DNSSEC

Domain Name System Security Extensions

Introduction

DNSSEC issues

The security of any system depends on both the security of its various components and the interactions between them. This observation is also valid for the DNS (Domain Name System), a key component in the operation of the Internet, in that almost all online services use domain names at one time or another.

The DNS was developed in the '80s, however, in a context during which its ability to meet the needs in terms of performance and resilience prevailed over security. Over time, and in particular since 2008 when the "Kaminsky flaw" was disclosed, the need to improve DNS security has become a top priority for all the industry.

While it cannot counter all of the attacks possible against the DNS, the DNSSEC protocol can provide protection that is viable – and will be vital tomorrow – against attacks known as "cache poisoning", designed to substitute false information for true data in the process of resolving domain names. The attacker can thus hope to lure users and divert them to its site without their knowledge. But the growing capacity of machines and networks in terms of speed now enables attacks that until today were only theoretical, given their low probability of success.

This issue paper continues on from that which AFNIC has already devoted to DNS security[1] , by exploring the more specific ins and outs of DNSSEC. The paper is designed to enable greater understanding of the issues involved and how DNS security operates, so that the reader can better grasp the developments in question, which are liable to change the face of the DNS in the years to come.

[1] www.afnic.fr/data/divers/public/afnic-dns-attacks-security-guide-2009-06.pdf

1 Organisation and operation of the DNS

The DNS is organised in the form of an inverted tree, with a "root" from which stem the various "branches". At the first level of the tree are the "Top Level Domains" such as *.fr*, *.com* etc. At the second level lie the "classic" domain names such as "afnic.fr".

Operating as a distributed database on millions of machines, the DNS is based on interactions between these machines in order to identify the one most likely to respond to the request of a visitor.

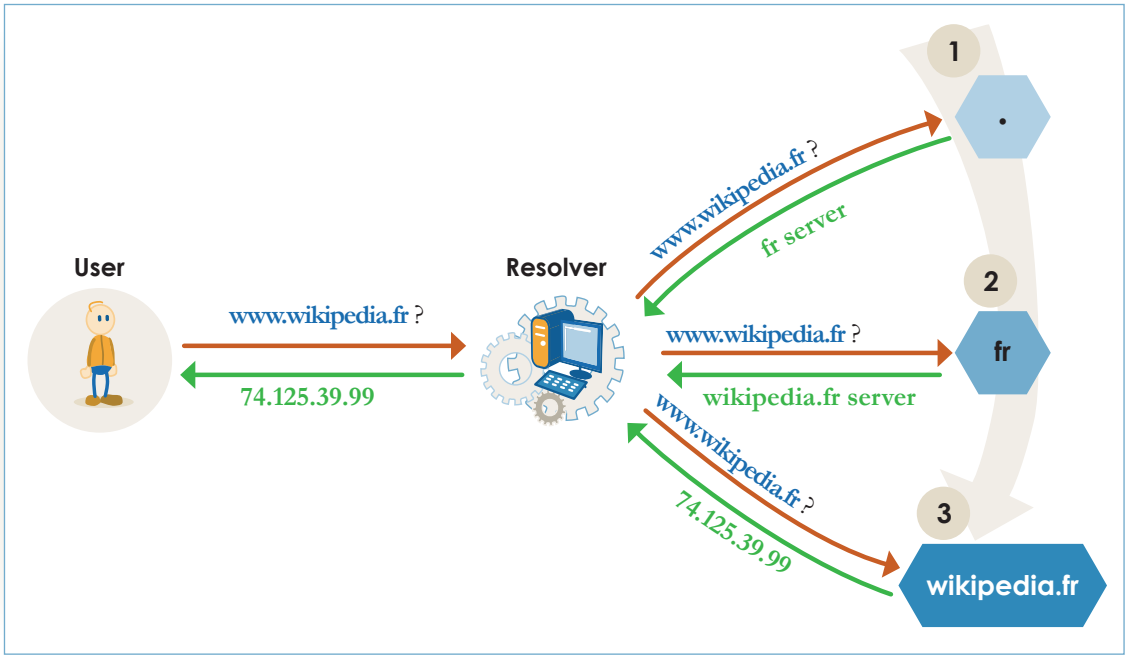

DNS Resolution

In the example above, the user wishes to connect to the site http://www.wikipedia.fr, and sends the request via a browser. The request is received by a server, known as a "resolver", the primary task of which is to identify the machine on which the domain name wikipedia.fr is installed. The resolver initially addresses the "root" of the DNS, which specifies which servers are "authoritative" (i.e. competent) for *.fr* since the domain name is in *.fr*. Secondly, the *.fr* servers in turn indicate to the resolver that the domain name wikipedia.fr is hosted on a given server. The resolver is then able to indicate to the user's browser the IP address of the web server hosting the content of the website www.wikipedia.fr.

