

Lazy Eye Inspection: Capturing the State of Happy Eyeballs Implementations

Patrick Sattler
sattler@net.in.tum.de
Technical University of Munich
Munich, Germany

Matthias Kirstein
kirstein@net.in.tum.de
Technical University of Munich
Munich, Germany

Lars Wüstrich
wuestrich@net.in.tum.de
Technical University of Munich
Munich, Germany

Johannes Zirngibl
jzirngibl@mpi-inf.mpg.de
Max Planck Institute for Informatics
Saarbrücken, Germany

Georg Carle
carle@net.in.tum.de
Technical University of Munich
Munich, Germany

Abstract

While transitioning to an IPv6-only communication, many devices settled on a dual-stack setup. IPv4 and IPv6 are available to these hosts for new connections. Happy Eyeballs (HE) describes a mechanism to prefer IPv6 for such hosts while ensuring a fast fallback to IPv4 when IPv6 fails. The IETF is currently working on the third version of HE. While the standards include recommendations for HE parameter choices, it is up to the client and OS to implement HE. In this paper, we investigate the state of HE in various clients, particularly web browsers and recursive resolvers. We introduce a framework to analyze and measure clients' HE implementations and parameter choices. According to our evaluation, only Safari supports all HE features. Safari is also the only client implementation in our study that uses a dynamic IPv4 connection attempt delay, a resolution delay, and interlaces addresses. We further show that problems with the DNS A record lookup can even delay and interrupt the network connectivity despite a fully functional IPv6 setup with Chrome and Firefox. We operate a publicly available website (www.happy-eyeballs.net) which measures the browser's HE behavior, and we publish our testbed measurement framework.

CCS Concepts

• Networks → Network protocols.

Keywords

Happy Eyeballs; IPv6; Browser Measurement Tool

ACM Reference Format:

Patrick Sattler, Matthias Kirstein, Lars Wüstrich, Johannes Zirngibl, and Georg Carle. 2025. Lazy Eye Inspection: Capturing the State of Happy Eyeballs Implementations. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25)*, October 28–31, 2025, Madison, WI, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3730567.3732925>

1 Introduction

The standardization of IPv6 [14] as a non-compatible protocol to IPv4 resulted in a separation of the IP layer. Due to the IPv4 address space exhaustion, IPv6 needs to be deployed and used more widely. While the general development is towards dual-stack devices, the non-compatibility imposes a decision on clients: Which address family should be used? However, the Internet has yet to fully transform [2, 13, 41], thus, the decision is nontrivial. To improve IPv6 usage without drastically influencing user experience, Happy Eyeballs (HE) was designed [44]. Across two versions [34, 44], HE provides suggestions for different parts of connections on how to prefer IPv6 but fall back on IPv4 if necessary.

Since IPv6, new transport layer protocols have been standardized, resulting in further separations on additional ISO/OSI stack layers. QUIC [20] resulted in another transport protocol besides TCP and the foundation for a new version of HTTP. Like IPv6, QUIC and HTTP/3 are incompatible with existing alternatives, and clients must decide which protocol to use. Therefore, as of November 2024, a new version of HE [28] and a new IETF working group [26] are proposed to provide guidance given the new complexity.

However, the existing HE ecosystem is not yet understood. All major browsers have implemented HE [15] in theory. Nevertheless, both RFCs only provide suggestions, and how browsers and other clients, e.g., resolvers, implement HE has yet to be documented. With the increasing complexity of the Internet and the attempt to standardize a new, more complex version of HE, it is essential to understand the current RFC implementations. Clients' decisions highly impact the user experience, the visibility of network issues, and potential deployment errors.

Our main contributions in this paper are:

(i) An evaluation of HE implementations in nine browsers on seven Operating Systems (OSes). We show that current implementations vary widely and are mostly limited to the first version of HE (RFC6555 [44]).

(ii) A measurement of resolver IP selection and fallback behavior. We observe that this critical network component generally does not rely on HE style approaches in a dual-stack setup. Our results show a wide variety of behavior with almost no operators using a strict IPv6 version preference.

(iii) A local test framework that tests the HE implementations of multiple clients across different versions.



This work is licensed under a Creative Commons Attribution 4.0 International License. *IMC '25, Madison, WI, USA*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1860-1/2025/10
<https://doi.org/10.1145/3730567.3732925>

(iv) A web-based testing tool that emulates network delays for all protocols HE considers. It is helpful to test browsers and resolvers on different OSes. Moreover, anyone can use our public tool to test the browser of their choice.

The website to access the tool, the source code for the test framework, and the web-based tool are available at:

<https://www.happy-eyeballs.net/>

2 Happy Eyeballs

Happy Eyeballs (HE) aims to enhance the user experience by allowing a fast fallback to IPv4 in cases of impaired IPv6 connectivity. The IETF has standardized two HE versions [34, 44], and version 3 [28] is in draft. This section introduces their algorithms (see Figure 1) and explains their configuration parameters (see Table 1).

HEv1 (2012) [44], focuses on basic connection attempts from a client to a server (gray area in Figure 1). The idea is to race the connection establishment between IPv6 and IPv4 addresses while prioritizing one protocol family based on the host’s address selection policy. If DNS returns multiple records for each address family, the host orders them according to its address selection policy. HEv1 recommends falling back to IPv4 after the first IPv6 attempt, even if multiple IPv6 addresses are available. If the initial IPv6 and IPv4 connections fail, HEv1 does not provide further instructions on the order of subsequent connection attempts, leaving this decision up to the application developers.

The core parameter of HEv1 is the *Connection Attempt Delay* (CAD), which is recommended to be within 150-250 ms (see Table 1). If an application’s IPv6 connection is not successfully established within this delay, the application should initiate a parallel IPv4 connection. HE largely avoids harming the network or servers, as the second connection is only initiated if IPv6 appears slow or broken. Once one connection attempt succeeds, the client discards the others and should cache the outcome for “the order of 10 min” [44].




