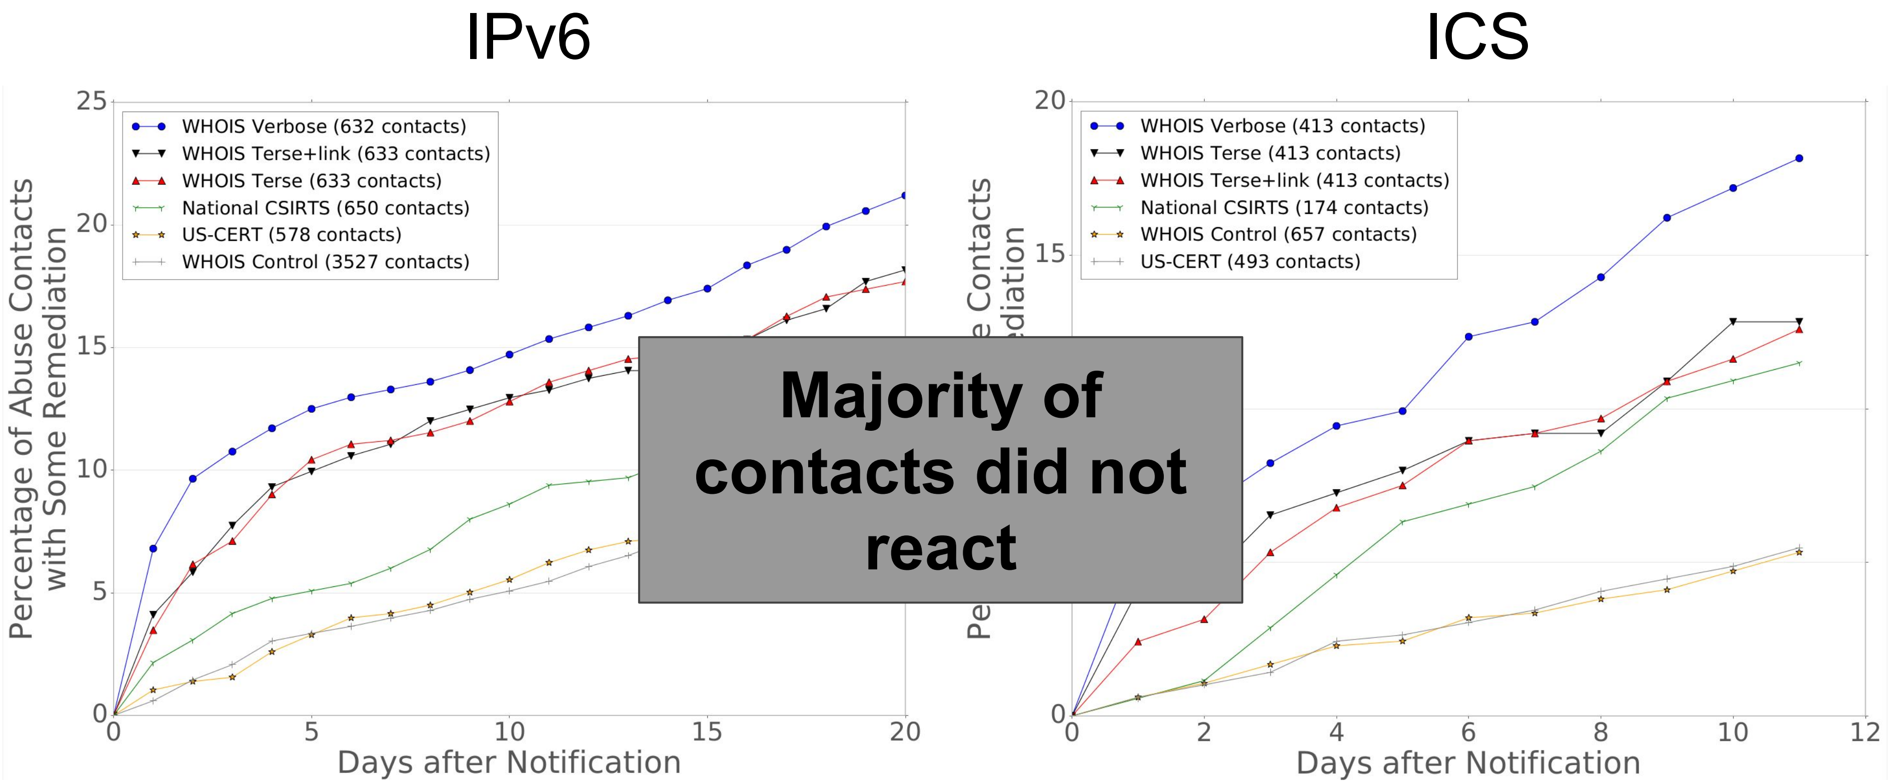

Majority of contacts did not react

Remediation Rates

IPv6

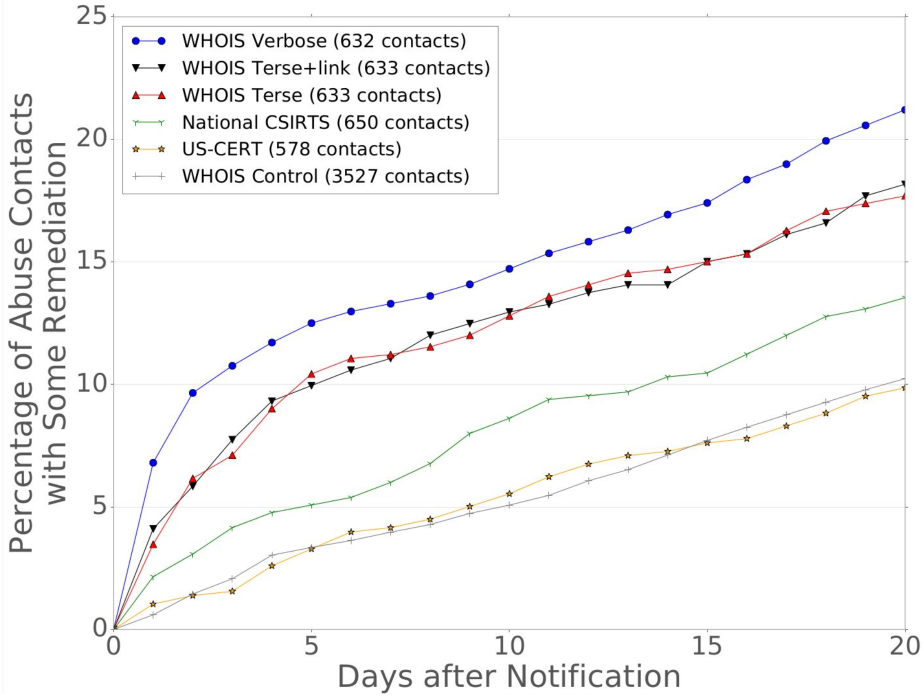

ICS

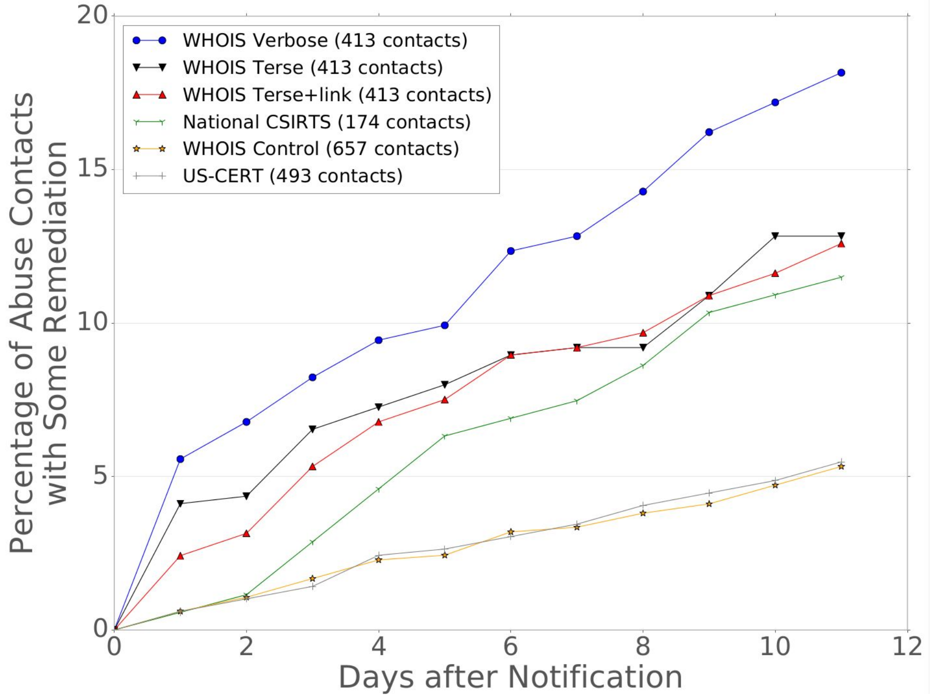

Remediation Rates

IPv6

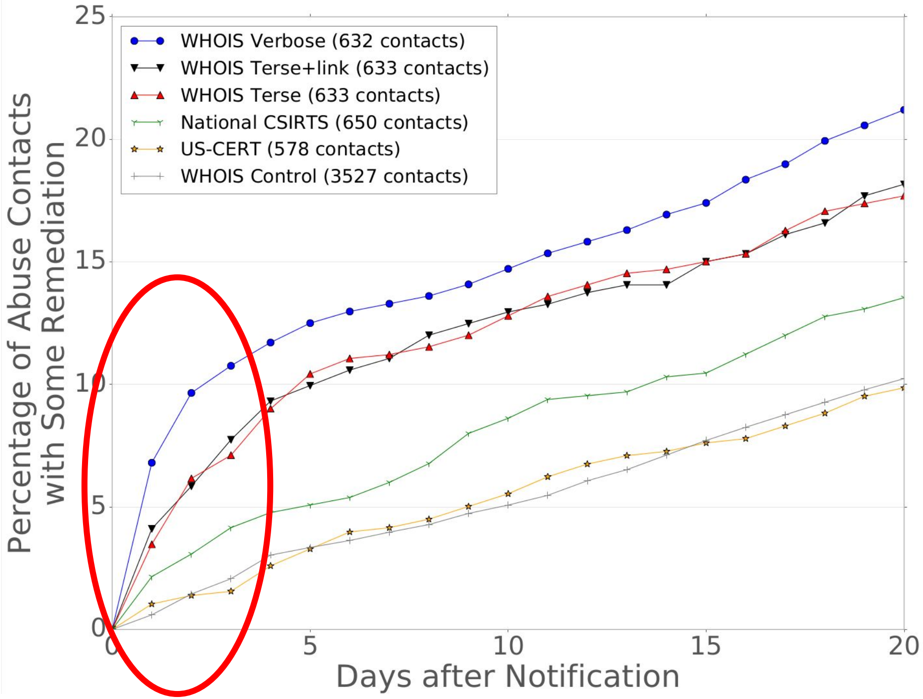

ICS

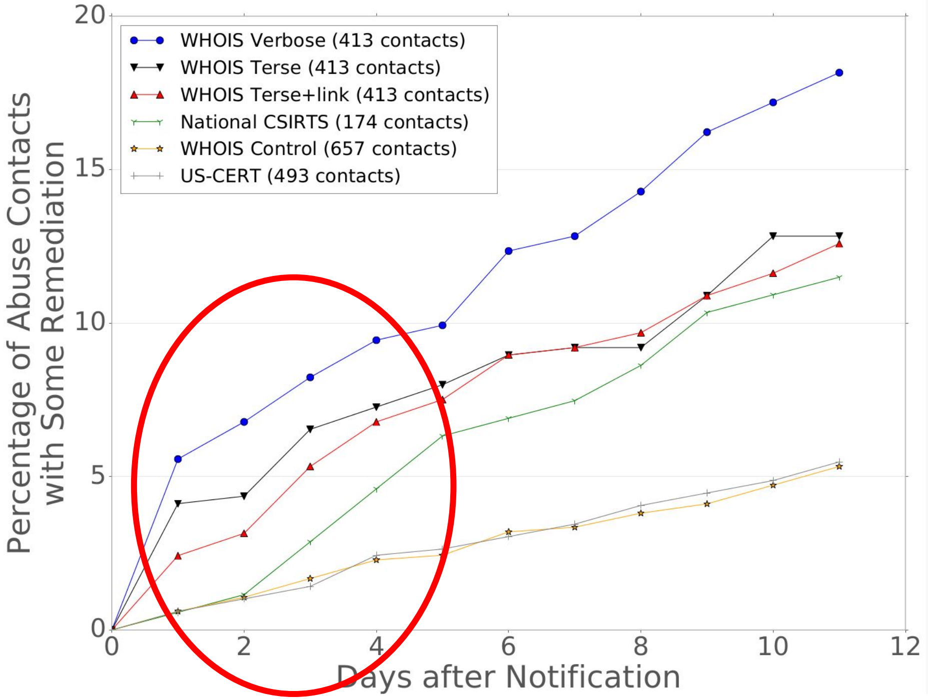

Remediation Rates

IPv6

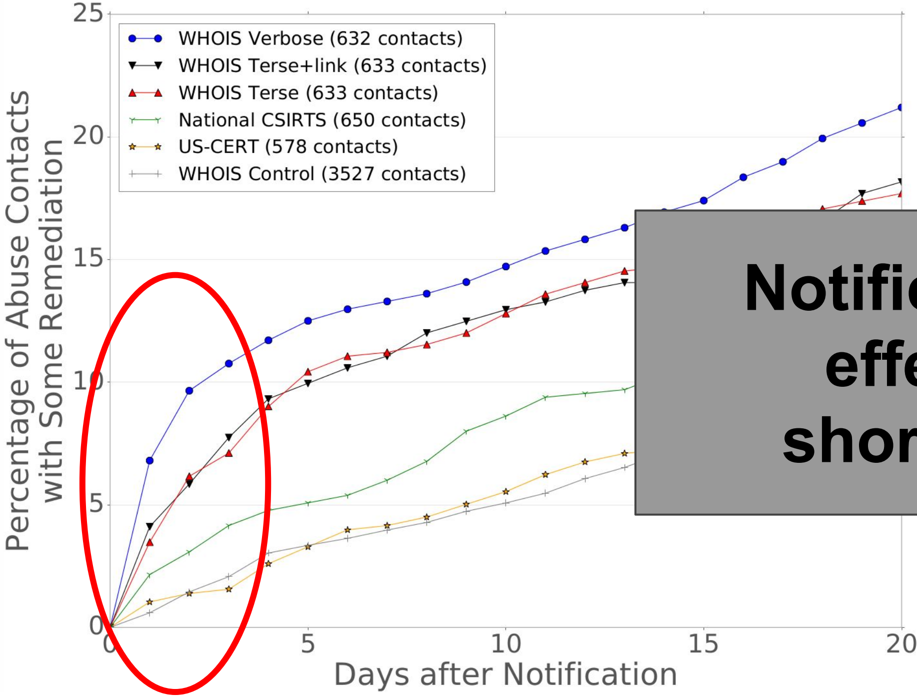

ICS

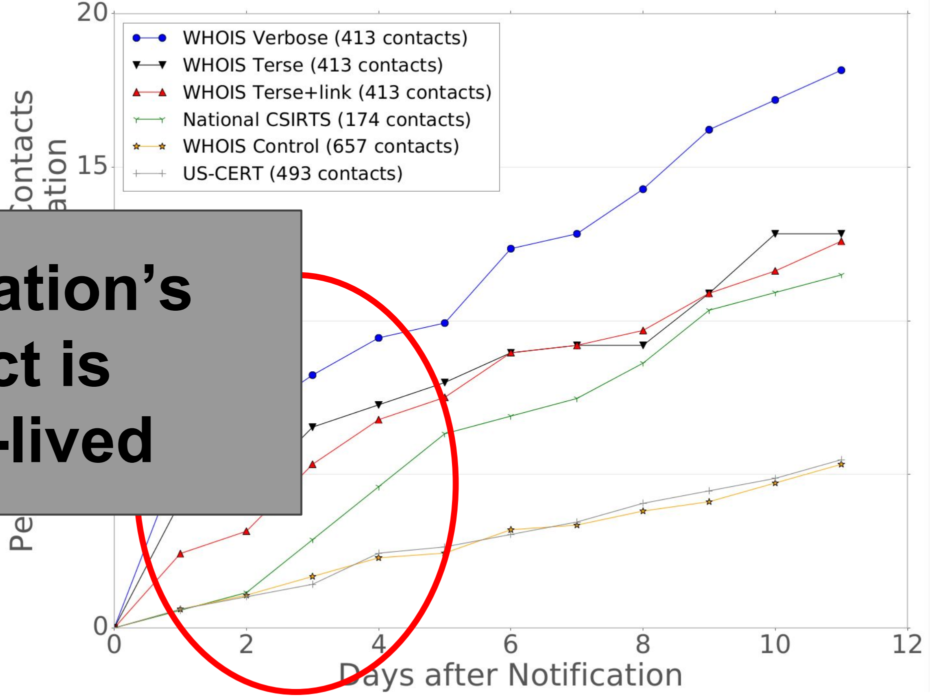

Notification’s effect is short-lived

Staying Power of Notification's Effect

IPv6

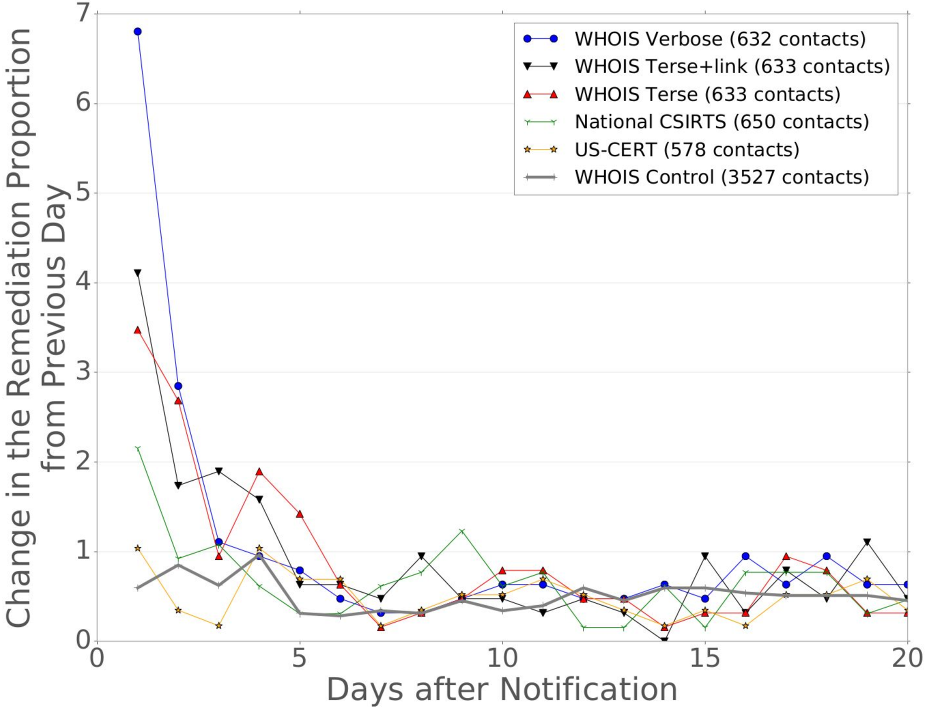

ICS

