Sieve Notification Mechanism: mailto
draft-ietf-sieve-notify-mailto-10

Status of this Memo

Abstract

This document describes a profile of the Sieve extension for notifications, to allow notifications to be sent by electronic mail.

Table of Contents

---

1. Introduction

1.1. Overview

The [Notify] (Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and T. Martin, "Sieve Extension: Notifications," December 2007.) extension to the [Sieve] (Guenther, P., Ed. and T. Showalter, Ed., "Sieve: An Email Filtering Language," January 2008.) mail filtering language is a framework for providing notifications by employing URIs to specify the notification mechanism. This document defines how [mailto] (Hoffman, P., Masinter, L., and J. Zawinski, "The mailto URL scheme," July 1998.) URIs are used to generate notifications by e-mail.

1.2. Conventions used in this document

Conventions for notations are as in [Sieve] (Guenther, P., Ed. and T. Showalter, Ed., "Sieve: An Email Filtering Language," January 2008.) section 1.1, including the use of [Kwds] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.). 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 [Kwds] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.).

---

2. Definition

The mailto mechanism results in the sending of a new email message (a "notification message") to notify a recipient about a "triggering message".

2.1. Notify parameter "method"

The mailto notification mechanism uses standard mailto URIs as specified in [mailto] (Hoffman, P., Masinter, L., and J. Zawinski, "The mailto URL scheme," July 1998.). mailto URIs may contain header fields consisting of a header name and value. These header fields are called "URI headers" to distinguish them from "message headers".

2.2. Test notify_method_capability

The notify_method_capability test for "online" may return "yes" or "no" only if the Sieve processor can determine with certainty whether or not the recipients of the notification message are online and logged in. Otherwise, the test returns "maybe" for this notification method.

2.3. Notify tag ":from"

The :from tag overrides the default sender of the notification message. "Sender", here, refers to the value used in the [RFC5322] (Resnick, P., Ed., "Internet Message Format," October 2008.) "From" header. Implementations MAY also use this value in the [RFC5321] (Klensin, J.,

[Ed., "Simple Mail Transfer Protocol," October 2008.)](#) "MAIL FROM"
command (the "envelope sender"), or they may prefer to establish a
mailbox that receives bounces from notification messages.

2.4. Notify tag ":importance"

The :importance tag has no special meaning for this notification
mechanism, and this specification puts no restriction on its use.
Implementations MAY use the value of :importance to set a priority or
importance indication on the notification message (perhaps a visual
indication, or perhaps making use of one of the non-standard but
commonly used message headers).

2.5. Notify tag ":options"

This tag is not used by the mailto method.

2.6. Notify tag ":message"

The value of this tag, if it is present, is used as the subject of the
notification message, and overrides all other mechanisms for
determining the subject (as described below). Its value SHOULD NOT
normally be truncated, though it may be sensible to truncate an
excessively long value.

2.7. Other Definitions

Because the receipt of an email message is generating another email
message, implementations MUST take steps to avoid mail loops. The
REQUIRED inclusion of an "Auto-Submitted:" field, as described in the
message composition guidelines, will also help in loop detection and
avoidance.
Implementations SHOULD NOT trigger notifications for messages
containing "Auto-Submitted:" header fields with any value other than
"No".
Implementations MUST allow messages with empty envelope senders to
trigger notifications.
Because this notification method uses a store-and-forward system for
delivery of the notification message, the Sieve processor should not

have a need to retry notifications. Therefore, implementations of this method SHOULD use normal mechanisms for submitting SMTP messages and for retrying the initial submission. Once the notification message is submitted, implementations MUST NOT resubmit it, as this is likely to result in multiple notifications, and increases the danger of message loops.

The overall notification message is composed using the following guidelines (see [RFC5322] (Resnick, P., Ed., "Internet Message Format," October 2008.) for references to message header fields):

*If the envelope sender of the triggering message is empty, the envelope sender of the notification message MUST be empty as well, to avoid message loops. Otherwise, the envelope sender of the notification message SHOULD be set to the value of the ":from" parameter to the notify action, if one is specified, has email address syntax and is valid according to the implementation specific security checks (see Section 3.3 of [Notify] (Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and T. Martin, "Sieve Extension: Notifications," December 2007.)). If ":from" is not specified or is not valid, the envelope sender of the notification message SHOULD be set either to the envelope "to" field from the triggering message, as used by Sieve, or to an email address associated with the notification system, at the discretion of the implementation. This MUST NOT be overridden by a "from" URI header, and any such URI header MUST be ignored.

*The envelope recipient(s) of the notification message SHOULD be set to the address(es) specified in the URI (including any URI headers where the hname is "to" or "cc").

