

Heuristics and Algorithms to Prioritize Protocol deployment B. Trammell
Internet-Draft Google Switzerland GmbH
Intended status: Informational 7 November 2025
Expires: 11 May 2026

Slow Alternate Detection for Happy Eyeballs
draft-trammell-happy-sad-00

Abstract

This document specifies Slow Alternate Detection (SAD) for Happy Eyeballs, an ICMP-based advisory path signal [RFC8558] for exposing information about path non-selection on-path devices in order to aid debugging and measurement of Happy Eyeballs.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at
<https://datatracker.ietf.org/doc/draft-trammell-happy-sad/>.

Discussion of this document takes place on the Heuristics and Algorithms to Prioritize Protocol deployment Working Group mailing list (<mailto:happy@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/happy/>. Subscribe at <https://www.ietf.org/mailman/listinfo/happy/>.

Source for this draft and an issue tracker can be found at
<https://github.com/britram/happy-sad>.

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 <https://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 11 May 2026.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions and Definitions	3
3. Message Format (ICMP)	3
3.1. Code	3
3.1.1. Not Selected	4
3.2. Hash Algorithm (HAlg)	4
3.3. Approximate Sample Rate	4
3.4. 5-tuple fields	5
3.5. Additional Data	5
4. Protocol Behaviors	5
4.1. Not Selected	5
5. Security Considerations	5
6. IANA Considerations	6
7. References	6
7.1. Normative References	6
7.2. Informative References	6
Acknowledgments	7
Author's Address	7

1. Introduction

Happy Eyeballs [I-D.ietf-happy-happyeyeballs-v3] encourages new protocol deployment by reducing the availability risk associated with attempting to use them. However, in doing so, it masks configuration and deployment failures of these very protocols. There are potential causes of such failures, with potential root causes at the end user terminal, CPE, access network, CDN, DNS configuration, and end server. Given the diffusion of root causes, debugging these errors can be difficult.

This document presents Slow Alternate Detection (SAD), an ICMP-based advisory path signal [RFC8558] to devices along the path of a non-selected candidate designed as part of an array of approaches to this problem. It is intended to be used together with complementary monitoring and logging information at each point along the potential failure chain of a non-selected candidate.

This design uses hashes of data relevant to a failure to allow the correlation of nonselection events to on path devices and actors who already have access to that data, while allowing only aggregate analysis by other on-path actors.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Message Format (ICMP)

The message format for SAD is identical for both ICMPv4 and ICMPv6, and is depicted in Figure 1

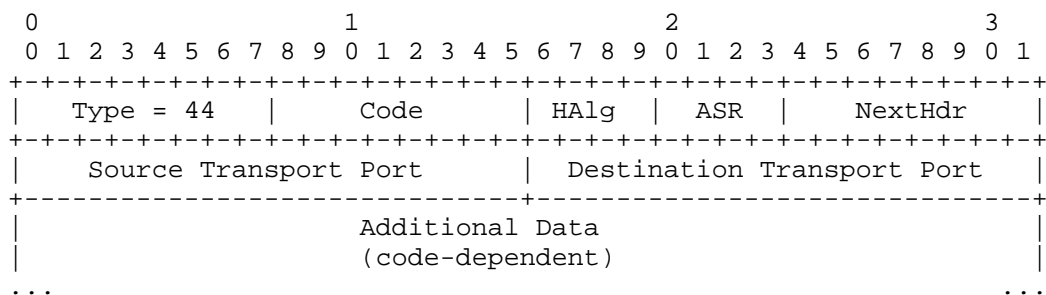


Figure 1: SAD Message Format

The Code, HAlg (Hash Algorithm), ASR (Approximate Sample Rate), Next Header and Source and Destination Transport Port, and Additional Data fields are described in the subsections below.

3.1. Code

The ICMP Code for a SAD message takes one of the following values, and specifies the message's semantics as well as the meaning of the Additional Data field. It can take the following values:

Code	Description	Additional Data
0	Not Selected	hashed DNS Answer
1-255	Reserved	not present

Table 1

3.1.1. Not Selected

Not Selected indicates that is sent to a non-selected candidate by the client, after it has made the decision not to use that candidate. See Section 4.1 for the use of this message.