Figure 1: Happy Eyeballs (v2): A client first resolves the server’s name. Then it prioritizes the addresses based on a local policy and proceeds to attempt to establish a connection to the server. HEv1 only focuses on the connection establishment, HEv2 additionally includes the name resolution.

HEv2 (2017) [34] additionally considers DNS (the complete process in Figure 1) and extends the description of the *address selection*. In HEv2, a client should first issue a AAAA query, immediately followed by an A query. If the client receives a response with AAAA records first, it immediately starts a connection attempt to the target host using IPv6. If an A response arrives first, the application should wait for a *Resolution Delay* (RD) of 50 ms that starts with the reception of the A record (see Table 1). If the client receives a AAAA response during the RD, a connection attempt via IPv6 is immediately initiated. The client starts a connection attempt via IPv4 after the RD if there is no AAAA answer. This behavior prioritizes IPv6 while minimizing the user-visible delay if it is impaired.

The next step in HEv2 is the address selection of the retrieved addresses by using a predefined policy. In addition to IP address family, clients can leverage knowledge about historical TCP round-trip times and previously used addresses. Finally, the addresses in the list are interlaced to alternate between IP address families. The first address at the top of the list should be an IPv6 address to maintain the IPv6 preference. The list may start with multiple IPv6 addresses. The number of those addresses is called the *First Address Family Count*, which is recommended to be 1 or 2 for aggressively favoring one family (see Table 1). Like in HEv1, the client then proceeds to establish a connection using the previously selected addresses. It uses the CAD between subsequent attempts to prevent network congestion. HEv2 recommends 250 ms for a fixed CAD. For history-inferred implementations, it defines an absolute minimum CAD of 10 ms, recommends a minimum of 100 ms, and a maximum CAD of 2 s (see Table 1).

HEv3 As of 2025, the IETF is working on HEv3 [28]. HEv3 will process existing SVCB or HTTPS DNS resource records [35, 47] to support protocol discovery. The HEv3 address selection should favor IP addresses with available TLS Encrypted ClientHello (ECH) over QUIC over TCP. Besides those changes, HEv3 is currently similar to HEv2, using the same parameters and recommendations.

3 Related Work

To the best of our knowledge, the actual implementation of HE in clients has not been researched, except for initial articles during the standardization of HEv1, e.g., by Aben [1] or Huston [15]. The effects of HE have often been analyzed based on the performance of IPv6 and IPv4 deployments on the Internet and the suggestions from the HE RFCs [34, 44] (see Table 1). Bajpai and Schönwälder [4, 5] evaluated the latency towards services based on the protocol used and inferred which protocol would be used by clients based on the delays from RFCs. Pujol et al. [29] evaluated client and service IPv6 capabilities. Based on network metrics, they argued that Happy Eyeballs is visible, but most clients should rely on IPv6 based on suggested delays. However, our work shows that the suggestions from RFCs are rarely used, and HE is only partially implemented. Different studies [2, 13, 41] have shown a trend towards an IPv6-only Internet, but essential components still lack full IPv6 support. Therefore, HEv1 and v2, focusing on IPv6, are still important mechanisms for the quality of experience on the Internet.

While IPv6 deployment started slowly, Zirngibl et al. [45, 46] have shown early and widespread QUIC deployment. Dong et al. [9] have shown that HTTPS DNS resource records, including the first

Table 1: Comparison of parameters defined for HEv1 [44], v2 [34] and the draft for v3 [28] as of April 2025.

Parameter	HEv1 (2012)	HEv2 (2017)	HEv3 (2025-ongoing)	Local tests	Web tests
Considered protocols	IPv4, IPv6	IPv4, IPv6, DNS	IPv4, IPv6, DNS, QUIC	✓	✓
DNS Records	-	AAAA, A	SVCB, HTTPS, AAAA, A	✓	✓
Resolution Delay	-	50 ms	50 ms	✓	✓
Address selection	IPv6 once, then IPv4	alternating	IP family and L4 protocol	✓	-
Fixed Conn. Attempt Delay	150-250 ms	250 ms	250 ms	✓	✓
→Min/Rec./Max when dynamic	-	10 ms / 100 ms / 2 s	10 ms / 100 ms / 2 s	✓	✓

ECH information, are available. Therefore, an early standardization of HEv3 considering these possibilities is reasonable. However, Dong et al. [9] show that most browsers do not fully use information from HTTPS resource records as of 2024. Our work shows that most browsers currently do not implement HEv2 yet. Foremski et al. [11] analyze the effect of small negative caching values on clients using HE and find domains with up to 90 % empty AAAA responses due to HE. Similar to HEv3 drafts, Papastergiou et al. [24] suggested racing different transport protocols alongside the IP layer. They implement an example to show potential advantages, but do not consider current HE implementations and whether they can be combined. Our findings should be considered during the standardization and evaluation of the Internet and the potential effects of HE in the future. Especially if the overall algorithm gets more complex, clients might drift farther from suggestions.

4 Measurement Methodology

We set up two test environments to get insight into how different clients use HE: a local, fully controlled one and a web-based tool. Our approach treats the software under test as a black box. Therefore, it is usable with any software that can target arbitrary targets (e.g., command line tools).

4.1 Measurement Targets

First, we define the measurement targets and our approach to evaluate these. HE defines *Configurable Values*. We evaluate HE clients along these values. Table 1 shows these variables and the environment in which we measured them. The following describes how we target the CAD, the RD, the address selection approach, and the resolver behavior. The observation method depends on the measurement setup and is described in Section 4.3.

