

Internet Intrusions: Global Characteristics and Prevalence

Vinod Yegneswaran, Paul Barford and Johannes Ullrich
{vinod, pb}@cs.wisc.edu, jullrich@sans.org

ABSTRACT

Network intrusions have been a fact of life in the Internet for many years. However, as is the case with many other types of Internet-wide phenomena, gaining insight into the *global* characteristics of intrusions is challenging. In this paper we address this problem by systematically analyzing a set of firewall logs collected over four months from over 1600 different networks world wide. The first part of our study is a general analysis focused on the issues of distribution, categorization and prevalence of intrusions. Our data shows both a large quantity and wide variety of intrusion attempts on a daily basis. We also find that worms like CodeRed, Nimda and SQL Snake persist long after their original release. By projecting intrusion activity as seen in our data sets to the entire Internet we determine that there are typically on the order of 25B intrusion attempts per day and that there is an increasing trend over our measurement period. We further find that sources of intrusions are uniformly spread across the Autonomous System space. However, deeper investigation reveals that a very small collection of sources are responsible for a significant fraction of intrusion attempts in any given month and their on/off patterns exhibit cliques of correlated behavior. We show that the distribution of source IP addresses of the non-worm intrusions as a function of the number of attempts follows Zipf's law. We also find that at daily timescales, intrusion targets often depict significant spatial trends that blur patterns observed from individual "IP telescopes"; this underscores the necessity for a more global approach to intrusion detection. Finally, we investigate the benefits of shared information, and the potential for using this as a foundation for an automated, global intrusion detection framework that would identify and isolate intrusions with greater precision and robustness than systems with limited perspective.

Categories and Subject Descriptors

c.2.0 [Computer Communication Networks]: Security and

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.

SIGMETRICS'03, June 10–14, 2003, San Diego, California, USA.

Copyright 2003 ACM 1-58113-664-1/03/0006 ...\$5.00.

Protection; c.2.5 [Computer Communication Networks]: Network Monitoring; c.4 [Performance of Systems]: Measurement Techniques, Modeling Techniques, Performance Attributes

General Terms

Security, Measurement, Performance

Keywords

Network Security, Wide Area Measurement, Internet Performance and Monitoring

1. INTRODUCTION

Defending wide area networks from intrusion in the form of port scans and attacks poses a significant, on-going challenge for network operators. Using *backscatter analysis* to characterize Denial-of-Service (DoS) activity in the Internet, Moore *et al.* show that these intrusions are numerous and on the rise [17]. In 2001, two widely reported Internet worms (Code Red and Nimda) each infected hundreds of thousands of nodes in less than a day and required countless hours to eradicate from systems. Recent work by Staniford *et al.* suggests that the intrusion activity we have seen to date might only be the tip of a very large iceberg and that significant steps are necessary to counter these risks [23].

While methods and technology for securing networks against intrusions continue to evolve, the basic problems are extremely challenging for a number of reasons. First, the Blackhats who perpetrate intrusions continue to find ingenious ways to compromise remote hosts and frequently make their tools publicly available. Second, the size and complexity of the Internet, including end host operating systems, make it likely that there will continue to be vulnerabilities for a long time to come. Third, sharing of information on intrusion activity between networks is complicated by privacy issues, and while there are certainly anecdotal reports of specific port scanning methods and attacks, there is very little broad understanding of intrusion activity on a global basis. Because of these challenges, current best practices for Internet security rely heavily on word-of-mouth reports of new intrusions and security holes through entities such as CERT [4] and DSFIELD [24].

The focus of this paper is the development of a quantitative characterization of intrusion activity in the global Internet. To our knowledge, this is the first study to broadly address this problem using intrusion logs from firewalls at sites distributed

throughout the Internet. Specifically, our data was collected in over 1600 networks world wide over a 4 month period by DSHIELD.ORG. Entries in these logs consist of packets rejected by firewalls and portscan logs recorded by Network Intrusion Detections Systems (NIDS -primarily Snort [19]). This data set provides us with a unique perspective on global intrusion activity.

We investigated a range of fundamental features of intrusion activity by evaluating our data along a number of dimensions. Specifically, we assess the daily volume of intrusion attempts, the sources and destinations of intrusion attempts, and specific types of intrusion attempts.

Our results show the following:

- **Volume:** Daily intrusion attempts take place on an massive scale - as many as 3 million scans in our logs on a single day - and in a bursty fashion.
- **Distribution:** Sources and destinations of intrusions from an Autonomous System (AS) perspective are nearly uniformly distributed around the Internet. Furthermore, the distribution of the number of scans per source IP versus their overall rank follows a power law (Zipf's Law [27]).
- **Types:** Worm activity (intrusion attempts on port 80) varies between approximately 20% and 60% of all intrusion attempts and worm signatures from Code Red and Nimda remain prominent today (over a year since their original release). Non-port 80 intrusion attempts make up a surprisingly large percentage of the daily volume. Decomposition of these indicates that while common scanning methods (*eg.* vertical and horizontal) are indeed prevalent, other methods such as coordinated and stealthy scans are also widely used.

Our next step was to use these results to project intrusion activity to the entire Internet. We do this using multiple perspectives so that, in addition to getting a perspective on global intrusion activity, we can assess the extent to which the intrusion activity seen by a single network is representative of intrusion throughout the Internet. Our projection method consists of using all of our data, only data from /16 networks and only data from /24 networks. Resulting projections have roughly an order of magnitude difference respectively (decreasing). The projection using the entire data set shows Internet wide intrusion attempts to be on the order of 25B per day over all ports. A simple least squares fit through the data shows an increasing trend for all projections.

