On the Mismanagement and Maliciousness of Networks

Jing Zhang[†], Zakir Durumeric[†], Michael Bailey[†], Mingyan Liu[†], and Manish Karir[‡]

[†]Computer Science and Engineering
University of Michigan

{*jingzj, zakir, mibailey, mingyan*}*@umich.edu*

[‡]Department of Homeland Security
Science and Technology Directorate
Cyber Security Division

manish.karir@hq.dhs.gov

***Abstract*—In this paper, we systematically explore the widely held, anecdotal belief that mismanaged networks are responsible for a wide range of security incidents. Utilizing Internet-scale measurements of DNS resolvers, BGP routers, and SMTP, HTTP, and DNS-name servers, we find there are thousands of networks where a large fraction of network services are misconfigured. Combining global feeds of malicious activities including spam, phishing, malware, and scanning, we find a statistically significant correlation between networks that are mismanaged and networks that are responsible for maliciousness.**

I. Introduction

Misconfigured networks have long been attractive resources for hackers [1], and anecdotal evidence suggests that mismanaged networks are often taken advantage of for launching external attacks, posing a risk not only to themselves, but to the Internet as a whole. One example of this can be seen in DNS amplification attacks in which attackers utilize open DNS resolvers to flood target hosts with a large number of DNS responses. These amplification attacks have long been observed in the wild and continue to occur with increasing scale and impact [2], [3]. These attacks are innately dependent on both widely-distributed misconfigured open DNS resolvers and the ability of attackers to forge request packets. In spite of calls by the Internet security community to address both of these issues by following standard deployment practices [4], [5], serious attacks continue to occur [6], [7]. As a result, these events are frequently described in terms of economic externalities: "a situation where a party could efficiently prevent harm to others—that is, a dollars worth of harm could be prevented by spending less than a dollar on prevention—but the harm is not prevented because the party has little or no incentive to prevent harm to strangers [8]."

Our study complements these anecdotes of individual incidents with a macroscopic, systematic study of one such externality—network mismanagement. For the purpose of this study, we define mismanagement as the failure to adopt commonly accepted guidelines or policies when administrating and operating networks. We explore the relationship of such misconfiguration with apparent network maliciousness, defined as the fraction of IP addresses in the network that are listed by 12 reputation blacklists. We seek to understand the relationship between different types of network mismanagement and global Internet security.

Rather than focusing on how individual vulnerabilities influence the likelihood of a host becoming compromised (e.g., CVE-2008-4250 resulting in Conficker infections), we instead investigate how symptoms of network mismanagement, such as the presence of open recursive resolvers or instances of observed BGP misconfiguration, relate to organizational reputation built from externally observed malicious behavior, such as malware hosting and SPAM. While these features are merely proxies, ultimately, we hope to answer the question of what relationships exist between poor network management and apparent network maliciousness, or reputation on the Internet.

We begin by measuring and analyzing the prevalence of eight varied network mismanagement symptoms ranging from open recursive DNS resolvers and open SMTP relays to misconfigured routing, naming, and web infrastructure. By leveraging Internet-scale measurements, we show that misconfigured systems and servers are pervasive, including over 27 million open recursive DNS resolvers, 22 thousand open SMTP relays, and 227 thousand DNS resolvers that do not utilize source port randomization.

Next we aggregate these misconfigured systems at the autonomous system (AS) level. We explore the distribution of these symptoms across ASes and the relationship between symptoms at the AS level. We find that a small fraction of the total ASes exhibit a significant amount of misconfiguration for any given symptom. For example, for several ASes, all of the mail servers within the AS are configured as open relays, and in hundreds of ASes, none of the DNS resolvers within the AS perform source port randomization. We further find that there are statistically significant correlations between symptoms within an AS. For example, we see moderate, positive correlations between the frequency of occurrence of open DNS resolvers within an AS and both the lack of port randomization as well as the prevalence of self-signed HTTPS certificates.

NDSS '14, 23-26 February 2014, San Diego, CA, USA

http://dx.doi.org/10.14722/ndss.2014.23057

We combine these symptoms into an overall mismanagement metric. This enables us to explore how different regions, topological locations, and business relationships between ASes relate to a network's mismanagement. For example, we find that networks in Latin and South America and Africa are more likely to be mismanaged than those in North America. Having constructed a plausible network mismanagement metric, we then explore whether the mismanagement and apparent maliciousness of networks are correlated. By leveraging 12 global blacklists based on spam, phishing, malware and scanning activity to infer network maliciousness, our results show a statistically significant, strong positive correlation (0.64 correlation coefficient $<$0.01 p-value) between mismanagement and apparent maliciousness.

Correlation does not necessarily indicate causality and ideally, a controlled experiment would allow firm causal inference. However, it is neither ethical nor feasible to perform such an experiment. Instead, we assume network management and security can both be impacted by common social and economic factors, and use graph-based, causal inference algorithms [9] to determine causality. Ultimately, we find a causal relationship between mismanagement and maliciousness while controlling for social and economic considerations.

Our study supports the intuition that network mismanagement influences network security. We hope that this understanding will prompt the security community to develop proactive approaches for network security rather than primarily relying on reactive metrics. With this new, formal understanding of the relationship between mismanagement and maliciousness, we hope to draw attention to these networks and ultimately we hope to reduce the attacks by proactively correcting these mismanaged networks.

II. Symptoms of Mismanagement

There are many symptoms that externally reflect poor network management. We analyze eight of these symptoms, which we list in Table I. While these symptoms do not necessarily comprehensively describe all manners in which a network could be mismanaged, we choose to focus on these particular symptoms because they are well-documented in published Request for Comments (RFCs) and Best Current Practices (BCPs) [10], and are part of the security community's best practices. We attempt to focus on characteristics that are symptomatic of overall network mismanagement rather than on specific vulnerabilities that could be used for mounting an attack. This is intended to reduce any bias between mismangement symtoms and maliciousness metrics we consider later in this work (e.g., CVE-2008-4250 resulting in Conficker infections).

We choose a range of symptoms ranging from BGP routing stability to the management and patching of SMTP, HTTPS, and DNS servers. This range of symptoms has several merits. First, it provides a global perspective of an organization's network management. For example, different teams potentially manage different services and by analyzing a range of different symptoms, we focus on the overall organizational network mismanagement rather than a single misconfigured service. Second, the analysis of multiple symptoms allows us to analyze the relationships between different symptoms. Although care was taken in choosing these symptoms, we make no claim that they are complete or without bias. We discuss potential drawbacks of these symptoms invidually in the follwing subsections and reflect on them at the end of this section.

In the following subsections, we discuss each of the mismanagement symptoms with respect to their security implications, associated best practices, and our data collection methodology.

A. Open Recursive Resolvers

Open DNS resolvers respond to recursive queries for any domain and pose a direct threat to the Internet due to their role in DNS amplification attacks. In an amplification attack, an attacker sends simple DNS queries to an open resolver with a spoofed source IP address. While the DNS lookup request itself is small, the response to the victim is much larger and, as a result, the responses overwhelm the victim. BCP 140 [4] provides several recommendations for how to configure open resolvers to mitigate these threats. Ultimately, recursive lookups should be disabled unless specifically required and, when enabled, limited to intended customers.

In order to analyze the misconfiguration of open resolvers, we utilize data provided by the Open Resolver Project [11], which conducts active scans of the public IPv4 address space by sending a DNS query to every public address on port 53 and capturing the responses. The project has been performing these scans weekly since April, 2013, and has identified more than 30 million open resolvers. Detailed data collection methodology and preliminary results can be found in their recent presentation at NANOG [12].

We specifically consider the scan from June 2, 2013, which found 34.2 millions open resolvers in total. We consider the hosts that support open recursive queries as misconfigured, given their potential risk to the Internet and their failure to implement even the simplest best practices. Ultimately, we find 27.1 million open recursive resolvers on the Internet.

B. DNS Source Port Randomization

DNS cache poisoning is a well-known attack in which an attacker injects bogus DNS entries into a recursive name server's local cache. Traditionally, DNS resolvers used a randomized query ID in order to prevent cache poisoning attacks. However, in 2008, Dan Kaminsky presented a new subdomain DNS cache poisoning attack that has two new advantages [13]. First, it extends the window of attack because there is no valid reply from the authoritative name server with which to compete. Second, the multiple identical queries allow attackers to brute-force the 16-bit transaction ID that was previously relied upon for preventing these types of attacks.

