

Toward Understanding Distributed Blackhole Placement

Evan Cooke, Michael Bailey, Z. Morley Mao,
David Watson, Farnam Jahanian
University of Michigan
{emcooke, mibailey, zmao, dwatson, farnam}@umich.edu

Danny McPherson
Arbor Networks
danny@arbor.net

ABSTRACT

The monitoring of unused Internet address space has been shown to be an effective method for characterizing Internet threats including Internet worms and DDOS attacks. Because there are no legitimate hosts in an unused address block, traffic must be the result of misconfiguration, backscatter from spoofed source addresses, or scanning from worms and other probing. This paper extends previous work characterizing traffic seen at specific unused address blocks by examining differences observed between these blocks. While past research has attempted to extrapolate the results from a small number of blocks to represent global Internet traffic, we present evidence that distributed address blocks observe dramatically different traffic patterns. This work uses a network of blackhole sensors which are part of the Internet Motion Sensor (IMS) collection infrastructure. These sensors are deployed in networks belonging to service providers, large enterprises, and academic institutions representing a diverse sample of the IPv4 address space. We demonstrate differences in traffic observed along three dimensions: over all protocols and services, over a specific protocol and service, and over a particular worm signature. This evidence is then combined with additional experimentation to build a list of sensor properties providing plausible explanations for these differences. Using these properties, we conclude with recommendations for better understanding the implications of sensor placement.

Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection—*Invasive Software*

General Terms

Security Measurement

Keywords

network security, computer worms, globally scoped threats, blackhole monitoring, blackhole placement, Internet Motion Sensor

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'04, October 29, 2004, Washington, DC, USA.
Copyright 2004 ACM 1-58113-970-5/04/0010 ...\$5.00.

1. INTRODUCTION

As network-based threats become increasingly prominent, characterizing, monitoring, and tracking these threats is critical to the smooth running of individual organizations and to the Internet as a whole. To increase their view of these threats, researchers and network operators are instrumenting unused address space. While tools to monitor unused address space are gaining in popularity, there is still an open question as to how applicable results from one address block are to the Internet as a whole. In this paper we demonstrate that achieving a representative sample may not be as simple as monitoring a few unused address blocks. We present empirical evidence that different address blocks observe significantly different traffic volumes and patterns. This evidence is then combined with additional experimentation to build a list of sensor properties providing plausible explanations for these differences. Using these properties, we conclude with recommendations for understanding the implications of sensor placement.

Monitoring network traffic for malicious content takes many different forms. By far the most common technique is passive measurement of live networks, which falls into three main categories: data from security or policy enforcement devices, data from traffic characterization mechanisms, and direct sensing or sniffing infrastructure. By either watching firewall logs, looking for policy violations, or by aggregating IDS alerts across multiple enterprises [20, 30], one can infer information regarding a worm's spread. Other policy enforcement mechanisms, such as router ACLs provide coarse-grained information about blocked packets. Instead of using ACLs to monitor dropped packets, CenterTrack [25] leverages the existing routing infrastructure to collect denial of service traffic for analysis. Data collection techniques from traffic planning tools offer another rich area of pre-existing network instrumentation useful in characterizing threats. Coarse-grained interface counters and more fine-grained flow analysis tools such as NetFlow [5] offer another readily available source of information.

Another compelling alternative to measuring live networks is monitoring blocks of unused address space. Because there are no legitimate hosts in an unused address block, traffic must be the result of misconfiguration, backscatter from spoofed source addresses, or scanning from worms and other network probing. This pre-filtering of the traffic eliminates many of the false positive and scaling issues of other monitoring approaches. There have been various names used to describe this technique such as network telescopes [14], blackholes [22, 15], and darknets [27]. The most common application of this technique is the global announcement and routing of unused space to a collection infrastructure that records incoming packets. This technique has been used both by host-based honeypot tools [24] and by wide address space monitors [11, 14, 22]. More recently, researchers have combined ideas from host-based honeypots with wide address measurement to elicit a behav-

ior that is only visible by participating in a network or application session. For example, software like honeypot [18] and the iSink [31] bring up a network of virtual honeypots over a single address space.

Using these techniques, researchers have successfully characterized and classified the traffic observed at unused blocks [12, 16]. The investigation in [16] utilized blocks located in three Class A networks and was able to see a large number of Internet threats. One interesting feature of the plots presented in that analysis were the differences in magnitude and composition of traffic between the different blocks. This is an important observation because blackholes sample a small portion of the total used address space, it is difficult to know if that sample is generalizable.

One approach for obtaining representative data is to increase the number and size of unused address blocks [11]. However, without understanding how the *placement* of the monitoring blocks relates to the traffic observed, there is no way to know if that sample is representative.

To better understand how observed traffic is affected by sensor placement, we use data from the Internet Motion Sensor [6] to present evidence that distributed unused address blocks observe significantly different traffic patterns. The Internet Motion Sensor (IMS) is a distributed collection of blackhole sensors. These sensors are deployed in networks belonging to service providers, large enterprises, and academic institutions representing a diverse sample of the IPv4 address space.

This paper is divided into three major sections. § 2 details the IMS collection infrastructure and describes the sensor deployments. Next, § 3 demonstrates differences between blackholes using three successively more specific views. Finally, § 4 discusses the features of sensor placement that affect what a sensor observes.

The main contributions of this work are:

- Deployment of ten distributed blackhole sensors at major service providers, large enterprises, and academic networks with address blocks that range in size from a /25 to a /8
- Identification of sensor placement as an important factor in understanding and generalizing measurements from unused address space
- Strong empirical results showing differences between traffic observed on diverse distributed blackhole sensors
- Definition and application of sensor properties to help explain differences in traffic measurement on blackhole sensors

2. IMS ARCHITECTURE OVERVIEW

The data collected and analyzed in this paper was collected using the Internet Motion Sensor distributed blackhole infrastructure. It is important to understand this architecture because blackhole sensors have a wide range of measurement fidelities. More specifically, the data collected depends on the extent to which a sensor emulates the services and characteristics of a real host, similar to the interaction spectrum of honeypots [23]. The IMS is designed to provide a consistent and comparable level of interaction across all sensors. The level of sensor emulation was chosen to:

- Maintain a level of interactivity that can differentiate traffic on the same service.
- Characterize emerging threats.
- Provide visibility into Internet threats beyond geographical and operational boundaries.

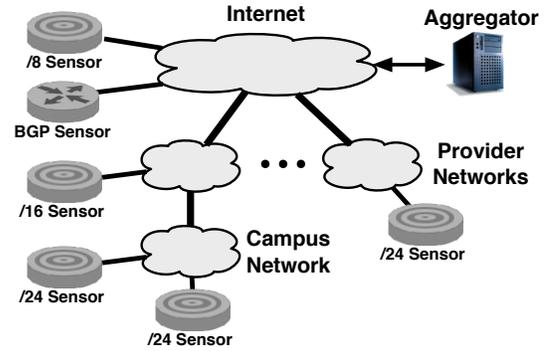


Figure 1: Internet Motion Sensor Architecture

The IMS architecture consists a set of heterogeneous sensors and data aggregators as depicted in Figure 1. The sensors deployed in the IMS can be divided into two basic categories: blackhole and topology sensors. The blackhole sensors form the core of the IMS by collecting threat data while the topology sensors provide context for that information.

