November 2014

Geoff Huston

The Resolvers We Use

The Internet's Domain Name System is a modern day miracle. It may not represent the largest database that has ever been built, but nevertheless it's truly massive. And even if it's not the largest database that's ever been built, it's perhaps one of the more intensively used. The DNS is consulted every time we head to a web page, every time we send an email message, or in fact every time we initiate almost any transaction on the Internet. It's the essential bridge between a world of human names and the underlying world of binary protocol addresses. And it's fast. Fast enough that it's still largely invisible as part of the user experience, despite continued growth in size. Given the fragmentation of the IPv4 address space with the widespread use of various forms of address sharing, then it increasingly looks as if the DNS is the only remaining common glue that binds the Internet together as a single network.

Not only is the DNS a miracle, we want the miracle to keep on getting better! We now want to improve the level of privacy in the DNS by minimizing the unnecessary exposure of information in our queries. We want it to be secure so that we can validate the responses that we get to ensure that they are the genuine article. And in a world replete with replicated content because of burgeoning content distribution networks, we would now like the DNS to be location-sensitive, so that when there are multiple possible responses to our query we would like it to provide us with the "best" answer.

It's this last aspect I'd like to examine here. Location sensitivity has always been important in the Internet, but as we voice a common requirement for greater performance, and greater levels of responsiveness and immediacy from applications, then we need to attack latency. We need to move content closer to users, and to do so notions of "locality" are becoming increasingly critical in today's Internet. For example, content distribution networks are observed to make use of the assumed location of the resolver that is querying the DNS in order to provide a DNS response that is intended to direct the user to a content distribution point that is assumed to be as close as possible to the user, to assist in the performance of subsequent content delivery. How robust is this assumption of co-locality of users and their resolvers? Are users always located "close" to their resolvers?

More generally, what is the relationship between the end user, and the DNS resolvers that they use? Are they in fact closely related? Or is there widespread use of distant resolvers?

Mapping Users to Resolvers

At APNIC Labs we've been using an online ad network to measure the rate of adoption of IPv6 and secure DNS (DNSSEC). We do this by embedding within the advertisement material a small code segment that requests the user's browser to retrieve a test set of URLs. In the case of IPv6 one of these URLs is only accessible if the user can perform the retrieval using IPv6. In the case of DNSSEC one of the DNS names is DNSSEC signed. That way we can use the impressed advertisement as another sample point in the measurement.

If you do it enough times across a large enough random sample of end users, the resultant picture is a good reflection of the capabilities of the end user population of the larger Internet in its entirety. However, it's possible to extract more information from these simple experiments. The IP address resolver that queries our authoritative name server for a DNS name associated with the measurement

experiment can be matched against the IP address of the end user who subsequently performs a retrieval of the web object from the experiment's web server. This allows us to match users to the resolvers that they use.

This ability to match users to the resolvers that they use probably needs a little further clarification.

The DNS uses a deceptively simple protocol. You ask a query and you get a response. It's intentionally very simple and potentially very fast. But while the protocol is simple, its implementation is not.

But the DNS resolver that you are querying is not necessarily the font of all knowledge about domain names. It does not maintain a comprehensive local store of the entire DNS database. When it receives your query it has to resolve the name. If you, or someone else who uses the same resolver, has queried this resolver for the same DNS name in the recent past then the local resolver may have the answer stored in a local cache, and it can answer immediately. Otherwise its going to have to discover the answer on your behalf by making its own queries, using the same DNS query protocol. Either this resolver is configured to use another resolver as a "forwarder" or it's configured to find out for itself. If this resolver uses a forwarder then your query may trigger your resolver to ask a query from the forwarder resolver, which may in turn cascade along a path of forwarders. It the name has never been used in the past then its details will not be held in any of these resolvers' caches, so at some point, at the end of this chain of forwarders, a resolver will ask one of the authoritative name servers for the name value.

By the time the query arrives at an authoritative name server this record of forwarders who have handled the query is not visible in the query. The query is a simple query that contains only the identity of the resolver who is posing the question. There is no trace information in the query as to the original entity who posed the original question, or any trace as to the chain of forwarding resolvers who may have passed the query on.

