

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 15, 2014

V. Bajpai
J. Schoenwaelder
Jacobs University Bremen
July 14, 2013

Measuring the Effects of Happy Eyeballs
draft-bajpai-happy-01.txt

Abstract

The IETF has developed solutions that promote a healthy IPv4 and IPv6 co-existence. The happy eyeballs algorithm for instance, provides recommendations to application developers to help prevent bad user experience in situations where IPv6 connectivity is broken. This document describes a metric used to measure the effects of the happy eyeballs algorithm. The insights uncovered by analysing the data from multiple locations is discussed.

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 January 15, 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

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

Table of Contents

1.	Introduction	2
2.	IPv6 Upgrade Policy	2
3.	Happy Eyeballs	3
4.	Related Work	3
5.	Metric	4
6.	Implementation	4
7.	Measurement Trials	5
8.	Data Analysis Insights	5
9.	Conclusions	7
10.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

The function `getaddrinfo(...)` resolves a service name to a list of endpoints in an order that prioritizes an IPv6-upgrade path [[RFC6724](#)]. The order can dramatically reduce the application's responsiveness when IPv6 connectivity is broken. The degraded user experience can be subverted by implementing the happy eyeballs algorithm [[RFC6555](#)]. The algorithm recommends that a host, after resolving the service name, tries a `TCP connect(...)` to the first endpoint. However, instead of waiting for a timeout, it waits for 300ms, after which it must initiate another `TCP connect(...)` to an endpoint with a different address family and start a competition to pick the one that completes first.

This document describes a metric used to measure the effects of the happy eyeballs algorithm. The insights uncovered by analysing the data from multiple locations is discussed.

[2.](#) IPv6 Upgrade Policy

The happy eyeballs algorithm as defined in [[RFC6555](#)] biases its path selection in favor of IPv6 by design. The connection establishment race has been handicapped for the following reasons:

- o Carrier-grade NATs (CGNs) establish a binding for each connection request. Dual-stack hosts by preferring IPv6 connection routes, reduce their contention towards the critical IPv4 address space.
- o The IPv4 traffic may be billed by Operation Support Systems (OSS) in some networks. Techniques that help move this traffic to IPv6 networks reduce costs.

- o Middleboxes maintain state for each incoming connection request. If the dual-stacked hosts prefer IPv6 path, the load on load balancers and peering links reduces automatically. This reduces the investment on IPv4, and encourages IPv6 migration.

3. Happy Eyeballs

The happy eyeballs algorithm defined in [[RFC6555](#)] honors the IPv6 upgrade policy. It is therefore not designed to encourage aggressive connection requests over IPv4 and IPv6, but instead to satisfy the following goals:

- o The connection requests must be made in an order that honors the destination-address selection policy as defined in [[RFC6724](#)], unless overridden by user or network configuration. The client must prefer IPv6 over IPv4 whenever the policy is not known.
- o The connection initiation must quickly fallback to IPv4 to reduce the wait times for a dual-stack host in situations where the IPv6 path is broken.
- o The network path and destination servers must not be thrashed by mere doubling of traffic by making simultaneous connection requests over IPv4 and IPv6. The connection requests over IPv6 must be given a fair chance to succeed to reduce load on IPv4, before a connection over IPv4 is attempted.

However, applications on top of TCP will not be happy eyeballed only in scenarios where IPv6 connectivity is broken, but also in scenarios where the dual-stack host enjoys comparable IPv6 connectivity. We want to measure how much imposition does such a user experience in reality by measuring the effects of the happy eyeballs timer value.

The recommended timer value is 150-250ms [[RFC6555](#)]. However, Chrome uses 300ms. Firefox appears to be using 250ms while an early open-source implementation of happy eyeballs seems to recommend 100ms [[Perreault](#)]. We want to affirm the right value by measuring TCP connection establishment times experienced by dual-stacked hosts in real environments over IPv4 and IPv6.