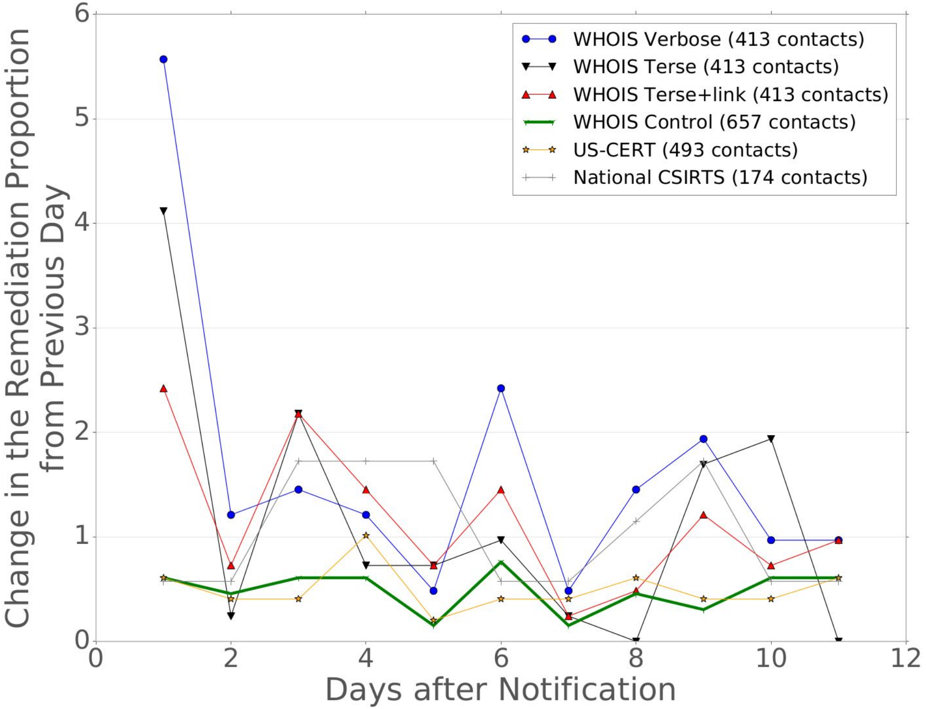

Notification Response

- Received 685 emails
- 13.6% were human, 77.4% were automated responses, and 9.1% were bounces

Notification Response

- Received 685 emails
- 13.6% were human, 77.4% were automated responses, and 9.1% were bounces
- Of human responses:
    - 77% were positive
    - 19% neutral
    - 4% negative

Insights

- Verbose messages to WHOIS contacts can be relatively effective.

- However, overall effectiveness is limited.

- Notification’s effect is short-lived, partly due to lack of reliable points of contact.