Finally, we investigate the utility of sharing information between networks as a basis for a distributed intrusion detection infrastructure. For example, we are interested in assessing the number of logs required to establish particular source IPs as "worst offenders". We use an information theoretic approach outlined in [2] to assess the *marginal utility* of intrusion information collected from multiple sites. In addition to worst offenders, we also assess the utility of additional logs in understanding intrusion prevalence with respect to port targets. We find that small random collections of intrusion logs are not sufficient for gaining a consistent view of worst offenders or port targets. This result combined with the Zipf result for worst offenders indicate that while there is certainly potential




Figure 1: Global Distribution of Provider Autonomous Systems. The dark and light squares respectively correspond to ASes that are represented and not represented. The grey lines or shading indicates direct connectivity between ASes

for sharing information between networks as a means for improving security, a thoughtful approach to log aggregation is likely to be required.

This paper is organized as follows. We describe studies related to this work in Section 2. In Section 3 we describe the details of our data and our methods of analysis and Section 4 provides background material on worms and scans. We present results of our characterizations of fundamental features of intrusions in Section 5 and extend those results to the entire Internet in Section 6. Our discussion of the use of shared information in a distributed intrusion detection system is presented in Section 7. We summarize and conclude in Section 8.

2. RELATED WORK

The work by Moore *et al.* is motivated by the question, "how prevalent are denial-of-service attacks in the Internet today?" [17]. Our work is similar in spirit although we address the general question of intrusions and are not specifically focused on DoS activity. Staniford *et al.* report on recent worm activity (Code Red, Nimda) in [23] and project the possibilities of much more serious worm threats in the future. Cowie *et al.* present a different perspective on the same work by examining hour long periods of "widespread instabilities" in global BGP system in July and September of 2001 [5]. They describe the idea of "worm induced traffic diversity" that is unlike other normal traffic experienced by routers and is the primary cause of the BGP instabilities.

Our work has implications in development and configuration of network intrusion detection systems. Many such systems have been developed and deployed (*eg.* [18, 19]). The standard approach for recognizing an intrusion is to create

Table 1: Sample log entries from DSHIELD portscan logs

Date	Time	Sub. Hash	No: Scans	Src IP	Src Port	Target IP	Target Port	TCP Flags
2002-03-19	18:35:18	provider2323	3	211.10.7.73	1227	10.3.23.12	21	S
2002-03-19	18:35:19	provider2323	16	211.10.7.73	1327	10.3.23.12	53	SF
2002-03-19	18:35:20	provider2323	1	211.10.7.73	1231	10.3.23.12	111	F
2002-03-19	18:35:21	provider2323	1	211.10.7.73	1331	10.3.23.12	22	SA

a rule-based description which is then used to configure the NIDS. However, the task of accurately identifying new types of intrusions remains quite challenging.

Other well known studies related to ours are the DARPA Intrusion Data Sets from 1998 and 1999 [13]. These data sets are typically used as training sets to test NIDS. They suffer from being both old and synthetically generated, thus their relevance to intrusion attempts of today should be questioned. Our data (or a set like ours) could serve as a benchmark for creating or validating training sets in the future.

The HoneyNet Infrastructure is a unique “network” dedicated to understanding the tools and activities of the Blackhat community [10]. Their network consists of machines deployed in the Internet with NIDS that are virtually unused. Their lack of network use significantly reduces the possibility of false positives in the data they generate - every packet received is considered suspect. The project publishes weekly reports of recorded attack traffic and successful exploits.

Finally, there have been a number of recent papers on router-based techniques for IP traceback [20, 21]. These techniques all face considerable challenges in gaining operational deployment however they offer interesting possibilities for identifying sources of intrusions.

3. INTRUSION DATA

We use a set of firewall logs of portscans collected over a 4 month period from over 1600 firewall administrators distributed throughout the globe as the basis for our study. The logs provide a condensed summary (lowest common denominator) of portscan activity obtained from various firewall/IDS platforms. Some of the platforms supported include BlackIce Defender, CISCO PIX Firewalls, ZoneAlarm, Linux IPchains, Portsentry and Snort. This approach significantly increases the coverage and reduces reliance on individual IDS’s interpretation of events. Table 1 illustrates the format of a typical lowest common denominator log entry. The date and time fields are standardized to GMT and the provider hash allows for aggregation of destination IP addresses that belong to the same administrative network.

Table 2 provides a high level summary of the data that was used in our study. The dataset was obtained from DSHIELD.-ORG – a research effort funded by SANS Institute as part of its Internet Storm Center. DSHIELD’s objectives include detection and analysis of new worms and vulnerabilities, notification to ISPs of exploited systems, publishing blacklists of worst offenders and feedback to submitters to improve firewall configuration. The data is comprised of logs submitted by a diverse set of networks and includes 5 full Class B networks, over 45 full Class C sized networks and many smaller

Table 2: Monthly summary of studied DSHIELD logs

Month	Number of Scans	Number of Dest IPs
Aug. 2001	30 million	260,726
May. 2002	48 million	375,323
June. 2002	61 million	382,224
July. 2002	68 million	402,050

subnetworks. One of the highlights of this data source was its utility and contribution in the detection and early analysis of CodeRed. Figure 1 is a Skitter-based AS level graph that shows the global distribution of the providers that submit to DSHIELD. A Skitter graph provides a unique way of visualizing Autonomous Systems (ASes) based on their connectivity and geography without compromising provider identities. There are three distinct regions in the graph and they correspond to autonomous systems based in Europe, North America and Asia Pacific in clockwise order. The grey lines or shading indicates which ASes are directly connected to one another. The dark and light squares correspond to ASes with representative participants and non-participant ASes respectively. The ASes closer to the center constitute the larger backbone providers and telecommunication companies that have maximum connectivity while the stub ASes populate the periphery. [11].

