

Internet-Scale Malware Mitigation: Combining Intelligence of the Control and Data Plane

Ying Zhang
University of Michigan
wingying@umich.edu

Evan Cooke
University of Michigan
emcooke@umich.edu

Z. Morley Mao
University of Michigan
zmao@umich.edu

ABSTRACT

Security on the Internet today is treated mostly as a *data plane* problem. IDS's, firewalls, and spam filters all operate on the simple principle of detecting malicious data plane behavior and erecting data plane filters. In this paper we explore how breaking down the barrier between the control and data plane can significantly enhance our understanding of how to detect and filter Internet threats like worms and botnets. Our investigation is guided by two specific goals: using information and anomalies detected on the data plane to inform control plane decision support and using anomalies detected on the control plane to inform data plane filtering. We begin by analyzing the source of persistent worms and other persistent malicious and misconfigured data plane traffic to understand the scope of this behavior on the control plane. We then analyze how anomalies on the control plane associated with poorly managed networks correlate with the sources of malicious and misconfigured traffic detected on the data plane. Our results show that malicious and misconfigured data plane behavior is widely spread across the control plane suggesting that constructing only a few control plane filters to block the most infected organizations is ineffective. We demonstrate that networks with data plane anomalies tend to exhibit more routing misconfigurations. Finally, we discuss how these correlations could be used to reject or filter routes and help stop recurring threats like persistent worms.

Categories and Subject Descriptors: C.2.0 COMPUTER-COMMUNICATION NETWORKS: Security and protection

General Terms: Management, Reliability, Security

Keywords: network security, BGP, routing anomaly, computer worms, Internet Motion Sensor

1. INTRODUCTION

The Internet routing infrastructure today places no distinction between the delivery of malicious traffic designed to disrupt or compromise systems and the delivery of le-

gitimate traffic. While the separation between routing and packet delivery has important benefits, it has also limited the ability of operators to filter malicious packets and contributed to many of the serious security problems facing the Internet today.

Most systems and methods designed to detect and mitigate critical security problems like Internet worms, bots, and spam ignore the *control plane* (routing data), and focus on the Internet *data plane* (packet headers and payloads). Internet security systems typically operate on the simple principle of detecting malicious data plane behavior and erecting data plane filters. For example, most firewalls, intrusion detection systems (IDS's), and spam filters operate independently of the inter and intra-domain routing infrastructure.

In this paper we explore how breaking down the barrier between the control and data plane can significantly enhance our understanding of how to detect and filter Internet threats like worms and botnets. Our investigation is guided by two specific goals: using information and anomalies detected on the data plane to inform control plane decision support and using anomalies detected on the control plane to inform data plane filtering.

- **Using data plane intelligence to inform control plane security:** Anomalies detected on the data plane can be used to identify the organizations, autonomous systems (ASes), and advertised prefixes on the control plane that are associated with malicious or misconfigured behavior. This information can be used to deploy control plane filters or inform other control plane decisions such as route selection.
- **Using Control Plane intelligence to inform data plane security:** Anomalies detected on the control plane can be used to detect and mitigate threats on the data plane. For example, a recent study [1] found that control plane attacks (*e.g.*, IP hijacking) may be exploited for malicious data plane activities such as sending spam. If we could identify abnormal control plane behavior then we might be able to deploy data plane filtering rules to more strictly detect and block persistent infections, spam, and other unwanted traffic from certain prefixes or autonomous systems (ASes).

To date, cross-layer analysis of control and data plane detection and mitigation approaches have been limited to a small set of specific applications. For example, iBGP routes are used to off-ramp traffic destined to certain addresses. Our goal in this paper is to take a broader approach and answer more fundamental questions about the utility of cross-layer information sharing.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WORM'06, November 3, 2006, Alexandria, Virginia, USA.
Copyright 2006 ACM 1-59593-551-7/06/0011 ...\$5.00.

All Prefixes	All ASes	Darknet Prefixes	Darknet ASes	Darknet Distinct /24s
193838	36982	63772	11686	418482

Table 1: Overall Statistics for All and Darknet

We begin by analyzing the source of persistent worms and other persistent malicious and misconfigured data plane traffic to understand the scope of this behavior on the control plane. Second, we separate out four specific classes of misconfigured behavior and analyze how widespread each behavior is on the control plane. Finally, we analyze anomalies on the control plane associated with poorly managed networks and attempt to correlate these anomalies with the sources of malicious and misconfigured traffic detected on the data plane.

We use Border Gateway Protocol (BGP) data as a source of the control plane information, and darknet data as a source of the data plane information. BGP is the inter-domain routing protocol on the Internet today. We use data from a /8 (16 million IP address) darknet or unused prefix. Darknet data provides an excellent source of data plane anomalies because every packet detected in the darknet is anomalous [2, 3] due to misconfiguration or malicious activities.

