

Computer Sciences Department

**A Framework for Malicious Workload
Generation**

Joel Sommers
Vinod Yegneswaran
Paul Barford

Technical Report #1508

June 2004

UNIVERSITY OF
WISCONSIN

M A D I S O N

A Framework for Malicious Workload Generation

Joel Sommers, Vinod Yegneswaran and Paul Barford

Computer Sciences Department
University of Wisconsin, Madison
{jsommers,vinod,pb}@cs.wisc.edu

1508

Abstract—Malicious traffic from self-propagating worms and denial-of-service attacks constantly threatens the everyday operation of Internet systems. Defending networks from these threats demands appropriate tools to conduct comprehensive vulnerability assessments of networked systems. This paper describes MACE, a unique environment for recreating a wide range of malicious packet traffic in laboratory testbeds. MACE defines a model for flexible composition of malicious traffic that enables both known attacks (such as the *Welchia* worm) and new attack variants to be created. We implement this model in an extensible library for attack traffic specification and generation. To demonstrate the capability of MACE, we provide an analysis of stress tests conducted on a popular firewall and two popular network intrusion detection systems. Our results expose potential weaknesses of these systems and reveal that modern firewalls and network intrusion detection systems could be easily overwhelmed by simple attacks launched from a small number of hosts.

I. INTRODUCTION

Network outages due to self-propagating worms and denial-of-service attacks have been widely reported over the past few years. Despite efforts by the research and operational communities to mitigate these threats, many Internet systems remain vulnerable. One reason for this insecurity is that systems and protocols are often not designed with deliberate consideration of threat models. An example that has received recent attention is the TCP initial sequence number protocol vulnerability that has the potential to inflict chaos by disrupting BGP sessions in the backbone of the Internet [4]. Another reason is that system behavior under heavy load is often unpredictable. Although scaling properties of commercial software systems and routing hardware are quite impressive, they are typically not developed with malicious traffic conditions in mind. Ogielsky *et al.* report how software bugs combined with load introduced by the unprecedented spatial diversity of Code-Red worm traffic and elevated BGP activity led to widespread cascading outages [10].

The potency of Internet worms and viruses has continued to evolve since Code-Red. Examples include high-speed worms such as SQL-Slammer and more recently multi-modal worms/viruses such as Agobot which package exploits for many known vulnerabilities. We have also witnessed the Witty worm which exploits vulnerable network stacks in firewalls. Witty targets all destination ports equally, hence it cannot be neutralized by simple techniques like port blocking at network gateways. These attacks underscore the need for scalable intrusion detection systems to protect large networks by performing signature matching at Internet gateways. They also galvanize the need for better tools for evaluating the resiliency routers, middleboxes and intrusion detection systems.

This work has two primary objectives. The first is to create a performance benchmarking tool that enables assessment of quality of service degradation (the effect of maltraffic on good traffic) and resilience of middleboxes and network intrusion de-

tection systems (NIDS) over a range of malicious traffic volume. The state of the art for benchmarking firewalls and NIDS is still considered nascent, and there are no universally accepted standards. Current best practice is to gather and use malware itself. This approach is limited in its ability to scale traffic volume and compose as yet unseen attack traffic. A flexible research-oriented toolkit for producing maltraffic is therefore a critical component for the development and standardization of benchmarks for NIDS and firewalls. Our second objective is to recreate attack traffic scenarios in a laboratory setting to test protocol suites, software, and hardware systems. This capability will enable us to move toward quantitative measurement and assessment of claims regarding the degree to which individual network devices contribute to large-scale correlated failures.

This paper describes the design and implementation of a flexible, extensible toolkit that called the Malicious Traffic Composition Environment (MACE). MACE provides the basic building blocks for recreating a large set of known attacks. MACE satisfies the following requirements:

1. Generate a large, diverse set of attacks,
2. Generate and control benign background traffic,
3. Compose attacks in a high-level language.

The MACE model decomposes attacks into three components: exploit, propagation and obfuscation. To our knowledge, MACE is the first tool to adopt an extensible approach to systematic attack synthesis through a modular attack composition framework. MACE enables resiliency of systems to resource exploits and traffic mix to be evaluated. MACE is distinguished from related efforts like Thor which focus on obfuscation methods and individual session morphing [5]. Results from such work can be easily folded into MACE.

