

Individual submission M.  
Kucherawy  
Internet-Draft Cloudmark,  
Inc.  
Intended status: Standards Track February 15,  
2012  
Expires: August 18, 2012

**Email Greylisting: An Applicability Statement for SMTP  
draft-ietf-appsawg-greylisting-03**

Abstract

This memo describes the art of email greylisting, the practice of providing temporarily degraded service to unknown email clients as an anti-abuse mechanism.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 18, 2012.

Copyright Notice

Copyright (c) 2012 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.

Kucherawy  
1]

Expires August 18, 2012

[Page

Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">1.1.</a>	<a href="#">Background . . . . .</a>	<a href="#">3</a>
<a href="#">1.2.</a>	<a href="#">Definitions . . . . .</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Types of Greylisting . . . . .</a>	<a href="#">4</a>
<a href="#">2.1.</a>	<a href="#">Connection-Level Greylisting . . . . .</a>	<a href="#">4</a>
<a href="#">2.2.</a>	<a href="#">SMTP HELO/EHLO Greylisting . . . . .</a>	<a href="#">4</a>
<a href="#">2.3.</a>	<a href="#">SMTP MAIL Greylisting . . . . .</a>	<a href="#">5</a>
<a href="#">2.4.</a>	<a href="#">SMTP RCPT Greylisting . . . . .</a>	<a href="#">5</a>
<a href="#">2.5.</a>	<a href="#">SMTP DATA Greylisting . . . . .</a>	<a href="#">6</a>
<a href="#">2.6.</a>	<a href="#">Exceptions . . . . .</a>	<a href="#">7</a>
<a href="#">3.</a>	<a href="#">Benefits and Costs . . . . .</a>	<a href="#">7</a>
<a href="#">4.</a>	<a href="#">Unintended Consequences . . . . .</a>	<a href="#">7</a>
<a href="#">4.1.</a>	<a href="#">Unintended Delivery Failures . . . . .</a>	<a href="#">8</a>
<a href="#">4.2.</a>	<a href="#">Unintended Client Failures . . . . .</a>	<a href="#">8</a>
<a href="#">4.3.</a>	<a href="#">Address Space Saturation . . . . .</a>	<a href="#">10</a>
<a href="#">5.</a>	<a href="#">Recommendations . . . . .</a>	<a href="#">10</a>
<a href="#">6.</a>	<a href="#">Measuring Effectiveness . . . . .</a>	<a href="#">11</a>
<a href="#">7.</a>	<a href="#">Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">7.1.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">7.2.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">8.</a>	<a href="#">References . . . . .</a>	<a href="#">12</a>
<a href="#">8.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">12</a>
<a href="#">8.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">12</a>
<a href="#">Appendix A.</a>	<a href="#">Acknowledgments . . . . .</a>	<a href="#">13</a>
	<a href="#">Author's Address . . . . .</a>	<a href="#">13</a>

Kucherawy  
2]

Expires August 18, 2012

[Page

## **1. Introduction**

Preferred techniques for handling email abuse explicitly identify good actors and bad actors, giving each significantly differential service. In some cases an actor does not have a known reputation; this can justify providing degraded service, until there is a basis for provider better service. This latter approach known as "greylisting". Broadly, the term refers to any degradation of service for an unknown or suspect source, over a period of time.

The

narrow use of the term refers to generation of an SMTP temporary failure reply code for traffic from such sources. There are diverse implementations of this basic concept, and, predictably therefore, some blurred terminology.

This memo documents common greylisting techniques and discusses their

benefits and costs. It also defines terminology to enable clear distinction and discussion of these techniques.

### **1.1. Background**

For many years, large amounts of spam have been sent through purpose-

built software, or "spamware", that supports only a constrained version of SMTP. In particular, such software does not perform retransmission attempts after receiving an SMTP temporary failure. That is, if the spamware cannot deliver a message, it just goes on

to

the next address in its list since, in spamming, volume counts for far more than reliability. Greylisting exploits this by rejecting mail from unfamiliar sources with a "soft fail" (4xx) [[SMTP](#)] error code, on the theory that standard MTAs will retry, and spamware won't. Another, less well-developed, application of greylisting is to delay mail from newly seen IP addresses on the theory that, if it's a spam source, then by the time it retries, it will appear in a list of sources to be filtered, and the mail will not be accepted.