Things can be further complicated by the use of so-called "resolver farms" where a query is passed to a "farm" of DNS resolvers, and a query may be passed to any one or more of the resolvers in the farm. Understanding that a set of queries seen at an authoritative name server is actually the result of a single query by a user is often a challenging exercise in such cases.

All this implies that we do not necessarily have visibility on the resolvers that are used by end users. What we can see is the resolver at the end of the query forwarding chain, which is the resolver that ends up asking the query from the authoritative name server.

In this measurement experiment a uniquely named URL is loaded into a user's browser, and then the measurement system collects the DNS queries that arrive at the authoritative name server, and matches the address of the resolver that posed these queries to the authoritative name server to the IP address of the user's browser that subsequently retrieves the named URL.

Resolvers

Across 2014 we've does this some 84,060,879 times, and we've collected some 539,011 unique resolver IP addresses who have passed to the authoritative name server some 1,837,294,171 queries. On average it looks like each resolver is used by 155 clients. But in this case the average is hiding some considerable detail. The experiment also includes a mix of DNS names that are unsigned, signed with DNSSEC and badly signed with DNSSEC. In this case it may be useful to look only at the resolvers that are involved in querying for an unsigned domain name, in order to filter out the additional queries associated with DNSSEC-validating resolvers.

If we look at the top 10 visible resolvers in terms of the number users who pass queries through forwarder chains to these resolvers then we can gain some insight as to the most heavily used resolvers – those resolvers that handle the largest population of end users. Here's the ten largest resolvers when using that metric.

Rank	Resolver	ASN	User Share	Network Name
1	58.217.249.160	4134	0.23%	CHINANET-BACKBONE, CN
2	58.217.249.138	4134	0.23%	CHINANET-BACKBONE, CN
3	58.217.249.137	4134	0.22%	CHINANET-BACKBONE, CN
4	58.217.249.161	4134	0.21%	CHINANET-BACKBONE, CN
5	123.125.81.212	4808	0.17%	CHINA169-BJ Beijing Province Network, CN
6	74.125.189.17	15169	0.14%	GOOGLE - Google Inc., US
7	74.125.189.20	15169	0.14%	GOOGLE - Google Inc., US
8	74.125.189.18	15169	0.14%	GOOGLE - Google Inc., US
9	74.125.189.16	15169	0.14%	GOOGLE - Google Inc., US
10	74.125.189.23	15169	0.14%	GOOGLE - Google Inc., US

Table 1 – The Top 10 DNS resolvers

Within this list of 10 visible resolvers there are only three distinct /24 subnets being used. It's a reasonable assumption that the resolvers that are addressed from a common subnet are likely to be part of a common resolver server farm. What if we look at a slightly different definition of a visible resolver that unites all resolvers that sit on a common /24 subnet, and call the result a "visible resolver system". What are the 10 largest visible resolver systems in that case?

Rank	Resolver	ASN	User Share	Network Name
1	74.125.41.0	15169	1.28%	GOOGLE – Google Inc., US
2	58.217.249.0	4134	1.02%	CHINANET-BACKBONE, CN
3	74.125.189.0	15169	0.98%	GOOGLE - Google Inc., US
4	74.125.16.0	15169	0.94%	GOOGLE - Google Inc., US
5	74.125.190.0	15169	0.90%	GOOGLE - Google Inc., US
6	74.125.181.0	15169	0.87%	GOOGLE - Google Inc., US
7	61.140.11.0	4813	0.80%	GUANGDONG-AP China Telecom Group, CN
8	60.215.138.0	4837	0.80%	CHINA169, CNCGROUP China169 Backbone, CN
9	74.125.73.0	15169	0.78%	GOOGLE – Google Inc., US
10	74.125.17.0	15169	0.74%	GOOGLE – Google Inc., US

Table 2 – The Top 10 DNS resolver systems (by /24 subnet)

In this picture of resolver systems, 6 of these 10 largest DNS resolvers are operated by Google as part of its Public DNS system. What if we group together all these Google Public DNS systems into a single resolver system? At the same time, let's also group together the resolvers in Comcast (AS7922), OpenDNS (AS36692) and ChinaNet (AS4134), and categorize each of these as a single visible resolver system as well.

