

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 1, 2010

N. Chrobok, Ed.
University of Otago, NZ
January 28, 2010

General Delivery Service Extension for SMTP

Abstract

This memo defines an extension to the SMTP service that changes the classical push based delivery of email-messages to a pull-based system. Instead of simply forwarding an email message to the receiving mail-host, a sending host notifies the receiving host instead. The receiving host decides if it wants to receive the message and retrieves the message using a pull-based technique or simply drops the notification. This extension is effective against spam email-messages that are distributed by botnet PC's.

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 1, 2010.

Table of Contents

- 1. Introduction 3
 - 1.1. Requirements Language 3
- 2. Framework for the General Delivery extension 3
- 3. The General Delivery service extension 3
- 4. Definition 5
- 5. The unique identifier 5
- 6. The GDEL command 5
 - 6.1. RMTA action on receipt of the GDEL command 6
 - 6.2. SMTA action on receiving response to the GDEL command . . 6
- 7. The RETR Command 7
 - 7.1. SMTA action on receipt of the RETR command 8
 - 7.2. RMTA action on receiving response to the RETR command . . 8
- 8. General Delivery replacing classical SMTP 9
- 9. General Delivery and Messages with a Null Reverse-Path 9
- 10. Minimal Usage 9
- 11. Example 9
- 12. Security Considerations 11
- 13. References 12
 - 13.1. Normative References 12
 - 13.2. Informative References 12
- Author's Address 12
- Intellectual Property and Copyright Statements 14

1. Introduction

Most of the emails sent nowadays are spam-messages originating from the botnets. These hijacked PC's are used to massively distribute spam worldwide. The SMTP-service provides no effective way of preventing bots from sending their unwanted junk-mail.

Therefore this memo introduces the General Delivery extension. Unlike classical SMTP it is not a push-based but a pull based protocol. As described in [x] the sending mail transfer agent (SMTA) should notify the RMTA about a message before disconnecting instead of simply delivering it. The RMTA may reconnect to the SMTA and retrieve the message.

1.1. Requirements Language

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 RFC 2119.

2. Framework for the General Delivery extension

The following service extension is therefore defined:

- (1) the name of the SMTP service extension is "General Delivery";
- (2) the EHLO keyword values associated with this extension are "GDEL" and "RETR" with no associated parameters;
- (3) the additional verb GEDEL with a single parameter that specifies an unique message identifier is used to start a notification;
- (4) the additional verb RETR with a single parameter specifying a unique message identifier is used to start the retrieval process;

The remainder of this memo specifies how support for the extension affects the behavior of an SMTP client and server.

3. The General Delivery service extension

SMTP is still one of the most used services of the internet. However, out-of-the-box it does not provide any means to prevent the abuse the internet mail service. One of the greatest problems nowadays is spam. These unsolicited email messages flood the inboxes of internet users. Most of the spam originates from the botnets, huge networks of hijacked PC's that are used for

distributing spam.

Unfortunately the SMTP-service provides no effective way of stemming the spread of these unwanted mails. Among the many ways to prevent spam the most effective way is filtering. However this method forces a receiving mail transfer agent (RMTA) to accept all incoming messages, even if they will be afterwards deleted. The responsibility for an email message lies on the receiving side which allows spammers to distribute their fraudulent messages very cheap.

This problem has been addressed in the design of the General Delivery service extension. Instead of using the classical functionality of SMTP, which pushes all email messages from the SMTA to the RMTP, the General Delivery service extension uses a pull based approach.

In a first step, the SMTA generates an unique identifier for an outgoing email-message, stores both message and unique identifier and instead sends a notification to the RMTA using the unique identifier. Then the connection is dropped.




Figure 1

In a second step, the RMTA might decide to retrieve the message from the SMTA. If so it reconnects and retrieves the message using the unique identifier. The SMTP then forwards the email as in classical SMTP.




Figure 2

Using this method the responsibility for storing an email message is shifted from the receiver to the sender. Especially bots cannot simply forward thousands of email-messages anymore. They have to provide a service for messages to be retrieved. So the classical

'hit-and-run'-tactic spammers use cannot be used which makes it harder for them to distribute their spam.