### **1.2. Definitions**

#### **1.2.1. Keywords**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[KEYWORDS](#)].

#### **1.2.2. E-Mail Architecture Terminology**

Readers need to be familiar with the material and terminology discussed in [[MAIL](#)] and [[EMAIL-ARCH](#)].

Kucherawy  
3]

Expires August 18, 2012

[Page

## **2. Types of Greylisting**

Greylisting is primarily performed at some phase during an SMTP session. A set of attributes about the client-side SMTP server are used for assessing whether to perform greylisting. At its simplest, the attribute is the IP Address of the client and the assessment is whether it has previously connected, recently. More elaborate attribute combinations and more sophisticated assessment, can be performed. The following discussion covers the most common combinations.

### **2.1. Connection-Level Greylisting**

Connection-level greylisting decides whether to accept the (TCP) connection from a "new" [SMTP] client. At this point in the communication between the client and the server, the only information

known to the receiving server is the incoming IP address. This, of course, is often (but not always) translatable into a host name.

The typical application of greylisting here is to keep a record of SMTP client IP addresses and/or host names (collectively, "sources") that have been seen. Such a database acts as a cache of known senders and might or might not expire records after some period. If the source is not in the database, or the record of the source has not reached some required minimum age -- such as 30 minutes since the initial connection attempt -- the server does one of the following, inviting a later retry:

- o returns a 421 SMTP reply, and closes the connection;
- o returns a different 4yz SMTP reply to all further commands in this SMTP session

A useful variant of the basic known/unknown policy is to limit greylisting to those addresses that are on some list of IP addresses known to be affiliated with bad actors. Whereas the simpler policy affects all new connections, including those from good actors, the constrained policy applies greylisting actions only to sites that already have a negative reputation.

### **2.2. SMTP HELO/EHLO Greylisting**

HELO/EHLO greylisting refers to the first [SMTP] command verb in an SMTP session. It includes a single, required parameter that is supposed to contain the client's fully-qualified host name or its literal IP address.

Greylisting implemented at this phase retains a record of sources

Kucherawy  
4]

Expires August 18, 2012

[Page



coupled with HELO/EHLO parameters. It returns 4yz SMTP replies to all commands until the end of the SMTP session if that tuple has not previously been recorded or if the record exists but has not reached some configured minimum age.

### **2.3. SMTP MAIL Greylisting**

MAIL command greylisting refers to the [\[SMTP\]](#) command verb in an SMTP

session that initiates a new transaction. It includes at least one required parameter that indicates the return email address of the message being relayed from the client to the server.

Greylisting implemented at this phase retains a record of sources coupled with return email addresses. It returns 4yz SMTP replies to all commands for the remainder of the SMTP session if that tuple has not previously been recorded or if the record exists but has not met some configured minimum age.

### **2.4. SMTP RCPT Greylisting**

RCPT greylisting refers to the [\[SMTP\]](#) command verb in an SMTP session

that specifies intended recipients of an email transaction. It includes at least one required parameter that indicates the email address of an intended recipient of the message being relayed from the client to the server.

Greylisting implemented at this phase retains a record of tuples that

combines the provided recipient address with any combination of the following:

- o the source, as described above;
- o the return email address;
- o other recipients of the message

If the selected tuple is not found in the database, or if the record is present but has not reached some configured minimum age, the greylisting MTA returns 4yz SMTP replies to all commands for the remainder of the SMTP session.

Note that often a match on a tuple involving the first valid RCPT is sufficient to identify a retry correctly, and further checks can be omitted.

Kucherawy  
5]

Expires August 18, 2012

[Page

## **2.5. SMTP DATA Greylisting**

DATA greylisting refers to the [\[SMTP\]](#) command verb in an SMTP session that transmits the actual message content, as opposed to its envelope details (see [\[MAIL\]](#)).

