February 2022
Geoff Huston

DNS4EU

The last few decades have not been a story of unqualified success for European technology enterprises. The European industrial giants of the old telephone world, such as the former stalwarts Alcatel, Siemens, Philips, Ericsson and Nokia, have found it to be extraordinarily difficult to translate their former dominant positions in the telco world into the Internet world. To be brutally frank, none of the current generation of major players in the digital environment are European. It looks like most semiconductor chip fabrication happens now happens in Taiwan, Korea, the US, China, and Japan. The supply chains for smart devices are even more restricted and they appear on the whole to be designed in the US and manufactured in China. Application and service innovation seems to be an activity that is dominated by US enterprises. European innovations, and there have been many important innovations within the Internet environment, such as the Web at CERN, or Skype from Estonia, has not directly led to the emergence of European enterprises with global reach. Many of these innovations have turned to the US venture capital markets to develop their ideas, and this has resulted in their further development and commercial exploitation in the context of the US business sector in so many cases.

Yes, this is a gross simplification of a more complex picture of the global technology landscape, and European enterprises probably contribute as much into the global technology space as the US, China or India. Our collective need for skilled and innovative contributions to the collective effort transcends the capacity of any nation or any single region, and the net contribution from the European sector is as significant as any other. However, it must be observed that Amazon, Apple, Alphabet, Meta and Microsoft are all US companies, and the current top ten largest publicly traded companies, as measured by market capitalisation, has eight US enterprises and one each from Taiwan and China (https://en.wikipedia.org/wiki/List_of_public_corporations_by_market_capitalization). Yes, some of these multinational enterprises may have taken advantage of Ireland as a relative form of corporate tax haven from time to time, but that was about it. The last time European-domiciled enterprises were included in this top ten list was in late 2015, when the Swiss Pharmaceutical corporation, Novartis, and the Swiss food and drink enterprise Nestlé were both listed in this top 10. The most recent time a European telecommunications and technology sector enterprise was included in this list was more than 20 years ago, when Vodaphone and Deutsche Telekom had brief periods of being listed. Whatever is going on here, it looks as if European enterprises are finding it hard to remain domiciled in Europe and keep up with their international competitors, particularly in today's technology sector.

The concern is that if today's technology world equates to the previous world of far-flung colonial empires, or that of the world of the industrial revolution, then relative national wealth and prosperity appear to be linked to the ability to master, or preferably dominate, critical aspects of the sector. And in this respect Europe appears to have been left behind. It still feels to many Europeans as if Europe is over on the exploited side of the techno-colonial landscape, rather than being one of the exploiters. And no doubt that prospect is a particularly concerning one to EU political leaders and within the EU bureaucracy. What should or could the EU do to avoid further decline in this area?

Before looking at the EU response to the questions posed by this situation, there is probably more to this than just keeping up with international competitors and maintaining a visible position in the set of leading enterprises. As tough as this sounds, I'm not sure that this issue of decline in perceived

importance of role of European digital enterprises in the global technology sector is the full story. It's more than this. It's also the combination of the increasing level of reliance on the goods and services produced by this sector and the source of these digital goods and services. The past twenty years has seen the progression of many of society's activities onto the Internet's ubiquitous digital platform. These days all forms of retail banking, shopping, and entertainment are all largely Internet-based. However, it's deeper and more pervasive than these simple examples might infer, as we find out from time to time when things break. From oil pipelines in the US, to critical infrastructure systems such as electricity distribution, we all now use various forms of digital cloud command and control structures within the framework of a common Internet. Few services now operate in a manner that is completely independent from the Internet, and perhaps more significantly, most services are critically reliant on the Internet. This reliance question can be re-cast with in more nationalistic tone. For any national society, to what extent is that national economy critically reliant on the continued access to digital services provided by entities who are domiciled in foreign jurisdictions, and even delivered across national borders in a completely seamless fashion?