Current best practices (RFC 5452 [14]) recommend randomizing the source port when performing DNS lookups in order to prevent these brute force attacks. In this configuration, a DNS server will use a large range of source ports instead of a single preset port, which significantly increases the search space for an attacker. For example, if a DNS server utilizes 2,000 source ports, the search space would increase from 64,000 to more than 100 million possibilities. Most popular DNS packages have already issued patches that implement source port randomization [15], [16].

Symptoms	Best Current Practices	Functions	Attacks	Dataset
Open Recursive Resolvers	BCP 140/RFC 5358	Naming Infrastructure	DNS Amplification	Global
DNS Source Port Randomization	RFC 5452	Naming Infrastructure	DNS Cache Poisoning	Global
Consistent A and PTR records	RFC 1912	Naming Infrastructure	-	Partial
BGP Misconfiguration	RFC 1918, RFC 6598	Routing Infrastructure	-	Global
Egress Filtering	BCP 38/RFC 2827	Transit	-	Partial
Untrusted HTTPS Certificates	RFC 5246, RFC 2459	Web Application	Man-in-the-middle	Global
Open SMTP Mail Relays	RFC 2505	Mail Application	SPAM	Global
Publicly Available out-of-band Management Devices	Manufacturer's Guideline	Server	Compromising Hosts	Global

TABLE I. SUMMARY OF MISMANAGEMENT METRICS AND THE THIRD-PARTY, PUBLIC DATA SOURCES USED FOR VALIDATION

In order to determine whether networks have patched their DNS resolvers with source port randomization, we analyze the set of DNS queries made against VeriSign's [17] .com and .net TLD name servers on February 26, 2013. In total, we observed approximately 5 billion queries from 4.7 million DNS resolvers.

In this experiment, we track the source ports utilized to make DNS queries against these TLD servers and infer that resolvers that only utilize the default source port without implementing source port randomization are misconfigured. We find that 226,976 resolvers, which account for 4.8% of total resolvers seen in the data, do not utilize source port randomization.

C. Consistent A and PTR records

DNS supports several types of records, of which Address (A) and Pointer (PTR) records are two of the most common. An A record is used to map a hostname to an IP address. A PTR record resolves an IP address to a canonical name.

One merit of PTR records is that they facilitate the validation of connecting clients and are widely used for detecting and blocking malicious IP addresses. For example, SMTP servers often discard messages from IP addresses without a matching PTR or MX record. The DNS operational and configuration guidelines (RFC1912 [18]) dictate that every A record should have a corresponding PTR record[19].

In our study, we utilize two datasets in order to estimate the global status of DNS records: the .com and .net second level domains stored in the VeriSign zone files and the domains in the Alexa Top 1 Million popular websites [20].

In order to determine which A records have associated PTR records, we perform a DNS query for each domain in our two datasets, finding 116 million A records. We then perform a reverse DNS lookup of the IP addresses appearing on these 116 million A records. We find that 27.4 million A records, which account for 23.4% of A records we queried, do not have a matching PTR record.

We note that our dataset is biased toward domains within North America and Europe. However, given that .com and .net domains account for more than half of all domains on the Internet [21] and that Alexa includes the most popular sites in the world, we believe our results still provide insights into the status of DNS records management.

D. BGP Misconfiguration

Publicly routed networks utilize Border Gateway Protocol (BGP) in order to exchange advertised routes. A router can announce a new route for a prefix or withdraw a route when it is no longer available. Routers are expected to not send updates unless there are topological changes that cause its advertised routes to change. However, misconfiguration and human error can result in unnecessary updates, which can potentially lead to both security vulnerabilities (e.g., Bogons [22], [23]) and downtime (e.g., AS7007 incidents [24]).

Mahajan et al. note that 90% of short-lived announcements (less than 24 hours) are caused by misconfiguration [25]. This is because policy changes typically operate on human timescales, while changes due to misconfiguration typically last for a much shorter time.

In order to measure BGP misconfigurations, we use this simple heuristic in coordination with BGP updates from 12 BGP listeners in the Route Views project [26]. In our experiment, we track the time period for every new route announcement during the first two weeks of June, 2013 and infer that routes that last less than a day were likely caused by misconfiguration. We detect 42.4 million short-lived routes, which account for 7.8% of announced routes during the period of two weeks. We note that the Mahajan methodology is dated, and a fruitful area of future work would be to validate this methodology in the context of current routing practice (i.e., the current practice of fine-graned routing announcements).

E. Egress Filtering

Attackers often spoof source IP addresses to achieve anonymity or as part of DDoS attacks [27], [28]. In order to counter these attacks, it has been a best practice since 2000, to perform egress filtering as documented in BCP 38 [5].

In order to measure which networks have implemented egress filtering, we consider data from the Spoofer Project [29], which utilizes approximately 18,000 active clients to send probes to test for the presence of egress filtering. We specifically analyze data from April 29, 2013 and check in which netblocks an arbitrary routable source IP address can be spoofed. Because spoofed IP addresses are primarily used by attackers, we consider netblocks that do not implement address filtering to be misconfigured. The dataset from April 29th contained results for 7,861 netblocks, of which 35.6% have not implemented egress filtering. Unfortunately, the status of the remaining 195,000 netblocks is unknown.

F. Untrusted HTTPS Certificates

HTTPS sites present X.509 certificates as part of the TLS handshake in order to prove their identity to clients. When properly configured, these certificates are signed by a browser-trusted certificate authority.

Now that browser-trusted certificates are available for free from several major providers, the best practice is for public websites to use browser-trusted certificates. As such, we consider the presence of untrusted certificates as a potential symptom of misconfiguration. However, a large number of sites utilize self-signed certificates or certificates that have not been validated by a trusted authority.

In order to understand the state of HTTPS certificate utilization, we consider a scan of the HTTPS ecosystem that was completed as part of the ZMap network scanner project [30]. In this scan, Durumeric et al. performed a TCP SYN scan on port 443 of the public IPv4 address space on March 22, 2013 using the ZMap network scanner. It then performed a follow-up TLS handshake with hosts that responded on port 443, and collected and parsed the presented certificate chains using libevent and OpenSSL.

Using this dataset, we consider whether presented certificates are rooted in a browser-trusted certificate authority or are not browser trusted (i.e. self-signed or signed by an unknown certificate authority). We found 33 million hosts with port 443 open, 21.4 million hosts who successfully completed a TLS handshake, and 8.4 million distinct X.509 certificates. Among these certificates, only 3.2 million (38%) were browser-trusted, and only 10.3 million (48%) of the hosts presented browser-trusted certificates.

G. SMTP server relaying

Open mail relays are SMTP servers that do not perform any filtering on message source or destination and will relay e-mail messages to any destination. These servers are frequently abused by spammers in order to avoid detection or to offload traffic onto third parties. Given their consistent abuse, the Internet community strongly recommends against their use (RFC 2505 [31], RFC 5321[32]).

In order to investigate the prevalence of open mail relays, we performed a TCP SYN scan of the IPv4 address space for port 25 using ZMap on July 23, 2013 and attempted the initial steps of an SMTP handshake in order to determine whether the server would reject the sender or receiver. After determining whether the server would accept the message, we terminated the connection without sending any mail.

Our scan identified 10.7 million servers with port 25 open of which 7.0 million identified themselves as SMTP servers. Of the responsive SMTP servers, 6.2 million explicitly rejected our sender, 433,482 terminated the connection or timed out, and 22,284 SMTP servers accepted the message, identifying them as open mail relays.

H. Publicly Available Out-of-Band Management Cards

Out-of-band management cards that allow remote control of power, boot media, and in some cases, remote KVM capabilities, are now commonplace on servers. Most of these management cards are implementations of the Intelligent Platform Management Interface (IPMI) industry standard, but come under a variety of names, including Dell's Remote Access Card (iDRAC), HP Integrated Lights Out (iLO), and Super Micro's Base Management Card (BMC).

While these interfaces are a valuable tool for systems administrators, they also pose a severe security risk if publicly available on the Internet [33]. These devices have recently been found to be riddled with vulnerabilities, and manufacturers explicitly recommend that the devices be isolated on a private management network and not be made available on the public Internet [34], [33], [35]. As such, we consider any publicly-available management card to be a misconfiguration.

In order to measure the public availability of these IPMI cards, we consider the TLS certificate data set collected by Durumeric et al. by searching for known default certificates presented by IPMI cards manufactured by Dell, HP, and Super Micro. In this dataset, we found IPMI cards hosted on 98,274 IP addresses.

