Analysis of the Effects of COVID-19-Related Lockdowns on IMRS Traffic

ICANN Office of the Chief Technology Officer

Roy Arends

OCTO-008
15 April 2020

TABLE OF CONTENTS

EXECUTIVE SUMMARY 3

1 INTRODUCTION 3

2 METHODOLOGY 5

2.1 Classification 5
2.1.1 Chrome Queries 5
2.1.2 Jumbo Queries 7
2.1.3 Popular Nonexistent TLDs 7
2.1.4 Others 7

3 OBSERVATIONS 7

3.1 Chromium Queries 8
3.2 Jumbo Queries 9
3.3 Popular Nonexistent TLDs 9

4 CONCLUSION 9

This document is part of the OCTO document series. Please see https://www.icann.org/resources/pages/octo-publications-2019-05-24-en for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

Executive Summary

Restrictions during COVID-19-related lockdowns and school closures are expected to have a limited, but noticeable effect on the Domain Name System (DNS) traffic at ICANN Managed Root Servers (IMRS). ICANN's Office of the Chief Technology Officer (OCTO) has studied the impact of a nationwide lockdown in France on changes in both traffic volume and composition to the four IMRS instances in France.

The Réseaux IP Européens Network Coordination Centre (RIPE NCC) Atlas probes showed that traffic for the French IMRS instances mostly originated in France. The lockdown in France started 17 March 2020 (week 12 of 2020). Traffic statistics for this week showed a 28% increase compared to the average of the previous 6 weeks. A comparative analysis was done between week 6 and week 12, and the following categories were compared:

- Queries for existing top-level domains (TLDs)
- Queries that originate from Chromium-based browsers
- Queries for large TLDs
- Queries for popular TLDs (.home, .lan, .corp, and .local)
- All other queries

Most categories had an increase in traffic, contributing to the overall increase. The largest category of queries originated from Chromium browsers, which remained at about one third of all received requests. Some categories grew faster than others. The largest percentage increase came from the four categories of popular nonexistent TLDs (.corp, .home, .lan and .local). This is likely due to people working more from home, as normally workers are congregated in offices using a set of resolvers that understand how to respond to .corp, .lan, and .local domains. Now, they are now more dispersed and working from home using resolvers that may not understand how to respond to these domains. This would also explain the increase in .home queries: more people using the Internet more often from their homes.

The effects of nationwide lockdowns have had a limited, but noticeable effect on the DNS traffic at IMRS instances when observed at a country level. This increase in DNS traffic can be observed overall and the fact that no issues have arisen suggests that the DNS architecture is well suited to scale during remote work and increased use at home.

1 Introduction

The effects of nationwide lockdowns, restrictions on activity, and school closures are expected to have a limited, but detectable effect on the DNS traffic at IMRS servers. Generally speaking, the bulk of the DNS traffic seen at the IMRS stems from resolvers that submit DNS requests on behalf of clients such as mobile phones, tablets, personal computers (laptops and desktops), game consoles, etc. These resolvers have the ability to temporarily cache information which dampens the load on the root servers. For example, when a resolver has cached information about the nameservers for the .com namespace, it does not need to contact the root servers for information about example.com, it only needs to ask the .com name servers.

At the time of writing (31 March 2020), the IMRS consists of 167 instances located in 83 countries. This study is focused on the four IMRS instances in France. The motivation for the focus on these instances is that in France, school closures, restrictions on activity, and a nationwide lockdown were announced by the government in quick succession. On 12 March, the government announced that schools and universities would close from Monday, 16 March. On 13 March, gatherings of more than 100 people were banned. On 14 March, the closure of all nonessential public places, including restaurants, cafés, cinemas, and discothèques was ordered. On 16 March, a national lockdown was ordered starting the following day.

Traffic to IMRS servers originates from a wide variety of sources that do not necessarily reside in the same country as the instances queried. Using RIPE Atlas[1] probes as a proxy for resolver clients, we can visualize which individual probes use the four IMRS instances currently located in France.