Each blackhole sensor monitors a dedicated range of unused IP address space. The blackhole sensors in the IMS have an active and passive component. The passive component records packets sent to the sensor's address space and the active component responds to specific packets to illicit more data from the source.

The active component in each blackhole sensor only responds to TCP connection requests. UDP and ICMP packets do not need an active response because the sensors can collect all the information available from the initial packet without taking on the personality of specific services. For example, the Witty [21] and Slammer [12] worms were based on UDP in which the entire worm payload was transmitted in the first packet. TCP, on the other hand, is a connection based protocol and requires an active response to elicit payload data. Using an active responder allows the IMS to capture the payloads of TCP worms like Blaster [4] and Sasser [10].

The data storage mechanism in the IMS blackhole sensors is an integral part of data analysis. One of the main design goals for the IMS was to support real-time trending and analysis, necessitating a novel approach to data processing. The problem is that raw traces gathered on each sensor can approach a gigabyte of data per day. Attempting to transmit and store that data for a large number of sensors while supporting real-time data analysis requires very significant infrastructure. The obvious solution is to distribute the workload. Each blackhole sensor is responsible for gathering and archiving data, performing queries on its local data store, and generating alerts that are sent to the aggregator.

Storing the full payload for every packet has significant space requirements, so the IMS uses a novel payload storage approach. When a blackhole sensor receives a packet with a payload, it first computes the MD5 checksum of the payload (without network headers) and compares it against the checksum of all the other packets it has seen in the past day. If the checksum (signature), has already been recorded, the capture component logs the signature but does not store the payload. If the signature is new, the payload is stored and the signature is added to the database of signatures seen in that day.

The results presented in this paper are based on data collected on an IMS deployment of ten blackhole sensors. These deployments include major service providers, a large enterprise, and academic

Label	Organization	Size
A	ISP	/23
B	Academic Network	/24
C	Academic Network	/24
D,E,F	ISP	/20, /21, /22
G	ISP	/25
H	Large Enterprise	/18
I	National ISP	/17
J	ISP	/8

Table 1: IMS Blackhole Deployments

networks with address blocks that range in size from a /25 to a /8. Table 1 indicates the different deployments and their associated anonymized labels which are used to annotate the plots shown in the next section. These sensors represent a range of organizations and a diverse sample of the routable IPv4 space including seven of all routable /8 address ranges. This organizational and address space diversity allows us to compare sensors placed in widely different locations. In addition, we are also able to compare finer grained differences between sensors because several deployments at a single organization share the same /16 prefix.

3. OBSERVATIONS FROM A DISTRIBUTED BLACKHOLE DEPLOYMENT

The following investigation is motivated by the observation that there are significant differences in the traffic observed in the distributed IMS blackhole sensors. Because many Internet threats today, like many worms, are globally scoped, one might expect somewhat similar traffic on equally sized blackholes. Recall that unlike live networks, traffic must be the result of misconfiguration, backscatter from spoofed source addresses, or scanning from worms and other probing. Despite the global nature of these threats and the lack of legitimate traffic, the IMS distributed sensors see widely different amounts of traffic along several important dimensions.

This section probes these differences using three successively more specific views of traffic to a network of distributed blackhole sensors. The data was recorded over a one month period with SYN responders on TCP port 135, 445, 4444, and 9996 across all sensors. These ports were chosen in order to monitor specific worms. The first view looks at the traffic observed at each sensor over all protocols and services. The second view looks at traffic to a specific protocol and port. Finally, the third view looks at the signature of a known worm over all sensors.

3.1 All protocols and services

We begin by looking at the packet rate observed by each sensor. Figure 2 shows the average amount of traffic over all protocols and services observed by ten blackhole sensors. Packets are normalized by the size of a /24 so sensors covering different sized blocks can be compared. Normalization is performed such that the magnitude seen in a /23 would be divided by two and traffic in a /25 multiplied by two. The granularity of a /24 was chosen to focus the discussion on larger differences between blocks rather than individually targeted addresses. While differences likely exist within individual /24 blocks, this is beyond the focus of this paper.

Figure 2 clearly shows that the amount of traffic varies dramatically and can differ by more than two orders of magnitude between sensors. Of note, the larger blocks typically observe less traffic per /24 than the smaller blocks. One possible explanation is that the smaller blocks are closer in the IPv4 address space to live hosts than the large blocks. There are many reasons a person or program may

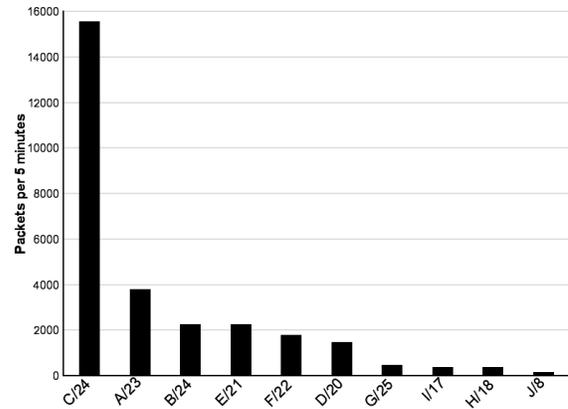


Figure 2: Packet rate as seen by each sensor normalized by /24

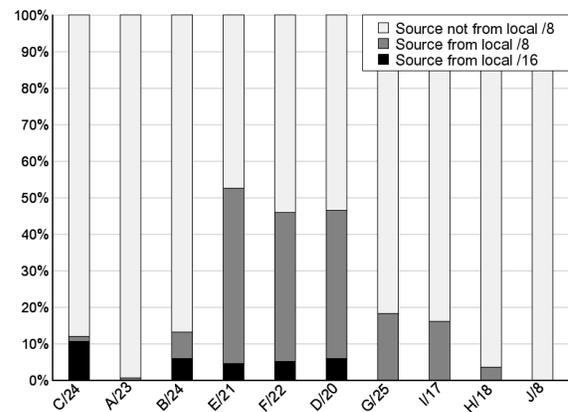


Figure 3: Distribution of local preference per sensor

prefer to scan local addresses. Take for example someone on a university network attempting to find open file shares. This person will likely scan other blocks in the same /16, assuming these address blocks are also in the university space. Another example is Internet worms which have a local address preference in their scanning algorithms, such as Code Red II [2], Nimda [1], and Blaster [4].