(i) *Connection Attempt Delay*: The CAD is the time the client waits for an IPv6 response before initiating a parallel connection attempt via IPv4. To obtain the CAD, we configure a domain name to resolve to an IPv4 and IPv6 address. We simulate a slower IPv6 connection by delaying IPv6 packets using `tc-netem` [21] on the server side. If the introduced delay exceeds the client’s CAD, it should initiate an IPv4 connection. We can observe these attempts and infer the CAD.

(ii) *Resolution Delay*: The RD is the delay the client waits for an AAAA (IPv6 address record) query response after it obtained the A (IPv4 address record) query response. We use a custom authoritative name server implementation to delay responses according to encoded test parameters. These parameters include the delay,

the resource record type to delay, and a nonce to prevent caching effects. This approach allows a variety of different test setups with a single server deployment.

(iii) *Address Selection*: Our DNS-based measurements show whether both record types are queried and if the client has any HE CAD configured. We extend this test case to include multiple addresses per IP family to evaluate the destination address selection by clients. We use addresses that do not respond at all, and thus should trigger the fallback mechanism to other addresses.

4.2 Recursive Resolver:

The HE standardization primarily targets end-user client software, and as part of these mainly browser maintainers are involved. While the specifics of the HE standardization (e.g., recommended delay values) are not generalizable, we can still apply our measurement setup to other protocols and applications to learn about their IP version selection and fallback behavior. The results can then be interpreted in the context of the analyzed use case. Our measurement approach can be expanded to include transport protocols in the future. The DELEG working group [16] discusses the indication of these protocols to resolvers.

In this work, we will look at the behavior of DNS resolver software. Recursive resolvers play a central role in Internet communication. They must also decide which IP version to use during the iterative resolution process. At each delegation to a new zone, and thus usually a new authoritative name server, the recursive resolver decides if and how to prefer IPv6 over IPv4 to connect to the zone’s name server.

Therefore, we extended our measurement setup and applied it to the resolver software to check their IP preference behavior in general and if they use HE style approaches. We need to apply our traffic shaping to the authoritative name servers serving our zones to enable such measurements. Instead of different domain names inside a single zone, we created entirely different zones for each measured delay. Our traffic shaping is applied to the name server records (in case of RD measurements) and the corresponding IP addresses (for CAD measurements). Recursive resolvers often include performance metrics when deciding which name server and IP address to use. Our authoritative IP addresses are not used by any zone outside our measurement campaign. Additionally, we use unique zone apexes and unique authoritative name server names to reduce the impact of such caching effects.

4.3 Measurement Setups

We implement these test cases in two different setups:

(i) *Local Testbed Framework*: We set up a local testbed of two directly connected hosts (client and server) to eliminate any network effects in our measurements. This setup is an ideal scenario where IPv4 and IPv6 connections between two dual-stacked hosts experience equal delays. We use a separate interface to set up and send instructions to the hosts.

A controller process runs our framework process to initiate measurement runs. Our framework is extendable, as the test cases and clients to test are defined outside the framework’s code (see App. Figure 3). They consist of a config file instructing the framework which executables to run when. For example, to set up our CAD test, we start an NGINX instance [22] and an authoritative DNS server and configure delays using `tc-netem` on the server host. On the client host, we start a packet capture, the client application, and initiate the request. Setup stages can be executed at each test run configuration (e.g., delay), or only at the start and end of a test case. The test configuration also supports defining the ranges and intervals for the test run configuration variable to execute coarse initial runs and fine-grained follow-ups. The clients are defined separately from the test cases and start via predefined script calls, which each client must implement independently. Our setup uses containers to start with a clean client state and use Selenium [36] to control browsers. After each test case configuration, we reset the client to a predefined state (e.g., drop and create a new container) to prevent any caching effects.

We determine the CAD by measuring the time between the first IPv6 packet and the first IPv4 packet observed in the client’s packet capture. The local testbed’s accuracy depends on the packet capture time stamping accuracy, which is usually below a millisecond [12]. Therefore, this method gives us the needed accuracy for the CAD value as real-world connections usually expose larger jitter than our measurement accuracy.

(ii) *Web-based Testing Tool*: We created a web-based tool to enable tests from a broader set of browsers in combination with different operating systems. This tool also allows testing on devices that cannot be directly attached to a server and remotely instructed to perform specific tasks. The results can be used to validate our local testbed measurements regarding real-world network conditions.

We evaluate results solely from the client side. Our web server returns the client’s source address in its response, and the client uses the response to determine the used IP version. The nature of this web deployment does not allow resetting client and server configurations after each measurement. Therefore, we use a fixed set of 18 delays between 0 and 5 s (see App. Figure 4a). Each delay has dedicated IPv4 and IPv6 addresses assigned, and we delay IPv6 traffic accordingly. Therefore, the CAD can only be determined to be in the interval of the last delay using IPv6 and the first delay using IPv4, e.g., the CAD for Safari in App. Figure 4a is $CAD \in (200, 250]$. Nevertheless, it is valuable to obtain results in actual network conditions. Some clients adjust their behavior in non-lab conditions, significantly influencing observed results (see Section 5.1). Furthermore, we associate a dedicated domain to each delay-address pair to prevent caching. We reuse our custom DNS setup with a publicly resolvable zone.

Table 2: HE feature evaluation of client applications with our local and web-based measurement tools and their consistency between the measurement methods.

	Prefers IPv6	CAD Impl.	AAAA first ¹	RD Impl.	IPv4 Addr. Used	IPv6 Addr. Used	Addr. Selection ²	Consistency
Chrome 130.0	●	●	●	○ ⁴	1	1	○	●
Chromium 130.0	●	●	●	○ ⁴	1	1	○	●
Edge 130.0	●	●	●	○	1	1	○	●
Firefox 132.0	●	●	○	○	1	1	○	●
Safari 17.6	●	●	●	●	10	10	●	○
Mobile Safari	●	●	●	●	-	-	-	-
Chrome Mobile	●	●	●	○	-	-	-	-
curl 7.88.1	●	●	●	○	1	1	○	-
wget ³ 1.21.3	●	○	○	○	-	1	○	-

● Observed as defined; ● Observed with RFC deviation; ○ Not observed

¹ Can also depend on the operating systems stub resolver except for Chromium-based browsers which use their own stub resolver.