I. Summary and Limitations of Symptoms

In this work, we choose to focus on eight symptoms that we believe expose mismanaged networks and, for the most part, are not vulnerabilities that will directly influence the the blacklists we consider later in this work. We further focus on symptoms that have clear and accepted best practices, which have been documented by the security community.

We note that these symptoms are not the only externally visible metrics for network mismanagement—there most likely exist networks that contain other mismanaged services, which may correlate to the symptoms we present. Additionally, we acknowledge that biases may exist between the symptoms that we select that cannot be discerned without operational details of an organization (e.g., open recursive DNS resovers and open SMTP relays).

Regardless, we observe pervasive failures in implementing common security practices in the symptoms that we do consider, several of which can, by themselves, result in easily exploitable vulnerabilities. Specifically, we find that there exist (1) 27 million open recursive resolvers, (2) 226,976 DNS resolvers that have not been patched to use source port randomization, (3) 27.4 million A records that do not have matching PTR records, (4) 42.4 million short-lived BGP routes, (5) 35.6% of the tested netblocks that have not implemented egress filtering, (6) 10.2 million HTTPS servers using untrusted certificates, (7) 22,284 SMTP servers that allow open mail relays, and (8) 98,274 public accessible IPMI cards.

III. Mismanagement Symptoms at Autonomous System Level

In this section, we analyze the previously discussed symptoms at the AS level in order to determine the global misconfiguration of different networks and to measure the relationships between different types of misconfiguration.

A. Abstracting Networks

While it would be ideal to measure the correlation between mismanagement and maliciousness at the organizational level, there exist no easily visible or authoritative network boundaries from an external perspective—it is often difficult or impossible to detect what sociopolitical organizations own or manage network blocks or specific hosts within a network block.

Several methodologies have emerged for aggregating networks ranging from AS-level aggregation, to BGP routed prefix [36], to aggregating hosts by adminstrative domains defined by authoritative name server [37], [38], [39]. We choose to aggregate hosts at the AS level because several of our metrics are only available at this granularity and because as we move forward, we ultimately hope to send information to owning organizations.

We make no claim that this choice of administrative boundary is ideal. For example, several uniquely managed organizations make exist within a single AS (e.g., customers of a large provider). Strictly speaking, we do not show in all cases a correlation between mismanaged organizations and malicious networks, but rather between mismanaged ASes and ASes that have been the source of malicious traffic.

B. Distribution of Misconfigured Systems

We hypothesize that the security postures of networks will differ based on the varied effort placed in management and security. To validate this hypothesis, we consider the distribution of each of type of misconfiguration based on host IP addresses in each AS.

We rank networks by the normalized number of misconfigured systems, and show the breakdown of vulnerabilities in Figure 1. In line with our hypothesis, mismanagement is different between different networks—symptoms of misconfiguration are typically concentrated in a small number of networks.

In the remainder of this section, we discuss how we normalized each metric and the results of aggregating specific vulnerabilities by AS.

1) Open recursive resolvers: We normalize the number of open recursive resolvers by total number of IP addresses announced by the AS. In Figure 1a, we show the normalized number of open recursive resolvers (i.e., fraction of IP addresses that are running open recursive resolvers) for each AS, ranked by a decreasing order. We find that in the top 10 most misconfigured ASes, close to 100% of the ASes' advertised addresses are running misconfigured open resolvers. While we do not know for sure why this is occurring, we suspect that these networks are centrally managed and hosts are similarly configured. Beyond these several cases, 477 ASes (1.2%) have more than 10% of IPs running misconfigured open recursive resolvers. The long-tail distribution shows that approximately 95% of all ASes are well-managed, with a small number of no open recursive resolvers.

2) DNS source port randomization: We normalize the number of DNS resolvers without source port randomization by the total number of unique resolvers in the AS. The results are shown in Figure 1b. There are 14,102 ASes (33%) with at least one misconfigured DNS server. Among these, the top 584 most misconfigured ASes have 100% of their resolvers misconfigured, and more than 50% of the resolvers do not implement source port randomization in the top 1,762 ASes.

3) Consistent A and PTR records: We define the normalized number of unmatched PTR records as the fraction of the AS' A records that do not have a corresponding PTR record. We show the results of this normalization in Figure 1c. At least one A record is mismatched in 21,418 ASes (49%). A large number of ASes have a disproportionally higher fraction of their A records mismatched: none of the A records in the top 5,929 ASes have corresponding PTR records and more than half of the A records are mismatched in the top 10,863 ASes.

4) BGP misconfiguration: In order to normalize BGP misconfigurations, we consider the fraction of routing announcements originating from an AS that is misconfigured. Results are shown in Figure 1d. Unlike the previously discussed metrics, we do not find clearly divided groups of ASes. Instead, we find many ASes that announce a similar number of short-lived routes. Only 37 ASes have more than half of their updated routes as short-lived, and only a few ASes have less than 5% of their updates that are caused by misconfiguration. We suspect that this is because the causes of BGP misconfiguration are numerous and complex [25].

5) Egress Filtering: Ideally, the number of netblocks without egress filtering would be normalized by the total number of netblocks in an AS. However, our dataset only includes information for a fraction of the netblocks in 2,987 ASes. Therefore, we estimate the normalized number by calculating the fraction of known netblocks that are spoofable in these 2,987 ASes. As shown in Figure 1e, approximately half of these ASes do not have any netblocks that allow address spoofing, while all of the tested netblocks in the top 638 ASes do not implement egress filtering and are spoofable.

We note that this particular metric may not accurately represent the distribution of networks without egress filtering. First, we can only estimate the deployment of source address validation in 6% of all ASes. Secondly, the results may be biased given that the tested netblocks in a particular AS may not accurately represent the behavior of the entire AS. However, even with these limitations, we believe that the existence of egress filtering is a symptom worth considering when discussing mismanaged networks due to the potential abuse for attacks.

6) Untrusted HTTPS certificates: We normalize the servers that present untrusted certificates with the total number of HTTPS servers seen in each AS. The results are plotted in Figure 1f. While there is less risk associated with using self-signed certificates, we find that a large number of ASes contain servers with a self-signed certificate. Specifically, more than 36,000 ASes (82%) have at least one mismanaged HTTPS server. In 8,042 ASes, all hosts serving HTTPS on port 443 use a self-signed or an otherwise untrusted certificate.

7) Open SMTP mail relays: We normalize open mail relays with the total number of SMTP servers in each AS, and we show the per-AS normalized number of open mail relays in Figure 1g. In comparison to other mismanagement symptoms, we find that mail servers are relatively well maintained. Only 1,328 ASes hosted open mail relays and only 135 ASes contained more than 10% of mail servers that are misconfigured.

8) Publicly available IPMI devices: We find relatively few publicly available IPMI cards in comparison to the previously listed metrics; in total we find IPMI cards in 5,648 ASes. Normalized by the total number of IP addresses of the ASes, the number is tiny (Figure 1h). But, a few ASes are relatively poorly managed—2% of IP addresses have been detected with IPMI cards in the top 44 ASes.