We discuss the MACE framework in greater detail and provide example attack traffic configurations in Section III. Section IV demonstrates MACE by providing a case study of a popular firewall (a Cisco PIX) and two NIDS (Snort and Bro) under varying traffic mixes. Our results reveal that modern firewalls and NIDS could be easily overwhelmed by simple attacks launched from a small number of hosts. In most instances, the resilience of the devices to particular attacks varies with respect to the degree of connection state maintained by each device. The case study also illustrates the capabilities of MACE as a performance benchmarking tool.

II. RELATED WORK

Internet traffic generation for measuring application and network device performance has been well-studied. Harpoon reproduces network traffic in an application-oblivious manner [15]. Numerous application-aware traffic generators like

SURGE produce workloads to stress-test web servers [9]. These tools are complementary to MACE and can be used as plug-ins that generate legitimate (benign) background traffic.

Taxonomies of malicious traffic inform the design and development of the MACE attack types and propagation models. One such taxonomy of DDoS attack characteristics is provided in [13]. A similar taxonomy of Internet worms based upon target discovery, carriers, activations, payloads and attackers is proposed in [16]. Our study is also related to the work by Lippman *et al.* which provides a dataset for evaluation of NIDS [12].

Our laboratory measurement of the performance of a firewall middlebox is complementary to the work of Allman, who measured the performance of an operational middlebox infrastructure in [7].

Two other analogs to our work are Mucus [14] and Nessus [3]. Mucus is a tool for black-box testing of NIDS systems. As a benchmarking tool for NIDS systems, Mucus' evaluation criteria are *alert correctness* and *alert quality*. Nessus is a tool for penetration testing of network hosts. MACE differs considerably from both of these tools in its objectives. The MACE framework for malicious traffic generation accommodates its use both for penetration testing or as a NIDS benchmarking tool. Our goals of performance benchmarking NIDS and middleboxes do not require a complete database or accurate replication of specific attacks. While it is unnecessary, for example, to faithfully replicate all variants of the Beagle worm, it is desirable to have representatives from all classes of attacks. Finally, since the MACE model isolates obfuscation and propagation from exploit signatures, it can generate attacks that are not present in current vulnerability databases.

III. MACE FRAMEWORK

We begin by defining the abstract model for MACE. We then describe implementation details of the MACE toolkit with examples of attack instances and show how they can be composed.

A. Abstract Model Definition

The conceptual model for MACE is illustrated in Figure 1. It provides flexibility in specifying the base characteristics of malicious traffic which we define as the following:

- **Exploit Model** — set of vulnerabilities that are part of the attack sequence.
- **Obfuscation Model** — morphs in the header or payload to enable the exploit to elude NIDS. These could either be at the network layer or at the application layer.
- **Propagation Model** — order in which victims are chosen to be attacked.
- **Background Traffic Model** — legitimate traffic flow in the network.

B. Model Realization

The modular MACE library consists of the exploit, obfuscation, and propagation components defined by the model, as well as a number of functions to support interpretation, execution, and exception handling of attack profiles. Our objective is not to provide a complete attack database, but rather to support a set of basic building blocks that can be used to create both known and custom attack vectors.

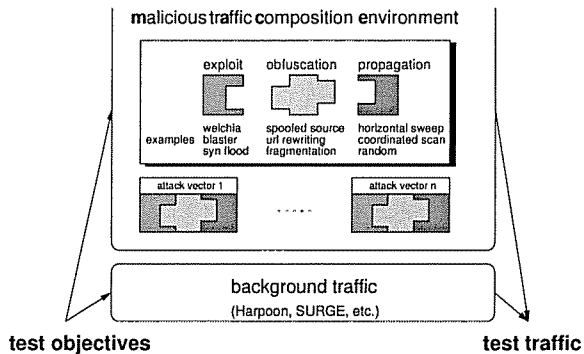


Fig. 1. The MACE architecture. Test objectives inform the selection and composition of exploit, obfuscation, and propagation components into a series of attack vectors. Existing tools such as Harpoon produce the desired benign background traffic.

MACE is currently implemented in the Python programming language. The dynamicism of a language like Python is important for MACE as it enables sequences of exploits or obfuscations to be fabricated on-the-fly. Python has libraries that support many application layer protocols, such as HTTP, NNTP, and SMTP, allowing application-specific exploits to be easily created. For evolving scalability requirements, we plan to migrate key components of the library to C, extending the capabilities of the Python interpreter much in the way ns-2 [6] extends Tcl/Tk.