We can add to this picture of international dependence the perils of cyber-hostility. How can a national or regional community defend itself from digital attack, be it attacks on the provision of the service or access to it by the users? This topic raises a whole set of uncomfortable questions about the level of interdependence within the digital landscape and the vulnerabilities presented by this. To what extent is the resilience of a national digital infrastructure reliant on services provided by foreign entities? And when this interdependence is abused in a hostile context then how can nations respond? Unlike the national responses to the ongoing COVID–19 pandemic we can't simply seal up all movement across the border! At best our current actions are looking to mitigate, to some very small extent, this level of foreign dependence in our digital infrastructure. We saw this thinking exposed in various countries with the construction of national 5G mobile infrastructure, where a number of countries have taken steps to exclude various Chinese enterprises from central roles in these projects. We saw this again in 2018 with the efforts in Russia to construct a DNS infrastructure within Russia that could operate only on domestically controlled infrastructure.

As uncomfortable as this interdependence may be, doing something about it in more meaningful ways can be very challenging. For many national communities the issue is simply one of relative size: many nations may have already adopted the position that they are too small to take on today's digital behemoths and declare independence and self-sufficiency (in the sense of eliminating their dependence on them). The in-country data retention measures seem like a relatively poor second choice substitute to address such fundamental concerns. No matter how uncomfortable it may be to observe that national communities are now critically dependent on these digital giants, they also have been forced to acknowledge that it is just not feasible to contemplate alternatives that have domestic ownership and control. Other national communities are not so willing to embrace a future that includes such critical dependencies on services provided by foreign enterprises at a fundamental level. I would suppose that they feel that they are large enough to take on these enterprises and use their own resources to decrease this level of foreign dependence for critical services. And it is in this situation that the EU community finds itself today.

I should hasten to add at this point that this situation is not the outcome of any chosen strategy on the part of today's digital giants. In designing an Internet architecture that was based on stateless packet forwarding and eschewing the traditional control points of network state as was used in the circuit switched telephone network not only did we get a new system that could scale its infrastructure and services to the size of today's Internet, but we also built a network and a service platform that paid no heed to concepts such as national or regional political boundaries, network control points, bilateral infrastructure and traffic agreements, transaction-based accounting practices and various forms of international financial and regulatory agreements was inevitable. The internet was not constructed as an amalgam of various national networks but was conceived and constructed as a single artefact that had never integrated such geopolitical concepts into its internal architecture. The result was somewhat inevitable in that a large enterprise in this environment could reach across the entire span of the network without any technical requirement to negotiate national boundaries. In retrospect, where we find

ourselves today, as discomforting as it is to many, is a natural consequence of the technology choices made in the basic architecture of the packet-switched Internet.

We can take this macro view of national and regional interests and the modes of participation in the technology of the digital environment and apply it at a finer level of granularity to individual activities within this sector. What I want to look at here is the very particular issue of the Domain Name System (DNS) and the market for name resolution and the European perspective.

The DNS really is Everything!

This choice of the DNS here is not a random choice. The Internet's name system is an important topic of conversation in today's Internet, as it appears that the DNS is the only remaining part that is left of the "glue" that hold the Internet together and is now the defining medium of what is "the Internet". IP addresses, the other part of the Internet's original common infrastructure, appear to have become a for more amorphous and fractured concept. We've passed all the heavy lifting of service identification and rendezvous over to the name system, and passed the issue of endpoint identification over to the applications and service environment, that in turn rely on the underlying name system.

This central role of the DNS is reflected in the way we use the DNS and related services:

- Our concerns with privacy and trustworthiness are reflected in our efforts to improve the privacy and integrity of DNS resolution transactions.

- Our collective obsession with faster performance of digital service is reflected in our efforts to improve the speed of DNS transactions through the use of ever-larger multi-headed anycast server clouds and continual tuning of the protocol and servers to shave delays out of processing transactions.

- Service rendezvous is a role that increasingly is being undertaken by the DNS, such as in the SVCB and HTTPS resource records. Instead of asking the DNS for the IP address associated with a DNS name we can now ask the DNS to inform the client of where to connect, what port to use, what encryption protocol is needed and even details of the public key information to support this encrypted channel.

- Content filtering is a role executed by filtering in DNS resolvers. If the DNS does not resolve a name, then that name and the associated service simply does not exist for any practical purpose.

Because of the role of the DNS as an essential facilitator in every network transaction the DNS really is the most critical component of the Internet's infrastructure these days.

