DNS: Types of attack and security techniques

- Presentation of the DNS (Domain Name System)
- Major types of attack targeting the DNS and domain names
- Main security techniques

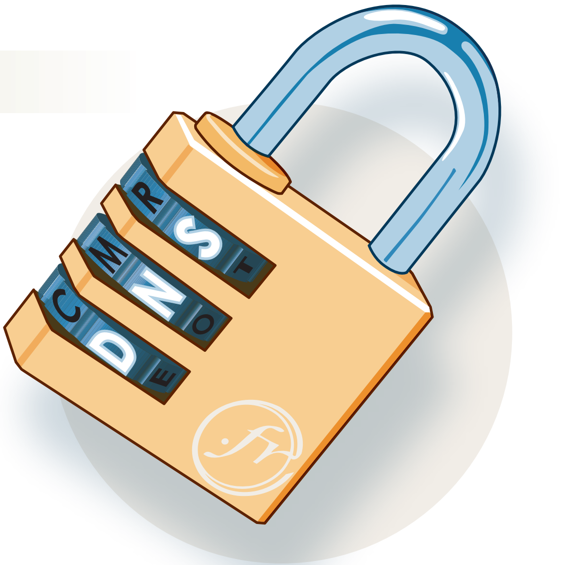

The DNS (Domain Name System), a key component of the Internet infrastructure

In today's world, the Internet is critical to both the economy and society. The infrastructure is highly distributed and home to a wide range of players, such as ISPs, Internet exchange points and network access points, telecoms carriers, hosting providers and registrars. They all play an essential role in how the Internet operates, and each is faced with a specific range of threats.

A key building block of the Internet is the DNS or Domain Name System. This acronym actually conceals a whole range of technical infrastructures, software and hardware required for the domain name system to function correctly, which in turn allows users to access websites and exchange e-mails.

Since its launch in the 1980s, this particularly rugged system has not run into any major problems. However, it is prone to certain weaknesses inherent in its design, the development of new forms of attack and known vulnerabilities.

This feature publication provides an overview of the challenges in ensuring the correct operation of the DNS, the most frequently encountered attacks and the main security techniques.

I Presentation of the DNS (Domain Name System)

I.1 – Tree structure

The DNS is organised into an upside-down tree structure, with a "root" on which the different "branches" depend. The first level of the tree contains top-level domains, such as *.fr* and *.com*. The second level contains "conventional" domain names like "afnic.fr".

The DNS functions like a database distributed over millions of machines and relies on interactions between those machines for identifying which one is most likely to respond to the user's query.