The lowest common denominator approach by DSHIELD provides us with a unique, globally diverse and stable data source. The simplicity and generality of this approach also makes analysis straightforward. There are however some pitfalls that need to be considered. The logs do not provide information about packet headers, or what happens during active connections. There is also a certain degree of vulnerability to flooding by malicious users and by misconfigured firewalls. For example local broadcast traffic and network games like Half-life can result in many false positives. These instances were filtered out from the dataset before analysis. Finally NIDS systems may not be able to capture all packets during a Denial of Service attack. However, spikes before and after an attack should get recorded.

Finally, we need to consider portscans from spoofed sources. Normally there is little benefit to an individual spoofed portscan as the source does not get a response. A few specialized situations where a spoofed portscan by itself could prove beneficial to a portscanner, like *spoof-bounce*, have been documented [7]. However, the best known motivation for spoofing portscans is to create spurious background noise to hide the real sources. By considering and correlating reports from mul-

multiple sources, the effect of the spoofed sources are marginalized and real sources rise to the top. This works because in reality, spoofed portscans form only a small fraction of all recorded portscans. Thus, analyzing source IPs of portscans can provide valuable insight into the geographic distribution of malicious host subnets or *stepping stones* [26].

4. BACKGROUND

4.1 The Worms

In this section we provide background information on the major Internet worms released over the last two years. We first describe the major port 80 worms CodeRed I/II and Nimda. This is important because port 80 scans still form the single most dominant group of scans accounting for nearly 20% - 60% of all scans in any given day. Most port 80 scans observed in the 3 month period between May 2002-June 2002 can be attributed to either CodeRed or Nimda. The release date for Nimda was Sep 18, 2001 and so the port 80 scans in the August 2001 dataset are exclusively CodeRed. We also describe the SQL-Snake—a worm which affects Microsoft SQL Servers.

4.1.1 CodeRed I

CodeRed I exploited a well known Windows Internet Information Server (IIS) buffer overflow vulnerability [12] and was released on July 12, 2001. The worm was so named because it defaced some web pages with the words “hacked by Chinese”. The worm operates in two distinct phases. In its first phase, the worm uses a random IP generator to search for vulnerable targets. In the second phase (20-28th of every month), the worm stops propagation and launches Denial of Service attacks against the <http://www1.whitehouse.gov> website [3].

4.1.2 CodeRed II

Despite superficial similarities, CodeRed II is a completely different worm that uses the same IIS vulnerability. The name “CodeRed II” is derived from a string in the worm’s source code. Unlike CodeRed I, this worm is not memory resident and hence a reboot does not disinfect the machine. CodeRed II’s propagation mechanism generates a random IP address and a mask whose size determines the similarity between the infected IP and probed IP. About 1/2 of the time CodeRed II probes an IP in the same /24, and about 3/8 of the time CodeRed II probes an IP in the same /16 and a random IP remainder of the time. Finally CodeRed II installs a root level backdoor that allows any other code to be remotely executed [3].

4.1.3 Nimda

Nimda stands for *admin* spelled backwards. The algorithm for target detection is not well known, but seems to follow these rough probabilities: 50% an IP address with same first 2 octets, 25% an IP address with matching first octet and 25% a completely random IP address [4].

Figure 2 shows the number of distinct sources scanning port 80 during the 3 months in 2002 from our dataset. The graph shows the predictable nature of the port 80 scans and also very strong day of the month characteristics. The sharp drop around the 19th of each month confirms that CodeRed I is still very much alive. CodeRed I is supposed to self destruct

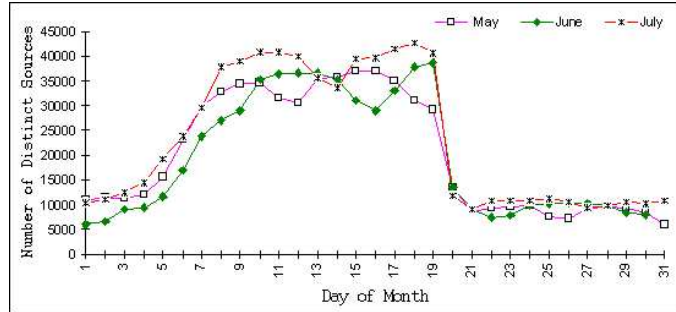


Figure 2: Daily number of distinct sources of port 80 scans, May 2002-July 2002

at the end of each month. The hypothesis is that however, there are enough machines with a sufficient clock offset that the vulnerable machines get reinfected producing this slow start like growth between days 1-7.