Our results show that malicious and misconfigured data plane behavior is widespread across the control plane. We observe 31% of ASes sending at least one packet to the darknet over a week long period though 80% of the traffic is sent from 20% of ASes. Interestingly, we find that most of the ASes sending traffic to the darknet announce smaller prefixes. It indicates that the source of persistent malicious and misconfigured behavior originates from many advertised prefixes and organizations. This suggests that constructing a few control plane filters to block the most infected organizations will not have significant impact. Finally, we also find that even when isolating specific behaviors the sources are still widespread and so targeted filtering may not be effective. We observe stability in the darknet data from two time periods: March 2006 and July 2006. For example, 30% top 30 /24s sorted by packet counts are the same while 90% top 30 ASes sorted by packet counts are the same.

Finally, we discuss both the implications of the collected data and further explore how correlating Internet control plane anomalies with data plane anomalies can be used to detect and mitigate Internet threats. From the evaluation of filtering, we find that a significant fraction of darknet traffic can be filtered with a few false positives using control plane intelligence. Filtering top 20 ASes sorted by probing duration help reduce 51% of darknet traffic while only affecting 4.5% legitimate traffic. We propose a hybrid approach by using anomalies detected on the control plane to classify packets as suspicious based on their source address to decrease false positives and increase performance of data plane filtering systems.

2. CONTROL PLANE CHARACTERIZATION OF DARKNET DATA

One of the most important problems facing the Internet is the *persistently infected computers*. A recent report by Microsoft characterizing the malicious software removed by their anti-malware tool found that the 10th most prevalent malware removed from Windows computers between January 2005 and March 2006 was the Blaster worm [4]. The Blaster was originally released in 2003 and yet in 2006 it




Figure 1: Distribution of packets received/source /24 in each prefix and AS




Figure 2: Origin AS grouping in darknet

still the 10th most removed pieces of malware [5]. In fact, half of the top 10 malware families removed by the tool were originally detected during 2003 or earlier.

Recurring malcode is an Internet-wide problem and in this section we explore how analyzing the source of persistent worms and other persistent malicious and misconfigured traffic detected on the data plane can help us understand the scope of infected behavior on the control plane. We first investigate the number and diversity of persistently infected organizations and compare their behavior to the overall control plane behavior of all organizations on the Internet. We then analyze the stability of this behavior.

The data plane darknet data were collected over a contiguous five day period from March 1 to March 5 2006 and also a contiguous seven day period from July 9 to July 15 2006. The BGP data was gathered from the Oregon RouteViews project [6] and from RIPE NCC [7]. We also use additional BGP data of longer time periods as history information. The darknet data was collected at a /8 (16 million IP address) darknet that is part of the Internet Motion Sensor project [8]. To reduce the impact of spoofing we removed all TCP or UDP backscatter including TCP SYN-ACK and ICMP port-unreachable packets [9]. We discuss the impact of spoofing later.

2.1 Source analysis

The number of ASes and advertised prefixes observed in BGP data and observed as the source of packets detected in the darknet are shown in Table 1. We observe at least one packet from 32% of all the prefixes advertised through BGP. These prefixes are located in 11686 different ASes representing 31% of the all ASes observed in BGP data. Interestingly, we never observe a packet from 69% of ASes suggesting that a large number of ASes are not persistently infected or are not leaking infections onto the wider Internet.

ASN	Packets Received	Distinct /24	AS Name	Tier	Primary Country
4131	24624262(6.3%)	144(0.04%)	CHINANET-BACKBONE No.31,Jin-rong Street	2	China
4837	14441537(3.67%)	87(0.03%)	CHINA169-BACKBONE CNCGROUP China169 Backbone	3	China
22909	3527351(0.89%)	16(0.004%)	DNEO-OSP1 - Comcast Cable Communications, Inc.	4	United States
22773	3442703(0.87%)	31(0.007%)	CCINET-2 - Cox Communications Inc.	2	United States
11486	2752848(0.69%)	2(0.0005%)	WAN - Worldcom Advance Networks	3	United States
3320	2747187(0.69%)	82(0.019%)	DTAG Deutsche Telekom AG	1	Germany
36193	2151384(0.54%)	1(0.0002%)	FASTCOLOCATION - FAST COLOCATION SERVICES	5	United States
20115	2062938(0.52%)	29(0.007%)	CHARTER-NET-HKY-NC - Charter Communications	3	United States
17506	1869566(0.47%)	3(0.0007%)	JPNIC-JP-ASN-BLOCK Japan Network Information Center	3	Japan
6478	1679491(0.43%)	28(0.007%)	ATT-INTERNET3 - AT&T WorldNet Services	3	United States

