

Complaint Feedback Loop Operational Recommendations

Abstract

Complaint Feedback Loops similar to those described herein have existed for more than a decade, resulting in many de facto standards and best practices. This document is an attempt to codify, and thus clarify, the ways that both providers and consumers of these feedback mechanisms intend to use the feedback, describing some already common industry practices.

This document is the result of cooperative efforts within the Messaging Anti-Abuse Working Group, a trade organization separate from the IETF. The original MAAWG document upon which this document is based was published in April, 2010. This document does not represent the consensus of the IETF; rather it is being published as an Informational RFC to make it widely available to the Internet community and simplify reference to this material from IETF work.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6449>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may not be modified, and derivative works of it may not be created, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Overview	4
2. Glossary of Standard Terms	5
3. Mailbox Providers and Feedback Providers	9
3.1. Benefits of Providing Feedback	9
3.2. Collecting Complaints	10
3.3. Creating Reports	11
3.4. Policy Concerns	11
3.4.1. Privacy and Regulatory Compliance	11
3.4.2. Terms of Use	12
3.5. Handling Requests to Receive Feedback	12
3.5.1. Application Web Site	13
3.5.2. Saying No	14
3.5.3. Automation	14
3.6. Ongoing Maintenance	15
3.6.1. IP Validation	15
3.6.2. Email Address Validation	16
3.6.3. Feedback Production Changes	16
4. Feedback Consumers	16
4.1. Preparation	17
4.2. What You'll Receive	18
4.2.1. Feedback Reports	18
4.2.2. Administrative Messages	18
4.2.3. Report Cards	18
4.3. Handling Feedback Messages	19
4.3.1. Unsubscription or Suppression	20
4.3.2. Trending and Reporting	21
4.4. Automatically Handling an Incoming Feedback Stream	22
5. Conclusion	25
6. Acknowledgments	26
6.1. About MAAWG	26
7. Security Considerations	26
8. Informative References	26
Appendix A. Abuse Reporting Format (ARF)	28
A.1. A Brief History	28
A.2. Structure of an ARF Message	28
Appendix B. Using DKIM to Route Feedback	29
Appendix C. Unsolicited Feedback	30
C.1. Guidelines	30
C.2. Pros	30
C.3. Cons	31

1. Overview

The intent of a Complaint Feedback Loop is to provide Feedback Consumers with information necessary to mitigate Spam or the perception of Spam. Thus, feedback was originally only offered to mailbox, access, and network providers -- in other words, to ISPs -- who would use the feedback to identify network compromises and fraudulent accounts or to notify their downstream customer that there may be a problem.

Senders of bulk, transactional, social, or other types of email can also use this feedback to adjust their mailing practices, using Spam Complaints as an indicator of whether the Recipient wishes to continue receiving email. Common reactions often include refining opt-in practices, mailing frequency, list management, message content, and other measures. Over time, this has become the Feedback Consumer use case most often discussed at MAAWG meetings and other industry events -- but readers are cautioned that it is not the sole use for feedback.




Figure 1

When an End User of a Mailbox Provider issues a Spam Complaint, the Feedback Provider sends a report to the Feedback Consumer. This report may include the Full Body of the original email or (less commonly) only the full header of the original email. Some Feedback Providers will redact information deemed private, such as the Message Recipient's Email Address.

Ensuring that Feedback Messages are only sent to authorized Feedback Consumers is the responsibility of the Feedback Provider, with the identity of each message Sender generally determined from the SMTP session's connecting IP address or a message's DomainKeys Identified Mail (DKIM) signature domain, both of which are hard to forge. This is important because Spammers and other miscreants may also attempt to apply for Feedback Loops on networks not belonging to them, in an attempt to steal Email Addresses and other private personal or corporate information.

It is the responsibility of the Feedback Consumer to identify the source and nature of the original message in the reports they receive and take any appropriate action. The Feedback Provider does not make any claims or judgments about the validity of the complaint, beyond whatever technical data the Feedback Provider has themselves included. Every complaint is forwarded to the Feedback Consumer without human review, without any additional application of filters; thus, some individual reports may prove not to be actionable.

The Feedback Consumer and the Feedback Provider will each evaluate a Spam Complaint for validity and take whatever action deemed necessary from their own perspective and, in most cases, will not communicate with each other which actions were (or were not) taken. Similarly, it is rare for any party to communicate further with the End User who initiated the complaint.

2. Glossary of Standard Terms

Wherever possible, these terms are derived from [RFC5598].

- o Abuse Reporting Format - The standard format for Feedback Messages, defined in Appendix A and [MARF].
- o Access Provider - Any company or organization that provides End Users with access to the Internet. It may or may not be the same entity that the End User uses as a Mailbox Provider.
- o Application for Feedback Loop - the process, manual or online, by which a prospective Feedback Consumer requests to receive a Feedback Loop from a particular Feedback Provider.
- o ARF -- See "Abuse Reporting Format".
- o ARF Report -- See "Feedback Message".
- o Body - See "Full Body".
- o Complaint or Complaint Message - See "Feedback Message".
- o Complaint Feedback Loop - See Overview and Taxonomy section.
- o Complaint Stream - See "Feedback Stream".
- o Delivery - See "Message Delivery".
- o DKIM - DomainKeys Identified Mail, further described in the MAAWG email authentication white paper "Trust in Email Begins with Authentication" [Trust] and [DKIM].