Another context: Hijacked Websites

Another context: Hijacked Websites

World Wide Web Conference (WWW) 2016

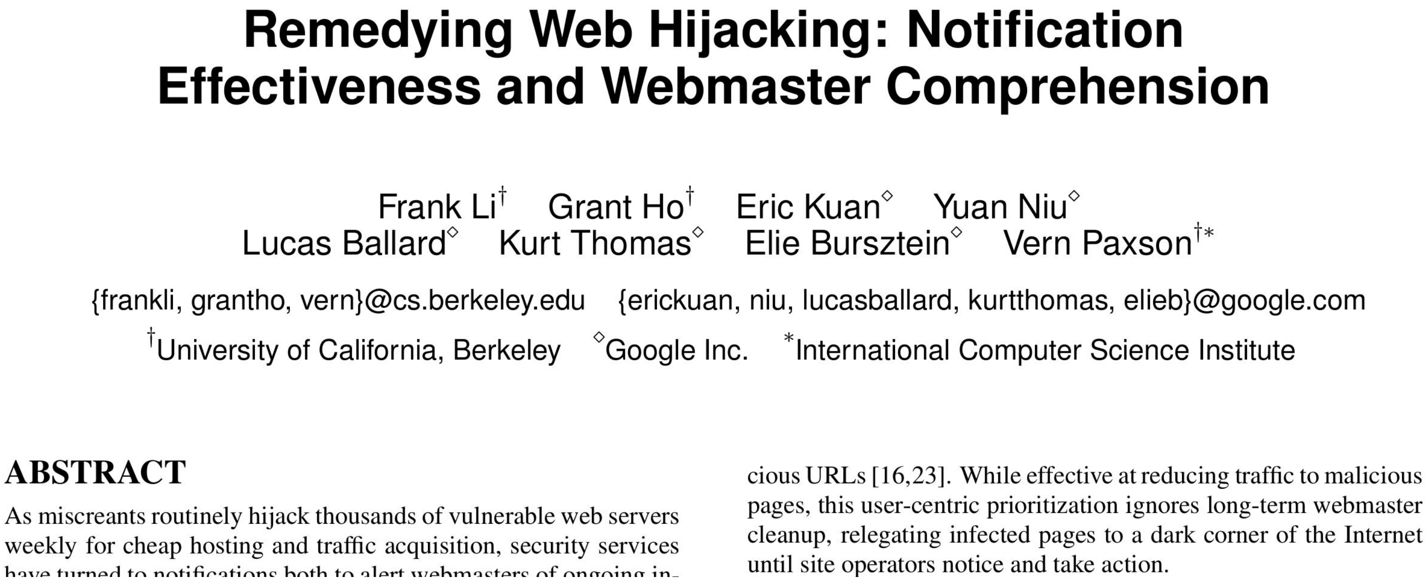

Remedying Web Hijacking: Notification Effectiveness and Webmaster Comprehension

Frank Li[†] Grant Ho[†] Eric Kuan[◇] Yuan Niu[◇]
Lucas Ballard[◇] Kurt Thomas[◇] Elie Bursztein[◇] Vern Paxson[†*]

{frankli, grantho, vern}@cs.berkeley.edu {erickuan, niu, lucasballard, kurtthomas, elieb}@google.com

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

ABSTRACT

As miscreants routinely hijack thousands of vulnerable web servers weekly for cheap hosting and traffic acquisition, security services have turned to notifications both to alert webmasters of ongoing in-

cious URLs [16,23]. While effective at reducing traffic to malicious pages, this user-centric prioritization ignores long-term webmaster cleanup, relegating infected pages to a dark corner of the Internet until site operators notice and take action.

Websites are constantly hijacked...

sanfranciscobaycoffee.com

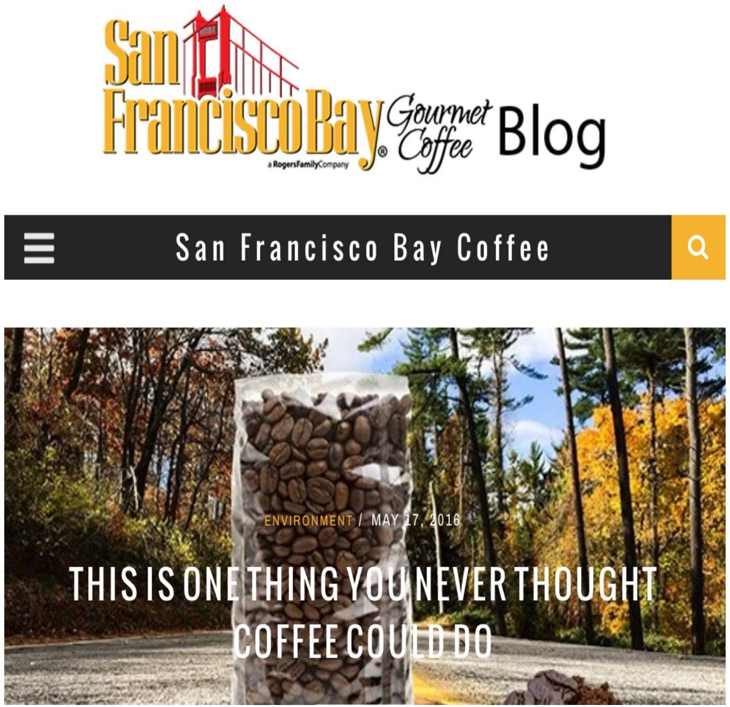
San Francisco Bay® Gourmet Coffee Blog
a RogersFamilyCompany
San Francisco Bay Coffee
ENVIRONMENT / MAY 17, 2016
THIS IS ONE THING YOU NEVER THOUGHT COFFEE COULD DO

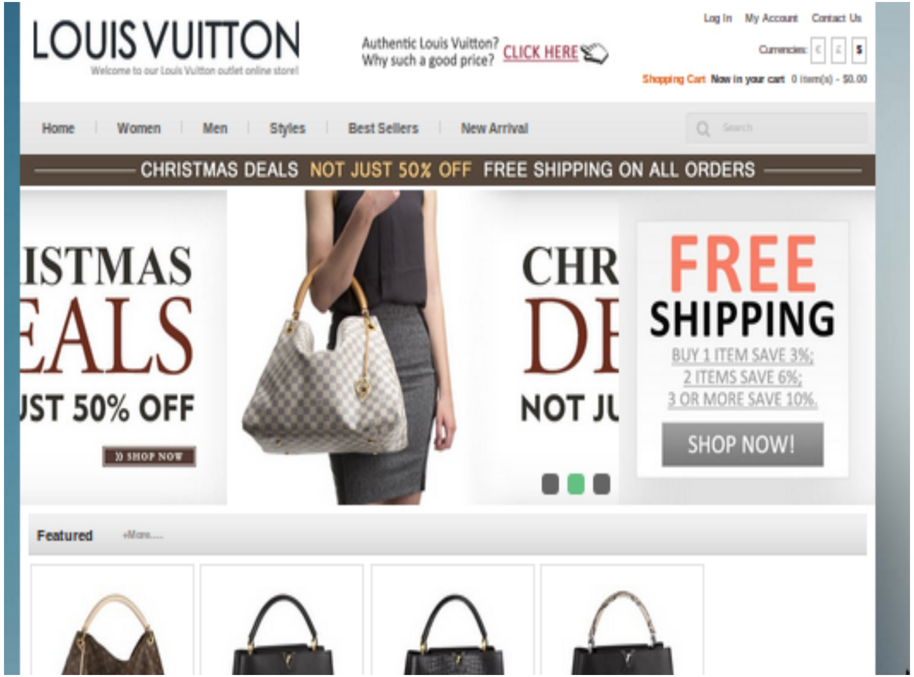
LOUIS VUITTON
Welcome to our Louis Vuitton outlet online store!
Authentic Louis Vuitton? Why such a good price? CLICK HERE
Log In
My Account
Contact Us
Currencies: € £ $
Shopping Cart Now in your cart 0 item(s) - $0.00
Home
Women
Men
Styles
Best Sellers
New Arrival
Search
CHRISTMAS DEALS NOT JUST 50% OFF FREE SHIPPING ON ALL ORDERS
ISTMAS EALS JST 50% OFF
» SHOP NOW
CHR DE NOT JU
FREE SHIPPING
BUY 1 ITEM SAVE 3%; 2 ITEMS SAVE 6%; 3 OR MORE SAVE 10%.
SHOP NOW!
Featured
+More....

Websites are constantly hijacked...

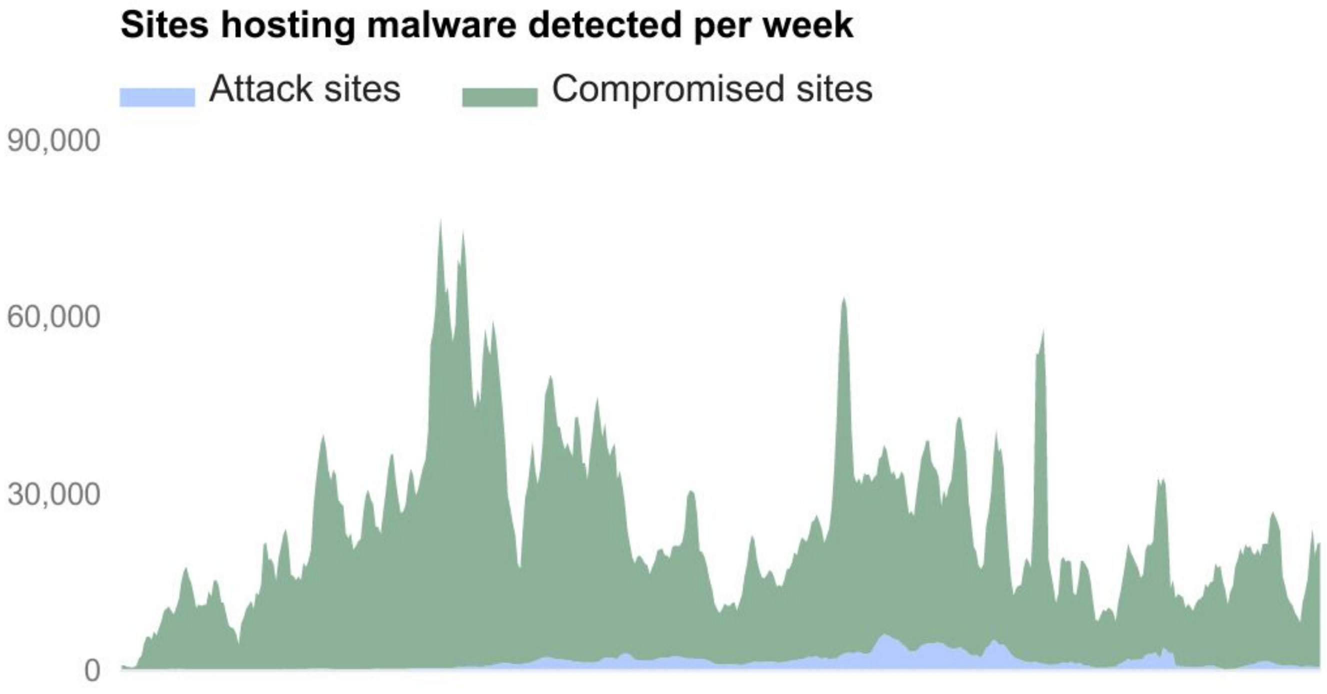

Google Safe Browsing Transparency Report

Compromised sites lead to...

- Drive-by downloads
- Cloaked redirections
- Scams
- Phishing
- Defacements

This Study: Analysis of ~1 Year of Google Webmaster Notifications

This Study: Analysis of ~1 Year of Google Webmaster Notifications

What works effectively for notifying webmasters?

This Study: Analysis of ~1 Year of Google Webmaster Notifications

What works effectively for notifying webmasters?

What factors affect remediation behavior?

This Study: Analysis of ~1 Year of Google Webmaster Notifications

What works effectively for notifying webmasters?

What factors affect remediation behavior?

How well are webmasters able to comprehend the remediation process?