Table 2: Amount of packets received from darknet in the top 10 ASes

ASN	Prefixes Announced	Prefixes in Darknet	Avg. Mask Length	Avg. Update Count	Avg. Update Interval Duration (sec)
4131	995	633	19	5417	2547
4837	257	170	19	489	828
22909	441	306	20	973	876
22773	677	432	21	4677	2564
11486	277	44	23	295	932
3320	279	59	22	1405	1616
36193	1	1	22	69	785
20115	661	507	20	7834	933
17506	28	16	17	63	4915
6478	501	341	20	1047	869

Table 3: Routing statistics in the top 10 ASes sorted by received packets

To analyze the contribution of each prefix and AS we plotted the fraction of packets observed at the darknet from each prefix and from each AS. Figure 1(a) shows that more than 80% of the packets are from 20% of all the distinct origin ASes. Similarly, Figure 1(b) shows that 80% of packets are from 20% of prefixes. We also analyzed the number of distinct source /24 networks (*i.e.*, unique based on the first three octets of the IP address) that were the source of at least one packet detected in the darknet. The results show that of 70% distinct source /24s are within 20% prefixes and that 80% of distinct source /24s are located in 18% ASes. Thus, a small fraction of prefixes and ASes contribute to most of the packets we observe in the darknet.

The next step was to take a closer look at the top sources of malicious and misconfigured traffic detected in the darknet. We found that nearly 9% of probing packets originated from IP addresses in just two ASes, both of which were from China. The top 10 ASes as shown in Table 2 are responsible for 10% of the received packets. Among the top 10 ASes, 6 ASes are primarily based in the United States and half of them are tier-3 ASes [10]¹.

We then analyzed business types of the ASes detected in the darknet and in BGP data. We used keywords in the AS names of each AS to get a rough classification of educational networks, broadband and cable modem networks, corporation networks, other Internet communication provider networks and country classifications. The results are shown in Figure 2 and indicate that the sources of the traffic observed in the darknet cover a broad range of organizations.

Besides AS business types, the average size of the advertised prefixes can be another indication of AS size. We measured the average size of advertising prefixes by the average mask length of all prefixes announced by a given AS. Figure 3 shows that most of the prefixes associated with the darknet are small prefixes with mask length less than 19. Among all the darknet prefixes, 74.5% have mask length longer than 19.

¹Lower tier or edge networks have larger tier or rank number. The top tier providers are tier-1 or tier-2 ISPs.

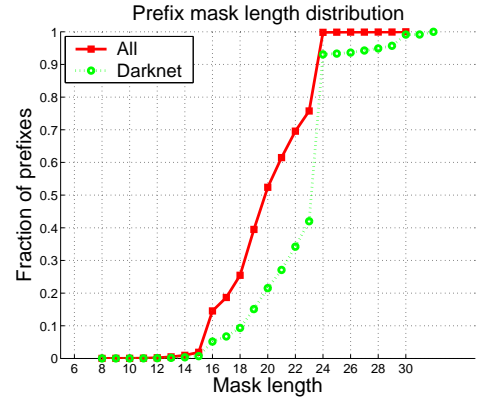


Figure 3: Distribution of mask length of prefixes in All and darknet

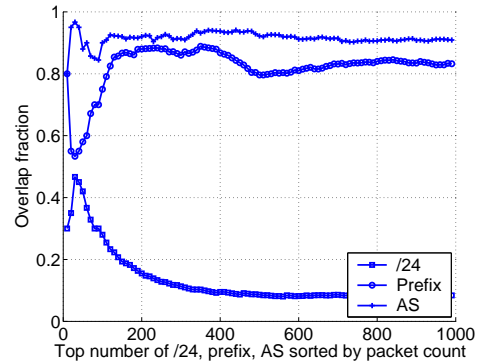


Figure 4: The overlap fraction between two weeks in March and July

2.2 Stability analysis

So far we have shown that the source of persistent malicious and misconfigured activity on the data plane is spread over a broad range of network prefixes and organizations.

Interestingly, we also found that the prefixes advertised by these organizations cover less address space than other advertised prefixes on the Internet. These results are only for one week of data in March 2006. We now analyze the stability of our observations by examining another data set from seven days in July 2006.

We now compare the number of overlapping distinct /24s, prefixes, and ASes between these two data sets. We find almost 30% of the top 30 IPs are the same sorted by packet counts, as shown in Figure 4. Comparing the observed ASes, we observe more overlap: almost 98% top 30 ASes sorted by packet counts are the same. The analysis shows persistent properties in darknet data across time. Moreover, we analyze the network characteristics of the second data set and confirm the conclusions drawn previously.