- o End User - A customer of a Mailbox Provider or Access Provider.
- o Envelope Sender - The Email Address included as the argument to the [SMTP] "MAIL" command during transfer of a message.
- o Email Address - A string of the form user@domain, where the domain (after the @ symbol) is used to determine where to transfer an email message so that it may be delivered to the mailbox specified by the username (before the @ symbol). The precise technical format of an Email Address is defined in [SMTP]. Email delivery can be a complex process and is not described further in this document.
- o Email Service Provider (ESP) - A provider of email sending services; the ESP is often a Message Originator working on behalf of a Message Author. MAAWG uses the term "ESP" solely for this definition and does not refer to a Mailbox Provider for End Users as ESPs.
- o FBL - The acronym "FBL" (Feedback Loop) is intentionally not used in this document.
- o Feedback or Feedback Stream - A set (often a continuous stream) of Feedback Messages sent from a single Feedback Provider to a single Feedback Consumer.
- o Feedback Consumer - A Recipient of the Feedback Messages, almost always on behalf of or otherwise associated with the Message Originator. Often the Message Originator and Feedback Consumer are the same entity, but we describe them separately in this document because they are each responsible for different parts of the Complaint Feedback Loop process, as demonstrated in the flowchart in the Overview section.
- o Feedback Loop - See Complaint Feedback Loop.
- o Feedback Message - A single message, often using the Abuse Reporting Format defined above and outlined in Appendix 1, which is part of a Feedback Stream.
- o Feedback Provider - The Sender of the Feedback Messages, almost always on behalf of or associated with the Mailbox Provider. Often the Mailbox Provider and Feedback Provider are the same entity, but we describe them separately in this document because they are each responsible for different parts of the Complaint Feedback Loop process. In some instances, the Feedback Provider may be operating solely on behalf of the Message Recipient, without any direct participation from their Mailbox Provider.

- o Full Body - An email message (the "DATA" portion of the [SMTP] conversation) consists of two parts: the header and the body. The "Full Body" is simply the entirety of the body of the message, without modification or truncation. Note that images or other so-called "attachments" are actually part of the body, designated in accordance with the [MIME] standard.
- o Full Header Section - An email message (the "DATA" portion of the [SMTP] conversation) consists of two parts: the header and the body. The header contains multiple header fields, each formatted as "Header-Name: header contents". Although most Mail User Agents (MUAs) only show the basic four header fields (From, To, Date, and Subject), every message includes additional header fields that primarily contain diagnostic information or data intended to assist automatic processing. Often informally called "Full Headers". These fields are fully defined in [RFC5322]
- o Header - See "Full Header Section" above.
- o ISP - Internet Service Provider, usually referred to as either an Access Provider or a Mailbox Provider in this paper.
- o Mail Abuse Reporting Format (MARF) - See "Abuse Reporting Format" above.
- o Mailbox Provider - A company or organization that provides email mailbox hosting services for End Users and/or organizations. Many Mailbox Providers are also Access Providers.
- o Mailing List - A set of Email Addresses that will receive specific messages in accordance with the policies of that particular list.
- o Message-ID Header Field - One of the diagnostic header fields included in every email message (see "Full Header Section" above) is the Message-ID. Theoretically, it is a unique identifier for that individual message.
- o Message Delivery - The process of transferring a message from one mail transfer agent (MTA) to another. Once the message has been accepted by the MTA operating on behalf of the Recipient, it is considered to be "delivered" regardless of further processing or filtering that may take place after that point.
- o Message Originator - The Sender, but not necessarily the author or creator, of a message.
- o Message Recipient - The person or mailbox that receives a message as final point of delivery.

- o MIME - Multipurpose Internet Mail Extensions refers to a set of standards permitting non-plaintext data to be embedded in the body of a message. Concepts such as file attachments and formatted or "rich" text are all accomplished solely through [MIME].
- o MUA - Mail User Agent; loosely referring to the software used by an End User to access, interact with, or send email messages.
- o Provider - See "Feedback Provider" above.
- o Received Header Field - Diagnostic header fields included in an email message (see "Full Header Section" above) that start with "Received:" and document (from bottom to top) the path a message traversed from the originator to its current position.
- o Recipient - See "Message Recipient" above.
- o Return-Path - An optional message header field (see "Full Header Section" above) that indicates the Envelope Sender of the message.
- o Reverse DNS - The [DNS] name of an IP address, called "reverse" because it is the inverse of the more user-visible query that returns the IP address of a DNS name. Further, a Reverse DNS query returns a PTR record rather than an A record.
- o Sender - see "Message Originator" above.
- o SMTP - Simple Mail Transfer Protocol, the mechanism and language for transferring an email message from one place to another as defined in RFC 5321 [SMTP].
- o Spam - For the purposes of this document (and for most Complaint Feedback Loops), "spam" is defined as any message that the Recipient chooses to complain about, regardless of the intent of the message's author or Sender.
- o Spam Complaint - See "Complaint" above.
- o Spammer - An entity that knowingly, intentionally sends Spam messages (see "Spam" above).
- o Terms of Use - A legal document describing how a particular system or service is to be used.
- o VERP - Variable Envelope Return Path [VERP], an informally standardized method for encoding information about the Message Recipient into the return path while delivering a message in order to ensure that any non-delivery notices are processed correctly.

3. Mailbox Providers and Feedback Providers

In practice, a Mailbox Provider receives complaints from their End Users, and is often also the Feedback Provider for those complaints and is a consumer of feedback from other providers. In this document, we separate the Mailbox Provider and Feedback Provider functions to reduce possible confusion over those cases where they are separate, and we also urge Mailbox Providers to read the "Feedback Consumer" section later in this document.

3.1. Benefits of Providing Feedback

The decision to provide a Complaint Feedback Loop service should not be taken lightly. The benefits of a Feedback Loop are great, but success depends on a sound plan, organized implementation, and dedication to upkeep.

What are some benefits of providing feedback to fellow Mailbox Providers and Access Providers? Primarily, other industry actors are quickly alerted to Spam outbreaks on their networks.

End Users are becoming more aware of and comfortable with mechanisms to report Spam, and a Feedback Loop does just what it implies; it closes the loop. The End User's complaint makes its way back to the Message Originator (not necessarily the message Sender, who may be a Spammer), allowing the originator to take appropriate action. In this process, the mail system operator is just a messenger, relieved of the responsibility of reviewing and forwarding complaints manually.

Further, because every complaint is sent immediately -- without any review or analysis by the Feedback Provider -- the complaint is received by the Feedback Consumer in near real time. If the Feedback Consumer is paying attention to their Feedback Stream and taking appropriate action on it, the receiving Mailbox Provider receives less Spam, blocks less legitimate mail, and does not have to assign staff to follow up with the originating network. If the Mailbox Provider does not pay attention to its Feedback Stream, and does not take appropriate action, the Feedback Provider may block or otherwise filter the email from that Message Originator, considering the Feedback Messages to be sufficient notice.

