

A SIMPLE FRAMEWORK FOR FILTERING QUEUED SMTP MAIL (CYBERWAR COUNTERMEASURES)

Tim Bass
Science Applications International Corporation
Center for Information Protection
McLean, Virginia

Lt. Col. Glenn Watt
United States Air Force, Air Combat Command
Langley AFB, Virginia

ABSTRACT

Pre-information age military battlefields are based on the traditional land, sea, air, and space paradigm. Global internetworking is causal to the creation of a dangerously real *5th Dimension of Warfare - Cyberspace*. This paper describes an Internet based assault, commonly referred to as *e-mail spam*, on the Langley AFB internetworking infrastructure. We discuss the cyber-attack, a framework for defending against the attack, and the results of the campaign. The countermeasure was accomplished by running the MTA in a mode which accepts and queues SMTP mail; processes the messages with a rules-based filter; and then forwards mail after filtering. The filtering framework is simple and effective for a large subset of e-mail bombs. The prototype filter scripts may be obtained from the authors.

INTRODUCTION

E-mail is extensively used for the exchange of numerous types of messages in the military environment. From low value questions and staff summaries to messages bordering on command and control, in every case the recipient assumes the apparent originator is the actual person sending the message. It is often overlooked, however, how trivial it is to impersonate a user or masquerade behind mail relays to send forged SMTP messages and mail bombs.

Our cyberwar began when we assumed the moral high-ground and began stopping hackers from using Air Force SMTP relays for the distribution of pornographic and bigoted hate mail. The hackers extracted their revenge by launching an e-mail attack of epic proportions; and statistics indicated that over 70 percent of the e-mail imprisoned during the cyber-campaign was spam. During many periods on the Internet battlefield, approximately 30,000 e-mail messages were captured per day. All imprisoned messages were either pornographic, malicious, or bigoted hate-mail in nature.

This paper summarizes the virtual battle and the successful countermeasures implemented during the SMTP cyber-attack. Our solution was a rules-based filtering

utility that, within 48 hours of implementation, shunted the attack and frustrated our international Internet opponents. A summary chronology follows:

Jan 5	Director receives first forged e-mail
Jan 21	Sendmail logging level increased
Jan 23	SMTP prototype filter completed
Jan 27	Filter report identifies large problem
Feb 14	USAF assigned configuration control
Mar 04	SMTP mail relay crashed
Mar 12	SMTP mail relay crashed by DOS attack
	ACC forms Tiger Team
	Commander Lt Col Watt
	Technical Management Maj Gruber
	Chief Scientist Mr. Bass
	Software Engineering Capt Fish
	Engineering Support Lt Baker
Mar 12	Tiger Team repels first wave
Mar 14	AFCERT Team visits Langley AFB
Mar 18	Analysis of jail queue and logs
Apr 03	Prototype filter enhanced by USAF
Apr 09	USAF coins phrase BOMBSHELTER
May 05	Hackers remove AFB from attack list

PRELUDE TO WAR: A FORGED E-MAIL

On Monday, 5 January 1997, 0830 hours, one of the authors received a phone call ordering him to report to his director's office. From the tone of the conversation something serious happened. The director was the latest recipient of malicious e-mail impersonating the sender, "*clinton@whitehouse.gov*."

The forged e-mail was inflammatory in nature and bordering on threatening, igniting an immediate quest to determine the source of the attack. The perpetrator used a widely exploited hole in the SMTP protocol requiring open access to port 25. The appendix hosts a brief summary of the well known SMTP spoofing technique.

At Langley AFB, a covert design effort was initiated to develop an e-mail filter to capture and process malicious and criminal spam. The chief scientist of the team developed a prototype SMTP filtering program and architecture (Fig. 1) which initially identified 586 malicious

or bogus SMTP messages in the first 41 hour sampling period. This, however, was only a small indication of the severity of the e-mail spam problem uncovered on SMTP host installations.

The initial basis for identifying spam was a rules-based filter which detected the absence of the '@' character in the SMTP sender address field. Below are examples of the first e-mail captured by the filter, arranged by the forged [sender], [content], and [number of messages]:

- alena* sexually explicit material, 138 messages;
- borwig* no content, 11 messages;
- aaaa* a test, 4 messages;
- goldie* a test, 5 messages;
- webmaster* prank message, 7 messages;
- doody* prank message, 20 messages;
- f__ you __hole* prank message, 1 message;
- steve case* forged AOL prank, 86 messages;
- Hockey God* prank message, 165 messages;
- organizer* sexually explicit material, 53 messages
- Concerned* student politics, 10 messages; and
- Hoo Wah!* prank message, 85 messages.