2.3 Signature analysis

The analysis thus far has characterized the control plane properties of *all* traffic observed in the darknet however, darknet traffic contains a broad range of behaviors [3] and certain malicious or misconfigured behaviors can significantly bias the results. To explore how the results might be impacted we now characterize the sources of four specific behaviors detected in the darknet. We analyze two types of malicious behavior and two types of misconfigured behavior:

- **DNS Misconfiguration:** isolate all packets destined to UDP port 53. A UDP packet designated to UDP port 53 in unused IP address space typically indicates an attempt to perform DNS lookup to a server that doesn't exist.
- **P2P Misconfiguration:** isolate all UDP packets containing the string "LIME". This signatures typically indicates a misconfigured request from a LimeWire peer-to-peer client for a file with a specific hash.
- **PopUp Spam:** isolate all UDP packets containing the string "ALERT". This signatures typically indicates an attempt to cause a Windows popup message to appear which is a type of spam as described in [3]. Note that this is a single packet attack thus the source could be spoofed.
- **MyDoom Backdoor:** isolate all packets destined to TCP port 3127. The MyDoom worm leaves a backdoor which enables a remote attacker to execute arbitrary code. Packets destined to this port typically indicate an attempt by other worms or bots to infect the target with more malware.

The overall number of prefixes and ASes of each type of darknet traffic is shown in Table 4. We find that the source addresses in P2P type to be distributed more widely compared to other three types. The malicious probing MyDoom backdoor concentrates within a smaller number of ASes but located in a large number of distinct /24s within these prefixes. As compared to the aggregate darknet traffic, the number of ASes and prefixes is significantly smaller but is still quite large in absolute terms. For example, the sources of the MyDoom behavior covered only 2.7% of ASes however this represents 991 total ASes, still very large.

Figure 5 shows the relative number of packets from each prefix that is the source of one of the four behaviors. The distribution for each of the four behaviors is similar to the aggregate darknet traffic, *i.e.*, there are some prefixes that

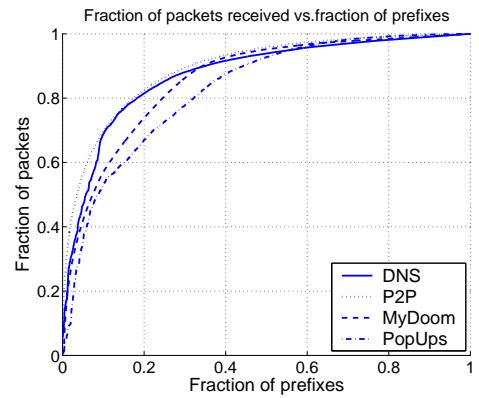


Figure 5: Distribution of packets received/source IPs in each prefix

contribute significantly more traffic. There some differences between the different types, only 60% of the MyDoom packets are from the top 20% prefixes while in P2P more than 80% of the packet are from the same fraction of prefixes.

We also find similar distributions between the specific behavior and the aggregate traffic when looking at routing update counts and inter-arrival time. There is no significant difference in the distribution of both the number of updates and update inter-arrival time among these four types.

We do however find some differences when we look at the types of source networks. Table 5 shows the top three ASes for each type of traffic. We find that the MyDoom backdoor is mostly from one country and small networks with lower tiers while the P2P misconfigurations come mostly from higher tier big networks.

In summary, we found that the two types of misconfiguration and two types of malicious behavior spanned a large number of different ASes and prefixes and had similar control plane behavior compared to the aggregate darknet traffic. While certain behaviors (such as the outbreak of a new worm) may have significantly different behavior, these results suggest persistent threats are widespread and any control plane mitigation must be broadly scoped. We discuss the implications further in Section 4.

3. JOINT ANALYSIS OF ROUTING ANOMALIES AND DARKNET TRAFFIC

In this section, we analyze anomalies on the control plane and identify correlation between networks experiencing control plane anomalies and networks sending traffic to the darknet. Control plane or *routing anomalies* can be broadly defined as unexpected behavior on the control plane, including network failures, routing misconfigurations, and security attacks targeted at routing protocols. In this work we focus on anomalies in the BGP protocol, as they can be more easily monitored using publicly available BGP feeds and BGP is the critical routing protocol directly affecting the global Internet. We suspect that networks that send traffic to darknet are likely to contain compromised hosts which scan other networks for vulnerable hosts to spread malware. Thus, such networks are "mismanaged", as they do not readily correct compromised hosts and are conjectured to be more likely to experience routing anomalies. One reason is that networks with host mismanagement are likely to also have improper configuration of network elements such as routers, potentially causing routing anomalies. Another reason is that

Type	DNS	P2P	MyDoom	PopUps
Prefix	2200(3.4%)	21569(33.8%)	2121(3.3%)	2765(4.3%)
AS	557(1.5%)	3497(9.4%)	991(2.7%)	652(1.8%)
Packets	2038308(0.5%)	1488402(0.38%)	2198838(0.56%)	13029763(3.3%)
Distinct /24	23229(5.6%)	49324(11.7%)	27446(6.6%)	62777(15%)

