

Network Working Group  
Internet-Draft  
Updates: [5321](#) (if approved)  
Intended status: Standards Track  
Expires: June 08, 2014

F. Martin, Ed.  
LinkedIn  
A. Peterson, Ed.  
Message Systems  
December 05, 2013

SMTP target host selection in Mixed IPv4/IPv6 environments  
draft-martin-smtp-target-host-selection-ipv4-ipv6-01

## Abstract

The Simple Mail Transfer Protocol (SMTP) is defined in [[RFC5321](#)]. [Section 5](#) of that document describes the process of host selection. While locating the target host in IPv4 is well defined, this process is unclear for systems in IPv4 and IPv6 environments. This specification extends [[RFC5321](#)] to provide a standard way of selecting the target host in mixed environments.

## 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 June 08, 2014.

## Copyright Notice

Copyright (c) 2013 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

Internet-Draft

SMTP host selection with IPv4/v6

December 2013

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">1.1.</a>	Requirements Language . . . . .	<a href="#">2</a>
<a href="#">1.2.</a>	Rationale . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Locating the target host . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Acknowledgements . . . . .	<a href="#">4</a>
<a href="#">4.</a>	IANA Considerations . . . . .	<a href="#">4</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">4</a>
<a href="#">6.</a>	References . . . . .	<a href="#">4</a>
<a href="#">6.1.</a>	Normative References . . . . .	<a href="#">4</a>
<a href="#">6.2.</a>	Informative References . . . . .	<a href="#">5</a>
<a href="#">Appendix A.</a>	Examples . . . . .	<a href="#">5</a>
<a href="#">A.1.</a>	Target list . . . . .	<a href="#">5</a>
<a href="#">A.1.1.</a>	Example 1: dual stacked servers . . . . .	<a href="#">5</a>
<a href="#">A.1.2.</a>	Example 2: single stacked servers . . . . .	<a href="#">5</a>
<a href="#">A.1.3.</a>	Example 3: single and dual stacked servers . . . . .	<a href="#">6</a>
<a href="#">A.1.4.</a>	Example 4: SMTP-Sender limiting the number of IP tried per MX . . . . .	<a href="#">6</a>
Authors' Addresses	. . . . .	<a href="#">8</a>

## [1.](#) Introduction

Target host selection is discussed in [\[RFC5321\] section 5](#) and some operational experience is described in [\[RFC3974\]](#). However the process is unclear when selecting the target host in a mixed environment. Several widely deployed open source MTA do not make this selection the same and sometimes do not find all the hosts. This specification addresses the issue of host selection in mixed environments.

### [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](#) [[RFC2119](#)].

### [1.2.](#) Rationale

[RFC5321] [section 5.2](#) specifically discusses SMTP implementations

taking into account 'local circumstances' when deciding which network addresses to use. However, neither it nor [\[RFC3974\]](#) explore what 'local circumstances' should be evaluated or how they should be evaluated. The result has thus far been an inconsistent set of SMTP client implementations. Providing a guideline for how SMTP client

implementations should handle various dual-stack situations will improve the stability of the network while we operate in a dual-stack mode (which will likely be for an extended period of time).

## [2.](#) Locating the target host

This part applies if one or more MX Resource Records (RR) [\[RFC1035\]](#) are found for a given name and if the sender-SMTP server is IPV4 and/or IPv6 capable.

When a domain name associated with an MX RR is looked up and the associated data field obtained, the data field of that response MUST contain a domain name. That domain name is queried for A RR if the sender-SMTP is IPv4 capable and is queried for AAAA RR if the sender-SMTP is IPv6 capable. A list of IPv4 and IPv6 addresses of the SMTP servers to which the message should be directed is therefore built.

When there are more than one MX RR associated to the given domain, the above process is done for each MX RR. An ordered list of IPv4 and IPv6 addresses is then established using the MX RR preference. Lower numbers are more preferred than higher ones.

For the same MX RR preference, when there is more than one IP address (IPv4/IPv6) to be tried, the sender-SMTP SHOULD sort this list either following the application configuration, recognition of an easily reached or well performing address, or the operating system's address selection [\[RFC6724\]](#) or when there is no clear reason to favor one, then the sender-SMTP MUST randomize them to spread the load across multiple mail exchangers for a specific domain name.

Most sender-SMTP limit the number of IP addresses that will be tried for each MX record. For the same MX RR preference, the sender-SMTP SHOULD ensure that at least two IP addresses of the other address family are tried before that limit.

If the sender-SMTP orders the IP list for the same MX RR preference

based on some reachability criteria, it will not prefer IP addresses for instance it cannot connect to, or IP addresses with more rejection or latency. In some ways, this strategy is similar to the Happy Evedalls [[RFC6555](#)] strategy where web browsers initiate parallel connections to web servers using IPv4 and IPv6 and select the better connection. It is however not similar, because the sender-SMTP MUST NOT try to establish a connection to all the MX RR in parallel but remembers for a period of time which connection, for IP addresses with same MX RR preference, had better response.

Network paths and servers are constantly breaking and being repaired. Thus, implementations MUST occasionally age out their IPv6 or IPv4

preferences. It is RECOMMENDED that this information be aged out several times before the maximum retry time of a message while also following the DNS TTL [[RFC1035](#)] of the MX RR.

### [3.](#) Acknowledgements

Thanks to Murray Kucheraway for guidance in getting this draft out.

### [4.](#) IANA Considerations

This section has no action requested of IANA.

[RFC Editor: Please remove this section before publication.]

### [5.](#) Security Considerations

Rogue sender-SMTP do not follow MX RR preferences, this existed over IPv4 and common defense can be used over IPv6

### [6.](#) References

#### [6.1.](#) Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.

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

- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), July 2003.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", [RFC 5321](#), October 2008.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", [RFC 6555](#), April 2012.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", [RFC 6724](#), September 2012.