Unknown to most Internet e-mail administrators, SMTP mail servers are covertly used as platforms to relay malicious and criminal e-mail to users outside the intended domain. The developed SMTP filter, *smtpfilter.pl*, captured or copied all suspected mail during the relay process and stored complete message content, including SMTP header information. We discovered that the number of spoofed e-mail addressed directly to recipients at Langley was small relative to mail covertly relayed to the rest of the world by hackers via the Langley SMTP servers.

The original spam count, prior to the cyberwar escalation, was approximately 700 messages per day. Most of these e-mails contained sexually explicit, anti-Semitic or other unacceptable message content. E-mail hackers were relaying mail via unprotected USAF SMTP servers; in essence, creating Grand Central Stations for pornography and hate-mail. Pedestrian attempts to trace the e-mail back to the 'originator' would falsely point to military establishments, creating numerous opportunities for media-based perception warfare.

Forty-eight hours after the original SMTP prototype filter was operational, hacker bulletin boards were reporting problems with the targeted mail relay. E-mail spam and mail bombs were no longer being successfully relayed through Langley as intended by hackers. At this point, our adversaries launched numerous mail bomb attacks at the SMTP relays. At one point over 30,000 captured messages per day were received. This number would have been higher; however the steady-state congestion of the DoD Internet served as a pseudo cyber-buffer.

The countermeasure was to enhance the prototype SMTP filter to process a broader technical range of e-mail spam. A virtual Internet cat-and-mouse game of countermeasures and counter-countermeasure occurred as different rule-sets were implemented. The cyberwar had begun.

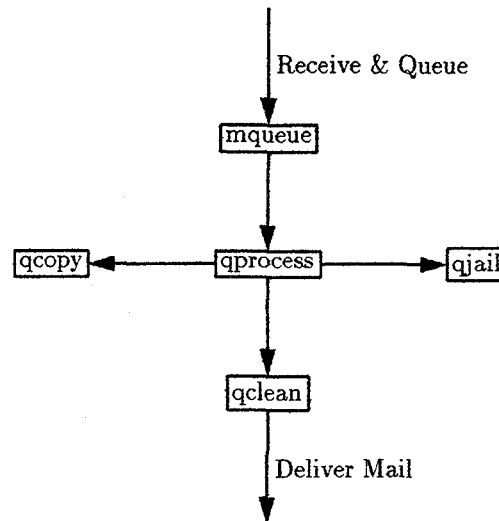


Fig. 1. Process Flow Diagram for SMTP Filter

CYBERWARS

Defensive information warfare has traditionally engaged a strategy advocating intrusion detection and limited prevention in response to covert activities designed to affect information systems. These activities generally include malicious software designed to activate after penetration and destroy, manipulate or retrieve official information.

Our invisible Internet adversaries employed a far more simple and effective technique during Langley's cyber attack. In this battle, opponents used legitimate network resources as a precision guided weapon. E-mail was the specific weapon of choice and SMTP was the delivery agent. These SMTP relays receive mail from other Internet mail sites with minimal, if any, restrictions. Sendmail configuration file restrictions often delete forensic evidence or generate SMTP reject messages. Neither of these two results are acceptable options for highly successful and covert cyber countermeasures.

During the initial filter prototyping phase, we copied and delivered all mail with the keyword 'whitehouse' in the header fields because it was theoretically possible that valid mail could come from "whitehouse.gov." This was the prototype of a queue which would become *qcopy*, which would be utilized to help refine additional filter rule-sets. All captured mail that was *taken prisoner* was stored in the jail queue, *qjail*.

The prototype filter design also provided valuable information which assisted investigators discover the content type and origin of the spoofed mail. Trapping and jailing politically charged mail became a high priority as we uncovered pornography being relayed from servers on bases to pedestrian users in the commercial Internet.

Concerns also mounted about the potential for media based information warfare. If the news media begin publishing inaccurate articles with headlines including forged AFB e-mail addresses (and unknown to the public or the media) with controversial or politically charged content,

the reputation of USAF could be severely damaged:

”USAF Distributing Pornography to AOL”

– Hypothetical Media Hype

Covertly distributing pornography could cause unintended adverse side-effects via the news media which may or may not have been the primary objective of the information warrior. Early recognition of the potential damage by adverse public opinion created the overarching strategy as the first heavy wave of Internet spam was shunted.

