

Computer Sciences Department

On the Design and Use of Internet Sinks for Network Abuse Monitoring

Vinod Yegneswaran
Paul Barford
Dave Plonka

Technical Report #1497

February 2004

UNIVERSITY OF
WISCONSIN
MADISON

On the Design and Use of Internet Sinks for Network Abuse Monitoring

Vinod Yegneswaran, Paul Barford and Dave Plonka*

Abstract

Network abuse monitoring (*eg.*, for intrusions and denial of service attacks) is an important component in security architecture. Monitoring *unused* IP addresses offers opportunities to significantly improve perspective on abuse activity without many of the problems associated with typical network intrusion detection and firewall systems. In this paper, we describe a scalable architecture for an IP traffic monitoring system called an Internet Sink (iSink). The objective of this system is to measure abuse activity on unused or “dark” IP addresses in an efficient and scalable fashion. A distinguishing feature of an iSink in contrast to traditional intrusion detection systems or firewalls, is that it includes a stateless *active* component that generates response packets to incoming traffic. This gives the iSink an important advantage in discriminating between different types of attacks (through examination of the response payloads). In the second part of the paper, we report a case study of live deployment and performance results of our iSink implementation in controlled laboratory experiments. The case study demonstrates the utility of iSink by revealing interesting network phenomena such as periodic probing and SMTP hot-spots. The laboratory results demonstrate the efficiency and scalability of our implementation.

1 Introduction

Network abuse in the form of intrusions by port scanning or self propagating worms is a significant, on-going threat in the Internet. Clever new scanning methods are constantly being developed to thwart identification by standard network intrusion detection systems (NIDS). Work by Staniford *et al.* [32] and by Moore *et al.* [23] project and evaluate the magnitude of the threat of new classes of worms and the difficulty of containing such worms. The conclusions of both papers is that addressing these threats presents the research community with serious challenges. An important step in protecting networks and systems against malicious intrusions is to im-

prove measurement and detection tools.

One means for improving the perspective and effectiveness of detection tools is to monitor both used *and* unused address space in a given network. Monitoring unused addresses is not typically done since packets to those addresses are often dropped by a network’s gateway or border router. However, tracking packets destined for unused addresses offers a number of important advantages. First, other than misconfiguration, packets destined to unused addresses are almost always malicious, thus false positives - a significant problem in NIDS - are minimized. Second, unlike NIDS that monitor traffic passively, a detection tool that monitors unused addresses can actively respond to connection requests, thus enabling the capture of data packets with attack-specific information. Third, a detection tool that monitors unused address space can masquerade in different ways, thus offering the possibility of slowing attacks (so-called tarpitting). The possibility for unused address space monitoring is perhaps most significant in class A and class B networks, where the number of unused addresses is often substantial. The idea of monitoring unused address space has been adopted in a number of different studies and on-going projects including the DOMINO project [37], Honeyd [13], the Honeynet project [34], LaBrea tarpits [17] and in the backscatter analysis conducted by Moore *et al.* in [24].

This paper makes two contributions. The first is our description of a new system architecture and implementation for measuring IP packet traffic. An *Internet Sink* or iSink, is a system we have developed for monitoring abuse traffic by both passive and active means. An iSink falls in a unique design space including capabilities of both high speed monitoring systems and flexible routing systems. The key design requirements of an iSink are extensibility of features and scalability of performance since it is meant to be used to monitor potentially large amounts of address space.

Our design of an iSink includes capabilities to trace packets, to actively respond to connection requests, to masquerade as several different application types, to fingerprint source hosts and to sample packets for increased scalability. The passive component of our implementation (which we call Passive Monitor) is based on Ar-

*V. Yegneswaran and P. Barford are members of the Computer Science Department at the University of Wisconsin, Madison. E-mail: vinod.pb@cs.wisc.edu. D. Plonka is a member of the Division of Information Technology at the University of Wisconsin, Madison. E-mail: plonka@doit.wisc.edu

gus [4] - a freely available IP flow measurement tool. The active component of our implementation (which we call Active Sink) is based on the Click modular router platform [15]. Click is open source software designed to be extensible and to provide high performance on commodity hardware. Click was a natural choice for our implementation since it directly addresses our design requirements. Development with Click required us to write modules for each of iSink's active capabilities. Our focus in this regard was on developing a set of "responders" which generate response packets for connections targeted at different network services. The resulting implementation was deployed and tested to assess its performance capabilities.

The second contribution of this paper is a measurement and evaluation case study of our iSink implementation. We use results from the case study to demonstrate the scale and diversity of traffic characteristics exposed by iSink based monitoring. These results provide validation for our architectural requirements and rationale for subsequent evaluation criteria. We deployed the iSink *in situ* to monitor four class B address spaces within our campus for a period of 4 months and one entire class A address space to which we have access. From these data sets we report results that demonstrate the iSink's capabilities and the unique information that can be extracted from this measurement tool. One example is that traffic characteristics from our class B monitor are substantially different from those on the class A monitor. This suggests that the *location* of the iSink in IP address space is important. Another example is the identification of periodic probing in our class A monitor which we were able to isolate to the LovGate worm [1]. We also uncover an SMTP hot-spot within the class A network that has been unreported prior to our study. We were able to attribute this anomaly to misconfigured wireless routers from a major vendor. Finally, we assess basic performance of the iSink in controlled laboratory experiments, and show our implementation to have a highly scalable active response capability.

These results demonstrate that our iSink architecture is able to support a range of capabilities and to provide scalable performance. The results also demonstrate that Active Sinks are a simple and very useful way to extend basic intrusion monitoring capabilities in individual networks or in the Internet as a whole.

2 Related Work

The notion of monitoring unused IP addresses as a source of information on intrusions has been in use in various forms for some time. While we coin the terms "Internet Sink" and "iSink", these monitors have variously been referred to as "Internet Sinkholes" [10], "Blackhole Routers" [11] and "Network

Telescopes" [20]. Traditional *Honeypots* are defined as systems with no authorized activity that are deployed with the sole purpose of monitoring intrusions. *Honeynets* are network of honeypots (typically set up as VMware hosts). Their deployment is often associated with significant management and scalability challenges [34]. The systems that are perhaps most similar to the Active Sink have been developed in the Honeyd and LaBrea Tarpit projects [13, 17]. Active Sink's design differs in significant ways from these two systems. Much like the Active Sink, Honeyd is designed to simulate virtual honeypots over unused IP addresses, with the potential for a diverse set of interactive response capabilities. However, Honeyd's stateful active responder design has significant scalability constraints that make it inappropriate for monitoring large IP address ranges which is one of iSink's primary objectives. LaBrea's primary design objective is to slow the propagation of Internet worms (*i.e.*, a sticky honeypot), and as such, it lacks the richness of interaction capabilities that is required to gather important response information. In addition to a richer response set, our Active Sink's performance greatly exceeds that of LaBrea as will be seen in Section 5.

There are a number of empirical studies of intrusion and attack activity that motivate and inform our work. In [38], the authors explore the statistical characteristics of Internet intrusion activity from a global perspective. That study is based on the use of intrusion logs from NIDS and firewalls located broadly across the Internet. Moore *et al.* examined the global prevalence of denial-of-service attacks using backscatter analysis in [24]. That work was conducted by gathering packet traces from a relatively quiescent class A network. Characteristics of the Code Red worm have been analyzed in a number of studies. In [22] the authors investigate the details of the Code Red outbreak and provide important perspective on the speed of worm propagation. The authors do a similar post-mortem analysis on the recent SQL-Slammer worm in [21]. Staniford *et al.* extend prior work on Code Red evaluation by providing a model for its propagation based on epidemiological analysis in [32]. Moore *et al.* provide further insights on the speed at which countermeasures would have to be installed to inhibit worms propagation [23]. While the prospects for successful containment are rather grim, it is clear that rapid detection will be a key component in any quarantine strategy.