If local traffic was a significant component of the overall traffic, it might explain the differences seen between the sensors. In particular, if the blocks which see the most traffic also receive a large amount of local traffic, that might explain the overall inequalities. Figure 3 shows the percentage of traffic to all protocols and services from the same local /16 and local /8 as the sensor where the traffic was observed. There are two important implications of this graph. First, there are very different relative amounts of local /8 traffic seen at the various sensors. Interestingly, the three blocks belonging to a single ISP (D/20, E/21, F/22) observe close to 50% local traffic and the other blocks each see less than 20%. One might expect more local /8 traffic in the academic networks given less central control, but clearly other organizations also show significant percentages of local traffic. Second, although some sensors see a very significant amount of normalized local /8 traffic, those blocks do not correlate with the sensors with the greatest magnitude of overall traffic. For example, C/24 observes by far the greatest amount of traffic but less than 10% of that traffic is from within the same /8 as the

Sensor	80	135	137	139	445	1025	4444	Other
A/23	1.2	9.5		1.3	16			4.9
B/24		11		0.2	68		0.7	1.1
C/24		104	0.3		30		0.6	3.9
D/20		4.1		0.7	4.9	0.7		2.2
E/21		5.4		1.0	9.5	0.7		3.1
F/22		4.5		0.9	7.0	0.7		2.4
G/25		1.4	0.2		1.2		0.5	0.6
H/18		0.5		0.1	1.8		0.1	0.4
I/17		1.3		0.1	1.1		0.1	0.5

Table 2: Packets (1,000,000s) to Top 4 TCP destination ports normalized by /24

sensor. So, even though local traffic can be significant, the major traffic differences are not due to local preference.

The next step is to break down the traffic by protocol and port. We start by looking at the protocol distribution at each sensor. When the global (non-local /8) traffic is broken down by protocol, almost 99% of that traffic is dominated by TCP. This is logical because large portion of malware and scans are targeted toward TCP and because the IMS sensors actively respond to TCP packets. One SYN packet response typically generates three or more followup packets not counting connection retries. Thus, the differences between sensors are dominated by differences in TCP traffic.

Given that TCP traffic dominates at all sensors, the next question is what TCP destination ports are being targeted. Table 2 shows the top four TCP ports at each sensor normalized by a /24. Notice that the distribution is different across sensors and only TCP port 445 and 135 are consistent across all sensors. Despite a preference toward actively responded ports, it is interesting that other TCP ports like 137 and 139 show up as top ports. The similarity in the top ports across all sensors implies that the differences we see cannot be explained by traffic targeted at ports observed at some sensors but not at others.

3.2 TCP port 135

Thus far we have shown that the differences between sensors are not due to a simple port or protocol bias. To the contrary, TCP accounts for almost all traffic and TCP port 135 and 445 are major contributors to the traffic at all sensors. If these ports account for such a large portion of the overall traffic, an interesting question is whether this traffic is correlated with the differences in overall traffic. TCP port 135 is used by the Windows DCOM RPC service which has been the target of several vulnerabilities and exploits. In particular, TCP port 135 is the infection vector for the well known and still prevalent Blaster worm.

Figure 4 shows the number of unique source IPs observed on TCP port 135 across the individual /24s that make up the sensor blocks. Each bar in the figure represents the number of unique source IPs seen at a /24. Thus, large address blocks like I/17 are composed of many component /24s that are all shown individually. Blocks are labeled on the horizontal axis and the separation between blocks is denoted with a small amount of white space. Sensor blocks smaller than a /22 and the J/8 block are not shown for legibility reasons.

Figure 4 has two important implications. First, there are significant differences between the component /24s within the each block. This means that even the destination addresses within a sensor block observe different sources. Second, there are even larger disparities between the unique sources seen across sensors, and these differences correlate with differences in overall traffic (Figure 2).

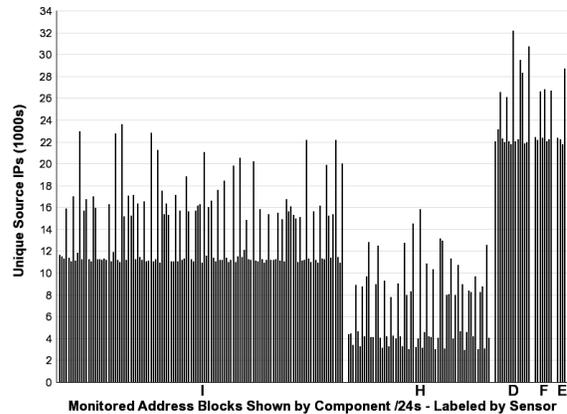


Figure 4: Unique source IPs to TCP 135 by /24 excluding local /8 traffic

Intra-Sensor Standard Deviation

D/20, E/21, F/22	3188
H/18	3564
I/17	3478

Inter-Sensor Standard Deviation

All Sensors	21116
-------------	-------

*Single /24s are not shown because variance is zero.

Table 3: Inter-sensor and Intra-sensor variance in unique source IPs between /24s

In order to quantify the inter-sensor and intra-sensor differences we analyzed the variances between the /24s within a given sensor and between the sensors. Table 3 summarizes the results. As we saw in the graphical representation, the inter-sensor variance is much greater than the intra-sensor variance. To support this hypothesis we used an Analysis of Variance (ANOVA) to show that the inter-sensor differences are indeed significant. We consider seven groups of sensors (D/20, E/21, and F/22 are grouped together) including those with only a single /24 block. Using F distribution with 7 degrees of freedom in the numerator, and 217 degrees of freedom in the denominator, we calculate an F statistic of 145.1. The calculated P-value is extremely close to zero and much smaller than the typical α value of 0.05. Therefore, we can reject the null hypothesis and conclude the inter-sensor variance cannot be due to chance and represents some other process.

Table 3 clearly demonstrates that significant differences exist in the number of source IPs observed within and between sensors. Another metric to analyze the differences within and between blocks is the rate at which packets are received. While the number of sources indicates the number of senders, the packet rate indicates how fast a particular address block is receiving packets. Figure 5 shows the average packet rate on TCP port 135 to each /24 labeled by address block. There are two important conclusions that can be drawn from this graph. First, there appears to be a significant bias toward the beginning of each block which tails off toward the end of the blocks (reading left to right). The intra-sensor variance is quite small so the tails do not appear to be due to a random process.