² No web test validation. ³ Does not implement any type of HE.

⁴ Offers a flag to enable feature (HEv3 flag)

5 Evaluation

We evaluate our measurement targets based on results obtained via the local testbed and the web-based testing tool.

5.1 Connection Attempt Delay

Our local testbed includes Mozilla Firefox, Google Chrome, Microsoft Edge, and Chromium. All browsers were tested on the latest Ubuntu LTS image (24.04.1). We included browser versions as old as January 2021. Additionally, we measured Safari (17.5) on macOS. This is a set of the most used browsers [37, 43] and, thus, a solid base for our web-based evaluation. Moreover, most alternative browsers are Chromium-based, and they behave the same. In addition to browsers, we included tests for curl [38] and wget [30] as popular command-line tools.

Initial measurements using coarse delay configurations indicate the rough time interval of the configured CAD. Except for Safari, all clients stayed within 300 ms. Our fine-grained delay configurations perform measurements from 0 to 400 ms in 5 ms steps. It is important to note that any local measurement that uses a delay larger than the client’s CAD also observes the CAD value in the packet capture. Our coarse evaluation measures the CAD at least 20 times per client and testbed run.

Additionally, we collected 161 web-based measurement results with at least ten repetitions per delay configuration. Results cover nine browsers in 22 versions run on seven operating systems. That amounts to a total of 33 different combinations (see Table 5).

Table 2 provides an overview of our results. While all client applications prefer IPv6 if both versions are offered, wget (1.21.3) does not implement any fallback mechanism and thus does not implement HE. wget fails without using the provided IPv4 addresses




Figure 2: IP address family of established connection as measured on our local testbed setup. We omitted Safari, which has a CAD of 2 s to improve the visibility of the remaining results.

for delays larger than its timeout. Figure 2 shows the address family of the established connection at each configured delay for all browsers and versions included in our local testbed measurement. The observed CAD obtained from the packet captures overlaps with the largest delay on an established connection via IPv6. Only Firefox has a few outliers, but the median and standard deviation are within a ms of the obtained value.

All Chromium-based browsers use a CAD of 300 ms. This value can also be found in its code [3]. Measurement results collected via our web-based tool for these browsers (including mobile variants) validate these observations. The command-line tool `curl` uses the smallest CAD of 200 ms (see code [39, 40]). Firefox follows the RFC recommendation of 250 ms, but has a few outliers where it waits longer than 250 ms for the IPv6 connection establishment.

In contrast to other browsers that matched our local measurements, we observed a wide range of CAD values for Safari [32] in our web-based measurements. While local measurements were consistent at 2 s, they ranged from 50 ms up to 5 s on the web, indicating a dynamic application of CAD. Moreover, Safari also showed a unique behavior. While other web-based measurements usually show a clear picture of IPv6 usage up until the CAD value, the measurement results with Safari did not. This behavior manifested in repetitions where it used IPv4 at a smaller delay and IPv6 again for larger delays. Firefox and Chrome also had a few test runs with such inconsistencies, but these only manifested in a maximum of two out of ten test repetitions. At the same time, Safari exposed them between six and ten times. To investigate this behavior we conducted web-based measurements with different network conditions. Neither the network context (cable vs wireless, the access network, parallel network activity), nor the focus of the application window (MacOS and Safari can throttle background tabs and windows), nor the power supply did have any noticeable impact on the

reproducibility of results. Mobile Safari measurements report varying CAD values and inconsistency. Nevertheless, the CAD never rose beyond 1 s for IPv6 connections on mobile phones with iOS. We assume a preference towards lower values on mobile devices. Our web-based measurements can help to evaluate these effects in more detail with different clients and conditions.

Moreover, Apple announced and deployed iCloud Private Relay (iCPR), a privacy-preserving proxying service, in December 2021 [17, 31]. As part of iCPR, Safari can relay traffic via an egress operator using MASQUE proxying [27, 33]. Connections via iCPR show completely different HE behavior compared to Safari. The reason is that Safari does not create an IP tunnel via the relay network but only informs the egress node about the server name it wants to connect to. The egress node handles the DNS resolution and the stack up to the transport protocol headers of TCP and UDP. Therefore, measurements via iCPR show how the egress operators implement HE. Akamai and Cloudflare egress nodes use a CAD of 150 ms and 200 ms, respectively. We cannot actively select the egress node provider and Fastly did not appear as a provider in our measurements. This example shows that local testbed measurements can provide an overview of which features are deployed, but a web-based measurement campaign must be used to validate the results to cover complex real-world scenarios.

5.2 Resolution Delay

The RD test case delays the AAAA record to observe how clients implement the RD. Already, the testbed measurement results show that only Safari actually implements it. We found Safari to be using the RFC recommendation of 50 ms. In their default state, Chromium-based browsers and Firefox depend on the resolver’s timeout. They do not apply any DNS resolution timeout on their own.

The missing timeout leads to the delegation of timeouts to resolvers. Only after the AAAA query hits the resolver’s timeout do these browsers initiate the IPv4 connection. Similar observations are valid for `curl` and `wget`.

Usually, in the context of HE, IPv4 is considered the stable fallback, while IPv6 is the *new* variant that can break. In reality, IPv4 is not 100 % reliable, and queries for A records are also sent via IPv6. Therefore, we test how clients react to delayed A query responses. To our astonishment, while all clients continued to prefer IPv6, all but Safari always waited for the A response to arrive. Only then would they initiate the IPv6 connection to the target host, although the AAAA record was not delayed. This behavior resulted in Chrome and Firefox completely failing connections in case of high delays with some resolver configurations. Since April 2024, Chromium provides a feature flag (HEv3) that adds RD and gets rid of this problematic behavior [18, 19].