*The header field "Auto-Submitted: auto-notified" MUST be included in the notification message (see Section 2.7.1 (The Auto-Submitted header field)). This is to reduce the likelihood of message loops, by tagging this as an automatically generated message. Among other results, it will inform other notification systems not to generate further notifications. mailto URI headers with hname "auto-submitted" are considered unsafe and MUST be ignored.

*The "From:" header field of the notification message SHOULD be set to the value of the ":from" parameter to the notify action, if one is specified, has email address syntax and is valid according to the implementation specific security checks (see Section 3.3 of [Notify] (Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and T. Martin, "Sieve Extension: Notifications," December 2007.)). If ":from" is not specified or is not valid, the "From:" header field of the notification message SHOULD be set either to the envelope "to" field from the triggering message, as used by Sieve, or to an email address associated with

the notification system, at the discretion of the implementation.
This MUST NOT be overridden by a "from" URI header, and any such
URI header MUST be ignored.

*The "To:" header field of the notification message SHOULD be set
 to the address(es) specified in the URI (including any URI
 headers where the hname is "to").

*The "Subject:" field of the notification message SHOULD contain
 the value defined by the :message notify tag, as described in
 [Notify] (Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and
 T. Martin, "Sieve Extension: Notifications," December 2007.). If
 there is no :message tag and there is a "subject" header on the
 URI, then that value SHOULD be used. If that is also absent, the
 subject SHOULD be retained from the triggering message. Note that
 Sieve [Variables] (Homme, K., "Sieve Extension: Variables,"
 January 2008.) can be used to advantage here, as shown in the
 example in Section 3 (Examples).

*The "References:" field of the notification message MAY be set to
 refer to the triggering message, and MAY include references from
 the triggering message.

*If the mailto URI contains a "body" header, the value of that
 header SHOULD be used as the body of the notification message. If
 there is no "body" header, it is up to the implementation whether
 to leave the body empty or to use an excerpt of the original
 message.

*The "Received:" fields from the triggering message MAY be
 retained in the notification message, as these could provide
 useful trace/history/diagnostic information. The "Auto-Submitted"
 header field MUST be placed above these (see Section 2.7.1 (The
 Auto-Submitted header field)). URI headers with hname "received"
 are considered unsafe, and MUST be ignored.

*Other header fields of the notification message that are normally
 related to an individual new message (such as "Message-ID" and
 "Date") are generated for the notification message in the normal
 manner, and MUST NOT be copied from the triggering message. Any
 URI headers with those names MUST be ignored. Further, the "Date"
 header serves as the notification timestamp defined in [Notify]
 (Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and T. Martin,
 "Sieve Extension: Notifications," December 2007.).

*All other header fields of the notification message either are as
 specified by URI headers, or have implementation-specific values;
 their values are not defined here. It is suggested that the
 implementation capitalize the first letter of URI headers and add

a space character after the colon between the mail header name
and value when adding URI headers to the message, to be
consistent with common practice in email headers.

2.7.1. The Auto-Submitted header field

The header field "Auto-Submitted: auto-notified" MUST be included in
the notification message (see [RFC3834] (Moore, K., "Recommendations
for Automatic Responses to Electronic Mail," August 2004.)). The "Auto-
Submitted" header field is considered a "trace field", similar to
"Received" header fields (see [RFC5321] (Klensin, J., Ed., "Simple Mail
Transfer Protocol," October 2008.)). If the implementation retains the
"Received" fields from the triggering message (see above), the "Auto-
Submitted" field MUST be placed above those "Received" fields, serving
as a boundary between the ones from the triggering message and those
that will be part of the notification message.
The auto-notified Auto-Submitted field MUST include one or both of the
following parameters:

   *owner-email - specifies an email address of the owner of the
    Sieve script that generated this notification. If specified, it
    might be used to identify or contact the script's owner. The
    parameter attribute is "owner-email", and the parameter value is
    a quoted string containing an email address, as defined by "addr-
    spec" in [RFC5322] (Resnick, P., Ed., "Internet Message Format,"
    October 2008.). Example:

      Auto-Submitted: auto-notified; owner-email="me@example.com"

   *owner-token - specifies an opaque token that the administrative
    domain of the owner of the Sieve script that generated this
    notification can identify the owner with. This might be used to
    allow identification of the owner while protecting the owner's
    privacy. The parameter attribute is "owner-token", and the
    parameter value is as defined by "token" in [RFC3834] (Moore, K.,
    "Recommendations for Automatic Responses to Electronic Mail,"
    August 2004.). Example:

      Auto-Submitted: auto-notified; owner-token=af3NN2pK5dDXI0W

See Section 5 (Security Considerations) for discussion of possible uses
of these parameters.

