Report of the Afnic Scientific Council

Consequences of DNS-based Internet filtering

Consequences of DNS-based Internet filtering[1]

In France and around the world, Internet filtering has been a central issue in several recent events[2] and has been addressed in numerous decrees or bills (ARJEL[3], LOPPSI2[4], SOPA[5], PIPA[6], etc.). These legislative or regulatory initiatives involve various areas, from protecting against viruses and malware, to protecting children against pedo-pornography, prohibiting access to unauthorized online gaming sites, or the fight against piracy.

Filtering, or blocking, can be done at IP address level as well as at domain name level. In this paper, the Afnic Scientific Council provides a situation report on the techniques that can be used when filtering is used at the naming level in the DNS (*Domain Name System)*, and the foreseeable consequences of such measures.

1. Filtering of domain names

Filtering consists in modifying the normal resolution of a name on a server to an IP address, either by blocking the response, or by replying with an error message, or by returning the address of another server generally indicating that access to the website in question is prohibited.

Unlike telephone networks, which rely on a dedicated infrastructure, the DNS service, like the other services that use it (Web, Mail, etc.), relies on the Internet itself to exchange data. It should be noted that the DNS arrived relatively late in the history of the Internet, when the names of the devices connected could no longer be stored in a single file[7].

[1] This text is based on excerpts from a reference document produced by the Council of European National Top-Level Domain Registries (CENTR), http://centr.org/system/files/agenda/attachment/centr-paper-blocking-20120302.pdf

[2] For information purposes only, see the following affairs : “Pirate Bay” (http://money.cnn.com/2012/01/20/technology/pirate_bay/index.htm), wikileaks (http://fr.wikipedia.org/wiki/WikiLeaks), copwatch (http://abonnes.lemonde.fr/societe/article/2011/10/14/la-justice-interdit-le-site-web-copwatch_1588162_3224.html)...

[3] http://www.arjel.fr/

[4] http://fr.wikipedia.org/wiki/Loi_d'orientation_et_de_programmation_pour_la_performance_de_la_s%C3%A9curit%C3%A9_int%C3%A9rieure

[5] http://en.wikipedia.org/wiki/Stop_Online_Piracy_Act

[6] http://en.wikipedia.org/wiki/PROTECT_IP_Act

[7] Historically, this information was stored in a hosts.txt file that is given priority over the DNS, even in the

The DNS works in a distributed manner, according to an architecture in which the reference information is delegated to a very large number of stakeholders. Registries such as Afnic refer to servers that perform the mapping between domain names and IP addresses. To make the system both more robust and more efficient, three mechanisms are used: structuring the information hierarchically, the use of multiple server replicas, and the caching of the responses obtained.
Information hierarchy is based on the structure of the domain name. For example, in order to resolve the name `www.wikipedia.fr`, a server, known as the root, is queried in order to refer to the name servers of the *.fr* zone, called a "Top-level" zone. One of the *.fr* servers in queried in turn in order to refer to the name servers of the `wikipedia.fr` zone (a zone delegated under the *.fr* zone). Finally, one of the name servers of `wikipedia.fr` in queried in turn. The latter then returns the IP address being sought (that of `www.wikipedia.fr`).

Resolvers are used to iteratively process the queries described above. They require a minimal configuration, in particular having the list of the root servers IP addresses, and can be used to cache the responses in order to avoid overloading the DNS architecture. Public ISPs implement resolvers to be used by their customers. The addresses of those resolvers are provided to customers alongwith the other configuration settings. But the customers can easily configure another resolver address either on their home gateway (aka "box"), or on their computer. The number of resolvers in the world is very large (around ten million) and is steadily growing[8]. Google[9] and OpenDNS currently provide this service. Users can even install their own resolvers on their computers.

Blocking the mapping between domain names and IP addresses can therefore be done in two places, either at the level of the registry by no longer publishing information which will therefore gradually disappear from caches, or at the level of the resolvers. On the latter, filtering can be done by Black Lists provided by governmental or judicial authorities. This system depends on the cooperation of ISPs and other resolver operators to maintain the list in the configuration of the resolvers. In technical terms, blocking a resolution is available in some software distributions such as BIND (one of the most used), but this feature is not standardized[10]. Note that we have very little hindsight with which to assess the impact of filtering at the country level.

most modern operating systems..

[8] ICANN estimates their number to be about 10 million worldwide (http://blog.icann.org/2012/03/ten-million-dns-resolvers-on-the-internet/)

[9] https://en.wikipedia.org/wiki/Google_Public_DNS

[10] Since 2011 this software contains the RPZ option (*Response Policy Zone*) allowing a server to lie during the resolution process, either by returning an error message or another value.

2. Filtering is relatively ineffective

As long as the content remains online[11], filtering a domain name does not completely fulfill its objective. Users can still access the required content, either by directly entering the IP address of the server, or by putting the name-to-address mapping in a file on their machine. The site may also create alternative names that are not part of the list of blocked sites[12].

If name resolution is filtered in a country, anonymous proxy https can be used to access (usually free of charge) content from another country. Search engines offer lists of proxies.

The use of VPN (virtual private networks) is more complex, but offers more stability and fewer risks than the anonymous proxy http mechanism. This is because VPNs allow users to appear to be connected from another location. They often require a fee-paying account, but it is easily conceivable that a commercial site such as an online casino will pay this in exchange for its customers' subscriptions.