The General Delivery service extension, being an extension to SMTP requires no end-user interaction for the service to work. However to be effective against botnet-spam it, all mail-services must use these extension.

4. Definition

General Delivery means that using an SMTP or ESMTP connection, the RMTA may accept no pushed messages from an SMTA. Instead it accepts only notifications that a message could be retrieved at a later time from the SMTA. As General Delivery uses two connections, one for notifying, the other one for retrieving the terms client and server are not used. Instead the terms SMTA and RMTA are used.

The RMTA is the node that is responding to a GDEL command and initiates the RETR command.

The SMTA is the node that is responding to a RETR command and initiates the GDEL command.

5. The unique identifier

In order for General Delivery to work, a unique identifier must be provided for every email-message an SMTA wants to send. This unique identifier must contain a randomly generated sequence, the character '@' and the domain name for which the SMTA provides the mail service. The specified domain name is REQUIRED to be a fully qualified domain name, which may refer to a A or AAAAA pointer in the DNS.

```
8f1b0fb0a9f67ffaa43a83cad28435ca@somedomain.com
```

The structure of the randomly generated sequence is free, there is no minimum or maximum number of characters it should contain. However it is recommended that the randomly generated sequence is long enough to avoid name-collisions.

6. The GDEL command

The GDEL command is issued by the SMTA when it wishes to notify a RMTA of an email-message that it stores. The syntax of the command is as follows:

GDEL <unique identifier><CRLF>

This command may be issued at any time once a session is established, as long as there is not a transaction occurring. Therefore issuing this command is illegal between the MAIL FROM: command and the end of the DATA command and responses.

Preceding this notification the SMTA must store the email message and generate a unique identifier. The SMTA is required to ensure that the email-message can be retrieved using the unique identifier.

The SMTA is required to generate one unique identifier for every message. It must not store several messages using only one unique identifier, even if there are more messages for the same RMTA or the same recipient. In this case it has to generate a unique identifier for every message.

6.1. RMTA action on receipt of the GDEL command

When the RMTA receives the GDEL command, it should make a local decision if it honors the request. If not, it should return the correct error codes to the SMTA. In the case the RMTA accepts the notification, it must store the unique identifier and the IP-address of the connecting SMTA and inform the agent that retrieves the message. During this process the RMTA may also store the name of the connecting host (the name it identified itself using the HELO or EHLO command).

The valid return codes for this command are:

```
250 OK, queuing <unique identifier> for retrieval
500 Syntax Error
501 Syntax Error in Parameters
503 Bad sequence of commands
```

6.2. SMTA action on receiving response to the GDEL command

If the RMTA responded with a successful reply (250), the SMTA should close the connection using the QUIT command. In the case it has several messages for recipients at the domain the RMTA provides mail-service it may initialize another GDEL command.

In the case of the 500 level error codes (500, 501 or 503) the SMTA should assume that the RMTA does not support the General Delivery service extension. Thus it should try to send the email using classical SMTP. However it should be considered that a RMTA

supporting General Delivery might not accept classical SMTP connections.

7. The RETR Command

The RETR command is issued by the RMTA if it wishes to retrieve an email-message from a SMTA using a unique identifier. The syntax of the command is as follows:

```
RETR <unique identifier><CRLF>
```

This command may be issued at any time once a session is established, as long as there is not a transaction occurring. Like for the GDEL command issuing RETR is illegal between the MAIL FROM: command and the end of the DATA command and responses.

Preceding this command the RMTA has to connect to the SMTA. Based on the unique identifier, which holds a full qualified domain name and the sender's IP-address which was stored during the notification process, it should establish a connection to the SMTA using it's IP-address. It may use the name of the SMTA (HELO/EHLO) if it has stored it, however there is no guarantee that this name leads to a valid mail-service.