What are some benefits of providing Feedback Loops to bulk Feedback Consumers? As Message Recipients become more aware of and comfortable with Spam reporting mechanisms, they often prefer this method over the often-confusing and inconsistent "unsubscribe" or "opt-out" mechanisms provided by most legitimate Message Originators or Senders.

End Users often do not remember what lists they signed up for or are otherwise not confident in the established relationship they may have with a message Sender. As such, they often choose to report messages as Spam to their Mailbox Providers, considering that to be sufficient notification of their desire not to receive such email in the future.

If the Message Originator is paying attention to and taking appropriate action on their Feedback Stream, it will have a happier set of Message Recipients and should receive fewer Spam Complaints (assuming their opt-in processes are sound). If the Message Originator is not paying attention to Feedback and not taking appropriate action, the Mailbox Provider may consider the Feedback Stream sufficient notice that messages from that originator may no longer be accepted in the future.

3.2. Collecting Complaints

To produce Feedback Messages and to ensure they are useful, the Feedback Provider needs to obtain near real-time complaints from the Mailbox Provider's users. This is typically done by integrating the feedback mechanism with the collection of Spam reports from its users.

These reports are typically made using the "Report Spam" buttons integrated into Webmail interfaces, or a proprietary desktop client provided to users. Mailbox Providers may also look at deploying a toolbar or MUA plug-in that provides a "Report Spam" button in the MUA interface.

Usability studies with average users should be performed on all interface changes before implementation. A "help" interface should also be available to educate users about how the Spam button should be used and what it does.

If the Mailbox Provider does not offer its customers a mail client with this button, then the Feedback Provider's chances for providing an effective Feedback Loop are slim. While it is possible for the Mailbox Provider to instruct its customers to forward unwanted mail to a central location and for the Mailbox Provider to explain how to ensure the report includes headers and bodies, the success rate of customers doing so tends to be low. Even those complaints that do contain all required information might prove difficult to parse, as variations in formatting and content types will lead to automated tools being consistently updated with new logic blocks for each variation that occurs.

3.3. Creating Reports

It is recommended that Feedback Messages be sent using the standard Abuse Reporting Format, to facilitate uniformity and ease of processing for all consumers of feedback. This will also enable the Feedback Provider to extensively automate the processes of generating and sending Feedback Messages and of analyzing complaint statistics. This format is described further in Appendix 1.

Feedback Loops are usually (but not always) keyed to the "last hop" IP address (i.e., the IP address that passed the unwanted message to the Mailbox Provider's servers). Consequently, the Feedback Provider must be able to process the header from each complaint to determine the IP address for the complaint.

A Feedback Provider may wish to provide, as part of its Feedback Loop, other information beyond Spam Complaints that Feedback Consumers may find valuable. It might include summary delivery statistics (volume, inbox delivery rate, Spam trap hits, etc.) or other data that the Feedback Provider may deem pertinent to Feedback Consumers.

Any mature Feedback Loop system will produce situations in which the Feedback Consumer may have follow-up questions or have other information to provide in regard to the feedback. Feedback Messages should include contact information (typically an Email Address) for the Feedback Consumer to use for such questions, and ideally the contact Email Address will feed into a ticket system or other automated tool used by the Mailbox Provider's postmaster and/or anti-abuse staff for handling general email delivery issues.

3.4. Policy Concerns

3.4.1. Privacy and Regulatory Compliance

Feedback Messages provide information relayed by Feedback Providers from a Mailbox Provider's End Users to the Feedback Consumer. There might not be any concerns with relaying non-private data to a third party. However, the information provided in the complaints generated by the user must be evaluated and any data deemed private may need to be removed before distributing to a third party, per local policy. For example, the Recipient's or reporter's Email Address and IP address may be categorized as private data and removed from the feedback report that is provided to the Feedback Consumer. Privacy laws and corporate data classification standards should be consulted when determining what information should be considered private.

Information provided by the Feedback Consumer to the Feedback Provider for the purpose of enrolling in the Feedback Loop should also be kept private. It should only be shared or used for the purposes explicitly agreed to during the enrollment process (see the "Terms of Use" section below).

Feedback Loops inevitably span country borders. Local laws and regulations regarding distribution of information domestically and internationally need to be considered when implementing a Feedback Loop program. For example, in some European countries, data exchange requires permission from governing bodies. The terms and circumstances surrounding the exchange of data need to be clearly defined and approved.

3.4.2. Terms of Use

A written Terms of Use agreement should be provided by the Feedback Provider and agreed to by the Feedback Consumer before any feedback is provided. The following concepts should be considered when drafting the terms of use agreement:

- o Data provided in Feedback Messages are provided to a specific, approved entity. Information should not be transmitted outside of the intended, approved Recipient. Any inappropriate use of the information can lead to immediate termination from the feedback program.
- o Consumers of Feedback have a responsibility to keep the information they provide for Feedback Loop purposes -- such as abuse contact information, IP addresses, and other records -- accurate and up to date.
- o The providing of Feedback information is a privilege and needs to be treated appropriately. It does not entitle the consumer of the feedback to any special sending privileges.
- o Approval and continued enrollment in the program is a privilege that can be denied or revoked for any reason and at any time.

3.5. Handling Requests to Receive Feedback

There should be a streamlined application process for receiving feedback and the vetting of such applications. This vetting may be stringent in cases where the Mailbox Provider chooses to tie its Complaint Feedback Loop program to a whitelist. Criteria may involve the following:

- o Cross-checking that the requestor is indeed authorized to receive feedback for the IP addresses concerned.
- o Gathering other information such as whether the IPs are an ISP smarthost network, a webhosting farm, an email marketing or Mailing List service, or other entity.
- o Requesting information such as a link to the policies of the requestor, contacts to send Feedback Messages, and escalation points of contact.

Ideally, enrollment will be a two-step process, with the applicant filling out a form and being required to receive and acknowledge a confirmation email (best sent to abuse@ or postmaster@ the domain in question) before the applicant's request is even put into the queue for the Feedback Provider to process.

Ownership of IP addresses can and should be cross-checked by means of origin Autonomous System Number (ASN), WHOIS/RWHOIS records, Reverse DNS of the sending hosts, and other sources. This can be automated to some extent, but it often requires some manual processing.