The Langley AFB mail system averages 5000-6000 files of legitimate SMTP mail daily, most arriving over a 10 hour period. On the 11th of March and again on the 12th of March 1997, one mail server received approximately 25,000 messages within 8 hours. The overwhelming mass of spam brought the legacy 486 based SMTP mail relay to a complete halt. The Internet based cyber-attack was in full force and spam was targeted to an America On Line (AOL) customer, using a USAF post office as a relay point. A casual observer would come to the wrong conclusion that the spam originated from Langley AFB and not the true foreign origin.

In the early hours of the campaign, we observed that a majority of e-mail spam contained spoofed addresses without the '@' sign similar to the original SMTP spoof, described earlier. Filtering and jailing messages with invalid SMTP address format fortified our defenses and repelled the first wave of the attack.

The hackers quickly adapted, created bogus address with random '@' symbols and were actively probing Langley AFB mail relays, to understand what countermeasures were being implementing. An analysis on the 11th of March indicated the majority of the spam came from sites in Estonia and Australia. Considering this information, we shunted all network activity from those sites by a packet-filtering router located in front of the mail server.

Again, our adversary adapted within 24 hours and begin relaying spam from other DoD, US Government, and commercial sites with a technique we refer to as *chain bombing*. We could not shunt this attack using packer-filtering routers without denying critical SMTP services to legitimate users. The hackers also started employing more sophisticated techniques not generally associated with amateurs, including:

1. Direct E-mail Spam,
2. Indirect E-mail Spam,
3. Rejected E-mail Spam,
4. Chain Relay E-mail Spam, and
5. Mailing List Spam.

Direct spam came from the hacker's host system to the target via the port 25 hack described earlier. Indirect spam is e-mail bounced off another site needed to conduct business. Reject spam results from massive volumes of mail sent to undeliverable addresses at Internet Service Providers (ISPs). The undeliverable messages generated rejection notices that were returned to the forged sender, Langley AFB, or the reverse.

Chain relay spam, or chain bombing, works by linking together the host addresses of SMTP relays, telling SMTP the exact path or chain to follow in transferring e-mail. This is analogous to locust swarming from bush to bush, destroying vegetation, along the path. In fact, some of the popular mail bomb tools, e. g. *Avalanche*, *KaBoom*, and *UpYours*, had Langley's mail relay hard coded into the software. The final technique involved spam via electronic mailing lists. Hackers simply sign up a site to numerous electronic mailing list exploders in order to increase the flood of e-mail traffic.

A collective strategy session produced three critical decisions that successfully mitigated the cyberwar:

1. Acquire Processor Maneuvering Room
2. Continue to Filter and Jail Mail
3. Train Operators On Attack ID and Response

First, we replaced a legacy 486 based relay with a high powered Pentium server. The success of the remaining strategy hinged on the requirement to process all the spam the available network bandwidth could deliver with CPU cycles remaining to implement our filters and traps. Our software development team simultaneously enhanced the original prototype SMTP filter, later called *BOMB-SHELTER*.

The technical strategy of the countermeasures against mail spoofing was simply to queue incoming mail messages, filter the mail based on developed rules-sets, and forward the *clean* mail. Rule-sets triggered on information in the header control files of the mail messages. The message content was not used in the filtering process. All filtered mail was processed via one of two paths. Mail was either sent to jail, *qjail*, and not delivered, or copied into *qcopy* for further analysis. Denying direct feedback to hackers was our cornerstone strategy, referred to by the team as *Black Holing* spam.

Trained system operators formed our third line of defense. System operators received crash training on recognizing when an attack was under way and what manual actions were necessary to avoid total system shutdown. Fortunately, our automated defenses worked and negated the need to put the final line of defense, carbon-system intervention, to the test.

The first engagement in this campaign ended with the team decision to deny feedback to the enemy. Our *Black Hole* strategy expanded; requiring zero reject notices for bogus e-mail, thus creating a situation where hacker e-mail destined for Langley AFB, or anywhere else, terminated in the *Black Hole*.

THE HACKERS STRIKE BACK

Battlefields in 5th dimensional warfare are different than traditional military theaters. Information warriors can coordinate attacks from the global virtual battlefield by communicating globally and immediately via the Internet. High grade encryption is universally available to

information warriors, further complicating effective countermeasures.

"Hackers are analogous to global nats. They swarm together and really bug the victim."

– Bass

Langley AFB became a virtual information warfare battlefield. The hackers quickly discovered simple filter rule-sets, for example the '@' symbol rule, and adapted with e-mail spam messages including pseudo random '@' characters in the forged addresses. However, the defensive team always responded, establishing a covert cyberfront with the intangible group of invisible global hackers.

