



Messaging Anti-Abuse Working Group (MAAWG)

Best Common Practices for Mitigating Abuse of Web Messaging Systems

August 2010

Table of Contents

INTRODUCTION 1

TYPICAL ATTACKS 2

MONITORING AND ALERTING..... 2

PROACTIVE DEFENSE 3

 UI ACCESS 3

 CONTENT FILTERING 5

 DISTRIBUTION 6

RESPONSES TO ABUSE 6

 CHALLENGE 6

 QUARANTINE 6

 REJECT 7

CONCLUSION..... 7

Introduction

As spam filters have improved at blocking direct connections from spammers, cyber criminals are increasingly turning to Web-based messaging systems to transmit their content. By volume, the preponderance of this spam still occurs directly on traditional "Web Mail" sites, where criminals write automated scripts that mimic user actions and send outbound spam. However, these same abusive techniques have been used outside the email domain with spammers entering content on blogs and discussion boards; abusing instant messaging and SMS portals; and misusing notification, share-with-a-friend, instant messaging and "status" features.

Many techniques are employed to prevent or mitigate these attacks and this document details the Best Common Practices for protecting Web-based systems. Its intended audience is technical professionals responsible for Web messaging feature development working in engineering, information technology (IT), network operations or a more general abuse prevention role.



Messaging Anti-Abuse Working Group

P.O. Box 29920 ■ San Francisco, CA 94129-0920 ■ www.MAAWG.org ■ info@MAAWG.org

This paper is the result of cooperative efforts within the Messaging Anti-Abuse Working Group. MAAWG is the largest global industry association working against spam, viruses, denial-of-service attacks and other online exploitation. Its members include ISPs, network and mobile operators, key technology providers and volume sender organizations. It represents over one billion mailboxes worldwide and its membership contributed their expertise in developing this description of Web messaging abuse prevention practices.

Typical Attacks

The unifying concept of Web messaging abuse is that the abuser interacts with a Web (i.e. HTTP-based) user interface, through which he submits content that is either displayed online or distributed to target users. In addition to direct spam email from a Web messaging product, several other examples of this type of abuse are:

1. "Spammy" content added to event invitations or e-cards
2. "Ridealong spam," where spammy content is added as a preamble to a "share this article" or "invitation to connect" system-generated message
3. Spammy comments added to a blog, discussion board or media sharing site
4. Spammy "status" spam on an instant messenger or a social network profile
5. SMS spam sent through a Web->SMS gateway

As more and more notification channels are added online – including RSS readers, newsfeeds, status updates, Twitter clients and instant messaging – the temptation for spammers continues to increase.

Monitoring and Alerting

The first stage in protecting a site lies in understanding the degree and frequency of abuse. Lacking adequate monitoring, many sites are caught unaware, only realizing an attack has been brewing when it reaches a critical level.

Unfortunately, there is no generic approach to monitoring. Instead, administrators must conduct an audit of the resources attackers may wish to abuse and implement a tracking solution for each. To facilitate this audit, several example areas that most properties should monitor on an ongoing basis are:

- User/Account Metrics
 - New user registrations
 - User logins
 - User password changes
 - User account deactivations
- Abusive Transactions
 - Visits to non-existent pages
 - Suspicious query strings
 - Visits to uncommon ports
- Potentially Suspicious Transactions
 - Total actions completed (e.g., messages sent, comments posted)
 - Identities who reached any of the limits or heuristics suggested above

Proactive Defense

In conjunction with the operational monitoring techniques detailed above, the bulk of service providers' time and efforts must be spent on proactive defense with systems put in place to limit the number and types of attacks. There are three key areas to consider in defending against Web messaging abuse:

1. A potential abusers' access to the UI (user interface)
2. The content that a user is allowed to submit
3. The distribution breadth and format of the content

Each of these stages represents an area that must be both monitored and protected against abuse since a vulnerability in any can be exploited by spammers. The sections below address each area.

UI Access

It almost goes without saying, but if the abuser is prevented from accessing the website by various technologies like those discussed below, his ability to send spam ends there — yet many providers leave vulnerabilities in this area, instead focusing entirely on content filtering.

Authentication

The most common technique used to limit access to a particular Web feature is authentication. While far from a panacea, requiring a user to create an account before completing an action increases the cost to abusers, thus reducing the amount of abuse. The impact of authentication can be considered in simple, economic terms: Each "step" the abuser needs to take to accomplish his goal has a "cost," so requiring him to create an account and log in increases the setup cost, which in turn decreases the "return on investment" of this attack. Obviously the cost is not infinite, so if the value of the attack remains greater than the cost, abuse will continue.

Registration/provisioning and account compromise are at the heart of authentication: The difficulty of creating or obtaining an account is a major factor in the "cost" to the abuser. Nearly all of the messaging-specific techniques in this document can and should be used to secure account registration and provisioning as well. While account compromise and password cracking are outside the scope of this document, one possible solution for some types of service is to rely on an outside provider. Rather than building a registration and authentication system, companies can accept credentials via API from AOL, Microsoft, Yahoo!, Facebook, Twitter or other services and avoid the pains and perils of building a secure implementation from scratch.

