You've Got Vulnerability: Exploring Effective Vulnerability Notifications

Frank Li, *University of California, Berkeley;* **Zakir Durumeric,** *University of Michigan, University of Illinois at Urbana–Champaign, and International Computer Science Institute;* **Jakub Czyz,** *University of Michigan;* **Mohammad Karami,** *George Mason University;* **Michael Bailey,** *University of Illinois at Urbana–Champaign;* **Damon McCoy,** *New York University;* **Stefan Savage,** *University of California, San Diego;* **Vern Paxson,** *University of California, Berkeley, and International Computer Science Institute*

https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li

This paper is included in the Proceedings of the 25th USENIX Security Symposium

August 10–12, 2016 • Austin, TX

ISBN 978-1-931971-32-4

Open access to the Proceedings of the 25th USENIX Security Symposium is sponsored by USENIX

You've Got Vulnerability: Exploring Effective Vulnerability Notifications

Frank Li[†] Zakir Durumeric[*‡⋆] Jakub Czyz[*] Mohammad Karami[◇]
Michael Bailey[‡] Damon McCoy[▷] Stefan Savage[○] Vern Paxson[†⋆]

[†]*University of California Berkeley* [*]*University of Michigan* [◇]*George Mason University*
[‡]*University of Illinois Urbana-Champaign* [▷]*New York University*
[○]*University of California San Diego* [⋆]*International Computer Science Institute*

Abstract

Security researchers can send vulnerability notifications to take proactive measures in securing systems at scale. However, the factors affecting a notification's efficacy have not been deeply explored. In this paper, we report on an extensive study of notifying thousands of parties of security issues present within their networks, with an aim of illuminating which fundamental aspects of notifications have the greatest impact on efficacy. The vulnerabilities used to drive our study span a range of protocols and considerations: exposure of industrial control systems; apparent firewall omissions for IPv6-based services; and exploitation of local systems in DDoS amplification attacks. We monitored vulnerable systems for several weeks to determine their rate of remediation. By comparing with experimental controls, we analyze the impact of a number of variables: choice of party to contact (WHOIS abuse contacts versus national CERTs versus US-CERT), message verbosity, hosting an information website linked to in the message, and translating the message into the notified party's local language. We also assess the outcome of the emailing process itself (bounces, automated replies, human replies, silence) and characterize the sentiments and perspectives expressed in both the human replies and an optional anonymous survey that accompanied our notifications.

We find that various notification regimens do result in different outcomes. The best observed process was directly notifying WHOIS contacts with detailed information in the message itself. These notifications had a statistically significant impact on improving remediation, and human replies were largely positive. However, the majority of notified contacts did not take action, and even when they did, remediation was often only partial. Repeat notifications did not further patching. These results are promising but ultimately modest, behooving the security community to more deeply investigate ways to improve the effectiveness of vulnerability notifications.

1 Introduction

A secure Internet ecosystem requires continual discovery and remediation of software vulnerabilities and critical misconfigurations. Security researchers discover thousands of such issues each year, across a myriad of platforms [1]. This process consists of four key phases: (1) discovering new security problems, (2) identifying remedies, (3) determining affected parties, and (4) reaching out to promote remediation among those affected.

The security community has decades of experience with the first two phases, and developments in high-speed scanning [10, 11] and network monitoring [23, 25] have significantly advanced the ease of the third phase for many security issues. However, the process of outreach remains today at best ad hoc. Unlike the public health community, which has carefully studied and developed best practices for patient notification (e.g., [4, 19]), the security community lacks significant insight into the kinds of notification procedures that produce the best outcomes.[1] Instead, for most software, the modern practice of vulnerability notification remains broadcasting messages via well-known mailing lists or websites that administrators must periodically poll and triage.

Given the relative ease with which investigators can today often determine the affected parties, the question then arises of how they should best utilize that information. In the past, performing large-scale notifications was often seen as both ineffective and impractical [2, 7, 12, 18]. However, several recent case studies have provided clear evidence to the contrary. For example, to promote patching of the 2014 OpenSSL Heartbleed vulnerability, Durumeric et al. emailed notices to operators of hosts detected as vulnerable via scanning

[1] An exception concerns the development of online software update systems that explicitly tie together notification and remediation, allowing precise and automated updating targeting the affected parties. Unfortunately, the vast majority of software lacks such systems; even for those that do, operators may disable it in some contexts (critical servers, embedded systems) to avoid unplanned downtime.

and found that notified operators patched at a rate almost 50% greater than a control group [9]. Similarly, Li et al. analyzed the efforts of Google Safe Browsing and Search Quality in reaching out to operators of compromised websites, and found that direct communication with webmasters increased the likelihood of cleanup by over 50%, and reduced infection durations by more than 60% [14].

With these clear indications that notifications *can* drive positive security outcomes, it behooves the security community to determine how to best conduct the outreach efforts. At the same time, we must balance the benefits to the ecosystem (and the associated ethical responsibilities to notify) against the burden this imposes on the reporter, which calls for determining notification regimens that will not prove unduly taxing.

In this work, we strive to lay the foundations for systematically determining the most effective notification regimens, seeking to inform and drive the development of "best practices" for the community. The solution space has many more dimensions than we can hope to methodically explore in a single study. Here, we aim to develop soundly supported results for the most salient basic issues, with an eye towards then facilitating follow-on work that builds on these findings to further map out additional considerations. The issues we address include (1) who to notify (e.g., WHOIS contacts versus national CERTs versus US-CERT), (2) the role of notification content (e.g., do reporters need to devise detailed messages or do short ones suffice), (3) the importance of localization (e.g., what role does native language play in notification response rates), and (4) how these considerations vary with the nature of the vulnerability (including whether for some vulnerabilities notification appears hopeless).

We evaluate these questions empirically in the context of notification campaigns spanning three different vulnerability categories: publicly accessible industrial control systems, misconfigured IPv6 firewalls, and DDoS amplifiers. Using large-scale Internet scanning to identify vulnerable hosts and then monitor their behavior over time post-notification, we infer the effects of different notification regimes as revealed by the proportion and timeliness of contacts remediating their vulnerable hosts.

Our results indicate that notifications can have a significant positive effect on patching, with the best messaging regimen being directly notifying WHOIS contacts with detailed information within the message itself. An additional 11% of contacts addressed the security issue when notified in this fashion, compared to a control. However, we failed to push the majority of contacts to take action, and even when they did, remediation was often only partial. Repeat notifications did not further patching. We additionally characterize the responses we received through our notification campaigns, of which 96% of human-sent responses were positive or neutral. Given these promising yet modest findings, it behooves the security community to more deeply investigate vulnerability notifications and ways to improve their efficacy. Our methodology and results form the basis for establishing initial guidelines to help drive future efforts.

2 Related Work

Several recent studies have found that large-scale security notifications increase patching and remediation—particularly for infected websites.

Vasek et al. notified 161 infected websites [24] and found that after 16 days, 55% of notified sites cleaned up compared to 45% of unnotified sites. They further note that more detailed notifications outperformed reports with minimal information by 13%, resulting in a 62% cleanup rate. Cetin et al. performed a similar study, measuring the role of sender reputation when notifying the owners of hijacked websites [5]. They emailed the WHOIS contacts of 240 infected sites from email addresses belonging to an individual independent researcher (low reputation), a university research group (medium reputation), and an anti-malware organization (high reputation). While nearly twice as many notified sites cleaned up within 16 days compared to unnotified ones, they found no significant differences across the various senders.

On a larger scale, Li et al. investigated the life cycles of 761 K website hijacking incidents identified by Google Safe Browsing and Search Quality [14]. They found that direct notifications to webmasters increased the likelihood of cleanup by over 50% and reduced infection lengths by 60% on average. Absent this communication, they observed that browser interstitials—while intended to protect browser users—correlated with faster remediation.

Most similar to the vulnerabilities we investigate, Durumeric et al. used Internet-wide scanning to track the Heartbleed vulnerability and notified system owners two weeks after public disclosure [9]. Their notifications drove a nearly 50% increase in patching compared to a control: 39.5% versus 26.8%.

Concurrent to this work, Stock et al. investigated the feasibility of large-scale notifications for web vulnerabilities [22]. Similar to our study, they experimentally evaluated the effectiveness of different communication channels, including WHOIS email contacts and CERTs. Additionally, they analyzed the reachability and viewing behavior of their messages. Their results largely accord with ours, providing a complementary study of notifications in a separate context (namely, vulnerable websites). Notably, they likewise observed that while notifications

Dataset	Hosts	WHOIS Abuse Contacts	Hosts with WHOIS Contacts
ICS	45,770	2,563	79.7%
IPv6	180,611	3,536	99.8%
Ampl.	83,846	5.960	92.4%

Table 1: Vulnerable Hosts—We notified network operators about three classes of vulnerabilities found in recent studies: publicly accessible industrial control systems (ICS), hosts with misaligned IPv4 and IPv6 firewall policies, and DDoS amplifiers (NTP, DNS, and Chargen).

could induce a statistically significant increase in patching, the raw impact was small. In the best case, only an additional 15% of the population patched compared with a control group.

