

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 05, 2014

V. Bajpai
J. Schoenwaelder
Jacobs University Bremen
July 04, 2013

Measuring the Effectiveness of Happy Eyeballs
draft-bajpai-happy-00.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 effectiveness 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 05, 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

Measuring Happy Eyeballs

July 2013

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

Table of Contents

1.	Introduction	2
2.	Related Work	2
3.	Metric	3
4.	Implementation	3
5.	Measurement Trials	4
6.	Data Analysis Insights	4
7.	Conclusions	6
8.	Informative References	6
	Authors' Addresses	6

[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 effectiveness of the happy eyeballs algorithm. The insights uncovered by analysing the data from multiple locations is discussed.

[2.](#) 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.

[3.](#) Metric

We have defined a metric that uses the TCP connection establishment times as a parameter to measure the algorithm's effectiveness. 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.

[4.](#) Implementation

We have developed happy, a simple TCP happy eyeballs probing tool that conforms to the definition of our metric. It uses non-blocking

connect(...) calls to concurrently establish connections to all endpoints of a service and measures the elapsed time. The tool enforces a small delay between concurrent connect(...) calls to avoid bursty TCP SYN traffic. The initially performed service name resolution is not accounted in the connection establishment elapsed time.

5. Measurement Trials

We use Alexa's top 1M service names as input to prepare a top 100 dual-stacked service names list. We run happy on our internal test-bed of multiple measurement agents with different flavors of connectivity ranging from native IPv4, native IPv6, IPv6 tunnel broker endpoints, Teredo and tunnelled IPv4. The list of Measurement Agents (MAs) is shown in Table 1.

MA #	IPv4 AS	IPv6 AS	City	Country	Platform
1	AS680	AS680	Bremen	Germany	Mac OS X
2	AS680	AS680	Braunschweig	Germany	GNU/Linux
3	AS13237	Teredo	Berlin	Germany	GNU/Linux
4	AS31334	AS6939	Bremen	Germany	OpenWrt
5	AS680	AS680	Bremen	Germany	SamKnows
6	AS31334	AS6939	Bremen	Germany	SamKnows
7	AS24956	AS24956	Braunschweig	Germany	SamKnows
8	AS3320	AS3320	Bremen	Germany	SamKnows
9	AS5607	AS5607	London	England	SamKnows
10	AS3269	AS3269	Torino	Italy	SamKnows
11	AS8903	AS8903	Madrid	Spain	SamKnows
12	AS2614	AS2614	Timisoara	Romania	SamKnows
13	AS13030	AS13030	Olten	Switzerland	SamKnows
14	AS2856	AS2856	Ipswich	England	SamKnows

Table 1: A List of Measurement Agents (MAs)

6. Data Analysis Insights

The initial results show higher connection times and variations over IPv6 as shown in Figure 1.






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 when IPv4 connectivity 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.

7. Conclusions

We have performed a preliminary study on measuring the effectiveness 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.

8. Informative References

- [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