4.1.4 SQL-Snake

The SQL-Snake was detected on May 20th 2002¹. This worm scans for open MS-SQL 7 Servers which run on port 1433 by default and exploits machines that have the default SA (Admin) account without an associated password. The worm scans for IPs of the form A.B.C.D randomly on the following IP ranges where: A = random² number not equal to 10,127,172 or 192, B = 0-255, C = 1-255 and D = 1-254. The primary function of the worm is to email passwords and related system information to ixltd@postone.com [1, 14, 25].

4.2 Scan Types

We broadly categorize scans into four well known types [22].

1. **Vertical Scan** - is defined as a sequential or random scan of multiple (more than 5) ports of a single IP address from the same source during a one hour period. These are usually an attempt to survey which of several well known vulnerabilities applies to this host and are also known as *strobe* scans, based on one of the original script-kiddie tools.
2. **Horizontal Scan** - is a scan from a single source of several machines (5 or more) in a subnet aimed at the same target port, ie. the same vulnerability. In this case the attacker is searching for any machine that is running specific service and does not care about any single machine in particular. The attacker could be just recruiting peers for launching larger distributed attacks.
3. **Coordinated Scans** - are scans from multiple sources (5 or more) aimed at a particular port of destinations in the same /24 subnet within a one hour window. These

¹Not to be confused with the SQL Slammer worm released in Jan 2003.

²Aside from avoiding certain class A-s, the first byte is also weighted with some networks being scanned more frequently than others.




Figure 3: Source AS distribution of Nimda/CodeRed(left), SQL-Snake(middle) and Non-Worm(right) scans.

are also called Distributed Scans [23]. These scans usually come from the more aggressive/active sources that comprise several collaborative peers working in tandem.

4. **Stealth Scans** - are horizontal or vertical scans initiated with a very low frequency to avoid detection. The key parameters in this definition include the maximum threshold (1 hour) and the minimum threshold for the average interscan distance. An average interscan distance below the minimum threshold indicates that the scan was not stealthy, *ie.* not intended to evade NIDS systems. Two successive scans from the same source that are separated by more than the maximum interscan distance are considered to be unrelated or parts of different scanning episodes.

5. INTRUSION CHARACTERISTICS

In this section, we first provide a high level summary of intrusion distribution in terms of the destination ports and the attack sources. We illustrate instances where the top sources depict striking correlated behavior. We then project the observed scan rates over the entire IP space and try to identify temporal trends. We identify the predictable and persistent behavior of the port 80 sources. Finally we investigate prevalence of spatial trends in the scanning behavior.

5.1 Port Distribution

Monitoring the destination port of intrusion attempts in the Internet proves to be an effective method for detecting exploits for new vulnerabilities and dissemination of new worms. The cases of CodeRed I/II, Nimda and recent Opaserv (port 137) are instances where the heightened scanning rates were observed for several days before they were identified.

The daily scanning volume of the top 10 destination categories for the 3 months from May 2002 - July 2002 is shown in Figure 4. Obviously the most significant component of the graphs are the port 80 scans from Nimda and CodeRedI/II. The Linux Slapper worm (released later this summer) is not part of this data. The visible spike in the port 1433 scans around the 4th week of May is directly attributable to the release of SQL-Snake worm [25]. The P2P entry comprises of scans from Gnutella, Kazaa and EDonkey peers and the port 137




Figure 4: Scan rates for top 10 destination port categories between May-July, 2002.

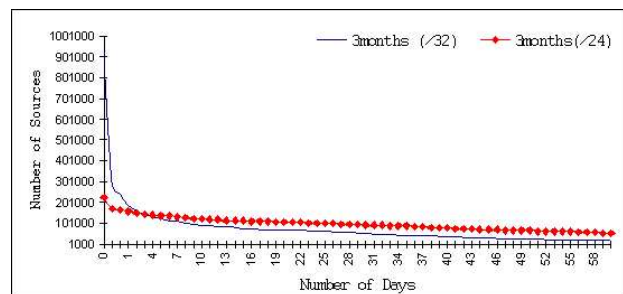


Figure 5: Persistence of port 80 scans between May-July, 2002 in /32 and /24 aggregates.

scans are dominated by Windows NetBios misconfiguration. Of the non-worm scans, ssh (22), ftp (21), rpc (111), dns(53) and icmp(0) tend to dominate. The scans from the Subseven trojan (leaves worm) and scans to port 3128 (misconfigured proxy servers used to redirect/hide surfing behavior) also figure in the top ports for the 3 months [24].

5.2 Source Distribution

One of the challenges in a study of the sources of the intrusions is to account for the disparities in the scanning behavior of self propagating and non-worm traffic. To simplify we classify scans into three categories: port 80, 1433 and the remaining non worm traffic. We filtered out traffic from Sub-

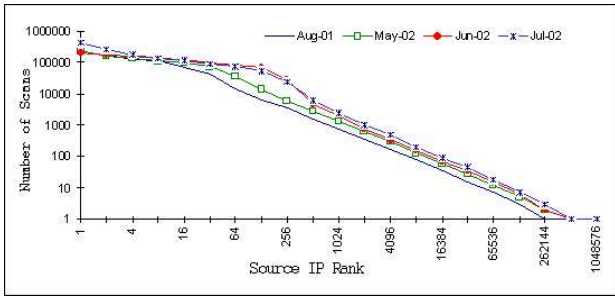


Figure 6: Log-log plot of source IP rank versus number of monthly scans for each of the four months.