Compromise Life Cycle

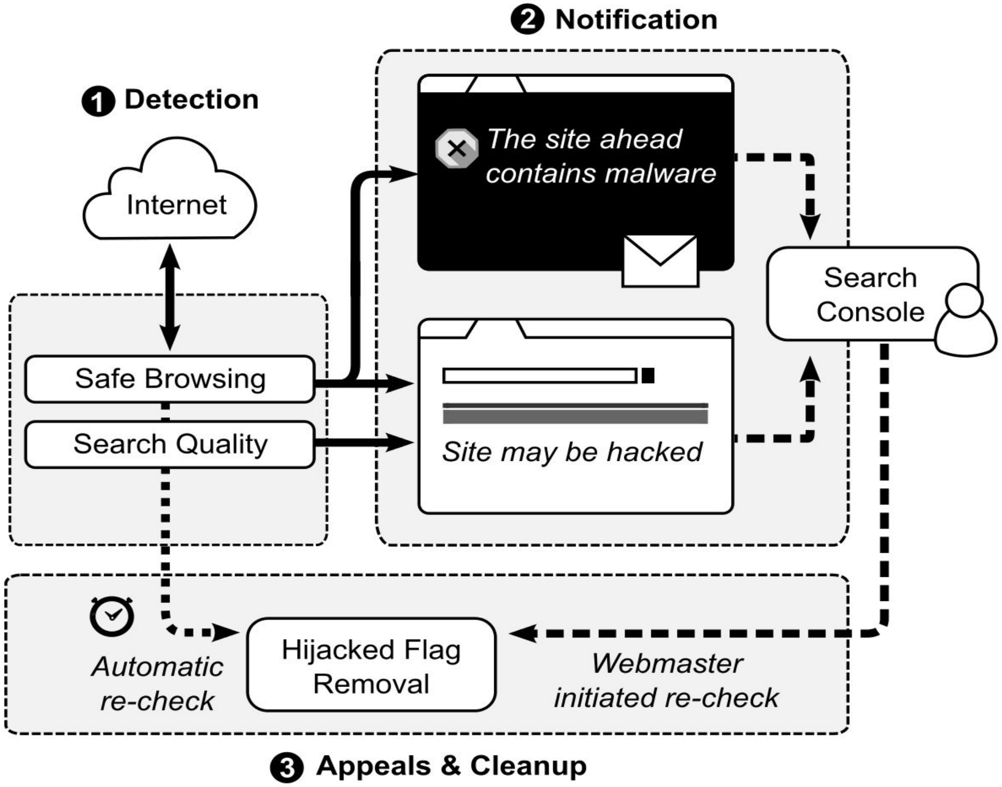

Compromise Life Cycle

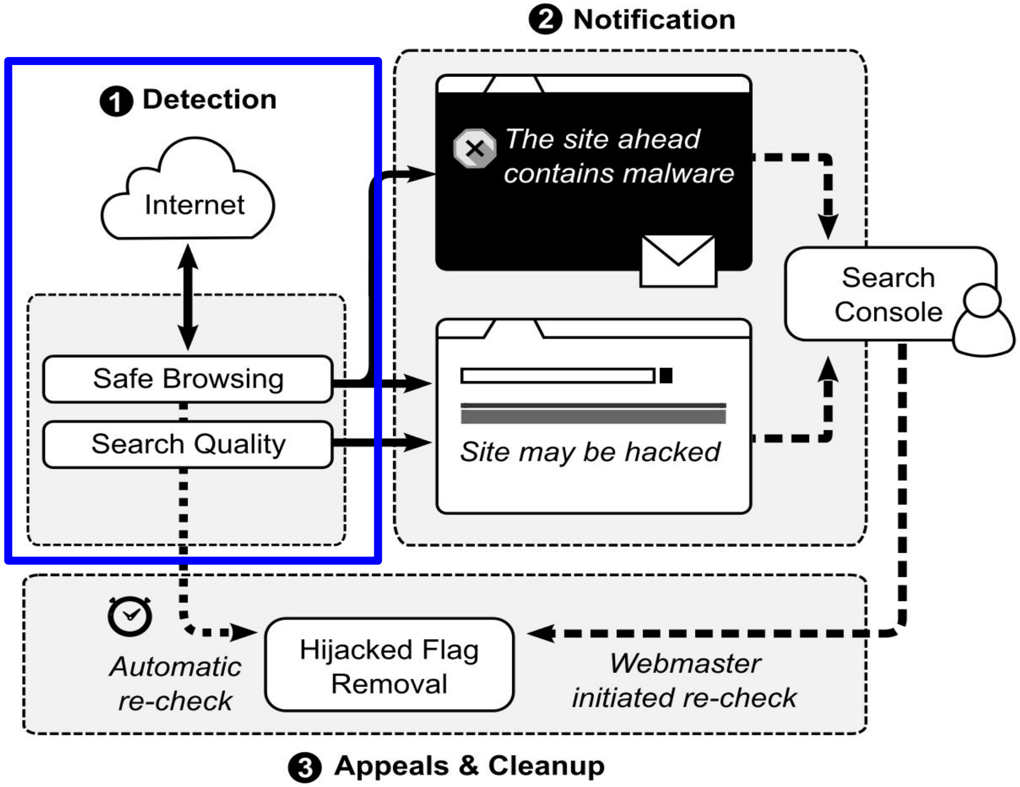

Compromise Life Cycle

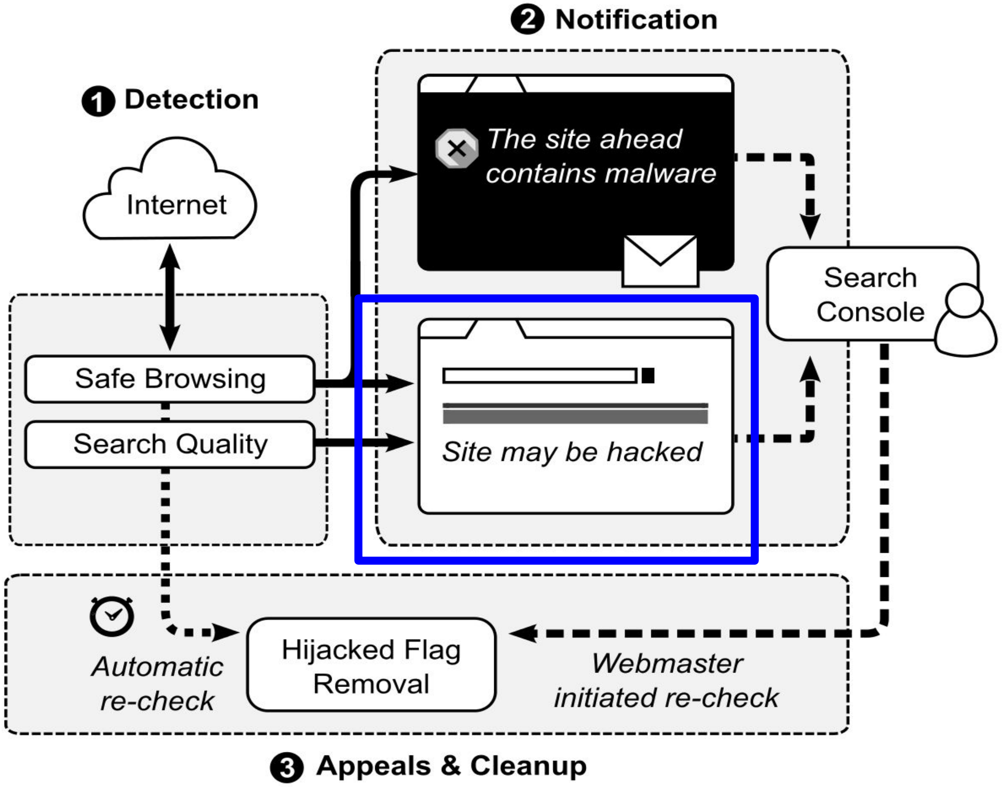

Compromise Life Cycle

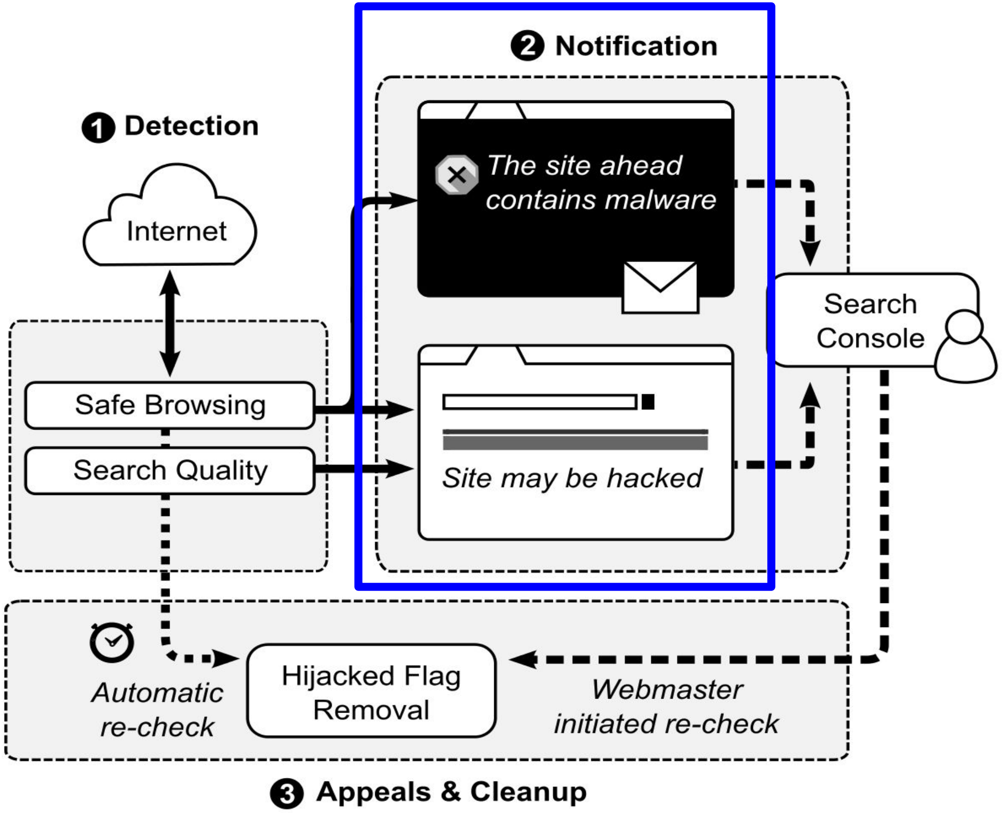

Compromise Life Cycle

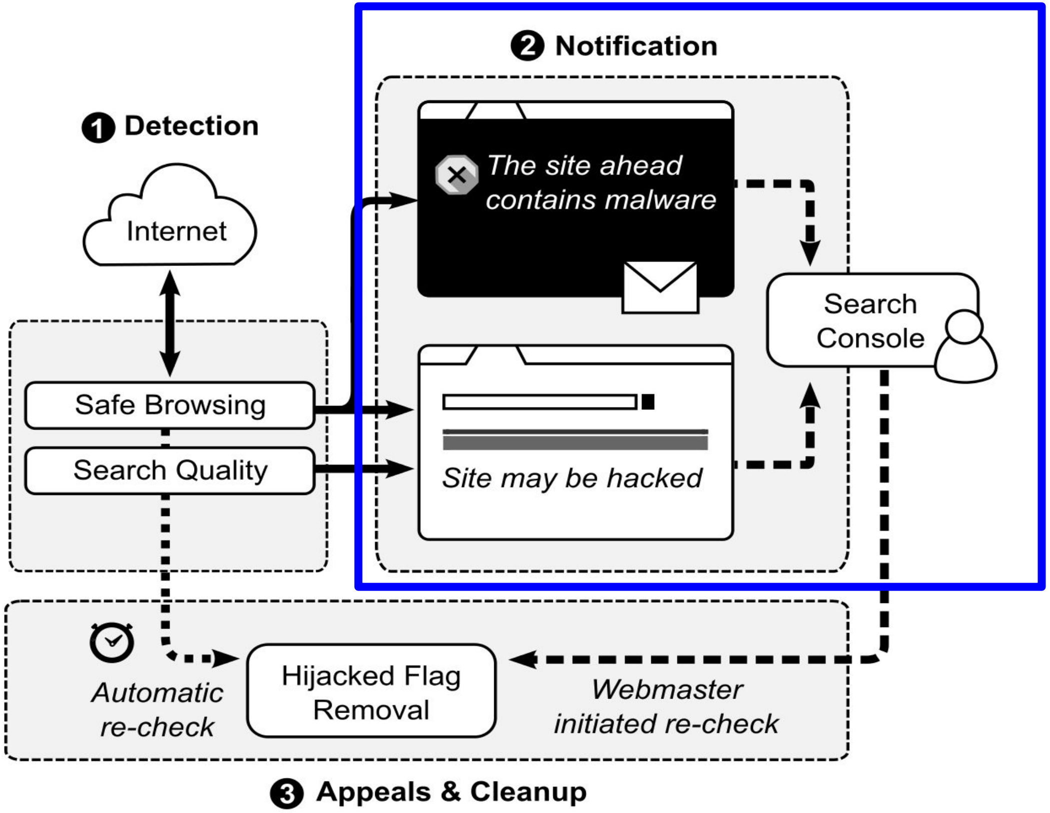

Compromise Life Cycle

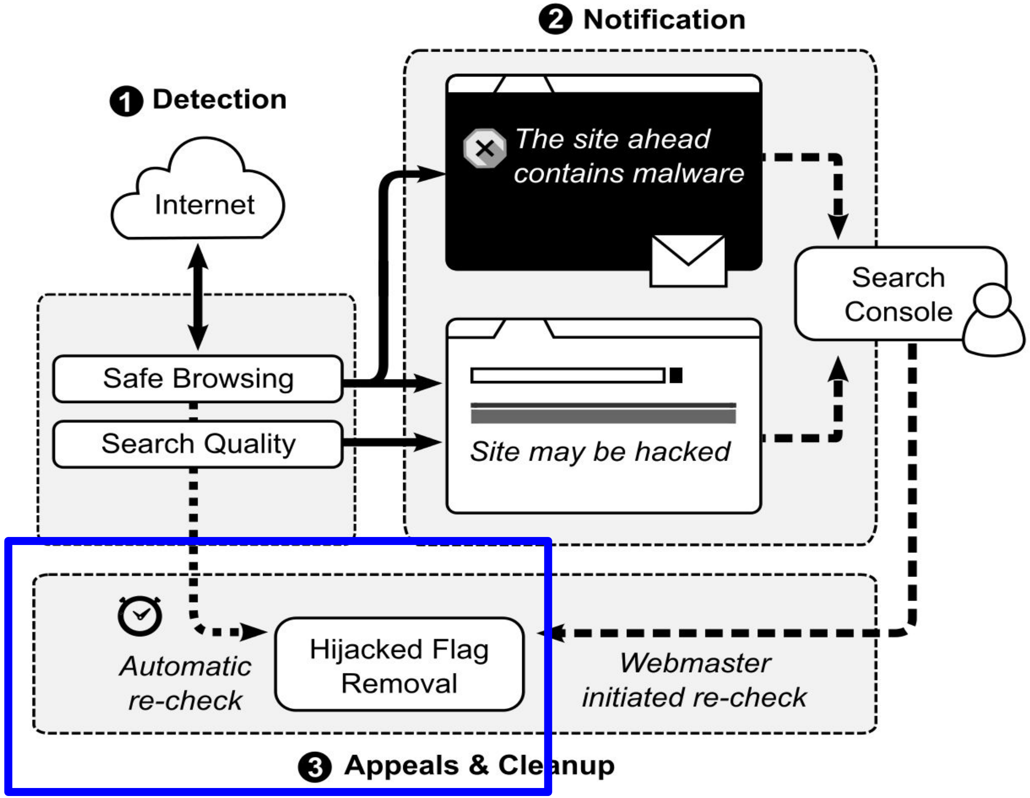

Compromise Life Cycle

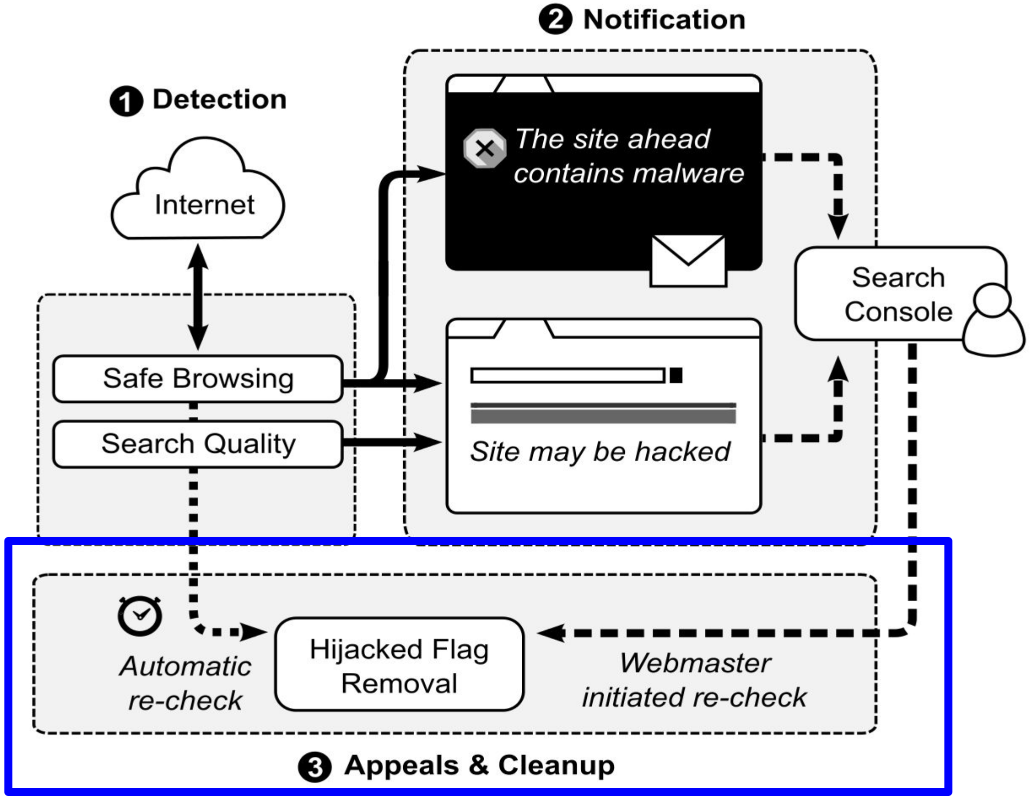

Compromise Life Cycle

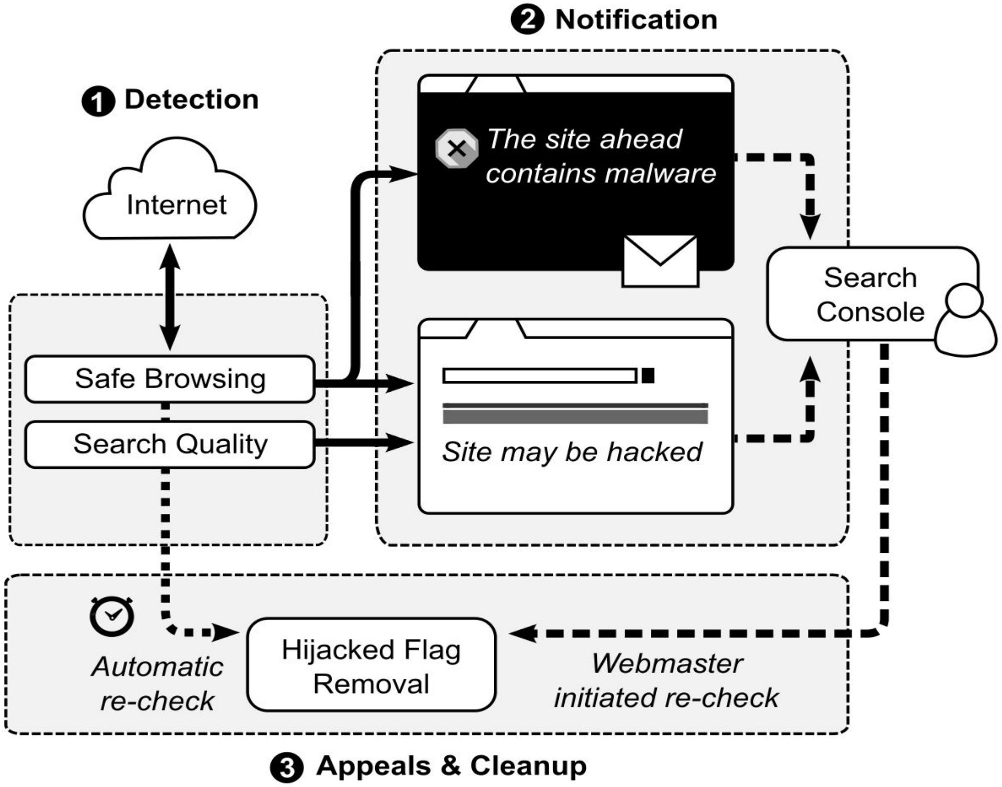

Data Sources

1. Compromised *incidents* detected by Safe Browsing (drive-bys) and Search Quality (blackhat SEO)