When present, the Additional Data field of a Not Selected contains the DNS Message [RFC1035] associated with the answer that

3.2. Hash Algorithm (HAlg)

The Hash Algorithm field determines both the length of the Additional Data field and the hash algorithm used to hash it. The following hash algorithms are available:

Value	Length	Algorithm
0	0	Additional Data Omitted
1	256	SHA256 [RFC6234]
2-15	undef	Reserved for Future Use

Table 2

3.3. Approximate Sample Rate

TODO: allows the sender to expose a sample rate, if applied. Design an encoding for common useful sample rates.

3.4. 5-tuple fields

The Next Header (or Protocol) field contains the IP protocol (e.g. TCP, UDP) of the non-selected path. The Source and Destination Transport Port fields contain the source and destination port of the non-selected path. While this will not allow NAT transparency of the SAD message, it does allow analysis and correlation across the NAT by the operator thereof. If the candidate transport protocol does not have port numbers, or if the client chooses not to expose them, these fields are set to 0.

3.5. Additional Data

The Additional Data contains code-specific additional data. If present, it MUST be hashed with a hash algorithm specified by the Hash Algorithm field.

4. Protocol Behaviors

Each of the (currently one) message code(s) associated with the SAD message has an associated protocol behavior, defined below.

4.1. Not Selected

A client sends a Not Selected message after it has decided not to use a given candidate identified by the 5-tuple in the Not Selected message. There is no guarantee of the relative timing of the SAD Not Selected message and the transport layer shutdown datagrams associated with this nonselection.

The client may send Not Selected messages on only a sampled portion of its non-selected candidates. In this case, the client SHOULD expose the selected sample rate in the Approximate Sample Rate field.

As this is an advisory path signal, forwarding elements and servers MUST NOT take any action on the receipt of a Not Selected message beyond logging them for later analysis.

If the client knows that its DNS Answer message was retrieved from an off-path resolver (e.g., via DoH [RFC8484]), it MUST NOT include the hashed DNS Answer in the Not Selected message.

5. Security Considerations

TODO analyze various attacks against this approach.

6. IANA Considerations

This document has two actions for IANA:

- * It requests the allocation of Type 44 in the ICMP Type Numbers registry for Slow Alternative Detection, with this document as reference on publication as an RFC.
- * It requests the establishment the Slow Alternative Detection ICMP Code Field subregistry initialized with the contents of Section 3.1.

7. References

7.1. Normative References

- [I-D.ietf-happy-happyeyeballs-v3]
Pauly, T., Schinazi, D., Jaju, N., and K. Ishibashi,
"Happy Eyeballs Version 3: Better Connectivity Using
Concurrency", Work in Progress, Internet-Draft, draft-
ietf-happy-happyeyeballs-v3-02, 20 October 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-happy-happyeyeballs-v3-02>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and
specification", STD 13, RFC 1035, DOI 10.17487/RFC1035,
November 1987, <<https://www.rfc-editor.org/rfc/rfc1035>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119,
DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms
(SHA and SHA-based HMAC and HKDF)", RFC 6234,
DOI 10.17487/RFC6234, May 2011,
<<https://www.rfc-editor.org/rfc/rfc6234>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

7.2. Informative References

- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS
(DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018,
<<https://www.rfc-editor.org/rfc/rfc8484>>.

[RFC8558] Hardie, T., Ed., "Transport Protocol Path Signals",
RFC 8558, DOI 10.17487/RFC8558, April 2019,
<<https://www.rfc-editor.org/rfc/rfc8558>>.

Acknowledgments

Thanks to the participants in the discussions on error reporting at the HAPPY WG meetings at IETF 123 in Madrid, IETF 124 in Montreal, and on the mailing list in between, to which this document is an answer. Special thanks to Martin Duke for backronyming the name of this extension.

Author's Address

Brian Trammell
Google Switzerland GmbH
Email: ietf@trammell.ch