The building blocks of MACE include the following:

Payload Construction Elements - Payload elements in MACE are defined as character arrays. In practice, these are often payloads from various higher level protocols such as HTTP, Net-BIOS/SMB or DCE RPC. The following example defines an HTTP GET request for the file index.html.

```
[httprequest, ( 'method':'GET', 'absolute_path':'/index.html' )]
```

Header Construction Elements - Attack traffic often requires the use of raw sockets to construct custom TCP, UDP and IP headers. The header construction elements in MACE modify specific attributes without exposing the entire header to the user. The following example illustrates the definition of a TCP SYN packet. The source and destination IP addresses and ports are defined by the propagation elements. Other unspecified fields are appropriately filled with default or random values.

```
[rawtcp, ( 'th_flags':'TH_SYN' )]
```

Obfuscation Elements - The obfuscation elements can be specified at the network layer or at the application layer. An example of network layer obfuscation is IP fragmentation. Examples of application layer obfuscation include HTTP URL encoding techniques such as Bare-byte or Unicode encoding used by worms such as Nimda [8]. The following example applies a Unicode obfuscator to the example HTTP request given above:

```
[ httprequest, ( 'method':'GET', 'absolute_path':URLObfuscator.uencode('/index.html') ) ]
```

Propagation Elements - MACE supports different propagation models via AddressPool objects. Each AddressPool is instantiated with a list of CIDR prefixes and port ranges, along with an indication of how to generate addresses from the given pool. Address pools may be traversed randomly, horizontally (sequentially across IP addresses, sequentially across ports),

vertically (sequentially across ports, sequentially across IP addresses), or by a more complex methodology definable by the user. The following example illustrates an address pool defined to perform a vertical sweep of the target IP address 10.42.1.1.

```
target_pool = AddressPool(AddressPool.Vertical, '10.42.1.1:1-65536')
```

C. Example Test Scenarios

We focus on four attack vectors:

- *SYN flood*: A standard denial-of-service attack. TCP SYN packets are sent from the attacker but no more packets are sent.
- *Welchia*: A ping followed by a series of HTTP requests designed to exploit a buffer overflow in the WebDAV module of Microsoft’s IIS web server [1].
- *Rose*: An attack that exploits poorly implemented handling of IP fragments. Two small packet fragments are sent, one with an offset of 0 and one with a large offset. Some network stacks reserve memory for the fragment hole, so a series of Rose packets can exhaust memory [11].
- *Blaster*: An attack that exploits a buffer overflow in Microsoft Windows RPC service (epmapper) [2].

The four attack profiles are realized in the example MACE code fragment shown in Figure 2. Attack profiles are defined by their vulnerability exploit and propagation method. A full exploit is a sequence of *generator* and *validator* steps, along with parameters. A generator builds packet traffic using the payload and header construction and obfuscation building blocks in MACE. A validator collects and processes responses from the attack target, verifying that the generated traffic evoked an appropriate response. Each step in an attack vector is executed as long as they are successful. An exploit step may simply be “create a TCP packet with the SYN flag set” (line 1). Other steps might be as complicated as “create an HTTP GET request for the document ‘/’ and validate that the HTTP response contained a 200 (success) code and that it was produced by a Microsoft web server” (lines 6–7). Some generator steps are not followed by validation, *e.g.*, the Rose attack where validation is unnecessary (lines 12–15). Since generators and validators are simply Python functions, it is easy to define new exploit steps. Propagation models for the four attack profiles are defined by the address pools in lines 23–25. Finally, attacks are sent using the `send_once` or `send_periodic` functions (lines 27–28). `send_periodic` is a convenience function used to produce attack traffic for a given duration (delaying a specified amount of time between successive call to `send_once`).

IV. SECURITY SYSTEM PERFORMANCE EVALUATION

To demonstrate MACE’s capability, we examine performance characteristics of three standard network security systems: a firewall middlebox and two network intrusion detection systems running on commodity hardware.