Similar to the CAD evaluation, we collected web-based measurements with Safari via iCPR. We observed that Cloudflare egress nodes use IPv6 up until a delay of 1.75 s. Akamai egress nodes use a timeout of 400 ms. Both operators use the same timeout for A and AAAA record queries. These properties are far from the behavior of Safari without iCPR. Therefore, iCPR has the potential to negatively influence the user experience if issues on the resolution path exist.

The address selection is another measure standardized as part of HE. We measured it with a domain name with 10 A and AAAA records.

Table 3: Evaluation of resolver IPv6 usage as observed on the authoritative name server. For open resolvers, the observed IPv6 delay is a lower bound.

Service	AAAA Query	IPv6 Share	Max. IPv6 Delay Used	# IPv6 Packets
BIND	●	100.0 %	800 ms	1
Unbound	●	43.8 %	376 ms	2
Knot Resolver	◐	27.9 %	400 ms	2
DNS.sb	●	0.0 %	-	-
Google P. DNS	●	0.0 %	-	-
DNS0.EU	●	9.5 %	- ¹	2
NextDNS	●	8.9 %	200 ms	1
Quad 101	●	10.0 %	400 ms	1
114DNS	●	11.1 %	600 ms	1
Cloudflare	●	11.1 %	500 ms	2
Verisign P. DNS	●	15.3 %	250 ms	1
Yandex	●	17.4 %	300 ms	6
H-MSK-IX	●	20.5 %	600 ms	2
MSK-IX	●	22.1 %	600 ms	2
Quad9 DNS	●	34.2 %	1250 ms	2
OpenDNS	●	100.0 %	50 ms	1

● Sends AAAA before A; ● Sends AAAA after A; ● Sends AAAA after query to IPv4 auth. name server; ◐ Sends either AAAA or A but not both

¹ Cannot determine the delay due to parallel queries on IPv4 and IPv6.

All records pointed to unresponsive addresses. In Table 2, we list the number of addresses per IP family used for connection attempts by the clients. Again, only Safari shows a visible address selection behavior (see App. Figure 5). Others only perform a simple fallback to IPv4 and stop after that, ignoring other records' addresses.

5.3 Recursive Resolvers

Most browsers do not implement RD and rely on the resolver's timeout. To understand how and if resolvers implement some version of IPv6 version preference, we evaluate the behavior towards authoritative name servers.

We focus on three popular recursive resolvers for our local measurements: BIND 9 [6], Unbound [23], and Knot resolver [7]. Furthermore, we issued queries to 82 IP addresses of 17 popular open resolvers (see App. Table 4). While our sample does not cover ISP-operated resolvers, it shows the differences that resolvers can expose. Moreover, we provide a web-based testing tool that allows users to check their configured resolver. Our analysis only includes resolvers that can resolve zones with IPv6-only authoritative name servers. Hurricane Electric, Lumen (Level3), Dyn, and G-Core are not able to resolve domain names with an IPv6-only delegation, leaving us with 13 to analyze. For Quad 101, only the IPv6 addresses successfully connect to our IPv6-only name servers. One dns0.eu IP address did not provide reliable IPv6-only resolution.

To connect to our IPv6 addresses, the resolvers need to obtain the address through glue records or name resolution. All operators sent AAAA queries for the name server names. Even when we include glue records, 12 out of 13 public services and the BIND resolver requested the AAAA record before issuing the query to the authoritative name server. Only Google Public DNS does not issue any AAAA query before sending the query to the authoritative name server. The Knot

resolver does send out either A or AAAA queries for the delegated name server name but never both.

11 open resolvers and Unbound sent the AAAA record before the A record, a behavior defined in RFC 8305 [34]. DNS.sb, Google Public DNS, and the BIND resolver query the A record first. While packet reordering could influence our results, the results were consistent over different test runs and delay configurations.

When the resolver has both addresses available, it needs to select an address for the next iterative query. BIND performs classic HE IP version preference and always prefers IPv6. Knot and Unbound use IPv6 in 25 % and 50 % respectively (see Table 3). We observe a CAD of 800 ms for BIND, 400 ms for Knot, and 376 ms for Unbound. While BIND and Knot are consistently falling back to IPv4 addresses, Unbound retries the IPv6 address in 44 % of the times and increases the CAD to 1128 ms due to an exponential backoff [42].

Except for Google Public DNS and DNS.sb, all public operators use the IPv6 name server address at least once. Only OpenDNS always tries to connect via IPv6 first and falls back to IPv4 after 50 ms. The other operators do not perform any type of HE style behavior. Nine operators do not use IP protocol version interleaving and send up to six queries to the IPv6 address (Yandex). We never see DNS0.EU switching IP protocol version when retrying, but it sticks to the IP version initially chosen and fails at some point.

6 Discussion and Conclusion

We systematically approach the measurement of the client's Happy Eyeballs feature implementations. A local testbed allows us to perform in-depth measurements in a controlled setup. The web-based measurements can validate local results and cover more clients in actual network conditions. All widely used browsers prefer IPv6 connections when possible. Except for Safari, which employs a dynamic, unpredictable approach, the observed CAD values are within a 50 ms interval of the 250 ms RFC suggestion.

Only Safari implements a RD and address selection mechanisms, thus, HEv2. In contrast, Chromium-based browsers and Firefox delay IPv6 connection initialization until the corresponding DNS A record arrives, even though IPv6 address information is available and preferred. We suggest implementing a timeout for DNS queries for all clients, even when HE is not implemented. The current situation is even worse from an IPv6 deployment perspective, as slow A queries also slow down IPv6, even if it is not at fault.