The second important conclusion that can be drawn from Figure 5 is that the inter-sensor variations in this plot are very similar to the inter-sensor variations in the unique IPs plot 4. The implication is that, on average, each unique source sends a simi-

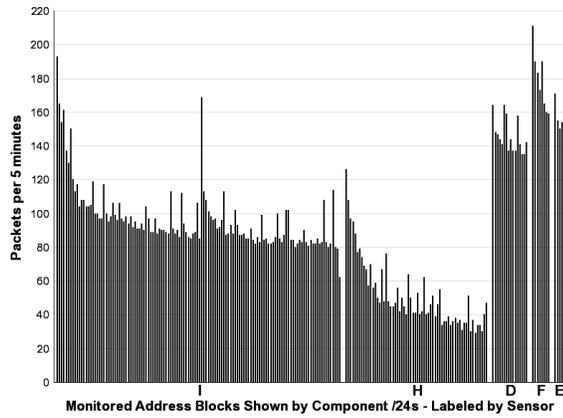


Figure 5: Packet rate of TCP 135 by /24 excluding local /8 traffic

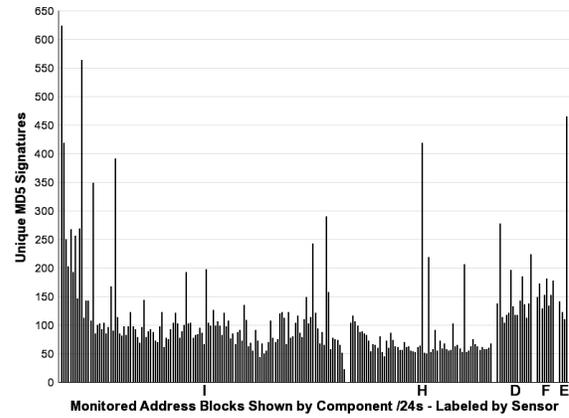


Figure 7: Number of unique signatures to TCP 135 by /24 excluding local /8 traffic

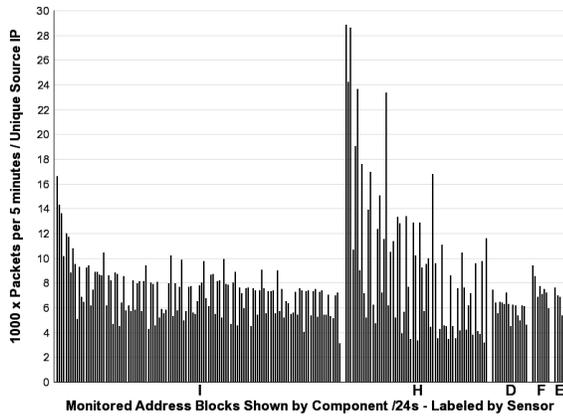


Figure 6: Ratio of packet rate to the number of unique sources observed on TCP 135 by /24 excluding local /8 traffic

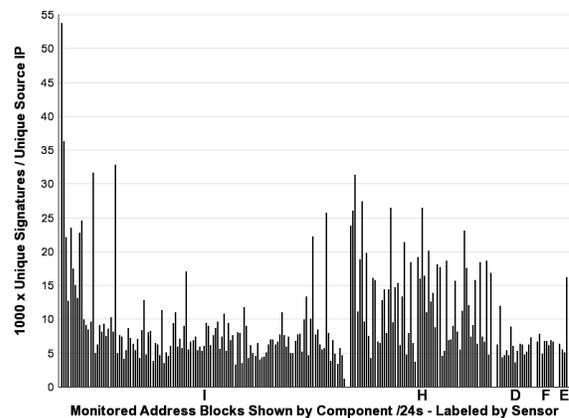


Figure 8: Ratio of the number of unique signatures to the number of unique sources observed on TCP 135 by /24 excluding local /8 traffic

lar amount of packets regardless which block the source is sending to. This relationship is shown clearly in Figure 6 which depicts the average packet rate per unique source IP on TCP port 135 to each /24. Thus, no block appears to be the target of long lasting directed attacks by a small number of sources. Recall that the data presented is the aggregate over a month so short attacks are mostly likely masked. The important message is that the average packet rate per unique source is similar across blocks.

Another measure of the difference between blocks is the kinds of payloads each block observes. When a packet is received by an IMS blackhole sensor, the capture mechanism computes the MD5 hash of the payload and only stores that payload if the hash has not been seen before. While the MD5 hash (also referred to as a payload *signature*) does not say anything about the contents of the payload, it is a simple means of differentiating payloads. Figure 7 shows the number of unique MD5 hash signatures observed on TCP port 135 to each /24. Recall that a lightweight SYN-ACK responder was running on TCP port 135 for the duration of the experiment over all IPs. Hence, the payloads received all result from one SYN-ACK packet.

The interesting message in Figure 7 is that all blocks see a similar number of unique signatures with a large intra-sensor variance.

While Figure 4, depicting the number of unique source IPs by /24, showed a large inter-sensor difference, Figure 7 indicates very little difference between blocks. The implication is that many unique sources are sending exactly the same payloads. Figure 8 shows the average number of unique signatures per unique IP on TCP port 135 by /24. Even at the /24 that sees the most new signatures, it only sees 5 new unique signatures per 100 new unique hosts on average. The conclusion is that most the differences between blocks on TCP port 135 cannot be due to certain payloads being targeted specifically at one or a few blocks.

The analysis of TCP port 135 demonstrates that the number of unique source IPs and number of packets per time observed varies significantly between blocks. In contrast, the number of unique signatures and the overall behavior per unique source (e.g. unique signatures per unique host) is similar between blocks. It appears there is no dominate behavior which is directed at a specific block. In fact, it appears certain blocks are simply more heavily targeted and thus see more unique source IPs and more packets. It must be remembered that the data being presented is the aggregate over a longer time scale so smaller features are lost. For example, it is very difficult to isolate specific behaviors like linear scans from

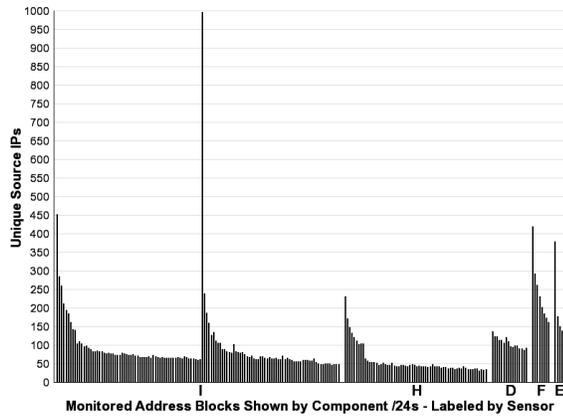


Figure 9: Unique Source IPs of Blaster Infection Attempts by /24 with local /16 traffic removed

the Blaster worm or random infection attempts from the Slammer worm. The next section picks up on this idea and isolates one particular propagation strategy and attempts analyze the differences seen between blocks.

3.3 Blaster signature

To test the propagation hypothesis, we isolated a particular signature which has well known propagation strategy. The signature we choose was that of the Blaster worm. Blaster has a simple propagation mechanism which is based on a sequential scan through IPv4 space. When the Blaster worm is launched due to a new infection or a rebooted computer, the worm choses an initial target address that is in the same local /16 as the source 40% of the time and a completely random address the other 60%. The Blaster worm will then scan sequentially from that initial address attempting to infect IPs in blocks of 20 at a time.