A. Test Environment

The firewall we test is a Cisco PIX 515e. It contains three Fast Ethernet interfaces, 64MB of RAM, a 433 MHz Intel Pentium II, and is running version 6.2(2) of the PIX Firewall operating system. It is a typical device deployed as the first line of defense

```
synflood = [ [ rawtcp, ( 'th_flags':TH_SYN ) ] ]
bad_string = '...' # the buffer overflow string - not defined here
welchia = [
    ( ping ),
    ( httprequest, ( 'method':'GET' ),
      httpvalidate, ( 'Server':'Microsoft-IIS/5.0' ),
    ( httprequest, ( 'method':'SEARCH' ),
      httpvalidate, ( 'code':411, 'Server':'Microsoft-IIS/5.0' ),
    ( httprequest, ( 'method':'SEARCH', 'absolute_path':bad_string ) ] ]
rosepayload = '\0' * 32 # just a small fragment
rose = [
    ( rawudp, ( 'frag_offset':0, 'frag_flag':IP_MF, 'payload':payload ),
    ( rawudp, ( 'frag_offset':8100, 'payload':payload ) ) ] # 64800 byte offset
bindreq = '...' # the DCE RPC bind request - not defined here
overflow = '...' # the buffer overflow exploit request - not defined here
blaster = [
    ( dcerpcbind, ( 'bindreq':bindreq ), dcerpcbindack, ( 'bindreq':bindreq ),
    ( dcerpcreq, ( 'payload':overflow ) ] ]
src_pool = AddressPool(AddressPool.Vertical, '10.42.0.0/16:1-65536')
dst_pool = AddressPool(AddressPool.Random, '10.52.128.0/20:1024-65536')
http_pool = AddressPool(AddressPool.Horizontal, '10.52.128.0/20:80')
send_once(src_pool, dst_pool, rose)
send_periodic(src_pool, http_pool, welchia, duration=60, delay=0.001)
```

Fig. 2. A MACE code fragment, similar to what was used in our experiments.

for edge networks, implementing basic packet filtering and performing network address translation (NAT). The network intrusion detection systems are Bro (version 0.8a79) and Snort (version 2.1.1). Bro generally maintains significant connection state, while Snort does not. Each NIDS runs on a separate workstation with a 2 GHz Intel Pentium 4 processor and 1 GB of RAM. FreeBSD 5.1 is installed on each machine¹. For each NIDS, we use a default set of rules. Our Snort instance uses the default `snort.conf` included with the software distribution and Bro uses the `mt` policy.

The PIX and NIDS are configured in a testbed as shown in Figure 3. The setup mimics an edge network connected to an ISP, with legitimate background traffic and attacks focused on a remote network. The PIX resides between the edge (“internal”) network and the ISP (“external”) network. The internal network contains traffic generators for MACE and background traffic, and the two NIDS. All network traffic received from or sent to the PIX is duplicated on the links connected to the NIDS. In the external network, we use a hardware propagation delay emulator (Adtech SX-14) between two backbone-class routers (Cisco 12000) to create a round-trip time of roughly 100 milliseconds between the traffic generators and the target servers. We use popular enterprise-scale switches (Cisco 6500) to aggregate traffic at the endpoints.

On each host in the internal network, we create 2^{12} alias addresses and on the remote hosts we create 2^{24} aliases. The PIX performs network and port address translation between hosts on the internal network and a pool of 2^{24} addresses routable across the external network. It is important to note that the PIX performs an implicit packet filtering based on its NAT configuration: it only performs NAT or port address translation (PAT) for local addresses that are part of its configuration. Packets from any other source address are dropped.

Using two levels of benign background traffic, we generate attack traffic using a set of five exploits and six levels of offered load. Two levels of background traffic, “low” and “high”,

¹On each host we modify the kernel parameter `debug.bpf_bufsiz` from 4096 bytes to 524,288 bytes, as suggested in the Bro documentation. Snort presumably can benefit from this change as well so we apply the change to each NIDS host.

are generated using Harpoon [15] and are tuned to averages of 20 Mbps and 70 Mbps, respectively. Source and destination addresses for legitimate traffic are chosen randomly from the pool of 2^{12} source addresses and the pool of 2^{24} destination addresses. The six levels of attack load are generated by using one to five hosts running MACE. For each exploit, the MACE processes on a single host are configured to generate roughly 1 Mbps of traffic, regardless of background traffic level.

The exploits we use are (1) Welchia worm traffic, (2) SYN flood denial-of-service attack, (3) a SYN flood with spoofed source addresses, (4) the Rose fragment attack, and (5) a multimodal attack consisting of the previous four exploits plus Blaster worm traffic. Each host running MACE uses a source address pool of 2^{12} , as described above. For the SYN floods and Rose attack, we horizontally traverse the source address space. For the two worm exploits, we randomly (uniformly) sample the source address space. All attack traffic is directed toward a single address on the remote network.