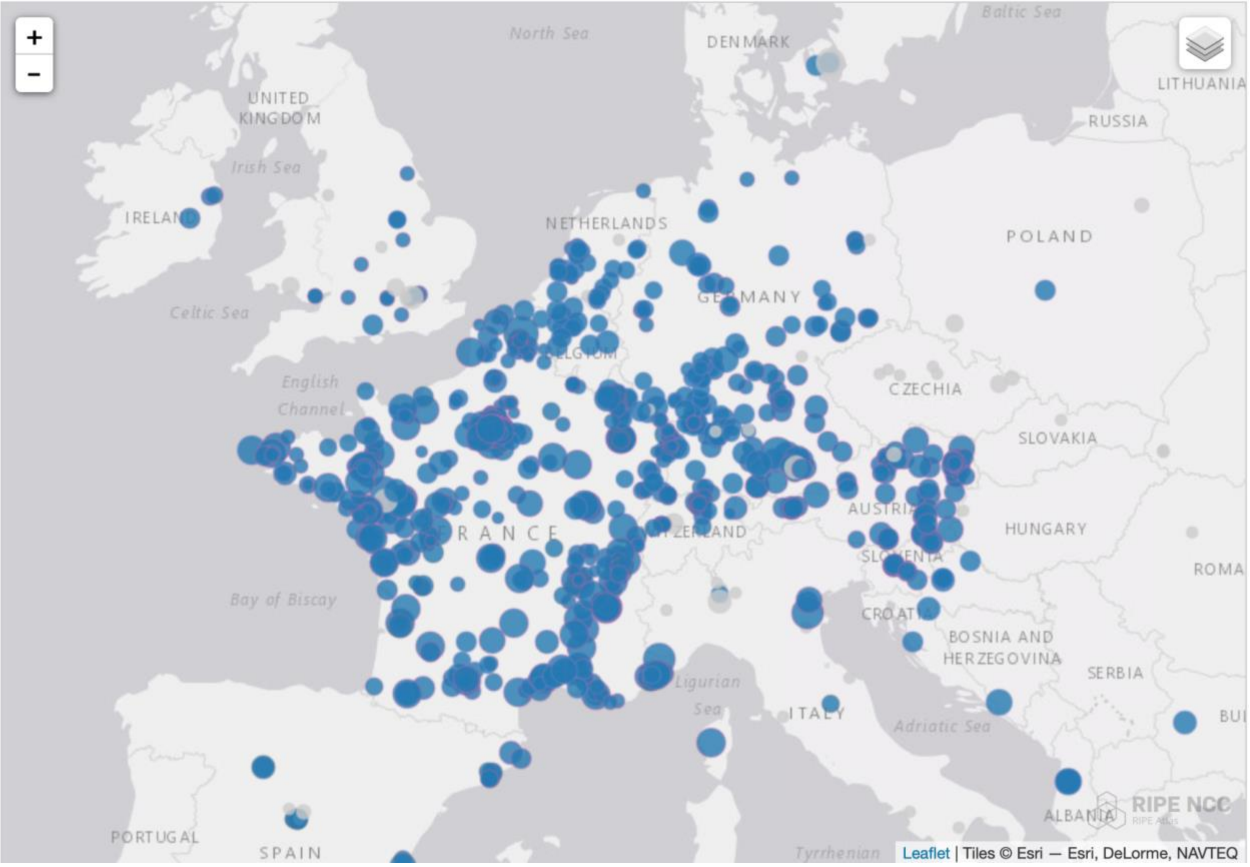

Leaflet | Tiles © Esri — Esri, DeLorme, NAVTEQ

Figure 1. Distribution of Atlas probes in the catchments of IMRS instances located in France

As seen in Figure 1, while there are a fair amount of probes located outside of France that have seen a response from aforementioned IMRS instances, a significant amount of traffic for the IMRS instances in France originates in France.

[1] RIPE Atlas is a global, open, distributed Internet measurement platform, consisting of thousands of measurement devices that measure Internet connectivity in real time.