2. Search Console + WHOIS alerts sent for hijacked sites
3. Webmaster appeals (requests for re-check)

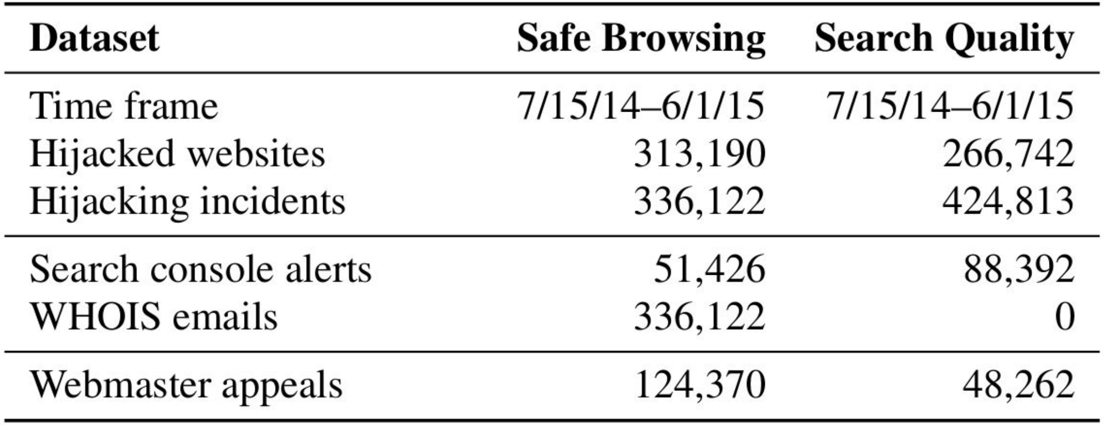

Dataset	Safe Browsing	Search Quality
Time frame	7/15/14–6/1/15	7/15/14–6/1/15
Hijacked websites	313,190	266,742
Hijacking incidents	336,122	424,813
Search console alerts	51,426	88,392
WHOIS emails	336,122	0
Webmaster appeals	124,370	48,262

Notification Effectiveness: Remediation Likelihood

Notification Effectiveness: Remediation Likelihood

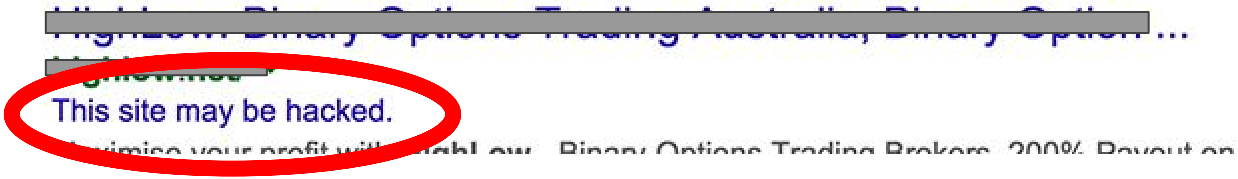

[redacted]

[redacted]

This site may be hacked.