Intrusion detection systems are a standard component in network security architectures. These tools typically monitor packet traffic at network ingress/egress points and identify potential intrusions using a variety of techniques. Standard methods for intrusion identification include misuse detection (*eg.* [26, 30]), statistical anomaly detection (*eg.* [31]), information retrieval (*eg.* [2]), data

Table 1: Design Space of Sink-Hole Responders

	Configurability	Modularity	Flexibility	Interactivity	Scalability
Active Sink	High	High	High	Low-Medium	High
Honeyd	High	Low-Medium	High	Low-Medium	Low-Medium
Honeynet	Low	Medium	Medium	High	Low-Medium
LaBrea	Low	Low	Low	Limited	High

mining (eg. [16]), and inductive learning (eg. [33]). Our work is distinguished from general NIDS in that they operate on active IP addresses and must deal with the problem of identifying the nefarious traffic mixed in with all of the legitimate traffic. We expect iSinks and NIDS to complement each other in future operational environment.

High performance packet monitors have been used for collecting packet traces in the Internet for years. These systems relate directly to our iSink design in that they must scale to reliably log packets on very high speed links. Examples of these include systems that have been developed with a variety of commodity and special purpose hardware such as [5, 7, 14]. Our iSink differs significantly from these systems (as well as the NIDS mentioned above) in that it not only passively monitors and logs packets, but it also *actively responds* to incoming TCP connection requests and has application level response capability.

3 Internet Sink Architecture

In this section we describe the iSink requirements, architecture and implementation. The implementation is described within the context of deployments on two different sets of address spaces.

3.1 Design Requirements

The general requirements for an iSink system are that it be secure and possess scalable capability for both passive and active monitoring. We discuss the issues of security in more detail in Section 7.

Passive monitoring capability must be able to capture packet header and payload information accurately. While there are many standard tools and method for packet capture, if either these or new tools are employed, they should be flexible and efficient in the ways in which data is captured and logged.

Active response capability is included in iSink’s design as a means to gather more detailed information on abuse activity. This capability is enabled by generating appropriate response packets (at both transport and application levels) to a wide range of intrusion traffic. While active responses also have the potential to interfere with malicious traffic in beneficial ways such as tarpitting, this is not a focus of iSink’s design.

We expect Internet Sinks to measure abuse activity over potentially vast unused IP address spaces. For ex-

ample, in our experimental setup, we needed the ability to scale to an entire class A network (16 million addresses). With the continued growth in malicious Internet traffic, and transition to IPv6, we expect the scalability needs to grow significantly for both the active and passive components of our system. Our basic approach to scalability is to maintain as little state as possible in our active responders. Another means for increasing scalability is through the use of sampling techniques in both active and passive components of the system. If sampling is employed, then the measurement results must not be substantially altered through their use.

Finally, our intent is to develop iSink as an open platform, thus any systems that are used as foundational components must be open source.

3.2 Active Response: The Design Space

In this section we explore the architectural alternatives for sink-hole response systems. The choices we consider are LaBrea, Honeyd, Honeynets and Active Sink (iSink’s active response system) as shown in Table 1. We compare these systems based on the following characteristics.

1. **Configurability** describes the ability of the configuration language to define the layout and components of response networks. Honeyd’s strengths are in fine-grained control of virtual network topologies and network protocol stacks. However Honeyd’s language does not provide support for assigning large blocks of IP addresses to templates (except for the default template)¹. Active Sink’s configuration language (inherited from Click) uses a BPF like language and provides excellent support for both fine-grained and coarse-grained control of a virtual network topology. Active Sink’s design is stateless and hence does not replicate network stack retransmission timers. LaBrea and Honeynets only allow for limited configurability.
2. **Flexibility** relates to the ability to mix and match services with operating systems. For example, the ability to define two types of Windows Servers: one with a telnet service and FTP service and another with NetBIOS Service and a Web server. The design of Honeyd and Active Sink both provide a high

¹This feature is particularly necessary for large network sinks

degree of flexibility. It is somewhat harder to do the same with Honeynets. LaBrea's flexibility in this regard is limited as it was designed with a different objective.

3. **Modularity** describes the ability to compose and layer services on top of one another. For example, layering Server Message Block (SMB) service over NetBIOS or layering Web services over SSL. Active Sink's design is inherently modular which directly facilitates service composition. In contrast, Honeyd's design is more monolithic and hence less straightforward to layer services.
4. **Interactivity** refers to the scope of response capability. The levels of interactivity of Honeyd and Active Sink are comparable. Obviously, Honeynets could provide more complete response capabilities. However, to mitigate the risk of Honeynets being used as a stepping-stone for additional attacks, data controls are required to be placed which limit interactivity. There are other practical configuration issues that also could limit interactivity. For example, Active Sink's NetBIOS responder grants session requests for all NetBIOS names and all user/password combinations, while a Honeynet Windows monitor would only allow NetBIOS session requests if it matches its list of valid names. Hence, the realized degree of interaction could be lower than low-interaction responders in Honeyd or Active Sink.
5. **Scalability** refers to the number of connections that can be handled in a given time period. In our monitoring environment we typically see hundreds of thousands of connection attempts per minute. Active Sink's stateless kernel module design provides high degree of scalability by eliminating unnecessary system calls and interrupt handling overheads. LaBrea's stateless design also provides reasonable scaling properties, however its user level implementation makes it inferior to the Active Sink. A weakness of Honeyd's design is its inherent statefulness that limits its scalability (Honeyd forks a process per connection attempt). Our experience suggests that Honeyd works well in environments that see tens of connection attempts per minute. The scalability of Honeynet systems vary from low to medium depending on the service and licensing issues.

3.3 Implementation

The objective of our monitoring infrastructure implementation was to create a highly scalable backplane with sufficient interactivity to filter out known worms, attacks and misconfigurations. To accomplish this, the iSink design includes a Passive Monitor, an Active Sink and a

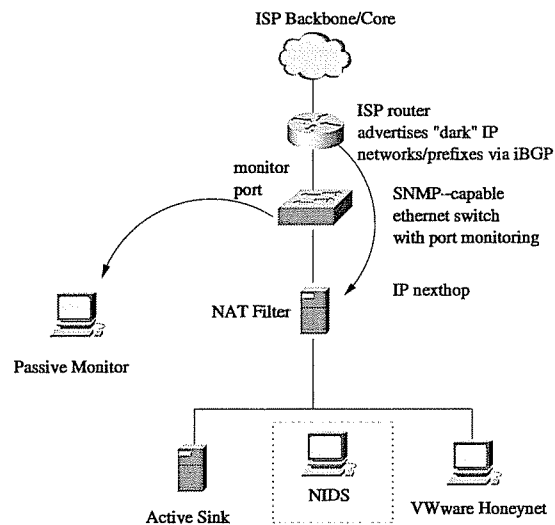


Figure 1: Internet Sink Implementation. In our current implementation the NIDS is run offline.