4. Related Work

Fred Baker in [[RFC6556](#)] describes metrics and testbed configurations to measure how quickly an application can reliably establish connections from a dual-stacked environment. The metrics measure whether the communication establishment time is same regardless of the address family and the routing viability available to a dual-stacked host. The metrics defined in [[RFC6556](#)] is different in three ways:

- o DNS is accounted in connection establishment time. Our metric does not take this into account. Accounting DNS resolution may invite multiple input factors (slow resolvers) that may bias our TCP connection establishment time results. In addition, according to [[RFC6555](#)], the 300ms advantage applies to the first address family after the `getaddrinfo(...)` call. From a programming perspective, an application calls `getaddrinfo(...)` and that does its job, regardless of which address family is used.
- o The testbed configuration in [[RFC6556](#)] is more passive than active. An external analyser is used to passively observe the client's traffic using `tcpdump`. There is no active measurement test, instead the routers along the path are configured to control what connectivity route is taken. We on the other hand, have an active measurement test running on the client. The test is agnostic to network path configuration since it independently tries a TCP connection to each connectivity route. It also actively measures the time taken instead of relying on an external analyser program.
- o The testbed setup in [[RFC6556](#)] is designed for a controlled environment. The router in the path is configured to disrupt all but one routes to control the prefix used in the connection. As such, the test is repeated N times with different router configurations to try all possible permutations of route connectivity. Our measurement test is agnostic to the network path and does not require path configuration changes.

5. Metric

We have defined a metric that uses the TCP connection establishment times as a parameter to measure the algorithm's effects. The methodology also helps examine the impact of tunneling mechanisms employed by early adopters. The input parameter of the metric is a (IP address, port number) tuple and the output is the connection establishment time, typically measured in microseconds.

6. Implementation

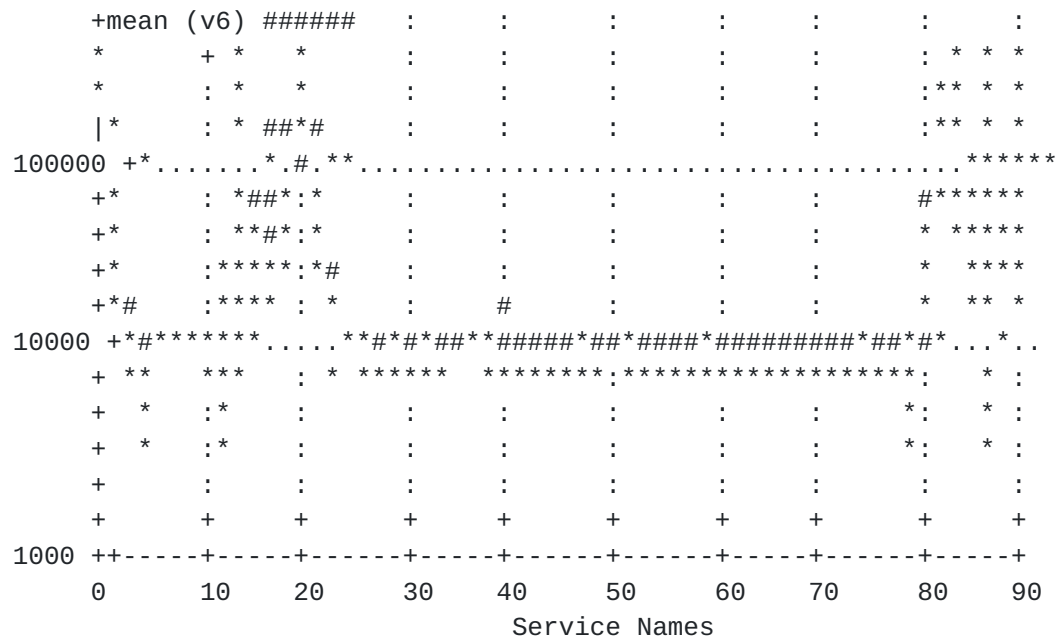


Figure 1: service vs {mean_v4, mean_v6}: samsbox1 (30 days, 300ms)

Fig. 1. shows the average TCP connection establishment times for both IPv4 and IPv6. The Measurement Agent (MA) is a SamKnows probe connected at Jacobs University Bremen. It receives IPv4 and IPv6 connectivity via German Research Network (DFN) [AS 680]. A PDF rendering of the plot is available at [[mean](#)].