An RMTA that got a notification from an SMTA is free to decide if it wants to retrieve the message from this host. The decision might be based on the use of services like blackhole-lists, whitelists, the sender policy framework [also see X.] or any other means. The algorithm used to make this decision is up to the implementation and therefore not part of this RFC. However it is strongly suggested to use methods to identify if an SMTA is a legitimate mailing host or not.

The time that passes between notification and retrieval is up to the RMTA. The RMTA may retrieve the message immediately or it may let pass a certain period of time until it starts the retrieval process. For end-user convenience it is suggested that this period is appropriate to keep up a reliable internet mail-service.

An SMTA has to provide a service for a message to be retrieved. However as it is possible that a message might not be retrieved at all, a maximum period of time that a message needs to be available should be defined. During this period the SMTA must be able to provide the message for retrieval. As it is possible that an RMTA cannot retrieve a message due to technical difficulties an SMTA must provide the message for the minimum period of 48 hours (2 days). After this period the SMTA should delete the received message and may

send a notification

The valid return codes for this command are:

```
250 OK, preparing to send message for <unique identifier>
490 Error, Invalid unique identifier, nothing to retrieve
500 Syntax Error
501 Syntax Error in Parameters
503 Bad sequence of commands
```

7.1. SMTA action on receipt of the RETR command

When an SMTA receives the RETR command with a unique identifier, it searches for this unique identifier in its local database. If the the unique identifier is valid and the corresponding email-message is available, it sends a positive reply code and switches it's internal state-table into sending mode.

From now on an SMTA must use the SMTP-sending procedure (RFC 5321) starting with MAIL FROM: followed by RCPT TO: and DATA and finishing with <CRLF>.<CRLF>.

After receiving the reply code from the RMTA, the SMTA changes it's internal state back into receiving mode. An SMTA must be able to successively process several RETR requests during a session.

Getting a successful reply code (250 OK) from the RMTA, an SMTA must make sure, that the message cannot be retrieved again by using it's corresponding unique identifier. Thus it should delete the stored message and all references (i.e. unique identifiers) to this message.

In the case of a an reply code that indicates an error, the SMTA should not try to send the message again. Instead it should wait for the RMTA to send another RETR command. In any case a not yet retrieved message should be provided until the end of the defined period of time (see above).

7.2. RMTA action on receiving response to the RETR command

After getting a positive reply code (250 OK) to the RETR command, an RMTA must immediately change it's internal state table into receiving mode. The following communication between the SMTA and the RMTA follows the SMTP sending procedure (RFC 5321). All the commands and reply codes must be used according to the description in this RFC.

Following the reply code to the SMTA's <CRLF>.<CRLF> the RMTA must

change it's internal state table back into sending mode. The RMTA should be able to successively send several RETR requests during a session.

8. General Delivery replacing classical SMTP

To be effective against botnet spam, the General Delivery service extension should eventually replace classical email-delivery using the push-based method of SMTP. However it might take a long time until all mail-services adopt a pull-based approach. Being a extension to SMTP, the General Delivery service extension allows mail-hosts to provide both, push and pull-based email-delivery. This makes a transitional phase for the introduction of General Delivery very easy.

It could be possible, that at a yet undefined date in the future the pull-based approach replaces the classical push-based SMTP. From this time mail-host may deny mail-delivery using push-based SMTP. RMTA therefore should react reasonable to SMTAs trying to use this legacy system by sending correct reply codes. These error codes should be in the 400-range as they are then permanent errors. A valid response for denying push-based email could be return code 450 Requested action not taken.

9. General Delivery and Messages with a Null Reverse-Path

Messages with a Null Reverse-Path such as non-delivery notifications, other status delivery notifications (DSNs, RFC 3461) or Message Disposition Notifications (MDNs, RFC 3798) which are all notifications about a previous message should be processed like other messages. This makes sure that the General Delivery service extension does not break functionality of internet-mail in regard to it's reliability and robustness.

10. Minimal Usage

There is no minimal usage for an SMTP-host implementing the General Delivery service extension. All hosts supporting this extension must implement both the GDEL and the RETR command.