The DNS Resolver Landscape

In the early days of the Internet when mainframe computers were the only thing around, the name system was a far more rudimentary service. Every host had a local copy of a simple text file, *hosts.txt* and applications who wanted to translate a name to an IP address to use on IP packets consulted this file for a matching entry. If you look hard on the platform you are using to read this, you will probably still find a remnant of this *hosts.txt* file. The task at the time was to coordinate all these independent versions of this file so that the same name was recorded with the same address on all the Internet's hosts. As the Internet grew, this task became harder. The first step was to augment this local host file with a lookup into a shared distributed database. If the name was not defined in the local host file, then the platform would pass a query to the local implementation of a DNS database front end, which would then perform a directed query through the distributed database.

The problem is that this database query could be extremely slow, as the local agent first must find which database server holds the authoritative information for the name being queries, and then pose the query that server. The design response to increase the efficiency and speed of the DNS was to use local caches.

The name-to-address binding changed infrequently, so once a local implementation learned of a binding of a name to a service address it could store this in a local cache and reuse it for subsequent queries without further consultation into the database. When the caches ran "hot" the performance of this database query was as quick as a local hosts file, but with far better consistency of the overall resolution of names.

We distinguished between the DNS servers that handled queries for applications running in end hosts, so called *stub resolvers* at the edge of the network, and *recursive resolvers,* which are DNS servers that assist a collection of stub resolvers by acting as their agent and performing the distributed database queries for them. Not only did this offload a set of database navigation functions from the stub resolver to the recursive resolver, but it allowed these recursive resolver middle-agents to cache the answers for a larger collection of stub resolver clients, further increasing the effectiveness of caching in the DNS.

For many decades these resolvers were integrated into the Internet's service landscape by assigning the role of operating these recursive resolvers to the local Internet Service Provider (ISP). The ISP not only provided its clients with access to the Internet, and IP addresses for these clients to use, but also provided access to the common name system through the provision of a DNS recursive resolver service for its clients. This was a relatively stable arrangement for many years, but at the same time there was a lot of churn lurking just below the seemingly placid surface of the DNS. It became increasingly apparent that operators of these recursive resolvers were privy to large volumes of useful and timely information about user behaviour, and in an Internet economy that is increasingly defined by surveillance capitalism this is extremely valuable information. It was also apparent that operators of these resolvers were in a unique position to control the visible content that was accessible for their users.

This was an enticing temptation for some ISPs. In this era of the Internet's surveillance-based economics, a real-time stream of data about what users are doing has considerable market value, and the DNS resolvers' query logs had considerable value, despite the somewhat disturbing privacy issues. Given that the ISP was unable to convert the costs of operating its recursive resolver service into a revenue stream by charging its user base, and the ISP business has been squeezing its margins for many years, any additional revenue stream must be an interesting proposition. There is also the possibility of monetising the DNS service by performing NXDOMAIN substitution. Here, instead of responding that the name does not exist, the ISP can instead respond with a sponsored referral to a search engine.

It's not just ISPs who are exposed to the temptation to play fast and loose in the DNS. The DNS has become fodder for various national regimes to both observe their citizens and to impose controls on their online activities. These days it is common for governments to proscribe the resolution of certain DNS name, and phrase this as a legal obligation for ISPs and other domestic service providers. The motives for these blocking lists are varied, and include attempting to curtail the propagation of malware, disrupt the command-and-control channels of co-opted zombie attack bot armies, censor offensive content and protect rights holders from efforts to infringe their intellectual property rights.

This latter aspect of DNS censorship and the EU is already an active topic.

IPR interests associated with Sony Music Germany bought a suit against the open DNS resolver provider Quad9 in a German court. The court ruled that Quad9 must block resolution of a domain name of a website in the Ukraine that itself does not hold copyright infringing material, but instead contains pointers to another web site that is reported to hold alleged copyright infringements.

Quad9's interpretation of this ruling is that queries received from IP addresses that can be geolocated to Germany must generate a SERVFAIL response from Quad9's recursive resolvers.