This type of greylisting can be performed at two places in the SMTP sequence:

1. on receipt of the DATA command, because at that point the entire envelope has been received (i.e., all MAIL and RCPT commands have been issued);
2. on completion of the DATA command, i.e., after the "." that terminates transmission of the message body, since at that point a digest of the message could be computed.

Some implementations do filtering here because there are clients that don't bother checking SMTP reply codes to commands other than DATA.

Numerous greylisting policies are possible at this point. All of them retain a record of tuples that combine the various parts of the SMTP transaction in some combination, including:

- o the source, as described above;
- o the return email address;
- o the recipients of the message, as a set or individually;
- o identifiers in the message, such as the contents of the [RFC5322.From](#) or [RFC5322.To](#) fields;
- o other prominent parts of the content, such as the [RFC5322.Subject](#) field;
- o a digest of some or all of the message content, as a test for uniqueness;

(The last three items in that list are only possible at the end of DATA, not on receipt of the DATA command.)

If the selected tuple is not found in the database, or if the record exists but has not reached some configured minimum age, the greylisting MTA returns 4yz SMTP replies to all commands for the remainder of the SMTP session.

Kucherawy  
6]

Expires August 18, 2012

[Page

## **2.6. Exceptions**

Most greylisting systems provide for an exception mechanism, allowing one to specify IP Addresses, IP Address CIDR blocks, hostnames or domain names that are exempt from greylisting checks and thus whose SMTP client sessions are not subject to such interference.

## **3. Benefits and Costs**

The most obvious benefit with any of the above techniques is that spamware does not retry, and is therefore less likely to succeed, absent a record of a previous delivery attempts.

The most obvious detriment to implementing greylisting is the imposition of delay on legitimate mail. Some popular MTAs do not retry failed delivery attempts for an hour or more, which can cause expensive delays when delivery of mail is timely. Worse, some legitimate MTAs do not retry at all(!) The counterargument to this "false positive" problem is that email has always been a "best-effort" mechanism, and thus this cost is ultimately low in comparison to the cost of dealing with high volumes of unwanted mail. Still, the actual effect of such delays can be significant, such as altering the tone of a multi-participant discussion to a mailing list.

The cache of information stored about client history does not benefit legitimate clients that are already listed for acceptance, when the clients are subjected to any kind of reconfiguration, especially such as network renumbering. To the greylisting implementation such clients are once again unknown, and they will once again be subjected to the delay.

Another obvious cost is for the required database. It has to be large enough to keep the necessary history, and fast enough to avoid excessive inefficiencies in the server's operations. The primary consideration is the maximum age of records in the database. If records age out too soon, then hosts that do retry per [\[SMTP\]](#) will be periodically subjected to greylisting even though they are well-behaved; if records age out after too long a period, then eventually spamware that launches a new campaign will not be identified as "unknown" in this manner, and will not be required to retry.

## **4. Unintended Consequences**

Kucherawy  
7]

Expires August 18, 2012

[Page

#### **4.1. Unintended Delivery Failures**

There are a few failure modes of greylisting that are worth considering. For example, consider an email message intended for user@example.com. The example.com domain is served by two receiving mail servers, one called mail1.example.com and one called mail2.example.com. On the first delivery attempt, mail1.example.com greylists the client, and thus the client places the message in its outgoing queue for later retry. Later, when a retry is attempted, mail2.example.com is selected for the delivery, either because mail1.example.com is unavailable or because a round-robin [\[DNS\]](#) evaluation produces that result. However, the two example.com hosts do not share greylisting databases, so the second host again denies the attempt. Thus, although example.com has sought to improve its email throughput by having two servers, it has in fact amplified the problem of legitimate mail delay introduced by greylisting.

Similarly, consider a site with multiple outbound MTAs that share a common queue. On a first outbound delivery attempt to example.com, the attempt is grey listed. On a later retry, a different outbound MTA is selected, which means example.com sees a different source, and once again greylisting occurs on the same message.

For systems that do DATA-level greylisting, if any part of the message has changed since the first attempt, the tuple constructed might be different than the one for the first attempt, and the delivery is again greylisted. Some MTAs do reformulate portions of the message at submission time and this can produce visible differences for each attempt.