3.5.1. Application Web Site

Applications for Feedback Loops can be accepted on a stand-alone web site or can be part of the Mailbox Provider's postmaster site. Regardless, the web site for the Complaint Feedback Loop program should contain other content specific to the Feedback Loop, including FAQs for the Feedback Loop program, the Terms of Service for the Feedback Loop, and perhaps a method for enrolled parties to modify their existing enrollments.

The web site should also provide the Feedback Consumer with general information on how the feedback will be sent, including:

- o Report Format (ARF or otherwise)
- o Sending IP addresses and/or DKIM "d=" string
- o "From" Email Address

3.5.2. Saying No

Denial of a Feedback Loop application may be appropriate in certain cases such as:

- o Where the Feedback Provider suspects "gaming" of delivery policies via the Feedback received, with attempts to pollute Feedback Loop metrics by, for example, creating bogus accounts and reporting false negatives with these, to offset the negative reputation caused by high complaint rates.
- o Where the Feedback Provider has decided to block the Message Originator's IP space for which feedback has been requested on the grounds that email from that originator has a sufficiently negative reputation that it will not be delivered at all. This is somewhat on the lines of a global unsubscribe of the Message Provider's users from the originator's lists, which would make rendering additional feedback unnecessary.

It is recommended that the Feedback Provider send notification if an application is denied. Additionally, they should maintain a documented, clear, and transparent appeals process for denial of requests. This process can be as simple as the prospective Feedback Consumer replying to the denial email requesting review or escalation to a team lead, which also cites reasons the application should be reviewed.

3.5.3. Automation

For a Feedback Loop to be cost-effective and usable for large Feedback Consumers and Feedback Providers, it must be possible for reports to be generated and processed automatically without any human interaction. On the other hand, it should be possible for small Feedback Consumers to handle a low volume of reports manually, without requiring any automation.

In automating the feedback process, the consumer of the Feedback Stream must receive enough information about the report that it can take appropriate action, typically to remove the Recipient from the Mailing List about which it is sending a report. The Recipient's Email Address is not enough, as the Recipient may be on several Mailing Lists managed by the Feedback Loop consumer and only need to be removed from the particular list reported.

Also, some producers of Feedback Loops might redact the Recipient's Email Address for privacy reasons. Effective implementation of a Complaint Feedback Loop requires that the Feedback Provider put in

place as many automated processes and tools as feasible to handle all aspects of the process. Feedback Providers should seek to automate or script the following:

- o Accepting and validating Feedback Loop Applications from prospective Feedback Consumers.
- o Processing requests to determine whether or not they meet the Feedback Provider's criteria for enrollment in the program.
- o Accepting Spam Complaints from End Users; this will form the bulk (and perhaps sole) component of the feedback sent by the Feedback Provider.
- o Production of Feedback Messages from Spam Complaints.
- o Production of other Feedback Loop artifacts as chosen by the Feedback Provider.
- o Optionally, provision of a mechanism for Feedback Consumers to further engage a Feedback Provider about a given Feedback Message.
- o Ongoing validation of Feedback Loop enrollments to determine if a currently enrolled IP address or network merits continued inclusion in the Feedback Loop.
- o Optional periodic emails to Feedback Consumers to determine if their enrolled Email Addresses are still valid.

3.6. Ongoing Maintenance

It is recommended that self-service maintenance be offered to Feedback Consumers, to the extent practicable. The more they can do themselves, the less you have to do.

3.6.1. IP Validation

The criteria that a Feedback Provider uses to validate a Feedback Loop application may change over time. It is a near certainty at least some subset of Feedback Consumers enrolled to receive feedback will at some point after enrollment fail to meet those criteria, regardless of whether or not the criteria change.

The Feedback Provider should put in place tools to periodically re-validate all Feedback Consumers enrolled in its Feedback Loop system against its current criteria. Additionally, the Feedback

Provider will likely have objective criteria for remaining in the Feedback Loop for enrolled Feedback Consumers; the enrolled consumers should be validated against those criteria as well.

3.6.2. Email Address Validation

Just as some Mailing List software has the built-in ability to send periodic "probe" emails to subscribed addresses to validate them, so too should the Feedback Provider develop tools to send similar emails to the addresses receiving Feedback Messages to ensure that they are valid. This is especially true for the addresses that are not the abuse@ and postmaster@ addresses originally used as part of the enrollment acknowledgment step. Over time, people may change employers, or at least roles, and validating the Email Addresses associated with an IP is one way for the Feedback Provider to ensure that Feedback Messages are still being accepted and acted upon by the Feedback Consumer.

3.6.3. Feedback Production Changes

Updating Feedback Consumers when one's own IP addresses are changing is an important aspect of Feedback Loop maintenance. The exact format, automation, and other considerations of these updates are outside the scope of this document, but are topics worthy of further discussion and eventual documentation.

4. Feedback Consumers

A Feedback Consumer receives its Feedback Messages after its submitted Application for a Complaint Feedback Loop is approved. A Feedback Consumer will usually have Complaint Feedback Loop subscriptions set up with multiple Feedback Providers. Different Feedback Streams may be in different formats or include different information, and the Feedback Consumer should identify a process to organize the data received and take appropriate action.

A Feedback Consumer, Mailbox Provider, or Access Provider (i.e., a hosting company or ISP) will use this Feedback to identify network compromises, fraudulent accounts, policy violations, and other concerns. The Feedback Loop provides real-time visibility into Spam Complaints from Message Recipients, greatly enabling these Mailbox Providers to mitigate Spam propagating from their networks.

Senders of bulk email should use the complaints to make decisions regarding future mailings. Such decisions may include one or more of the following: modification of email frequency, branding, opt-in practices, or list management.

The authors of this document urge those who are solely Feedback Consumers to also read the previous sections for Mailbox Providers and Feedback Providers. This will provide the proper context of the recommendations included below.

Further recommendations for bulk senders may be found in the MAAWG Sender Best Communications Practices [MAAWG-BCP].