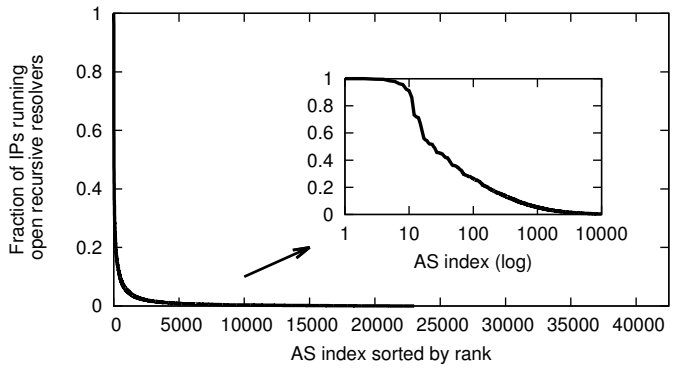

(a) Open recursive resolvers

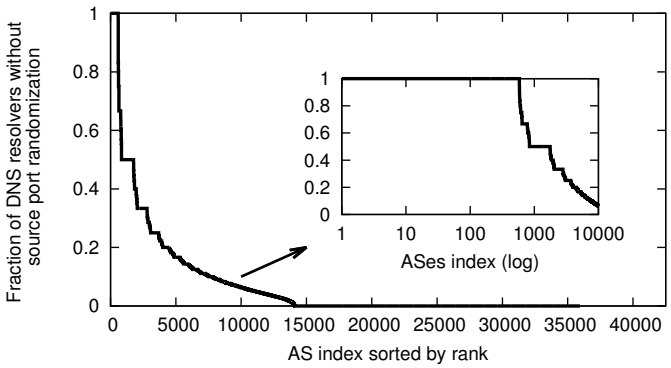

(b) DNS source port randomization

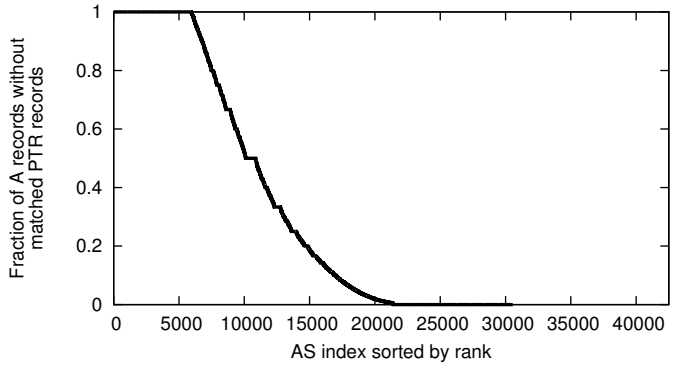

(c) PTR records

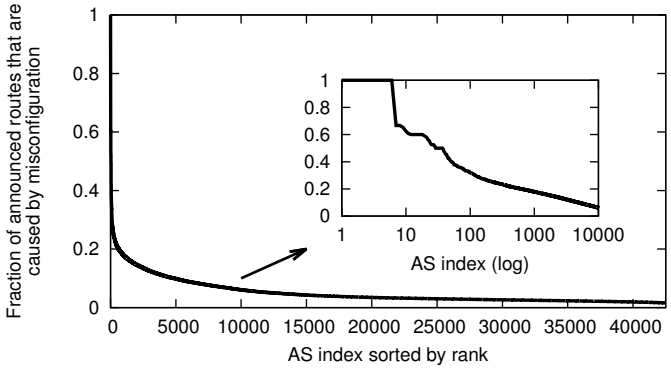

(d) BGP misconfiguration

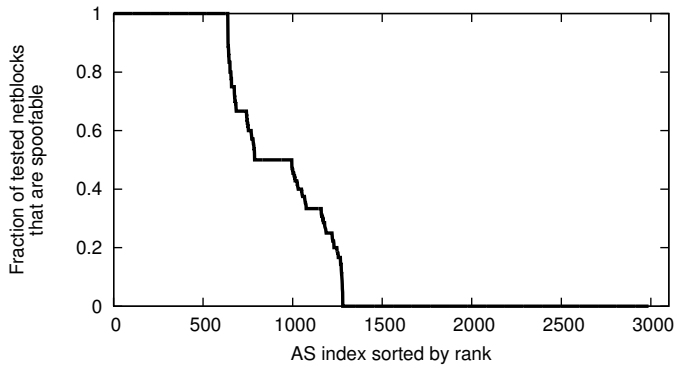

(e) Egress filtering

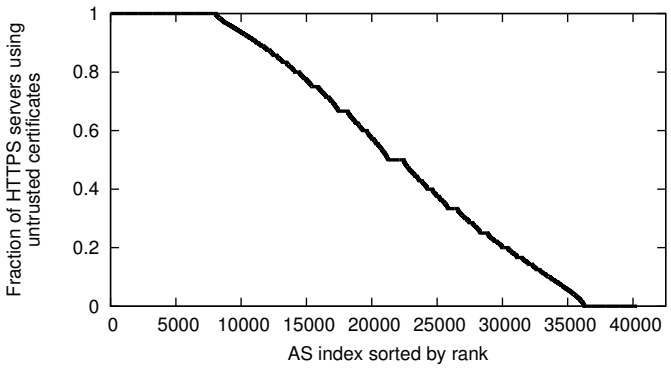

(f) HTTPS certificates

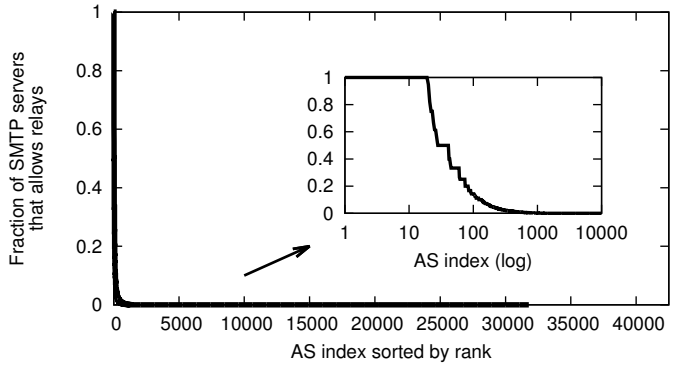

(g) Open mail relays

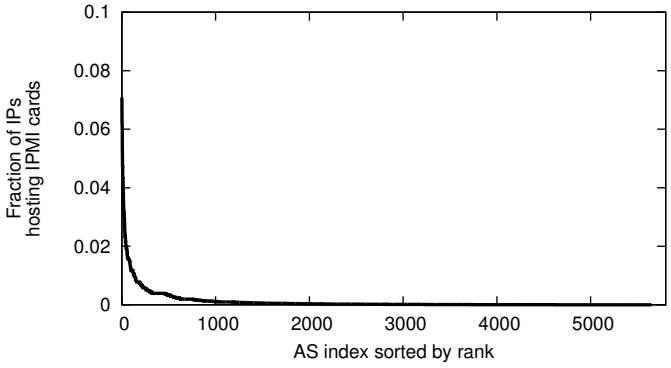

(h) IPMI cards

Fig. 1. Normalized distribution of misconfigured systems in autonomous systems. All of the symptoms show that there are a few ASes that have a disproportional number of misconfigured systems.

	Open resolver	Port randomization	PTR record	BGP misconfig.	Egress filtering	HTTPS certificate	SMTP relay	IPMI cards
Open resolver	-	**0.35** (< 0.01)	0.09 (< 0.01)	**0.17** (< 0.01)	0.09 (< 0.01)	**0.46** (< 0.01)	**0.14** (< 0.01)	**0.26** (< 0.01)
Port randomization	**0.35** (< 0.01)	-	**0.14** (< 0.01)	0.07 (< 0.01)	0.04 ($= 0.02$)	**0.23** (< 0.01)	**0.16** (< 0.01)	**0.26** (< 0.01)
PTR record	**0.10** (< 0.01)	**0.15** (< 0.01)	-	0.03 (< 0.01)	0.01 ($= 0.46$)	0.00 ($= 0.37$)	**0.11** (< 0.01)	**0.15** (< 0.01)
BGP misconfig.	**0.17** (< 0.01)	0.07 (< 0.01)	0.03 (< 0.01)	-	0.04 ($= 0.04$)	**0.16** (< 0.01)	0.02 (< 0.01)	0.03 (< 0.01)
Anti-spoofing	0.09 (< 0.01)	0.04 ($= 0.02$)	0.01 ($= 0.46$)	0.04 ($= 0.04$)	-	-0.02 ($= 0.32$)	**0.14** (< 0.01)	**0.10** (< 0.01)
HTTPS certificate	**0.46** (< 0.01)	**0.23** (< 0.01)	0.00 ($= 0.37$)	**0.16** (< 0.01)	-0.02 ($= 0.32$)	-	0.06 (< 0.01)	**0.15** (< 0.01)
SMTP relay	**0.14** (< 0.01)	**0.16** (< 0.01)	**0.10** (< 0.01)	0.02 (< 0.01)	**0.14** (< 0.01)	0.06 (< 0.01)	-	**0.26** (< 0.01)
IPMI cards	**0.26** (< 0.01)	**0.26** (< 0.01)	**0.15** (< 0.01)	0.03 (< 0.01)	**0.10** (< 0.01)	**0.15** (< 0.01)	**0.26** (< 0.01)	-

TABLE II. CORRELATION COEFFICIENTS AND P-VALUES BETWEEN DIFFERENT MISMANAGEMENT SYMPTOMS. THERE ARE SIGNIFICANT CORRELATIONS BETWEEN DIFFERENT SYMPTOMS. (RED: MODERATE CORRELATION; BLUE: WEAK CORRELATION.)