Figure 7: Daily scan rate of top 100 non-worm sources in May, 2002 as compared with all sources.




Figure 8: Cluster of 8 out of top 20 sources in August, 2001 with very similar on-off behavior.




Figure 9: Cluster of 8 out of top 20 sources in June, 2002 with very similar on-off behavior.




Figure 10: Cluster of 4 out of top 20 sources in May, 2002 with similar on-off patterns.

seven and noise due to Windows NetBios and the peer-to-peer scans. All future references to non-worm scans use this filtered dataset. Figure 3 shows the global distribution of autonomous system level sources of port 80, port 1433 and the non-worm scans for the month of June 2002. A snapshot from the Oregon route views table was used to convert source IP addresses to AS numbers [15]. These graphs serve to illustrate the global reach and uniformity of attack sources (also victims in the former two cases) in each category. It is apparent that the distribution of the SQL Snake victims even at its peak, was less dense in comparison to port 80 or other non-worm traffic. Nevertheless it was responsible for a significant portion of all recorded scans.

5.2.1 Persistence of Worm Activity

Another interesting aspect of worm behavior is the persistence of attacks. In particular, we wanted to understand how long a subnet remained infected. We compute the duration of port 80 scans during the 3 month period based on aggregates of /24 and /32. The motivation behind looking for /24 matches comes from CodeRed II's and Nimda's affinity towards local targets. The half lives of the 3 categories are 18 days for /24 matches and 6 hours for /32 match. The /24 matches are biased by individual sources that send only a single portscan and do not really indicate quick disinfection. The graph in Figure 5 reveals an almost linear relationship between duration and number of infected /24 subnets with a very small slope. This indicates that while a majority of subnets become Nimda/CodeRed free in about 18 days, significant number of subnets continue to be infected for extended periods.

5.3 Top Sources

Isolating and understanding behavior of worst offenders (in terms of source IP) is crucial to defending networks. We focus primarily on the non-worm intrusions in this section. The self-propagating nature and local affinity of worm traffic makes a similar analysis less meaningful and more complicated. Figure 6 plots the rank of each source IP against its number of scans (*ie.* popularity) on a log-log scale. The decreasing linear slope of this plot indicates a power-law distribution which leads us to conclude that "worst offender" IPs follow Zipf's Law [8, 27]. Many other phenomenon in the Internet have a similar characteristic. This indicates that very few sources are in fact responsible for generating a significant fraction of all non-worm scans that are observed. Figure 7 plots the daily

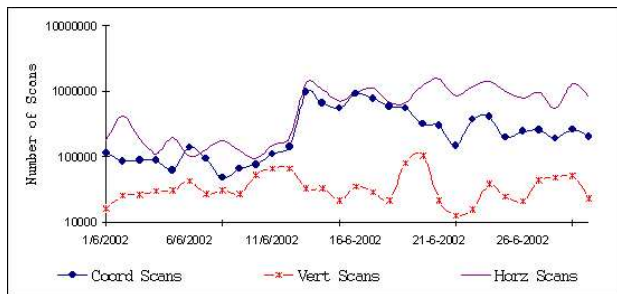


Figure 11: Distribution of coordinated, horizontal and vertical scans for the month of June, 2002.

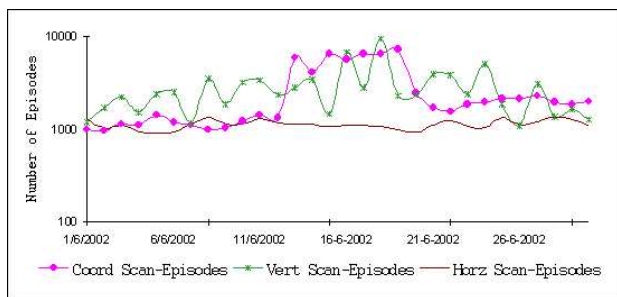


Figure 12: Distribution of coordinated, horizontal and vertical scan episodes for the month of June, 2002.

scan volume from the top 100 sources (out of 261K) along with the total scans and further supports this results. The graph shows that the top sources which are responsible for roughly half of all non-worm scans, also account for most of the variability that is exhibited.

Another way to consider worst offender behavior is to focus on the scans emanating from the top 20 sources for each month. Observing the daily on/off patterns of these top 20 sources during the course of each month reveals clusters of correlated behavior. Such clusters were observed prominently in the top 20 sources of each of the four months under consideration. Figures 8 and 9 show clusters from the top 20 sources from August 2001 and June 2002. These sources are “on” during the same days of the month with similar levels of activity and bear little locality in IP space. Figure 10 captures a set of four IPs that display similar staggering behavior but on different days. The similarity in the actual number of scans during the on periods is striking. This points to an identical attack or an attack using the same tool launched from disparate sources on different days.

These results lead us to conclude that such attacks are in fact fairly common, and that blacklisting worst offenders would be an effective mechanism of defending against non-port 80 intrusions. We further conclude that instances of collaborative clusters can be effectively isolated and should be investigated with greater vigor.

5.3.1 Identification of Scan Types

Figures 11 and 12 show the daily distribution of the three scan types during the month of June 2002. This indicates that horizontal scans account for 60-70% of all non-worm scans.