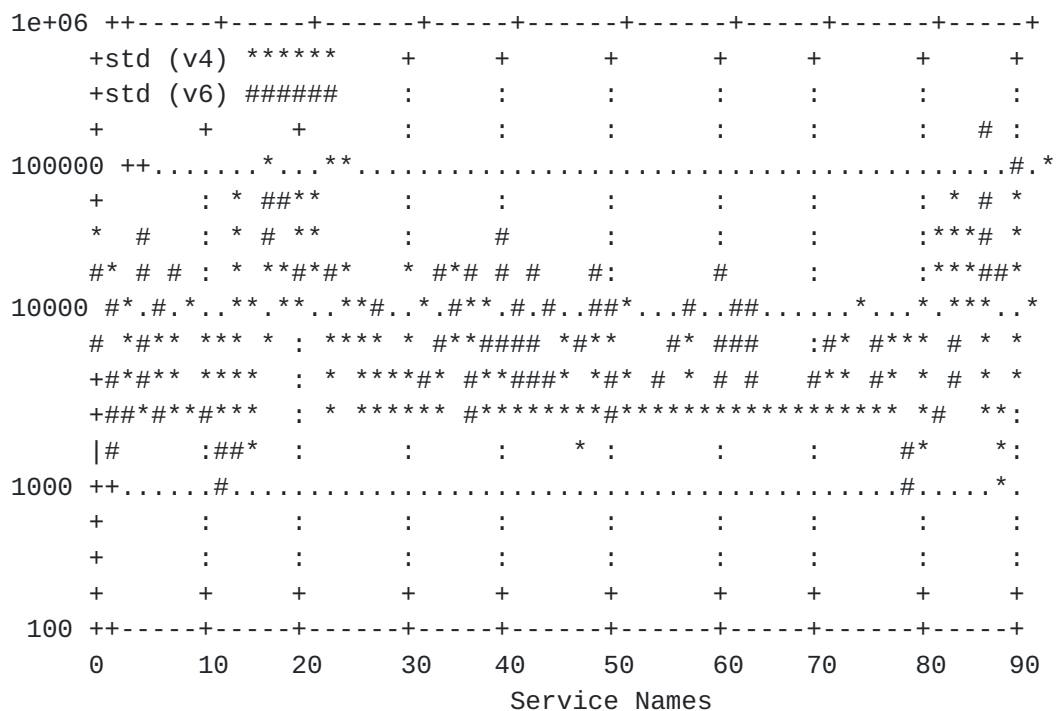


Figure 2: service vs {std_v4, std_v6}: samsbox1 (30 days, 300ms)

Figure 2 shows the standard deviation of the TCP connection establishment times for both IPv4 and IPv6. The Measurement Agent (MA) is a SamKnows probe connected at Jacobs University Bremen. It receives IPv4 and IPv6 connectivity via German Research Network (DFN) [AS 680]. A PDF rendering of the plot is available at [[std](#)].

It appears that an application never uses IPv6 using Teredo except in situations where IPv4 reachability of the destination service is broken. We noticed, that a 300ms advantage leaves a dual-stacked host only 1% chance to prefer a IPv4 route even though it may be significantly faster than IPv6. We also measured the margin by which happy eyeballs is inhibiting the fastest available route by comparing the slowness of a happy eyeballed winner to that of the loser.

9. Conclusions

We have performed a preliminary study on measuring the effects of happy eyeballs. We noticed several cases where the algorithm does not select the best route and instead hampers the user experience. We are working towards running this test on a large-scale measurement platform to develop a more comprehensive picture to help improve the algorithm.

10. Informative References

[I-D.ietf-6man-addr-select-opt]

Matsumoto, A., Fujisaki, T., and T. Chown, "Distributing Address Selection Policy using DHCPv6", [draft-ietf-6man-addr-select-opt-10](#) (work in progress), April 2013.

[Perreault]

Perreault, S., "Happy Eyeballs in Erlang", July 2013, <http://www.viagenie.ca/news/index.html#happy_eyeballs_erlang>.

[RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", [RFC 6555](#), April 2012.

[RFC6556] Baker, F., "Testing Eyeball Happiness", [RFC 6556](#), 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.

[mean] Bajpai, V., "IPv4 and IPv6 Average Connection Establishment Times", July 2013, <<http://cnds.eecs.jacobs-university.de/users/vbajpai/ietf87-v6ops/samsbox1-mean.pdf>>.

[std] Bajpai, V., "IPv4 and IPv6 Connection Establishment Times Variations", July 2013, <<http://cnds.eecs.jacobs-university.de/users/vbajpai/ietf87-v6ops/samsbox1-std.pdf>>.

Authors' Addresses

Vaibhav Bajpai
Jacobs University Bremen
Campus Ring 1
28759 Bremen
Germany

Phone: +49 421 200 3112

Email: v.bajpai@jacobs-university.de

Juergen Schoenwaelder
Jacobs University Bremen
Campus Ring 1
28759 Bremen
Germany

Phone: +49 421 200 3587

Email: j.schoenwaelder@jacobs-university.de