Honeynet component. Unsolicited traffic can be directed to each of these components which provide unique measurement capabilities. These components, in addition to MRTG [25] and FlowScan [28], were run on Linux-based commodity PCs. Details of our implementation as illustrated in Figure 1 and include:

1. **Passive Monitor** - This component is based on Argus which is a generic libpcap based IP network auditing tool. It allows for flow level monitoring of sink traffic and can be interfaced with FlowScan which is a flow level network traffic visualization tool.
2. **Active Sink** - The standard collection of elements provided with Click enabled most of the basic capabilities required for building active responses in iSink. Figure 2 illustrates iSink's configuration based on Click's modular design. Some of the important elements include: Poll Device which constantly polls the interface for new packets; IP Classifier which routes ARP packets to the ARP Responder, ICMP ping packets to the Ping Responder and TCP packets to the Windows Responder (all other packets are discarded); Windows Responder which responds to connection attempts on open ports and forwards HTTP requests to the Web Responder and SMB data packets to the NetBIOS Responder. As far as we know, *we are the first non-commercial Honeypot system to provide emulation capabilities for Windows Networking(NetBIOS/SMB/CIFS).*
3. **NAT Filter** - This module serves two purposes. It routes requests to appropriate responders (Active Sink or Honeynets) through network address

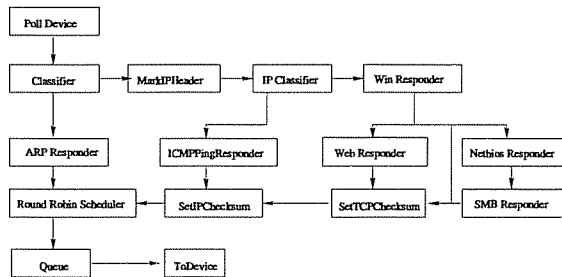


Figure 2: Active Sink Configuration based on Click's modular design

translation. It also filters requests that attempt to exploit known vulnerabilities or misconfigurations. The motivation behind filtering is to reduce the volume of traffic generated by active responders. This makes mapping of iSinks more difficult and increases scalability of analysis daemons that have to process large volumes of data. We experimented with several filtering strategies:

For each source IP allow only:

- (a) first N connections
- (b) first N connections per <destination port>
- (c) connections to first N destinations IPs targeted by the source

Of the three strategies, **option (c)** [N destination IPs per source IP] seemed the most attractive. The performance of options (a) and (c) were comparable. They both provided two orders of reduction in the volume of packets and bytes and were significantly better than option(b). We chose **option (c)** because it has the additional advantage of providing a *consistent view* of the network to the scan sources thus allowing the iSink to appear as if it were a subnet with N live hosts ².

4. **VMware Honeynets** - These are, quite simply, commodity operating systems running on VMware. Currently, we route packets of services for which we don't have complete responders to fully patched Windows systems.
5. **NIDS** - This system can be used to evaluate the packet logs collected at the filter. We plan to implement support for NIDS rules that can communicate with the filter and implement real time filtering decisions. For example, the decision to route packets or migrate connection to VMware Honeynet could be triggered upon the absence of a signature in the NIDS ruleset for the connection.

²The set of N destination hosts varies with each source depending on the order in which the source scans the address space.

For this study, we built and deployed two separate iSinks: a "campus-enterprise" iSink and a "service-provider" iSink. These were used to assess our iSink design and demonstrate its capabilities.

3.4 Deployment: Campus-Enterprise Sink

The campus iSink received unsolicited traffic destined for approximately 100,000 unused IPv4 addresses within 4 sparsely-to-moderately utilized class-B networks that are in use at our campus. Essentially, these unused addresses are in the "holes" between active subnets, each of which typically contains 128 to 1024 contiguous host addresses (*i.e.*, 25 through 22-bit netmasks, respectively).

A so called "black-hole" intra-campus router was configured to also advertise the class B aggregate /16 routes into the intra-campus OSPF. The result was that there were persistent less-specific (16 bit netmask) routes for every campus address. Unsolicited traffic, whether from campus or outside sources, destined for unused campus IP addresses always "falls through" to those less-specific /16 routes, and therefore is routed to the iSink and measured. Furthermore, *occasionally* traffic destined for campus addresses that are normally in use can fall through to the iSink if its subnet's more specific route disappears. Typically, this only happens during network outages, making the iSink a potential warning system of problems because it can passively detect routing failures. Whenever traffic that was destined for a campus IP address known to be in use reaches the iSink instead, the operators know that there is a problem.

It was important in our environment that the iSink machine was not capable of actively participating in the intra-campus routing, other than to respond via ARP as the IP nexthop on its transit link. The iSink is not an OSPF router, but instead is the destination of a static route. This limits the possible damage that could be caused if ever the iSink system was compromised and was attempted to be used maliciously.

3.5 Deployment: Service-Provider Sink

The service-provider iSink received unsolicited traffic destined for 16 million IPv4 addresses in one class A network. An ISP router, located at our campus' service-provider, served as the gateway for the service-provider iSink. The service-provider was responsible for advertising the class A network via BGP to our service provider's commercial transit providers, Internet2's Abilene network, and to various other peers. SNMP-based measurements at the Ethernet switch's ports were used to compute any packet loss by the libpcap-based Argus software.

4 Experiences with Internet Sink

This section demonstrates iSink's capabilities and illustrates the complementary roles of the Passive Monitor and the Active Sink using results from our two iSink deployments. We first discuss issues of perspective by comparing the passive-monitoring results observed in the campus-enterprise sink with that of the service-provider sink. We then demonstrate the utility of the Active Sink in investigating network phenomenon revealed by the Passive Monitor including periodic probing and SMTP hot-spots. Finally, we measure the impact of active response in attracting additional flows and provide results from passive fingerprinting of network scanners.

4.1 Campus Enterprise iSink Case Study

Because the campus iSink is located inside one autonomous system and advertised via the local interior routing protocol, this system sees traffic from local sources in addition to traffic from sources in remote networks. Traffic observed from local sources included:

- Enterprise network management traffic attempting to discover network topology and address utilization (such as ping sweeps and SNMP query attempts)
- Traffic from misconfigured hosts. For instance, a few hosts continually send domain queries to what is now an unused campus IP address. Presumably, an operational DNS server used to be at that address. We also see traffic from misconfigured AFS clients and NetBIOS name registration requests from local windows hosts with incorrect WINS address.
- Malicious probes and worm traffic that has an affinity for hosts within their classful network.

Figure 3 shows the traffic observed from only remote sources in a typical week at the campus-enterprise iSink. There are several notable features. The dominant protocol is TCP since the campus border routers filter scans to port 1434 (ms-sql-m) that was exploited by the SQL-Slammer worm [21]. The peak rate of traffic is about 1Mb/s and 1500 packets per second. There is no obvious periodicity in this dataset. Finally, because TCP is the dominant protocol, the packet sizes are relatively constant and the number of bytes, packets, and flows follow a predictable ratio. Hence, the graphs of flow, byte, and packet rate show very similar trends.

4.2 Service Provider iSink Case Study

The volume of unsolicited inbound traffic to the class A network varied between average rates of 5,000 packets-per-second (pps) when we brought the system on line to over 20,000pps six months later at the end of our study.

Table 2: Top Services (Service Provider Sink)

Service:	Inbound flows per second
udp_netbios-ns_dst	1932
udp_ms-sql-m_dst	1187
http_dst	197
netbios-ssn_dst	133
microsoft-ds_dst	115
smtp_dst	67
http_src	44
https_dst	11
ms-sql-s_dst	10
telnet_dst	2

Table 3: Active backscatter sources (victims) in service provider sink (12 hours - 5 minute averages)

Type	Num IPs	% IPs
TCP_RST	295	38%
TCP_SYN_RST	105	14%
TCP_ACK	81	10%
TCP_ACK_RST	80	10%
ICMP_INTRANS.TIME_EXCEEDED	58	7%
ICMP_PORT_UNREACH	29	4%
ICMP_PKT_FILTERED_UNREACH	23	3%
TCP_SYN_ACK	10	1%
ICMP_HOST_UNREACH	6	1%
OTHER	87	11%

One consequence that was relayed to us by experienced network operators is that it is not possible to effectively operate even this relatively quiescent class A network at the end of a 1.5 megabit-per-second T1 link because the link becomes completely saturated by this unsolicited traffic.

To operate the service-provider iSink continuously, we originally assumed that we could safely introduce the class A least-specific /16 route for the iSink and still allow operators to occasionally introduce more-specific routes to draw the network's traffic elsewhere in the Internet when need-be. While sound in theory (according to "CIDR and Classful Routing" [29]), it didn't work in practice. Because today's Internet is bifurcated into commercial/commodity networks and research/education networks (Internet2's Abilene), some institutions connected to both types employ creative routing policies. We found that some sites prefer less-specific routes over more-specific *when* the less-specific route is seen on what is likely to be a higher-performance (or fixed cost) service such as Internet2.