B. Test Measurements

For the 52 distinct experiments, we measured CPU and memory utilization at all three systems every five seconds. We also measured packet counts in and out of the PIX every five seconds and the number of reported packet drops using SNMP, and took packet traces on either side of the PIX. At the two NIDS hosts, we verified and used the capabilities of each software package to report received packet volume and the number of dropped packets. Packets are dropped by each system due to overflow of the queue of incoming packets received by the packet filter (each NIDS system uses the Berkeley Packet Filter and libpcap for packet capture). For packet drops at each NIDS, we did not discriminate between benign and malicious traffic. For the PIX, we used the packet traces to measure benign packet drops. Each experiment was run for six minutes, including a one minute warm-up phase from which measurements are discarded.

C. Results

Figure 4 shows average CPU utilization and Figure 5 shows packet loss measurements for each experiment. The two rows in each figure correspond to low and high background traffic levels, and the three columns display results for each device. The first feature to notice in the plots is the diversity of responses of each system to distinct MACE attack profiles. Except for a few cases, there is also a noticeable, and sometimes very large, divergence between the first two data points. These points correspond to zero malicious traffic and a single MACE host. We discuss detailed results for each device class (firewall and NIDS) below.

C.1 Effects on the PIX firewall

For the PIX, the Rose attack has the least effect on performance. When processing fragmented packets, the PIX keeps a queue of (by default) 200 fragments in order to reassemble them before forwarding them to the remote network. If the missing fragments do not arrive in a configurable amount of time (the default is 5 seconds), the fragments are dropped. For the Rose attack, the fragment queue becomes full shortly after starting MACE. When fragments arrive that cannot be queued, they are

dropped. Although there is no fragmented legitimate traffic in our setup, these packets would very likely be dropped even with an attack rate of just over 40 Rose packets per second. Also, queues for each interface and the fragment reassembly queue share a common buffer pool. Since there are fewer buffers available during a Rose attack, interface queues are more likely to fill, causing additional packet drops².

The non-spoofed SYN flood is the attack with the most impact on the PIX. Since the PIX is performing NAT, it must maintain state for each connection. Even with a single MACE host, all 64 MB of system memory is used after a short while because of this state requirement³. While there is still some memory available for buffering packets as they flow through the system, this memory pool is now much smaller and consequently, the PIX has diminished ability to absorb bursts of packets. This situation does not occur with the spoofed SYN flood, since the source addresses do not conform to the NAT configuration at the PIX and are dropped. In our traces, we see persistent dropping of legitimate packets during the non-spoofed SYN flood and multimodal experiments (in both low and high background traffic regimes) and aggregate traffic rates through the PIX clearly show the well-known poor performance of TCP in the face of such high packet loss: for the low background SYN flood experiment using only one MACE host, the inbound (external to internal) packet rate through the PIX for the spoofed SYN flood is nearly twice that of the non-spoofed SYN flood. For experiments with more than one MACE process, the results are more dramatic. In summary, maintaining state and managing system resources under a low rate non-spoofed SYN flood is difficult even for a specialized device. Considering the rate at which resources consumed by the embryonic connections are reclaimed by the PIX, there is probably a “sweet spot” at which SYNs can be sent at a low enough rate to cause problems for good traffic, but are not at a high enough rate to easily detect.

The Welchia and Blaster worms, as with other worms, are typically short flows, so even with a low attack rate the primary effect on the PIX is an increased rate of connection initiations. To the PIX (without any special packet filters installed) these worms look like benign traffic and are treated the same as all other legitimate packets. In our experiments, the PIX appears to be sufficiently provisioned to handle the increased volume.

Finally, it is interesting to note that while all Rose and spoofed SYN flood packets are dropped by the PIX, these attacks have peculiar effects on CPU usage. Without detailed internal information of the PIX, we can only surmise the path for handling IP fragments is significantly faster than the process of matching a source address to the NAT configuration at the PIX (though, as pointed out above, there is a potential denial-of-service problem related to fragment processing).

Summary: (1) Non-spoofed low-rate SYN floods are effective

²The PIX documentation notes that it is possible to set the fragment reassembly maximum queue length to be equal to the total number of buffers available. The documentation warns that such a configuration would enable fragment attacks to be effective denial-of-service attacks.