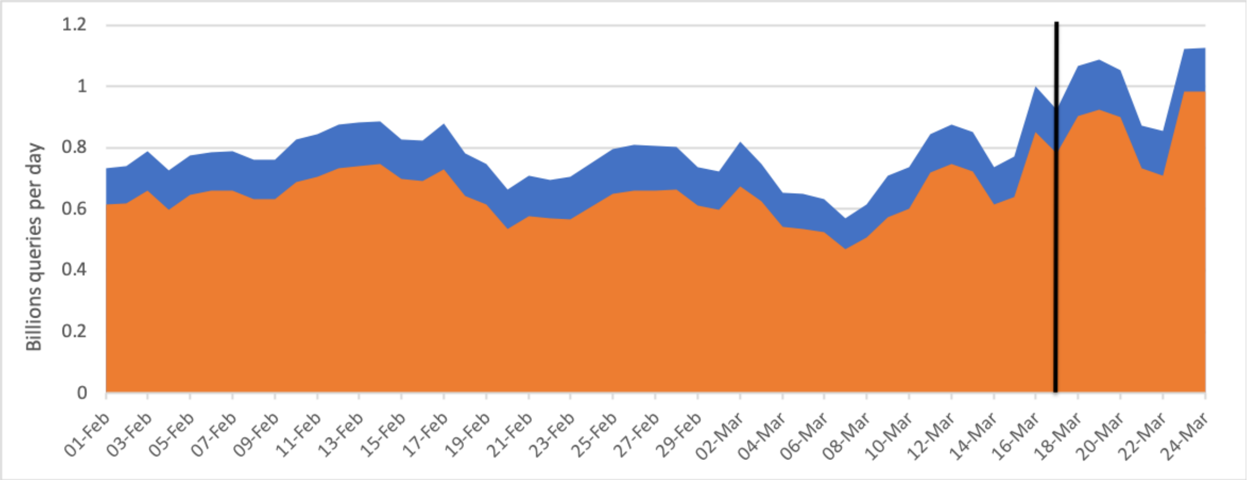

Figure 2. Daily query volume (blue) and daily NXDOMAIN response volume (orange) observed at the 4 IMRS instances in France. The black vertical line indicates the start of March 17.

As seen in Figure 2, there was a rise in traffic volume after 16 March. To understand what drove this increase, we examined the composition of the traffic. We will compare the composition before and after 16 March to see if we can correlate the changes in composition to the lockdowns.

2 Methodology

We will compare two weeks of traffic. The first week of February (week 6, starting February 3) against the week of March 16 (week 12) which was the first week of the lockdown. We will then subsequently classify parts of that traffic and show which classification has the most significant changes.

2.1 Classification

Traffic is grouped into several categories based on the TLD that is queried for:

- **Exists**: Queries for TLDs that are currently delegated from the root zone
- **Chrome**: Queries for nonexistent TLDs between 7 and 15 characters long
- **Jumbo**: Queries for nonexistent TLDs longer than 15 characters long
- **.home**: Queries for domains that end in .home
- **.lan**: Queries for domains that end in .lan
- **.local**: Queries for domains that end in .local
- **.corp**: Queries for domains that end in .corp
- **Others**: Queries for all other domains

2.1.1 Chrome Queries

The Chromium web browser and derivatives (such as Google Chrome, recent versions of Microsoft Edge, Amazon Silk, and Opera’s web browser) issue three DNS requests with a random label to detect if the resolver in use on the local network redirects nonexistent domains, e.g., if the query returns the address of a “helper” search website for domains that do not exist.

The label consists of random letters and is between 7 and 15 characters long.[2] Since the domain queried is random, the receiving resolver will not have it cached and will issue a query to a root server. In networks without redirection, the expected response of that random query would be an NXDOMAIN error code.