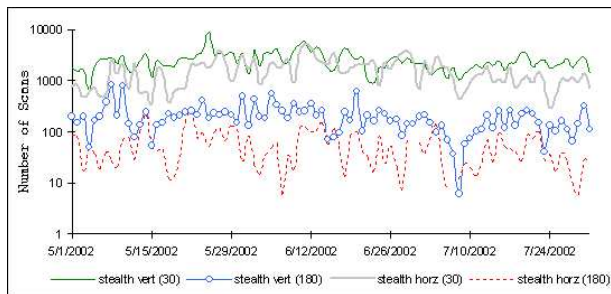


Figure 13: Distribution of Stealth scans in June, 2002 with minimum interscan thresholds of 30 and 180 seconds.

Another surprising revelation is that a large proportion of the daily scans are coordinated³ or come from distributed sources. The coordinated scan rate also seems much more tightly tied to the number of coordinated scan episodes with the ratio of 100 scans per episode. The most common ports scanned included port 111 (RPC), port 53 (DNS) and scans for alternate web-server ports like 8000 and 8080. Although horizontal scans are more common than vertical scans, there are fewer horizontal scan episodes than vertical scan episodes. Figure 13 shows daily distribution of observed stealth scans with minimum thresholds of 30 and 180 seconds. Stealthy scans are not uncommon, however only makeup a small percentage of all vertical and horizontal scans. Vertical scans seem to be much more likely to exhibit stealthy behavior than horizontal scans.

5.4 Network Telescopes

Network Telescopes are defined as large chunks of unused but globally routable IP space. Passive monitoring of telescopes serve as a useful mechanism for measuring and understanding Internet attack behavior [16], especially worms like CodeRed and Nimda. We examine their potential for characterizing the patterns of non-worm scans in the larger Internet. Figures 14 and 15 show the observed daily scan rates from an entire Class B with respect to the overall scan rates for both non-worm and port 80 scans. They reconfirm that for modeling port 80 scans even a single Class B telescope works reasonably well. However the non-worm traffic modeled from the same telescope exhibits significant variability unlike the global rates. This suggests that non-worm traffic has inherent spatial components that may only be able to be captured effectively only from a globally distributed set of telescopes.

6. GLOBAL PREVALENCE

6.1 Projection

Scanning patterns seen in the Internet have been highly dynamic especially over the last couple of years with the emergence of novel and more sophisticated worms. One of the intriguing questions that we would like to answer is how the volume of scans have changed over the last year? To address this question, we project the daily scans observed in our dataset to the larger Internet. We do this by simply taking the “average scans per IP” for our set of destination IPs and then multi-

³as per our definition in section 4.2.

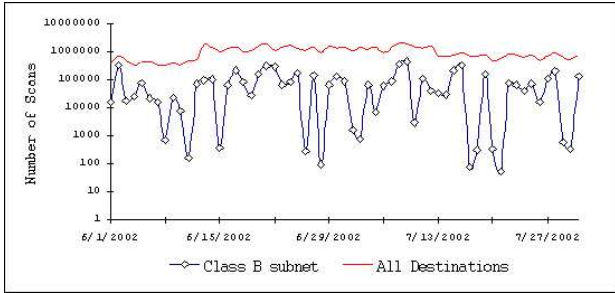


Figure 14: Daily scan rate of non-worm scans observed from a Class B telescope with respect to overall rate, June 2002-August 2002.

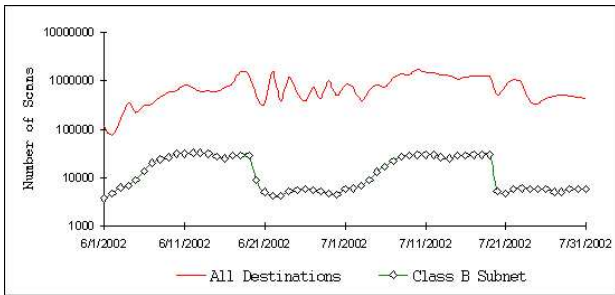


Figure 15: Daily scan rate of port 80 scans observed from a Class B telescope with respect to overall rate, June 2002-August 2002.

plying that by the number of IPs in the entire IP space⁴. We assume uniformity, but do not test for it. That is, we assume that since our set of provider networks are reasonably well distributed (both geographically and over the IP space), our perspective reflects what is seen over the larger Internet. This projection indicates daily scan rates as high as 25B/day.

We perform similar projections using /24 and /16 aggregates and try to discern trends via linear fits. The /16 and /24 aggregates provide very conservative estimates of the observed scan rates due to sparse network representations. One motivation behind them is to account for possible IPs that may not receive any scans during an entire day. Figures 16 and 17 show the results for port 80 and non-worm scans. The port 80 scan rates show a decreasing trend which is due to high levels of CodeRed incidences in Aug 2001. The rates between May and August 2002 (not shown explicitly in the figure) are relatively steady with a very small upward slope. The non-worm scan rates however show an increasing trend. The average daily number of scans over the IP space jumps from 6.5 B scans to 8.2 B, an increase of over 25%.

7. IMPLICATIONS OF SHARED INFORMATION

A number of recent papers and proposals address the con-

⁴ Alternately we could have obtained the portion of routable IP space from a BGP table. While the gross volume of scans might be slightly reduced, we expect the trends to be preserved.