There are a number of curious aspects of this situation. It appears that the other significant open DNS resolver providers (Google, Cloudflare, and Cisco's OpenDNS) have not been similarly targeted by legal action in Germany by Sony. Perhaps the Swiss domicile of Quad9 made Quad9 a more appealing target for German legal action. Or perhaps there are some involved issues in attempting to compel a non-EU provider to take certain actions with respect to blocking content. Open DNS providers do not sell their services in a conventional manner. There are no paying clients. No contracts. Nothing. Clients of these service make their own decision to use these open DNS services and do so without any form of payment and without any formal commitment. Possibly in terms of enforcement mechanisms this becomes an issue for the individual clients of this service and not necessarily an issue for the non-EU DNS resolver service operators.

See "The Curious Court Case of Quad9" from my recent write up of the 2021 ICANN DNS Resolver Symposium (https://www.potaroo.net/ispcol/2021-12/dns-sym.html).

Obviously, these developments in co-opting the DNS for such purposes has not gone unnoticed. Some clients, both consumer and enterprise clients, may feel that the DNS filtering being performed by their ISP is unwarranted. Clients may also be uncomfortable with their ISP being capable of performing detailed surveillance of their activities through the DNS. No matter what the level of assurance that their information is held in a way that preserves their privacy, there is a lingering doubt that this this is really the case, particularly when duly executed warrants are served on the service provider.

One potential answer for such clients is to operate a recursive resolver completely within the client network. That measure can circumvent any DNS filtering that is being performed by the ISP's recursive resolver, and the measure also stops providing a direct feed of client activities to the ISP's recursive resolver. However, that is also an additional role that the client has to perform, and the open unencrypted nature of the DNS makes any and all traffic from these locally operated recursive resolvers easy to detect, inspect and potentially manipulate in any case. It seems to require a higher level of expertise on the part of the client with little in the way of net benefit to the client in terms of privacy and integrity protection.

The Open Resolver model is an alternative here. The idea is that the open resolver may not be operating in the same regulatory or legal framework as the client and the client's ISP and may be able to resolve DNS names that would otherwise be proscribed. The Open Resolver may be in a different legal regime and may not necessarily be subject to domestic law enforcement processes of discovery of DNS queries. Again, the consideration of the open unencrypted nature of the DNS means that this does not substantially change the net privacy benefit to the user here, but in this case there is no effort on the part of the user to run local DNS services.

In December 2009 Google entered this space with its public resolver offering, on 8.8.8.8. Google's reasons for entering this market were couched in terms of better performance and better security in the handling of queries (https://developers.google.com/speed/public-dns/docs/intro?csw=1). However, it also should be observed that Google had a strong commercial motive to enter this space. Their major commercial asset is their search engine. If the DNS lookup could be transformed into a search engine, then this would represent a direct threat to their business, and in performing NXDOMAIN substitution this was exactly what some ISPs were doing. If the ISPs were performing this pseudo-search in the DNS as a revenue raising activity, then Google's DNS resolver represented an alternative that did not attempt to raise revenue from the ISP-operated DNS but eliminated the need for the ISP to operate any general DNS resolver infrastructure for its clients. All it needed to do was to forward all client queries to Google's service. From Google's perspective I would guess that this open resolver project represented a relatively small outlay to further protect their core business asset.

Open Resolvers represent a major shift in the DNS landscape, and Google plays a major role these days. Figure 1 shows the "market share" of the three largest open DNS resolvers, as a day-by-day time series since July 2019, based on measurements conducted by APNIC Labs (https://www.potaroo.net/ispcol/2019-09/centrality.html).