ChinaNet (AS4134) uses a different DNS server deployment framework than many other large ISPs. Instead of funneling all DNS queries into one or two very large forwarding resolver farms that then pass the queries to the authoritative name servers, the Chinanet DNS resolution architecture appears to use a highly distributed DNS resolver deployment without concentrating the forwarded queries into a smaller set of caching recursive resolvers. In this experiment some 5,000 different DNS resolvers in AS4134 were seen to query the experiment's authoritative name server, and 1,031 of these were observed to be acting for two or more different end clients. For the ChinaNet numbers these 1,031 resolvers have been grouped into a single report line, but it is not a true like-for-like comparison with the other resolvers listed here, as these resolvers within ChinaNet appear to be acting autonomously, and are not part of a larger common cache DNS resolver system. Although earlier observations appear to indicate that these resolvers operate with a common name filter set.

Rank	Resolver	AS	User Share	Cum.	Network Name
1	x.x.x.x	4134	10.55%	10.55%	CHINANET-BACKBONE, CN
2	8.8.8.8	15169	10.52%	21.07%	GOOGLE - Google Inc., US
3	69.252.66.0	7922	1.76%	22.83%	COMCAST-7922 - Comcast Cable Communications, US
4	61.140.11.0	4813	0.80%	23.63%	BACKBONE-GUANGDONG-AP China Telecom, CN
5	60.215.138.0	4837	0.80%	24.43%	CHINA169-BACKBONE CNCGROUP China169 Backbone, CN
6	208.69.34.0	36692	0.71%	25.14%	OPENDNS - OpenDNS, LLC, US
7	218.248.255.0	9829	0.70%	25.84%	BSNL-NIB National Internet Backbone, IN
8	195.175.255.0	9121	0.59%	26.43%	TTNET Turk Telekomunikasyon Anonim Sirketi, TR
9	200.33.146.0	8151	0.46%	26.89%	Uninet S.A. de C.V., MX
10	217.237.150.0	3320	0.43%	27.32%	DTAG Deutsche Telekom AG, DE

Table 3 – The Top 10 DNS resolver systems (by provider)

Those first two lines of this table are significant. ChinaNet appears to operate with a 50% market share within China, which is equivalent to serving the DNS queries of some 270 million users, or 11% of the entire Internet user population. The overall majority of the DNS queries from this group of users are passed through ChinaNet's diverse set of DNS resolvers.

Google's Public DNS system is different, in that this is not a DNS resolver that is sitting within the service architecture of a particular ISP. It's a distributed service that anyone can opt-in and use. And many Internet users have done precisely that. In terms of the population of Internet users who use Google's Public DNS, it's certainly is a very significant actor in today's DNS resolution environment. Across all of the anycast instances of the Google Public DNS service, some 10.5% of all users, or an estimated 270 million users, have their DNS queries passed through Google's DNS service for name resolution.

These 10 largest DNS visible resolver systems account for the DNS queries of almost 30% of the Internet's user population. What about the remaining 70% of the Internet's users? What's the overall distribution of resolver use? Figure 1 shows the cumulative distribution of users to these visible resolvers.