Figure 16: Projection of port 80 scans Aug 2001-July 2002, over /32, /24 and /16 aggregates




Figure 17: Projection of non-worm scans Aug 2001-July 2002, over /32, /24 and /16 aggregates

cept of developing an infrastructure that would pool resources in order to more rapidly and more effectively respond to attacks and intrusions [6, 10, 22]. There are many issues involved in the creation of such an infrastructure, not the least of which is understanding its potential for success. Given the fact that there is likely to be little synchronization of timestamps between daily firewall logs in our data set, we did not attempt to evaluate how rapidly attacks and intrusions could be identified if data were collected in central repository in real time from sites distributed across the Internet – we leave that for future work. Our data does, however, lend itself to evaluating other aspects of developing composite views of intrusion activity and we explore two examples of these in this section.

7.1 An Information Theoretic Approach

Exploring the extent of refinement of perspective provided by additional data is a standard notion in information theory. *Relative entropy* is a measure of the distributional similarity between two variables [9]. This measure is commonly estimated using the Kullback-Leibler distance metric which was extended in [2] to measure the marginal utility of adding additional experimental results to an aggregate data set of network topology measurements. Our interest is in understanding how the addition of intrusion logs to an aggregate data set improves the resolution of identifying “worst offenders” and the prevalence of scans of particular ports. The marginal utility metric will quantitatively express the information gained by the additional logs.

A framework for marginal utility evaluation is presented in




Figure 18: Utility of additional subnets for detecting worst offenders

[2]. The framework considers a set of n identical (*ie.* aimed at discovering a common property) experiments S^1, S^2, \dots, S^n . In our case, these experiments would be intrusion logs submitted by distinct sites. Marginal utility is defined as the reduction in uncertainty resulting from the next experiment added to the aggregate set. Two alternatives for calculating marginal utility are presented: one which considers the reduction in uncertainty in an *online* manner and the other in an *offline* manner. The essential difference between the two is that the *offline* metric considers marginal utility from the perspective of the set of all n experiments. We select the offline metric for analysis since, as stated above, we are not considering issues related to the order in which individual logs are submitted.

The formal definition of marginal utility of an experiment S^n is defined to be $U^m(S^n)$ and is given by the following equation:

$$U^m(S^n) = \sum_{s_i} Pr(s_i^m) \log\left(\frac{Pr(s_i^m)}{Pr(s_i^n)}\right) \quad (1)$$

where i ranges over all possible outcomes and $Pr(s_i^j)$ is the probability associated with outcome s_i after the conclusion of experiments S^1, S^2, \dots, S^j , and m is the total number of experiments conducted.

7.2 Identification of Worst Offenders and Prevalence of Target Ports

Our experiments to evaluate the marginal utility of intrusion log sharing focuses on two issues: the identification of worst offenders and the identification of ports (non-port 80) that are most frequently scanned. Our intent is to examine this issue in a general sense. To that end we conducted the experiments by selecting a single day at random from our data set. We then select logs from 100 /16's at random and 100 /24's at random to determine how many logs are required to get a consistent perspective on offenders and intrusion targets. Our analysis of network telescopes in Section 6 gives us the intuition that the aggregation of logs from a non-trivial number of sources, will be required to gain a representative perspective on these two issues.

The graph in Figure 18 shows the marginal utility of additional logs for identifying worst offenders. For this analysis (and the analysis of port targets), we ordered the logs by the number of scan entries. If the distribution of information in the logs is relatively stable, then this ordering should provide a sequence of marginal utility measures that follow a decreasing




Figure 19: Utility of additional subnets for detecting top target ports

trend. It can be seen in the figure for the larger /16 networks, that this is indeed the case. Beyond the aggregation of about 35 intrusion logs, almost no additional information is added. The plot for the /24 networks tells a different story. Non-negligible marginal utility metrics exist over a great deal of the graph indicating that aggregations more than 100 /24's may be necessary to get a clear view of worst offender distributions.

The results for the marginal utility of identifying target ports is somewhat different. As can be seen in Figure 19, there is a good deal of variability in marginal utility metrics for both /16's and /24's for log aggregations under 30. However, beyond 40, both exhibit fairly small marginal utility metrics indicating that stable perspectives on port targets may be achieved with relatively small numbers of logs.

8. SUMMARY AND CONCLUSION

In this paper we present a broad, empirical analysis of Internet intrusion activity using a large set of NIDS and firewall logs collected over a four month period. We found daily intrusion activity as seen in our data to be highly variable ranging from between about 1M to 3M scans per day. Examination of source IPs for these scans shows that they are widely disbursed across the autonomous system space, and that the distribution of attempts per source IP follows a power law. Our breakdown of scan types shows not only the predictably large amount of worm activity, but also a large amount of scanning directed toward ports other than 80. We find that while 60-70% of all non-worm scans are horizontal scans, the daily number of horizontal scan episodes is typically lower than vertical scan episodes.

To gain insight into the global nature of intrusions we used our data to project activity across the Internet. We used three different methods in this regard; considering first our entire set of data, then just /16 networks then just /24 networks. We find total intrusion activity to be as high as 25B per day and that non-port 80 scans increased by approximately 25% over our measurement period.

We also presented a high level information theoretic evaluation of the potential of using data shared between networks as a foundation for a distributed intrusion detection infrastructure. Our analysis indicates that small collections of logs from smaller networks may not be sufficient to identify either worst offenders or most popular port targets for attacks.

