

# Hey, You Have a Problem:

## On the Feasibility of Large-Scale Web Vulnerability Notification

---

**Ben Stock**

29th Annual FIRST Conference



# Motivation of our Work

---

- Large-Scale (Web) Vulnerability Detection
  - Drupaggedon SQLi, Joomla! Object Deserialization, Client-Side XSS, Execute After Redirect on Ruby, ...
- Focus previously on Detection, not Notification
- Our work: understand how notifications can work at scale
  - What are suitable communication channels for such a campaign?
  - Does such a campaign affect the prevalence of the notified vulnerabilities?
  - What might inhibiting factors be?
- Today's talk: get insights from the CERT community






## Study Setup

---


# Types of Vulnerabilities

- Well-known vulnerabilities for WordPress (43,865 domains, Top 1M)
  - Reflected Cross-Site Scripting in PIUpload flash component (CVE-2013-0237)
  - Client-Side Cross-Site Scripting in Genericons Example Code (CVE-2015-3429)
  - XMLRPC Multicall Vulnerability
    - allows attacker to try multiple user/password combinations in a single HTTP request
  - Existing patches for all of them
- Previously-unknown Client-Side XSS vulnerabilities (925 domains, Top 10K)
  - Site-specific flaws
  - No existing patches



# Communication Channels

- Direct Communication Channels
  - ~~Web contact forms~~
  - Generic email addresses (info@, security@, webmaster@, abuse@)
  - Domain WHOIS information (registrant or technical contact)
  
- Indirect Communication Channels
  - ~~Vulnerability Reward Programs~~
  - Hosting providers (abuse contacts for the hosting IP range)
  - Trusted Third-Parties
    - regional CERTs (e.g., CERT US, CERT-Bund)
    - FIRST
    - trusted community Ops-Trust



# Notification Procedure

---

- Split up data set of vulnerable domains into five groups of equal size
  - Generic, WHOIS, Provider, TTP, and Control
- Notification via email with link to our Web interface
  - alternatively: access via email using token
- Aggregated Disclosure to providers and TTPs
- Bi-weekly emails
  - January 14th, January 28th, February 11th



# Web Interface

Vulnerability Notification   Legal Disclaimer

We, researchers from the [Center for IT-Security, Privacy and Accountability](#), Saarland University, are conducting a large-scale notification of vulnerable Web applications. To enable the affected parties to fix their sites, we aim to notify them about discovered vulnerabilities in their applications.



This page contains a list of distinct flaws discovered on the domain [REDACTED] and its subdomains. To access the technical details, please follow the *view report* link.

In case you have any questions, please do not hesitate to contact us at [contact@notify.mmci.uni-saarland.de](mailto:contact@notify.mmci.uni-saarland.de).

## Discovered vulnerabilities for [REDACTED]


Type	Subdomain	Vulnerability Id	Last verified at (GMT)	Details
domxss	[REDACTED]	1501	July 21, 2016, 8:56 a.m.	<a href="#">view report</a>





# Reachability Analysis

- Mailbox and accessed reports to classify domains
  - reached: report viewed or email acknowledged
  - bounced: all emails for this domain bounced
  - unreachable: no WHOIS contact, no provider abuse mail, or redirect to Web interface
  - unknown: all others
  - indirect channels: first step of the chain measured