Moreover, not only the browser software itself but also new proxying techniques can impact the HE behavior. This result shows the usefulness of our web-based measurement tool in understanding the ecosystem in different scenarios, how clients with different configurations currently behave, and how future updates (e.g., the Chrome feature flag) change behavior. For example, we determined that if iCPR is enabled, Safari users lose RD and address selection features through our web test.

Due to the reliance on resolvers and their significant role, we included them in our analysis. Only single instances and applications follow an HE-style approach. We suggest starting dedicated discussions to develop recommendations on the behavior of protocol preference for critical Internet infrastructure clients, such as DNS resolvers. You can find our web-based tool and our measurement tools' source code on: <https://www.happy-eyeballs.net/>.

Acknowledgments

We thank the anonymous reviewers and our shepherd Pawel Foremski for their valuable feedback. This work was partially funded by the German Federal Ministry of Education and Research under project PRIMENet (16KIS1370).

References

- [1] Emile Aben. 2011. Hampering Eyeballs - Observations on Two "Happy Eyeballs" Implementations. <https://labs.ripe.net/author/emileaben/hampering-eyeballs-observations-on-two-happy-eyeballs-implementations/>
- [2] APNIC. 2024. *IPv6 Capable Rate by country (%)*. Retrieved 2025-05-01 from <https://stats.labs.apnic.net/ipv6>
- [3] The Chromium Authors. 2024. *Chromium Code Connection Attempt Delay*. Retrieved 2025-05-01 from https://chromium.googlesource.com/chromium/src/+refs/heads/main/net/socket/transport_connect_job.h#117
- [4] Vaibhav Bajpai and Jürgen Schönwälder. 2016. Measuring the Effects of Happy Eyeballs. In *Proceedings of the 2016 Applied Networking Research Workshop*. <https://doi.org/10.1145/2959424.2959429>
- [5] Vaibhav Bajpai and Jürgen Schönwälder. 2019. A Longitudinal View of Dual-Stacked Websites—Failures, Latency and Happy Eyeballs. *IEEE/ACM Transactions on Networking* (2019). <https://doi.org/10.1109/TNET.2019.2895165>
- [6] Internet Systems Compositum. 2024. *BIND 9*. Retrieved 2025-05-01 from <https://www.isc.org/bind/>
- [7] cz.nic. 2024. *Knot Resolver*. Retrieved 2025-05-01 from <https://www.knot-resolver.cz>
- [8] David Dittrich, Erin Kenneally, et al. 2012. The Menlo Report: Ethical principles guiding information and communication technology research. *US Department of Homeland Security* (2012).
- [9] Hongying Dong, Yizhe Zhang, Hyeonmin Lee, Shumon Huque, and Yixin Sun. 2024. Exploring the Ecosystem of DNS HTTPS Resource Records: An End-to-End Perspective. In *Proc. ACM Internet Measurement Conference (IMC)*. <https://doi.org/10.1145/3646547.3688410>
- [10] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *Proc. USENIX Security Symposium*.
- [11] Pawel Foremski, Oliver Gasser, and Giovane C. M. Moura. 2019. DNS Observatory: The Big Picture of the DNS. In *Proc. ACM Internet Measurement Conference (IMC)*. <https://doi.org/10.1145/3355369.3355566>
- [12] Sebastian Gallemlücker, Johannes Naab, Iris Adam, and Georg Carle. 2020. 5G QoS: Impact of Security Functions on Latency. In *EEE/IFIP Network Operations and Management Symposium*. <https://doi.org/10.1109/NOMS47738.2020.9110422>
- [13] Google. 2024. *Statistics: IPv6 Adoption*. Retrieved 2025-05-01 from <https://www.google.com/intl/en/ipv6/statistics.html>
- [14] Bob Hinden and Dr. Steve E. Deering. 1998. Internet Protocol, Version 6 (IPv6) Specification. RFC 2460. <https://doi.org/10.17487/RFC2460>
- [15] Geoff Huston. 2012. Bemused Eyeballs: Tailoring Dual Stack Applications for a CGN Environment. <https://labs.apnic.net/index.php/2012/05/13/bemused-eyeballs-tailoring-dual-stack-applications-for-a-cgn-environment/>
- [16] IETF. [n. d.]. *DNS Delegation (deleg)*. Retrieved 2025-05-01 from <https://datatracker.ietf.org/wg/deleg/documents/>
- [17] Apple Inc. 2021. *iCloud Private Relay Overview*. (2021). https://www.apple.com/privacy/docs/iCloud_Private_Relay_Overview_Dec2021.PDF
- [18] Kenichi Ishibashi. 2024. *Add HostResolver::ServiceEndpointRequest*. Retrieved 2025-05-01 from https://chromium.googlesource.com/chromium/src/+1751e607bf852a71f305b2778e291c7ee246ccf1%5E%21/net/dns/host_resolver_dns_task.cc
- [19] Kenichi Ishibashi. 2024. *Rename UseServiceEndpointRequest feature to EnableHappyEyeballsV3*. Retrieved 2025-05-01 from <https://chromium.googlesource.com/chromium/src/+70457ad4d152e853fe9a5240449580ea021f807b>
- [20] Jana Iyengar and Martin Thomson. 2021. QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000. <https://doi.org/10.17487/RFC9000>
- [21] Fabio Ludovici and Hagen Paul Pfeifer. [n. d.]. *tc-netem - Linux manual page*. <https://man7.org/linux/man-pages/man8/tc-netem.8.html>
- [22] nginx. 2024. *nginx*. Retrieved 2025-05-01 from <https://nginx.org/en/>
- [23] NLNet. 2024. *Browser Statistics*. Retrieved 2025-05-01 from <https://www.nlnetlabs.nl/projects/unbound/about/>