Our analysis has a number of implications. First, intrusion activity takes place on a massive scale throughout the Internet

and it is on the rise –network administrators should beware. Second, the worst offenders typically depict coordinated behavior and are responsible for significant fraction of all scanning activity. This is a strong argument for developing better blacklists and employing appropriate ingress filtering. Third, while current firewall and NIDS systems provide useful clues about attack patterns their views are limited by their vantage points. There is significant benefit to be achieved by collaboration, however this benefit is sensitive to the size of the peering group and its diversity.

Next steps in this work will be to attempt to refine the means by which intrusion data used in a distributed coordinated infrastructure. We are also interested in how effectively intrusion data collected in real time can positively identify new intrusion exploits in the Internet.

Acknowledgements

The authors would like to acknowledge KC Claffy and Brad Huffaker of CAIDA for providing access to the Otter graphing tool. We would like to thank Jim Gast for the use of scripts used in converting IP to AS numbers. We are also grateful for helpful comments from Steve Bellovin, Vern Paxson, Stefan Savage and the anonymous reviewers. Finally we would like to thank all the DSHIELD providers whose contributions made this analysis possible.

9. REFERENCES

- [1] George Bakos. SQLsnake code analysis. <http://www.incidents.org/diary/diary.php?id=157>, 2002.
- [2] Paul Barford, Azer Bestavros, John Byers, and Mark Crovella. On the marginal utility of network topology measurements. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, San Francisco, CA, November 2001.
- [3] CAIDA. CodeRed Worms a Global Threat. <http://www.caida.org/analysis/security/code-red/>, 2001.
- [4] CERT Coordination Center. <http://www.cert.org>, 2001.
- [5] James Cowie, Andy T. Ogielski, BJ Premore, and Yougu Yuan. Global Routing Instabilities Triggered by CodeRed II and Nimda Worm Attacks. http://www.renesis.com/projects/bgp_instability, 2001.
- [6] Frederic Cuppens and Alexandre Miege. Alert correlation in a cooperative intrusion detection framework. In *Proceedings of IEEE Symposium on Security and Privacy*, 2002.
- [7] Kevin Van Dixon. Spoof bounce. <http://rr.sans.org/intrusion/spoof.php>, 2001.
- [8] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *Proceedings of ACM SIGCOMM*, 1999.
- [9] Robert Gray. *Entropy and Information Theory*. Springer-Verlag, 1990.
- [10] HoneyNet Project. Know Your Enemy: Honeynets. <http://project.honeynet.org>, 2001.
- [11] Brad Huffaker, Andre Broido, Kim Claffy, Marina Fomenkov, Sean McCreary, David Moore, and Oliver Jakubiec. Visualizing internet topology at a macroscopic scale. http://www.caida.org/analysis/topology/as_core_network/about.xml/, 2001.
- [12] Eeye Security Inc. Microsoft IIS Buffer Overflow Advisory. <http://www.eeye.com/html/Research/Advisories/AD20010618.html>, 2001.
- [13] Richard Lippmann, David J. Fried, Isaac Graf, Joshua W. Haines, Kristopher R. Kendall, David McClung, Dan Weber, Seth E. Webster, Dan Wyschogrod, Robert K. Cunningham, and Marc A. Zissman. Evaluating Intrusion Detection systems: 1998 DARPA Off-line Intrusion Detection Evaluation. In *Proceedings of IEEE Security Symposium*, 1998.
- [14] McAfee. Virus alert. http://vil.nai.com/vil/content/v_9949.htm, 2002.
- [15] David Meyer. University of Oregon Route Views Project. <http://antc.uoregon.edu/route-views/>, 2002.
- [16] David Moore. Network Telescopes: Observing Small or Distant Security Events. http://www.caida.org/outreach/presentations/2002/usenix_sec/, 2002.
- [17] David Moore, Goeffrey Voelker, and Stefan Savage. Inferring Internet Denial-of-Service Activity. In *Proceedings of the 10th USENIX Security Symposium*, 2001.
- [18] Vern Paxson. BRO: A System for Detecting Network Intruders in Real Time. In *Proceedings of the 7th USENIX Security Symposium*, 1998.
- [19] Marty Roesch. The SNORT Network Intrusion Detection System. <http://www.snort.org>, 2002.
- [20] Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson. Practical Network Support for IP Traceback. In *Proceedings of ACM SIGCOMM*, 2000.
- [21] Alex Snoeren, Craig Partridge, Luis Sanchez, Christine Jones, Fabrice Tchakountio, and Stephen Kent. Hash Based IP Traceback. In *Proceedings of ACM SIGCOMM*, 2001.
- [22] Stuart Staniford, James Hoagland, and Joseph McAlerney. Practical Automated Detection of Stealthy Portscans. In *Journal of Computer Security*, 2002.
- [23] Stuart Staniford, Vern Paxson, and Nicholas Weaver. How to Own the Internet in Your Spare Time. In *Proceedings of the 11th USENIX Security Symposium*, 2002.
- [24] Johannes Ullrich. DSHIELD. <http://www.dshield.org>, 2000.
- [25] Johannes Ullrich. MSSQL worm (sqlsnake) on the rise. <http://www.incidents.org/diary/diary.php?id=156>, 2002.
- [26] Yin Zhang and Vern Paxson. Detecting Stepping Stones. In *Proceedings of the 9th USENIX Security Symposium*, 2000.
- [27] G. Zipf. *Human Behavior and the Principle of Least-Effort*. Addison-Wesley, Cambridge, MA, 1949.