Figure 4 depicts the traffic observed in a typical week at the service-provider iSink. Unlike the campus-enterprise network, the dominant protocol is UDP, most of which can be attribute to Windows NetBIOS scans on port 137 and the ms-sql-m traffic from worm attempting to exploit the vulnerable MS-SQL monitor. Since UDP




Figure 3: Inbound Traffic for a Typical Week on Campus-Enterprise Sink (bits/pkts/flows per second)




Figure 4: Inbound Traffic for a Typical Week on Service Provider Sink (bits/pkts/flows per second)




Figure 5: Time-volume graph of backscatter packet types on service-provider sink over a typical 12 hour period

traffic with payloads of varying sizes dominates, there is no strong correspondence between the graphs for bytes, flows and packets. The most interesting feature is the striking periodic behavior of the TCP flows, discussed in more detail in the section 4.2.2. Table 2 provides a summary of the inbound per second flow rate of the top services.

4.2.1 Analysis of Backscatter Packets

Backscatter packets are responses to spoofed DoS attacks and have been effectively used to project Internet wide attack behavior [24]. Figure 5 provides a time series graph of the backscatter packet volume observed in our service-provider sink. Noteworthy features include the following:

1. TCP packets with ACK/RST dominate as might be expected. This would be the most common re-




Figure 6: Inbound flows (per second) observed at service-provider sink on ports 139 and 445 over a typical week

sponse to a SYN flood from forged sources.

2. Less common short duration spikes of SYN/ACK and SYN/ACK/RST.
3. ICMP TTL exceeded packets could be attributed to either routing loops or DoS floods with a low initial TTL.

Table 3 provides a summary of the number of active sources of backscatter traffic, *i.e.*, the estimated count of the victims of spoofed source attacks. These numbers are an average during the 12 hours shown in Figure 5 of the number of sources in each 5 minute sample. In terms of the distribution of the volumes of Backscatter scan types, our results are consistent with those published in [24]. Backscatter made up a small percentage (under 5%) of the overall traffic seen on our service-provider sink.




Figure 7: Volume/Count of *type-1* port 139 scanners: 24 hours, Dec 14, 2003, (no. sources per peak = 100, total sources = 2,177)

4.2.2 Investigating Unique Periodic Probes

The periodicity observed in the service provider iSink data is an excellent example of the perspective on intrusion traffic afforded by iSink. The first step in our analysis of this periodicity was to understand the services that contributed to this phenomenon. We found that most of the periodicity observed in the TCP flows could be isolated to sources scanning two services (port 139 and 445) simultaneously. Port 139 is SMB (Server Message Block protocol) over NetBIOS and port 445 is direct SMB. However, this did not help us isolate the attack vector because it is fairly common for NetBIOS scanners to probe for both these services. Passive logs provided three additional clues: 1) scans typically involve 256 IP successive addresses that span a /24 boundary, 2) the probes had a period of roughly 2.5 hours, 3) the small timescale periodicity seemed to be super imposed over a diurnal periodic behavior at larger timescales.

Figure 6 shows the number of flows scanning both services in a week. To simplify our analysis we then focused on a single day's data and classified scanners on these services based on their scan footprints. We defined scanners that match our profile (between 250-256 successive IP addresses spanning a /24 boundary) as *type-1* sources. We also defined sources that scan five or more subnets simultaneously as *type-5* sources. This includes processes that pick destination IP addresses randomly and others that are highly aggressive. Figure 7 shows a time-volume graph of the *type-1* scanners. **The interesting aspect of this figure is that the number of sources in each peak (around 100) is more than an order of magnitude smaller than the total number of participants observed in a day (2,177).** From Figure 8, we can see that most of the diurnal behavior could be attributed to *type-5* sources.

This mystery motivated our development of NetBIOS and SMB responders. By observing the packet logs gen-




Figure 8: Volume of *type-5* port 139 scanners: 24 hours, Dec 14, 2003

erated by the active response system shown in Figure 9, we concluded that the scanning process was the LovGate worm [1] which creates the file `NetServices.exe` among others.

We proceeded to setup a controlled experiment which began by trying to infect a Windows 2000 host running on VMware with LovGate. LovGate uses a dictionary attack, so we expected a machine with blank administrative password to be easily infected. However, the NetBIOS sessions were continually getting rejected due to NetBIOS name mismatches. So we modified the `lmhosts` file to accept the name `*SMBSERVER` enabling us to capture the worm.

We verified that LovGate's NetBIOS scanning process matched the profile of the *type-1* scanners³. To date, we have not been able to disassemble the binary as it is a compressed self-extracting executable. So we monitored the scans from the infected host. There were two relevant characteristics that provide insight into the periodicity: 1) The scanning process is deterministic, *i.e.*, after every reboot it repeats the same scanning order 2) During the course of a day there are several 5-10 minute intervals where it stops scanning. Our conjecture is that these gaps occur due to approximately synchronized clocks in the wide area thus producing the observed periodicity.

4.2.3 SMTP Hot-spot

Analysis of SMTP (Simple Mail Transfer Protocol) scans in the service provider sink is another important demonstration of active sink's capabilities. From passive measurements, we identified an SMTP hot-spot *i.e.*, there was one IP address that was attracting a disproportionately large number of SMTP scans (20-50 scans per second). Hot-spots in unused address space are typically good indicators of misconfigurations. During a 10 day period in December we observed over 4.5 million scans from around 14,000 unique IP addresses all bound to one

³Besides the NetBIOS scanning LovGate also sent SMTP probes to www.163.com

```

Client: 61.110.200.170
NB_SESSION_REQUEST
NB_SESSION_RESPONSE
SMB_COM_NEGOTIATE TYPE=REQUEST
SMB_COM_NEGOTIATE TYPE=RESPONSE
SMB_SESSION_SETUP_ANDX TYPE=REQUEST

\IPC$????

SMB_SESSION_ANDX TYPE=REPLY
SMB_NT_CREATE_ANDX TYPE=REQUEST

\svcctl

NB_SESSION_REQUEST
NB_SESSION_RESPONSE
SMB_COM_NEGOTIATE TYPE=REQUEST
SMB_COM_NEGOTIATE TYPE=RESPONSE
SMB_SESSION_SETUP_ANDX TYPE=REQUEST

\ADMIN$????

SMB_NT_CREATE_ANDX TYPE=REQUEST

\system32\NetServices.exe

```

Figure 9: Trace of LovGate interaction with Active Sink

destination IP within our monitor. A cursory analysis suggested that these scans were all from cable-modem and DSL subscribers. Finally, the scans also seemed to have an uncommon TCP SYN fingerprint (win 8192, mss 1456).

The possibility of spam software as a source of this anomaly was ruled out due to the non-standard TCP fingerprint. We then hypothesized that this could be from a specific cable-modem or DSL device. We set up an SMTP responder on the target IP address and captured the incoming email. **This revealed the source of the email to be misconfigured wireless-router/firewall systems from a major vendor⁴.** The emails are actual firewall logs. Figure 10 is an anonymized example where we have replaced original IP addresses with private addresses.