Table 4: Overall statistics for four types of darknet traffic

Type	Top Three ASNs	AS Names	Tier Level	Primary Country
DNS	2856	BT-UK-AS BTnet UK Regional network	3	United Kingdom
	29075	IELO-AS Ielo Network Operator.	5	United States
	4837	CHINA169-BACKBONE CNCGROUP China169 Backbone	3	China
P2P	7132	SBIS-AS - SBC Internet Services	2	United States
	19262	VZGNI-TRANSIT - Verizon Internet Services Inc.	2	United States
	5089	NTL NTL Group Limited	2	United Kingdom
MyDoom	17633	CHINATELECOM-SD-AS-AP ASN for Shandong Provincial Net of CT	4	China
	17964	16.125 DXTNET Beijing Dian-Xin-Tong Network Technologies Co., Ltd.	4	China
	9394	CRNET CHINA RAILWAY Internet(CRNET)	5	China
PopUps	3462	HINET Data Communication Business Group	3	Taiwan
	1668	AOL-ATDN - AOL Transit Data Network	2	United States
	19262	VZGNI-TRANSIT - Verizon Internet Services Inc.	2	United States

Table 5: Organization statistics of four types of darknet traffic

control plane attacks are often used to launch data plane attacks. For example, IP address hijacking is exploited by attackers to steal IP address blocks which are subsequently used to perform malicious activities such as sending spam or launching DDoS attacks to thwart trace back.

Our main finding is that networks associated with darknet traffic tend to exhibit more routing misconfiguration behavior in the following dimensions. There is slightly more routing instability associated with address prefixes observed in darknet data than other prefixes, as evidenced in the slightly larger number of routing updates for darknet prefixes. Examining evidence of routing misconfigurations in the form of private AS numbers in AS paths and bogon prefixes in advertised routes again indicate that networks sending traffic to darknet tend to have more mismanaged routers. Finally, we find preliminary evidence that routing updates associated with darknet educational networks are more likely to have MOAS conflicts indicating potential IP hijacking attacks.

3.1 Routing anomaly classification and detection

We broadly define two classes of routing anomalies. The first is associated with anomalies due to network mismanagement resulting in misconfigurations or routing instabilities. Routing misconfigurations, studied previously in measurement studies such as [11], include events such as leaking network’s internal address blocks, *i.e.*, deaggregation and announcing routes violating AS relationships causing traffic blackholes. Routing instabilities are usually a consequence of physical network failures or improper network configurations. Both misconfiguration and persistent instability are likely caused by unintentional improper management of networks. The second class of routing anomalies directly results from malicious intent and include attacks such as IP address hijacking or propagating incorrect routing information to blackhole traffic. We summarize these two classes and their basic detection techniques below.

- **Network mismanagement:**

- **Routing instabilities:** networks experience significant instability on the control plane can be characterized by the number of routing updates and their interarrival times.

- **AS path loops:** updates containing AS path loops directly violate the loop-free routing protocol semantics and can be easily detected.
- **Deaggregation:** improper configuration of route filters may cause internal prefixes often of longer mask length to be leaked to external networks. This may increase routing table sizes, causing router memory exhaustion. It can be detected by updates of small prefixes covered by currently available routes.
- **Incorrect routing information:** routing announcements containing semantically incorrect information such as private or unallocated AS numbers or bogon prefixes (prefixes that are known to belong to the private or unallocated IP space) are often a result of route filter misconfigurations. Such private information is used internally and should not be globally advertised.

- **Routing protocol attacks:**

We focus on the IP address hijacking attack only, as it has been known to be used by attackers to conduct malicious data plane activities such as spam [1]. This attack occurs when attackers intentionally announce routes to prefixes that they do not own. Networks receiving such routing updates subsequently may be misled into using them as their best routes and send traffic destined using these stolen prefixes towards the networks controlled by attackers.

Detecting such attacks often relies on identifying conflicting origin ASes in routing updates or MOAS (Multiple Origin AS) conflicts [12] as well as short-lived nature of such updates [13]. However, legitimate reasons for MOAS conflicts also exist due to multi-homing, thus significant false positives may result using such simple heuristics. Correlation with data plane anomalies can improve detection accuracies.

Next, we correlate different types of routing anomalies with darknet prefixes. We focus on analyzing several types of routing anomalies: routing instabilities, updates with AS paths containing private ASes, bogon prefixes, address deaggregation, and finally AS path routing loops. Table 6

Anomaly Type	ASes	Prefixes	% all Anomalous ASes	% Darknet ASes	% all Anomalous Prefixes	Packets	Distinct /24s	Routing Updates
Bogon IPs	139	474	83.3%	1.2%	15%	98324751(17.8%)	76682(18.3%)	1769 (12%)
Private AS	37	139	46.8%	0.31%	41.9%	147859 (0.02%)	216(0.05%)	420(2.8%)
Deaggregation	104	194	10.4%	0.28%	7.4%	37228709(6.1%)	27843(6.4%)	2240(15.2%)
AS Loop	21	1258	80.7%	0.18%	33.2%	15121202(2.1%)	14225(3.3%)	39(0.26%)

Table 6: Anomalies in the darknet




Figure 6: Routing update distribution (for one day of Feb 28th, 2006)

shows the statistics for observed routing anomalies in darknet prefixes and ASes. We choose these anomalies because they are more likely due to mismanagement of networks originating such anomalies. We also study suspicious IP hijacking attacks to identify possible correlation between control plane and data plane attacks. The routing anomalies in network mismanagement class defined before can precede or follow data plane anomalies. However, the routing protocol attacks, *e.g.*, IP address hijacking precede data plane anomalies. In this case, we can use hijacking anomaly observed in control plane to help filter data plane malicious traffic, while data plane information can be used to increase the confidence of routing plane anomaly detection.