## [6.2.](#) Informative References

- [RFC3974] Nakamura, M. and J. Hagino, "SMTP Operational Experience in Mixed IPv4/v6 Environments", [RFC 3974](#), January 2005.

## [Appendix A.](#) Examples

### [A.1.](#) Target list

#### [A.1.1.](#) Example 1: dual stacked servers

For example, a site with dual stacked servers may have the following DNS definitions:

```
example.org.      IN MX   1  mx1.example.org.
                  IN MX  10 mx10.example.org.
mx1.example.org.  IN AAAA 2001:db8:ffff::1
                  IN A   192.0.2.1
mx10.example.org. IN AAAA 2001:db8:ffff::2
                  IN A   192.0.2.2
```

The ordered preference list will look initially:

```
1 2001:db8:ffff::1
1 192.0.2.1
10 2001:db8:ffff::2
10 192.0.2.2
```

#### [A.1.2.](#) Example 2: single stacked servers

For example, a site with single stacked servers may have the following DNS definitions:

```
example.org.          IN MX  1  mx1-6.example.org.
                     IN MX  1  mx1.example.org
                     IN MX  10 mx10-6.example.org.
                     IN MX  10 mx10.example.org
mx1-6.example.org.    IN AAAA 2001:db8:ffff::1
mx1.example.org.     IN A    192.0.2.1
mx10-6.example.org.  IN AAAA 2001:db8:ffff::2
mx10.example.org     IN A    192.0.2.2
```

The ordered preference list will look initially:

```
1 2001:db8:ffff::1
1 192.0.2.1
10 2001:db8:ffff::2
10 192.0.2.2
```

#### [A.1.3.](#) Example 3: single and dual stacked servers

For example, a site with a mix of single and dual stacked servers may have the following DNS definitions:

```
example.org.          IN MX  1  mx1.example.org.
                     IN MX  1  mx1-6.example.org
                     IN MX  1  mx2.example.org
                     IN MX  10 mx10.example.org.
```

```
mx1.example.org.      IN A      192.0.2.1
mx1-6.example.org.   IN AAAA   2001:db8:ffff::1
mx2.example.org.     IN A      192.0.2.2
mx10.example.org.    IN AAAA   2001:db8:ffff::2
mx10.example.org     IN A      192.0.2.3
```

The ordered preference list will look initially:

```
1 192.0.2.1
1 2001:db8:ffff::1
1 192.0.2.2
10 2001:db8:ffff::2
10 192.0.2.3
```

#### [A.1.4.](#) Example 4: SMTP-Sender limiting the number of IP tried per MX

For example, a SMTP-server that prefers IPv6 over IPv4 as per [\[RFC6724\]](#) default address preference is trying to send a message to example.org with the following configuration:

```
example.org.      IN      MX      10  mail1.example.org.
example.org.      IN      MX      20  mail2.example.org.
mail1.example.org. IN      A        192.0.2.1
mail1.example.org. IN      A        192.0.2.2
mail1.example.org. IN      A        192.0.2.3
mail1.example.org. IN      A        192.0.2.4
mail1.example.org. IN      A        192.0.2.5
mail1.example.org. IN      A        192.0.2.6
```

```

mail1.example.org. IN      AAAA  2001:db8::1
mail1.example.org. IN      AAAA  2001:db8::2
mail1.example.org. IN      AAAA  2001:db8::3
mail1.example.org. IN      AAAA  2001:db8::4
mail1.example.org. IN      AAAA  2001:db8::5
mail1.example.org. IN      AAAA  2001:db8::6
mail2.example.org. IN      A      192.0.2.100
mail2.example.org. IN      AAAA  2001:db8::100

```

The sender-SMTP would prefer mail1 (preference 10) and then mail2 (preference 20). An SMTP-server implementation that does not implement this specification and has a 6 address limit per MX record and [RFC6724] default address preference, would attempt to connect to the 6 IPv6 addresses belonging to mail1. If the IPv6 path is broken, these connection attempts would all fail (after a timeout), as would the connection attempt to mail2's IPv6 address. Finally, after those connection timeouts, the IPv4 address for mail2 would succeed. In contrast, an SMTP-server which is following the normative procedure above would succeed in connecting to mail1 if the IPv4 path is working but the IPv6 path is broken. This SMTP-server implementation would attempt to connect to the four IPv6 addresses belonging to mail1 and then to the IPv4 addresses belonging to mail2 -- resulting in mail delivery to mail1 (which is preferable for the receiving domain) and faster mail delivery. Following the procedure above with a 6 address limit per MX record and [RFC6724] default address preference, the connections attempted would be:

```

2001:db8::1    ; mail1, connection attempt 1
2001:db8::2    ; mail1, connection attempt 2
2001:db8::3    ; mail1, connection attempt 3
2001:db8::4    ; mail1, connection attempt 4
192.0.2.1      ; mail1, connection attempt 5 (IPv4)
192.0.2.2      ; mail1, connection attempt 6 (IPv4)
; the other addresses for mail1 are not used by this particular
; sender-SMTP because of its 6 address limit per MX record.
2001:db8::100 ; mail2, connection attempt 1 (IPv6)
192.0.2.100   ; mail2, connection attempt 2 (IPv4)

```



Franck Martin (editor)  
LinkedIn  
Mountain View, CA  
US

Email: [fmartin@linkedin.com](mailto:fmartin@linkedin.com)

Alec Peterson (editor)  
Message Systems  
Columbia, MD  
US

Email: [alec@messagesystems.com](mailto:alec@messagesystems.com)