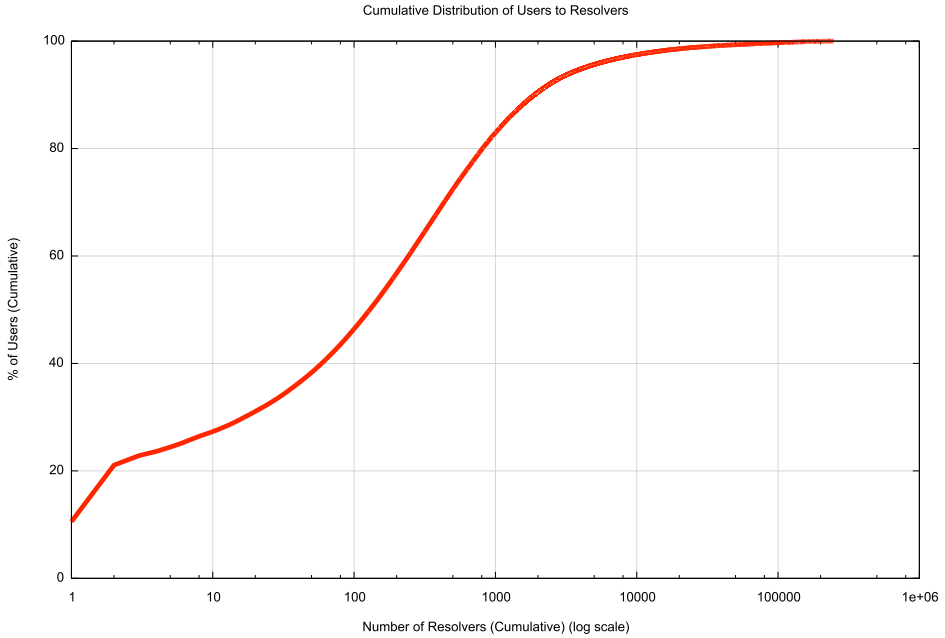

Figure 1: Cumulative Distribution of Users to Resolvers

Figure 1 indicates that the 1,000 largest visible resolver systems undertake queries for 83% of users. The largest 1,889 visible resolver systems, or 0.7% of the total count of resolver systems, handle the query load for 90% of the end users. Obviously, the DNS resolution "market" is a highly concentrated environment, where a relatively small number of resolvers handle the query load for a very significant proportion of users.

This observation of DNS resolver concentration has a number of implications. If most of the queries that are passed towards authoritative name servers come from intensively used resolvers, then the local caches in these resolvers would normally absorb the majority of queries from users. Authoritative name servers therefore would see a somewhat skewed set of queries that predominately come from the trailing part of the resolver set that serve very small user populations. The larger resolvers will cache the authoritative name server's response and serve a far larger population of users from the cached response. In other words, most authoritative name servers do not "see" most users' DNS behavior, as this is occluded by these large resolvers

It also implies that any orchestrated effort to introduce changes in this small set of DNS resolvers will affect a significant proportion of the end user population. This has been evident in the DNS with the introduction of support for introduced extensions to DNS, EDNS0, and with the introduction of validation of DNSSEC-signed responses, where some 12% of users exclusively use DNS resolvers that perform DNSSEC validation.

With the exceptions of Google's Public DNS and OpenDNS, the remaining 8 DNS resolver systems listed in Table 3 appear to be ISP-provided resolvers that service the ISP's user base, and the population of clients served by these resolvers appear to largely match the estimates of the size of the client base served by the ISP. That observation would lead one to assume that for most Internet users the DNS resolver they use is provided by their ISP. In other words, it looks like it would be a reasonable assumption that the resolver used to pass a query on to an authoritative name service is located "close" to the end user who initiated the query, at least at the level of granularity of the ISP itself.

To what extent is this assumption of locality of users and resolvers actually the case? How "local" are the resolvers we use?

Non-Local Resolution

While 8 of the 10 largest resolver systems in the Internet appear to be ISP-based systems that service the clients of that ISP, this is not always the case when one looks at the larger picture of all end users, their users and the DNS resolvers that they use. It appears that there is a significant fraction of the level of DNS resolution that is non-local, for some definition of "non-local".

Of the 84 million individual tests that were conducted as part of this experiment, it was observed that 34% of the time, or some 29 million tests, we observed DNS queries at the authoritative name server from resolvers that were located in a different network (by ASN) from that of the end user who initiated the query in the first place. This is a surprisingly high proportion.

But it's useful to recall the entry in Table 3, where Google's Public DNS service is used by around 10% of the world's user base. If we remove this single non-local factor of use of Google's Public DNS service from the picture then we are left with some 20 million users, or around 1 on 4 users who use DNS resolvers from a different ASN from that where they are located.