3.2 Routing instabilities

A previous study by Rexford *et al.* [14] shows that popular prefixes, or destinations that receive a large amount of traffic tend to have more stable routing, *i.e.*, there are fewer routing updates associated with these networks compared to less popular networks. Motivated by this work, here we examine the routing stability of networks containing darknet traffic sources as shown in Figure 6 compared with all prefixes. The cumulative distribution of the total updates for one day time period of Feb 28th, 2006 (within the 5 day time period over which the darknet data was collected) indicates that darknet prefixes have slight more updates. However, most prefixes are observed to be quite stable: with more than 80% of darknet prefixes have fewer than 60 updates during the entire day across many BGP feeds from RouteViews.

3.3 Private AS routing announcements

Private AS numbers according to RFC 1930 in the range between 64512 through 65535 are often used internally and should not be advertised externally. Similarly, unallocated AS numbers should not be externally used. However, mismanaged networks may use private AS numbers as their own AS number in announced routes. We found altogether 274 prefixes were announced with private AS numbers as the origin AS some time in the history data of 3 months. Although during the examined five day period, these prefixes have

changed the origin ASes to be 88 non-private ASNs, they may be mismanaged given the history. Based on such intuition, we correlate these 274 prefixes with darknet prefixes and found that 139 (41.9%) of them are observed in darknet. These darknet prefixes are announced by a total of 37 ASes (41.9% of all 88 anomalous ASes) with non-private ASN in the latest routes. These 37 ASes were possibly misconfigured to use private ASN some time in the history. Furthermore, we found that although these networks do not cause much traffic, they cause relatively many updates, shown in the last three columns of Table 6.

3.4 Bogon prefix announcements

Bogon prefixes are private (in RFC1918 space) or unallocated addresses. We use the bogons list in CIDR reports [15] to identify altogether 168 ASes announcing such routes within the previous four months from the darknet data time period. Among these ASes, we found most of them are observed in darknet: 139 ASes or 83.3% of the 168 ASes, shown in the first and third column in Table 6. These ASes are the origin ASes of 474 darknet prefixes, shown in the second column. Although these ASes correspond to a small portion in all darknet ASes (1.2%), they are responsible to relatively large amount of traffic – 17.8% of total packets received in darknet. Hence, we conclude that most of the networks announcing bogon prefixes are observed in darknet and are also relatively active in sending traffic.

3.5 Address deaggregation

Deaggregation behavior is defined to be advertising many small prefixes covered in larger prefixes already present. It can be caused by misconfiguration, leaking out many subnets from one’s internal network. Routers receiving such deaggregated announcements may experience resource exhaustion in router memory and CPU processing.

Keeping track of the prefixes announced by each AS in the routing table and all distinct prefixes each AS announces during all five days, we observe that altogether 1003 ASes have deaggregation behavior, 104 of which are observed in darknet. Again, these observed 104 ASes cause relatively large amount of traffic in darknet. Moreover, they are responsible for a large fraction of updates. They are possibly more mismanaged on the control plane.

3.6 AS path loops

It is recommended that routes containing AS loops should not be used due to possible forwarding loops. Each router should perform loop detection on received routes and exclude such routing updates. However, in practice we still observe AS routing loops in advertised routes. These are likely caused by ASes not performing the loop checking or incorrectly prepending their own AS numbers in the paths. Thus, the repeated AS in the AS path is likely responsible.

Since the responsible ASes in the routing loop are more likely to be mismanaged, we analyze all the AS path loops within 5 days period. Interestingly, we found that 21 ASes, 80.7% of all the responsible ASes, are observed in Darknet.

Although they are a very small fraction of Darknet ASes, they are responsible for relatively large amount of traffic. We further analyze the prefixes announced by these ASes and observe 33.2% of these prefixes appear in darknet, which do not cause many updates.