³There are timeouts defined in the PIX to reclaim resources used by idle and half-closed connections, and a feature called “floodguard” which essentially prioritizes which resources to reclaim in order to deal with resource shortages. Our experiments are shorter than either one of the default values for these timeouts (1 hour for idle connections, 10 minutes for half-closed connections).

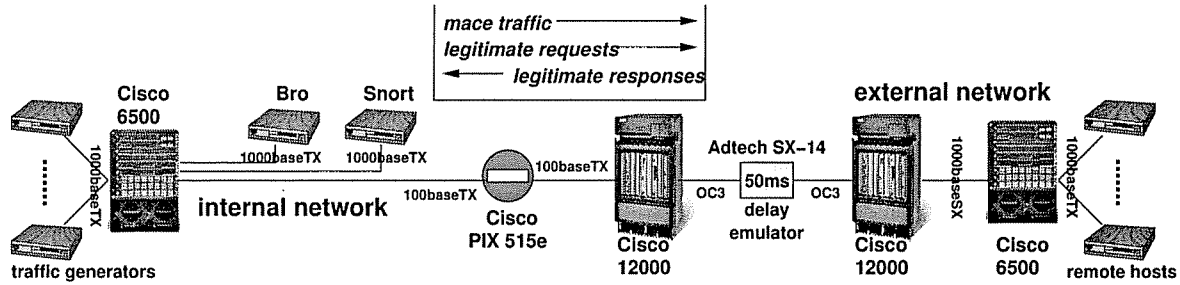


Fig. 3. Experimental Environment. Legitimate and malicious traffic originate from the internal network and are directed toward a host on the external network. The PIX firewall separates the inside and outside networks and performs network address translation.

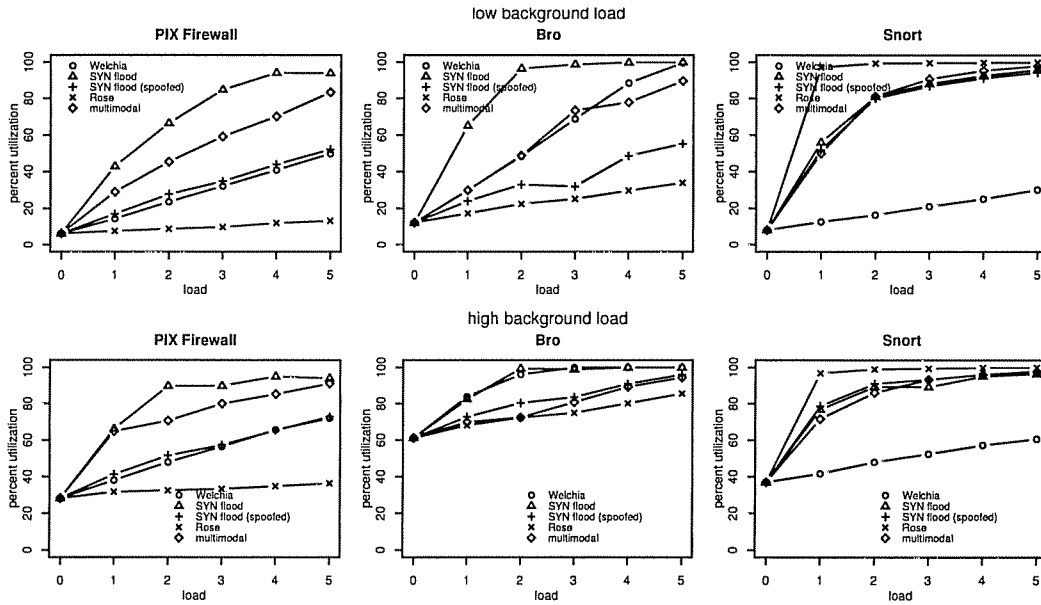


Fig. 4. CPU utilization measurements. Results for the PIX, Bro and Snort are in left, middle, and right columns, respectively. Top and bottom rows contain results for low and high background traffic loads, respectively. Load levels along the x-axis correspond to number of MACE hosts used in each test. Each MACE host generates roughly 1 Mbps of traffic regardless of attack or background traffic level.