However having your DNS resolver in a different ASN is not necessarily a distant relationship. It could be that the resolver is located in an "upstream" service provider and the user and the DNS resolver are not necessarily located far from each other. What is perhaps more notable is the case where the DNS resolver is located in an entirely different country. The numbers of these cases are still significant. Some 21% of users send their queries via DNS resolvers located in a foreign country. Even when we remove Google's Public DNS from the picture the numbers are still significant, with some 11% of users, or 9 million in the case of our experiment, using resolvers located in different countries.

Where are these users? Are there countries where the level of use of "foreign" DNS resolvers is significantly higher than others? There is certainly a very high variation, with 99% of users who were geo-located in Algeria using foreign DNS resolvers (half of this count in Algeria is due to use of Google's Public DNS, which the other half is dispersed across other countries). At the other end of the spectrum 97.7% of users in South Korea use DNS resolvers located in South Korea. A map of the world showing the proportion of "foreign" DNS resolution is shown below.

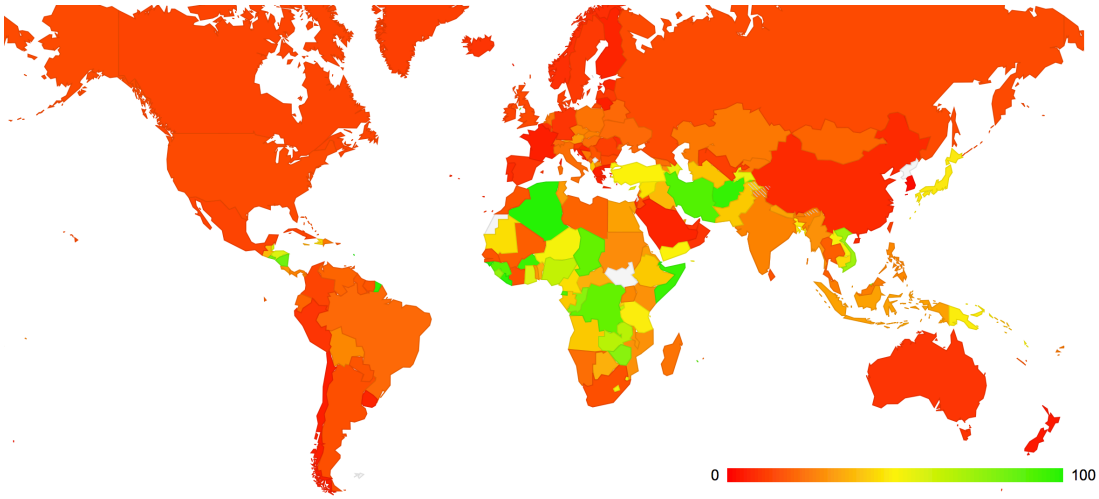

Figure 2: Relative use of "Foreign" DNS Resolvers by Country

It's evident that the highest levels of "foreign" DNS resolvers occur in Africa, parts of the Middle East, South East Asia and Central America. The lowest levels are found in South Korea, New Zealand, France, Greece and Finland.

However, Google's public DNS is a big factor here. A second map (Figure 3) removes the Google Public DNS component from this foreign use map, and looks at the proportion of "foreign" DNS resolution using resolvers other than Google's Public DNS.