A host that sends mail to a particular destination infrequently might not remain "known" in the receiving server's database and will therefore be greylisted for a high percentage of mail despite possibly being a legitimate sender.

All of these and other similar cases can cause greylisting to be applied improperly to legitimate MTAs multiple times, leading to long delays in delivery or ultimately the return of the message to its sender. Other side effects include out-of-order delivery of related, sequenced messages.

#### **4.2. Unintended Client Failures**

Atypical client behaviours also need to be considered when deploying greylisting.

Some clients do not retry messages for very long periods. Popular open source MTAs implement increasing backoff times when messages

Kucherawy  
8]

Expires August 18, 2012

[Page



receive temporary failure messages, and/or degrade queue priority for very large messages. This means greylisting introduces even more delay for MTAs implementing such schemes, and the delay can become large enough to become a nuisance to users.

Some clients do not retry messages at all. This means greylisting will cause outright delivery failure right away for sources, envelopes, or messages that it has not seen before, regardless of the client attempting the delivery, essentially treating legitimate mail and spam the same.

If a greylisting scheme requires a database record to have reached a certain age rather than merely testing for the presence of the record in the database, and the client has a retry schedule that is too aggressive, the client could be subjected to rate limiting by the MTA independent of the restrictions imposed by greylisting.

Some SMTP implementations make the error of treating all error codes as fatal; that is, a 4yz [SMTP] response is treated as if it were a 5yz response, and the message is returned to the sender as undeliverable. This can result in such things as inadvertent removal from mailing lists in response to the perceived rejections.

Some clients encode message-specific details in the address parameter to the [SMTP] MAIL command. If doing so causes the parameter to change between retry attempts, a greylisting implementation could see it as a new delivery rather than a retry, and disallow the delivery. In such cases, the mail will never be delivered, and will be returned to the sender after the retry timeout expires.

A client subjected to greylisting might improperly move to the next host found in the ordered [DNS] MX record set for the destination domain and re-attempt delivery. This can generate an increase in traffic to those servers. Moreover, it is likely that those servers are somehow related and thus will be able to bypass greylisting, either because the servers are in a bypass list at the final destination or due to the traffic their special relationship with the intended recipient implies.

There are some applications that connect to an SMTP server and simulate a transaction up to the point of sending the RCPT command in an attempt to confirm that an address is valid. Some of these are

legitimate applications (e.g., mailing list servers) and others are automated programs that attempt to ascertain valid addresses to which to send spam (a "directory harvesting" attack). Greylisting can interfere with both instances, with harmful effects on the former.

### 4.3. Address Space Saturation

Greylisting is obviously not a fool-proof solution to avoiding abusive traffic. Bad actors that send mail with just enough frequency to avoid having their records expire will never be caught by this mechanism after the first instance.

Where this is a concern, combining greylisting with some form of reputation service that estimates the likely behaviour for IP addresses that are not intercepted by the greylisting function would be a good choice.

## 5. Recommendations

The following practices are RECOMMENDED based on collected experience:

1. Implement greylisting based on each tuple (IP address, [RFC5321](#).MailFrom, [RFC5321](#).RcptTo). This means that a single message can produce multiple tuples for evaluation as the message can have multiple [RFC5321](#).RcptTo values. After a successful retry, allow all further [SMTP] traffic from the IP address in that tuple.
2. Include a time window within which a retry is considered, and ignored otherwise. The default window SHOULD be from one minute to 24 hours. Retries during the period of this window are refused and otherwise ignored, as they might be immediate retries by the spamware; retries after the window don't benefit the source in case it happens to retry a week later as part of a different mailing. Some sites use a longer minimum time to match common legitimate MTA retry timeouts, but additional benefit from doing so appears unlikely.
3. Include a timeout for database entries, after which entries are deleted even if the associated client has been active. This will subject otherwise-known clients to another round of greylisting. The default SHOULD be one week.
4. For an Administrative Domain (ADMD) all inbound border MTAs listed in the [DNS] SHOULD share a common greylisting database and common greylisting policies. This handles sequences in which a client's retry goes to a different server after the first 4yz reply, and it lets all servers share the list of hosts that did retry successfully.