Each of these studies has established that notifications can increase vulnerability patching and cleanup. We build on these works and explore the next critical step: understanding what factors influence patching and how to construct effective vulnerability notifications.

3 Methodology

To measure notification efficacy and to understand how to construct effective notifications, we notified network operators while varying aspects of the notification process. In this section, we detail the datasets of vulnerable hosts, the variables we tested, and how we tracked remediation.

3.1 Vulnerable Hosts

We notified operators about the three classes of vulnerabilities listed below. We show the population of vulnerable hosts in Table 1.

Publicly Accessible Industrial Control Systems Industrial control systems (ICS) are pervasive and control physical infrastructure ranging from manufacturing plants to environmental monitoring systems in commercial buildings. These systems communicate over a myriad of domain and manufacturer specific protocols, which were later layered on Ethernet and TCP/IP to facilitate long distance communication. Never designed to be publicly accessible on the Internet, these protocols lack important security features, such as basic authentication and encryption, but nonetheless are frequently found unsecured on the public Internet. To identify vulnerable ICS devices, Mirian et al. extended ZMap [10] and Censys [8] to complete full IPv4 scans for several ICS protocols: DNP3, Modbus, BACnet, Tridium Fox, and Siemens S7 [17]. In total, they found upwards of 46 K ICS hosts that were publicly accessible and inherently vulnerable.

We coordinated with Mirian et al. to complete daily scans for each protocol against the public IPv4 address space from January 22–24, 2016. We limited our study to the 45.8 K hosts that were present all three days to reduce the noise due to IP churn. To track the impact of our notifications, we continued the daily scans of these hosts using the same methodology.

Misconfigured IPv6 Firewall Policies Czyz et al. found that 26% of IPv4/IPv6 dual-stack servers and routers have more permissive IPv6 firewall policies compared to IPv4, including for BGP, DNS, FTP, HTTP, HTTPS, ICMP, MySQL, NTP, RDP, SMB, SNMPv2, SSH, and Telnet access [6]. For example, twice as many routers have SSH accessible over IPv6 compared to IPv4. Given the presumed rarity of IPv6-only services, this likely indicates a misconfiguration and potential security issue.

To identify dual-stack servers, Czyz et al. looked for hostnames in the Rapid7 DNS ANY dataset [20] that had both A and AAAA records. After filtering out automatically generated hostnames, they identified 520 K dual-stack servers. To find routers, the team performed reverse DNS lookups and subsequent A and AAAA lookups for hosts in the CAIDA Ark dataset [3], identifying 25 K routers. Czyz et al. then scanned these hosts using Scamper [15] to identify firewall inconsistencies.

We scanned the hosts that Czyz et al. identified over a 25 day period from December 31, 2015 to January 24, 2016. We limited our study to the 8.4 K routers and 172.2 K servers that were consistently available during that period. Similar to the ICS measurements, we continued to perform daily scans using the same methodology to track the impact of our notifications.

DDoS Amplifiers Several UDP protocols allow attackers to launch distributed denial of service attacks when improperly configured [21]. In this scenario, an attacker spoofs a small request to a misconfigured server, which then sends a large response to the victim. For example, an attacker can spoof a DNS lookup to a recursive DNS resolver, which will then send the full recursive lookup to the victim's machine. We identified 152 K misconfigured hosts that were actively being used to launch DDoS attacks over NTP, DNS, and Chargen by monitoring the sources of DDoS attacks against a university network between December 11–20, 2015.

We restricted our notifications to the vulnerable hosts that were consistently available during our daily scans from December 21, 2015 to January 26, 2016. In total, we discovered 5.9 K Chargen amplifiers, 6.4 K NTP amplifiers, and 71.5 K DNS amplifiers on 83.8 K distinct IP addresses. We continued to track these hosts by perform-

ing daily protocol scans (e.g., Chargen requests, NTP monlist commands, and DNS recursive lookups).

In each case, we coordinated with the studies' authors to ensure that they did not simultaneously notify operators. However, we do note that groups have previously sent notifications to DDoS amplifiers [13].

3.2 Experiment Variables

To understand how to best construct and route notification messages, we performed notifications using several methodologies and measured the differences in remediation. We specifically aimed to answer the following questions:

Who should researchers contact? Researchers have several options when deciding where they should report vulnerabilities, including directly contacting network operators, notifying national CERTs, and asking their own country's CERT to disseminate the data to other CERT groups. We tested three options: (1) notifying the abuse contact from the corresponding WHOIS record, (2) geolocating the host and contacting the associated national CERT, and (3) asking our regional CERT (US-CERT) to propagate the information.

How verbose do messages need to be? It is not clear how much information researchers need to include when notifying operators. For example, are notifications more effective if researchers include detailed remediation steps or will such instructions go unheeded? We sent three types of messages: (1) a terse message that briefly explained that we discovered the vulnerability with Internet-wide scanning, and the impact of the vulnerability (e.g., for ICS notifications, we wrote "These devices frequently have no built-in security and their public exposure may place physical equipment at risk for attack."), (2) a terse message with a link to a website with detailed information, and (3) a verbose email that included text on how we detected the problem, vulnerability details, and potential remediation steps. We provide the full text of our different messages in Appendix B–G.

Do messages need to be translated? We tested sending messages in English as well as messages translated by native technical speakers to several local languages.

3.3 Group Assignment

To test the impact of our experiment variables, we randomly formed experiment groups that received different notification regimens. Here we describe our process for constructing these groups.

For each IP address, we extracted the abuse contact from the most specific network allocation's WHOIS record. For the 16.7% of dual-stack hosts with different contacts extracted from IPv4 and IPv6 WHOIS records, we used the contact with the deepest level of allocation, and preferred IPv6 contacts when all else was equal (4.3% of dual-stack hosts).

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
WHOIS: Language – Terse			
Germany: German		71	
Germany: English		72	
Netherlands: Dutch		32	
Netherlands: English		32	
Poland: Polish			37
Poland: English			37
Russia: Russian			123
Russia: English			123
WHOIS: Language – Verbose			
Germany: German		70	
Germany: English		72	
Netherlands: Dutch		32	
Netherlands: English		29	
Poland: Polish			36
Poland: English			36
Russia: Russian			123
Russia: English			123

Table 2: Notification Groups—We aggregated vulnerable hosts by WHOIS abuse contacts and randomly assigned these contacts to notification groups. Here, we show the number of contacts notified in each group. Note that for the language experiments, we tested terse and verbose messages for several countries, both translated and in English.

To test each variable, we split the abuse contacts from each vulnerability into treatment groups (Table 2). For the ICS and amplifier experiments, we randomly allocated one quarter of abuse contacts to the control group (Group 1), one quarter to the CERT groups (half US-CERT, half national CERTs), and the remaining half to the WHOIS groups. For IPv6, to act in a responsible manner we needed to complete *some* form of notification for all hosts to ensure adequate disclosure prior to the release of the corresponding study [6] in February 2016. This prevented us from using a true control group. Instead, we approximate the behavior of the control group using the 25 days of daily scans prior to our notifications. We allocated a third of the IPv6 contacts to the CERT groups, and the remainder to the WHOIS groups.

For the vulnerable hosts assigned to the CERT groups,

we geolocated each IP using MaxMind [16] and identified the associated CERT. We note that not all countries have an established CERT organization. This was the case for 2,151 (17%) IPv6 hosts, 175 (8%) ICS devices, and 2,156 (19%) DDoS amplifiers. These hosts were located in 16 countries for IPv6, 26 countries for ICS, and 63 countries for DDoS. Many of these countries are in Africa or Central America (e.g., Botswana, Ethiopia, and Belize), or are smaller island states (e.g., American Samoa, Antigua and Barbuda, and the Bahamas). We did not include hosts without a CERT organization in the CERT experiment (although we later passed them along to US-CERT).

In total, 64 CERTs were responsible for IPv6 hosts, 57 for ICS, and 86 for amplifiers. To compare directly contacting national CERTs versus having US-CERT distribute information to them, we randomly divided the affected national CERTs into two halves. For national CERTs in the first half, we contacted them directly with vulnerable hosts in their region (Group 2). We sent the remaining hosts for CERTs in the second half to US-CERT (Group 3).

We obtained native translations of our WHOIS messages for several countries. We allocated contacts in the WHOIS groups that were in those countries (based on the WHOIS records) for our language experiment, further detailed in Section 4.3. The remaining contacts were randomly split into three groups based on message verbosity: terse (Group 4), terse with a link (Group 5), and verbose (Group 6).

3.4 Notification Process

We sent notification emails with the FROM and REPLY-TO header set to an institutional mailing list: *security-notifications@berkeley.edu*. In each message, we attached a CSV file that contained the list of vulnerable hosts along with the latest scan timestamp and the list of vulnerable protocols. We also included a link to an anonymous survey, which asked for the organization's perspective on the reported security issue and whether they found our detection and notification acceptable. The messages were sent from a server in UC Berkeley's network, which was listed as a valid mail server by UC Berkeley's SPF policy. We note that we also included a randomly generated identifier in each email subject that enabled us to match a reply to the original notification.