C. Correlations between Symptoms

We next explore the question of what relationship, if any, exists between the different mismanagement symptoms within an AS. To quantify the relationship between two symptoms, we use Spearman's rank correlation test, which measures the statistical dependence between two ranked variables. We use rank-based correlation rather than value-based tests because of the differences in scale between ASes and the varying implications of each mismanagement symptom. Further, rank-based correlation is a nonparametric measure that does not require data from a normal distribution.

The result of Spearman's test is a value between -1 and 1, where a negative value indicates a negative relationship between two metrics and positive value indicates a positive relationship. For any nonzero value, we perform a hypothesis test with a 95% confidence level in order to determine whether the observed correlation is significant (i.e., if $p-value < 0.05$). For a significant nonzero correlation coefficient, the larger the absolute value, the stronger the relationship. According to Cohen's guidelines [40], values with absolute correlation coefficients from 0.1 to 0.3 can be considered weakly correlated, 0.3 to 0.5 moderately correlated, and 0.5 to 1.0 to be strongly correlated.

The pair-wise correlation coefficients and p-values are shown in Table II. We find a statistically significant correlation between 25 of the 28 comparisons at a 95% confidence level. Of these, two of the pairs are moderately correlated, 14 pairs are weakly correlated, and the remaining correlations are trivial. Of the symptoms, we find the strongest correlation within vulnerability-related symptoms: open DNS resolvers, failure to implement source port randomization, and using untrusted HTTPS certificates.

Missing PTR records and BGP misconfiguration have the weakest correlation to other metrics. In the case of the PTR records, this may be caused by the biased dataset as discussed in Section II-C. For BGP misconfiguration, we expect to see little correlation with other metrics due the complexity and potential inaccuracy of measurements (see Section II-D).

We expected to find the lack of egress filtering significantly correlated with other symptoms, which we do not observe. However, we note that the relatively size sample size of this metric has skewed its results. Specifically, the measured ASes in our egress filtering dataset are biased toward fewer misconfigured systems as indicated by other metrics. As such, we do not draw any conclusions based on this metric.

One plausible explanation for the correlation between these technically disparate mismanagement metrics is that they are likely impacted by the organizational culture of security management. In other words, while we expect that disparate systems are managed by different groups within an organization, we suspect that members in an organization are influenced by its culture, including its hiring process, daily operating procedures, and general awareness of security vulnerabilities.

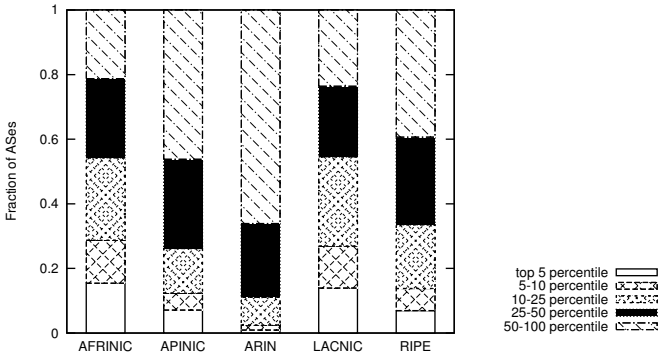

Fig. 2. Regional differences in mismanagement. Networks assigned by ARIN are relatively well managed, while a larger fraction of networks under AFRINIC and LACNIC are poorly managed.

D. Unified Network Mismanagement Metric

We next analyze the mismanagement of networks as a whole using the eight metrics we previously described. We first combine the individual symptoms into an overall mismanagement metric. Our rationale is that while each symptom may be an inaccurate measure of the AS' mismanagement, the combination of disparate metrics provides a more holistic view. Using this global metric, we consider different attributes of ASes including their geographic region and topological role.

1) Combining Symptoms: We combine different symptoms into a single metric using Borda's method [41], which is a linear combination algorithm for aggregating ranked results. This provides us with an overall score for each AS that is equivalent to an unweighted average of the AS' rank in each individual symptom. We exclude our metrics on ingress filtering and PTR records given that they only represent a small number of ASes. We rank ASes by their overall mismanagement scores from the worst to best managed.

2) Geographical Distribution: We first consider the geographic distribution of mismanagement by mapping ASes to their geographical regions using the WHOIS services provided by Team Cymru [42]. To compare mismanagement of ASes, we group ASes into five groups based on their rank percentile in the overall mismanagement metric. We show the distribution of ASes in these five groups in Figure 2.

RBL Type	RBL Name
Spam	BRBL[43], CBL[44] , SBL[45], SpamCop[46], WPBL[47], UCEPROTECT[48]
Phishing/Malware	SURBL[49], PhishTank[50], hpHosts[51]
Active attack	Darknet scanners list, Dshield[52], OpenBL[53]

TABLE III. SUMMARY OF SECURITY BLACKLISTS.

We find that networks allocated by ARIN are relatively well-managed, and that ASes in AFRINIC and LACNIC have a disproportionally large number of poorly-managed ASes. Approximately 15% of their ASes fall into the 5th percentile of mismanaged ASes, and 60% fall into the 25th percentile of mismanaged ASes.

One possible explanation for the regional differences is that less developed areas may devote less resources to network management. In addition, with different network operator groups being geographically based, the exposure to management regulations and best practices could potentially vary between geographic regions.

IV. MISMANAGEMENT AND MALICIOUSNESS

In this section, we explore whether there is a relationship between the eight mismanagement symptoms we measured and the apparent maliciousness of networks based on twelve IP reputation blacklists. We choose to consider IP blacklists that identify hosts based on sending SPAM messages, hosting phishing websites, and performing malicious port scans. In total, these blacklists contain approximately 160 million unique IP addresses. We list these blacklists in Table III.

We note that while we attempt to choose mismanagement symptoms that appear to be unrelated to the blacklists in question, there is potential for bias between some mismanagement symptoms and these blacklists. We specifically acknowledge that there is likely a bias between the SPAM blacklists we use and the existence of open SMTP relays on a network. However, we ultimately find only a weak positive correlation between the two, less than the correlation with many of the other mismanagement symptoms we investigated.

A. Maliciousness of Networks

We quantify an AS' maliciousness in three steps. First, we aggregate the blacklists in order to find the set of IP addresses that appear on any blacklist. Second, we aggregate these IP addresses by origin AS, and finally, we normalize the number of malicious IPs with the number of announced addresses in each AS. In this sense, we consider an IP address to be malicious based on its appearance on any blacklist; we do not consider an address to be any more or less malicious based on the number of blacklists on which it appears.

We find that 29,518 ASes (67%) have at least one blacklisted IP address. Figure 3 depicts the maliciousness of ASes sorted in descending order. Similar to the distribution of misconfigured systems, the maliciousness of ASes varies greatly: the top 350 ASes have more than 50% of their IP addresses blacklisted, while the bottom ASes have a negligible number of blacklisted IPs.

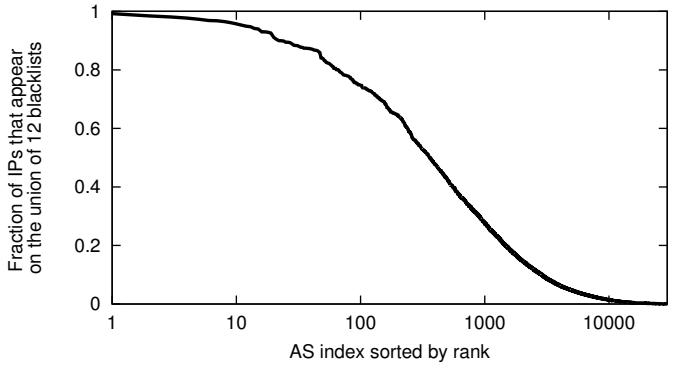

Fig. 3. Maliciousness of autonomous systems. Similar to the distribution of misconfigured systems, a few ASes have a disproportionally large number of malicious IP addresses.

Metric	Coefficient	P-value	Interpretation
Open recursive DNS resolvers	0.59	< 0.01	strong positive
DNS source port randomization	0.45	< 0.01	moderate positive
Consistent A and PTR records	0.20	< 0.01	weak positive
BGP misconfiguration	0.19	< 0.01	weak positive
Lack of Egress filtering	0.04	< 0.01	no correlation
Untrusted HTTPS certificates	0.44	< 0.01	moderate positive
Open SMTP mail relays	0.23	< 0.01	weak positive
Mismanaged IPMI cards	0.22	< 0.01	weak positive
Overall	0.64	< 0.01	strong positive

TABLE IV. CORRELATION COEFFICIENTS AND P-VALUES BETWEEN MISMANAGEMENT AND MALICIOUSNESS. THERE IS A STATISTICALLY SIGNIFICANT CORRELATION BETWEEN OUR MISMANAGEMENT SYMPTOMS AND MALICIOUSNESS.