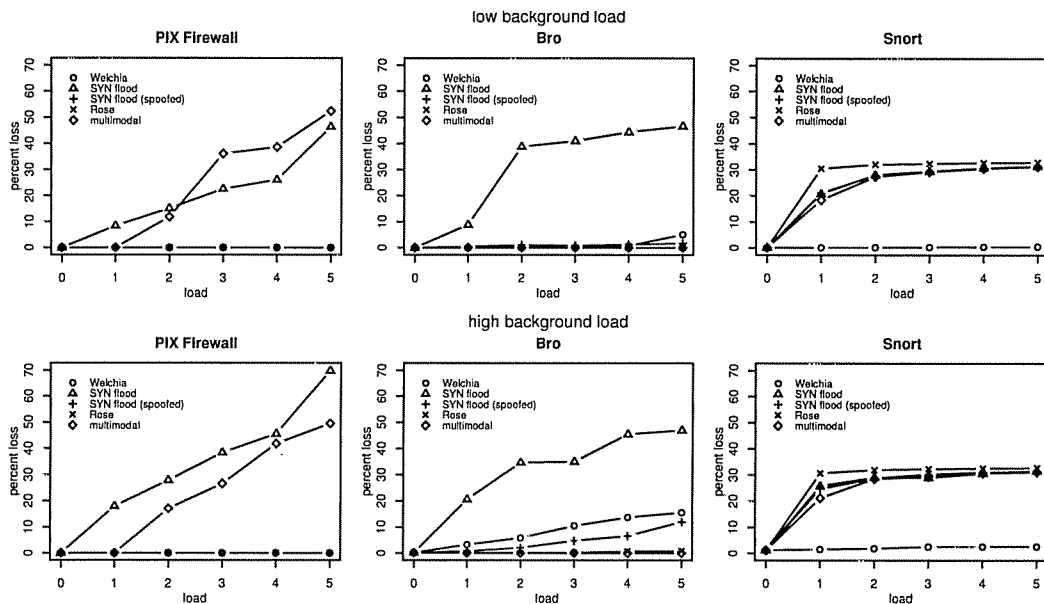


Fig. 5. Packet loss measurements. Results for the PIX, Bro and Snort are in left, middle, and right columns, respectively. Top and bottom rows contain results for low and high background traffic loads, respectively. Plots for the PIX show packet drops for benign traffic only, while plots for Bro and Snort show aggregate packet drops. Load levels along the x-axis correspond to number of MACE hosts used in each test. Each MACE host generates roughly 1 Mbps of traffic regardless of attack or background traffic level.

resource exploits leading to significant service degradation, and (2) Obfuscations via packet fragmentation are effective resource exploits against poorly configured systems.

C.2 Effects on Bro and Snort

There is a sharp contrast between behavior of Bro and Snort. For example, the Rose attack has little effect on Bro, but an enormous impact on Snort. Since Bro and Snort receive raw packets from the network interface, they must perform reassembly of fragments. Bro is clearly able to handle fragments more efficiently than Snort, even though Snort maintains much less overall state than Bro.

For Bro as with the PIX, the non-spoofed SYN flood has the greatest impact on CPU load and packet loss. For the non-spoofed SYN flood, Bro maintains state for all attack-induced connections. Although Bro periodically expires idle connection state, the rate of SYNs in our experiment was high enough that Bro eventually exhausted available memory⁴. Since the table of known connections continues to grow during this attack, connection state lookups are more costly. The resulting effect on packet drops is clearly shown in the center column graphs of Figure 5. The reason the spoofed SYN flood has relatively little impact on Bro is that the PIX silently blocks the spoofed SYNs so Bro will never see a SYN/ACK response. Apparently this lack of response allows Bro to flush the embryonic connection state in a more efficient manner.

Except for the Welchia attack, Snort performs similarly under all attacks. The SYN floods, Rose, and multimodal attacks each contain packet-level attacks in contrast to Welchia, which (at least from the perspectives of IP and TCP) looks like the legitimate background traffic to Snort⁵. Efficiently processing ill-formed packets and pathological packet sequences is clearly a requirement and a challenge for NIDS.

With respect to packet loss, Bro and Snort again exhibit contrasting behavior. Except for the Welchia attack, Snort consistently drops roughly 20-30% of all packets once MACE traffic is introduced. Bro, despite maintaining significant connection-level state, drops a relatively small portion of packets except for the non-spoofed SYN flood attack. For both NIDS, any significant level of packet dropping will clearly affect the ability of the tool to detect ongoing attacks. Additionally, knowledge of packet dropping behavior could be exploited by an attacker to launch a relatively benign low-rate attack in order to mask one that is more insidious.