CAPTCHA

A close second in common UI-level prevention techniques is the use of CAPTCHA – most often a visual puzzle difficult for computers to recognize, but also occasionally an audio puzzle, mathematical puzzle or trivia question – as a requirement before completing an action. Like authentication, CAPTCHA is not a silver bullet but undeniably reduces the incidence of abuse. That said, the "cost" to abusers of solving CAPTCHA has been declining rapidly due to advances in Optical Character Recognition (OCR) and the use of low-cost "sweatshops" employing human CAPTCHA-solvers, so it should always be used in conjunction with other techniques. As with authentication, it is recommended that implementers stay away from developing "home made" CAPTCHA solutions and should instead rely on a field-tested and actively maintained solution such as Google's reCAPTCHA.

Web Page "Hardening"

A third method of limiting abusers' access to a Web feature is by "hardening" the interface against automated scripting. Some techniques include randomizing or varying the structure of the page (e.g., changing the

HTML "ID" parameter of required input fields), obfuscating the structure of the page (e.g., using JavaScript to encrypt field names or AJAX to dynamically present fields), and laying traps (e.g., using CSS to create invisible fields with the expectation that humans will skip over them but robots will reveal themselves by entering data).

Unfortunately, many of these techniques have the unintended consequence of decreasing accessibility for blind users or others who do not rely solely on visual cues when completing a Web page. Additionally, purists will correctly claim that obfuscation provides no true security and that these techniques will, at best, present a temporary hassle for determined attackers. While true, every additional hassle increases the cost to abusers. Locking the door does not guarantee you will not be robbed, but leaving the door ajar does make you an inviting target for casual or opportunistic thieves.

Rate Limiting and Dynamic Blocking

The final topic in the UI Access category involves dynamically limiting whether and how often certain users, computers or browsers can access your Web service. "Rate limiting" – in which a system regulates the allowable quantity and frequency of a particular action – should be employed on almost all Web services that accept or relay user-generated content (UGC).

The determination of an attacker's identity is always a major challenge in Internet security. While clearly the broad topic of identity is outside the scope of this document, for rate limiting purposes it is necessary for a service provider to choose a stable "handle" upon which rate limiting can be performed. The most common attributes are a combination of one or more of:

- Remote IP address
- The browser UserAgent string
- Username (where applicable)
- Session ID or cookie value (where applicable)
- Computed machine identifier (e.g., a MAC address, CPU serial number or OS-provided identifier)

Several providers have successfully employed a solution based on the open-source "mod security" Web server plugin while others have accomplished this task with a simple counting server. With either tool, at its most basic, administrators maintain a count for each action clustered along one or more of the identity dimensions listed above. For example, counters might be configured to limit the number of messages to "50 per hour and 100 per day for each user ID," after which the Web server will redirect user requests to an error page. Care must be taken that all access points for a feature are protected (for example, all of the servers behind a load-balancer), and that the protection cannot be circumvented by a misbehaving Web browser (e.g., if the attacker ignores an HTTP redirect code).

One major difficulty with rate limiting is the definition of the rules. Care must be taken to handle accidental spikes in user behavior (for example, the case where a good user attempts to send "I just had a baby" to his whole list) and misinterpretation of clustering dimensions (e.g., the IP address of an Internet cafe or public library may represent 20 users, or a business traveler may log in from a large number of unusual locales).

Attention must also be paid to the user interface presented when the limit is reached. "False positives" can be completely mystifying to users, and in an effort to not divulge the limits to abusers (e.g., "Suspicious activity detected..." versus "You are not allowed to send more than 100 messages per hour"), these error pages can generate costly customer support volume. Oftentimes, the spammers either already know the limits or will continue sending even beyond the limits. However, legitimate users may be completely unaware they have reached the limit because they are sharing a NAT IP address with their roommate's laptop, which has a bot and is hemorrhaging spam.

Dynamic blocking can be considered a degenerate case of rate limiting, where a forbidden IP address is allowed exactly zero transactions per day. Some providers have used third-party data feeds to determine the blacklist – for example the Spamhaus XBL may be applicable – but care must be taken when adapting a blacklist to an unintended purpose. (For example, the Spamhaus PBL lists IP addresses that should not connect to port 25 but is completely inappropriate for blocking port 80.)

A special note on IPv6: In almost all Web messaging abuse scenarios, the remote user's IP address is a valuable identifier for rate limiting since it is fairly robust against spoofing or tampering. As the forthcoming migration from IPv4 to IPv6 radically expands the remote IP space, care must be taken in the design of rate limiting systems to accommodate this much larger signal space. Implementers should consult the upcoming MAAWG document on IPv6; several specific concerns to keep in mind around rate limiting are:

- Storage of remote IPs may exceed memory constraints
- Lack of remote IP scarcity may require lowering thresholds

Carrier-grade NAT routers – which hide dozens or hundreds of consumers behind a single IPv4 address – may necessitate raising thresholds.