4.1. Preparation

Feedback Consumers need to prepare to process and act on feedback before asking to receive it. At a minimum, make sure to have:

1. The "Role" Email Addresses such as abuse@ and postmaster@. The person who applies for the Feedback needs to make sure they have access to these Email Addresses. Feedback Providers often send a confirmation link to those accounts to prevent End Users, Spammers, or competitors from signing up for Feedback for which they are not authorized.
2. A dedicated Email Address to receive the Feedback Messages, such as fbl@example.com or isp-feedback@example.com. While not required, this will make it easier for to process the reports received.
3. A list of IP addresses for which you want to receive Feedback Messages, making sure you can prove the ownership of the IP addresses and associated domains. Feedback Providers often require that:
 - * Reverse DNS for each IP shares the domain of either the applicant's Email Address or the Email Address that will be receiving the Feedback Messages.
 - * WHOIS information for the IPs requested is obviously associated with the domain name.
4. Contact information such as name, Email Address, phone number, and other relevant information.
5. The knowledge that if the application form asks for your credit card number or other financial information, it is assuredly a scam.

4.2. What You'll Receive

Once a Feedback Consumer has signed up to receive feedback from a Feedback Provider, it may also receive several other sorts of delivery-related reports. This includes Feedback Messages, administrative messages, and other messages.

4.2.1. Feedback Reports

Feedback Messages are the main emails generally associated with a Feedback Loop. Each time a Recipient hits the "This Is Spam" button, the Feedback Loop system creates a boilerplate report with a copy of the original email attached and sends it to the consumer of the Feedback Loop.

The handling of feedback reports is discussed in the next section.

4.2.2. Administrative Messages

Administrative messages will typically be sent to the Email Address provided for contacting the person who originally applied for the Feedback Loop, rather than to the address provided for handling the Feedback Messages. These messages are likely to be sent infrequently and irregularly, but it is important they are seen by the person managing the Feedback Stream processor in a timely manner. It is usually a poor idea to have these sent to an individual's Email Address since they may be lost if that person is on vacation, changes position within the company, or leaves the company.

Instead, they should be sent to a role account that goes to a ticketing system or "exploded" to multiple responsible parties within the organization. If there is not already an appropriate role account such as support@ or noc@ that reaches the right team, it may be a good idea to set up a dedicated alias such as fblmaster@ to sign up for all Feedback Loops.

4.2.3. Report Cards

The detail in a report card can vary greatly. Feedback Providers might send a regular summary of traffic levels and complaint rates seen, perhaps just an overview or possibly broken down by source IP address or some other identifier. Sometimes these may be sent just when some metric (typically a complaint rate) reaches a level that causes the Mailbox Provider to notify the Feedback Consumer there may be a problem developing that needs to be investigated and addressed. At the other extreme, some report cards will contain almost no useful

data at all, just a warning that the Message Originator is causing complaints -- with the implication that its email will be blocked unless it is improved.

Report cards are human readable, since there are not currently any standard machine-readable formats and the information they include, both the provided metrics and their semantics, varies widely from one Mailbox Provider to another. They are useful reference overviews for a Message Originator to monitor the overall perceived quality of the email it sends and, in the case of ESPs, perhaps which customers are causing higher than expected rates of complaints. They can also be the only warning of serious problems prior to email being blocked altogether by the receiving Mailbox Provider. It is critical they be seen by someone handling delivery issues for the Message Originator, so again, they should be handled by an email alias that is always read.

Report cards also contain useful data to track mechanically and perhaps report on trends, though as their content varies, it is hard to generalize what use might be made of them. At the very least, the "warning" report cards are something that should be visible on an ESP's business intelligence or delivery dashboard.

4.3. Handling Feedback Messages

Mailbox Providers sending feedback may have published policies as to how they expect a Feedback Consumer to use Feedback Messages or may expect the Feedback Consumer to simply "make the problem stop". In practice, this mostly boils down to three things:

- o First, where the consumer of the feedback has some specific control over sending the email, it is expected not to send email of the same type to the same Recipient again.
- o Second, it should identify the underlying problem (if any) and fix it so that it receives fewer reports of that type in the future.
- o Third, it is not necessary to inform the Mailbox Provider or Feedback Provider, or their End User(s), of which actions have been or will be taken in response to automated complaint feedback.

If the Feedback Consumer is a separate entity from the Message Originator, the two entities are expected to work together to resolve any problem.

4.3.1. Unsubscription or Suppression

A Sender (whether author or originator) of commercial email should treat the Feedback Message similar to an unsubscribe request, ensuring that no further email from that list is sent to that Recipient, either by removing the email from that list or adding it to the associated suppression list. It needs to use its best judgment, keeping in mind the goal of reducing future complaints, as to how broadly to apply that unsubscribe. Suppressing the address across an entire ESP is likely too broad. However, if a single Feedback Consumer (or customer of an ESP) has multiple segmented lists, then suppressing them across all those lists is probably a good idea.

It is universally acknowledged that not all complaints are intentional; for example, Recipients might accidentally hit the wrong button or mark an entire mailbox as Spam. However, it is best for Feedback Consumers to assume the Recipient does not want more email and to suppress mail to the Recipient in all but fairly extreme cases such as a Mailing List the Recipients pay to receive, email from a genuine company to its valid employees, or email from an Access Provider or Mailbox Provider to its users.

This gets more complex in the case of transactional mail -- mail that is tied to some other service, such as ticket purchase confirmations or billing statements. In that case, the Feedback Consumer has to, again, use its best judgment based on the specific situation. In some cases, the right thing to do may be to communicate with the Recipient via another channel, such as a message on a web site used for the service; i.e., "You reported your notification mail as Spam so we are not going to send you any more messages unless you tell us otherwise".

In some cases, the best thing to do may be to ignore the Feedback Message. For example, if your customer has reported as Spam the airline tickets he purchased and you emailed him, he probably did not mean it and he is going to be very annoyed if you do not send him the other tickets he has ordered. In rare cases, it might be appropriate to suppress email to the Recipient, but also to suspend access to a service he or she uses until the Recipient confirms a desire to receive the associated email. In all these cases, the important goal is to keep the customer happy and reduce future complaints, even in the apparently paradoxical situations where the way to do that is to ignore their Feedback. In the real world, however, these are a small minority of cases.