Figure 9 shows Blaster infection attempts by unique source IPs as seen by a /24. In order to eliminate the possibility of certain blocks being biased by Blaster’s local preference, Blaster sources from the same /16 were eliminated. This figure again reveals large differences between sensors. This is a very surprising result. Even though we have attempted to control for propagation strategy, there are still significant differences in the number of unique sources between sensors. Another interesting observation is that the sensor blocks which observed more overall traffic (Figure 2) also observed relatively more Blaster sources. However, there are certain hotspots (for example, the large spike in I/17) in the middle of blocks which do not correlate with patterns in overall traffic. The regularity of the distribution indicates some other non-random process may be at work. For example, a poorly designed random number generator or a bad source of entropy may contribute to the results depicted in Figure 9.

This section has demonstrated differences between sensors using three successively more specific views of traffic to a network of distributed blackhole sensors. The first view showed differences in the traffic observed at each sensor over all protocols and services. The second view demonstrated that these differences persist on traffic to TCP port 135. It was also shown that inter-sensor differences dominated intra-sensor differences. The third view established that differences existed even when observing the Blaster worm which has a known propagation strategy. The result is that there is no definitive explanation for the differences between sensors and more investigation is required. In the next section, we enumerate the dis-

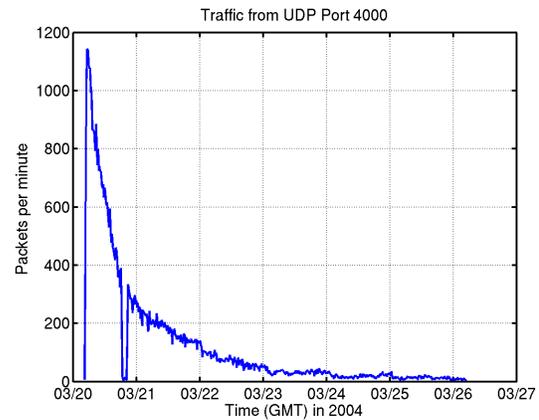


Figure 10: The Witty worm as recorded by three blackhole sensors

tinguishing properties of blackhole sensors in order to better understand how these differences arise.

4. EXPLAINING DIFFERENCES IN SENSOR OBSERVATIONS

There are several properties of blackhole sensors that influence the traffic they observe. These properties provide valuable insight into why differences exist between sensors and the importance of understanding sensor placement. Furthermore, this discussion provides a methodology for characterizing sensor properties which aids in interpreting blackhole sensor measurements. In addition, these properties can guide decisions about sensor placement, allowing maximization of limited resources. Our goal is to enable more accurate predictions of global attack trends. To illustrate the impact of these properties, we use concrete examples based on our measurement data as well as simple analysis.

4.1 Filtering policy

It is important to understand factors that affect the *reachability* from sources to the blackhole sensors. One reason why traffic cannot reach a sensor is due to filtering by routers or firewalls. Thus, policies applied at a router/appliance between the traffic source and the sensor can affect the visibility of the traffic at the sensor. These routing policies are in terms of what traffic to filter at routers via route or packet filters. When attacks become known and if they target or originate from a well-known port not used by other popular services, ISPs typically install access control lists to filter such traffic. More sophisticated packet filters based on packet content signatures can also be deployed. We differentiate between two types of filtering: at the core and at the edge.

4.1.1 Filtering at the core

If the network paths between end hosts and the sensors traverse through any such filters blocking traffic, then the sensors will not make complete observations. An extreme case would be all network paths traversing a common provider, most likely the immediate upstream provider of the sensors. This problem is illustrated in our own measurement data from sensors deployed in three diverse address blocks (two /16s, one /8) sharing the same upstream provider. That is, the sensors are address-diverse but topologically similar. Figure 10 shows the traffic seen from these three sensors during the Witty worm outbreak. It shows a sharp drop in UDP src

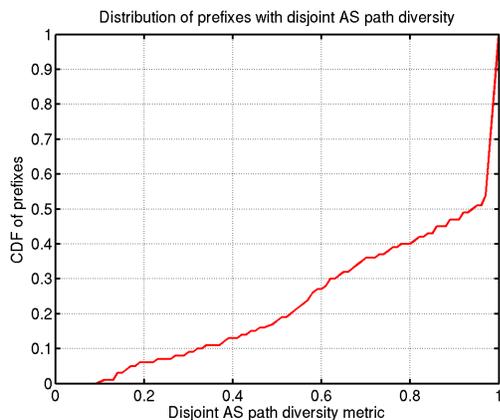


Figure 11: Distribution of disjoint AS path diversity metric for all routable prefixes

port 4000 traffic for a 20 minute period. Upon further investigation, it was revealed that the drop was due to a policy decision to filter such traffic imposed by the shared upstream provider. This type of filtering has been seen frequently during large worm outbreaks. In order to maximize visibility, a blackhole sensor’s immediate upstream provider should impose no filtering policies on the monitored address space.

Considering the possibility of being filtered at the AS level, a good placement strategy is to place sensors in ASes to which other ASes have maximally diverse set of AS paths. This strategy is highly resistant to filtering and has a high probability of finding scan traffic. We call such a property *disjoint AS path diversity*. For example, for AS X where a sensor resides, if all other ASes reach it via AS path $[YABX]$, with only Y being variable, then if either A or B instruments filtering, X will observe no scan traffic. This argues for multihomed sensors and sensors with a diverse set of upstream providers that are all used for live traffic (rather than only providing backup services). Hypothetically, if all tier-1 providers impose filtering at their edge routers, it becomes very difficult for sensors to achieve good coverage due the fact that most AS paths go through a tier-1 ISP. However, it is known that there is an increasing trend for lower tier providers to peer with each other to reduce transit costs [7]. In such cases, traffic does not necessarily have to transit through any tier-1 providers assuming that is where traffic is most likely to be filtered.

To illustrate that address blocks can have varying degrees of disjoint AS path diversity, we do a simple analysis using BGP data from the Oregon RouteViews project [28] which receives BGP feeds from 31 distinct Autonomous Systems (ASes). We assume scanning traffic originates from ASes where we have default-free BGP tables using RouteViews data. We now randomly pick some address blocks commonly announced by most of these 31 ASes and study the disjoint AS path diversity of these AS paths. We use a simple metric to characterize the overlap in these AS paths: we calculate the percentage of AS paths that the next popular AS (excluding the origin AS) occurs in. Popularity is measured by how often an AS appears in the AS paths. The higher the value of this metric, the less disjoint AS path diversity there is. For example, Table 4 shows four randomly selected prefixes with widely diverging degrees of disjoint AS path diversity ranging from 0.14 to 1.00. For the first prefix 4.17.225.0/24, besides the origin AS, AS11853 occurs in all AS paths. If it imposes filtering, no traffic

from the RouteViews ASes can reach the prefix. In contrast, for prefix 65.116.144.0/24, the most common AS AS2914 occurs in only 14% of the AS paths. Thus, the impact of filtering by a single AS can affect at most 14% of the traffic to this sensor.

Obviously, longer AS paths have a higher chance of a commonly shared AS. Upon closer inspection, the address block 65.116.144.0/24, which has the most disjoint AS path diversity, is directly allocated by Qwest, a tier-1 ISP. The other three blocks are allocated to an individual company. This explains why this block has such a diverse number of nexthop ASes and a shorter AS path from other ASes.