Maximise your profit with HighLow - Binary Options Trading Brokers. 200% Payout on

Search Warning Only (Search Quality sites):

43.4%

Notification Effectiveness: Remediation Likelihood

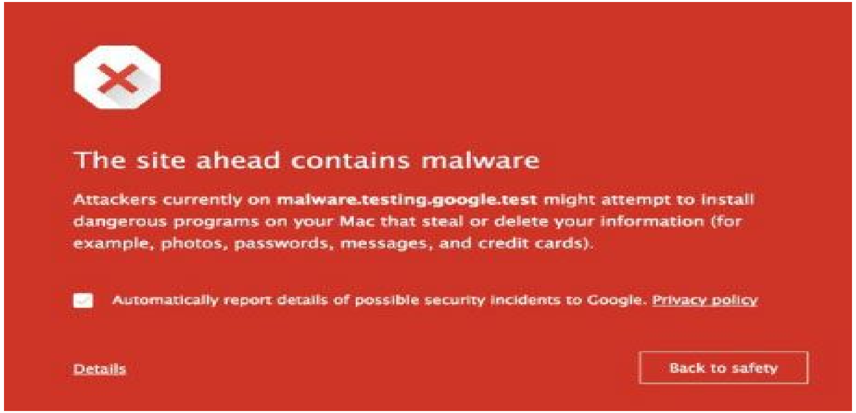

The site ahead contains malware

Attackers currently on **malware.testing.google.test** might attempt to install dangerous programs on your Mac that steal or delete your information (for example, photos, passwords, messages, and credit cards).

☑ Automatically report details of possible security incidents to Google. Privacy policy

Browser Warning + WHOIS alert (Safe Browsing sites):

54.6%

Notification Effectiveness: Remediation Likelihood

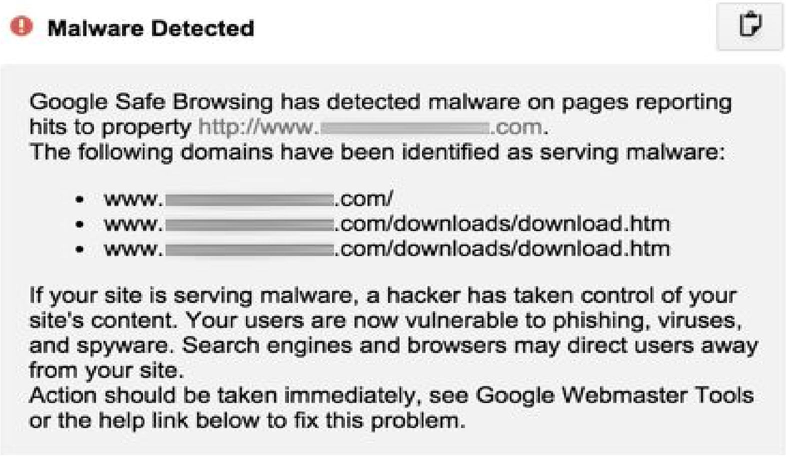

Malware Detected

Google Safe Browsing has detected malware on pages reporting hits to property http://www.[redacted].com.
The following domains have been identified as serving malware:

- www.[redacted].com/
- www.[redacted].com/downloads/download.htm
- www.[redacted].com/downloads/download.htm

If your site is serving malware, a hacker has taken control of your site's content. Your users are now vulnerable to phishing, viruses, and spyware. Search engines and browsers may direct users away from your site.
Action should be taken immediately, see Google Webmaster Tools or the help link below to fix this problem.

Search Console Alert:

82.4% - *Safe Browsing*
76.8% - *Search Quality*

Notification Effectiveness: Remediation Speed

Time for 50% of sites to remediate

Notification Effectiveness: Remediation Speed

Time for 50% of sites to remediate

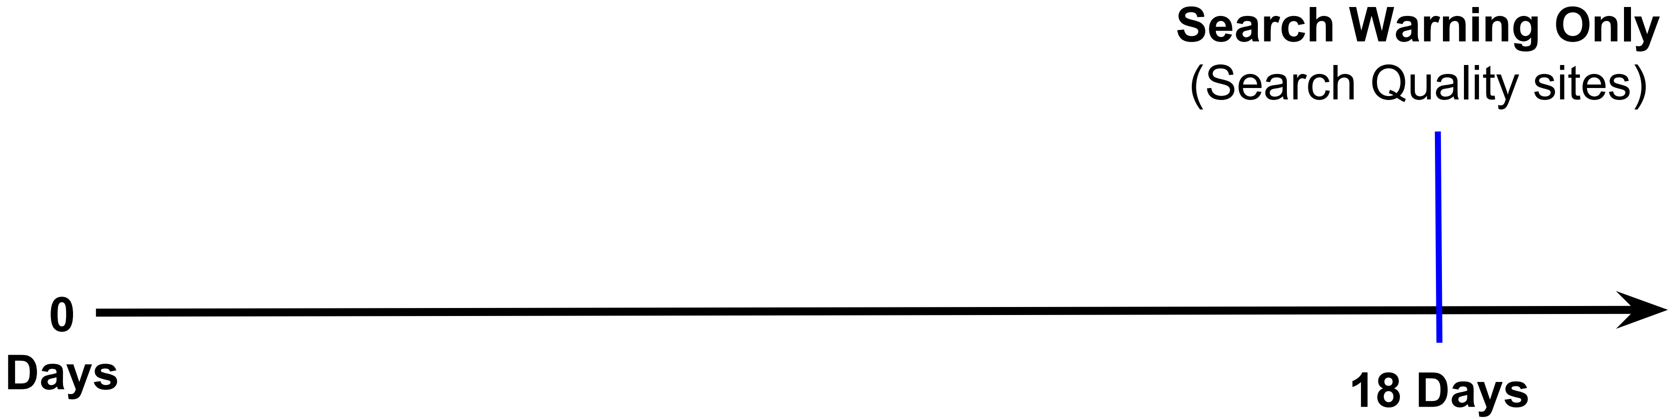

Notification Effectiveness: Remediation Speed

Time for 50% of sites to remediate

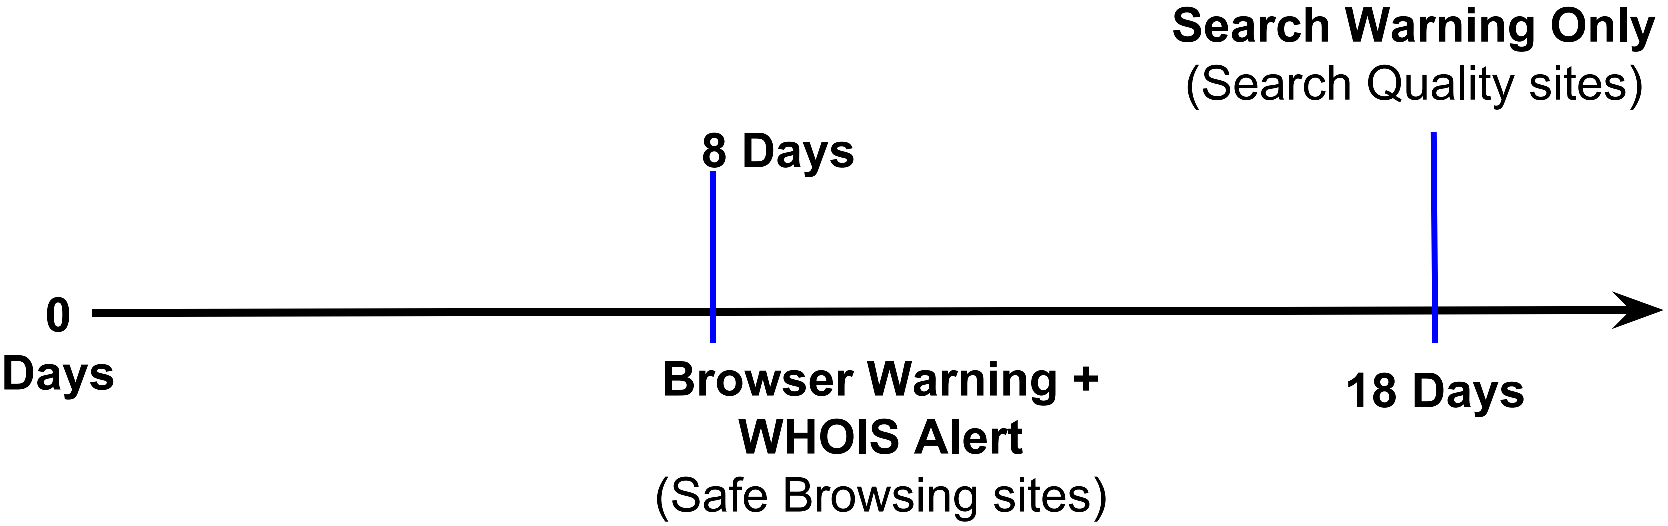

0 Days

8 Days

Browser Warning + WHOIS Alert
(Safe Browsing sites)

Search Warning Only
(Search Quality sites)