- [24] Giorgos Papastergiou, Karl-Johan Grinnemo, Anna Brunstrom, David Ros, Michael Tüxen, Naeem Khademi, and Per Hurtig. 2016. On the Cost of Using Happy Eyeballs for Transport Protocol Selection. In *Proceedings of the 2016 Applied Networking Research Workshop*. <https://doi.org/10.1145/2959424.2959437>
- [25] Craig Partridge and Mark Allman. 2016. Addressing Ethical Considerations in Network Measurement Papers. *Commun. ACM* 59, 10 (Oct. 2016).
- [26] Tommy Pauly. 2024. *HAPPY WG (Heuristics and Algorithms to Prioritize Protocol deployment)*. Retrieved 2025-05-01 from <https://github.com/tpauly/draft-happy-eyeballs-v3/blob/main/charter.md>
- [27] Tommy Pauly, David Schinazi, Alex Chernyakhovsky, Mirja Kühlewind, and Magnus Westerlund. 2023. Proxying IP in HTTP. RFC 9484. <https://doi.org/10.17487/RFC9484>
- [28] Tommy Pauly, David Schinazi, Nidhi Jaju, and Kenichi Ishibashi. 2025. *Happy Eyeballs Version 3: Better Connectivity Using Concurrency*. Internet-Draft draft-ietf-happy-happyeyeballs-v3-00. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-ietf-happy-happyeyeballs-v3/00/> Work in Progress.
- [29] Enric Pujol, Philipp Richter, and Anja Feldmann. 2017. Understanding the Share of IPv6 Traffic in a Dual-Stack ISP. In *Proc. Passive and Active Measurement (PAM)*.
- [30] Tim Rühse, Darshit Shah, and Giuseppe Scrivano. 2025. *GNU Wget*. Retrieved 2025-05-01 from <https://www.gnu.org/software/wget/>
- [31] Patrick Sattler, Juliane Aulbach, Johannes Zirngibl, and Georg Carle. 2022. Towards a Tectonic Traffic Shift? Investigating Apple's New Relay Network. In *Proc. ACM Internet Measurement Conference (IMC)*. <https://doi.org/10.1145/3517745.3561426>
- [32] David Schinazi. 2015. *[v6ops] Apple and IPv6 - Happy Eyeballs*. Retrieved 2025-05-01 from https://mailarchive.ietf.org/arch/msg/v6ops/DYil9v_O66RNbMJsx0NsatFkubQ/
- [33] David Schinazi. 2022. Proxying UDP in HTTP. RFC 9298. <https://doi.org/10.17487/RFC9298>
- [34] David Schinazi and Tommy Pauly. 2017. Happy Eyeballs Version 2: Better Connectivity Using Concurrency. RFC 8305. <https://doi.org/10.17487/RFC8305>
- [35] Benjamin M. Schwartz, Mike Bishop, and Erik Nygren. 2023. Service Binding and Parameter Specification via the DNS (SVCB and HTTPS Resource Records). RFC 9460. <https://doi.org/10.17487/RFC9460>
- [36] Selenium. 2024. *Selenium*. Retrieved 2025-05-01 from <https://www.selenium.dev>
- [37] Statista. 2024. *Global market share held by leading internet browsers from January 2012 to August 2024*. Retrieved 2025-05-01 from <https://www.statista.com/statistics/268254/market-share-of-internet-browsers-worldwide-since-2009/>
- [38] Daniel Stenberg. 2024. *curl*. Retrieved 2025-05-01 from <https://curl.se>
- [39] Daniel Stenberg. 2024. *curl - man page HE timeout*. Retrieved 2025-05-01 from <https://curl.se/docs/manpage.html#-happy-eyeballs-timeout-ms>
- [40] Daniel Stenberg. 2025. *curl - Code Connection Attempt Delay*. Retrieved 2025-05-01 from <https://github.com/curl/curl/blob/d9ca7ad5cb036c429faca2af81be8c8bc332236/include/curl/curl.h#L1975>
- [41] Florian Streibelt, Patrick Sattler, Franziska Lichtblau, Carlos H. Gañán, Anja Feldmann, Oliver Gasser, and Tobias Fiebig. 2023. How Ready is DNS for an IPv6-Only World?. In *Proc. Passive and Active Measurement (PAM)*.
- [42] Unbound. 2024. *Unbound Timeout Information*. Retrieved 2025-05-01 from <https://www.nlnetlabs.nl/documentation/unbound/info-timeout/>
- [43] W3Schools. 2024. *Browser Statistics*. Retrieved 2025-05-01 from <https://www.w3schools.com/browsers/>
- [44] Dan Wing and Andrew Yourtchenko. 2012. Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555. <https://doi.org/10.17487/RFC6555>
- [45] Johannes Zirngibl, Philippe Buschmann, Patrick Sattler, Benedikt Jaeger, Juliane Aulbach, and Georg Carle. 2021. It's Over 9000: Analyzing Early QUIC Deployments with the Standardization on the Horizon. In *Proc. ACM Internet Measurement Conference (IMC)*. <https://doi.org/10.1145/3487552.3487826>
- [46] Johannes Zirngibl, Florian Gebauer, Patrick Sattler, Markus Sosnowski, and Georg Carle. 2024. QUIC Hunter: Finding QUIC Deployments and Identifying Server Libraries Across the Internet. In *Proc. Passive and Active Measurement (PAM)*.
- [47] Johannes Zirngibl, Patrick Sattler, and Georg Carle. 2023. A First Look at SVCB and HTTPS DNS Resource Records in the Wild. In *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. <https://doi.org/10.1109/EuroSPW59978.2023.00058>

A Ethics

We follow strict ethical measures in our recursive resolver active measurements and our web-based measurement campaign. These measures are mainly based on informed consent [8] and well known best practices [25]. We apply measures described by Durumeric et al. [10] to our active recursive resolver measurements. We limit the rate of our scans and use a blocklist based on requests to be excluded from our scans. The scanning prefix IRR information points to an email contact on our control to react quickly to all requests. Furthermore, scanning machines operate a website pointing to information about our scans. We did not receive any complains about this research.