Figure 11 shows the distribution of disjoint AS path diversity for all routable prefixes visible in the RouteViews BGP feeds. We observe that less than 20% of prefixes have a diversity value below 50%. It is preferable to select such prefixes in order to be resilient to AS level filtering. In general, it is important for blackhole sensors to have *diverse routing topology*, so that the network paths between end hosts and the sensors have a diverse set of network providers.

It is well known that some countries such as China impose very strict traffic filtering to prevent access to certain Web sites. To overcome such traffic restrictions, sensors need to have sufficient *geographic diversity*. For instance, if sensors are only located in a single country with restrictive traffic filters to the external network, these sensors will have a highly biased view of traffic behavior and may not be able to observe much activity.

4.1.2 Filtering at the edge

So far we have discussed filtering instrumented at the ISP’s networks or provider networks. Filtering can also be done at the edge or stub networks where a sensor resides. Given relatively smaller amounts of traffic, the edge networks can afford more expensive filtering using additional state information. Network intrusion detection systems such as Bro [17] and Snort [19] as well as firewalls are often used by enterprise networks to protect against intrusion attempts. Of interest here are the security policies which influence the filtering of incoming traffic. Imagine an enterprise network hosting the blackhole sensor that disallows all incoming traffic to port 80, 443, 2002, 1978, and 4156, which are ports used by the Slapper worm [26]. In such cases, a sensor residing in this network cannot detect any infection attempts from external networks. If there are compromised hosts inside the enterprise network, scan traffic from internal local hosts can still hit the sensor. In this case, the traffic sources will be biased towards local hosts.

In summary, given service-specific filtering instrumented at the edge networks, it is important for sensors to have *organizational diversity* to achieve good coverage. This means it is best to have sensors placed in different types of organizations, e.g., university networks, large company enterprise networks, small business site, government sites, ISP networks, etc.

4.2 Propagation strategy

Worm target selection algorithms or propagation strategies often have a bias towards local addresses [29], e.g., Code Red II [2], Nimda [1], and Blaster [3] all prefer to scan nearby addresses. This effectively enables a worm to increase the speed of the infection due to smaller network distance to local hosts. Nearby addresses of vulnerable hosts are also target-rich due to common administrative practices. Local preferences also allows a worm to take advantage of breaches in firewalls and other security mechanisms. The easiest way to identify local hosts is to select nearby addresses. A slightly smarter strategy is to use BGP routing information to identify other address blocks, potentially non-contiguous, which originate from the same organization. This idea is similar to taking advantage

Prefix	4.17.225.0/24	12.29.162.0/24	64.106.248.0/21	65.116.144.0/24
Metric	1.00	0.67	0.44	0.14
Shared AS	AS11853	AS1239	AS2914	AS2914
AS-paths	11608 6461 3561 11853 6496 293 3561 11853 6496 6453 701 11853 6496 13237 2914 11853 6496 16150 8434 8210 4200 2914 11853 6496 7018 3561 11853 6496 2905 701 11853 6496 15290 701 11853 6496 6762 701 11853 6496 267 2914 11853 6496 6453 701 11853 6496 3130 1239 3561 11853 6496 10876 1239 3561 11853 6496	11608 2914 7018 12163 293 7018 12163 6453 1239 5778 12163 13237 174 7018 12163 16150 8434 8210 3356 1239 5778 12163 7018 12163 2905 701 1239 5778 12163 15290 7018 12163 6762 701 7018 12163 6453 7018 12163 3130 1239 5778 12163 286 209 1239 5778 12163 6939 7911 1239 5778 12163	11608 3491 22205 6453 2914 22205 293 3356 22205 13237 2914 22205 16150 8434 3257 2914 22205 7018 12182 22205 2905 701 12182 22205 15290 3491 22205 6762 3491 22205 267 2914 22205 6453 2914 22205 3130 2914 22205 286 3491 22205	11608 2914 209 6453 209 293 209 13237 174 209 16150 8434 8210 4200 209 7018 209 2905 701 209 15290 7018 209 6762 701 209 6453 209 3130 2914 209 286 209 6939 6453 209

Table 4: A subset of BGP paths from RouteViews ASes to four randomly selected prefixes as potential candidates for blackhole sensors: showing varying degrees of disjoint AS path diversity.

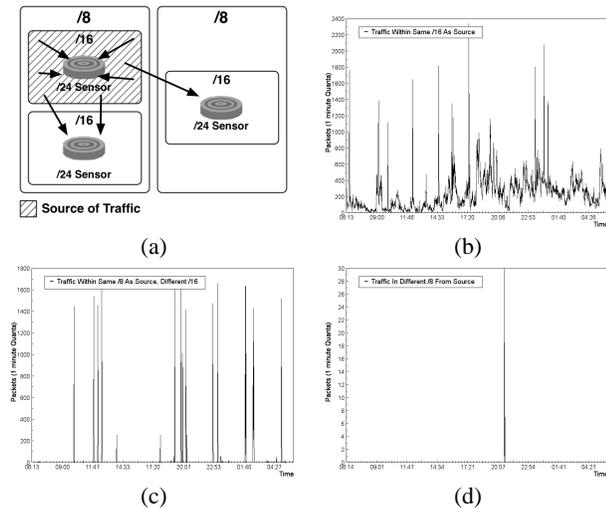


Figure 12: Traffic from a single /16 captured from 3 different perspectives

of BGP routing information by scanning only routable addresses rather than the entire address space [32].

Figures 12 clearly illustrates that there is a strong local preference in traffic directed towards blackhole sensors. Figure 12(a) shows the placement of the sensors. There are three /24 sensors, two within two separate /16s within the same /8, and another in a different /8. Figures (b), (c), and (d) show traffic to the three sensors from hosts within the same /16 as the first sensor. We see that traffic originating from the single /16 network favors sensors in nearby address space. In this case, the sensor in the same /16 sees the most traffic from the source. The sensor in the same /8 sees slightly less traffic, and the sensor in a different /8 rarely sees any traffic. Thus, sensors placed near less centrally managed hosts e.g., cable modem networks and open university campus networks, will more likely see more local traffic.

In addition to local preference, certain networks are more likely to be targeted. This can include highly visible targets (e.g., www.whitehouse.gov), hosts with high traffic volume, and highly active or popular hosts. This implies that blackhole sensors placed in nearby address space to highly targeted networks are more likely to see attack traffic. Overall, proximity to both live hosts and to targeted networks will affect the traffic observed at a blackhole sensor.

4.3 Sensor address visibility

Previous research [8] has shown that 5% of the routed Internet address space is not globally visible. The lack of global reachability can be accounted for by reasons such as misconfiguration, policies, network failures, or even malicious intent. It is important that blackhole address blocks be globally reachable. Instability of the routes to the sensor address space can also result in reachability problems, especially given that route flap damping can be triggered during convergence to suppress unstable routes [9].