3. Examples

Triggering message (received by recipient@example.org):

```
Return-Path: <knitting-bounces@example.com>
Received: from mail.example.com by mail.example.org
  for <recipient@example.org>; Wed, 7 Dec 2005 05:08:02 -0500
Received: from hobbies.example.com by mail.example.com
  for <knitting@example.com>; Wed, 7 Dec 2005 02:00:26 -0800
Message-ID: <1234567.89ABCDEF@example.com>
Date: Wed, 07 Dec 2005 10:59:19 +0100
Precedence: list
List-Id: Knitting Mailing List <knitting.example.com>
Sender: knitting-bounces@example.com
Errors-To: knitting-bounces@example.com
From: "Jeff Smith" <jeff@hobbies.example.com>
To: "Knitting Mailing List" <knitting@example.com>
Subject: [Knitting] A new sweater

I just finished a great new sweater!
```

Sieve script (run on behalf of recipient@example.org):

```
require ["notify", "variables"];

if header :contains "list-id" "knitting.example.com" {
  if header :matches "Subject" "[*] *" {
    notify :message "From ${1} list: ${2}"
         :importance "3"
         "mailto:0123456789@sms.example.net?to=backup@example.com";
  }
}
```

Notification message:

```
Auto-Submitted: auto-notified; owner-email="recipient@example.org"
Received: from mail.example.com by mail.example.org
  for <recipient@example.org>; Wed, 7 Dec 2005 05:08:02 -0500
Received: from hobbies.example.com by mail.example.com
  for <knitting@example.com>; Wed, 7 Dec 2005 02:00:26 -0800
Date: Wed, 7 Dec 2005 05:08:55 -0500
Message-ID: <A2299BB.FF7788@example.org>
From: recipient@example.org
To: 0123456789@sms.example.net, backup@example.com
Subject: From Knitting list: A new sweater
```

Note that:

> *Fields such as "Message-ID:" and "Date:" were generated afresh
> for the notification message, and do not relate to the triggering
> message.

> *Additional "Received:" fields will be added to the notification
> message in transit; the ones shown were copied from the
> triggering message. New ones will be added above the "Auto-
> Submitted:" field.

> *If this message should appear at the mail.example.org server
> again, the server can use the presence of a "mail.example.org"
> received line to recognize that. The Auto-Submitted header field
> is also present to tell the server to avoid sending another
> notification, and it includes an optional owner-email parameter
> for identification.

---

4. Internationalization Considerations

This specification introduces no specific internationalization issues
that are not already addressed in [Sieve] (Guenther, P., Ed. and T.
Showalter, Ed., "Sieve: An Email Filtering Language," January 2008.)
and in [Notify] (Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and
T. Martin, "Sieve Extension: Notifications," December 2007.).

---

5. Security Considerations

Sending a notification is comparable with forwarding mail to the
notification recipient. Care must be taken when forwarding mail
automatically, to ensure that confidential information is not sent into
an insecure environment.
The automated sending of email messages exposes the system to mail
loops, which can cause operational problems. Implementations of this
specification MUST protect themselves against mail loops; see
Section 2.7 (Other Definitions) for discussion of this and some
suggestions. Other possible mitigations for mail loops involve types of
service limitations. For example, the number of notifications generated
for a single user might be limited to no more than, say, 30 in a 60-
minute period. Of course, this technique presents its own problems, in
that the actual rate limit must be selected carefully, to allow most
legitimate situations in the given environment, and even with careful

selection it's inevitable that there will be false positives -- and false negatives.

Ultimately, human intervention may be necessary to re-enable notifications that have been disabled because a loop was detected, or to terminate a very slow loop that's under the automatic-detection radar. Administrative mechanisms MUST be available to handle these sorts of situations.

Email addresses specified as recipients of notifications might not be owned by the entity that owns the Sieve script. As a result, a notification recipient could wind up as the target of unwanted notifications, either through intent (using scripts to mount a mail-bomb attack) or by accident (an address was mistyped or has been reassigned). The situation is arguably no worse than any other in which a recipient gets unwanted email, and some of the same mechanisms can be used in this case. But those deploying this extension have to be aware of the potential extra problems here, where scripts might be created through means that do not adequately validate email addresses, and such scripts might then be forgotten and left to run indefinitely.

In particular, note that the Auto-Submitted header field is required to include a value that a recipient can use when contacting the source domain of the notification message (see Section 2.7.1 (The Auto-Submitted header field)). That value will allow the domain to track down the script's owner and have the script corrected or disabled. Domains that enable this extension MUST be prepared to respond to such complaints, in order to limit the damage caused by a faulty script.

Problems can also show up if notification messages are sent to a gateway into another service, such as SMS. Information from the email message is often lost in the gateway translation, and in this case critical information needed to avoid loops, to contact the script owner, and to resolve other problems might be lost. Developers of email gateways should consider these issues, and try to preseve as much information as possible, including what appears in email trace headers and Auto-Submitted.