To better understand the reasons behind this SMTP hot-spot, we examined to examine the firewall system’s firmware. The `unarj` utility was used to extract the compressed binary. However, searching for the hot-spot IP address string in the binary proved fruitless. Examination of the firmware “application” revealed that there was an entry for SMTP server that was left blank by default. This led us to conjecture that the target IP address was the result of an uninitialized garbage value that was converted to a network ordered IP address. It also turns out that every byte in our hot-spot address is a printable ASCII character. So we searched for this four byte ASCII string and found a match in *almost all versions of firmware for this device*. The string occurred in both the extracted and compressed versions of the firmware. As a

⁴We are in the process of notifying the manufacturer and plan to reveal the name of the vendor once this is completed

Table 4: Passive OS Fingerprinting (Discriminant Fields)

OS	IP TTL	Window	TCP Options
Linux 2.2	64	32120	mss, sack, ts, ws, nop
Linux 2.4	64	5840	mss, sack, ts, ws, nop
FreeBSD-4.0	64	16384	mss
Solaris 7	255	8760	mss
Windows 2000/NT	128	(16384) variable	mss, sack, 2 nops mss, wscale, 3 nops, sack

Table 5: Passive fingerprinting summary for a typical day on the service provider iSink

OS	Num Scans
All	97,197,009
Window 2000/NT	86,054,187
Linux 2.4	5,373,467
Linux 2.2	26,054
Solaris 7	1,106,032
FreeBSD	112,760
Other ⁵	4,524,509

sanity check, we looked for other similar ASCII strings, but did not find them. These kind of hot-spots can have very serious ramifications in network operations. For example, one the authors discovered a similar problem with Netgear routers that inadvertently flood our campus NTP servers [27].

4.3 Attack Source Fingerprinting

The ability to fingerprint operating systems often proves valuable in characterizing attack sources and identifying exploited vulnerabilities in attack drones. iSink naturally lends itself to this activity. We identify three different techniques that can be used in combination to categorize offending hosts.

- **Passive Fingerprinting** techniques try to identify the operating system or attack tools without sending any packets. There are several fields in the TCP and IP headers that enable identification of attackers. Table 4 provides an illustration of the notable discriminants among contemporary protocol stacks [19]. These techniques are all vulnerable to spoofed headers, however by combining multiple techniques it is usually possible to narrow candidates down to one or two systems.

We performed passive OS fingerprinting on one day’s data from the service-provider iSink. We were able to fingerprint 95% of the nearly 100 million TCP scans that were observed. The summary of our results is given in Table 5.

Passive techniques can also be used to discriminate between various tools used by the blackhats. In [3] Arkin describes techniques to distinguish between five popular ICMP tools.

```

Received: from XXX
        by localhost (8.9.3/8.9.3) with id h242udeJ41951
        for <>; Sun, 4 Jan 2004 08:22:24 -0600 (CST)

From:
Subject: XXX Log
Sender: XXX
To:

Jan/04/2004 00:20:38
Drop TCP packet from WAN src 192.168.10.10:2853 dst:192.168.10.19:80 Rule: Default deny
Jan/04/2004 00:20:32
Drop TCP packet from WAN src:192.168.10.10:2853 dst:192.168.10.19:80 Rule: Default deny
Jan/04/2004 00:20:29
Drop TCP packet from WAN src:192.168.10.10:2853 dst:192.168.10.19:80 Rule: Default deny
Jan/04/2004 00:20:25
Drop TCP packet from WAN src:10.42.42.10:1825 dst:192.168.10.19:25 Rule: Default deny
...

```

Figure 10: Sample email captured at the SMTP Hot-spot

- **Proactive Fingerprinting** refers to techniques that proactively establish one or more new connections for the purpose of network or host reconnaissance. These techniques can provide highly accurate information, for instance, minor version of host operating systems and services. However, they are limited by ethical and practical considerations. They also do not provide information on the tools used by the attacker.
- **Reactive Fingerprinting** techniques try to gather more information about the intruder, by exchanging packets within the context of the existing connection. In the case of a tarpit, the intervals between retransmission timeouts and window probes in persist mode can yield valuable clues about the source TCP stack. For example, in our measurements the Windows 2000 and Linux TCP stacks have completely different timeout intervals for resending the response packet for the tarpit's SYN-ACK. These numbers are repeatable to 10 ms granularity.

- Windows 2000: $r_1 = 2.92^6$, $r_2 = 6.02$, $r_3 = 12.04$, $r_4 = 24.08$, $r_5 = 48.15$ and then stops.
- Linux 2.4: $r_1 = .21$, $r_2 = .42$, $r_3 = .99$, $r_4 = 1.58$, $r_5 = 3.36$, $r_6 = 6.72$, $r_7 = 13.44$, $r_8 = 26.88$, $r_9 = 53.76$, $r_{10} = 107.52$, $r_{11} \dots r_{18} = 2$ minutes.

4.4 Differentiated Responses

We investigated how probe traffic is affected by network responses. For this purpose, we dedicated 10 of the /16's from the service-provider iSink to emulate responses from systems running various commodity operating systems and a few popular applications. Table 6 provides an outline of our designated partitions. These active subnets

⁶Here r_1 refers to interval in seconds between original and 1st duplicate packet, r_2 refers to interval between 1st duplicate and 2nd duplicate packet and so on.

Table 6: Differentiated Response Networks (service provider sink)

ID:	Name
.1.0/16	Synacker
.3.0/16	Synacker + Ping Resp
.5.0/16	Windows
.7.0/16	Windows + Kazaa (1214)
.9.0/16	Windows + Gnutella (6346)
.11.0/16	RedHat Linux
.13.0/16	RedHat Linux + HTTP(80)
.15.0/16	Solaris
.17.0/16	Solaris + HTTP
.19.0/16	ICMP Unreachable

were interleaved with passive control group networks (even numbered partitions of same sizes, omitted from the table).

The observed rate of inbound flows for the active and passive (control) subnets is shown in Figure 11. The average rates for each type are provided in the legend. If probe behavior were unaffected by responses, we would expect the active subnets to see the same number of inbound flows (although different number of packets) as the passive nets. However, that is clearly not the case. The passive networks all see very similar scanning rates except for the rare event spikes. *Our results suggest that responding to all connection attempts (synacker) or responding to Windows ports and in particular the Kazaa port, significantly increases the network's affinity to additional probing.* Responding to ICMP pings, Gnutella and HTTP also seem to introduce additional scanning. Finally, responding with ICMP unreachable reduced the incoming flow rate to 13 flows per second in comparison with passive subnets which averaged 15 flows per second. This suggests that firewall rules might be best served by responding with ICMP unreachable to anomalous traffic (rather than dropping packets).

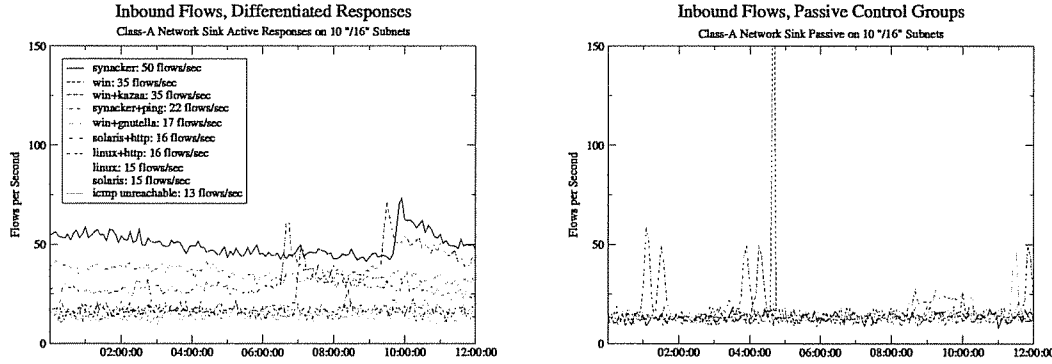


Figure 11: Service Provider Sink: Differentiated Responses (inbound flows per second) active (left) and passive-control(right) subnets for a 12 hour period

5 Basic Performance

One of the primary objectives of the iSink's design is scalability. We performed scalability tests on our Active Sink implementation using both TCP and UDP packet streams. The experimental setup involved four 2GHz Pentium 4 PCs connected in a common local area network. Three of the PCs were designated as load generators and the fourth was the iSink system that promiscuously responded to all ARP requests destined to any address within one class A network. Figures 12 demonstrates the scalability under of LaBrea⁷ and Active Sink under TCP and UDP stress tests. The primary difference between the TCP and UDP tests is that the TCP connection requests cause the iSink machine to respond with acknowledgments, while the UDP packets do not elicit a response. Ideally, we would expect the number of outbound packets to equal the number of inbound packets. The Click-based Active Sink scales well to TCP load with virtually no loss up to about 20,000 packets (connection attempts) per second. LaBrea performance starts to degrade at about 2,000 packets. The UDP test used 300 byte UDP packets (much like the SQL-Slammer worm). In this case, both the LaBrea and Active Sink perform admirably well. LaBrea starts to experience a 2% loss rate at about 15,000 packets/sec.