4.3.2. Trending and Reporting

Counting the Feedback Messages received over regular time periods can provide much useful information to ISPs, ESPs, and other Feedback Consumers, especially when broken down appropriately.

An ISP (Mailbox Provider or Access Provider) might want to count the number of Feedback Messages a particular customer or IP address causes in a given day. If there is a sudden increase from a particular customer or server, it may be a sign that a Spammer has signed up or a system has been compromised. If there is a high level of complaints about a particular customer, it may be worth investigating to see if there is a reason for that. For example, 10 Feedback Messages a day would be a sign of serious problems in some cases, but might be perfectly reasonable "background" levels for a Message Originator that sends 300,000 emails a month. If the count shows there may be a problem, the ISP can dig down and look at the emails that are being reported to determine the underlying cause.

An ESP can do similar things but can also break the data down in more ways: by customer, by Mailing List, by campaign. An ESP also has access to more information; it knows how many emails were delivered to the receiving Mailbox Provider over a given time period. As a result, it can estimate the number of complaints divided by the number of emails sent, which is often a more useful metric than the absolute number of reports. This is critical data for ESPs to track over time because it can help identify and quantify problem customers.

An individual Feedback Consumer, whether sending their own email or using an ESP, can acquire at least some information from complaint rates. A spike in complaints on an otherwise stable list might be a sign there is a problem with address acquisition, if the spike is due to reports from new subscribers. If it came from older subscribers, it might be attributable to content of a particular mailing that was not well received. Perhaps the branding was not recognized or the content was offensive or inappropriate for the list.

The complaint rate is determined by the number of Feedback Messages received over a given time period divided by the number of emails delivered to the associated Mailbox Provider over the same period. It is an obvious and useful metric to track, but there are a few subtle issues to be aware of.

One issue is that Feedback Messages tend to be counted on the day the complaint was sent, which is the day the original message was read by the Recipient. That may not be the same day that the message was sent. A simple example is the fact that a Message Originator that

sends email regularly Monday through Friday will often see a high complaint rate on Saturday. The absolute number of Feedback Messages sent by people catching up with the week's email over the weekend may not be that high. However, since hardly any email is sent on Saturday, a fairly reasonable number of complaints end up being divided by a very small number of total sent emails, possibly even zero, which would break the reporting engine. This can lead to a complaint rate that seems to range anywhere from suspicious to ridiculous. Consequently, large Mailing Lists that are virtually silent on the weekend could end up receiving more complaints on a Saturday than email they sent that day, leading to complaint rates of well over 100%.

Another arithmetic issue to consider is the interaction between the inbox, the bulk folder, and the "This Is Spam" button. If an organization sends a high volume of email that has a terrible reputation, it may end up with perhaps 500 of its 10,000 mails in the inbox and the remaining 9,500 in the bulk folder. If it gets 10 Feedback Messages and divides that by the 10,000 emails it sent, it will get a very respectable 0.1% complaint rate. However, the Mailbox Provider is probably going to calculate the complaint rate by dividing the number of emails delivered to the inbox instead -- giving a 2% complaint rate, which is probably grounds for immediate blocking. So, if one sees a large difference between a complaint rate as reported by a Mailbox Provider or other reputation system and the rate calculated from raw delivery numbers, it is important to look closely at the data.

4.4. Automatically Handling an Incoming Feedback Stream

Even when signing up for a Feedback Loop is partly automated, modifications to it tend to be handled manually. Even something as trivial as changing the Email Address that the Feedback Messages are sent to can be time-consuming and can cause significant overhead to the Feedback Provider. Multiply that by a dozen Feedback Loops, and getting it right the first time can save a lot of time and energy.

Even the smallest of users should create a unique email alias for each Feedback Loop. There are several advantages to this, even if they all deliver to the same person's inbox at first. Sending each Feedback Loop to a unique address makes it immediately clear which Feedback Provider was the source of any given report, even if it is sent from an inconsistent From address. It makes it easy to put lightweight pre-processing in place for a particular Feedback Stream, if needed. It makes it easy to discard Feedback Messages if needed (though only temporarily, as it could be very bad for one's

reputation to miss a changing trend). If a Feedback Consumer needs to scale up, it is easy to point the existing aliases at a Feedback Loop processing engine.

If an organization might possibly scale up appreciably in the future or consider outsourcing its Feedback Loop processing to a third-party Feedback Consumer, it may be even better to create a subdomain for handling Feedback Streams. For example, example.com might use fbl-aol@fbl.example.com to accept its AOL Feedback Loop, allowing it to delegate the whole of @fbl.example.com to a Feedback Loop handling appliance or service, should the need arise.

Small Feedback Consumers, with lists of no more than a few thousand Recipients, or small ISPs with no particular history of problems, should be able to handle feedback reports with little or no automation, as an ARF message should be readable in most mail clients. It may be worthwhile to add some very lightweight processing to the inbound Feedback Messages to make them easier to triage from other email client. For example, arffilter.c [Wise] can annotate the Subject line of inbound Feedback Messages with the IP address being reported, making it easier to see patterns of problems by sorting the messages by Subject line in the mail client. To identify which Recipient is causing the feedback to be sent, small Feedback Consumers should add some of the automation mentioned below that is intended for larger Feedback Consumers.

Larger Feedback Consumers need to be able to automate the handling of Feedback, as it scales beyond the ability of someone to manage manually quite quickly. The main capability a Feedback Loop processor needs is to extract some relevant data from the report, reliably. The most important bits of data tend to be the following:

- o The Recipient of the original email
- o The Mailbox Provider originating sending the Feedback Message (some Feedback Providers operate on behalf of multiple Mailbox Providers)
- o The customer who sent the original email (in the case of an ESP or Mailbox Provider)
- o The campaign and Mailing List that the original email belonged to, if any
- o (Possibly) the IP address from which the original email was sent and any [DKIM] signature domain

The last isn't vital, but may be a useful piece of data in diagnosing delivery problems.

It can be very difficult to extract some of this data without some upfront work before email is sent. Some Feedback Providers will redact the Email Address in the To: header or Recipient Email Addresses anywhere within the message. Some will delete any identifying information they can. It may be possible to identify the End User based on the Message-ID, Subject line, and Received header timestamps, if there is access to the mail server logs, but at best it is painful and time-consuming, and only worth doing in an exceptional case.