The web-based measurements in this paper were only collected by research group members and devices owned by them. We collect information on the client and its state during the web clients






Figure 3: Overview of the local testbed and its components and configurations (see Section 4).

Happy Eyeballs Connection Attempt Delay Webtester



(a) Connection Attempt Delay testing tool website.

Happy Eyeballs Resolution Delay Webtester



(b) Resolution Delay testing tool website.

Figure 4: Screenshots of our web-based testing tool website taken from measurements with Safari (see Section 4). We make the tool publicly available on accept to allow individuals to easily test and understand their client behavior.

measurement to get a better picture. This includes the client’s user agent and platform as provided by the browser and the client network’s Autonomous System (AS) number which are not considered personal identifiable information. We need this information to correctly attribute results to applications, OSes, and access networks. The findings for Safari with iCloud Private Relay were only possible by including the AS information. Nevertheless, we offer users the option to not submit results to our study and hence no data is collected while using the tool.

B Framework Details

Figure 3 provides an overview about our *Local Testbed Framework* as described in Section 4. Test cases, the host setup and tested clients are specified and adapted in one configuration. A test runner sets up the client and server node accordingly and iterates through all test cases, configurations and clients.

Figure 4 shows an example result of our CAD and RD web-based testing tool as described in Section 4. It allows users testing their individual setup (browser + resolver + OS) and their HE capabilities. The source code for the test framework and the web-based tool are publicly available. Both tools include a complete Ansible setup script. We will make the tool publicly accessible if the paper is accepted.

C Open Resolvers

We tested the HE capabilities of open resolvers listed in Table 4. We selected these from a publicly available list of open resolvers¹. Results regarding the capabilities of these resolvers are described in Section 5.3.

¹<https://gist.github.com/mutin-sa/5dcbd35ee436eb629db7872581093bc5>




Figure 5: Address selection based on an internal priority list. The plot shows the address family used by browsers at the n-th connection attempt.

Table 5: Operating systems and browsers for which we obtained web-based measurements. This information was extracted from the user agent. Therefore, some entries do not contain an OS version.

OS		Browser	
Name	Version	Browser	Browser Version
Android	10	Chrome Mobile	127.0.0
Android	10	Chrome Mobile	130.0.0
Android	10	Firefox Mobile	131.0
Android	10	Samsung Internet	26.0
Android	14	Firefox Mobile	125.0
Android	14	Firefox Mobile	128.0
Android	14	Firefox Mobile	131.0
Chrome OS	14541.0.0	Chrome	129.0.0
Linux		Chrome	130.0.0
Linux		Firefox	128.0
Linux		Firefox	130.0
Linux		Firefox	131.0
Linux		Firefox	132.0
Mac OS X	10.15	Firefox	128.0
Mac OS X	10.15	Firefox	131.0
Mac OS X	10.15	Firefox	132.0
Mac OS X	10.15.7	Chrome	127.0.0
Mac OS X	10.15.7	Chrome	129.0.0
Mac OS X	10.15.7	Chrome	130.0.0
Mac OS X	10.15.7	Opera	114.0.0
Mac OS X	10.15.7	Safari	17.4.1
Mac OS X	10.15.7	Safari	17.5
Mac OS X	10.15.7	Safari	17.6
Mac OS X	10.15.7	Safari	18.0.1
Ubuntu		Firefox	128.0
Ubuntu		Firefox	131.0
Windows	10	Chrome	127.0.0
Windows	10	Edge	130.0.0
Windows	10	Firefox	130.0
iOS	17.5.1	Mobile Safari	17.5
iOS	17.6	Mobile Safari	17.6
iOS	17.6.1	Mobile Safari	17.6
iOS	18.1	Mobile Safari	18.1

Table 4: Tested recursive resolvers. Hurricane Electric, Lumen (Level3), Dyn, and G-Core do not support IPv6 on the resolution path and are therefore not included in our evaluation presented in Section 5.3.

Service	# Addr.	
	IPv4	IPv6
DNS.sb	2	2
Google P. DNS	2	2
DNS0.EU	2	2
114DNS	2	0
Cloudflare	2	2
H-MSK-IX	2	2
MSK-IX	2	2
NextDNS	2	2
OpenDNS	6	6
Quad 101	2	2
Quad9 DNS	6	6
Verisign P. DNS	2	2
Yandex	2	2
G-Core	2	2
DYN	2	0
Lumen (Level3)	4	0
HE	4	4

114DNS does only offer IPv4 resolver addresses but is capable of resolving domain names on IPv6-only resolution paths. The WhoAmI tool by Akamai (whoami.ds.akahelp.net) shows an address in a different AS. Therefore, we assume it is a forwarder and uses an IPv6-capable recursive resolver.

D Address Selection

The main findings regarding address selection are described in Section 5.2. The following provides further information on the actually visible selection from different applications. Figure 5 plots the address family used by browsers on parallel connection attempts when the browser hits its configured CAD. We configured a total of ten addresses per family which do not offer a service but all time out. Only Safari retries as often as there are addresses. All others which implement the CAD only try one IPv6 and one IPv4 address. This is the described behavior in HEv1.

Safari does use a *First Address Family Count* of two and prefers IPv6. Safari’s interleaving strategy is to attempt one IPv4 address after the two IPv6 addresses. Then it continues with all remaining IPv6 addresses and only after the remaining IPv4 addresses are tried.

E Measured Browser and OS Versions

Table 5 contains the browser versions and the operating systems running these browsers used in our web-based measurements. We extract this information from the browser provided user agent string. Linux and Ubuntu do not provide the OS version in the user agent.