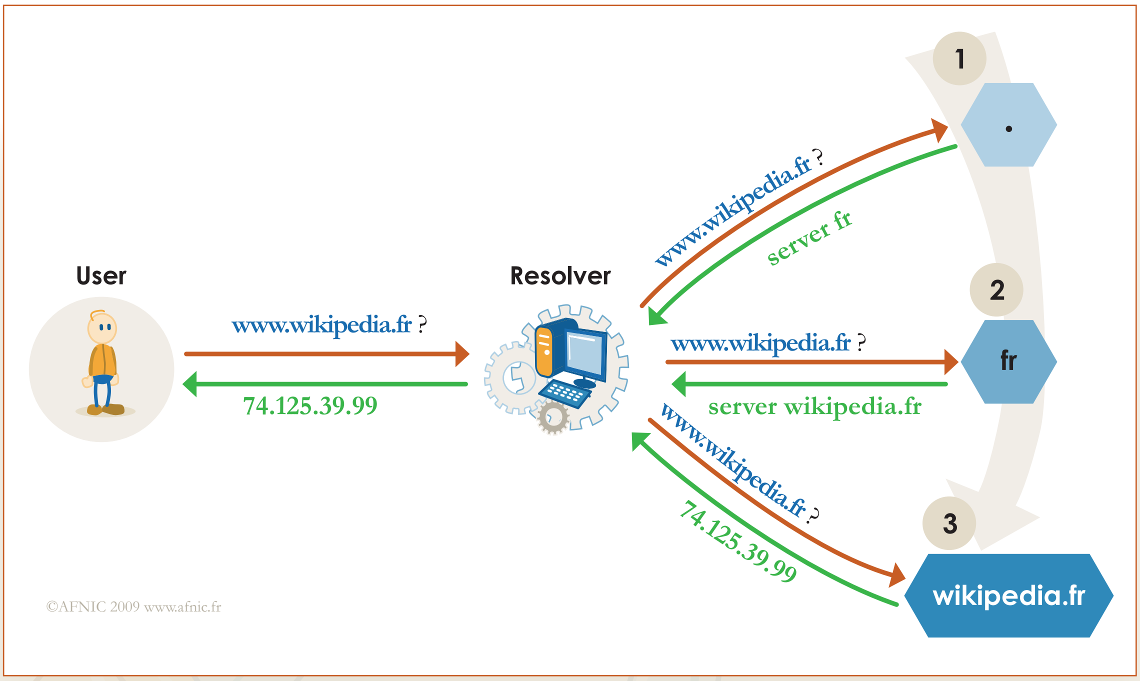

©AFNIC 2009 www.afnic.fr

In the example above, the user wishes to connect to http://www.wikipedia.fr. He sends his query via his browser. The query is received by a server called a "resolver", whose main task is to identify the machine hosting the wikipedia.fr domain name. The resolver first queries the root server to find the servers "authoritative" (meaning responsible) for *.fr*, since the domain name contains *.fr*. The *.fr* servers then inform the resolver which server the wikipedia.fr domain name is hosted on. The resolver can then give the browser the IP address of the web server hosting the content of the www.wikipedia.fr website.

This mechanism always holds true, irrespective of the desired website and regardless of the e-mail address to which users wish to write. That is why a secure DNS system and understanding of the attacks that could harm its operation are challenges for all Internet users.

I.2 – Software

The DNS works with specific software; some applications are marketed, while others are licence-free. The most commonly used software is BIND, developed and maintained by the Internet Systems Consortium (ISC). AFNIC and other registries are committed to supporting the development of BIND 10, the future version of the software.

In some cases, attacks target the actual infrastructures, which include servers hosting the domain names. In other cases, hackers look to exploit loopholes in the software to create abnormal situations from which they can profit. The strategies involved may be subtle, but they often follow relatively well-identified patterns.

II Major types of attack targeting the DNS and domain names

There are several different types of attack on domain names and the DNS.

II.1 – Attacks not specifically directed at the DNS

Some attacks may look to exploit the administrative side of domain names rather than directly targeting the infrastructures and DNS servers:

- **Cybersquatting** involves registering a domain name with the deliberate intent of undermining and profiting from a third party's rights or in some way harming that third party. There are many cybersquatting techniques, and the general aim is to steal the victim's identity and/or divert traffic away from the victim's website.

- **"Name-jacking" or theft** by appropriating the domain name (updating the holder's field and/or contacts) or taking control by technical means to divert traffic, such as by modifying the name servers hosting the site.

Other non-DNS attacks are "social" (convincing a careless employee to give their password to a stranger) or involve such techniques as SQL injections, which were used to attack several registries and registrars early 2009

web

For more information:
www.infoworld.com/t/authentication-and-authorization/google-blames-dns-insecurity-web-site-defacements-722

II.2 – Attacks specifically directed at the DNS

Attacks on DNS infrastructures are mainly technical, using mass attacks or techniques that corrupt the information exchanged between the resolvers and DNS servers:

- **DNS cache poisoning** dupes the resolver into believing that the "pirate" server is an authoritative server in place of the original server. These attacks capture and divert queries to another website unbeknownst to users, the danger being that users might divulge personal data on what they believe to be a bona fide site. The "Kaminsky flaw" discovered during the summer of 2008 is one such attack that poisons DNS resolvers.

- **Denial of service** (DoS) attacks are attempts to make a given service impossible or very hard to access. Attacks sometimes use brute force (saturating servers by flooding them with simultaneous queries) or go for a more subtle approach by exhausting a rare resource on the server. Attacks made against the DNS root system in February 2007 were mainly DoS attacks.

- **Distributed denial of service** (DDoS) attacks are an elaborate form of DoS that involve thousands of computers generally as part of a botnet or robot network: a network of zombie computers that the attacker commandeers from their unwitting owners by spreading malware from one machine to another.

- **Reflected attacks** send thousands of requests with the victim's name as the source address. When recipients answer, all replies converge on the official sender, whose infrastructures are then affected.

- **Reflective amplification DoS**: if the size of the answer is larger than the question, an amplification effect is caused. The same technique as reflected attacks is used, except that the difference in weight between the answer and question amplifies the extent of the attack. A variant can exploit the protective measures in place, which need time to decode the long replies; this may slow down query resolution.

- **Fast flux**: In addition to falsifying their IP address, attackers can hide their identity by using this technique, which relies on fast-changing location-related information to conceal where the attack is coming from. Variants exist, such as single flux (constantly changing the address of the web server) and double flux (constantly changing the address of the web server and the names of the DNS servers).