Summary: (1) Multiple attack vectors are effective resource exploits leading to packet loss, implying degraded detection rates, and (2) The marginal impact of resource exploits does not appear to be greater for NIDS that maintain connection state.

V. CONCLUSIONS AND FUTURE WORK

The escalation of malicious activity in the Internet motivates the need for better tools to measure the resiliency of routers and

⁴We also experimented with the `reduce-memory` policy script with Bro, which caused an increase in CPU usage and higher packet loss rates for all attacks.

⁵In our experiments we used the default configuration of Snort, in which there are no specialized rules for processing HTTP traffic. Enabling these rules causes higher CPU load and packet drops for all attack profiles.

middleboxes to malicious traffic. To address this need, we propose MACE, a framework for malicious network traffic generation. The MACE architecture is composed of three building blocks: exploits, obfuscators, and propagation elements. These components define and create malicious traffic for use in laboratory testing of routers and network security infrastructure. We provided experimental results of measurements conducted on a popular firewall and two network intrusion detection systems to document the varying responses of these systems to malicious traffic. Our results show that relatively low rates of attack traffic can exploit the overheads of maintaining connection state or inefficient processing of certain packets.

A tool like MACE can be used for testing and refining the operation of network systems, but if in the wrong hands, could be used for generating malicious traffic in the live Internet. Our plan for making MACE available to a wider community is to supply the code only to legitimate researchers and, to the best of our ability, keep careful documentation regarding who has copies of the code.

We plan to expand the list of exploit and obfuscation building blocks within MACE and to make improvements to the volume of exploit traffic that MACE is able to produce. We believe these enhancements will facilitate the laboratory emulation of large-scale failure scenarios using more elaborate physical and logical topologies and a greater diversity of network devices.

REFERENCES

- [1] Microsoft Security Bulletin MS03-007. <http://www.microsoft.com/technet/security/bulletin/MS03-007.mspa>, 2003.
- [2] Microsoft Security Bulletin MS03-026. <http://www.microsoft.com/technet/security/bulletin/MS03-026.mspa>, 2003.
- [3] Nessus. <http://www.nessus.org>, 2004.
- [4] NISCC Vulnerability Advisory 236929. <http://www.uniras.gov.uk/vuls/2004/236929/>, 2004.
- [5] THOR: A Tool to Test Intrusion Detection Systems by Variations of Attacks. <http://thor.cryptojail.net/>, 2004.
- [6] UCB/LBNL/VINT Network Simulator - ns (version 2). <http://www.isi.edu/nsnam/ns/>, 2004.
- [7] Mark Allman. On the Performance of Middleboxes. In *Proceedings of Sigcomm Internet Measurement Conference*, 2003.
- [8] Eugene J. Aronne. The Nimda worm: An overview. <http://www.sans.org/rr/papers/36/95.pdf>, 2001.
- [9] Paul Barford and Mark Crovella. Generating Representative Web Workloads for Network and Server Performance Evaluation. In *Proceedings of ACM SIGMETRICS*, 1998.
- [10] James Cowie, Andy Ogielsky, BJ Premore, and Yougu Yuan. Global Routing Instabilities Triggered by CodeRed II and Nimda Worm Attacks. <http://www.renesys.com/projects/bgpinstability>, 2001.
- [11] Gandalf. IP Fragmentation --> The Rose Attack. <http://www.securityfocus.com/archive/1/359144>, 2004.
- [12] Richard Lippmann, David J. Fried, Isaac Graf, Joshua W. Haines, Christopher 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.
- [13] Jelena Mirkovic and Peter Reiher. A Taxonomy of DDoS Attack and DDoS Defence Mechanisms. In *Proceedings of Computer Communications Review*, 2004.
- [14] Darren Mutz, Giovanni Vigna, and Richard Kemmerer. An Experience Developing an IDS Simulator for the Black-Box Testing of Network Intrusion Detection Systems. In *Proceedings of ACSAC*, 2003.
- [15] Joel Sommers, Hyungsuk Kim, and Paul Barford. Harpoon: A Flow-Level Traffic Generator for Router and Network Tests. In *Proceedings of ACM SIGMETRICS*, 2004.
- [16] Nicholas Weaver, Vern Paxson, Stuart Staniford, and Robert Cunningham. A Taxonomy of Computer Worms. In *Proceedings of CCS Worms*, 2003.