6 Sampling

There are three reasons why *connection sampling* can greatly benefit an iSink architecture: (i) *reduced bandwidth requirements*, (ii) *improved scalability*, (iii) *simplified data management and analysis*.

If sampling is to be used, the task then becomes selecting a sampling design that meets the requirements of reduced resource demands and accurate estimation of probe population characteristics. In our application, sampling could be applied either to the packets that are pas-

sively logged or to those for which active responses are sent. In the former case, the overhead reduction objective must be balanced with the error introduced in the measurements based on the selected sampling method. In the latter case, simple random sampling would be appropriate as a first step although we leave treatment of this topic to future work. While there is a large literature on sampling design and analysis (e.g., see [36]), recent results from several studies provide insight on how to approach the problem of sampling in the network measurement context including [8, 9, 18].

We considered two different resource constraint problems in the passive portion of the iSink and evaluated the use of sampling as a means for addressing these constraints. We first considered the problem of a fixed resource in the iSink itself. Estan and Varghese in [6] describe sampling methods aimed at monitoring "heavy hitters" in IP flows through routers with a limited amount of memory. We adapted one of these methods for use in iSink. Second, we considered the problem of bandwidth as the limited resource. In this case, the idea is to reduce the total amount of traffic routed to an iSink by selecting subnets within the total address space available for monitoring. These methods would be used in combination with the filtering methods described in Section 3.3.

6.0.1 Memory Constrained iSink Sampling

The method that forms the basis of our sampling approach with a memory constrained iSink is called *Sample and Hold* [6]. This method accurately identifies flows larger than a specified threshold (i.e., heavy hitters). Sample and hold is based on simple random sampling in conjunction with a hash table that is used to maintain flow ID's and byte counts. Specifically, incoming packets are randomly sampled and entries in the hash table are created for each new flow. After an entry has been created, *all* subsequent packets belonging to that flow are counted. While this approach can result in both false positives and false negatives, its accuracy is shown to be high in workloads with varied characteristics. We

⁷We compare Active Sink with LaBrea because unlike LaBrea, Honeyd is stateful(forks a process per connection), and hence is much less scalable. Since Honeyd also relies on a packet filter LaBrea's scalability bounds affect Honeyd as well.




Figure 12: Scalability of Click-based Internet Sink and LaBrea for TCP (left) and UDP (right) flows

apply sample and hold in iSink to the problem of identifying “heavy hitters”, which are the worst offending source addresses based on the observed number of scans.

Adapting the sample and hold method to the iSink required us to define the size of the hash table that maintains the data, and the sampling rate based on empirical observation of traffic at the iSink. In [6], the objective is identifying accurately the flows that take over $T\%$ of a link’s capacity. An oversampling factor O is then selected to reduce the possibility of false negatives in the results. These parameters result in allocating $HT_{len} = 1/T * O$ locations in each hash table. The packet sampling rate is then set to HT_{len}/C where C is the maximum packet transmission capacity of the incoming link over a specified measurement period t . At the end of each t , the hash table is sorted and results are produced.

6.0.2 Bandwidth Constrained iSink Sampling

In the bandwidth constrained scenario, the sampling design problem is to select a set of subnets from the total address space that is available for monitoring on the iSink. The selection of the number of subnets to monitor is based on the bandwidth constraints. In this case we assume that we know the mean and variance for traffic volume on a “typical” class B or class C address space. We then divide the available bandwidth by this value to get the number of these subnets that can be monitored. The next step is to select the specific subnets within the entire space that will minimize the error introduced in estimates of probe populations.

Our analysis in this paper is based on the use of random sampling as a means for subnet selection. Our rationale for this approach is based on the observation that overall traffic volumes across the service-provider class A address space that we monitor is quite uniform. The strengths of this approach are that it provides a simple method for subnet selection, it provides unbiased estimates and it lends itself directly to analysis. The drawback is that sampling designs that take advantage of additional information such as clustered or adaptive sampling could provide more accurate population estimates. We

leave exploration of these and other sampling methods to future work.

After selecting the sampling design, our analysis focused on the problem of *detectability*. Specifically, we were interested in understanding the accuracy of estimates of total probe populations from randomly selected subsets. If we consider $\hat{\tau}$ is an unbiased estimator of a population total τ then the estimated variance of $\hat{\tau}$ is given by:

$$var(\hat{\tau}) = N^2 \left[\left(\frac{N-n}{N} \right) \frac{\sigma^2}{n} + \left(\frac{1-p}{p} \right) \frac{\mu}{n} \right]$$

where N is the total number of units (in our case, subnets), n is the sampled number of units, μ is the population mean (in our case, the mean number of occurrences of a specific type of probe), σ^2 is the population variance and p is the probability of detection for a particular type of probe. In the analysis presented in Section 6.1, we evaluate the error in population estimates over a range of detection probabilities for different size samples. The samples consider dividing the class A address space into its component class B’s. The probabilities relate directly to detection of worst offenders (top sources of unsolicited traffic) as in the prior sampling analysis. The results provide a means for judging population estimation error rates as a function of network bandwidth consumption.

6.1 Sampling Evaluation

Our evaluation of the impact of sampling in an iSink was an *offline* analysis using traces gathered during one day selected at random from the service-provider iSink. Our objective was to empirically assess the accuracy of sampling under both memory constrained and bandwidth constrained conditions. In the memory constrained evaluation, we compare the ability to accurately generate the top 100 heavy hitter source list over four consecutive 1 hour periods using different hash table sizes and different sampling rates. For each hour in the data set, we compare the percentage difference in the number of scans generated by the “true” top 100 blacklist and sampled top 100 blacklist sources. In the bandwidth constrained evalu-

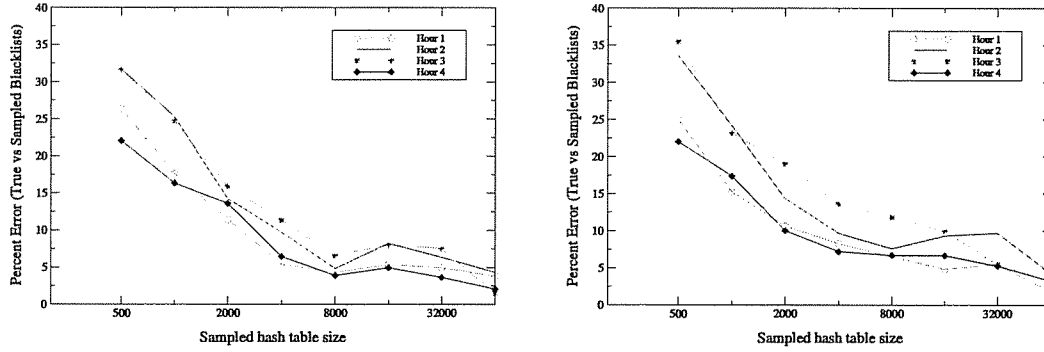


Figure 13: Error rates for different hash table sizes (x-axis is log scale) using Sample and Hold method with rates of 1/100 (left) and 1/300 (right).