B. Are Mismanaged Networks more Malicious?

We hypothesize that there is a positive correlation between mismanagement and maliciousness for two reasons. First, well-managed networks will expose fewer attack vectors, which will ultimately lead to fewer infected hosts and will prevent attackers from using well-managed networks as launch points for attacks. Second, well-managed networks are more likely to adopt other reactive approaches (e.g., anomaly detection, filtering/blocking) to mitigate the impact of compromise. Therefore, if compromise were to occur, hosts would not remain online long enough to be found in our scans or to be placed on a global blacklist.

In order to determine the relationship between mismanagement and maliciousness, we examine the correlation between the two metrics. We first calculate Spearman's correlation between each individual mismanagement symptom and maliciousness. All of the symptoms we examine have a statistically significant positive relationship with networks' apparent maliciousness at a 95% confidence level. We present the results in Table IV. In particular, the vulnerability-related symptoms (e.g., open DNS resolvers, DNS source port randomization, and HTTPS server certificates) have a moderate to strong correlation with maliciousness. We find that the correlation between anti-spoofing and maliciousness is negligible, which we believe is due to biased datasets.

Most interestingly, we find that our aggregated mismanagement metric has the strongest correlation with maliciousness. Given that our overall mismanagement metric is an approximation of the true management posture of a network, this observation shows that researchers need to consider a more

	Country GDP	GDP per capita
Rank of mismanagement	0.28 (< 0.01)	0.39 (< 0.01)
Rank of maliciousness	0.27 (< 0.01)	0.36 (< 0.01)

TABLE V. CORRELATION COEFFICIENT (P-VALUE) TO GDP/GDP PER CAPITA. THERE IS A STATISTICALLY SIGNIFICANT CORRELATION BETWEEN COUNTRY GDP/GDP PER CAPITA AND BOTH MISMANAGEMENT AND MALICIOUSNESS. THE HIGHER THE GDP/GDP PER CAPITA, THE BETTER THE MANAGEMENT AND SECURITY POSTURE OF THE NETWORKS.

	# of customers	# of peers
Rank of mismanagement	-0.30 (< 0.01)	-0.14 (< 0.01)
Rank of maliciousness	-0.27 (< 0.01)	-0.11 (< 0.01)

TABLE VI. CORRELATION COEFFICIENT (P-VALUE) TO NUMBER OF CUSTOMERS/PEERS. THERE IS A STATISTICALLY SIGNIFICANT CORRELATION BETWEEN NUMBER OF CUSTOMERS/PEERS AND BOTH MISMANAGEMENT AND MALICIOUSNESS. THE MORE THE CUSTOMERS/PEERS, THE WORSE THE MANAGEMENT AND SECURITY POSTURE OF THE NETWORKS.

holistic view of network health rather than only consider specific vulnerabilities or symptoms.

Correlation does not imply any cause-effect relationship; there could very well be a third variable that impacts both mismanagement and maliciousness [54]. For example, as we discussed in Section III-D, mismanagement differs between geographical and topological locations, which indicates that external social and economic factors influence mismanagement. Therefore, we need to further examine whether mismanagement causes maliciousness when controlling for social and economic factors. By utilizing a graph-based causal inference algorithm, we show that mismanagement is a cause of maliciousness when controlled by these social and economic factors.

We assume that the aforementioned differences in management within different geographic regions are caused by the differing economic development in these regions. Networks in a developed region might invest more in management and security than in less-developed countries. For each country, we use gross domestic product (GDP) and GDP per capita as economic indicators for the ASes located in the country. As shown in Table V, both maliciousness and mismanagement ranks are significantly correlated to these two economic factors at a 95% confidence level. As expected, the higher the GDP/GDP per capita, the better the management level and security posture (i.e., the lower the rank in mismanagement and maliciousness).

In addition, we look at the business relationship between ASes to infer their social and financial status. Specifically, we use two variables: number of customers and number of peers. The results in Table VI show that the number of customers or peers are negatively correlated with good management and security. This may be due to the diverse set of services offered by a network provider increasing the complexity of management and policy enforcement. Both gross domestic product (GDP) and business relationships have a similar correlation to mismanagement, and maliciousness and it is plausible that these are the common causes for both mismanagement and maliciousness.

In order to determine whether there exists a correlation between mismanagement and maliciousness when controlling for these factors, we use graphical model-based causal inference. We specifically choose to use the Fast Causal Inference (FCI) [9] algorithm, which is suitable for non-experimental data where there may be latent variables (factors influencing two or more measured variables). FCI estimates the equivalence class of maximal ancestral graphs (MAGs) that describe the conditional independence relationships between observed variables, ultimately producing a partial ancestral graph (PAG).

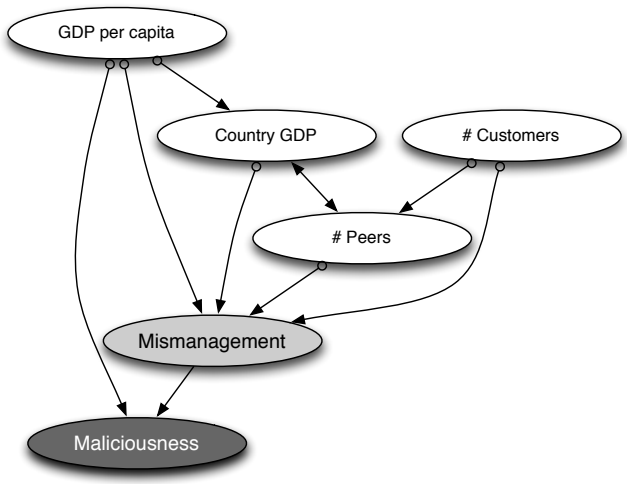

Fig. 4. Inferred causal relationships. Mismanagement is a cause for maliciousness when controlling for social and economic factors.

An edge in the PAG indicates that the two variables are conditionally dependent given S for all sets, S consisting of all variables and a subset of the observed variables. There are three types of edge-marks. A tail ($-$) or an arrowhead ($>$) indicates that this tail/arrowhead is present in all MAGs in the equivalence class. A circle (o) edge-mark is an uninformative or ambiguous mark, because it means there is at least one MAG in the equivalence class where the edge-mark is a tail, and at least one where the edge-mark is an arrowhead. Therefore, a bidirectional edge ($\longleftrightarrow$) indicates that there exists a latent variable where the correlation might be spurious. An o$\longrightarrow$ edge indicates possible latent variables, and $\longrightarrow$ shows a causal relationship.

Figure 4 depicts the inferred causal relationships within our dataset and indicates that *mismanagement is a cause for maliciousness when controlling for social and economic factors*. The edges between social and economic factors and mismanagement indicate correlations and probable causations, but there is no direct relationship between maliciousness and these factors (except BGP per capita) when controlled by mismanagement. Therefore, the possible causal chain indicates that economic factors are correlated to management level, which ultimately influences the security and apparent maliciousness of a network.

C. Impact of Aggregation Type on Maliciousness Correlations

In order to explore our choice to aggregate at the AS level instead of at a more granular level, such as by routed block or authoritative name server, we consider the correlation between three of the mismanagement symptoms and our global maliciousness metric at the routed block granularity. We find very slight differences between the level of correlation (e.g.,

Metric	Coefficient	P-value	Interpretation
Open recursive DNS resolvers	0.54	< 0.01	strong positive
DNS source port randomization	0.24	< 0.01	weak positive
Untrusted HTTPS certificates	0.39	< 0.01	moderate positive

TABLE VII. AGGREGATION AT BGP PREFIX LEVEL: CORRELATION COEFFICIENT AND P-VALUE BETWEEN MISMANAGEMENT AND MALICIOUSNESS.

the correlation between DNS port randomization moves from a moderate positive correlation to a weak positive correlation while other correlations remain unchanged). Ultimately, we find that there continues to be strong positive correlations for all of the mismanagement symptoms. We show the exact values in Table VII.

V. LIMITATIONS

There are several limitations in the data that we collect in this work. First, we utilize a large number of external data sources that were collected using disparate collection methodologies, from multiple networks with differing coverage, and during multiple time frames. While utilizing these datasets reduces the impact of active scanning on destination networks, these discrepancies could potentially impact the correlations we present. While we are not aware of any impact, there is future work to collect more consistent datasets.