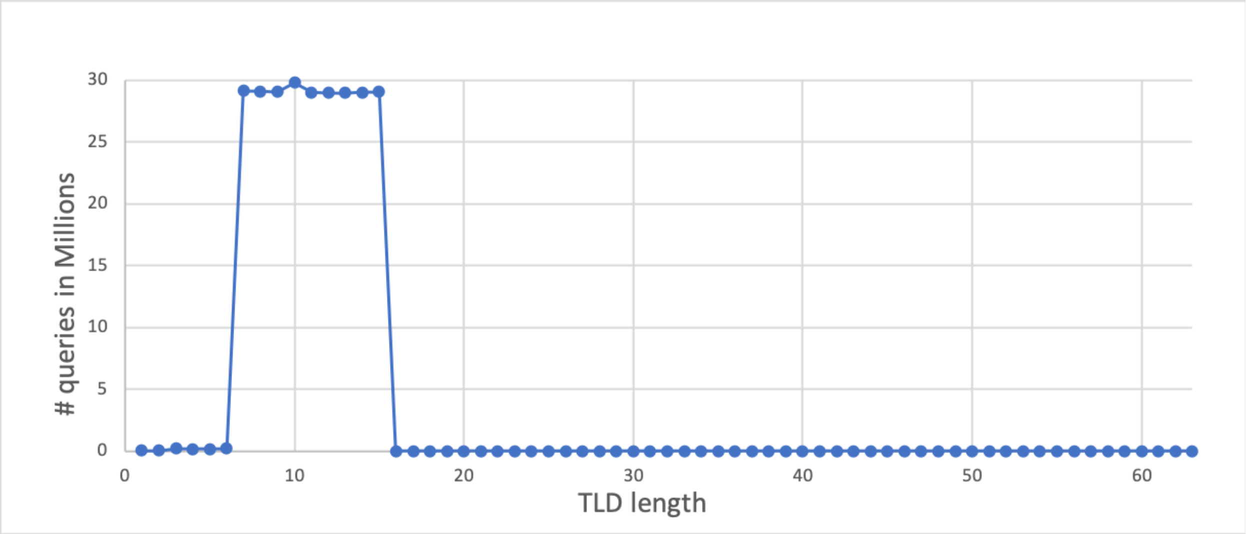

Figure 3. Histogram of number of queries for nonexistent TLD per TLD length.

The histogram in Figure 3, displaying data from 19 March shows the frequency distribution of queries per TLD length. The bulk of these queries are for domain names in the range between 7 and 15 characters. Figure 5 shows that these Chrome queries form 28% of all nonexistent domain queries.

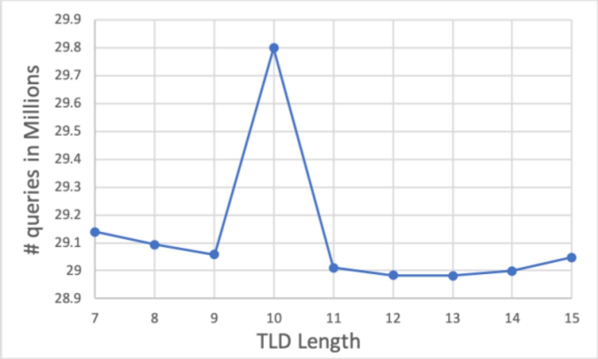

Figure 4. Detail of histogram of number of queries for nonexistent TLD per TLD length.

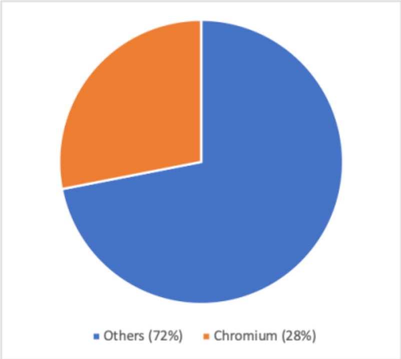

Figure 5. The number of Chromium queries over all nonexistent domain queries.

Other than 10 character TLDs, the distribution between 7 and 15 character TLDs is fairly uniform. The anomaly for labels of 10 characters can be ascribed to older versions of Chrome which issued random domains of 10 characters.[3]

[2] "We generate a random hostname with between 7 and 15 characters." https://chromium.googlesource.com/chromium/src/+/master/chrome/browser/intranet_redirect_detector.cc#150

[3] "Vary the length of DNS hijack detection names."

2.1.2 Jumbo Queries

These are queries for nonexistent TLDs longer than 15 characters. We're unaware of the sources or causes of these queries.

2.1.3 Popular Nonexistent TLDs

There is a range of popular labels that have not been delegated at the root and do not exist in the Internet's public DNS namespace. Among the most popular of these nonexistent TLDs are .home, .lan, .corp, and .local. These TLDs are categorized individually as they all increased in volume during our study.

2.1.4 Others

This category catches all the queries that cannot be categorized in any of the other categories previously described.

3 Observations

The four IMRS instances in France received 5.4 billion DNS requests per week on average between weeks 6 and 11 (see Figure 6). The same instances received 6.9 billion DNS requests in week 12. This represents a 28% increase in traffic to those four IMRS nodes.