II.3 – A real-life example: the February 2007 attack against the root system

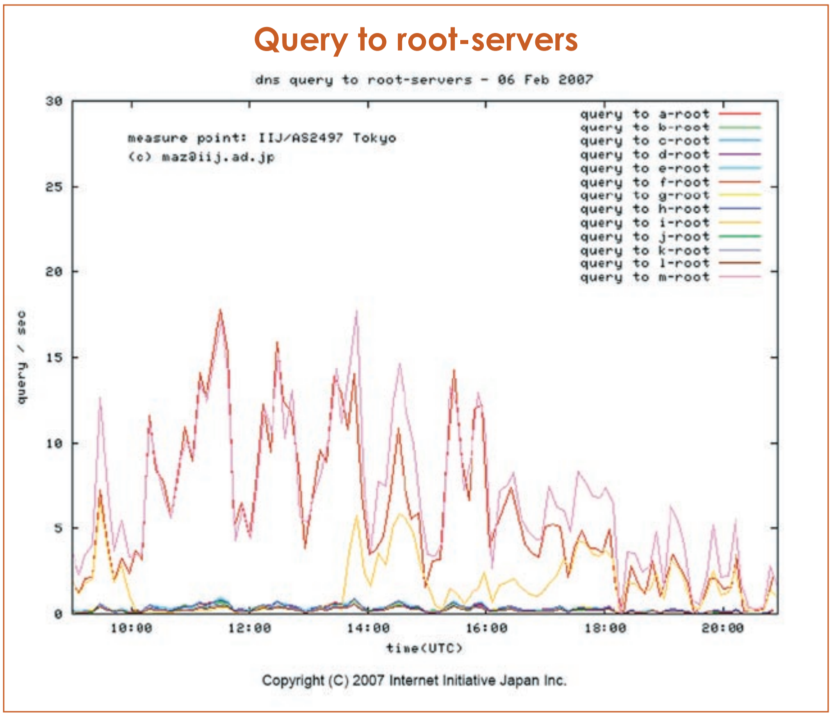

Source: Root Attack – end-user view – Matsuzaki Yoshinobu, 2007
http://www.nanog.org/mtg-0706/Presentations/lightning-maz.pdf

The following chart highlights the impact of the attack on 6 February 2007 against 13 servers hosting the DNS root system.

The chart clearly shows that some servers are much more affected than others, especially M, L and G. However, they only form a minority with the strong impact registered between 10 am and 4 pm, following which resolution times progressively return to normal. Despite the resources ploughed into the attack, the repercussions were actually not critical, since most web users around the world did not experience any poor service performance.

Although it is relatively easy to affect the DNS or server performance, trying to prolong the attack for any length of time without being spotted is much harder. As a result, infrastructures are designed to sustain considerable peak loads in activity over short periods.

The DNS system may be exposed to attacks, but on the whole it is robustly designed, not only capable of supporting a more intensive and varied use of the Internet, but also withstanding mass attacks. This does not mean that better security systems are not required, since the individual measures taken by each player could be easier to break than the DNS system as a whole. Any company on the web must ensure that its presence is not unwittingly built on overly fragile foundations.

web

Other real-life example:

- Attack against a Brazilian bank in April 2009
  (the only well-documented case of a Kaminsky attack)
  www.theregister.co.uk/2009/04/22/bandesco_cache_poisoning_attack/

III Main security techniques

Each entity on the Internet is a link in the value chain, where all links are interdependent. As such, this advice is not aimed at a specific category of users, but all those involved in the operation of the DNS: top-level domain managers (registry), registrars, businesses, ISPs and so on.

This feature publication does not set out to provide a detailed account of the measures that could or should be taken to guarantee an optimal level of security for a company's DNS system. A few guidelines, however, are worth mentioning.

- **Set up the best possible redundancy**, so that a server affected by an attack can be seamlessly replaced by other servers containing the same information, but connected to other networks. That is why registries such as AFNIC always require each domain name to be installed on no fewer than two name servers. Other more sophisticated techniques, like anycast schemes, take redundancy to even higher levels with clear improvements in terms of security and performance.

- **Use the latest DNS software versions**, especially BIND, and install the appropriate patches to prevent attacks exploiting well-known security loopholes.

- **Regularly keep an eye on the servers and their configuration**, preferably from several points across the Internet. Due to the robust nature of the DNS system, it often happens that a server failure is only detected when the last server in the line also fails. To check the configuration, freeware is available, such as ZoneCheck. To monitor the network from outside, companies not wishing to deploy a specific architecture can use existing commercial or community services.