3.5 Tracking Remediation

We tracked the impact of different notification methodologies by scanning all hosts for several weeks following our notifications. As our scanning methods tested the reachability of several services, we may have falsely identified a host as patched due to random packet loss or temporary network disruptions. To account for this, we only designated a host as patched if it did not appear vulnerable in any subsequent scans. We leveraged the last day's scan data for this correction, but did not otherwise use it in our analysis as it lacked subsequent data for validation.

One limitation in our tracking is the inability to distinguish true patching from network churn, where the host went offline or changed its IP address. While we can still conduct a comparative analysis against our control group, we acknowledge that our definition of patching is a mixture of true patching and churn. We investigated whether we could better approximate true remediation by distinguishing between RST packets and dropped packets. We compared the proportion of RSTs and drops between our control group and our notified groups two days after notification and two weeks after notification. At both times, we observed nearly identical proportions between the control and notified groups—in all cases less than 20% of hosts sent RST packets. This indicates that RST packets are not a reliable signal for remediation, as most hosts did not send RST packets even when truly fixed.

Unless stated otherwise, we consider a host as having taken remediation steps for a particular vulnerability if any of its affected protocols were detected as fixed. Likewise, we say a notification contact has taken remediation steps if any of its hosts have patched. We define the remediation rate as the percentage of notification contacts that have taken remediation steps. This definition is over contacts rather than hosts as we are measuring the impact of notifying these contacts, and contacts differ in the number of affected hosts.

3.6 Ethical Considerations

We followed the guidelines for ethical scanning behavior outlined by Durumeric et al. [10]: we signaled the benign intent of our scans through WHOIS entries and DNS records, and provided project details on a website on each scanning host. We respected scanning opt-out requests and extensively tested scanning methods prior to their deployment.

The ethics of performing vulnerability notifications have not been widely discussed in the security community. We argue that the potential good from informing vulnerable hosts outweighs the risks. To minimize potential harm, we only contacted abuse emails using addresses available in public databases. Additionally, we messaged all unnotified contacts at the conclusion of the study. We offered a channel for feedback through an anonymous survey with questions about the notified organization (described in Appendix A). We note that be-

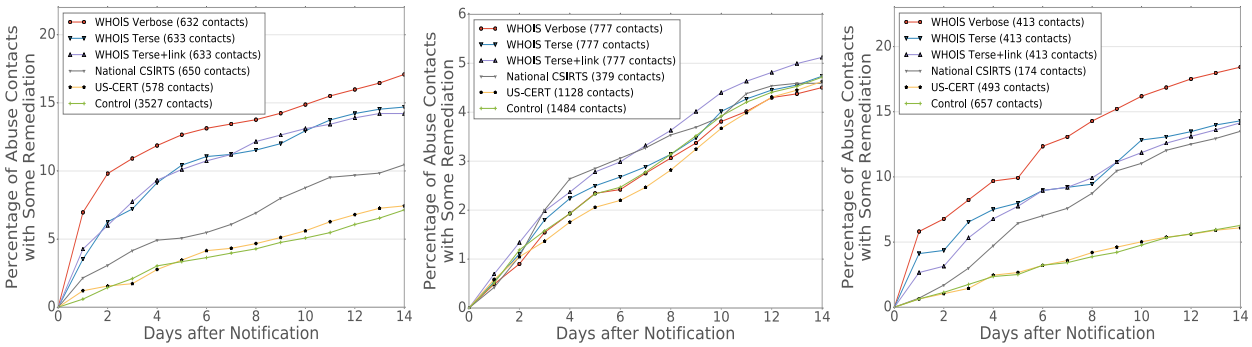

(a) Misconfigured IPv6 **(b)** DDoS Amplifiers **(c)** ICS Services

Figure 1: Remediation Rates—We show the remediation rate for each variable we tested. We find that verbose English notifications sent to network operators were most effective for IPv6 and ICS. Note the varying Y axes.

cause we only collected data about organizational decisions and not individuals, our study did not constitute human subjects research (confirmed by consulting the UC Berkeley IRB committee). Nevertheless, we followed best practices, e.g., our survey was anonymous and optional.

4 Results

For both ICS and IPv6, our notifications had a significant impact on patch rates. In our most successful trial—verbose English messages sent directly to operators—the patch rate for IPv6 contacts was 140% higher than in the control group after two weeks. For ICS, the patch rate was 200% higher. However, as can be seen in Figure 1b, none of our notifications had significant impact on DDoS amplifiers. This is likely due to the extensive attention DDoS amplifiers have already received in the network operator community, including several prior notification efforts [21]. In addition, these amplifiers were already previously abused in DDoS attacks without administrative responses, potentially indicating a population with poor security stances. It is also important to note that our best notification regimen resulted in at most 18% of the population remediating. Thus, while notifications can significantly improve patching, the raw impact is limited. In the remainder of this section, we discuss the impact of each experiment variable and how this informs how we should construct future notifications.

To characterize the performance of our trial groups, we measure the area under the survival curve for each group, which captures the cumulative effect of each treatment. To determine if observed differences have statistical significance, we perform permutation tests with 10,000 rounds. In each round of a permutation test, we randomly reassign group labels and recompute the area differences under the new assignments. The intuition is that if the null hypothesis is true and there is no significant difference between two groups, then this random reassignment will only reflect stochastic fluctuation in the area difference. We assess the empirical probability distribution of this measure after completing the permutation rounds, allowing us to determine the probability (and significance) of our observed values.

All reported p-values are computed via this permutation test. We use a significance threshold of $\alpha = 0.05$, corrected during multiple testing using the simple (although conservative) Bonferroni correction, where each test in a family of m tests is compared to a significance threshold of $\frac{\alpha}{m}$.

Ideally, we would have selected this procedure as part of our original experimental design. Unfortunately, we only identified its aptness *post facto*; thus, its selection could introduce a selection bias, a possible effect that we lack any practical means to assess.

4.1 Notification Contact

For both IPv6 and ICS notifications, directly notifying WHOIS abuse contacts was most effective—particularly early on. Two days after IPv6 disclosure, direct verbose notifications resulted in 9.8% of the population remediating, compared to 3.1% when contacting national CERTs and 1.4% by contacting US-CERT. For ICS, direct notifications promoted 6.8% of the population to patch, more than national CERTs (1.7%) and US-CERT (1.0%). In both cases, direct notifications were notably better than no notifications. As can be seen in Figures 1a and 1c, this gain was persistent. After two weeks, the patch rate of directly notified IPv6 contacts was 2.4 times as high as the control, and three times as high for ICS contacts.

To determine if these observations are statistically significant, we perform permutation tests using the Bonferroni correction. With six treatment groups, the family of

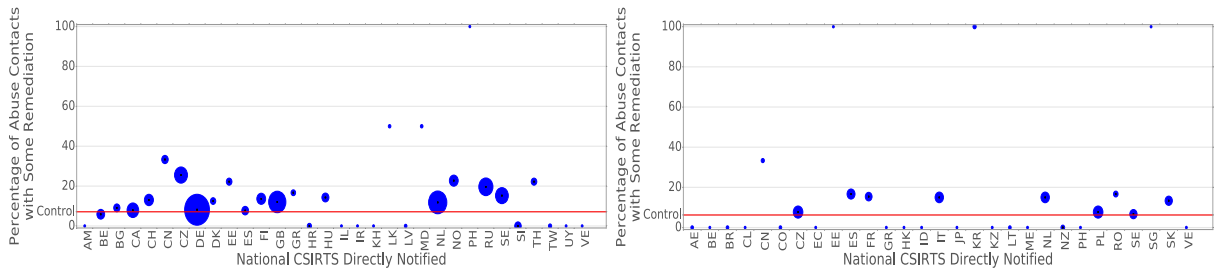

(a) Misconfigured IPv6 **(b)** ICS Services

Figure 2: Differences between National CERTs—We show the remediation rate for each directly notified national CERT after two weeks. The size of a data point is proportional to the number of abuse contacts in the country. We directly contacted 32 CERTs for IPv6, and 29 CERTs for ICS. We observe notable differences between CERT groups. However, none are statistically significantly different than the control group. This may be because there are too few hosts for some countries, and that the Bonferroni correction is conservative.

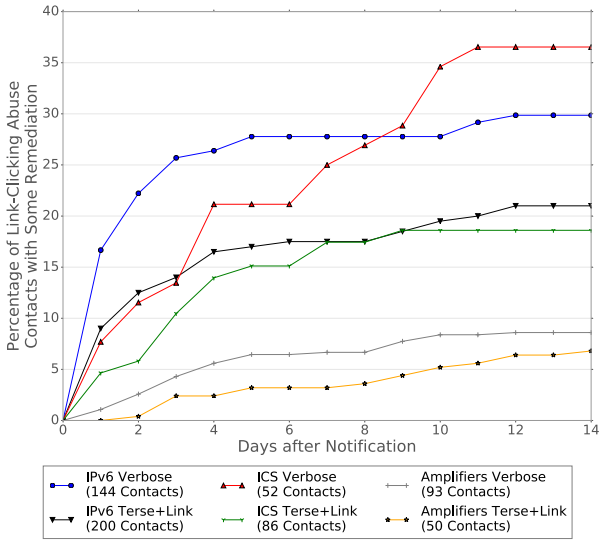

Figure 3: Remediation Rates for Website Visitors—The contacts who viewed our informational website remediated at a higher rate than those who received a verbose message. However, despite this, less than 40% of the contacts who visited the site fixed the vulnerability.

pairwise comparisons includes 15 tests, giving an individual test threshold of $\alpha = 0.0033$. Under the permutation test, the gains that direct verbose notifications had on the CERTs and the control group are statistically significant for both IPv6 and ICS, with all p-values less than 0.0001 except when comparing ICS verbose notifications with national CERTs ($p = 0.0027$).