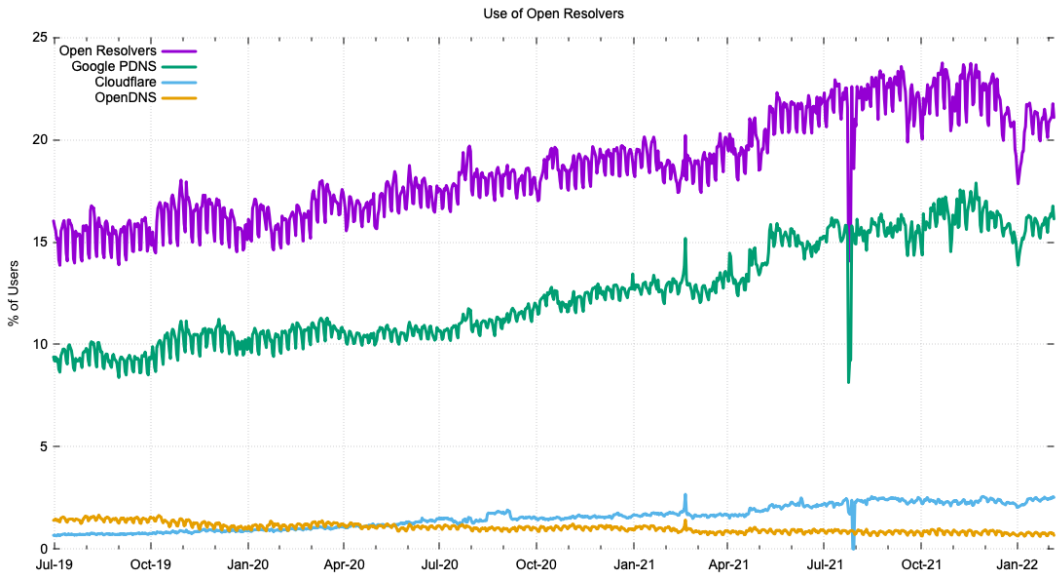

Figure 1 – Market Share of DNS Open Resolvers July 2019 – February 2022

Some 20% of the Internet's user population use an open resolver to resolve names, which is an unexpectedly high number. Of these open resolvers Google has the major share with its public resolver offering, and these days one in six (16%) of the world's users use Google's service. Cloudflare's 1.1.1.1 service is used by 2.5% of the world's users and OpenDNS has a 1% market share in this space.

It is also worth noting that the open resolver metrics have a visible weekday / weekend variance. The use of open resolvers is higher on weekdays, pointing to a likely preference for enterprise customers to eschew the ISP's DNS offering and prefer to use an open resolver service instead.

Now let's turn our attention to the EU and see if the same situation holds there.

Just how significant is this movement to use Open DNS resolvers in EU member states? Table 1 compares the data on use of public DNS resolvers in January between the internet-wide totals and the total in the EU.

January 2022	All	EU
Samples	455,721,405	41,635,975
Same AS (ISP)	67.38%	76.96%
Total Open Resolvers	20.44%	15.84%
Google 8.8.8.8	15.56%	12.65%
Cloudflare 1.1.1.1	2.35%	2.89%
OpenDNS	0.74%	0.65%
Quad9 9.9.9.9	0.14%	0.06%

Table 1 – Use of Open Resolvers in the EU, January 2022

The use of open DNS resolvers in the EU is slightly less than the internet-wide average. Google's service is 3% less common, and Cloudflare's service is slightly more (0.5%) common in the EU. In the use of open DNS resolvers the EU profile is not that far off the general profile.

Figure 1 also shows a steady growth in the proportion of users who have their queries passed to open DNS resolvers over the past 30 months. What is the trend data for the EU?