Kucherawy  
10]

Expires August 18, 2012

[Page

5. To accommodate those senders that have clusters of outgoing mail servers, greylisting servers MAY track CIDR blocks of a size of its own choosing, such as /24, rather than the full IP address.
6. Include a manual override capability for adding specific IP addresses or network blocks that always bypass checks. There

are

legitimate senders that simply don't respond well to greylisting for a variety of reasons, most of which do not conflict with [\[SMTP\]](#). There are also some highly visible online entities such as email service providers that will be certain to retry, and thus those that are known SHOULD be allowed to bypass the

filter.

There is no specific recommendation as to the specific choice of 4yz code to be returned as a result of a greylisting delay. It is possible that some clients treat different 4yz codes differently, but

no data are available on whether or not using 421 versus some other 4yz code is particularly advantageous.

## **6. Measuring Effectiveness**

A few techniques are common when measuring the effectiveness of greylisting in a particular installation:

- o Arrange to log the spam vs. legitimate determinations of messages and what the greylisting decision would have been if enabled;

then

determine whether there is a correlation (and, of course, whether too much legitimate email would also be affected)

- o Continuing from the previous point, query the set of IP addresses subjected to greylisting in any popular [\[DNSBL\]](#) to see if there

is

a strong correlation

## **7. Considerations**

### **7.1. IANA Considerations**

No actions are requested of IANA in this memo.

### **7.2. Security Considerations**

This section discusses potential security issues related to greylisting.

Kucherawy  
11]

Expires August 18, 2012

[Page

### **7.2.1. Exceptions**

The discussion above highlights the fact that, although greylisting provides some obvious and valuable defenses, it can introduce unintentional and detrimental consequences for delivery of legitimate mail. Where timely delivery of email is essential, especially for security-related applications, the possible consequences of such systems need to be carefully considered.

Specific sources can be exempted from greylisting, but of course that means they have elevated privilege in terms of access to the mailboxes on the greylisting system, and malefactors can seek to exploit this.

### **7.2.2. Database**

The database that has to be maintained as part of any greylisting system will grow as the diversity of its SMTP clients hosts grows, and of course is larger in general depending on the nature of the tuple stored about each delivery attempt. Even with a record aging policy in place, such a database can grow large enough to interfere with the system hosting it, or at least to a point at which greylisting service is degraded. Moreover, an attacker knowing which greylisting scheme is in use could rotate parameters of SMTP clients under its control, in an attempt to inflate the database to the point of denial-of-service.

Implementers could consider configuring an appropriate failure policy so that something locally acceptable happens when the database is attacked or otherwise unavailable.

## **8. References**

### **8.1. Normative References**

[KEYWORDS]

Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[SMTP]

Klensin, J., "Simple Mail Transfer Protocol", [RFC 5321](#), October 2008.

### **8.2. Informative References**

[DNS]

Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.

Kucherawy  
12]

Expires August 18, 2012

[Page



Internet-Draft  
2012

Greylisting

February

[DNSBL] Levine, J., "DNS Blacklists and Whitelists", [RFC 5782](#),  
February 2010.

[EMAIL-ARCH]  
Crocker, D., "Internet Mail Architecture", [RFC 5598](#),  
October 2008.

[MAIL] Resnick, P., Ed., "Internet Message Format", [RFC 5322](#),  
October 2008.

### [Appendix A](#). Acknowledgments

The author wishes to acknowledge Mike Adkins, Steve Atkins, Mihai Costea, Dave Crocker, Peter J. Holzer, John Levine, Chris Lewis, Jose-Marcio Martins da Cruz, S. Moonesamy, Suresh Ramasubramanian, Mark Risher, Jordan Rosenwald, Gregory Shapiro, Joe Sniderman, and Roland Turner for their contributions to this memo.

#### Author's Address

Murray S. Kucherawy  
Cloudmark, Inc.  
128 King St., 2nd Floor  
San Francisco, CA 94107  
US

Phone: +1 415 946 3800  
Email: [msk@cloudmark.com](mailto:msk@cloudmark.com)