Users can also configure their machine or their box to use an alternative DNS resolver such as those available from Google or OpenDNS. In this case, the filtering performed by the operator becomes ineffective, unless the DNS lookups are blocked as well[13]. In addition, for institutions using their own resolvers, blocking ISP resolvers has no effect.

3. Filtering increases the exposure of the user to threats

The basic principle behind the DNS is that users must obtain the information they seek and enable them to access the required site. This is the reason why many security messages advise the user to check the name of the site in the URL bar before entering confidential data. Filtering and redirection weaken this rule. Users lose one of the trust seals when they consult a web site or receive email.

Users may find that the content they want to access is still online, but the rules used to access it have changed. By changing the behavior of their browser or their machine (as indicated above), they can once again have access. These changes are likely to develop on a large scale.

A proxy configured in this way can, without the user knowing it, divert traffic to another site, capture user traffic, and obtain confidential data. In most cases, the use made of these data,

[11] It may even be replicated as shown by the case of Wikileaks or Copwatch.

[12] The list of these workarounds is long, such as the piratebay servers, which can be accessed via the URL http://www.baiedespirates.be/

[13] Several protocols are currently being defined to encrypt requests to another resolver either by using TLS or by using as a resolution protocol a currently more standard format such as XML rather than HTTP.

which are often personal, by the providers of these services is not controlled by the user.

The implementation of alternative resolvers can also increase the risk of phishing. Take the case of users configuring their systems to access an online gaming site. A few days later, wishing to connect to their bank, they could be directed to a fake site, without realizing the relationship between the two events.

Indirectly, blocking therefore expose the users to new threats.

Filtering can also have significant side effects. Domain name-based blocking using the DNS cannot block a subset of pages or URLs for a site[14]. On the contrary, blocking a domain name blocks all of the URLs using that name, i.e. a complete website. This is called overblocking. For example, one can easily imagine the disruption that would be caused by the complete blockage of Wikipedia, Facebook, or Dailymotion after the publication of specific illegal content (as was the case in the UK for an article on the music album "Virgin Killer"[15]). Such opportunities for denial of service by complete blockage could result in abusive use.

Numerous examples show that configuration errors or incorrect filtering can have adverse effects. For example, in 2006, bizar.dk, a Danish site, found itself blocked by error[16], which led the holder to threaten to sue the Danish police for libel, until the authorities apologized and removed the site from the block list.

4. The results of the AFNIC survey

In 2012 Afnic published the latest edition of its "Technology Backdrop"[17] survey on the foreseeable technological trends and developments over the next ten years for Internet professionals and experts.

Two of the survey's questions are particularly relevant to this document. The questions are set out in the form of predictive assertions, and require an answer expressing the respondent's level of agreement with the assertions.

The first question / assertion (Q70 of the survey) is as follows: "Local DNS resolvers (caches installed on user machines) will take a significant part (25% or more) compared with ISP resolvers or "open" resolvers of the Google DNS type". The results of the survey show a

[14] For a finer granularity, HTTP requests must be blocked, which results in significant additional cost.

[15] http://www.ecrans.fr/Wikipedia-victime-collaterale-de,5883.html

[16] http://translate.google.com/translate?hl=en&sl=da&tl=en&u=http%3A%2F%2Fwww.computerworld.dk%2Fart%2F33184%2Fpolitiet-erkender-censurbroeler-paa-nettet

[17] http://www.afnic.fr/en/about-afnic/news/general-news/6391/show/the-internet-in-10-years-professionals-answer-the-afnic-survey.html

divergence of forecasts between two "schools" (40% of the respondents agreed, 42% disagreed and 18% were undecided). We can nonetheless retain that 40% of the respondents still foresee the use of individual resolvers to replace common resolvers (whether they are those of the ISPs or those of an alternative provider such as Google DNS).

The second question / assertion (Q71 of the survey) is as follows: "In the case of DNS requests sent through a third party (ISPs or suppliers of alternative solvers), the use of alternative solvers will exceed the use of one’s own ISP resolver service". Here again, there is a divergence between "two schools" (41% of the respondents agreed, 39% disagreed and 20% were undecided). Again, it can be noted that almost half of those who made a prediction foresee a gradual erosion of the use of the ISP resolvers in favor of alternative suppliers, leading to a reversal in the balance of power.

In addition, the respondents who agreed with either of the two assertions above were questioned about the motivations for their response. A majority of them believe on the one hand that it will "allow a better guarantee of the integrity of responses," and on the other that it will "enable better performance." This illustrates the credit they give these alternative methods compared with the resolver of their own ISP.

5. Conclusion

The DNS is an essential component in the architecture of the Internet. Many efforts are made to secure it, such as the introduction of cryptography to validate the responses with DNSSEC. The technical security provided by DNSSEC to protect the end-to-end integrity of the DNS against attacks by malicious intermediaries risks being disturbed by the introduction of filtering by altering the responses, and may compromise the adoption of these security extensions.

In addition, for years, ISPs and e-commerce sites have educated users about common risks. By compromising the name resolution mechanism, and thus encouraging risky behavior, DNS filtering may reduce users’ confidence in e-commerce as a whole.

The DNS was not designed to filter content. If a decision on the use of blocking based on the DNS or other techniques were to be contemplated, it should be evaluated in terms of the proportionality between the objective and the relative effectiveness of the measure, and take into account the consequences outlined above[18].

[18] http://www.securityweek.com/dns-blocking-where-technical-considerations-meet-political-considerations