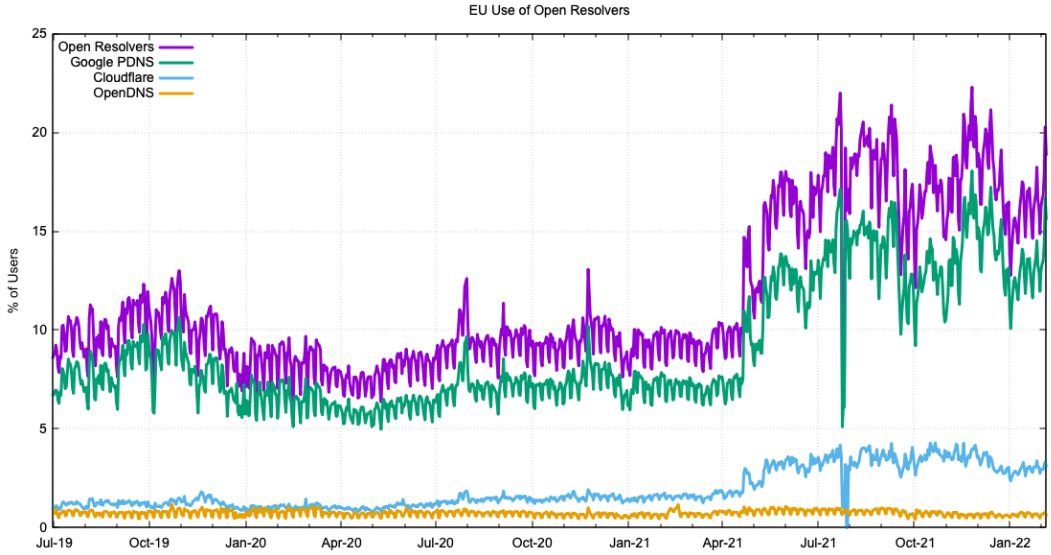

Figure 2 – Market Share of DNS Open Resolvers in the EU July 2019 – February 2022

As shown in Figure 2 the use of open resolvers has been growing over the past 30 months (the discontinuity in April 2021 is an artefact of the measurement system). The use level has almost doubled from mid 2019 to early 2022, which is a higher relative growth rate than the overall Internet-wide numbers.

We've been looking at the EU as a uniform collection of nations. To what extent do these member states differ amongst themselves? Table 2 shows this comparison.

CC	Name	Samples	Same AS	All Open Resolvers	Google	Cloudflare	OpenDNS
AT	Austria	950,035	74.0	8.8	6.3	2.0	0.4
BE	Belgium	1,192,973	94.6	4.1	2.9	0.8	0.3
BG	Bulgaria	549,103	67.3	16.2	13.8	2.4	0.3
CY	Cyprus	121,955	52.6	35.8	9.4	1.9	24.6
CZ	Czechia	1,000,906	80.1	15.9	12.8	2.6	0.6
DE	Germany	8,112,657	72.6	26.0	20.8	5.4	0.5
DK	Denmark	648,425	77.1	16.1	10.7	4.2	1.5
EE	Estonia	135,535	94.0	5.3	3.8	1.3	0.2
ES	Spain	4,820,556	79.3	13.3	11.9	1.9	0.0
FI	Finland	559,249	88.8	10.9	9.3	1.2	0.3
FR	France	6,310,803	69.5	21.0	16.0	4.7	0.4
GR	Greece	860,678	86.1	6.5	4.9	1.5	0.2
HR	Croatia	295,220	73.6	7.4	4.6	1.3	1.6
HU	Hungary	860,014	89.7	6.7	5.5	0.6	0.5
IE	Ireland	481,699	79.7	17.2	14.9	1.7	0.4
IT	Italy	4,370,226	90.7	7.9	6.3	0.6	1.0
LT	Lithuania	257,982	89.1	9.0	6.6	2.1	0.3
LU	Luxembourg	70,770	58.4	40.9	29.7	2.8	8.6
LV	Latvia	180,075	84.3	10.1	8.8	1.0	0.3
MT	Malta	42,435	34.0	33.9	9.8	0.9	23.1
NL	Netherlands	1,877,122	50.4	26.0	22.2	3.9	0.6
PL	Poland	3,547,175	72.2	12.5	10.5	1.5	0.6
PT	Portugal	916,876	88.2	5.8	4.6	0.7	0.5
RO	Romania	1,548,032	89.8	5.5	4.5	0.7	0.4
SE	Sweden	1,194,018	75.1	8.5	6.0	2.0	0.5
SI	Slovenia	199,861	94.0	5.7	4.4	0.7	0.5
SK	Slovakia	528,112	82.9	14.1	9.7	2.6	1.6
EU	**EU Total**	41,632,502	76.9	15.8	12.6	2.9	0.6
XA	**World**	455,721,600	67.4	20.4	15.6	2.3	0.7