Global Impact of Notification

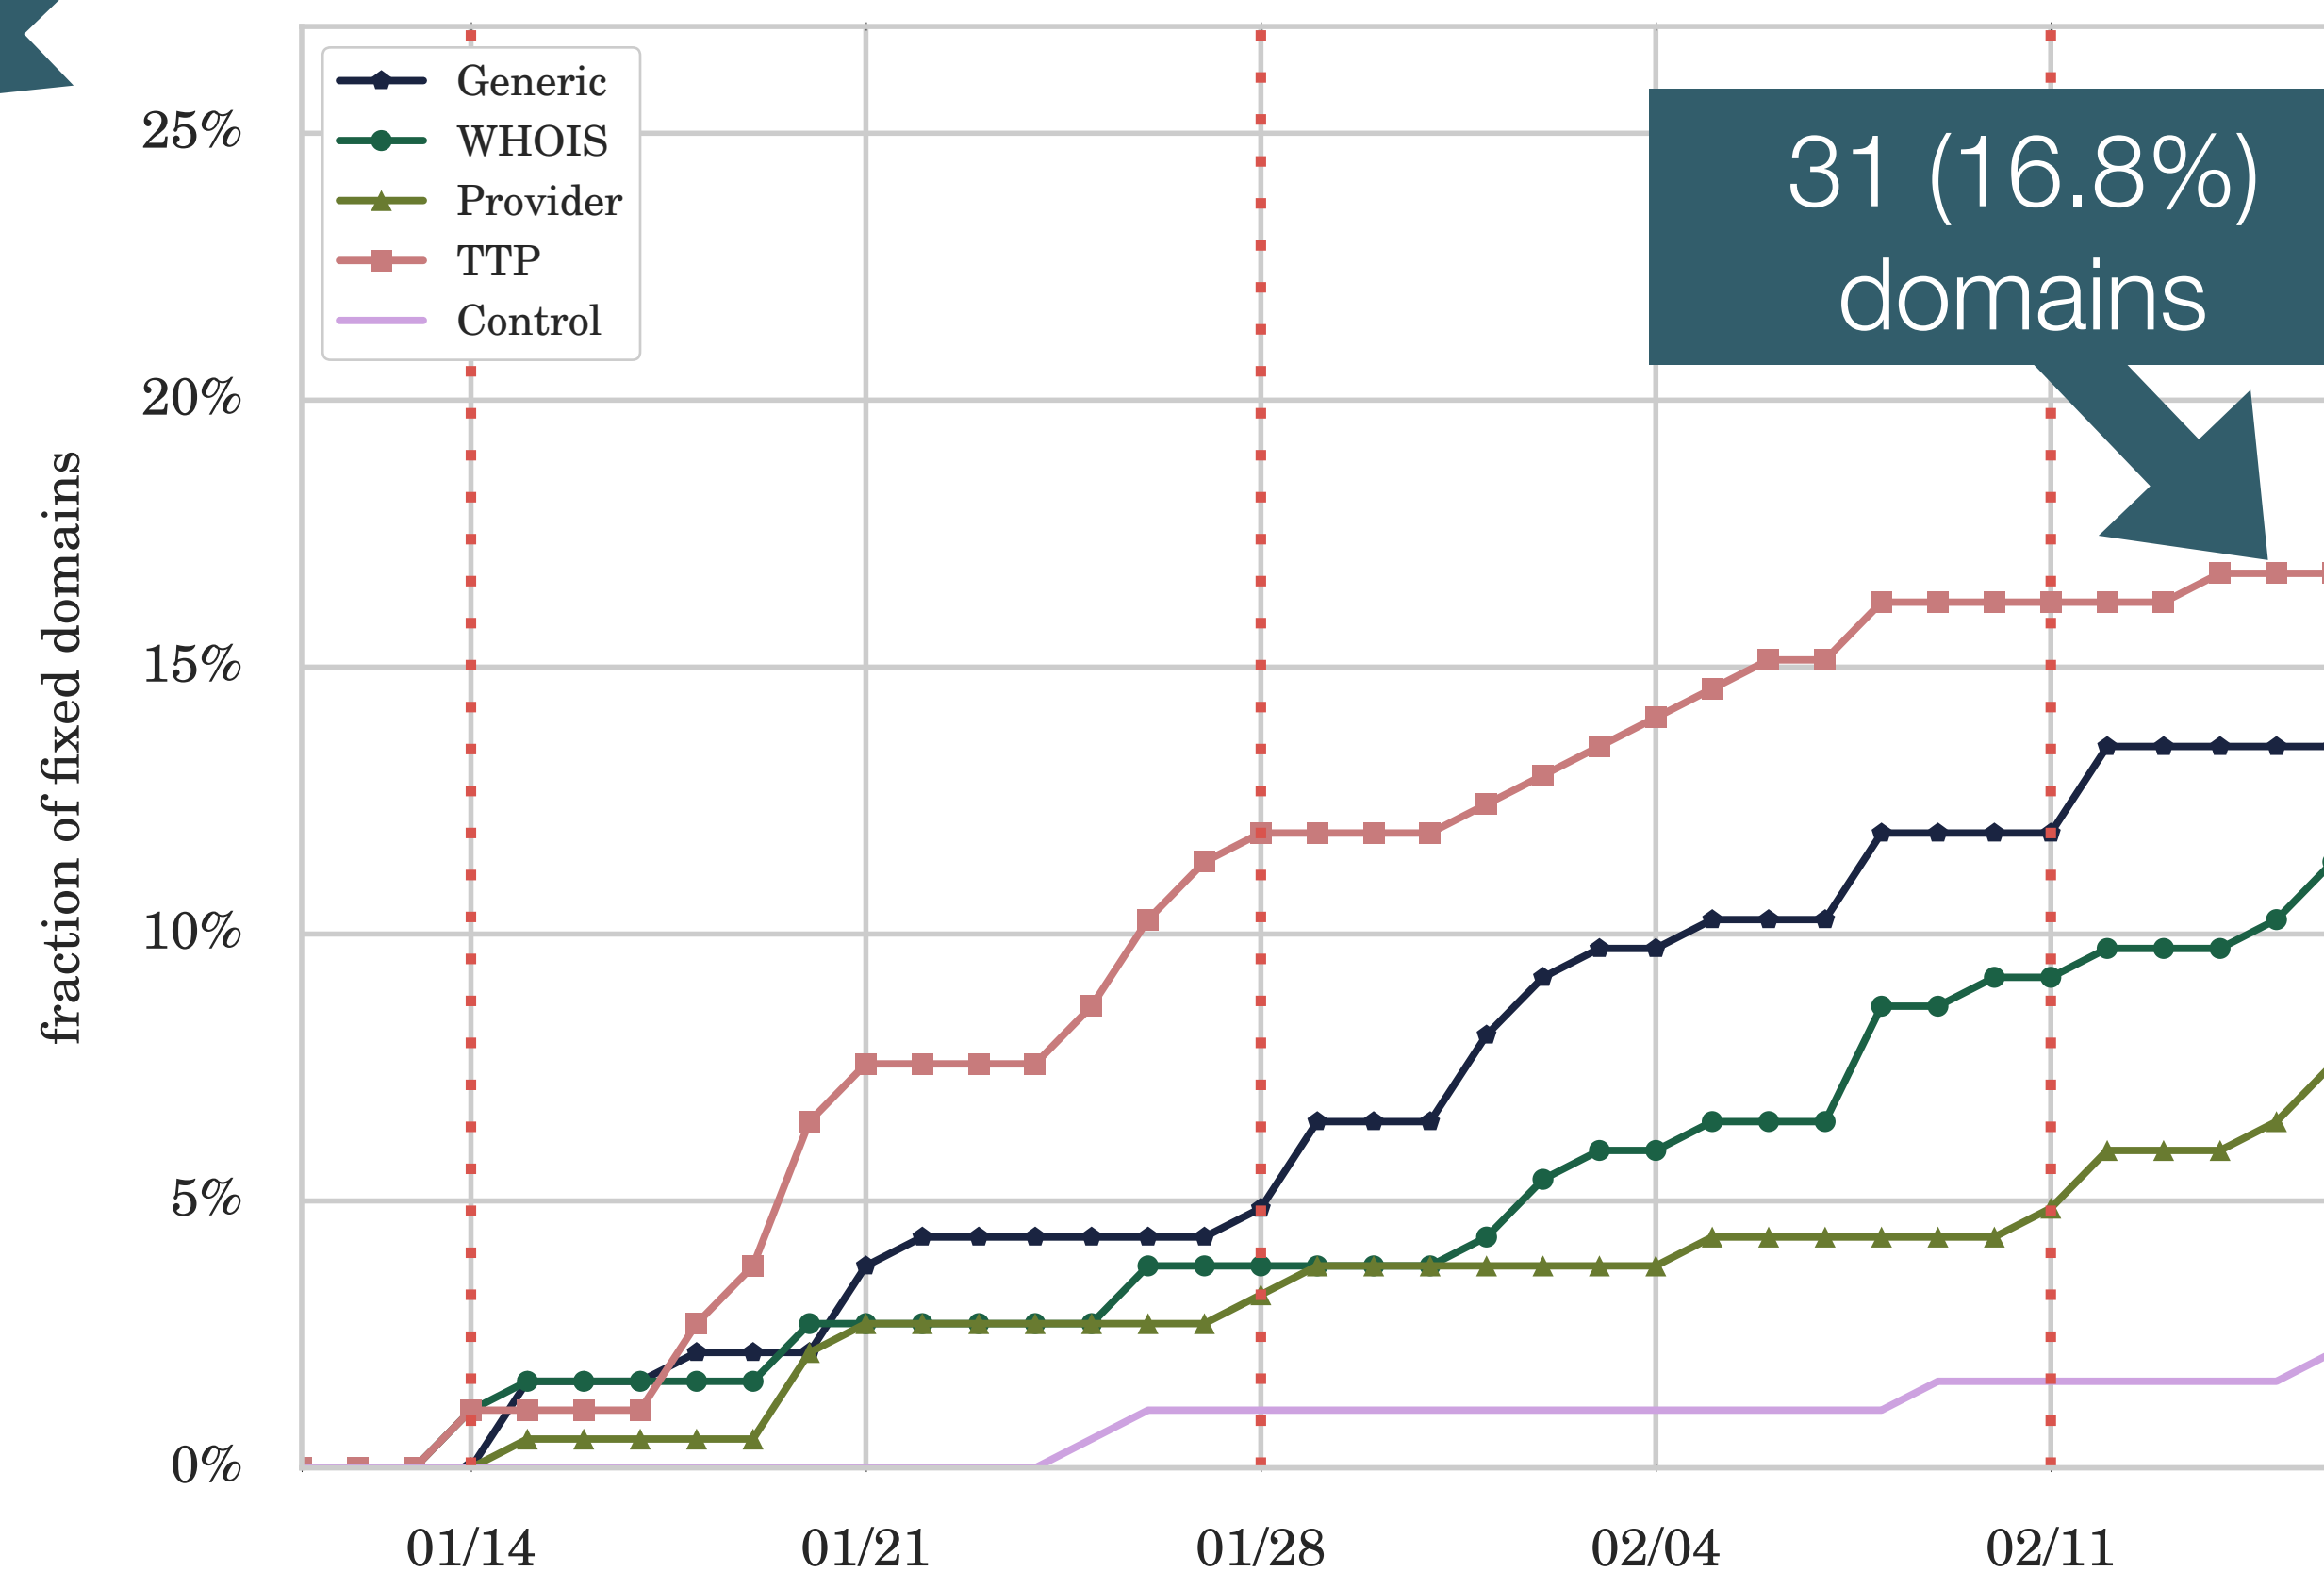

# Fixed Sites over Time

## WordPress


+280 (+3.1%)  
domains

## Client-Side XSS

31 (16.8%)  
domains




„Although the notifications for both WordPress and Client-Side XSS showed **significant improvements** over the **control group**, the number of domains which were fixed is **unsatisfactory** (25.8% and 12.6%, respectively).“