11. Example

The following example illustrates the use of General Delivery with some permanent and temporary failures.

Notification

```
RMTA: <wait for connection on TCP port 25>
SMTA: <open connection to RMTA>
RMTA: 220 mailhub6.otago.ac.nz ESMTP Sendmail 8.13.8/8.13.8
SMTA: EHLO mail.nattl.at
RMTA: 250-mailhub6.otago.ac.nz
RMTA: 250-HELP
RMTA: 250 RETR
SMTA: MAIL FROM: natascha@nattl.at
RMTA: 250 OK
SMTA: GDEL 1234567@nattl.at
RMTA: 503 Bad sequence of commands
SMTA: RSET
RMTA: 250 OK
SMTA: GDEL
RMTA: 500 Syntax Error
SMTA: GDEL 1234567
RMTA: 501 Syntax Error in Parameters
...

SMTA: GDEL 1234567@nattl.at
RMTA: 250 OK
...
SMTA: QUIT
RMTA: 250 Goodbye
```

Retrieval

```
SMTA: <wait for connection on TCP port 25>
RMTA: <open connection to RMTA>
SMTA: 220 mail.nattl.at, Envelope Server 1.1
RMTA: EHLO mailhub6.otago.ac.nz
SMTA: 250-mail.nattl.at
SMTA: 250-HELP
SMTA: 250 RETR
RMTA: RETR
SMTA: 500 Syntax Error
RMTA: RETR 1234567
SMTA: 501 Syntax Error in Parameters
RMTA: RETR 7654321@nattl.at
SMTA: 490 Error, nothing to retrieve
...

RMTA: RETR 1234567@nattl.at
SMTA: 250 OK
SMTA: MAIL FROM: natascha@nattl.at
RMTA: 250 OK
SMTA: RCPT TO: nat@cs.otago.ac.nz
RMTA: 250 OK
SMTA: DATA
RMTA: 354 Start mail input, end with <CRLF>.<CRLF>
SMTA: Hello World!
SMTA: <CRLF>
SMTA: .<CRLF>
RMTA: 250 OK
...

RMTA: QUIT
SMTA: 250 Goodbye
```

12. Security Considerations

It might be possible that an attacker tries to retrieve email-messages from a mail-server by simply guessing the unique identifiers. Although it is quite unlikely that such an attempt might be successful if the mail-host uses an appropriate unique identifier with sufficient length there still is the chance that messages could be retrieved. Therefore mail-hosts may implement methods to prevent this kind of attack. This could be achieved by storing the IP address of the notified host or by making queries to the DNS. The detailed method how to prevent the illegitimate

retrieval is up to the implementer of the mail-host.

Another vulnerability of the General Delivery service extension is to use the retrieval process to make a distributed denial of service (DDOS) attack to a mail-service. An attacker could notify a great number of mail-hosts using the GDEL command, pretending to be the target mail-host. As all those mail-hosts might try to retrieve the message using RETR from the target, this could eventually lead to a denial of service.

This vulnerability can be reduced to a minimum risk by using the Sender Policy Framework (RFC 4408) to identify notifying hosts. By providing an spf-record for valid mail-hosts, RMTAs could query the DNS if a notification is from legitimate SMTA. In case of a negative answer an RMTA could drop the notification. It is strongly suggested to use the General Delivery service extension in conjunction with the Sender Policy Framework. A positive side-effect of using the Sender Policy Framework is that it is harder for bots to send messages as they would need a valid spf-record for their messages being accepted.

13. References

13.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

13.2. Informative References

[RFC3798] Hansen, T. and G. Vaudreuil, "Message Disposition Notification", RFC 3798, May 2004.

[RFC4361] Lemon, T. and B. Sommerfeld, "Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)", RFC 4361, February 2006.

[RFC4408] Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", RFC 4408, April 2006.

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.

Author's Address

Natascha Chrobok (editor)
University of Otago, NZ
Dunedin,
NZ

Phone:
Email: nat@cs.otago.ac.nz

Full Copyright Statement

Copyright (C) The IETF Trust (2010).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.