Notably, US-CERT—our local CERT who we asked to disseminate data to other CERT groups—had the lowest patch rate, which is statistically indistinguishable from the control group that had no notifications. We suspect that US-CERT did not disseminate the data to any other CERT groups or notify any US operators. One national CERT included in the report to US-CERT informed us they had not received any notices from US-CERT. As seen in Figure 2, there were stark differences between CERT groups—some duly notified operators, while others appear to have ignored our disclosures.

Overall, this suggests that the most effective approach—in terms of both the number of hosts patched and the rate of patching—is to directly notify network operators rather than contact CERT groups.

4.2 Message Verbosity

To determine what information needs to be included in notification messages, we sent three types of emails: (1) verbose, (2) terse, and (3) terse with a link to a website with additional details. We observed the best remediation by contacts who received verbose messages. For IPv6, verbose messages were 56.5% more effective than either terse messages after two days and 55.5% more effective for ICS. However, as can be seen in Figure 1, the differences between verbose and terse messages decreased over time.

Using permutation testing and the Bonferroni correction, we find that the differences between the message types are not statistically significant for IPv6 and ICS. However, given the earlier benefits that verbose messages had for both data sets, we argue notifiers may still want to prefer verbose messages over terse ones. We discuss this effect further in Section 4.4 and note that further investigation of this variable is warranted.

We tracked the remediation rate of contacts who visited the linked website, as shown in Figure 3. We note that all of the information included in the verbose message was available on the linked website and that 16.8% of users who received an email with a link visited the site. This indicates that a sizable population of users engaged

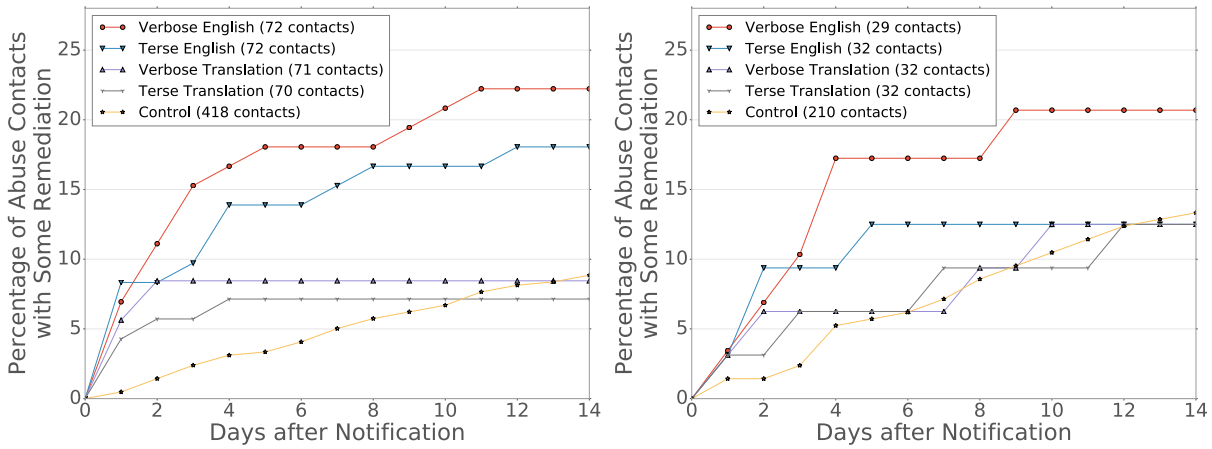

(a) Misconfigured IPv6 - German Contacts

(b) Misconfigured IPv6 - Dutch Contacts

Figure 4: Remediation Rates for Translated Messages—We find that sending verbose English messages was more effective than translating notifications into the language of the recipient. Note, though, that this observation is limited to the small set of languages we were able to evaluate.

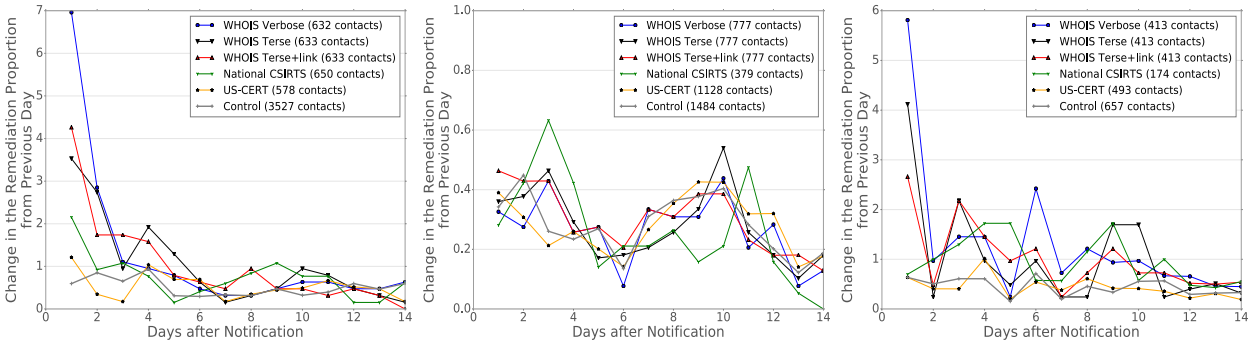

(a) Misconfigured IPv6

(b) DDoS Amplifiers

(c) ICS Services

Figure 5: Daily Changes in Remediation Proportions—We show the differences in the proportions of remediated contacts from one day to the next. We find that most contacts that remediated fixed the problem immediately after disclosure. After a few days, contacts returned to remediating at the same rate as the control group.

with our site, but many would not patch even after visiting the link. Specifically, no more than 40% of website visitors patched. Thus, even when our messages successfully reached contacts, the majority did not take action.

4.3 Message Language

To investigate whether notifications need to be translated into recipients' local languages or can be sent in English, we distributed translated messages for two countries for DDoS and IPv6 notifications. For DDoS amplifiers, we obtained native Russian and Polish translations—for the countries with the third and fourth largest number of vulnerable organizations. For IPv6, we translated messages into German and Dutch, for the second and third largest countries. The population of contacts in non-English speaking countries for the ICS dataset was too low to provide significant meaning. We randomly split the WHOIS contacts in each country into four groups that vary language and verbosity.

We observe no significant effect from language for DDoS notifications. This is unsurprising given our notifications' overall lack of effect on DDoS amplifiers. For IPv6, as seen in Figure 4, we observe that translated messages resulted in *worse* patching than when left in English. Several survey respondents were surprised at receiving translated messages from United States institutions and initially suspected our notifications were phishing messages or spam, which may explain the lower patch rate. The additional overhead of translating messages paired with less successful disclosure suggests that it may be most effective to send notifications in English.

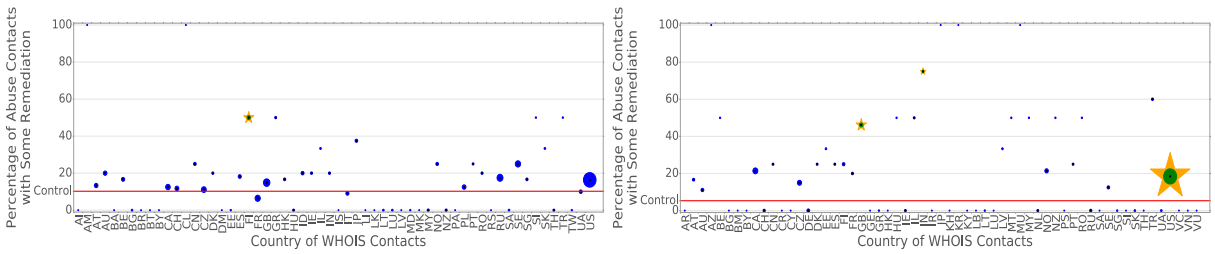

(a) Misconfigured IPv6 **(b)** ICS Services

Figure 6: Contact Remediation per Country—We show the percentage of contacts who remediated per country after two weeks. The data sizes are proportional to the number of contacts. Green data points surrounded by an orange star signify countries with a remediation rate statistically better than the control group's, under the permutation test using the Bonferroni correction.