## Communication Channel Analysis

# Reachability of Direct Channels

■ Bounced
 ■ Unreachable
 ■ Unknown
 ■ Reached




# Reachability of Direct Channels




# Reachability of Indirect Channels

■ Bounced
 ■ Unreachable
 ■ Unknown
 ■ Reached

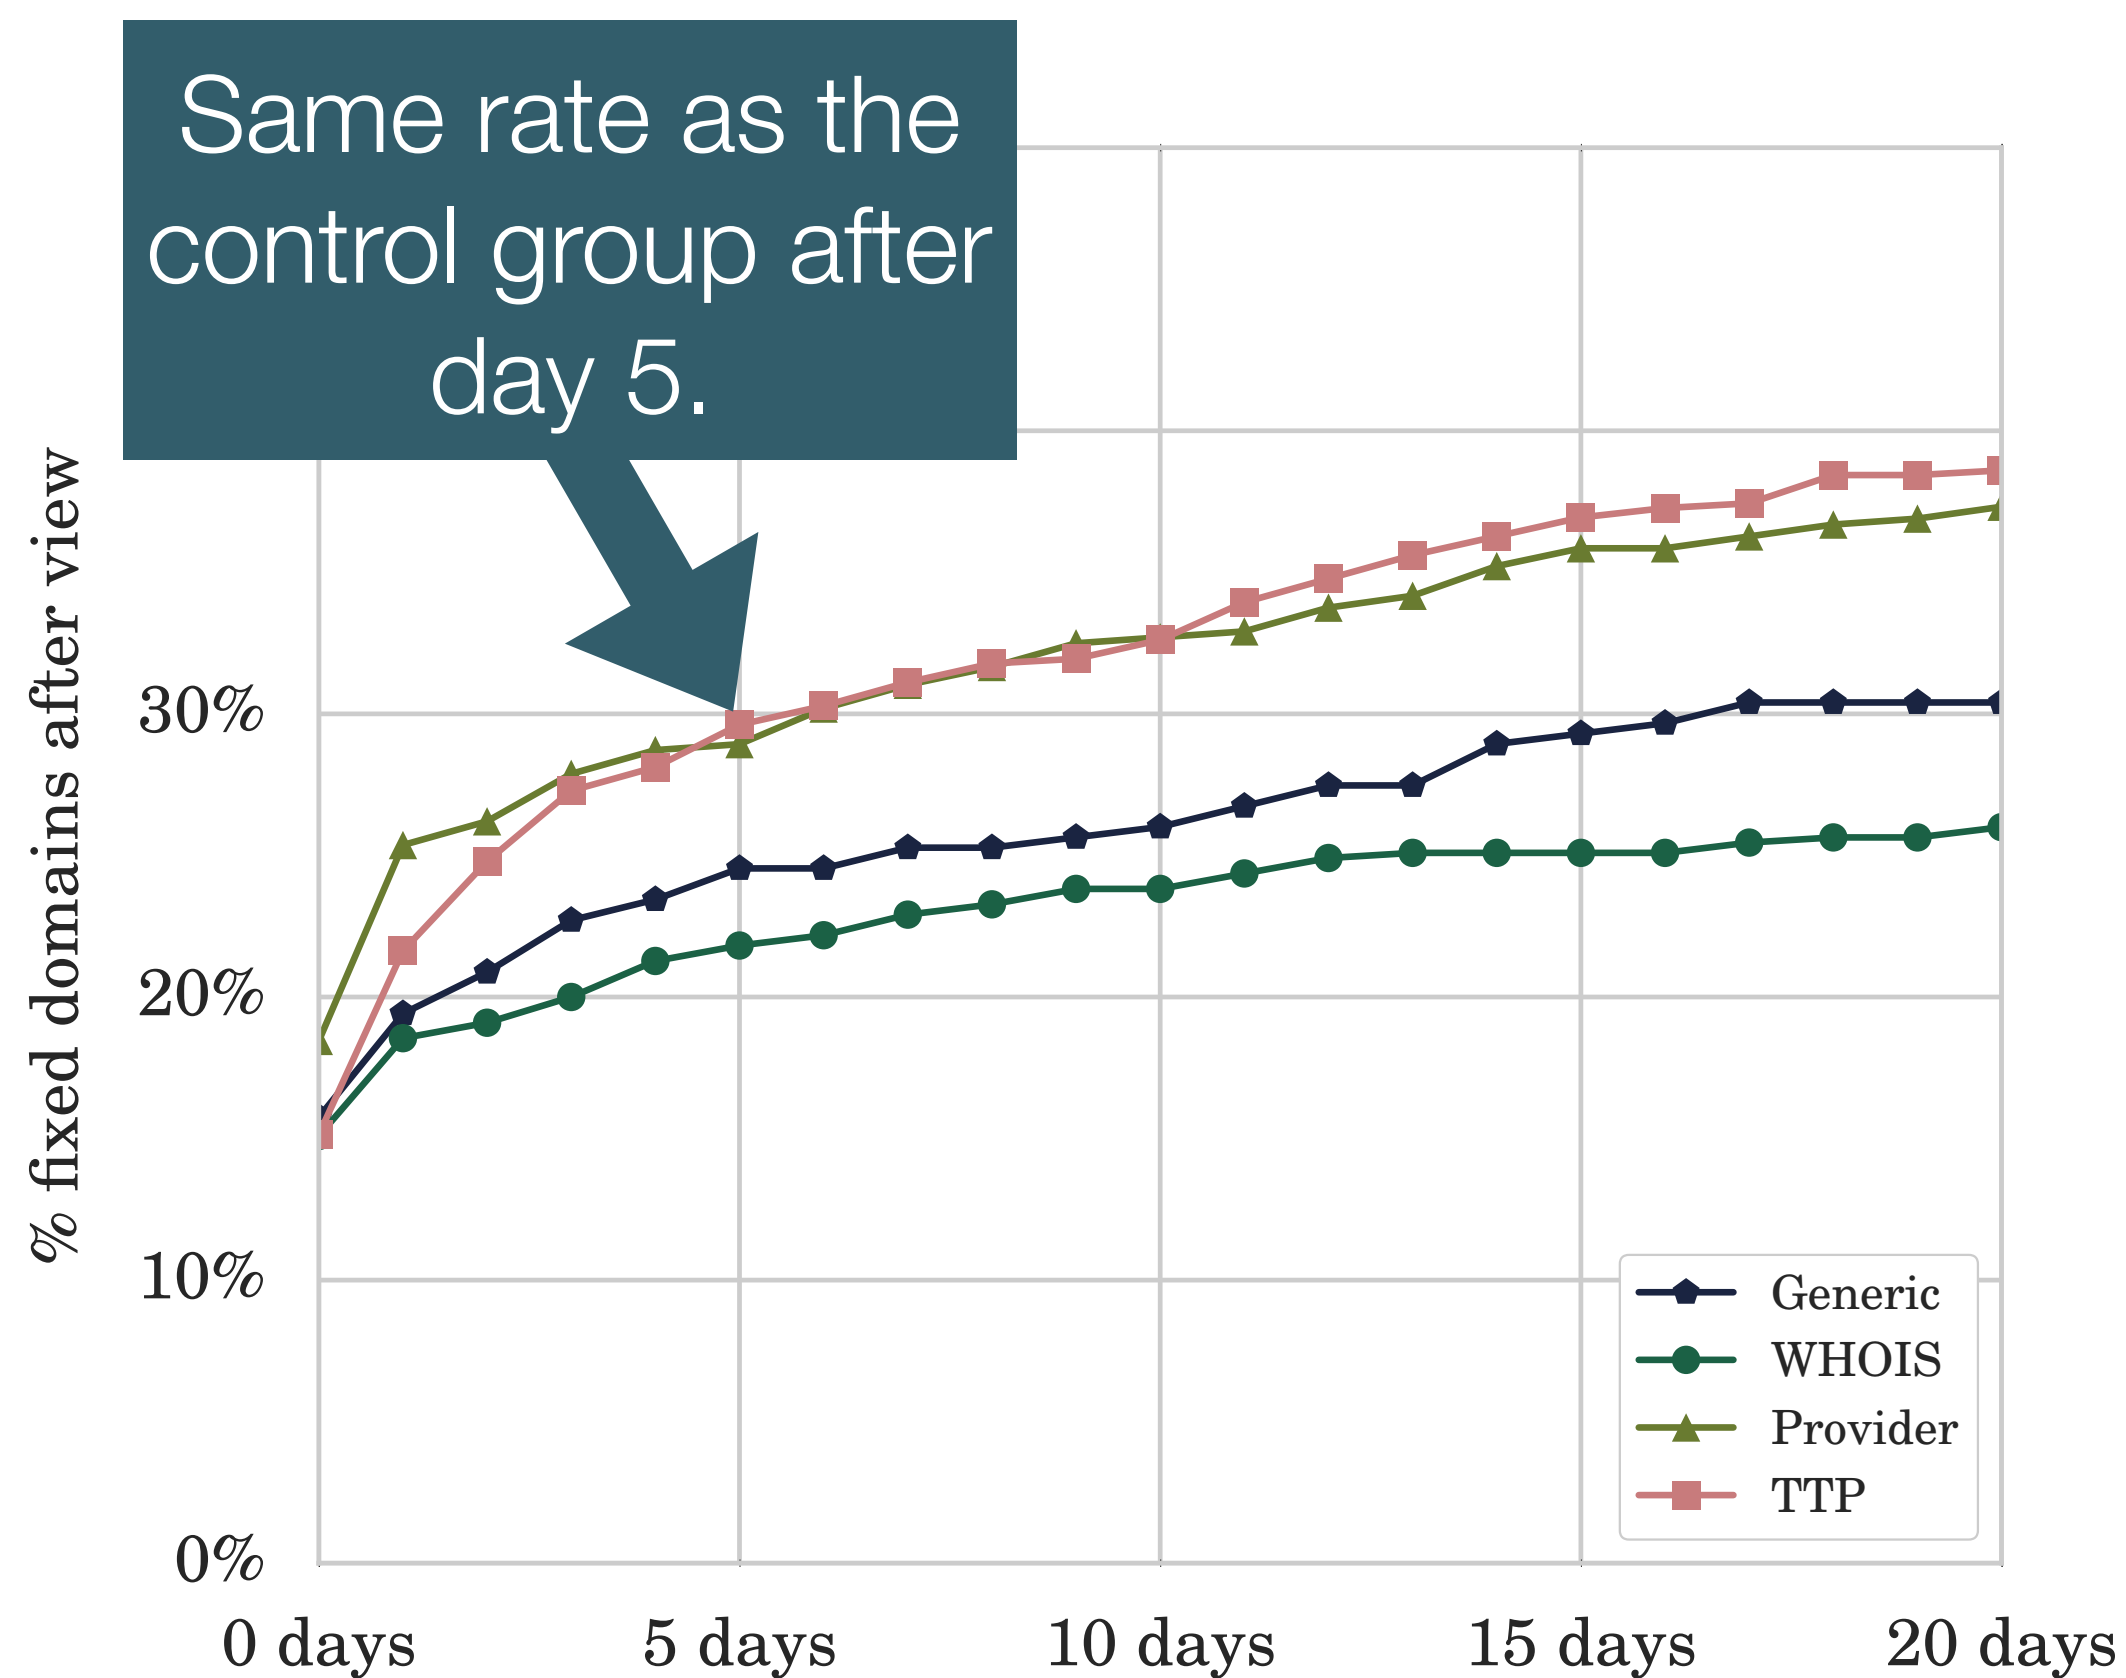





# Reachability of Indirect Channels



# Time to Fix after Report View



WordPress



Client-Side XSS




Key Insights and *Follow-Up Questions*

# Establishing Communication Channels

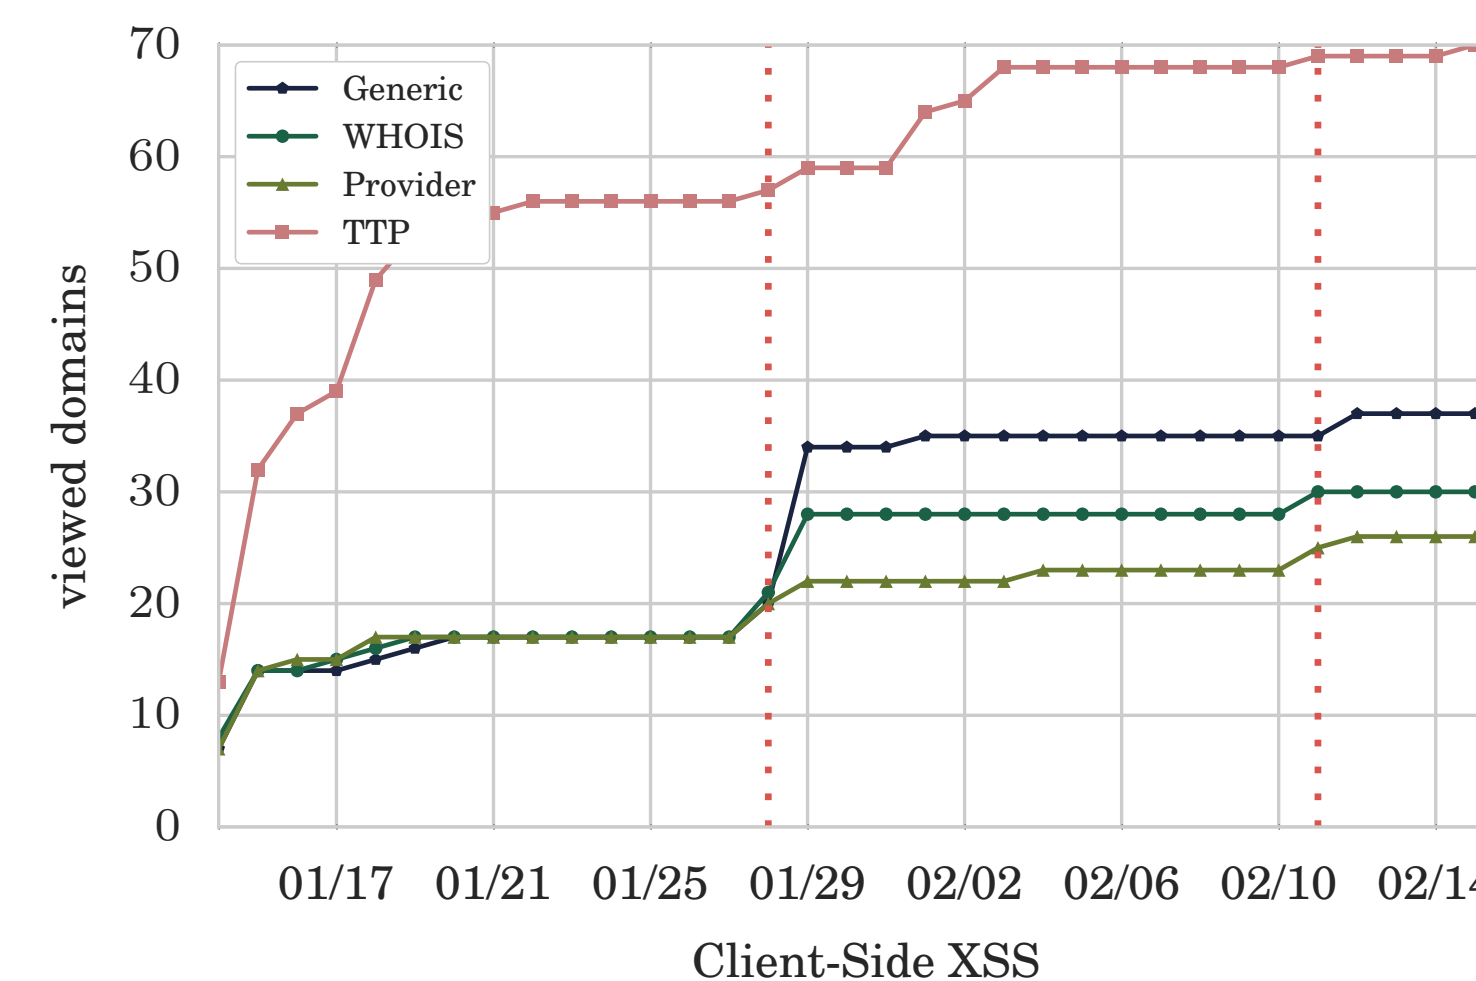

---

- Direct channels are hard to reach
  - generic emails perform really bad for average Web sites
  - WHOIS helps, but is incomplete (~18.5% without entry)
- Indirect channels are easier to „reach“
  - Often do not forward the information
  - top 5 providers (~25% of domains) did not react