As discussed throughout the paper, we aggregate networks at an AS level. Additional insight may be gained by grouping hosts at a more granular level or with the addition of organizational data. As well, there is future work required to determine whether hidden biases exist between the symptoms we select or between the symptoms and the mismanagement metrics we utilize, particularly between mail server mismanagement and SPAM metrics.

VI. DISCUSSION

Incentives to secure networks. We find that a large number of networks systemically fail to implement even the simplest best practices and ultimately pose a threat to the Internet as a whole. Unfortunately, in several cases (e.g., egress filtering), these misconfigurations pose a risk to the rest of the Internet, but pose little internal threat. As a result, organizations have little incentive to fix these services. Recent work has shown that providing social or financial incentives may be more effective than developing new technical solutions for improving overall security [55]. As such, if we are able to develop strategies in which edge networks are incentivized to better manage their systems, we may be to able increase the stability of the Internet as a whole. In one example, Gill et al. propose a strategy for increasing BGP security in which network operators assign a higher priority to routes that adopt appropriate security measures. This increased traffic translates to increased revenue, serving as a financial incentive for securing networks [56]. However, the strategy is limited to BGP security. We show that one future research direction is to develop additional incentives to encourage better organizational network management.

Shifting from Attacker-Focused to Defender-Focused Research. Numerous studies have attempted to understand attackers' motivations and technological capabilities by studying malware installed on compromised machines [57], the data malware collects [58], [59], and what data is sold on black markets [60]. While this research helps us to understand attackers' motivations, it does not necessarily improve defensive mechanisms. There are numerous questions surrounding defense mechanisms that have remained unresolved, including how to maximize a network's defense given minimal resources and how to prevent large scale DDoS attacks.

Proactive vs. Reactive reputation. While we have shown that mismanagement ultimately leads to maliciousness, how to utilize this information is an open question.

Traditional security reputation is of a reactive nature. However, reactive reputation is ineffective due to the latency between exploit and detection. In contrast, a proactive security reputation could ward off future damages by predicting future malicious sources. Is there a point at which a network becomes too dangerous to be allowed to remain connected to the public Internet? Is it appropriate to proactively blacklist open mail relays in SPAM filters or to drop DNS responses originating from known open recursive resolvers?

A proactive reputation which shows the management status of networks can also help the security community to target their efforts. For example, the security community should put more effort into understanding the causes of and develop solutions for the relatively severe symptoms of mismanagement.

VII. FUTURE WORK

A. Measuring Mismanagement

While we choose to measure eight specific symptoms of mismanagement, there exist a large number of other potentially indicative symptoms. Other open questions include how to most effectively measure mismanagement. What other metrics are available to researchers? Is it ethical to perform regular active scans of the Internet if it leads to a safer, more stable Internet? Is it possible to continuously measure the security posture of different organizations in a continuous or passive manner? Can these metrics be used to predict or prevent future attacks? There is also the potential for additional types of network aggregation. While we choose to aggregate at the AS-level, other results might be found through more granular aggregation levels.

B. Culture of Mismanagement

The correlation between different types of mismanagement indicates that there may exist a culture of mismanagement or lack of attention to security among poorly-managed networks. However, there is currently little understanding of the priorities of different teams within organizations that appear poorly managed and their impact on mismanagement and public-facing security. By better understanding these relationships, we may be able to improve this mismanagement and ultimately the security of the Internet.

C. Validating Causality.

In our study, we use four coarse estimates of organizational, social, and economic status to infer the causal relationship between mismanagement and maliciousness. There is potential for more fine-grained future research in determining exactly which social and economic factors most influence network management.

VIII. Related Work

A. Network Mismanagement

There exist a large number of best practices for specific services and for organizationally managing security, including ISO 17799 [61], the Information Security Forum [62], and Network Protection Practices [63], and there have been numerous studies on the adoption and efficacy of various best practices that we build upon. In 2009, Beverly et al. [27] performed an active measurement experiment from 12,000 clients in order to study the deployment of egress filtering. Their team showed that 31% of the clients are able to spoof any arbitrary routable source address and that 77% can forge an address within their /24 subnetwork. The results are consistent with our findings, and indicate a lack of anti-spoofing deployment improvement within the past 5 years.

In 2002, Mahajan et al. showed that 90% of short-lived routes are caused by misconfiguration and that 0.2%-1% of the global routing table consists of misconfigured routes [25]. The study also found that these misconfigured routes had a variety of causes, including human error, configuration errors, and software bugs. In our study, we use their heuristics to define updates caused by BGP misconfiguration.

In 2013, Durumeric et al. used active probing to measure misconfigurations in the deployment of HTTPS and to study the HTTPS certificate ecosystem [64]. Other studies have revealed the lack of adoption of various security technologies, for example, security-related HTTP headers [65], [30]. These projects discuss primarily the previlance of bad or the adoption of good practices. Nikiforakis et al. synthesized different misconfigurations and designed a metric to evaluate the maintenance quality of websites [66]. The metric includes availability, cookies, anti-XSS and anti-clickjacking, cache control, SSL/TLS implementation, and outdated web servers. This metric is used to inform the analysis of trust relationships between Internet sites and JavaScript providers.

Rather than focusing on a specific best practice, policy, or vulnerability, our study instead focuses on using a comprehensive view of network mismanagement that synthesizes and expands prior studies of specific incidents or ecosystems.

B. Network Maliciousness

Shue et al. [67] use a similar union approach to combine and aggregate IP-based reputation lists into reputation of autonomous systems. They examine the Internet connectivity properties of the malicious ASes and find that malicious ASes have more frequent changes with their BGP peers. However, they do not focus on types of mismanagement that might be present in these networks.

Stone-Gross et al. [68] developed FIRE, a project that aims to detect rouge networks—those that support malicious activities such as drive-by-downloads and phishing. In comparison to this work, FIRE is reactive, attempting to detect malicious networks after they begin to perform attacks. In this work, we consider symptoms of mismanagement and attempt to determine whether these are correlated to the type of malicious activities that FIRE attempts to detect. Ultimately, we hope that our work will allow projects similar to FIRE to be developed that proactively predict which networks will be used maliciously rather than waiting for them to behave maliciously.

While many works informally discuss the security implications of best practices and the presence of malicious networks, we are unaware of a systematic study on the relationship between overall network maliciousness and network management with real-world data.

IX. Conclusion

There is a widely held, anecdotal belief that mismanaged networks not only pose a risk to themselves, but to the Internet as a whole. In this paper, we systematically examine the relationship between mismanagement and maliciousness by analyzing eight Internet-scale mismanagement metrics and twelve commonly used global blacklists. Through this analysis, we find that different symptoms of mismanagement are highly correlated among themselves, and we ultimately find a causal relationship between mismanagement and maliciousness while controlling for social and economic considerations.

While security has primarily been reactionary, the understanding of the relationship between mismanagement and maliciousness is the first step in developing proactive security systems. We encourage the security community to switch some of their attention from studying attacks to researching defensive mechanisms and incentivizing organizations to implement even the simplest security best practices. Ultimately, we hope that networks can be secured proactively from such research intead of primarily reactively.

X. Acknowledgments

We wish to thank our shepherd Vyas Sekat and the anonymous reviewers for their assistance in improving this work. This work was supported in part by the Department of Homeland Security Science and Technology Directorate under contract numbers D08PC75388, FA8750-12-2-0314, and FA8750-12-2-0235; the National Science Foundation (NSF) under contract numbers CNS 1111699, CNS 091639, CNS 08311174, CNS 0751116, CNS 1330142, and CNS 1255153; and the Department of the Navy under contract N000.14-09-1-1042.

References

[1] "Hackers focus on misconfigured networks," http://forums.cnet.com/7726-6132_102-3366976.html.

[2] R. Vaughn and G. Evron, "DNS Amplification Attacks," http://www.isotf.org/news/DNS-Amplification-Attacks.pdf, 2006.

[3] "DoS Attack against DNS?" http://seclists.org/nanog/2006/Jan/294.

[4] J. Damas and F. Neves, "Preventing use of recursive nameservers in reflector attacks," RFC 5358 / BCP 140, 2008.

[5] P. Ferguson and D. Senie, "Network ingress filtering: Defeating denial of service attacks wich employ IP source address spoofing," RFC 2827 / BCP 38, 2000.

[6] "The DDoS that knocked Spamhaus offline," http://blog.cloudflare.com/the-ddos-that-knocked-spamhaus-offline-and-ho.