Table 2 – Use of Open Resolvers in EU member states for January 2022

There is a strong preference to use the ISP's provided DNS resolution service in Belgium, Estonia, Italy, and Slovenia, where more than 90% of the samples show that the local resolver is being used. Google's open DNS resolver is used in more than 20% of cases in Germany, Luxembourg, and the Netherlands. Cloudflare's open DNS service is used by more than 4% of users in Germany, Denmark, and France. OpenDNS is used extensively in Cyprus and Malta. It is not entirely clear if this outcome is the result of various DNS configurations performed by ISPs, by enterprise clients or by individual retail consumers, although there is a noted preference on the part of individual consumers not to alter the default configurations in their devices, so the outcome may well the result of ISP preferences and enterprise network configurations (such as AS12709, MelitaCable in Malta, preferring OpenDNS, and AS6866, CYTA-Network in Cyprus, also preferring to send client queries to the OpenDNS service)

Is the observation that some 16% of users in the EU have their DNS queries passed to open DNS resolvers a significant issue for the EU, or is it a number that really warrants no particular concern? Yes, it's a big number, and it is getting bigger over time. On the other hand, it's a smaller proportion than the world average. It also should be noted that Google have been clear in maintaining that their resolver service is a precise and accurate representation of the DNS. Nothing is omitted, added, or altered in responses from their recursive resolver. Google does not disclose data about the way its resolver is used other than what is required under various national jurisdictions. Google report some information on the requests for data in a Transparency Report (https://transparencyreport.google.com/user-data/overview).

The reporting for the "Same AS" resolver could be misleading to some extent. Even within the ISP industry the DNS function has been the subject of outsourcing, and Nominum became a major player in this service market. In 2017 Nominum was sold to Akamai, which means that today Akamai, is now a significant service provider to ISPs for DNS resolution. What this means is that the true extent to which the DNS has been outsourced to a small number of service providers, and the pace at which the DNS as a market is consolidating, is not only evident in the use of Open DNS resolvers, but also lies in the choices in outsourced DNS service provision made by ISPs. This latter behaviour is not readily measured by conventional DNS measurement techniques.

DNS4EU

DNS4EU is the name of a European Union initiative intended to exert more control over the DNS within Europe, aimed specifically at the current level of use of open resolvers in the EU. As Andrew Campling reported in January 2022, "The European Commission announced its intention to support the development of a new European DNS resolver ("DNS4EU") in December 2020. It has been in dialogue since then to refine its thinking, in particular placing much greater emphasis on the potential cybersecurity benefits that could arise from the deployment of the resolver." (https://419.consulting/encrypted-dns/f/dns4eu-update)

This program aims to address the consolidation of DNS resolution in the hands of few companies, which renders the resolution process itself vulnerable in case of significant events affecting one major provider, or at least that's the rationale provided in the EU documents. It appears that DNS4EU will provide EU funding to support part of the capital costs for EU enterprises to construct DNS resolver services in the EU.

The intended benefit is to provide a DNS resolution service that is able to comply with the various content regulations in the EU by blocking the resolution of certain DNS names. It is unclear in my reading of the proposals how the DNS query data is to be handled, and whether such financially supported DNS resolver services would be obligated to share the DNS query data with various EU law enforcement authorities and security agencies, although the reference to potential cybersecurity benefits tend to suggest that some form of data sharing is being contemplated.

Related DNS4EU material suggests an expectation of a "better" DNS resolver service, although given that many of the benchmarks of what constitutes a "best practice" DNS resolver seem to be based on measurements of Cloudflare's and Google's resolver services. Presumably then the interpretation of

"better" relates to the level of service provided by ISP-operated DNS service, but the implication that EU money would be used to provide competition in the DNS resolution service market by somehow highly directing funding to existing ISP-operated DNS resolvers seems to redefine the role of public funding in potentially anomalous ways.

Perhaps the EU folk have been looking at CIRA's Canadian Shield DNS resolver (https://www.cira.ca/cybersecurity-services/canadian-shield) where the .CA registry has launched an open DNS resolver service. The service appears to be fully funded by CIRA, and, like Quad9's service, appears to use active DNS filters that are informed by malware and threat feeds and conforms to Canadian policies. It's useful to note that CIRA is not a government body, but, like many other CC TLD registries is a private, not-for-profit, member-based organization that administers the .CA top level domain.