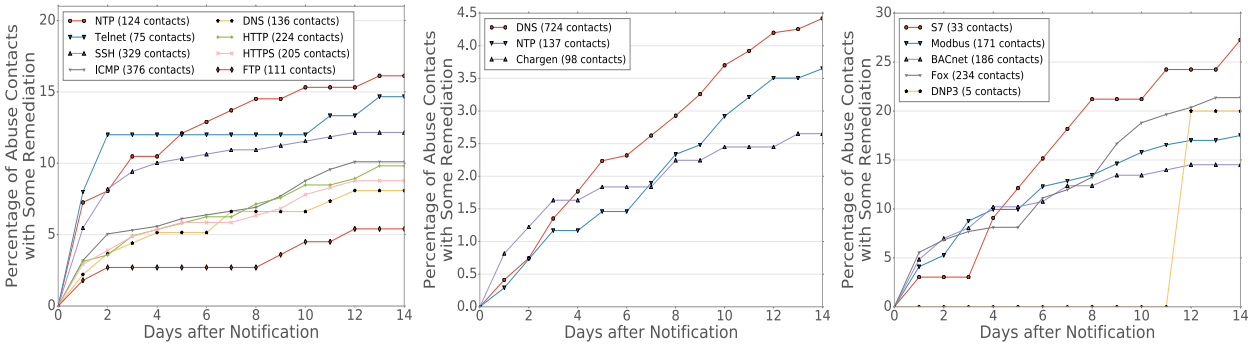

(a) Misconfigured IPv6 - WHOIS Verbose **(b)** DDoS Amplifiers - WHOIS Verbose **(c)** ICS Services - WHOIS Verbose

(d) Misconfigured IPv6 - Control **(e)** DDoS Amplifiers - Control **(f)** ICS Services - Control

Figure 7: Protocol Remediation Rates—We track the remediation rate for each specific protocol within the WHOIS verbose group and the control group. We note that operators patched some protocols significantly faster than others (e.g., Telnet versus FTP).

However, we note that our results are limited to the small set of languages we were able to obtain reliable translations for, and deeper investigation into the effects of message language is warranted.

4.4 Staying Power of Notification's Effect

As can be seen in Figure 5, our notifications caused a near immediate increase in patching. However, this increased patching velocity did not persist. In other words, we find that the effects of notifications were short-lived—on the order of several days. The day after notifications were sent, we observe large increases in the remediation proportions for IPv6 and ICS notified groups, as operators responded to our reports. However, we also see that the daily changes in remediation proportions drastically dropped by the second day.

For IPv6, the daily changes in remediation proportions

for all notified groups leveled off and matched that of the control group from the fifth day onward. We also witness a drop off in the daily remediation proportion changes for ICS, although a non-trivial amount of change continued throughout the first 10 days. Notably, the national CERTs first began accelerating remediation after two days, a delay compared to WHOIS experiment groups. For amplifiers, there was little change in the remediation rate over time, which is unsurprising given the limited effect of our notifications.

4.5 Geographic Variation

As with the national CERTs, we note variation in the patching rates between countries. This suggests that the geographic distribution of vulnerable contacts may influence a notification's outcome. As visible in Figure 6, the United States, Great Britain, India, and Finland were the only countries that patched significantly better than the control group. However, we note that some countries had too few hosts to be statistically significant, given the conservative nature of the Bonferroni correction.

4.6 Variation over Protocols

In Figure 7, we observe variation in the patch rates for different protocols within each vulnerability (e.g., Modbus versus S7 for ICS). As seen in Figure 7a, network administrators reacted most to open IPv6 NTP, Telnet, and SSH services, and least to FTP, with over a 200% difference in the remediation proportions. This variation is not reflected in the control group (Figure 7d), where all protocols exhibited similar behavior. This may reflect an increased likelihood that certain services were unintentionally left accessible, or that operators assessed different levels of risk for allowing different protocols to be reachable.

Operators also responded differently for the multiple ICS protocols (Figure 7c), but the variation is also reflected for contacts in the control group (Figure 7f). BACnet, Fox, and Modbus devices were fixed at similar rates. While the remediation of S7 systems initially lagged behind, there was a significant upswing in action after three days, with nearly 18% of contacts with vulnerable S7 systems patching after 8 days.

Surprisingly, no DNP3 systems had been patched within 10 days of notification (out of 5 contacts). We note that these five contact groups belonged to Internet service providers—not individual organizations. We similarly note that DNP3 differs from the other protocols and is specifically intended for power grid automation. These devices may be remote power stations which require more complex changes than local devices

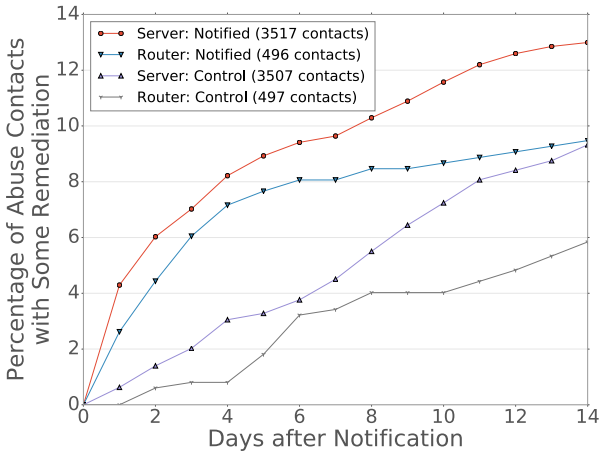

(a) Remediation rates

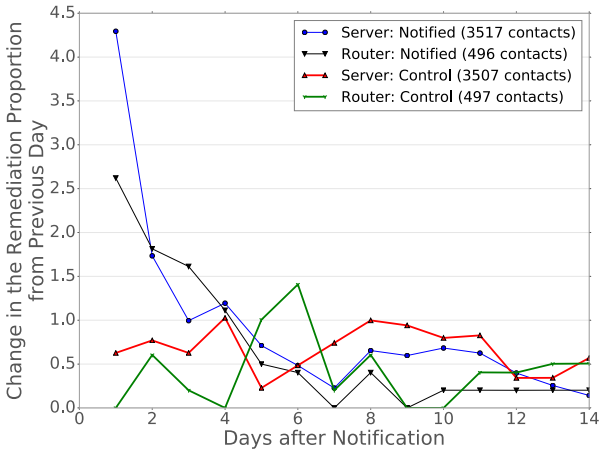

(b) Changes in the remediation proportions from one day to the next.

Figure 8: Remediation Rates by Host Type—We find no significant difference in the remediation rate between servers and routers.

(e.g., installation of new hardware versus a configuration change).

While we observe variation between amplifier protocols, these fluctuations are similar in both the notified and control group. Given the limited effect of our DDoS amplifier notifications, these differences likely reflect the varying natural churn rates of these hosts.

4.7 Host Type

When notifying IPv6 operators, we were able to distinguish between servers and routers. To assess the difference between device types, for each type, we only consider contacts with a vulnerable host of that type. We count a contact as having performed some remediation if that contact fixed at least one host of that type.

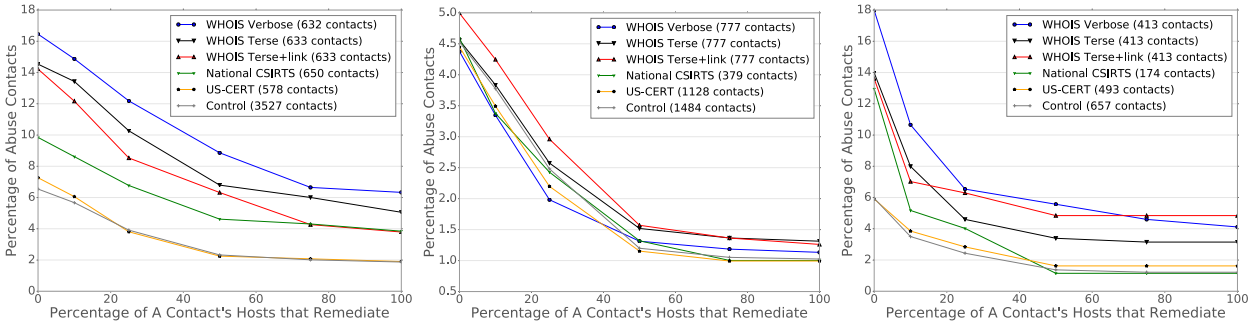

(a) Misconfigured IPv6 **(b)** DDoS Amplifiers **(c)** ICS Services

Figure 9: Remediation Completeness—We find that most operators only fixed a subset of their vulnerable hosts. For example, only 40% of the operators that fixed a single host fixed all hosts in their purview.

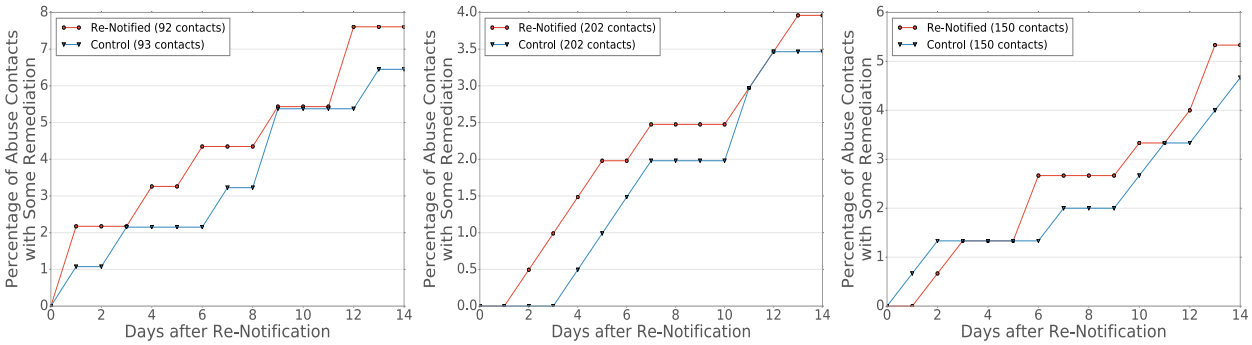

(a) Misconfigured IPv6 **(b)** DDoS Amplifiers **(c)** ICS Services

Figure 10: Re-Notifications—We find that a second round of notifications did not result in increased remediation.