Until we start to see large-scale deployments of IPv6, it is difficult to anticipate the impact on systems. This note is intended as a reminder to implementers of new systems to keep these concerns in mind and to continue to consult up-to-date information from MAAWG about evolving best practices.

Content Filtering

Least Privilege

The strongest suggestion for content filtering follows the security principle of "least privilege": Only allow the types of content that you absolutely must. Oftentimes this equates to stripping out all HTML from an input field, or allowing only the most basic formatting tags (i.e. , <i> and <u>). <script> tags should almost never be allowed in user-generated content and many have found that simply removing the <a> tag – forcing an abuser to type "visit <http://www.webcam.com>" instead of "visit [my web cam](#)" with a link embedded – can dramatically reduce the attractiveness of a particular channel. Attachments or file uploads, where permitted, should be restricted by file type or extension and ideally scanned for malware.

Another aspect of least privilege is restricting the character count of the message payload. In the note accompanying a shared news article, there is little need to accept 1000 characters of input; restricting the message to 200 characters supports the intended need while reducing the value of this field to spammers.

Heuristics

Heuristic filtering – for example, comparing submissions to a blacklist of forbidden URLs, text strings or other tokens – can further encourage spammers to look elsewhere. Unfortunately, heuristics must often be manually created and maintained, and yet they can often be easily evaded (for example, by writing "V.1agra" or using a URL redirector), making them a costly and complicated option for site administrators. These heuristic approaches must almost always be accompanied by a scalable process for populating and expiring rules.

Reputation Models

More flexible than heuristics filtering are online and offline reputation models which dynamically evolve based on patterns in outbound content and sender profiles. For example, a model could be constructed to examine user trust that considers the account creation date, the number of positive actions performed, and

the number of suspicious actions performed. Upon a questionable submission, the system could consult this model to determine the prior probability that the message is spammy.

Batch Analysis

The most flexible approach to message filtering includes examining a very broad feature vector for patterns that indicate suspicion or trustworthiness. To do so, clustering algorithms are run on a batch of recent messages to find unexpected patterns – perhaps accounts that registered on June 7 from an IP in 1.2.0.0/16 are all sending messages with 103 words. Those patterns are then loaded back onto the outbound filters, and in this case, subsequent messages matching this pattern might be flagged for further scrutiny.

Distribution

The final aspect to consider in mitigating Web-based abuse is the degree of distribution of the suspected message. The most compelling aspect for would-be spammers is how many potential customers will receive the message. Therefore, techniques that limit the distribution can be effective at combating the abuse itself.

Recipient Limiting

In an email or direct messaging system, administrators may choose to limit the number of recipients an untrusted user may include in a message without arousing suspicion. In systems with implicit distribution, such as a blog or a comment system, this may take the mode of required content moderation or a quarantine of new posts.

Responses to Abuse

Following the Content Filtering stage, a good message is "sent" (i.e., posted, stored, Tweeted, etc.) while a bad message should go through one of three paths:

- Challenge: A UI can be presented to the user for further validation
- Quarantine: The message can be delayed or forwarded for moderation
- Reject: The message is not delivered; it can be silently deleted or overtly rejected and sent back to the submitting user

Challenge

When Content Filtering determines the message falls into a grey area between "definitely good" and "flagrantly bad," administrators may wish to present a challenge to the submitting user for further validation. Oftentimes this takes the form of a CAPTCHA – e.g., "Before accepting this post, please solve this puzzle" – but increasingly, sites are using two-factor authentication methods like a code delivered by SMS.

Just as with the content submission itself, care should be taken that the challenge UI is similarly layered with security protections. Over-reliance on CAPTCHA or even SMS can create a vector for attackers to send flagrant messages for a reasonably low cost. A significant vulnerability exists if quasi-bad messages can be distributed for just the price of an automated CAPTCHA solution.

Quarantine

In certain circumstances, automated filters may need more time to examine a message – for example, to see whether an unusually large volume is being submitted – or in some cases, manual moderation may be required. This is most common with photo sharing sites.

Reject

The only decision around rejecting content is whether or not to alert the user. While doing so does convey information that spammers could use to adjust their tactics, silently deleting a message can generate customer support volume as well as tremendous difficulty troubleshooting any system bugs. Most providers have adopted a compromise position, displaying an error message back to the user but keeping its content fairly vague or generic.

Conclusion

With spammers and botnet herders increasingly moving towards scripting Web interfaces, the three sections above—Monitoring and Alerting, Proactive Defense, and Responses to Abuse – lay out a series of techniques that site administrators can use to protect against or mitigate the effects of abuse. Beginning with understanding the magnitude of the problems and leading through techniques for addressing them, this document lays out a comprehensive approach for site administrators to combat cyber abuse of their Web messaging systems.