The solution is similar to the one used for automated bounce handling using VERP -- embed the information in the email in a way that it is unlikely to be removed by Feedback Providers but is easy to recognize in the email. That information may already be there in a form such as VERP if the Return-Path header is included in the embedded email, or included in one-click unsubscribe links included in the body of the email. If it is not already there, a good place to add the information is in the local part of the Message-ID as that is often used to track the progress of email through delivery. It is often available from log files as well as in the headers of the original message included in the Feedback Message.

There are several good ways to store the mapping between Recipients and identifiers in mail. For a database-backed ESP or bulk sender, a synthesized database primary key can be used. It is very small, and very opaque, and it is not expensive to retrieve the associated data from the main database -- but it is impossible to read by hand. Therefore, it needs automation with access to the core database to map the key onto the actual data.

Recording the required information directly within the email but encrypting it with strong or weak encryption removes the need for database access to extract the data. However, it still does need some automation.

A hybrid approach with the various bits of data stored separately but having some pieces either encrypted or obfuscated is possible by just including a database ID. This can provide a good compromise where most of the data is not immediately obvious to third parties but patterns in it can be recognized by eye. For example, a Message-ID of "esp-423-27-42460@example.com" is opaque to a third party, but someone familiar with the format can tell that it is a Message-ID added by the system. In this case it starts with "esp" followed by three numbers separated by dashes, meaning it is from customer 423, campaign 27, and the Recipient has the database key 42460. Even

decoding this manually, while it may not be possible to identify customer number 423, it is easy to recognize that 10 Feedback Messages in a row relate to the same customer. From experience, it is not unusual for the vast majority of reports at an ESP to be about a very small number of customers, and one learns their customer IDs very quickly.

Once a Message Originator embeds Recipient identifiers in an easily recognizable format in all its mail, it is quite easy for a Feedback Message processor to extract that with a conventional expression match and possibly a couple of database queries. It can then suppress that Email Address and record the customer and campaign for future reporting. In the case where the Feedback Messages are recorded in a ticketing system, it can also annotate the tickets with that data (again, for reporting and trending analysis).

A Feedback Message processor is often bolted onto the side of an already complex bulk mail generator, making it difficult to reliably suppress mail to the Recipient. If the delivery data is stored in a way that makes it easy to convert into the same format as the VERP string used for bounce processing then the Feedback processor can create a "fake" hard bounce and send it to the existing bounce processor, suppressing mail to that address.

Mailbox Providers and Access Providers also need to automate Feedback processing. They are usually less interested in the details about the message and more interested in the IP address and which customer sent it. In most cases, the IP address can be extracted easily from ARF metadata; whereas, in other cases, it may need to be extracted from the Received headers embedded in the included original message. That data can then be used both for automated forwarding of Feedback Messages to the originating customer, if the ISP feels that is appropriate, and also for reporting on complaint levels across the ISP's customer base.

5. Conclusion

Whether you are acting as a Mailbox Provider or a Feedback Consumer, Complaint Feedback processing can be complex and scary -- or, with some intelligence and automation, simple and easy. In either case, it is an important and necessary tool for detecting messaging abuse and ensuring End User satisfaction.

MAAWG encourages all Mailbox Providers to offer Feedback of whatever form is appropriate for their local user base and legal framework, and it encourages all Senders of email to consume and act upon any Feedback available. An actively maintained list of known Feedback Loops can be found at [Wise].

6. Acknowledgments

This document was written within the MAAWG Collaboration Committee. The project was led by John Feaver and Kate Nowrouzi. The primary authors were Steve Atkins, Christine Murphy Borgia, J.D. Falk, John Feaver, Todd Herr, John Levine, Heather Lord, Kate Nowrouzi, and Suresh Ramasubramanian.

The document was edited by John Levine, J.D. Falk, and Linda Marcus. Further editing and formatting required for this version to be published by the IETF was performed by J.D. Falk, with advice from Barry Leiba and Murray Kucherawy.

6.1. About MAAWG

[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 current Feedback Loop practices.

7. Security Considerations

Security and privacy considerations are discussed in many sections of this document, most notably Sections 1, 3.4, and 3.5.

8. Informative References

- [DKIM] Crocker, D., Hansen, T., and M. Kucherawy, "DomainKeys Identified Mail (DKIM) Signatures", RFC 6376, September 2011.
- [DNS] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [DomainKeys] Delany, M., "Domain-Based Email Authentication Using Public Keys Advertised in the DNS (DomainKeys)", RFC 4870, May 2007.
- [MAAWG] Messaging Anit-Abuse Working Group, <<http://www.maawg.org/>>.

- [MAAWG-BCP] MAAWG, "MAAWG Sender Best Communications Practices Executive Summary and MAAWG Sender Best Communications Practices Version 2.0a-Updated", September 2011, <http://www.maawg.org/sites/maawg/files/news/MAAWG_Senders_BCP_Ver2.pdf>.
- [MARF] Shafranovich, Y., Levine, J., and M. Kucherawy, "An Extensible Format for Email Feedback Reports", RFC 5965, August 2010.
- [MIME] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.
- [RFC5598] Crocker, D., "Internet Mail Architecture", RFC 5598, July 2009.
- [SMTP] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.
- [Trust] Crocker, D., Ed., "Trust in Email Begins with Authentication", Issued by the Messaging Anti-Abuse Working Group (MAAWG), June 2008, <http://www.maawg.org/sites/maawg/files/news/MAAWG_Email_Authentication_Paper_2008-07.pdf>.
- [VERP] Wikipedia, "Variable Envelope Return Path", <https://secure.wikimedia.org/wikipedia/en/wiki/Variable_envelope_return_path>.
- [Wise] "arffilter - rewrite ARF reports", <<http://wordtothewise.com/products/arffilter.html>>.

Appendix A. Abuse Reporting Format (ARF)

A.1. A Brief History

The approach used by the first Feedback Loop to be deployed -- the "scomp" system at AOL -- was to send an entire copy of the message to the consumer of the Feedback Loop. It expected that large Feedback Consumers would embed sufficient information in the email so they could identify which Message Recipient had complained.