We observe that servers and routers remediated at similar rates for the first four days, after which router remediation dropped off and fell significantly below that of servers (Figure 8a). However, servers also naturally patched at a higher rate than routers in the control group. This difference accounts for the gap between notified servers and routers after four days. This is also visible in Figure 8b, where the daily changes in the remediation proportions converged after four days. After 14 days, notified contacts with servers fixed at a rate 44% higher than notified contacts with routers. The divergence in the control group was similar at 48%. This indicates that overall, network administrators respond to vulnerabilities in servers and routers about equally.

4.8 Degree of Remediation

Up to this point, we designated a contact as having patched if any host under its purview was patched. We now consider how well operators patched their hosts.

As can be seen in Figure 9, the majority of contacts did not patch all of their servers. Less than 60% secured all hosts and we note that 30% of groups with 100% remediation were only responsible for fixing one or two hosts. This highlights one of the challenges in the vulnerability notification process: even if our messages reach a designated contact, that contact may not have the capabilities or permissions to remediate all hosts. The multiple hops in a communication chain can be broken at any link.

4.9 Repeated Notifications

Given that our notifications resulted in improved patching, a natural question is whether repeat notifications promote further remediation. We conducted a second round of notifications for the contacts that were directly sent verbose messages in the first round since these proved to be the most effective. We randomly split contacts who had not remediated one month after our notifications into two groups, one as a control group and one to receive a second round of notifications.

As can be seen in Figure 10, the patch rates between the re-notified group and the control group were similar for all three vulnerabilities, indicating that repeat notifications are not effective. This suggests that contacts who did not remediate during the first round of notifications

either were not the appropriate points of contact, or chose (either intentionally or due to lack of capabilities) to not remediate. It is unlikely they simply missed or forgot about our original notification.

5 Notification Reactions

We included a link to an anonymous survey in all of our notification emails as well as monitored the email address from which we sent messages. In the two weeks following our disclosures, we received 57 survey submissions and 93 human email replies. In this section, we analyze these responses.

5.1 Email Responses

Of the 685 email responses we received, 530 (77%) were automated responses (e.g., acknowledgment of receipt), 62 (9%) were bounces, and 93 (14%) were human responses (Table 3). For all three vulnerabilities, over 70% of the human responses expressed positive sentiments. We received only four negative emails, all of which concerned IPv6. Two stated that we were incorrectly using the abuse contact; the other two noted that the open IPv6 services were intentional and asked to be excluded from notifications in the future. None of the emails were threatening. We detail the breakdown for each vulnerability type in Table 4.

Beyond expressing sentiments, 23 contacts requested additional information—primarily about how we detected the vulnerabilities; two requested remediation instructions. Of those 23 contacts, 15 (65%) received terse notifications without a link to additional information, while 3 contacts (13%) received verbose messages. We note that verbose messages both reduced follow-up communication and resulted in the highest patching rate.

Unexpectedly, all five contacts who requested information about DDoS amplifiers asked for evidence of DDoS attacks via network logs. This may be a result of the extensive attention amplifiers have received in the past, such that operators only respond to active abuse issues regarding amplifiers.

Twelve IPv6 contacts rebutted our claim of vulnerability. Six stated that the inconsistency was intentional; one was a honeypot; and five explained that the IP addresses we sent them no longer pointed to the same dual-stack host, likely due to network churn. Two amplifier contacts claimed we falsely notified, stating that their hosts were honeypots. However, we do note that these IPs were seen as part of an attack and were therefore likely misconfigured honeypots.

Most human responses were in English, with eight (9%) in other languages: 3 Russian, 1 German, 1 Czech, 1 Swedish, 1 French, and 1 Slovak. These non-English replies were in response to English notifications and expressed gratitude; none requested additional information.

Response Types	ICS	IPv6	Ampl.
Automated	143	214	173
Human	22	48	23
Bounces	10	34	18
Total	175	296	214
Contacts w/ No Reply	85.9%	87.2%	92.8%

Table 3: Email Responses—We received 685 email responses to our notifications, of which 14% were human replies.

Human Responses	ICS	IPv6	Ampl.
Positive Sentiments	17	35	19
Negative Sentiments	0	4	0
Neutral Sentiments	5	9	4
Request for Information	2	16	5
Taking Actions	12	17	15
False Positive Notification	0	12	2
Total	22	48	23

Table 4: Human Email Responses—We characterize the human email responses we received in reply to our notifications.

We note that the level of feedback we received regarding DDoS notifications was commensurate with our other efforts, yet the patch response was minimal. This could indicate that operators struggle with actually resolving the issue after encountering and responding to our messages, or have become desensitized enough to DDoS issues to not take real action.

5.2 Anonymous Survey Responses

All of our notification messages contained a link to an anonymous seven question survey (Appendix A), to which we received 57 submissions. We summarize the results in Table 5.

Interestingly, 46% of respondents indicated that they were aware of the vulnerability prior to notification, and 16% indicated that they had previously attempted to resolve the problem. This contrasts with the survey results in the Heartbleed study [9], where all 17 respondents indicated they were aware of the Heartbleed vulnerability and had previously attempted to resolve the problem. The widespread media attention regarding the Heartbleed bug may account for this discrepancy, highlighting the differences in the nature of various vulnerabilities.

For DDoS amplifiers and ICS vulnerabilities, the ma-

Survey Responses	ICS	IPv6	Ampl.
Aware of Issue	2/4	20/45	4/8
Taken Prior Actions	1/4	5/43	3/8
Now Taking Action	4/4	24/43	6/8
Acceptable to Detect	3/4	35/45	7/8
Acceptable to Notify	2/4	34/45	7/8
Would want Future Notifications	2/4	30/43	7/8
Correct Contact	1/3	37/43	6/8
Total	4	45	8

Table 5: Survey Responses—We included a link to a short, anonymous survey in all of our notifications. We find that most respondents (54%) weren't aware of the vulnerabilities, but found our scanning and notifications acceptable (over 75%). Further, 62% of respondents stated they were taking corrective actions and 71% of respondents requested future notifications.

jority of respondents expressed that they were now taking corrective action (75% for DDoS amplifiers, 100% for ICS). For IPv6, only 56% of respondents indicated they would fix the problem. Given the nature of the IPv6 notification, it is likely that some of the misaligned policies were intentional.

Over 80% of respondents indicated that we reached out to the correct contact, who found scanning and notifications acceptable and requested future vulnerability notifications. However, this is a population with whom we successfully established communication. The accuracy of the other contacts from whom we did not hear back could be lower.

Our survey also allowed respondents to enter free form comments. We received 17 IPv6 comments, 4 DDoS amplifier comments, and 1 ICS comment. Of the IPv6 respondents, 5 thanked us, 7 discussed how the misalignment could be intentional or that our detection was incorrect, 3 equated our messages to spam, and 2 noted that they initially thought our translated messages were phishing messages because they expected English messages from an institution in the United States. For amplifiers, we received four comments: two thanking us and two informing us not to notify unless there is a real attack. Finally, there was only one ICS commenter, who suggested contacting vendors instead of network operators, but thanked us for our notification.

The feedback we received from these survey answers and the email responses indicates an overall positive reception of our notifications. While it may be that those who provided feedback are more opinionated, these results suggest that further discourse on notifications is needed within our community.

6 Discussion

Here we summarize the main results developed during our study, and the primary avenues for further work that these suggest.

Effective Vulnerability Notifications Our results indicate that vulnerability notifications can improve remediation behavior and the feedback we received from network operators was largely positive. We conclude that notifications are most effective when detailed messages are sent directly to WHOIS abuse contacts. These notifications were most effective in our experiments and resulted in an additional 11% of contacts addressing a vulnerability in response to our message.

On the one hand, this result provides clear guidance on how to best notify network operators. On the other hand, the majority of organizations did not patch their hosts despite our notifications. Even among those who patched at least one host, most did not fix all of their vulnerable hosts. In the case of networks hosting DDoS amplifiers, *no* form of notification generated benefits statistically significant over the control.

The failures to remediate could signal a number of problems, including:

1. failure to contact the proper parties who could best instigate remediation;
2. a need for better education about the significance of the vulnerability;
3. a need for better education about the remediation process;
4. administrative or logistical hurdles that proved too difficult for those parties to overcome;
5. or a cost-benefit analysis by those parties that concluded remediation was not worth the effort.

Illuminating the role that each of these considerations plays, and the best steps to then address them, remains for future work.

In addition, we found the effects of our notification campaigns to be short-lived: if recipients did not act within the first couple days, they were unlikely to ever do so. Repeat notifications did not further improve remediation levels.

Thus, while we have developed initial guidance for conducting effective notifications, there remain many unanswered questions as to how to best encourage operators to patch vulnerable hosts.