- *How can the security community come up with reliable means of establishing communication channels between researchers and affected parties?*



# Need for Reminders and Time to Fix

- Reminders helped especially for direct channels
- Once report was viewed, fix ratio was ~25-30%
  - after five days, WordPress fix rate equaled control group
- Future notification campaigns should make frequent use of reminders
- *How can we improve on the fix ratio?*



# Sender Reputation

- Previous work found that sender reputation does not matter
- Our work begs to differ
  - German CERT more inclined to forward information
  - Providers more inclined to act upon German CERT info
- *What is the impact of the sender reputation, especially when using intermediaries, on the success of a notification campaign?*



# User Distrust

---


- Our experiments required users to click a link
  - or send an email with a token
- Community trains users not to click/react
- Notified control group with full disclosure email
  - results only differed significantly for WordPress
  - BUT: performed worse than with links!
- Potential issue in the message length
- *To what extent does the message tone, content, and length influence the success of notification campaigns?*



# Results Generality

---

- Results appear to be dependent on the domain
  - Providers worked best for Network vulns and Heartbleed
- Even within the same domain, results differ
  - e.g. Generic on WordPress v. Client-Side XSS
- *Are campaigns more successful if the vulnerabilities gained attention in the media (such as Heartbleed)?*
- *Does it matter who needs to fix the vulnerability, be it a network admin, Web site developer, or end-user?*





# Connecting with the FIRST Community

---

- Is it feasible to notify WordPress at scale?
- Should we use different formats, endpoints, ..?
- How could we make the reports more useful?
- *What else can the research community do to ensure that vulnerability notifications can work at scale?*



# Conclusion

---

- Conducted first analysis into notifications for Web vulnerabilities at scale
  - two data sets: well-known (WordPress) and previously-unknown (Client-Side XSS) flaws
  - four communication channels: direct (generic emails, WHOIS) and indirect (providers, TTPs)
- Results show statistically significant improvement caused by our campaign
  - WHOIS worked best for WordPress, TTP best for Client-Side XSS
- Overall improvement was unsatisfactory
  - 74.5% of all domains in data set vulnerable at the end of our experiments
- Main problem is reaching administrators in the first place
  - 30% fix rate within five days (WordPress) / 25% (Client-Side XSS)

Thank you!  
Questions?