18 Days

Notification Effectiveness: Remediation Speed

Time for 50% of sites to remediate

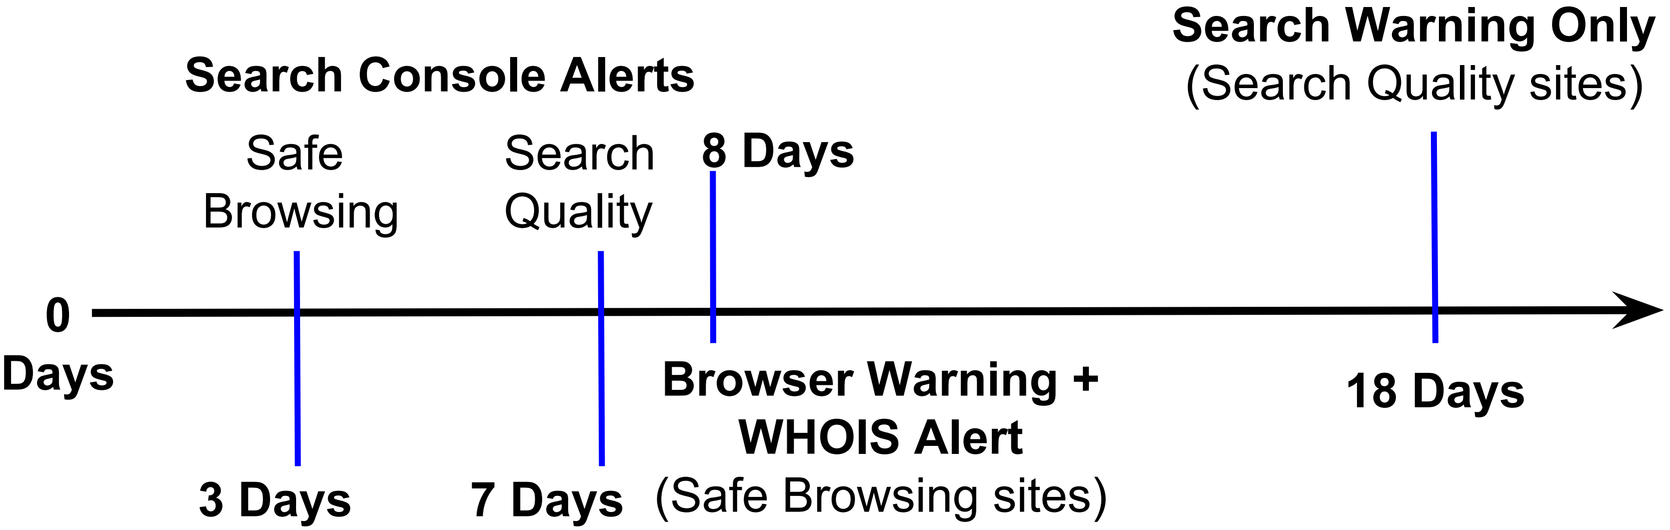

Notification Effectiveness: Remediation Speed

Time for 50% of sites to remediate

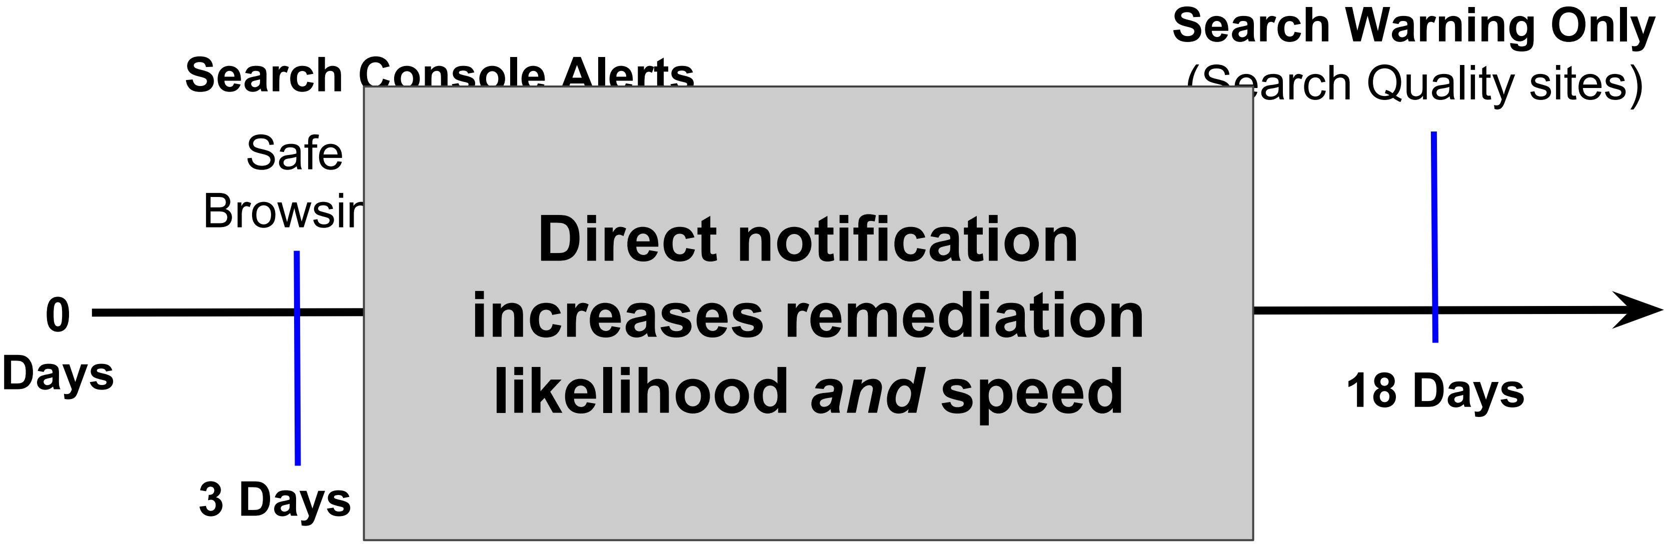

Direct notification increases remediation likelihood *and* speed

Appeals Performance before Success

Appeals Performance before Success

30.7% of Safe Browsing, 11.3% of Search Quality webmasters appeal

Number of Appeals Needed

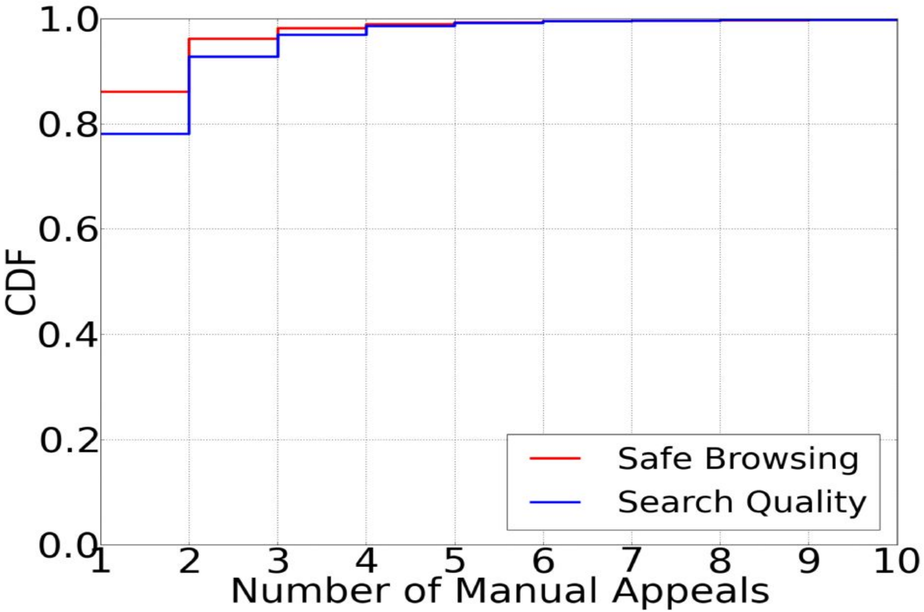

Appeals Performance before Success

30.7% of Safe Browsing, 11.3% of Search Quality webmasters appeal

Number of Appeals Needed

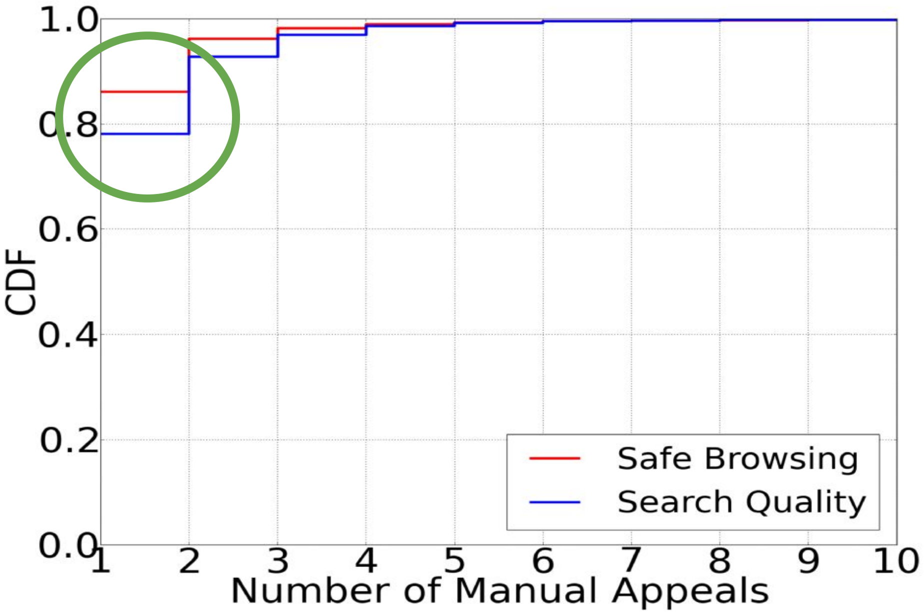

Appeals Performance before Success

30.7% of Safe Browsing, 11.3% of Search Quality webmasters appeal

Number of Appeals Needed

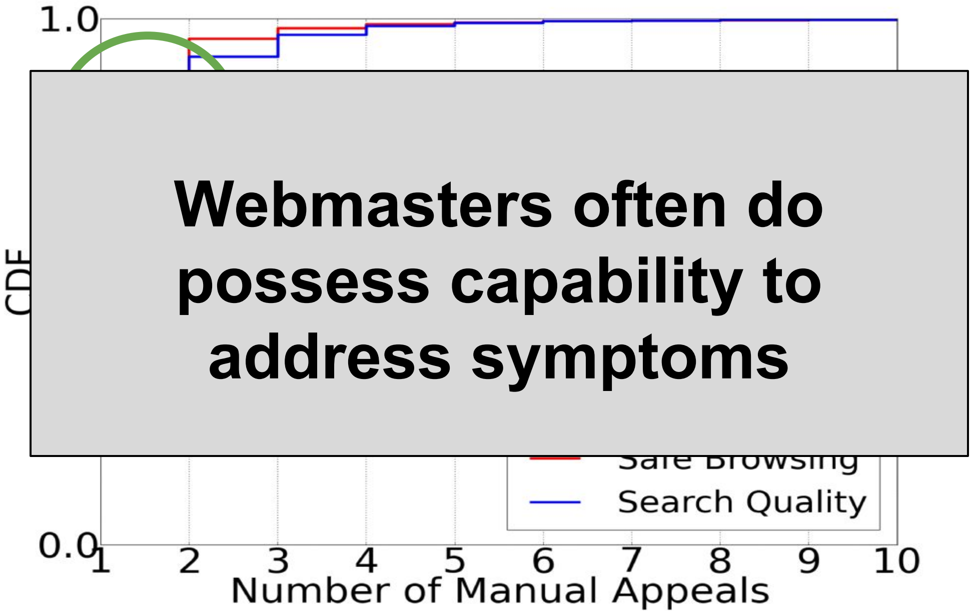

Webmasters often do possess capability to address symptoms

Reinfections

Reinfections

12% of remediated sites are reinfected within 30 days

Reinfections

12% of remediated sites are reinfected within 30 days

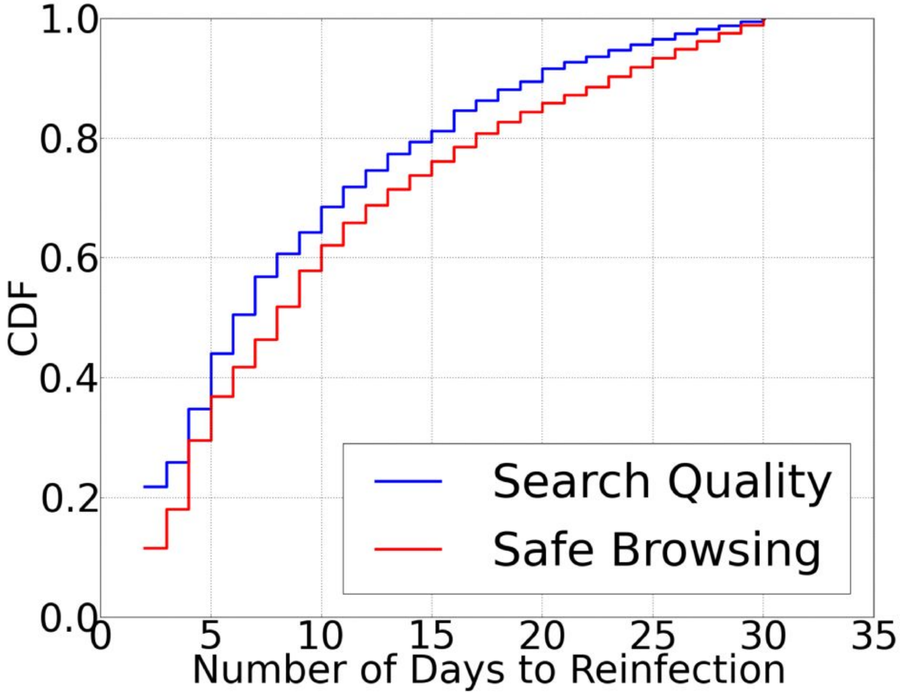

Reinfections

12% of remediated sites are reinfected within 30 days

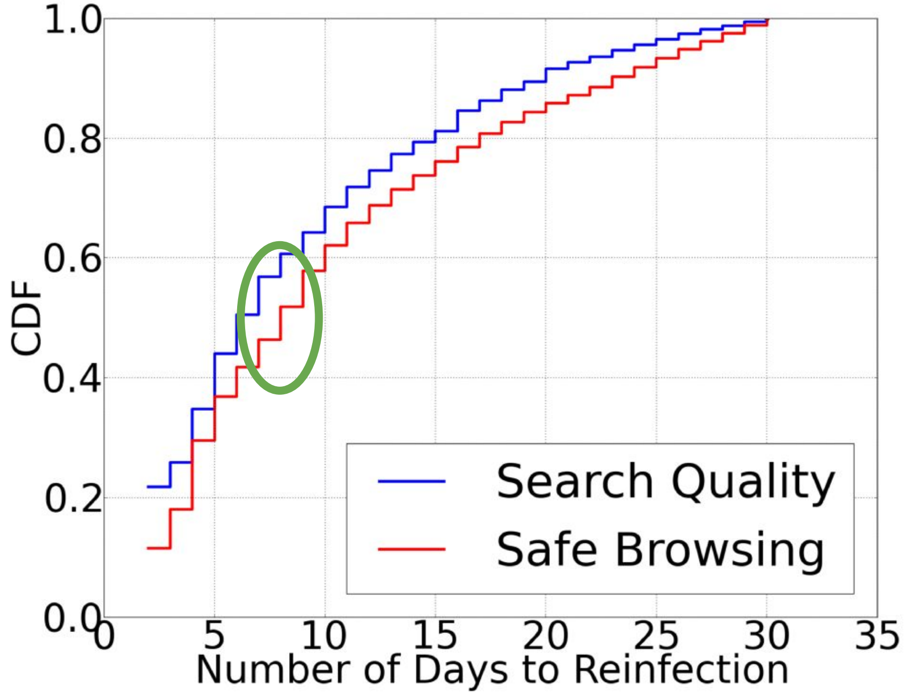

Reinfections

12% of remediated sites are reinfected within 30 days

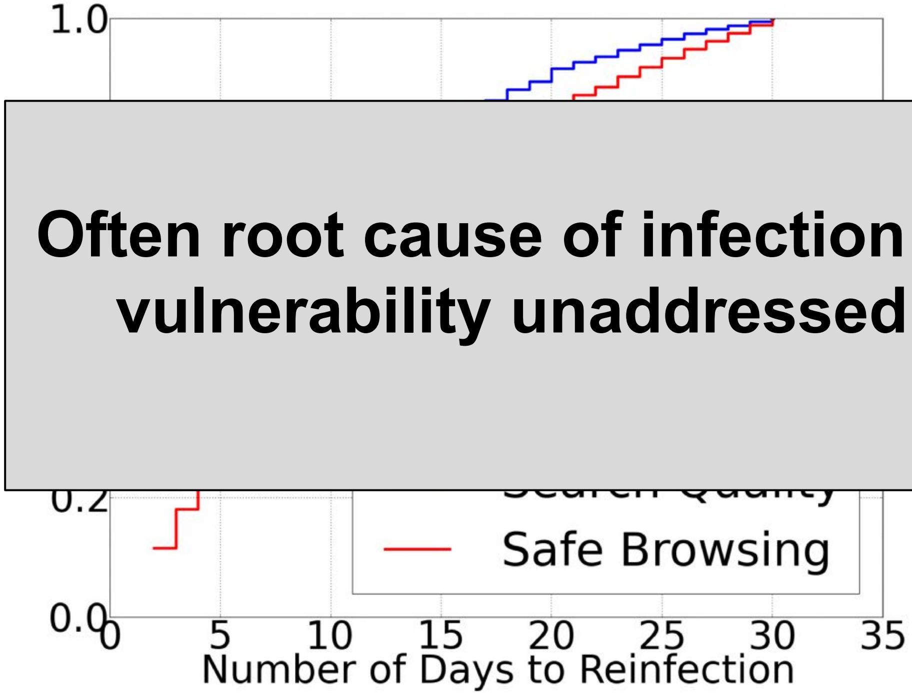

Often root cause of infection or vulnerability unaddressed

Insights

- Direct notifications help improve remediation.