Using the BGP updates data from RouteViews BGP monitor, we studied the availability of the routes to the sensor blocks in our deployment from a large set of ASes. There are occasional updates relevant to the sensor address blocks, most of them are announcements with very few withdrawal updates. We studied the duration of the withdrawal, i.e., the time period the withdrawal is in effect until the subsequent announcement, and found that they are of very short durations, typically on the order of 10s of seconds. Furthermore, during the period we captured the packet data, we never found any sequence of routing updates that affected a large number of views or ASes from which we receive the BGP data. This means that there was never a routing instability event related to these sensors with sufficient impact to affect a large portion of the Internet during our measurement time period.

One solution to these potential problems is to make every effort to ensure the stability of the monitored address blocks. For example, if there exist larger address blocks, i.e., a supernet, that cover the sensor address space, then the instability of the sensor address block can be contained and hidden. For instance, a /24 sensor block with a /16 supernet that is also announced and visible in most external routing tables will be less susceptible to flap damping and routing instability. Route announcements pertaining to smaller address blocks are more likely to be filtered than those for larger addresses. Thus, for both stability and reachability reasons, it is good practice to announce the supernet of the sensor address block.

Another way to enhance path and topological diversity and increase routing stability, is to *anycast* the sensor address blocks [13]. Anycast means that the address block is announced at topologically different locations. This translates to multiple diverse AS paths that reach to the same sensor address space. This has the added benefit of load-balancing the monitoring function at different servers. Route selection given a diverse set of AS paths also enables the traffic to flow along a shorter path in terms of AS hop count. Anycast has been used to increase the redundancy of the routes to Root DNS servers. Given that there is a limited number of unused address blocks available, anycast effectively increases coverage as redundant paths make it more resilient to filtering.

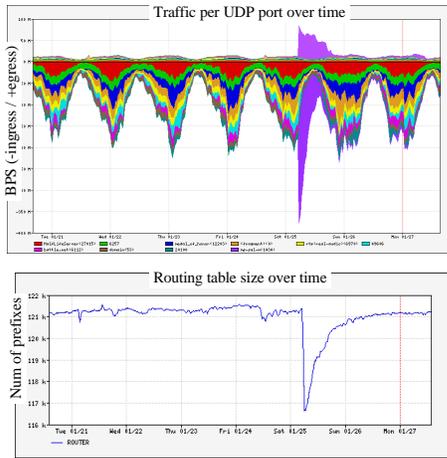


Figure 13: Effect of Slammer on traffic and routes at a tier-2 ISP (data and visualization courtesy of Arbor Networks)

While we didn't see any major routing instability during our own measurement, we still believe that routing instability is a significant potential cause of differences in sensor observation. For this reason it is advantageous to have more stable upstream providers with good connectivity for hosting blackhole sensors. It is essential to analyze the routing instability during the period of sensor observation to account for potential differences in observation.

Similar to the reachability from the external Internet to the blackhole address space block, it is also important to ensure good reachability from the site hosting the blackhole sensor to the rest of the Internet. Redundant upstream connectivity for the sensor address blocks to provide a diverse set of routes from each sensors will increase visibility. If the sensors are multihomed, failure at one upstream provider can be protected against by the other provider(s). Similarly, if one route is suppressed due to flap damping, there still exists an alternate route to reach the sensor. Such *routing redundancy* at the sensor site will improve the chances of a packet from the active responder reaching the original end host. The IMS topology sensors described in Section 2 are used to monitor connectivity. The enhancement of path and topological diversity is a topic of our ongoing work.

4.4 Resource constraints

So far we have discussed how reachability to blackhole sensors can be affected by: filtering policies along the path to the sensors, propagation strategies, as well as the availability of the routes to the sensors. Another important aspect that affects reachability is resource constraints which determine the performance and availability of the data paths to the sensors. The presence of a route in the routing table does not guarantee reachability, as the data path can have performance problems. For instance, if the network hosting the sensors is under a Denial of Service attack, causing the access link to the upstream provider(s) to be congested, scan traffic may not reach the sensors. Transient conditions due to link failures resulting in high packet drop rates inside the local network can also affect observations. Resource constraints are especially important during availability events like the power blackout in August of 2003 which affected the Northeastern US. Another significant availability event was the Slammer/Sapphire worm [12] which was a bandwidth-limited scanning worm. These worms use scanning

algorithms that are limited by the throughput of the sender rather than by latency to the target.

Figure 13 shows the impact of Slammer on both traffic levels and routing updates. The top graph shows the inbound traffic levels (negative numbers) as well as the outbound levels (positive numbers). The bottom graph shows the number of prefixes seen at a BGP sensor on the same network as the traffic sensor. At the same time that the traffic sensor saw a large increase in the amount of inbound and outbound traffic, the routing sensor detected a large decrease in the number of announced prefixes, mostly due to the loss of many /24 routes. With a single sensor, it is difficult to determine how this drop in routing table size affected the measurement of traffic. While increasing the number of sensors might help, it's still important to recognize that these resource constraints can greatly influence the observed traffic.

4.5 Statistical variations

Another important cause of differences between sensors arises from the fact we are sampling traffic from a wider pool of address space. Ignoring the other factors mentioned above, and assuming a perfectly uniformly random scanning algorithm, there may still be variance between sensors due to sampling error. As we have shown in Section 3, hypothesis testing for homogeneity can be used to help account for these errors. Thus, it is important to be aware of statistical variation when comparing results. One way to gain statistical significance is to increase the number of observations by sensors.

In addition to pure sampling error, there are also non-uniformities in many of the random propagation algorithms. For example the source of entropy used in many random address generation methods is based on a clock seed or other poor source of randomness. In addition the random number generator itself may be biased. Thus, the distribution of addresses produced may not uniform.

5. CONCLUSION AND FUTURE WORK

In this paper we have explored issues associated with inferring global threat properties from a set of distributed blackhole sensors. In particular we demonstrated that significant differences exist in observations made at sensors of equal size that are deployed at different locations. These differences were demonstrated to persist over a month's worth of data even when local scanning preference was removed. Furthermore, differences appeared not only in aggregate, but also with specific ports and protocols. This variance was demonstrated within each sensor and proved to be statistically significant between sensors. Finally, we showed these differences existed even controlling for propagation strategy.

We then examined several properties of blackhole sensors in order to explain the difference in sensor observations. First, we examined the effects of policy deployed at the core and the edge. Next, we analyzed the role of local scanning preferences and the proximity of a blackhole sensor to target hosts. We then discussed routing availability including both availability of source addresses and the availability of blackhole addresses to the rest of the world. In addition, we showed that resource constraints in the host, network, or routers can produce differences. Finally, variance due to sampling and random address can affect observations.

These results demonstrate that observations made at a certain address block are potentially influenced by a number factors that must be accounted for when trying to generalize the result to other address spaces. Specifically, we hope this will assist researchers in understanding data captured using blackhole sensors in order to generalize their results. We are continuing our data collection which will allow us to determine if these differences persist over

longer time frames and with different threats. We are also expanding the IMS blackhole sensor network to include additional diverse address blocks with locations around the globe. Commitment has been received for 20 new deployments, spanning universities, corporations, and ISPs.