Figure 14: Cumulative distribution of all traffic and TCP backscatter traffic across half of the class A address space monitor over a one hour period. The average number of probes per /16 is 1.6K packets per second.

ation, we consider accuracy along three dimensions: 1) estimating the worst offender population with partial visibility, 2) estimating black lists of different lengths, 3) estimating backscatter population.

Our memory constrained evaluation considers hash table sizes varying from 500 to 64K entries where each entry consists of a source IP and a access attempt count. Note that the hash table required to maintain the complete list from this data was on the order of 350K entries. We consider two different arbitrarily chosen sampling rates - 1 in 100 and 1 in 300 with uniform probability. In each case, once a source IP address has been entered into the table, all subsequent packets from that IP are counted. If tables become full during a given hour then entries with the lowest counts are evicted to make room for new entries. At the end of each hour, the top 100 from the true and sampled lists are compared. New lists are started for each hour. The results are shown in Figure 13. These results indicate that even coarse sampling rates (1/300) and relatively small hash tables enable fairly accurate black lists (between 5% - 10% error). The factor of improvement between sampling at 1/100 and

1/300 is about 1.5, and there is little benefit to increasing the hash table size from 5,000 to 20,000. Thus, from the perspective of heavy hitter analysis in a memory constrained system, sampling can be effectively employed in iSinks.

As discussed in the prior section in our bandwidth constrained evaluation we consider error introduced in population estimates when using simple random sampling over a portion of the available IP address space. We argue that simple random sampling is appropriate for some analyses given the uniform distribution of traffic over our class A monitor. The cumulative distribution of traffic over a one hour period for half of the /16 subnets in our class A monitor is shown in Figure 14. This figure shows that while traffic across all subnets is relatively uniform (at a rate of about 1.6K packets per second per /16), specific traffic subpopulations - TCP backscatter as an example - can show significant non-uniformity which can have a significant impact on sampling.

We use the mean normalized standard deviation (σ/μ) as an estimate of error in our analysis. In each case, we assess estimated error as a function of a randomly selected sample of /16 subnets. The results of this approach are shown in Figure 15. The leftmost graph shows the ability to accurately estimate the number of probes from the single worst offending IP source over a range of detection probabilities (*i.e.*, the probability of detecting a source in a selected /16). This graph indicates that worst offenders are detectable even with a small sample size and error-prone or incomplete measurements. The middle graph shows the ability to accurately estimate black lists from a selected sample of /16's. This graph indicates that it is easier to estimate larger rather than smaller black lists when sampling. We attribute this to the variability in black list ordering across the /16's. Finally, the rightmost graph shows the ability to accurately estimate TCP backscatter traffic over a range of detection probabilities. The graph indicates that while backscatter estimates are




Figure 15: Estimated error (σ/μ) for a given number of randomly selected /16's. Included are error estimates for total probes from worst offending source IP over a range of detection probabilities (left), estimates for total probes for worst offender lists of several sizes (middle), estimates for total TCP backscatter traffic (right).

robust in the face of error-prone or incomplete measurements, estimated error of total backscatter is quite high even with a reasonably large number of /16's. This can be attributed to the non-uniformity of backscatter traffic across the class A monitor shown in Figure 14 and suggests that alternative sampling methods for backscatter traffic should be explored. On a broader scale, this indicates that traditional backscatter methodologies that assumes uniformity could be error prone.

7 Discussion

In this section we discuss implementation and operational issues concerning iSink deployment and use. We also discuss extensions to the architecture which we are in the process of extending.

7.1 iSink Overhead

The use of iSinks will most certainly increase the amount of traffic in the Internet for two reasons. First, while unsolicited traffic destined for "Bogons" [35, 10] or other unused addresses is typically dropped by border routers, iSinks require this traffic to be forwarded. Without the iSink, such traffic could have been discarded by border routers. Second, since iSinks send active responses, they themselves contribute additional traffic. The question is, "is the additional cost for carrying this traffic and the potential congestion it causes worth the benefit of using iSinks?". We argue that the benefits outweigh the costs for most deployments since response packets sent by iSinks are of minimal size. For deployments on very large amounts of address space (e.g., multiple class A's) costs could be significant but could be justified based on the perspective gained from the data as well as the potential to expose and perhaps slow propagation of future worms.

7.2 iSink Security Model

As with any node on a network, iSinks are vulnerable to a number of potential attacks including denial of ser-

vice, break-in and obfuscation of data collected at a node. We assume that best practices for system management (patches, updates, etc.) will be employed at the nodes and out-of-band management used. If an iSink is not participating in a federated deployment, then it can be protected from external intrusions by firewalling any access to/from the Internet. If it is participating cooperatively with nodes in other networks, then we assume that the federation will have a security architecture that will protect individual nodes.

7.3 Legal Issues

Legal issues of *consent* potentially arise due to the active response component of our iSink. An iSink can be thought of as a "service" that is offered by unadvertised servers on a network. Thus, consent may be required from the Internet Service Provider. Another issue is that Honey pots may be considered a means of *entrapment* [34]. Finally, active response from an iSink to a connection request could be considered an act of "disruption or concealment of communication" which is illegal in some states. The LaBrea Project [17] recently encountered such problems. The issues are described in detail on their web site and is very much up in the air at this point in time, as laws vary from state to state. However, the result at this point is that LaBrea code is no longer available on their web site.

7.4 Advertising Unused IP space via BGP?

In [12], Greene proposes advertising specific CIDR blocks via BGP for "dark IP addresses". Similar to the Bogon list of unallocated and reserved prefixes, these dark IP advertisements would essentially form a blacklist in the global BGP table. Service providers could use this list to short-cut-route unsolicited traffic destined for blacklisted IP addresses to the nearest sink. This would have the potential benefit of relieving the Internet backbone from the responsibility of delivering this potentially harmful traffic.

However, one potential issue with this approach is that the blacklist will appear in the publicly available BGP routing tables. Would-be attackers could easily seed their probes or worms with this information thereby directing their attacks to address space that is more likely to have vulnerable hosts. Furthermore, with increased use of the sparsely populated IPv6 (but huge) address space, a BGP blacklist may effectively focus attacks to the slices of IPv6 address space where systems are deployed.

7.5 IPv6

The Internet Protocol version 6 (IPv6) address space is vast by comparison to the IPv4 address space. Consequently, unsolicited systematic sequential or random probing will take much more time (assuming the same amount of probing resources). Thus, the sparsely used address space presents a challenge for would-be abusers. However, it also presents a challenge for those intent on deploying iSinks. It seems likely that the overall amount of traffic in IPv6 iSinks will be substantially lower (assuming the same total amount of address space coverage) in the short term, thus reducing the effectiveness of their deployment.

7.6 Appliances and Middleboxes

Internet appliances such as Network/Port Address Translating (NAT/PAT) routers could act as either passive or active Internet sinks. An active sink in a PAT device could be effective with only one globally routable IP address, which is typical in residential and small-office/home-office applications. Today *teergrube*(tarpit) functionality is present in freely available SMTP Mail Transfer Agents (MTAs) to slow spam mail dissemination. Similarly, one could envision sink features provided in firewall systems and operating in conjunction with NIDS.

8 Conclusions and Future Work

In this paper we describe the architecture and implementation of an Internet Sink: a useful tool in a general network security architecture. iSinks have several general design objectives including scalability, the ability to passively monitor network traffic on unused IP addresses, and to actively respond to incoming connection requests. These features enable large scale monitoring of scanning activity as well as attack payload monitoring. The implementation of our iSink is based on a novel application of the Click modular router, NAT Filter and the Argus flow monitor. This platform provides an extensible, scalable foundation for our system and enables its deployment on commodity hardware. Our initial implementation includes basic monitoring and active response capability which we test in both laboratory and live envi-

ronments.