There is another interesting example with the .CZ registry, CZ.NIC, who have funded the development of the KNOT resolver (and server). One of the earlier concerns with the DNS infrastructure was the lack of diversity of implementations of the protocol standards. Most resolvers and servers ran the BIND software. There was a deliberate effort to increase the diversity of DNS implementations, and these days three of the major DNS implementations, NLNet's Unbound, CZ.NIC's KNOT and PowerDNS are all outcomes of European projects. Much of the DNS infrastructure runs on these implementations today. Not only does this provided much-needed diversity in DNS implementations to reduce the monoculture-related vulnerabilities, but it helps in increasing the level of subject-matter skills with DNS services within the EU.

In some ways the DNS4EU program is not all that different from these efforts, particularly with respect to the CIRA initiative. If you are unhappy with the collection of open resolver services and believe that you can do a better job, then perhaps the best option is to transform this sense of unease and discomfort into action and run your own. However, if the party wanting to prove that it can do a better job is the public sector itself, then this raises some quite predictable issues relating to public sector involvement in private sector activities. One of these issues is that of treading carefully, lest you scare away all private capital and leave the public-funded service as the last one standing in a supposedly deregulated private sector-led activity. Why would a private enterprise continue to invest in a service sector when it is competing on unequal terms with a public sector-operated service? How can a fair set of rules be enforced in the market when the rule-setting body is an active player as well?

What about ISPs? Why should they continue to spend their own money running a DNS resolution service for their clients when the EU is channelling funds to some third party to run an open DNS service? Why not just use a simple forwarder and pass all the ISP queries onto this same service? Is the level of funding from the EU to run this service truly at such an open-ended level where the successful bidder is in a position to build and operate a DNS resolution infrastructure that can cope with the demands posed by up to 500M users?

Now it could be argued that this is what Google are doing already, so there is an existence proof that this is not an infeasible ask. But Google is indeed special. Google is spending money and resources in defending its core business asset of search, and in running an open resolver that faithfully presents the contents of the DNS to its users it is helping to prevent the perversion of the DNS into a search engine. The issue here is that this is a relative unique motivation. Other DNS resolver operators do not share that motivation, given that they are not major players in the search space and have no existing business asset that they are attempting to defend. If a DNS resolver operator's operating resources are fixed, then the onset of larger query volumes results a degraded service, which tends to defeat the purpose of operating this service in the first place.

It is challenging to see how the DNS4EU program of partial-funding of the capital costs of setting up an open DNS resolution service and no operational funding would create a sustainable business model in the DNS resolution market that would have an impact on the market share of the existing open DNS resolver operators and the overall way in which DNS names are resolved in the EU.

The harsh truth here is that DNS resolution is a market failure, in that users don't pay to have their queries answered and information publishers don't pay recursive resolvers to have their answers served. The reason why ISPs run DNS resolvers is perhaps because this is what ISPs have always done. But DNS resolution is a cost centre for ISPs and there is no clear business motive to increase their investment in DNS infrastructure beyond the level of functional adequacy, particularly given that few, if any, users make their ISP selection on the basis of the quality of the ISP's DNS services.

So, on the one hand it's easy to understand that the situation the EU finds itself in, where significant parts of its internal digital infrastructure and being operated by foreign owned and controlled enterprises. It is not acceptable at a strategic level, and its entirely understandable that the EU would wish to change this picture of foreign dependence.

But having largely deregulated this industry and having dismantled many of the restrictions on international investment in digital services, the set of tools that are left to governments are at times somewhat inadequate, particularly when they contemplate forms of active intervention in the marketplace to redress what they perceive as strategic imbalance and vulnerability. The results of their various rule setting efforts can be judged as a mixed package that has both positive and negative outcomes. At worst, it could be judged as no more than placing a further brick in the wall of consolidation of the industry into the hands of the existing digital behemoths through imposing more overwhelming impediments in the path of emerging competitors. At best, its outcomes could be an expensive but merely palliative measure for EU users and member states.

So, what can the EU do? It seems that DNS4EU is an example of the line of thinking that if you can't throw rules at a problem, then try throwing money at it! Personally, I don't have any optimism that this approach will do any better than the previous rule-setting efforts. Creating a new set of enterprises based on dependence on government financial subsidies does not necessarily create a new set of competitors. More likely is the outcome that it merely creates a new set of dependants on the public purse!

Disclaimer

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

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net