- Webmasters can remedy hijacking symptoms.

- However, root causes are often unaddressed.

Next Steps:

- Increased direct communication coverage

Next Steps:

- Increased direct communication coverage
- Further investigation of notification factors

Next Steps:

- Increased direct communication coverage
- Further investigation of notification factors
- Better community coordination and organization

Next Steps:

- Increased direct communication coverage
- Further investigation of notification factors
- Better community coordination and organization
- Outreach + education

Next Steps:

- Increased direct communication coverage
- Further investigation of notification factors
- Better community coordination and organization
- Outreach + education
- Develop more automated or usable remediation tools

Next Steps:

- Increased direct communication coverage
- Further investigation of notification factors
- Better community coordination and organization
- Outreach + education
- Develop more automated or usable remediation tools

Thanks! frankli@cs.berkeley.edu

Extra Slides

Notification Responses + Reactions

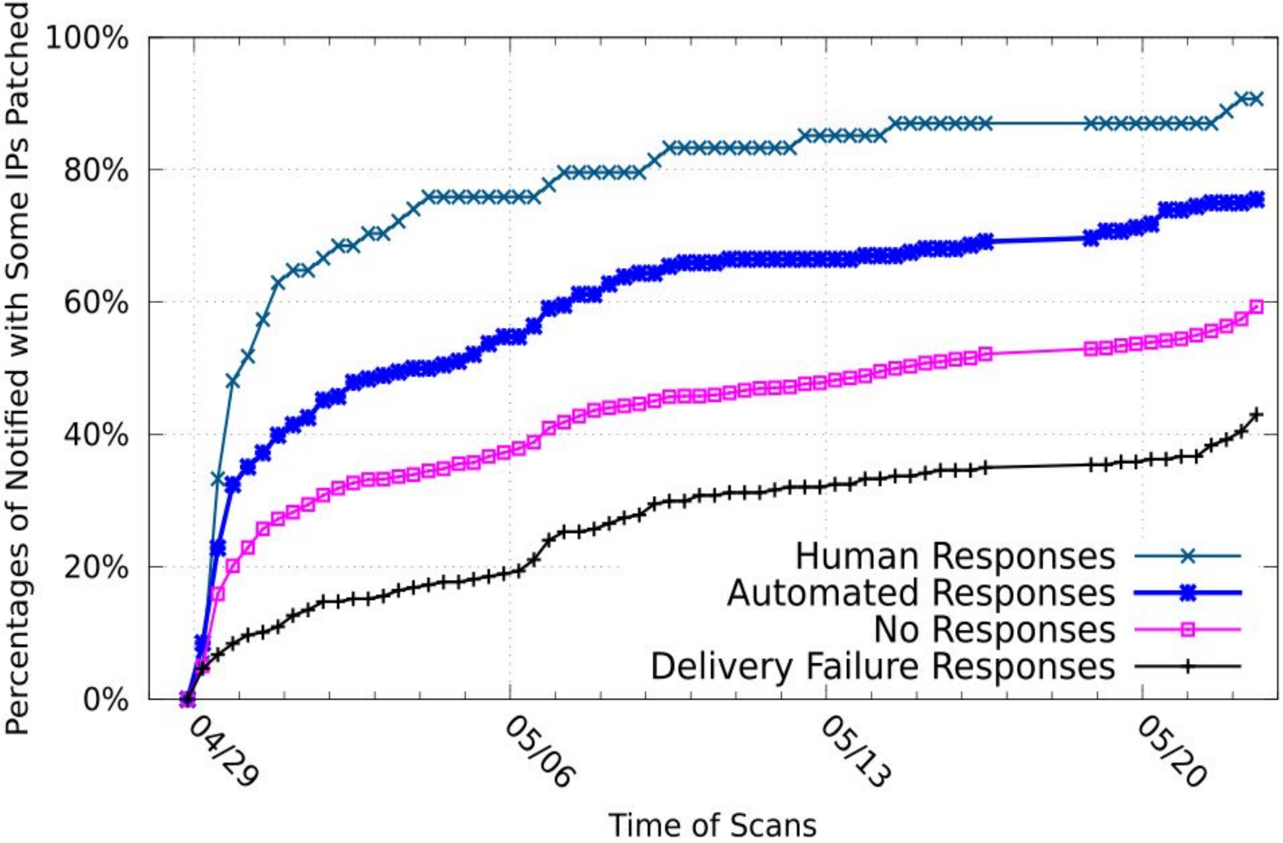

% of Responses:
2.6%
8.5%
77.7%
11.2%

Remediation Rates for CERTs

IPv6

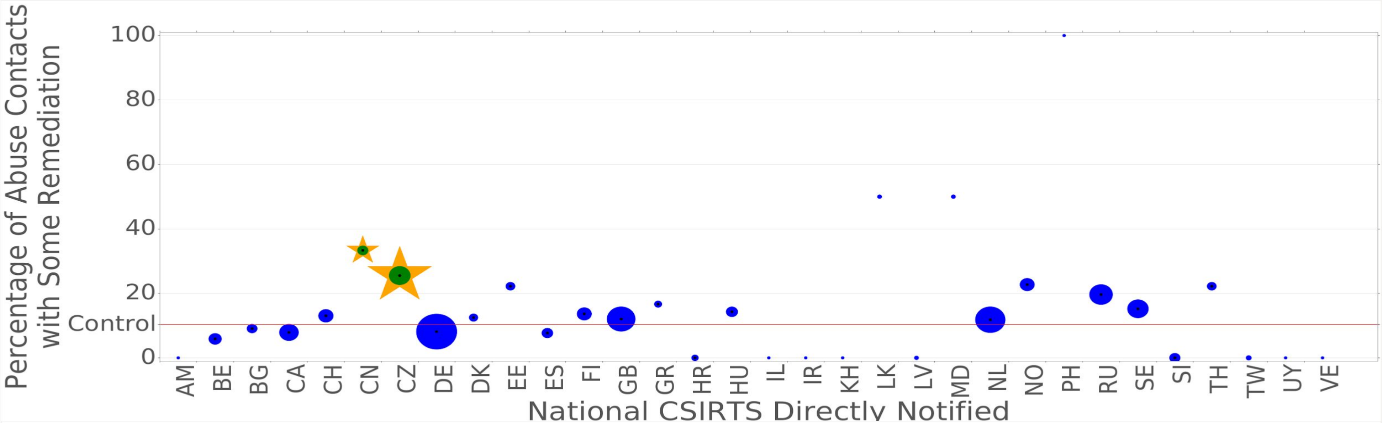

ICS

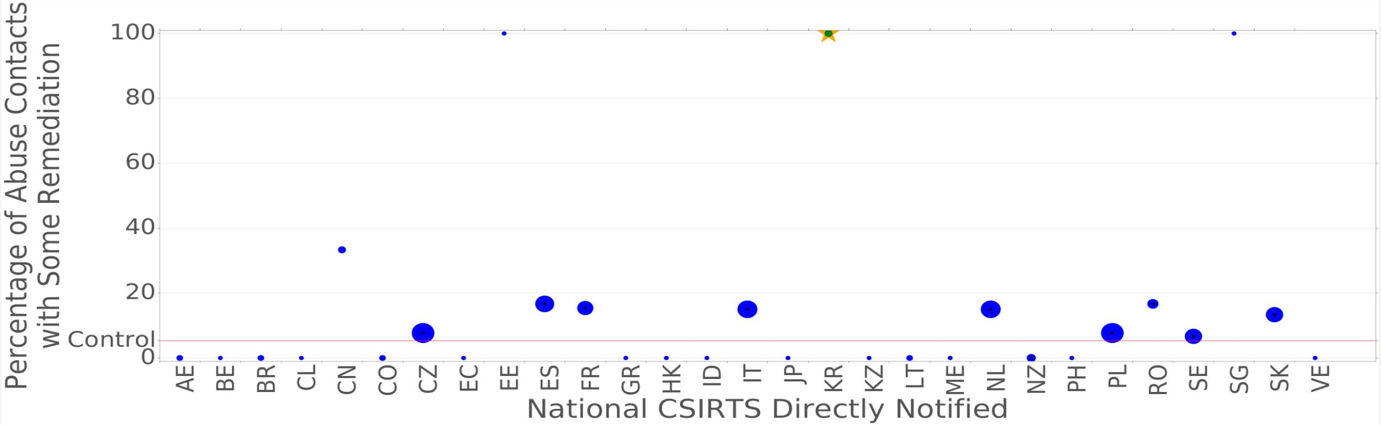
Percentage of Abuse Contacts with Some Remediation
Control
AE
BE
BR
CL
CN
CO
CZ
EC
EE
ES
FR
GR
HK
ID
IT
JP
KR
KZ
LT
ME
NL
NZ
PH
PL
RO
SE
SG
SK
VE
National CSIRTS Directly Notified