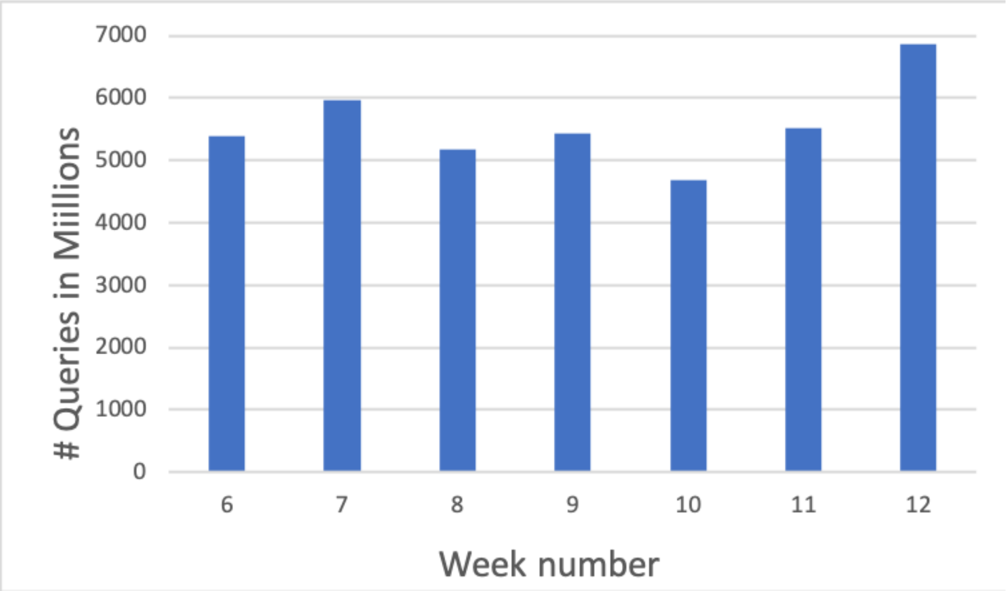

Figure 6: Query load on the 4 instances in France per week from week 6 to week 12.

We did capture some anomalies such as short bursts of traffic or instance maintenance outages over that time period, but these tended to be short-lived and we do not think they significantly

https://src.chromium.org/viewvc/chrome?view=revision&revision=249013

influence overall traffic. Other traffic patterns, such as diurnal patterns or weekends, are absorbed as well since the traffic was accumulated over a week. We are unaware of any other changes or events over this period of time that would influence the volume of DNS queries so significantly.

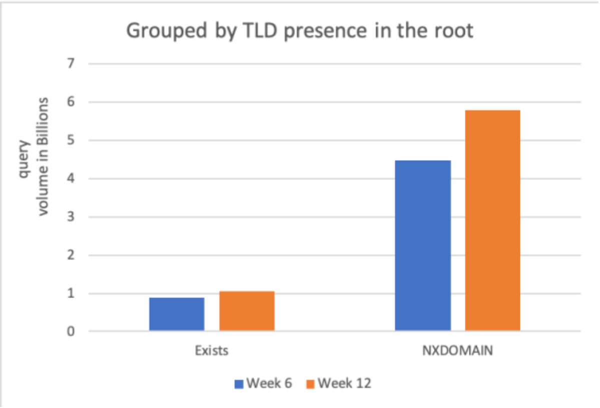

Figure 7: Volume of traffic for existing and nonexistent TLDs in weeks 6 and 12

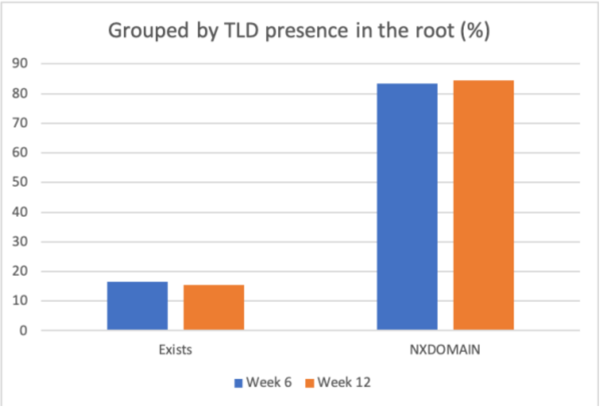

Figure 8: Volume of traffic for existing and nonexistentTLDs in weeks 6 and 12 as a percentage of the total volume for those weeks

Figure 7 shows the difference in query volume for existing and nonexistent domains in absolute numbers. Both groups have grown in volume. Figure 8 shows that there is a small shift in the composition of the traffic as well, since the percentage of queries for existing TLDs has declined compared to those for nonexistent domains. The traffic increase is mostly in queries for nonexistent domains.