Technically, our countermeasures proved effective. However, from an economic perspective the hackers were winning. At any given moment, 2 to 20 technical resources were engaged analyzing jailed messages, refining filter rules, or tuning the filter engine. USAF senior management also was continually engaged coordinating information and managing interested outside agencies.

Off-the-shelf modifications to the sendmail configuration file did not meet the *Black Hole* requirement set by the management team. The software development team made numerous changes to the sendmail configuration file; but in the final analysis, the process of queuing, filtering, and forwarding e-mail proved the more robust and flexible. In fact, many of the custom sendmail configuration rules actually deleted mail or created rejection messages. Neither option was acceptable to the management team under the *Black Hole* strategy. Deleting spam automatically was unacceptable because forensic evidence would be destroyed. Generating reject messages has two problems. First, hackers are given rapid feedback to countermeasures. Finally, reject messages can be used as another spamming technique, as earlier described.

The spam attacks continued, sometimes peaking at over 30,000 messages per day, without any denial-of-service to the AFB SMTP e-mail infrastructure. The constant barrage gave rise to the desire to "spam the hackers back!" but logic and discipline prevailed throughout the team. We redoubled our efforts to minimize all feedback in the cyber battlespace.

REVENGE OF THE JEDI

Brainstorming sessions between the team members for improving the filter algorithm yielded numerous excellent ideas. The hackers were adapting to new filter rule sets within 24 to 48 hours. The most common spam element was repeated e-mail spam with the same sender-receiver pairs. This fact provided a successful indicator of hostile e-mail. We refined the filtering algorithm based on this observation and provided the programming requirements to the development team. This algorithm became a key element in mitigating numerous types of mail bombs.

The software development team provided numerous enhancements to the original filter prototype which resulted in reduced countermeasures manpower and improved

filter granularity. The prototype version of the filter, *smt-filter.pl v0.0*, was conceived after numerous attempts to use the existing SMTP log files to look for forged e-mail and filter the mail in real time failed. The speed of the sendmail process receiving and forwarding mail in real-time made queuing the messages prior to filtering necessary. The additional latency in the mail processing was actually offset by performance improvements in the SMTP infrastructure.

In addition, by queuing the messages, the entire SMTP header file and control messages could be used in the filter process. This proved to be extremely valuable in the process of examining mail bombs, understanding the nature of the attacks, and simultaneously insuring all e-mail was correctly delivered with minimal delay.

Having secured the network, our phase II analysis began to look at ways to provide identification and warning. Statistical process control gave us our first tool in this phase of the campaign. By calculating ratios of good vs. bad mail, calculating message averages and establishing statistically based upper and lower control limits, trends began to emerge. If the amount of jailed mail exceeds the established control limit, Langley AFB was under attack. Likewise, if the number of jailed messages dropped too low, it could indicate that our adversaries had broken through our defenses and were on the verge of overwhelming the SMTP infrastructure.

The ratio of delivered vs. jailed e-mail led to another conclusion. Bogus e-mail accounts for almost three quarters of all e-mail entering Langley AFB. If the empirical data at Langley holds true for the rest of the DoD, substantial savings in bandwidth utilization might be possible by implementing similar filters throughout the DoD at strategic SMTP gateways. To verify this hypothesis, ACC plans to collect a large cross section of data and is planning on releasing a controlled ACC command wide version of the filter for study purposes.

CONCLUDING REMARKS

Traditionally, military strategies include continental, maritime and aerospace schools. Land-power advocates follow the Clausewitz strategies, maritime supporters the Mahan or Corbett theories. Douhet, considered the patriarch of the aerospace school, initiated strategies for the air environment. Each classic school emphasizes its unique decisive ability to win the war. 5th dimensional warfare makes similar claims, but the battlespace is different. Whereas land, sea, and air battles focus mainly on controlling a physical environment, information warfare is primarily political, social and psychological. Our example in this paper clearly demonstrates the exploitation of an indirect cumulative strategy versus a direct sequential one. Physical property is of minimal importance in 5th dimensional warfare.

The primary battlespace is the electronic manifestation of human mind-computer networks. Information dominance is a necessary new school of military strategy bet-

ter understood from campaigns like the one at Langley. Langley's 5th dimensional experience has proved beneficial to the USAF by helping the senior staff understand and develop future roles and responsibilities for the information warfare battlefield.