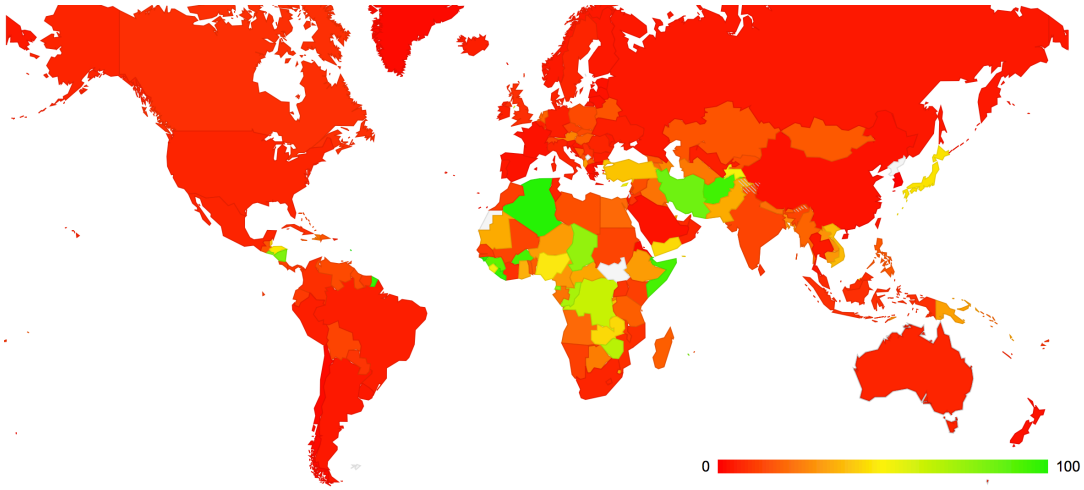

Figure 3: Relative use of "Foreign" DNS Resolvers by Country

Since so many users end up using Google's Public DNS service to query authoritative name servers, it's worth looking at the use of this service its own right. Figure 4 is the map of the world looking at the proportion of users in each country who pass their queries to Google's Public DNS service.

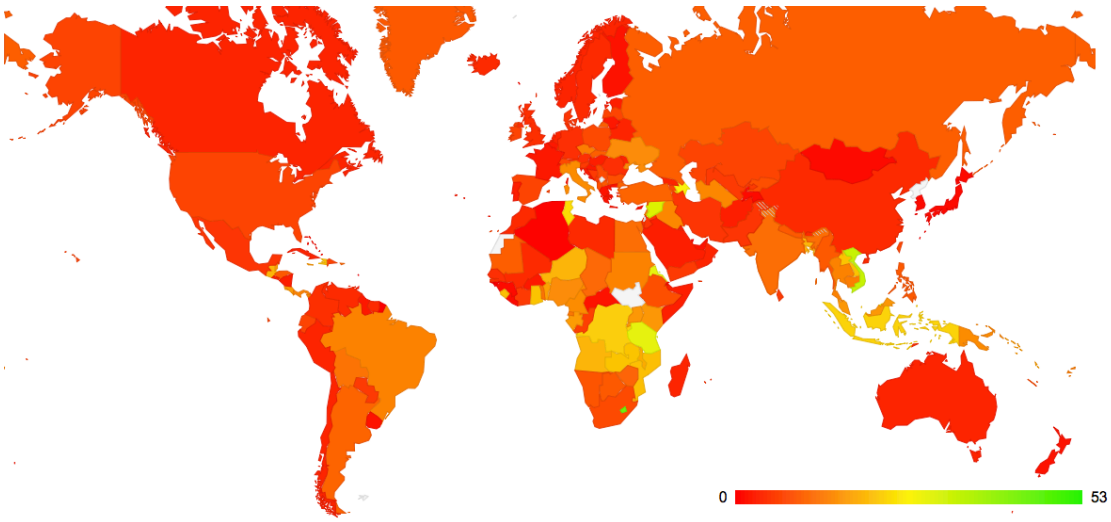

Figure 4: Relative use of Google Public DNS Resolvers by Country

Its clear that Google's Public DNS is widely used in central Africa, parts of the Middle East and across South East Asia. Perhaps not so obvious is 10% of Russian users, 14% of Brazilian users, 11% of Indian users, and 5% of users in both China and Great Britain also direct their DNS queries to Google's DNS service. (The highest relative level of Google's DNS usage was seen from the island nation of Kiribati, with 53% of the nation' users. The largest estimated pool of users comes from China, where an estimated 25 million users direct their queries to Google's PDNS.)

What about the reverse view? If 11% of the world's users direct their DNS queries to a foreign resolver other than Google, then where are these DNS resolvers located?