That worked well enough when there was only a single entity providing feedback, but as other Mailbox Providers started to offer Feedback, it became clear that it would be useful for the Feedback Provider to be able to add some additional information, both machine readable and human readable, to the report. This led to ARF, the Abuse Reporting Format, which quickly became the de facto standard for Feedback Messages.

Today, ARF is used by nearly all Feedback Providers, both within MAAWG and without, constituting the vast majority of all Feedback Messages generated worldwide. ARF is recognized by all MAAWG members that have developed software or services that consume and process Feedback Messages. There are no competing standards for reporting individual messages.

ARF has now been published by the IETF as RFC 5965 [MARF].

A.2. Structure of an ARF Message

An ARF report (Feedback Message) is sent by email, with one message sent for each Spam report made. It consists of three sections, in a standard [MIME] message format called multipart/report.

The first section contains human-readable plaintext, primarily for the benefit of small Feedback Consumers who are handling reports manually. It typically contains boilerplate text explaining that this is a Feedback Message and providing URLs to other data such as contact information for the Feedback Provider or Mailbox Provider that originated the Feedback Message.

The second section contains some machine-readable information, including the version of the ARF protocol used and the type of report it is ("abuse," "fraud," or other label). It also might include some optional information about the email being reported, such as the original Envelope Sender or the time the mail was received. In theory, the information in this section can be used to mechanically route and triage the report, though in current practice most Feedback

Messages are treated identically. As a result, this section is often ignored entirely by Feedback Consumers who prefer to process the third section themselves.

The third section of the report consists of a copy of the original email that the report is about, as a standard [MIME] message/rfc822 attachment. While ideally this would be an unmodified copy of the original email, it is likely that many producers of reports will modify or "redact" some elements of the report, especially the Email Address of the Recipient, due to privacy or other legal concerns.

The strict technical specifications of ARF, as well as some example reports and tools to handle the format, can be found at <http://mipassoc.org/arf/>, [Wise], and in [MARF]

Appendix B. Using DKIM to Route Feedback

Historically, the IP address of the "last hop" -- the MTA that transferred a message into the receiving Mailbox Provider's administrative domain -- was the sole reliable identifier used to denote the source of a message. With the emergence of authentication technologies such as [DKIM], another identifier can now be used; specifically, the authenticated domain associated with a message. This domain is the "d=" value in a DKIM-Signature header field.

In a social or policy context, applying a DKIM signature to a message is tantamount to stating, "I take responsibility for this message". The DKIM signature is most often applied by the author or originator of a message, which may be far upstream of the "last hop" MTA. This is true particularly in cases where the originator's intended Recipient Email Address is configured to forward to another Recipient Email Address. Stories of users who have strung together multiple forwarding accounts are not uncommon, and these users are unable to complain effectively about Spam because their Mailbox Providers cannot easily or reliably follow the path of a message back to the initial originator.

A single DKIM "d=" value may be used across multiple servers with multiple IP addresses. Servers may be added or removed at any time without changing the dynamics of the DKIM signature. When a Feedback Loop is based on the IP address, the Feedback Consumer must contact the Feedback Provider to change its subscription options every time an IP address needs to be added or removed. However, when a Feedback Loop uses DKIM, no reconfiguration is necessary because the signing domain does not change.

One recurring concern with DKIM, however, is that ESPs often send messages addressed with hundreds or thousands of customer domains, yet they want to receive Feedback Messages for all of these domains. This was particularly difficult with [DomainKeys] (the predecessor to DKIM), which tied the "d=" to the "From" header field. DKIM removed this tie, so it is simple for an ESP to use a domain of its own to sign the message and sign up for Feedback regarding all messages signed with that domain. Such a signature may be in addition to, or instead of, signatures from the various client domains. While there are still many unknowns related to reputation (which will be addressed in a future MAAWG document), this is clearly an appropriate use of DKIM to take responsibility (and receive Feedback) for a message.

Appendix C. Unsolicited Feedback

Is it always necessary for a Feedback Consumer to apply for a Feedback Loop or is it permissible for a Feedback Provider to configure a Feedback Loop for a Feedback Consumer without an explicit request? There is continuing debate about whether this is an acceptable practice, and MAAWG is neither endorsing nor condemning such activity at this time.

That said, if a Feedback Provider chooses to send Feedback without being asked first, certain guidelines should be followed. In general, it should make prudent decisions to minimize the negative impact on Mailbox Providers and Access Providers.

C.1. Guidelines

This should only be done for Mailbox and Access Providers.

This should only be done after attempting to contact the provider to ask if it is possible to set up a Feedback Loop via the normal practice.

These Feedback Loops should only be set up to send to the published abuse address from the provider's WHOIS record.

C.2. Pros

Feedback Consumers may not realize they have abuse problems until they begin receiving the spam complaints.

Feedback Consumers may not be aware of Feedback Loops and may appreciate the additional data feed.

Upstream providers have an additional information stream to help them identify problem customers.

Spam coming from a network is abuse; therefore it is appropriate to send reports of the abuse back to the Mailbox Provider or Access Provider. Setting up a Feedback Loop automates the process.

C.3. Cons

It creates confusion for Feedback Consumers if they did not apply and do not understand why they are suddenly receiving complaints.

It can conflict with existing Terms of Service because a new feed of information is available. For example, if a provider has a policy to terminate service after a certain number of abuse complaints, and it starts receiving unexpected Feedback Loop complaints, it may either be forced to terminate customers that did not have a previous issue or be required to update its Terms of Service and Acceptable Use Policy agreements.

Upstream providers do not have access to the mail being sent by their customers, so they cannot tell whether bulk mail complaints constitute a problem.

The listed abuse address may not be the correct place for automated spam complaints to be sent.

The listed abuse address may feed into a ticketing system that is not capable of correctly handling ARF messages.

Feedback Consumers may not be equipped to handle the volume or format of complaints without some warning and preparation.

Author's Address

J.D. Falk (editor)
Messaging Anti-Abuse Working Group
Presidio of San Francisco
P.O. Box 29920
572 B Rucker Street
San Francisco, CA 94129-0920
US

EMail: ietf@cybernothing.org
URI: <http://www.maawg.org/>