- **Look into deploying DNSSEC**, a DNS security protocol based on server authentication that reduces the threat of DNS cache poisoning. Opinions on DNSSEC changed strongly following the revelation of the Kaminsky flaw, which showed how to effectively exploit vulnerabilities that were already known on a theoretical level.

- **Define a "business continuity plan"** allowing the victim of an attack to continue or restore business with minimal downtime in the event of a major attack. This is a fairly essential precaution for all those that depend on the Internet – and therefore the DNS – for their revenues, particularly companies offering online services to their customers.

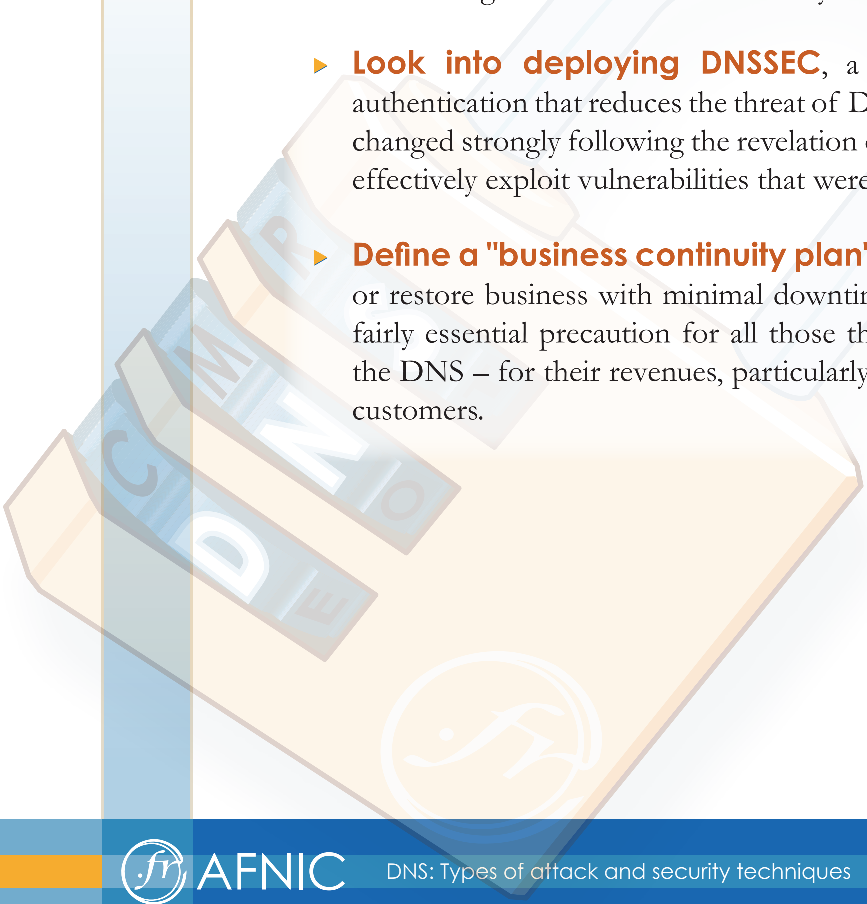

Conclusion

The security of the Internet infrastructure is based on roles being evenly distributed between the different players (service operators, ISPs, registries, registrars, hosting providers, Internet exchange points, public authorities, CERTs…). The wide range of structures, technologies and approaches represents one of the main guarantees for the Internet's resilience.

Every player in this ecosystem must apply the basic rules for effective security: **coordination**, **communication** and **cooperation** (the three Cs). In the case of the Internet, the variety and number of players involved raises a major challenge on both a national and international level.

Given the changing nature and growing scale of threats, isolated or uncoordinated responses are likely to be increasingly ineffective. In a similar vein, continuing to raise awareness on security issues among the different players is one of the underlying actions.

Registries have been strongly mobilised on such issues for several years, many of whom have already developed business continuity plans to overcome any unforeseen incidents beyond their control. This approach has also been adopted by service providers and end users managing their own infrastructures. Nevertheless, considerable progress needs to be made before all links in the security chain are compliant with the 3 Cs rule.

web

Further information

- DNS Resources Directory, a good directory of web resources concerning DNS: www.dns.net/dnsrd/

- AFNIC training aids available under free licence www.afnic.fr/doc/formations/supports

- A good summary report from the CERTA www.certa.ssi.gouv.fr/site/CERTA-2008-INF-002/

- Summary Report of the "Global DNS Security, Stability, and Resiliency Symposium, February 3-4, 2009" www.gtisc.gatech.edu/pdf/DNS_SSR_Symposium_Summary_Report.pdf

Read all of our issue papers:
http://www.afnic.fr/actu/presse/liens-utiles_en