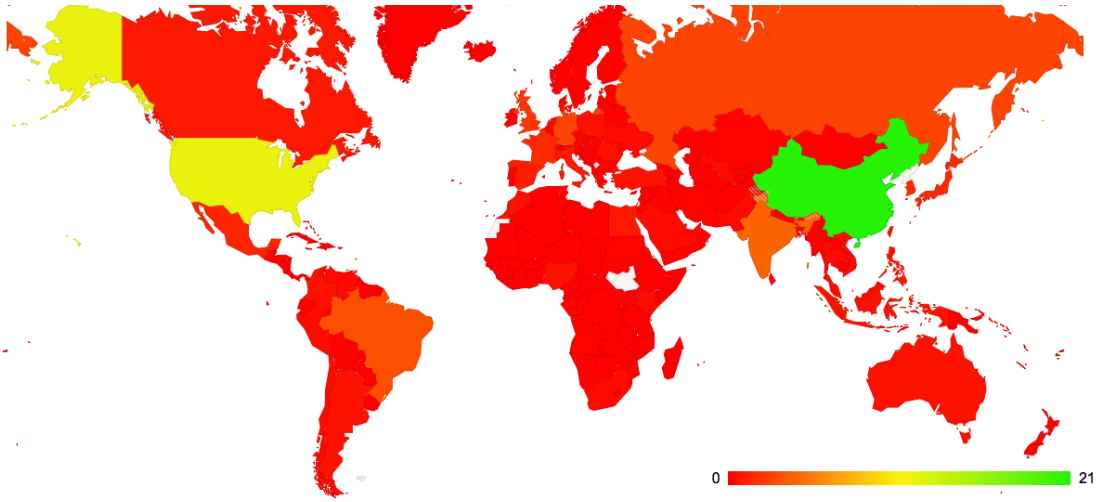

Figure 5: Location of "Foreign" DNS resolvers

Why?

There are a myriad of possible reasons why one in every 10 users of the Internet have their DNS resolvers located in another country to the one that they are located in.

The first potential reason is the inaccuracy of public geo-location tools. This exercise has used geo-location data published by Maxmind (https://www.maxmind.com/en/home), and augmented with the Regional Internet Registry data in those cases where Maxmind has not identified a country for the IP address. It's great that Maxmind is willing to publish a geolocation data set, and kudos to them for doing so, but at the same time we have to recognize that geo-locating addresses is not easy, and the published database is not perfect by any stretch of the imagination. In other words there is a distinct possibility that the geolocation of mapping an IP address to a particular country may well be incorrect.

The second potential reason is the widespread use of various hybrid forms of tunneling. Relocating a device's DNS resolvers to a different country often bypasses simple geo-blocking systems for restricted content. So its not surprising to see the US be the target of a large proportion of DNS relocation as one measure to provide the appearance of a US-located client of a content service.

Its also possible that we are seeing the outcomes of various forms of malware, where the infected system has its DNS resolution queries redirected to a compromised DNS resolver that will answer a critical subset of queries with some form of synthetic response, while answering all other queries using conventional DNS resolution queries.

But why the preponderance of use of Chinese resolvers? Some 20% of all foreign resolution cases involve the use of DNS resolvers located in China. One explanation could be in the adoption of off-shore web proxies in order to circumvent some form of content firewall or filter operation. In this case the DNS would be correctly locating the user into China, but the end-user's use of a foreign web proxy would mean that the web object retrieval would appear to originate from a different country. Other possible explanations involve the use of VPN tunneling and local DNS deep packet inspection and response spoofing, although such explanations are highly speculative.

The use of non-local DNS resolvers is so common in today's Internet, with 1 in 3 users using DNS resolvers located in a different network than the one in which they are located., and 1 in 5 users using DNS resolvers that are located in a different country than the apparent location of the user. Given such widespread use of non-local DNS resolvers its reasonable to speculate that all of the above explanations may be at play here, and doubtless a few more as well.

Reports

Reports from this experiment are available online at http://labs.apnic.net/cgi-bin/rmap.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.