Improving Centralized Notification Mechanisms We observed that relying on national and regional CERT organizations for vulnerability notifications had either a modest effect (compared to our direct notifications) or no effect (indistinguishable from our unnotified controls).

While certain national CERTs evinced improved levels of remediation, others either did not act upon the information we reported, or if they did so, recipients ignored their messages. Thus, the community should consider more effective mechanisms for facilitating centralized reporting, either within the existing CERT system, or using some separate organizational structure. This need is quite salient because the burden of locating and messaging thousands of individual contacts is high enough that many researchers will find it too burdensome to conduct notifications themselves.

Open Ethical Questions The process of notifying parties regarding security issues raises a number of ethical questions. The community has already discussed some of these in depth, as in the debates concerning "full disclosure." Contacting individual sites suffering from vulnerabilities, likewise, raises questions regarding appropriate notification procedures.

For example, WHOIS abuse emails are a point-of-contact that multiple notification efforts have relied on [5, 9, 13, 14, 22, 24]. However, these contacts are technically designated for reports of abusive, malicious behavior (a point noted in the feedback we received as detailed in Section 5). While vulnerability reports have a somewhat similar flavor, they do not serve the same purpose. It behooves the security community to establish a standardized and reliable point-of-contact for communicating security issues.

Another question concerns whether the benefits of repeated notifications for the same vulnerability outweigh the costs imposed on recipients. Some may derive no benefits from the additional messages due to having no means to effectively remediate, yet must spend time ingesting the notifications. From our results, we observed that repeat notifications did not promote further patching, which argues against performing re-notifications.

More provocative, and related to the full-disclosure debate mentioned above, is the notion of *threatening* recipients with publicly revealing their vulnerabilities if unaddressed after a given amount of time. Likely, the research community would find this (quite) unpalatable in general; however, one can imagine specific situations where the community might conclude that spurring vital action justifies such a harsh step, just as some have concluded regarding full disclosure.

Future Abuse of Notifications In a future with widespread notifications, we would hope that security issues could be rectified more extensively and quickly. However, this would provide a new avenue for abuse, as attackers could potentially leverage the open communication channel to target network operators. As a simple example, a malicious actor could notify operators about a real security issue, and inform the operators to install a malicious application to help hosts resolve the security gap. While existing techniques such as phishing detection and binary analysis can help limit these attacks, the problem domain likely will yield new challenges. It is important that the security community remain cognizant of these dangers as the state of security notifications evolves.

Effective Remediation Tools For contacts that do not remediate, our measurements cannot distinguish which of the underlying reasons sketched above came into play. However, while some operators may lack sufficient motivation to take action, it seems quite plausible that others wish to, but lack the technical capabilities, resources, or permissions to do. Accordingly, we see a need for investigation into the operational problems that operators encounter when considering or attempting remediation, as well as the development of effective and usable remediation tools that simplify the operators' tasks. By reducing the effort and resources required to address a vulnerability, such tools could also increase the likelihood that an operator would take the steps to react to vulnerability reports. Ultimately, automated systems would be ideal, but these face significant challenges, such as heterogeneous platforms, potential abusive or malicious behavior, and inadvertent disruption of mission-critical systems.

7 Conclusion

We have undertaken an extensive study of notifying thousands of network operators of security issues present within their networks, with the goal of illuminating which fundamental aspects of notifications have the greatest impact on efficacy. Our study investigated vulnerabilities that span a range of protocols and considerations: exposure of industrial control systems; apparent firewall omissions for IPv6-based services; and exploitation of local systems in DDoS amplification attacks.

Through controlled multivariate experiments, we studied the impact of a number of variables: choice of party to contact (WHOIS abuse contacts versus national CERTs versus US-CERT), message verbosity, hosting a website linked to in the message, and translating the message into the notified party's local language. We monitored the vulnerable systems for several weeks to determine their rate of remediation in response to changes to these variables.

We also assessed the outcome of the emailing process itself and characterized the sentiments and perspectives expressed in both the human replies and an optional anonymous survey that accompanied our notifications. The responses were largely positive, with 96% of human email responses expressing favorable or neutral sentiments.

Our findings indicate that notifications can have a significant positive effect on patching, with the best messaging regimen being directly notifying contacts with detailed information. An additional 11% of contacts addressed the security issue when notified in this fashion, compared to the control. However, we failed to prompt the majority of contacts to respond, and even when they did, remediation was often only partial. Repeat notifications did not further improve remediation. Given these positive yet unsatisfactory outcomes, we call on the security community to more deeply investigate notifications and establish standards and best practices that promote their effectiveness.

Acknowledgments

The authors thank L. Aaron Kaplan for insightful discussions regarding the CERT organizations and Philip Stark for providing statistical consultation. We similarly thank Jethro Beekman, Christian Kreibich, Kirill Levchenko, Philipp Moritz, Antonio Puglielli, and Matthias Vallentin for message translations. Additionally, we thank the reviewers and our shepherd Nicolas Christin for helpful feedback.

This work was supported in part by the National Science Foundation under contracts 1111672, 1111699, 1237264, 1237265, 1345254, 1409505, 1409758, 1518741, 1518888, 1518921, and 1619620. The first author is supported by a National Science Foundation Graduate Research Fellowship. The second author is supported by the Google Ph.D. Fellowship in Computer Security. The opinions in this paper are those of the authors and do not necessarily reflect the opinions of any funding sponsor.

References

[1] National Vulnerability Database. https://nvd.nist.gov/.

[2] Conficker Working Group: Lessons Learned, 2011. http://www.confickerworkinggroup.org/wiki/uploads/Conficker_Working_Group_Lessons_Learned_17_June_2010_final.pdf.

[3] CAIDA. Archipelago (Ark) Measurement Infrastructure. http://www.caida.org/projects/ark/.

[4] CENTERS FOR DISEASE CONTROL AND PREVENTION. Patient Notification Toolkit. http://www.cdc.gov/injectionsafety/pntoolkit/index.html.

[5] CETIN, O., JHAVERI, M. H., GANAN, C., EETEN, M., AND MOORE, T. Understanding the Role of Sender Reputation in Abuse Reporting and Cleanup. In *Workshop on the Economics of Information Security (WEIS)* (2015).

[6] CZYZ, J., LUCKIE, M., ALLMAN, M., AND BAILEY, M. Don't Forget to Lock the Back Door! A Characterization of IPv6 Network Security Policy. In *Symposium on Network and Distributed System Security (NDSS)* (2016).

[7] DITTRICH, D., BAILEY, M., AND DIETRICH, S. Towards Community Standards for Ethical Behavior in Computer Security Research. Tech. rep., 2009.

[8] DURUMERIC, Z., ADRIAN, D., MIRIAN, A., BAILEY, M., AND HALDERMAN, J. A. A Search Engine Backed by Internet-Wide Scanning. In *ACM Conference on Computer and Communications Security (CCS)* (2015).

[9] DURUMERIC, Z., LI, F., KASTEN, J., WEAVER, N., AMANN, J., BEEKMAN, J., PAYER, M., ADRIAN, D., PAXSON, V., BAILEY, M., AND HALDERMAN, J. A. The Matter of Heartbleed. In *ACM Internet Measurement Conference (IMC)* (2014).

[10] DURUMERIC, Z., WUSTROW, E., AND HALDERMAN, J. A. ZMap: Fast Internet-Wide Scanning and its Security Applications. In *USENIX Security Symposium* (2013).

[11] GRAHAM, R. Masscan: The Entire Internet in 3 Minutes. Errata Security Blog, 2013. http://blog.erratasec.com/2013/09/masscan-entire-internet-in-3-minutes.html.

[12] HOFMEYR, S., MOORE, T., FORREST, S., EDWARDS, B., AND STELLE, G. Modeling Internet-Scale Policies for Cleaning up Malware. In *Economics of Information Security and Privacy III* (2013).

[13] KUHRER, M., HUPPERICH, T., ROSSOW, C., AND HOLZ, T. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *USENIX Security Symposium* (2014).

[14] LI, F., HO, G., KUAN, E., NIU, Y., BALLARD, L., THOMAS, K., BURSZTEIN, E., AND PAXSON, V. Remedying Web Hijacking: Notification Effectiveness and Webmaster Comprehension. In *World Wide Web Conference (WWW)* (2016).

[15] LUCKIE, M. Scamper: A Scalable and Extensible Packet Prober for Active Measurement of the Internet. In *ACM Internet Measurement Conference (IMC)* (2010).

[16] MAXMIND, LLC. Geoip2 database.

[17] MIRIAN, A., MA, Z., ADRIAN, D., TISCHER, M., CHUENCHUJIT, T., MASON, J., YARDLEY, T., BERTHIER, R., DURUMERIC, Z., HALDERMAN, J. A., AND BAILEY, M. An Internet-Wide View of Publicly Accessible SCADA Devices. Unpublished Manuscript.

[18] MOORE, T., AND CLAYTON, R. Ethical Dilemmas in Takedown Research. In *International Conference on Financial Cryptography and Data Security (FC)* (2011).

[19] NATIONAL CENTER FOR BIOTECHNOLOGY INFORMATION. Clinical Effectiveness of Partner Notification. http://www.ncbi.nlm.nih.gov/books/NBK261439/.