The ACC Tiger Team effectively suppressed the covert distribution of pornographic material via on base mail relays by Internet distributors. This situation, if undetected, could have resulted in undermining the credibility of the USAF. We also gained a first-hand understanding of a very real and dangerous information warfare battlefield. The *Black Hole* strategy combined with the rapid prototype and deployment of a rules-based filter and detailed knowledge of the battlespace was an effective countermeasure to Internet mail bombs.

ACKNOWLEDGMENTS

Lt. Col. Dave Gruber's leadership and moral support during our late-night filter prototyping was the glue which bound our team together. A special thanks is in order for USAF AFNCC personnel at Langley AFB for providing access to SMTP mail servers to test the filter. In addition, the Computer Support Squadron (CSS) of HQ ACC added numerous enhancements and bug fixes to the original prototype filter code.

APPENDIX: SMTP MAIL BRIEF

SMTP mail requires higher level mechanisms if trust or privacy is a requirement [1]. Without these mechanisms it is impossible to explicitly determine who originated the SMTP message. This is because SMTP is a very simple protocol, as illustrated below. In 1994, Cheswick and Bellovin discussed SMTP spoofing, stating:

"... hackers learn these commands and occasionally type them by hand." [1]

SMTP exchanges 7-bit ASCII text characters using a very simple protocol. The sample session below illustrates this simplicity from an example SMTP session. Arrows pointing to the left represent information from the server and the right arrow represents the flow of information initiated from the client.

```
<--- 220 isp.isp.com ESMTP..Sun, 26 Jan 1997..
---> helo hacker.com
<--- 250 isp.isp.com Hello there!!
---> mail from: hacker@the.great
<--- 250 hacker@the.great... Sender ok
---> rcpt to: spy@hacker.club
<--- 250 spy@hacker.club...Rcpt ok (queue)
---> Data
<--- 354 Enter mail, end with "." ...
---> Execute plan B now. That's an order!
---> .
<--- 250 JAA16273 Message accepted for delivery
```

Any method of creating a TCP connection can accomplish the exchange above. One of the more simple ways

is to simply **telnet(1)** to an SMTP server and connect to the port of an SMTP server that is actively listening for incoming connections.

APPENDIX: FILTER ARCHITECTURE

The flow of events for the filtering process was previously illustrated in Fig. 1. The SMTP server is started with the **-odq** switch [2], instructing sendmail to receive and queue incoming mail, *mqueue*, only. The filter program is executed by **crond(8)** processing the sendmail queue, *mqueue*, by first copying all messages in the queue to another directory where the files may be processed. Incoming mail continues to arrive in *mqueue* and the filter program processes the staging queue, *qprocess*.

The remaining messages in *qprocess* are moved to another queue, *qclean*, and sendmail forwards the mail by executing with the **-q** switch, instructing the MTA to process the mail queue. Another switch, **-Q**, is used to instruct sendmail which queue to process. For more technical details on BOMBSHELTER, please contact either one of the authors via e-mail.

BRIEF LIST OF ACRONYMS

ACC	Air Combat Command
DNS	Domain Name System
MTA	Mail Transfer Agent
PERL	Pattern Extraction and Report Language
SMTP	Simple Mail Transfer Protocol
TCP	Transmission Control Protocol

REFERENCES

- [1] Cheswick, W., and Bellovin, S., *Firewalls and Internet Security*, Addison-Wesley, Reading, MA, 1994.
- [2] Costales, B., with Allman, E., and Rickert, N., *sendmail*, O'Reilly & Associates, Sebastopol, CA, 1993.

BIOGRAPHIES

Tim Bass (*bass@silkrad.com*) is a technical director with SAIC, Center for Information Protection. Mr. Bass graduated from Tulane University in 1987 with a B.S.E. with Departmental Honors in Electrical Engineering. He played a principal role in building the SprintLink Integrated Network Management Center and the Sprint Managed Router Network organization in 1993. He completed the design and implementation of the original Base Network Control Center (NCC) prototype for the USAF and installed over 20 Network Control Centers for USAF bases worldwide. His current interests are secure messaging, TCP/IP performance analysis, network security, and wireless networks.

LtCol. Glenn Watt (*wattg@hqaccsc.langley.af.mil*) is the branch chief for USAF HQACC/SCBP, Information Control and Protection. LtCol. Watt holds a B.S. in Mathematics from Kutztown State College and an M.S. in Computer Science from Lehigh University. He has spent the last 10 years involved with defensive information warfare at all levels from National Defense organizations to Air Force headquarters. He developed and implemented the original Secret-to-Unclassified Network Guard (SUNG) at USSTRATCOM.