3.7 Suspected IP hijacking attacks

The suspicious MOAS behavior is defined as the prefix announced as originated by an AS which rarely originates such a prefix in history. History information helps reduce false positives. We measure the rareness of one prefix with an origin AS by the estimated probability, which is defined to be the fraction of times such a prefix was announced by an AS compared to all announcements for such prefix in history data. Instead of analyzing the suspicious MOAS behavior in all prefixes, we group prefixes by different metrics. Then we analyze the anomalous fraction of prefixes in both darknet and all prefixes under each group. The first grouping metric is mask length. The intuition is that mismanaged networks are usually small edge networks compared to large ISP networks, and announce small prefixes. Mask length determines network size. Smaller prefixes (*e.g.*, $\text{len} \geq 22$) are observed to have a larger fraction of MOAS anomalies in darknet. We found that a larger fraction of darknet prefixes are announced by rarely found origin ASes. As shown in Figure 7, with strict definition of rareness (smaller probability value), there are more anomalous prefixes in darknet compared to the general case.

The second grouping metric is AS type. As shown in Figure 2, ASes can be classified into different geographic locations and business types. Similar to prefix mask length, business type also indicates the possibility of mismanagement. For example, Asian ASes, broadband ASes and educational ASes are more likely mismanaged. As shown in Figure 8, we analyze the suspicious MOAS behavior under different geographic locations and business types. Interestingly, we found that in educational network, darknet always has more suspicious MOAS incidences compared to the case for all prefixes.

Previous work [1] has shown that there is a strong correlation between prefixes originating spam activities and short-lived BGP announcements of such prefixes, an indication of possible IP hijacking. Besides spamming, other malicious activities can also be conducted by IP hijacking. In the following, we correlate darknet traffic with BGP updates. We identify in darknet data ASes using short-lived BGP announcement to perform scanning, propagate worm, or other malicious behavior, and study the duration of such short-lived BGP announcements.

Similar to [1], we consider routing announcements lasting less than one hour as short-lived. We define the time period in which we observe both darknet traffic and short-lived announcements to be the overlap period. We observe 0.02% prefixes announced by 183 ASes whose overlap period covers 80% of the entire period when traffic is sent. 90% of the short-lived announcements last less than 30 minutes. Moreover, we observe around 20 IPs only send traffic whenever there is a corresponding short-lived announcement, indicating possible IP hijacking for performing malicious scanning.

4. MALWARE MITIGATION BASED ON JOINT DATA AND CONTROL INFORMATION

ISPs control the Internet routing infrastructure and have traditionally been paid to carry packets regardless of their

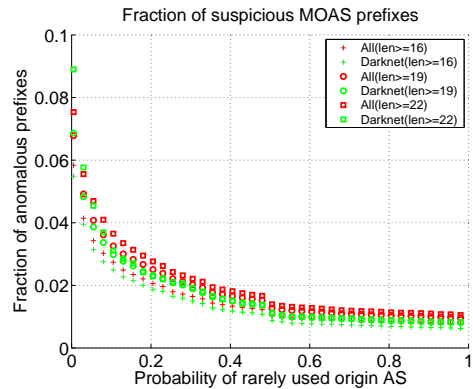


Figure 7: Suspicious MOAS prefixes with different mask lengths




Figure 8: Suspicious MOAS prefixes with different AS types

intent. The major implication of this arrangement is that Internet security is usually performed at the destination rather than at the source. We thus spend vast amount of resources to block traffic that attacks us rather than block attacks before they leave a network. On the other hand, recurring malware is an Internet-wide problem and one possible Internet-wide solution is simply to filter all misconfigured and malicious traffic before it reaches its destination. The critical problem is how to detect and then mitigate on an Internet-wide scale. We decompose filtering solutions into four main categories.

Filter the data plane using data plane information: This is the traditional approach used by firewalls, IDS's, and other common security systems. This approach is very effective at filtering malicious and misconfigured behavior when the behavior can be characterized using static signatures. However, malware activity is constantly changing and evolving so creating and deploying signatures is sometimes ineffective, expensive, and requires on-going effort. Furthermore, payload-level analysis is extremely processing intensive and thus expensive on high-speed links.

Filter the control plane using data plane information: Another filtering approach is to use information gathered by monitoring the data plane to filter the control plane. One simple approach is to have ISPs reject routes advertised by ASes that have been classified as "infected" using data plane intelligence such as packets captured by a honeypot. In order for such an approach to be feasible the number of infected ASes must be few. A more fine-grained approach is to filter only those prefixes classified as "infected". However,

we found that approximately 31% of all the ASes and 33% of all the prefixes are observed to send at least one packet to the darknet over the 5 days, suggesting that the sources for worms, bots, spam, and other anomalous data plane processes are widespread and cannot be eliminated by filtering the route advertisements from a few ASes.

Filter the control plane using control plane information: A third approach is to again filter the routes in the control plane but instead use the control plane as the source of detection data to classify certain routes as “infected”. We have shown that certain classes of mismanagement on the control plane are correlated with data plane anomalies. This is a very important result because the data suggests that *certain classes of data plane anomalies can be predicted by examining purely control plane information*. The challenge is that filtering only control plane data using control plane anomaly detection only would again produce very coarse-grained results as described above. It might be possible for stopping certain specific types of malicious behavior like route hijacking [16], however, it would be very difficult to deploy control plane filters to stop general malicious and misconfigured behavior.