[20] RAPID7. DNS Records (ANY) Dataset, 2015. https://scans.io/study/sonar.fdns.

[21] ROSSOW, C. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Symposium on Network and Distributed System Security (NDSS)* (2014).

[22] STOCK, B., PELLEGRINO, G., ROSSOW, C., JOHNS, M., AND BACKES, M. Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification. In *USENIX Security Symposium* (2016).

[23] STONE-GROSS, B., CAVALLARO, L., GILBERT, B., SZYDLOWSKI, M., KEMMERER, R., KRUEGEL, C., AND VIGNA, G. Your Botnet is My Botnet: Analysis of a Botnet Takeover. In *ACM Conference on Computer and Communications Security (CCS)* (2009).

[24] VASEK, M., AND MOORE, T. Do Malware Reports Expedite Cleanup? An Experimental Study. In *USENIX Workshop on Cyber Security Experimentation and Test (CSET)* (2012).

[25] YEGNESWARAN, V., BARFORD, P., AND PAXSON, V. Using Honeynets for Internet Situational Awareness. In *Hot Topics in Networks (HotNets)* (2005).

A Anonymous and Optional Security Notifications Survey

Help us better understand the factors surrounding security notifications by providing anonymous feedback in this survey. Each question is optional, so answer the ones you feel comfortable answering. Thank you!

1. Was your organization aware of the security issue prior to our notification?
2. Did your organization take prior actions to resolve the security issue before our notification?
3. Is your organization planning on resolving the security issue?
4. Do you feel it was acceptable for us to detect the security issue?
5. Do you feel it was acceptable for us to notify your organization?
6. Would your organization want to receive similar security vulnerability/misconfiguration notifications in the future?
7. Did we notify the correct contact?

B IPv6 Notification: Terse with Link

Subject: [RAND#] Potentially Misconfigured IPv6 Port Security Policies

Body: Computer scientists at the University of Michigan, the University of Illinois Urbana-Champaign, and the University of California Berkeley have been conducting Internet-wide scans to detect IPv4/IPv6 dual-stack hosts that allow access to services via IPv6, but not IPv4. This likely indicates a firewall misconfiguration and could be a security vulnerability if the services should not be publicly accessible. We have attached a list of hosts that are potentially vulnerable on your network.

[LINK: More information is available at https://security-notifications.cs.berkeley.edu/[RAND#]/ipv6.html.]

Thank you,

Berkeley Security Notifications Team

Help us improve notifications with anonymous feedback at: https://www.surveymonkey.com/r/Q2HLJ5D

C IPv6 Notification: Verbose

Subject: [RAND#] Potentially Misconfigured IPv6 Port Security Policies

Body: During a recent study on the network security policies of IPv4/IPv6 dual-stack hosts, computer scientists at the University of Michigan, the University of Illinois Urbana-Champaign, and the University of California Berkeley have been conducting Internet-wide scans to detect IPv4/IPv6 dual-stack hosts that allow access to services via IPv6, but not IPv4. This likely indicates a firewall misconfiguration and could be a security vulnerability if the services should not be publicly accessible. We have attached a list of hosts that are potentially vulnerable on your network (as determined by WHOIS information).

For each dual-stack host, we test whether popular services (e.g., SSH, Telnet, and NTP) are accessible via IPv4 and/or IPv6 using a standard protocol handshake. For ICMP this is an echo request, for TCP it is a SYN segment, and for UDP this is an application-specific request (e.g., DNS A query for 'www.google.com' or an NTP version query). We do not exploit any vulnerabilities, attempt to login, or access any non-public information.

The protocols we scanned are popular targets for attack and/or can be used to launch DDoS attacks when left publicly available to the Internet. We suspect they are misconfigured and are notifying you because hosts rarely offer services on IPv6 that are not offered on IPv4, and we believe these services may have been left exposed accidentally. This is a common occurrence when administrators forget to configure IPv6 firewall policies along with IPv4 policies.

If these IPv6-only accessible services should not be accessible to the public Internet, they can be restricted by updating your firewall or by disabling or removing the services. If none of your systems use IPv6, you can also disable IPv6 on your system. Make sure your changes are persistent and will not be undone by a system reboot.

More information is available at https://security-notifications.cs.berkeley.edu/[RAND#]/ipv6.html.

Thank you,

Berkeley Security Notifications Team

Help us improve notifications with anonymous feedback at: https://www.surveymonkey.com/r/Q2HLJ5D

D ICS Notification: Terse with Link

Subject: [RAND#] Vulnerable SCADA Devices

Body: Computer scientists at the University of Michigan and the University of California Berkeley have been conducting Internet-wide scans to detect publicly accessible industrial control (SCADA) devices. These devices frequently have no built-in security and their public exposure may place physical equipment at risk for attack. We have attached a list of SCADA devices on your network that are publicly accessible.

[LINK: More information is available at https://security-notifications.cs.berkeley.edu/[RAND#]/ics.html.]

Thank you,

Berkeley Security Notifications Team

Help us improve notifications with anonymous feedback at: https://www.surveymonkey.com/r/ZC7BVW5

E ICS Notification: Verbose

Subject: [RAND#] Vulnerable SCADA Devices

Body: During a recent study on the public exposure of industrial control systems, computer scientists at the University of Michigan and the University of California Berkeley have been conducting Internet-wide scans to detect publicly accessible industrial control (SCADA) devices. These devices frequently have no built-in security and their public exposure may place physical equipment at risk for attack. We have attached a list of SCADA devices on your network (as determined by WHOIS information) that are publicly accessible.

We scan for potentially vulnerable SCADA systems by scanning the full IPv4 address space and attempting protocol discovery handshakes (e.g., Modbus device ID query). We do not exploit any vulnerabilities or change any device state.

SCADA protocols including Modbus, S7, Bacnet, Tridium Fox, and DNP3 allow remote control and monitoring of physical infrastructure and equipment over IP. Unfortunately, these protocols lack critical security features, such as basic authentication and encryption, or have known security vulnerabilities. If left publicly accessible on the Internet, these protocols can be the target of attackers looking to monitor or damage physical equipment, such as power control, process automation, and HVAC control systems.

SCADA services are not designed to be publicly accessible on the Internet and should be maintained on an internal, segmented network, or otherwise protected by a firewall that limits who can interact with these hosts. Make sure your changes are persistent and will not be undone by a system reboot.

More information is available at https://security-notifications.cs.berkeley.edu/[RAND#]/ics.html.

Thank you,

Berkeley Security Notifications Team

Help us improve notifications with anonymous feedback at: https://www.surveymonkey.com/r/ZC7BVW5

F DDoS Amplification Notification: Terse with Link

Subject: [RAND#] Vulnerable DDoS Amplifiers

Body: Computer scientists at George Mason University and the University of California Berkeley have been detecting open and misconfigured services that serve as amplifiers for distributed denial-of-service (DDoS) attacks. Attackers abuse these amplifiers to launch powerful DDoS attacks while hiding the true attack source. We have attached a list of hosts that are potentially vulnerable on your network.

[LINK: More information is available at https://security-notifications.cs.berkeley.edu/[RAND#]/amplifiers.html.]

Thank you,

Berkeley Security Notifications Team

Help us improve notifications with anonymous feedback at: https://www.surveymonkey.com/r/Y99J8K8

G DDoS Amplification Notification: Verbose

Subject: [RAND#] Vulnerable DDoS Amplifiers

Body: During a recent study on distributed denial-of-service (DDoS) attacks, computer scientists at George Mason University and the University of California Berkeley have been conducting Internet-wide scans for open and misconfigured services that serve as amplifiers for DDoS attacks. Attackers abuse these amplifiers to launch powerful DDoS attacks while hiding the true at-

tack source. We have attached a list of hosts that are potentially vulnerable on your network (as determined by WHOIS information).

We detect amplifiers by monitoring hosts involved in recent DDoS attacks and checking whether these hosts support the features used for launching an attack (e.g., NTP monlist or recursive DNS resolution). We do not exploit any vulnerabilities or attempt to access any non-public data on these servers.

DDoS attacks are often conducted by directing an overwhelming amount of network traffic towards a target system, making it unresponsive. Amplifiers are services that send large amounts of data in response to small requests. Attackers leverage these in DDoS attacks by spoofing traffic to the amplifier, forging it to look as if it came from the attacker's target. Amplifiers then respond to the target with a large response that overwhelms the target. Publicly accessible amplifiers are constantly abused by attackers to conduct the DDoS attacks for them while hiding the tracks of the real attacker.

These amplifiers can be avoided by disabling the application or updating your firewall to block the application port or restrict the IP addresses that can access it. More specifically, Chargen should be closed as it is rarely useful and is inherently an amplifier. If left open, DNS should be configured to restrict who can make recursive requests, and NTP should be configured to disable the monlist functionality. Make sure your changes are persistent and will not be undone by a system reboot.

More information is available at https://security-notifications.cs.berkeley.edu/[RAND#]/amplifiers.html.

Thank you,

Berkeley Security Notifications Team

Help us improve notifications with anonymous feedback at: https://www.surveymonkey.com/r/Y99J8K8