Additional security considerations are discussed in [Sieve] (Guenther, P., Ed. and T. Showalter, Ed., "Sieve: An Email Filtering Language," January 2008.) and in [Notify] (Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and T. Martin, "Sieve Extension: Notifications," December 2007.).

---

6. IANA Considerations

6.1. Registration of notification mechanism

The following template specifies the IANA registration of the Sieve notification mechanism specified in this document:
To: iana@iana.org
Subject: Registration of new Sieve notification mechanism
Mechanism name: mailto
Mechanism URI: RFC2368
Mechanism-specific tags: none
Standards Track/IESG-approved experimental RFC number: this RFC
Person and email address to contact for further information:
    Michael Haardt <michael.haardt@freenet.ag>
This information should be added to the list of sieve notification mechanisms given on http://www.iana.org/assignments/sieve-notification.

6.2. New registry for Auto-Submitted header field keywords

Because [RFC3834] (Moore, K., "Recommendations for Automatic Responses to Electronic Mail," August 2004.) does not define a registry for new keywords used in the Auto-Submitted header field, we define one here, to be created as http://www.iana.org/assignments/auto-submitted-keywords. Keywords are registered using the "Specification Required" policy [IANA] (Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," May 2008.).
This defines the template to be used to register new keywords. Initial entries to this registry follow in Section 6.3 (Initial registration of Auto-Submitted header field keywords).
To: iana@iana.org
Subject: Registration of new auto-submitted header field keyword
Keyword value: [the text value of the field]
Description: [a brief explanation of the purpose of this value]
Parameters: [list any keyword-specific parameters, specify their meanings, specify whether they are required or optional; use "none" if there are none]
Standards Track/IESG-approved experimental RFC number: [identifies the specification that defines the value being registered]
Contact: [name and email address to contact for further information]

6.3. Initial registration of Auto-Submitted header field keywords

The following are the initial keywords to be registered for the Auto-Submitted header field, to be entered in http://www.iana.org/assignments/auto-submitted-keywords.

Keyword value: no
Description: Indicates that a message was NOT automatically generated,
but was created by a human. It is the equivalent to the absence of an
Auto-Submitted header altogether.
Parameters: none
Standards Track/IESG-approved experimental RFC number: RFC3834
Contact: Keith Moore <moore@network-heretics.com>
Keyword value: auto-generated
Description: Indicates that a message was generated by an automatic
process, and is not a direct response to another message.
Parameters: none
Standards Track/IESG-approved experimental RFC number: RFC3834
Contact: Keith Moore <moore@network-heretics.com>
Keyword value: auto-replied
Description: Indicates that a message was automatically generated as a
direct response to another message.
Parameters: none
Standards Track/IESG-approved experimental RFC number: RFC3834
Contact: Keith Moore <moore@network-heretics.com>
Keyword value: auto-notified
Description: Indicates that a message was generated by a Sieve
notification system.
Parameters: owner-email, owner-token. Both optional, both refer to the
owner of the Sieve script that generated this message. See the relevant
RFC for details.
Standards Track/IESG-approved experimental RFC number: this RFC
Contact: Michael Haardt <michael.haardt@freenet.ag>

---

7. References

7.1. Normative References

[IANA]	Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," BCP 26, RFC 5226, May 2008.
[Kwds]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119, March 1997.
[Notify]	Melnikov, A., Ed., Leiba, B., Ed., Segmuller, W., and T. Martin, "Sieve Extension: Notifications," work in progress, draft-ietf-sieve-notify, December 2007.
[RFC3834]	Moore, K., "Recommendations for Automatic Responses to Electronic Mail," RFC 3834, August 2004.
[RFC5322]	

	Resnick, P., Ed., "Internet Message Format," RFC 5322, October 2008.
[Sieve]	Guenther, P., Ed. and T. Showalter, Ed., "Sieve: An Email Filtering Language," RFC 5228, January 2008.
[mailto]	Hoffman, P., Masinter, L., and J. Zawinski, "The mailto URL scheme," RFC 2368, July 1998.

7.2. Non-Normative References

| [RFC5321] | Klensin, J., Ed., "Simple Mail Transfer Protocol," RFC 5321, October 2008. |
| [Variables] | Homme, K., "Sieve Extension: Variables," RFC 5229, January 2008. |

Authors' Addresses

	Barry Leiba
	IBM T.J. Watson Research Center
	19 Skyline Drive
	Hawthorne, NY 10532
	US
Phone:	+1 914 784 7941
Email:	leiba@watson.ibm.com
	Michael Haardt
	freenet.de GmbH
	Willstaetter Str. 13
	Duesseldorf, NRW 40549
	Germany
Phone:	+49 241 53087 520
Email:	michael.haardt@freenet.ag

Full Copyright Statement

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.