Filter the data plane using control plane information: The final filtering approach is to filter the data plane based on information collected on the control plane. For example, if a control plane IDS (CIDS) produced an alert indicating a specific prefix might be hijacked, a data plane filtering device could provide little value beyond what a control plane filter could have accomplished.

Hybrid control plane/data plane filtering: A powerful use of control and data plane information is to combine them. For example, a CIDS could be used to (1) detect misconfigured and suspicious prefixes (2) attach a confidence factor to each Internet address which is then fed to a data plane filtering device (3) the data plane filtering device then inspects each incoming packet and only performs more expensive operations like deep packet inspection on packets whose source Internet address has a low confidence factor (*i.e.*, are from a suspicious address block).

There are clearly caveats to this approach including the need to validate that the source address is not spoofed (to help avoid this problem one might take the confidence factor into account only after a TCP session is established) but there is still significant potential for reducing the load on expensive inline data plane filtering devices.

We show an example from the perspective of a local ISP the impact of filtering using joint control and data information. We analyze traffic using Netflow data from March 1 to March 5 collected in our local /16 network. Although prefixes contributing heavily to malicious activities indicated in darknet data also send significant amount of normal traffic found in Netflow data, we can still effectively filter darknet traffic with low false positives using the criteria based on the cumulative probe duration of traffic sources to the darknet. Specifically, filtering top 20 prefixes sorted by the probe duration helps reduce 12% of the darknet traffic while only affecting 1.5% of legitimate traffic. Similarly, filtering top 20 ASes sorted by probe duration eliminates 51% of darknet traffic while only affecting 4.5% of legitimate traffic. This shows a significant reduction of malicious traffic with low false positives.

5. CONCLUSION

This paper illustrates how to characterize the control plane data of anomalies detected in the data plane and

demonstrates how networks that are worm-infected, spam origins, and the source of other anomalous data plane behavior also exhibit more anomalous control plane behavior compared to other ASes. We use this information to evaluate the utility of combined data plane and control plane approaches to malware mitigation. Finally, we propose a hybrid approach by using anomalies detected on the control plane to classify packets as suspicious based on their source address to decrease false-positives and increase performance of data plane filtering systems.

6. REFERENCES

- [1] A. Ramachandran and N. Feamster, “Understanding the Network-Level Behavior of Spammers,” in *Proc. ACM SIGCOMM*, 2006.
- [2] P. Barford, R. Nowak, R. Willett, and V. Yegneswaran, “Toward a Model for Sources of Internet Background Radiation,” in *Proc. of the Passive and Active Measurement Conference (PAM ’06)*, March 2006.
- [3] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson, “Characteristics of Internet Background Radiation,” in *Proc. of ACM IMC*, 2004.
- [4] Matthew Braverman, “MSRT - Progress Made Lessons Learned.” <http://www.microsoft.com/>, 2006.
- [5] M. Bailey, E. Cooke, D. Watson, F. Jahanian, and J. Nazario, “The Blaster Worm: Then and Now,” *IEEE Security & Privacy*, vol. 3, no. 4, pp. 26–31, 2005.
- [6] “University of Oregon Route Views Archive Project.” www.routeviews.org.
- [7] “Ripe NCC.” <http://www.ripe.net>.
- [8] M. Bailey, E. Cooke, F. Jahanian, J. Nazario, and D. Watson, “The Internet Motion Sensor: A distributed blackhole monitoring system,” in *Proceedings of Network and Distributed System Security Symposium (NDSS ’05)*, (San Diego, CA), February 2005.
- [9] D. Moore, G. Voelker, and S. Savage, “Inferring Internet Denial of Service Activity,” in *Proc. USENIX Security Symposium*, August 2001.
- [10] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, “Characterizing the Internet hierarchy from multiple vantage points,” in *Proc. IEEE INFOCOM*, 2002.
- [11] R. Mahajan, D. Wetherall, and T. Anderson, “Understanding BGP misconfigurations,” in *Proc. ACM SIGCOMM*, August 2002.
- [12] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, “An Analysis of BGP Multiple Origin AS (MOAS) Conflicts,” in *Proc. ACM SIGCOMM Internet Measurement Workshop*, November 2001.
- [13] P. Boothe, J. Hiebert, and R. Bush, “How Prevalent is Prefix Hijacking on the Internet.” NANOG36 Talk, February 2006.
- [14] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, “BGP Routing Stability of Popular Destinations,” in *Proc. ACM SIGCOMM Internet Measurement Workshop*, November 2002.
- [15] G. Huston, “CIDR REPORT.” <http://www.cidr-report.org/>.
- [16] J. Karlin, “Pretty Good BGP and the Internet Alert Registry.” Nanog 37, June 2006, <http://www.nanog.org/mtg-0606/karlin.html>.