We report results from our iSink's deployment in a live environment comprising four class B networks and one entire class A network. The objectives of these case studies were to evaluate iSink's design choices, to demonstrate the breadth of information available from an iSink, and to assess the differences of perspective based on iSink location in IP address space. We show that the amount of traffic delivered to these iSinks can be large and quite variable. We see clear evidence of the well documented worm traffic as well as other easily explained traffic, the aggregate of which can be considered Internet background noise. While we expected overall volumes of traffic in the class B monitors and class A monitor to differ, we also found that the overall characteristics of scans in these networks were quite different. We also demonstrate the capability of iSinks to provide insights on interesting network phenomenon like periodic probing and SMTP hot-spots, and their ability gather information on sources of abuse through fingerprinting techniques.

The evaluation of our iSink implementation demonstrates both its performance capabilities and expectations for live deployment. From laboratory tests, we show that iSinks based on commodity PC hardware have the ability to monitor and respond to over 20,000 connection requests per second, which is approximately the peak traffic volume we observed on our class A monitor. This also exceeds the current version of LaBrea's performance by over 100%. Furthermore, we show that sampling techniques can be used effectively in an iSink to reduce system overhead while still providing accurate data on scanning activity.

We intend to pursue future work in a number of directions. First, we plan to expand the amount of IP address space we monitor by deploying iSinks in other networks. Next, we intend to supplement iSink by developing tools for datamining and automatic signature generation.

Acknowledgements

The authors would like to thank Jeff Bartig, Geoff Horne, Bill Jensen and Jim Martin for all of their help. The exploration of the filtering techniques grew out of fruitful discussions with Vern Paxson.

References

- [1] LovGate Virus Summary. <http://vil.nai.com/vil/content/Print100183.htm>, 2002.
- [2] R. Anderson and A. Khattak. The Use of Information Retrieval Techniques for Intrusion Detection. In *Proceedings of RAID*, September 1998.
- [3] Ofir Arkin. Identifying ICMP Hackery Tools Used In The Wild Today. *BlackHat*, 2000.

- [4] C. Bullard. Argus Open Project. [http : //www.qosient.com/argus/](http://www.qosient.com/argus/).
- [5] C. Cranor, Y. Gao, T. Johnson, V. Shkapenyuk, and O. Spatscheck. Gigascope: High Performance Network Monitoring with an SQL Interface. In *Poster Session of ACM SIGMOD '02*, Madison, WI, June 2002.
- [6] C. Estan and G. Varghese. New Directions in Traffic Measurement and Accounting. In *Proceedings of ACM SIGCOMM '02*, Pittsburgh, PA, August 2002.
- [7] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford. NetScope: Traffic Engineering for IP Networks. *IEEE Network Magazine*, Special Issue on Internet Traffic Engineering, 2000.
- [8] P. Gibbons and Y. Matias. New sampling-based summary statistics for improving approximate query answers. In *Proceedings of ACM SIGMOD*, June 1998.
- [9] A. Gilbert, Y. Kotidis, S. Muthukrishnan, and M. Strauss. QuickSAND: Quick Summary and Analysis of Network Data. Technical report, DIMACS Technical Report, 2001.
- [10] B. Greene. BGPv4 Security Risk Assessment, June 2002.
- [11] B. Greene. Remote Triggering Black Hole Filtering, August 2002.
- [12] B. Greene. Sink Hole Deployments: Is it time to change BGP BCPs?'. IEPG Meeting, San Francisco, March 2003.
- [13] Honeyd: Network Rhapsody for You. [http : //www.citi.umich.edu/u/provos/honeyd](http://www.citi.umich.edu/u/provos/honeyd).
- [14] G. Iannaccone, C. Diot, I. Graham, and N. McKeown. Monitoring very high speed links. In *SIGCOMM Internet Measurement Workshop*, November 2001.
- [15] E. Kohler, R. Morris, B. Chen, J. Jannotti, and F. Kaashoek. The Click modular router, journal = *ACM Transactions on Computer Systems*, volume = 18, number = 3, year = 2000, month = August.
- [16] W. Lee, S.J. Stolfo, and K.W. Mok. A Data Mining Framework for Building Intrusion Detection Models. In *IEEE Symposium on Security and Privacy*, 1999.
- [17] T. Liston. The Labrea Tarpit Homepage. [http : //www.hackbusters.net/LaBrea/](http://www.hackbusters.net/LaBrea/).
- [18] G. Manku and G. Motwani. Approximate Frequency Counts over Data Streams. In *Proceedings of VLDB '02*, August 2002.
- [19] Toby Miller. Passive OS Fingerprinting: Details and Techniques. [http : //www.sans.or/rr - /special/passiveos.php](http://www.sans.or/rr-special/passiveos.php), 2000.
- [20] D. Moore. Network Telescopes. DIMACS Large Scale Internet Attacks Workshop, 2003.
- [21] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. The Spread of the Sapphire/Slammer Worm. Technical report, CAIDA, 2003.
- [22] D. Moore, C. Shannon, and K. Claffy. Code Red: A Case Study on the Spread and Victims of an Internet Worm. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, Marseilles, France, November 2002.
- [23] D. Moore, C. Shannon, G. Voelker, and S. Savage. Internet Quarantine: Requirements for Containing Self-Propagating Code. In *Proceedings of IEEE INFOCOM '03*, San Francisco, CA, April 2003.
- [24] D. Moore, G. Voelker, and S. Savage. Inferring Internet Denial of Service Activity. In *Proceedings of the 2001 USENIX Security Symposium*, Washington D.C., August 2001.
- [25] T. Oetiker. The multi router traffic grapher. In *Proceedings of the USENIX Twelvth System Administration Conference LISA XII*, December 1998.
- [26] V. Paxson. BRO: A System for Detecting Network Intruders in Real Time. In *Proceedings of the 7th USENIX Security Symposium*, 1998.
- [27] D. Plonka. Flawed Routers Flood University of Wisconsin Internet Time Server. [http : //www.cs.wisc.edu/ plonka/netgear - sntp](http://www.cs.wisc.edu/plonka/netgear-sntp).
- [28] D. Plonka. Flowsnscan: A network traffic flow reporting and visualization tool. In *Proceedings of the USENIX Fourteenth System Administration Conference LISA XIV*, December 2000.
- [29] Y. Rekhter. RFC 1817: CIDR and Classful Routing, August 1995.
- [30] M. Roesch. The SNORT Network Intrusion Detection System. [http : //www.snort.org](http://www.snort.org).
- [31] S. Staniford, J. Hoagland, and J. McAlerney. Practical Automated Detection of Stealthy Portscans. In *Proceedings of the ACM CCS IDS Workshop*, November 2000.
- [32] S. Staniford, V. Paxson, and N. Weaver. How to Own the Internet in Your Spare Time. In *Proceedings of the 11th USENIX Security Symposium*, San Francisco, CA, August 2002.
- [33] H.S. Teng, K. Chen, and S. C-Y Lu. Adaptive Real-Time Anomaly Detection Using Inductively Generated Sequential Patterns. In *IEEE Symposium on Security and Privacy*, 1999.
- [34] The HoneyNet Project. [http : //project.honeynet.org](http://project.honeynet.org).
- [35] R. Thomas. Bogon List v2.0 08 April, 2003. [http : //www.cymru.com/Documents/bogonlist.html](http://www.cymru.com/Documents/bogonlist.html).
- [36] S. Thompson. *Sampling, Second Edition*. Probability and Statistics. John Wiley & Sons, New York, 2002.
- [37] V. Yegneswaran, P. Barford, and S. Jha. Global Intrusion Detection in the DOMINO Overlay System. In *Proceedings of NDSS*, San Diego, CA, 2004.
- [38] V. Yegneswaran, P. Barford, and J. Ullrich. Internet Intrusions: Global Characteristics and Prevalence. In *Proceedings of ACM SIGMETRICS*, San Diego, CA, June 2003.