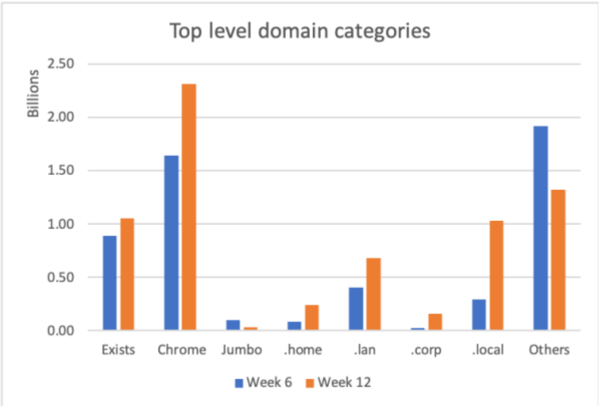

Figure 9: Breakdown of traffic in various categories, comparing weeks 6 and 12 in absolute numbers.

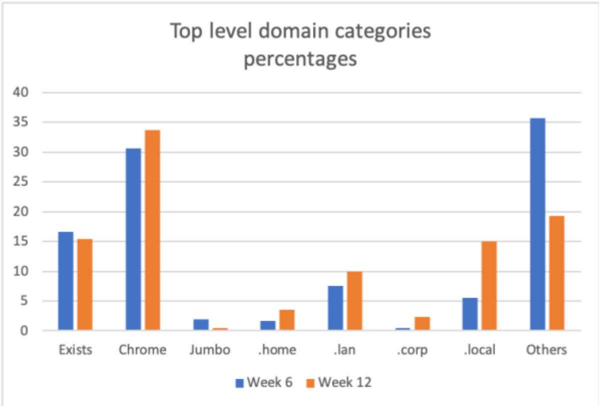

Figure 10: Breakdown of traffic in various categories, comparing weeks 6 and 12 as a percentage of total volume.

3.1 Chromium Queries

We observed that 31% of the received requests in week 6 and 34% in week 12 fall into the category of Chromium DNS requests. This remains a significant part of the overall traffic. After the lockdown, the total number of requests increased by 28%, while the Chromium part of those requests increased by 41%. The increase is likely due to more devices coming online more often due to stay-at-home mandates.

Due to Chromium's DNS redirection detection queries, there will be a higher rate of DNS requests for random strings between 7 and 15 characters in length visible in the traffic when more devices with Chromium-based browsers come online. Note that queries for Chromium DNS requests have not increased by the same percentage as the overall traffic. This indicates a slight shift in the composition of the overall traffic. Other categories have seen a higher increase than the Chromium queries.

The Chromium queries are the largest single cause of queries to root servers. Other IMRS instances often see over 50% of all incoming queries from Chromium. The purpose of these queries is to check if Chromium is behind a captive portal. Provisioning for root servers is often a function of the overall load on root servers to satisfy the scaling needs. While these queries are free for Chromium to make, the cost of provisioning for root-server instances is not. Google has been notified of this issue, but it remains outstanding.[4]

3.2 Jumbo Queries

We've observed that the volume of requests with large TLD domains (larger than 15 characters) has decreased. We have not looked into the reason behind this drop in traffic.

3.3 Popular Nonexistent TLDs

The four most popular nonexistent TLDs that saw an increase in volume were .corp, .home, .lan, and .local. Of these, .corp, .lan, and .local saw the most significant increase. This is likely due to people working more from home. Normally workers are congregated into offices using a set of resolvers that understand how to respond to .corp, .lan, and .local domains. Now they are more dispersed and working from home using resolvers that may not understand how to respond to these domains. This would also explain the increase in .home queries: more people using the Internet more often from their homes.

4 Conclusion

The effects of nationwide lockdowns to contain the global pandemic have had a limited but noticeable effect on the DNS traffic at IMRS instances when observed at a country level. This increase in DNS traffic can be observed overall. The fact that no issues have been seen suggests that the DNS architecture is well suited to scale during scenarios of remote work and increased use at home.

Authors: Adiel Akplogan, Roy Arends, David Conrad, Alain Durand, Paul Hoffman, David Huberman, Matt Larson, Sion Lloyd, Terry Manderson, David Soltero, Samaneh Tajalizadehkhoob, Mauricio Vergara Ereche.

[4] The three random probes by the intranet redirect detector do not have a TLD and thus reach the root servers. https://bugs.chromium.org/p/chromium/issues/detail?id=946450&q=intranet%20redirect&can=2