In addition to more deployments, we plan on continuing to evaluate blackhole sensor placement. We plan to study the efficacy of a given deployment using the following metrics.

- Coverage: the observed source addresses should be representative of the actual threat host population.
- Speed in detection: there is a high probability of capturing initial activity of a global event.
- Ability to identify scanning algorithm: sufficient breadth to determine the target selection function of an automated global threat.
- Resistance to filtering and resource bottlenecks: traffic observed by the sensor is unlikely to be blocked by filtering or constrained due to lack of resources.
- Scalability and cost: accomplishing the above while minimizing the number of deployed sensors and locations.

As a next step to sensor placement evaluation, we plan to develop a methodology for sensor placement that balance the above metrics. Finally, we hope to use this knowledge and our experience in IMS deployment to elucidate global trends in worm infection based on our sensor observations. We believe our work in identifying essential differences in sensor observations and providing an understanding of sensor properties is an important first step towards understanding distributed blackhole placement.

Acknowledgments

This work was supported by the Advanced Research and Development Activity (ARDA) under contract number NBCHC030104.

The authors would like to thank all the IMS participants for their help and suggestions. We would also like to thank Jose Nazario at Arbor Networks and Larry Blunk, Bert Rossi, and Manish Karir at Merit Network for their assistance and support through this paper. The IMS project had its roots in an earlier system created by Dug Song, Robert Stone, and G. Robert Malan.

6. REFERENCES

- [1] CERT. CERT Advisory CA-2001-26 Nimda Worm. <http://www.cert.org/advisories/CA-2001-26.html>, September 2001.
- [2] CERT. Code Red II: Another worm exploiting buffer overflow in IIS indexing service DLL. http://www.cert.org/incident_notes/IN-2001-09.html, August 2001.
- [3] CERT. CERT Advisory CA-2003-20 W32/Blaster worm. <http://www.cert.org/advisories/CA-2003-20.html>, August 2003.
- [4] CERT. CERT advisory CA-2003-20 W32/Blaster worm. <http://www.cert.org/advisories/CA-2003-20.html>, August 2003.
- [5] Cisco Systems. NetFlow services and applications, 1999.
- [6] Evan Cooke, Michael Bailey, David Watson, Farnam Jahanian, and Jose Nazario. The Internet motion sensor: A distributed global scoped Internet threat monitoring system. Technical Report CSE-TR-491-04, University of Michigan, Electrical Engineering and Computer Science, July 2004.
- [7] Dan Golding. Peering Evolution. Nanog Presentation, October 2002.
- [8] Craig Labovitz, Abha Ahuja, and Michael Bailey. Shining Light on Dark Address Space. http://www.arbornetworks.com/downloads/research38/dark_address_space.pdf, November 2001.
- [9] Zhuoqing Morley Mao, Ramesh Govindan, George Varghese, and Randy Katz. Route Flap Damping Exacerbates Internet Routing Convergence. In *Proc of ACM SIGCOMM*, 2002.
- [10] Microsoft Corporation. What you should know about the Sasser worm and its variants. <http://www.microsoft.com/security/incident/sasser.msp>, May 2004.
- [11] David Moore. Network telescopes: Observing small or distant security events. In *11th USENIX Security Symposium, Invited talk*, San Francisco, CA, August 5–9 2002.
- [12] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. Inside the slammer worm. *IEEE Security & Privacy*, 1(4):33–39, 2003.
- [13] David Moore, Colleen Shannon, Geoffrey M. Voelker, and Stefan Savage. Network Telescopes: Technical Report. Technical report, Cooperative Association for Internet Data Analysis - CAIDA, 2004.
- [14] David Moore, Geoffrey M. Voelker, and Stefan Savage. Inferring Internet denial-of-service activity. In *Proceedings of the Tenth USENIX Security Symposium*, pages 9–22, Washington, D.C., August 13–17 2001. USENIX.
- [15] Chris Morrow and Brian Gemberling. How to Allow your Customers to blackhole their own traffic. <http://www.secsup.org/CustomerBlackHole/>.
- [16] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of Internet Background Radiation. Available at <http://www.cs.princeton.edu/nsg/papers/telescope.pdf>.
- [17] Vern Paxson. Bro: A System for Detecting Network Intruders in Real-Time. *Computer Networks*, 31(23-24):2435–2463, 1999.
- [18] Niels Provos. Honeyd — A virtual honeypot daemon. In *10th DFN-CERT Workshop*, Hamburg, Germany, February 2003.
- [19] M. Roesch. Snort: Lightweight intrusion detection for networks. In *Proc. 13th Systems Administration Conference (LISA)*, pages 229–238, 1999.
- [20] SANS Institute. Internet storm center. <http://isc.incidents.org/>, June 2004.
- [21] Colleen Shannon and David Moore. The spread of the Witty worm. <http://www.caida.org/analysis/security/witty/>, June 2004.
- [22] Dug Song, Rob Malan, and Robert Stone. A snapshot of global Internet worm activity. Technical report, Arbor Networks, 2001.
- [23] Lance Spitzner. *Honeypots: Tracking Hackers*. Addison-Wesley, 2002.
- [24] Lance Spitzner et al. The honeynet project. <http://project.honeynet.org/>, June 2004.
- [25] Robert Stone. CenterTrack: An IP overlay network for tracking DoS floods. In *USENIX*, editor, *Proceedings of the 9th USENIX Security Symposium*, pages 199–212, Berkeley, CA, USA, August 14–17 2000. The USENIX Association.
- [26] Symantec Corp. Linux.Slapper.Worm. <http://securityresponse.symantec.com/avcenter/venc/data/linux.slapper.worm.html>.

- [27] Team CYMRU. The darknet project. <http://www.cymru.com/Darknet/index.html>, June 2004.
- [28] University of Oregon. RouteViews project. <http://www.routeviews.org/>.
- [29] Nicholas Weaver, Vern Paxson, Stuart Staniford, and Robert Cunningham. A taxonomy of computer worms. In *Proceedings of the 2003 ACM workshop on Rapid Malcode*, pages 11–18. ACM Press, 2003.
- [30] Vinod Yegneswaran, Paul Barford, and Somesh Jha. Global intrusion detection in the DOMINO overlay system. In *Proceedings of Network and Distributed System Security Symposium (NDSS '04)*, San Diego, CA, February 2004.
- [31] Vinod Yegneswaran, Paul Barford, and Dave Plonka. On the design and use of Internet sinks for network abuse monitoring. Technical Report 1497, University of Wisconsin, Computer Science Department, 2004.
- [32] Cliff C. Zou, Don Towsley, Weibo Gong, and Songlin Cai. *Routing Worm: A Fast, Selective Attack Worm based on IP Address Information*. UMass ECE Technical Report TR-03-CSE-06, November 2003.