[7] "Putting the Spamhaus DDoS attack into perspective," http://www.arbornetworks.com/corporate/corporate/blog/4813-putting-the-spamhouse-ddos-attack-in-perspective.

[8] E. Felten, "Security Lessons from the Big DDoS Attacks," https://freedom-to-tinker.com/blog/felten/security-lessons-from-the-big-ddos-attacks/.

[9] P. Spirtes, C. Meek, and T. S. Richardson, "Causal inference in the presence of latent variables and selection bias," *CoRR*, 2013.

[10] S. Bradner, "The Internet Standards Process – Revision 3," RFC 2026 / BCP 9, 1996.

[11] "Open Resolver Project," http://openresolverproject.org/.

[12] "Open Resolver Project — Results from 3 months of active scans," http://www.nanog.org/sites/default/files/tue.lightning3.open_resolver.mauch_.pdf, 2013.

[13] "Multiple DNS implementations vulnerable to cache poisoning," http://www.kb.cert.org/vuls/id/800113.

[14] A. Hubert and R. V. Mook, "Measures for making DNS more resilient against forged answers," RFC 5452, 2009.

[15] https://kb.isc.org/article/AA-00924/0.

[16] http://technet.microsoft.com/en-us/library/dd197515(v=ws.10).aspx.

[17] "Verisign. Inc," www.verisigninc.com.

[18] D. Barr, "Common DNS operational and configuration errors," RFC 1912, 1996.

[19] "Importance of PTR records for reliable mail delivery," http://www.mxpolice.com/email-security/importance-of-ptr-records-for-reliable-mail-delivery/.

[20] "Alexa - Top Sites," http://www.alexa.com/topsites.

[21] "VeriSign's future looks stable with .com and.net registries in the bag," http://www.forbes.com/sites/greatspeculations/2012/08/20/verisigns-future-looks-stable-with-com-/and-net-registries-in-the-bag/, 2012.

[22] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address allocation for private internets," BCP 5/RFC 1918, 1996.

[23] J. Weil, V. Kuarsingh, C. Donley, C. Liljenstolpe, and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space," BCP 153/RFC 6598, 2012.

[24] http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html.

[25] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *Proceedings of SIGCOMM '02*, 2002.

[26] U. of Oregon, "Route Views Project," http://www.routeviews.org/.

[27] R. Beverly, A. Berger, Y. Hyun, and k. claffy, "Understanding the efficacy of deployed Internet source address validation filtering," in *Proceedings of IMC '09*, 2009.

[28] "DDoS strike on Spamhaus highlights need to close DNS open resolvers," http://www.techrepublic.com/blog/security/ddos-strike-on-spamhaus-highlights-need-to-close-dns-open-resolvers/9296, 2013.

[29] "Spoofer project," http://spoofer.cmand.org/index.php.

[30] Z. Durumeric, E. Wustrow, and J. A. Halderman, "ZMap: Fast Internet-wide scanning and its security applications," in *Proceedings of the 22nd USENIX Security Symposium*, 2013.

[31] G. Lindberg, "Anti-Spam recommendations for SMTP MTAs," BCP 30/RFC 2505, 1999.

[32] J. Klensin, "Simple mail transfer protocol," RFC 5321, 2008.

[33] A. J. Bonkoski, R. Bielawski, and J. A. Halderman, "Illuminating the security issues surrounding lights-out server management," *Proceedings of the 7th USENIX Workshop on Offensive Technologies*, Aug. 2013.

[34] *HP Integrated Lights-Out security*, 7th ed., Hewlett-Packard, Dec. 2010, http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c00212796/c00212796.pdf.

[35] J. Ullrich, "IPMI: Hacking servers that are turned "off"," ISC Diary blog, Jun. 2012, https://isc.sans.edu/diary/IPMI%3Aminimal+Hacking+servers+that+are+turned+%22off%22/13399.

[36] B. Krishnamurthy and J. Wang, "On network-aware clustering of web clients," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '00. New York, NY, USA: ACM, 2000, pp. 97–110. [Online]. Available: http://doi.acm.org/10.1145/347059.347412

[37] S. Venkataraman, S. Sen, O. Spatscheck, P. Haffner, and D. Song, "Exploiting network structure for proactive spam mitigation," in *Proceedings of Usenix Security 2007*, 2007.

[38] Z. Qian, Z. M. Mao, Y. Xie, and F. Yu, "On network-level clusters for spam detection," in *Proceedings of NDSS'10*, 2010.

[39] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," in *Proceedings of SIGCOMM '06*, 2006.

[40] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*. Routledge Academic, 1988.

[41] J. C. de Borda, "Mémoire sur les élections au scrutin," *Histoire de l'Académie Royale des Sciences*, 1784.

[42] http://www.team-cymru.org/Services/ip-to-asn.html.

[43] "Barracuda Reputation Blocklist," http://www.barracudacentral.org/.

[44] "CBL: Composite Blocking List," http://cbl.abuseat.org/.

[45] "The SPAMHAUS project: SBL, XBL, PBL, ZEN Lists," http://www.spamhaus.org/.

[46] "SpamCop Blocking List," http://www.spamcop.net/.

[47] "WPBL: Weighted Private Block List," http://www.wpbl.info/.

[48] "UCEPROTECTOR Network," http://www.uceprotect.net/.

[49] "SURBL: URL REPUTATION DATA," http://www.surbl.org/.

[50] "PhishTank," http://www.phishtank.com/.

[51] "hpHosts for your pretection," http://hosts-file.net/.

[52] "DShield," http://www.dshield.org/.

[53] "OpenBL," http://www.openbl.org/.

[54] H. A. Simon, "Spurious Correlation: A Causal Interpretation," *Journal of the American Statistical Association*, no. 267, pp. 467–479, 1954.

[55] L. Jiang, V. Anantharam, and J. Walrand, "How bad are selfish investments in network security?" *IEEE/ACM Trans. Netw.*, vol. 19, no. 2, 2011.

[56] P. Gill, M. Schapira, and S. Goldberg, "Let the market drive deployment: a strategy for transitioning to bgp security," in *Proceedings of SIGCOMM '11*, 2011.

[57] J. Caballero, C. Grier, C. Kreibich, and V. Paxson, "Measuring Pay-per-Install: The Commoditization of Malware Distribution," in *Proceedings of the 20th USENIX Security Symposium*, 2011.

[58] T. Holz, M. Engelberth, and F. Freiling, "Learning more about the underground economy: a case-study of keyloggers and dropzones," in *Proceedings of ESORICS'09*, 2009.

[59] B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna, "Your botnet is my botnet: analysis of a botnet takeover," in *Proceedings of CCS '09*, 2009.

[60] J. Franklin, V. Paxson, A. Perrig, and S. Savage, "An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants," in *Proceedings of CCS '07*, 2007.

[61] "International standard ISO/IEC 17799:2000 code of practice for information security management," http://17799.denialinfo.com/whatisiso17799.htm.

[62] "Information security forum," https://www.securityforum.org/.

[63] R. Communication Security and I. Council, "Internet Service Provider (ISP) Network Protection Practices," http://transition.fcc.gov/pshs/docs/csric/CSRIC_WG8_FINAL_REPORT_ISP_NETWORK_PROTECTION_20101213.pdf, 2010.

[64] Z. Durumeric, J. Kasten, M. Bailey, and J. A. Halderman, "Analysis of the HTTPS certificate ecosystem," in *Proceedings of the 13th ACM SIGCOMM Conference on Internet Measurement (IMC'13)*. ACM, Oct. 2013.

[65] "HTTP Header Survey," http://www.shodanhq.com/research/infodisc/report.

[66] N. Nikiforakis, L. Invernizzi, A. Kapravelos, S. Van Acker, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna, "You Are What You Include: Large-scale Evaluation of Remote Javascript Inclusions," in *Proceedings of the 2012 ACM Conference on Computer and Communications Security (CCS '12)*, 2012.

[67] C. A. Shue, A. J. Kalafut, and M. Gupta, "Abnormally malicious autonomous systems and their Internet connectivity," *IEEE/ACM Trans. Netw.*, pp. 220–230, 2012.

[68] B. Stone-Gross, C. Kruegel, K. Almeroth, A. Moser, and E. Kirda, "FIRE: FInding Rogue nEtworks," in *Proceedings of the 2009 Annual Computer Security Applications Conference (ACSAC '09)*, 2009.