This pattern applies, regardless of the website that the user wishes to access, or the e-mail address to which s/he wishes to write.

2 Cache poisoning attacks

The issue paper by AFNIC devoted to the security of the DNS provides an overview of main types of possible attacks, from those that do not directly target the DNS (cybersquatting, theft of domain names, etc.) to those focusing on the DNS, such as Denial of Service or cache poisoning.

DNSSEC specifically addresses cache poisoning attacks designed to intoxicate the resolver into believing that the "hacker" is the legitimate server, instead of the original server. In particular, this allows the hacker to collect and divert requests to another website without the users realising this is the case, the issue at stake being the risk of users entrusting personal data to the hacker in the belief that they are on the legitimate site of the victim of the attack.

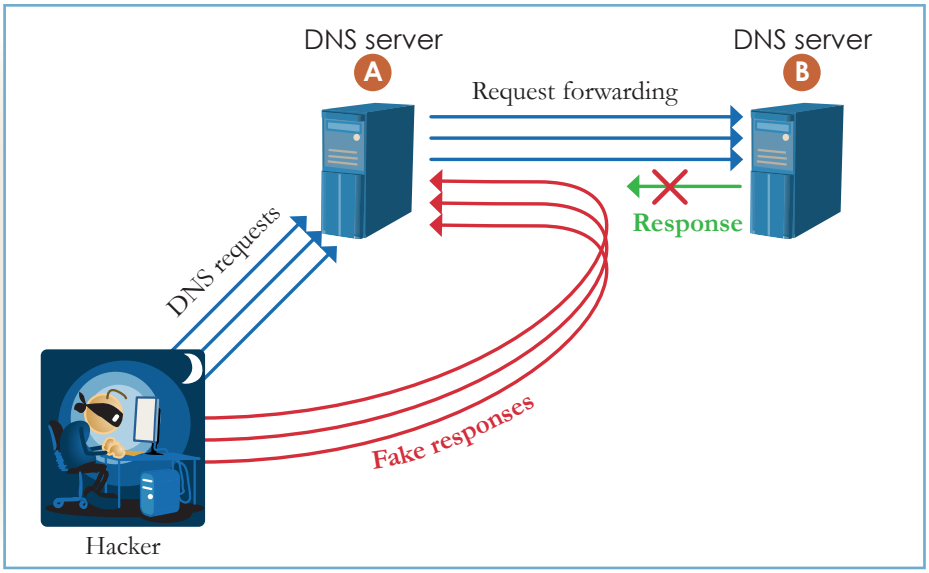

Cache poisoning attack

Correct operation of the DNS depends to a high degree on the reliability of the data transmitted at each step. DNS Security Extensions seek to address this constraint by ensuring the integrity of the data transmitted over the network, in particular between resolvers and the authoritative servers.

3 What DNSSEC can do

DNSSEC is the acronym for Domain Name System Security Extensions, and refers to a finite set of security extensions of the DNS protocol.

These extensions use the mechanisms of asymmetric cryptographic signature to authenticate records. The signatures and public keys are in the form of complementary new records in order to ensure their authentication.

These new records cause an increase in the size of messages and the number of exchanges required to verify signatures and keys. DNSSEC therefore requires additional machine resources, as well as recent, updated versions of DNS software.

The protocol was naturally designed to operate safely in an environment which, initially at least, will not be entirely composed of DNSSEC-validating resolvers. In the case where a non-DNSSEC validating resolver queries DNSSEC servers, they simply return the information normally exchanged without the signatures or records specific to DNSSEC.

4 What DNSSEC cannot do

Although rightly presented as a necessary evolution in terms of securing the DNS, DNSSEC does not purport to address all the types of attacks which may occur.

For example, it is not intended to encrypt DNS records, or ensure the confidentiality of the information exchanged over the network, or ensure the security of a transaction in the way that SSL certificates do. It does not protect against phishing and theft of domain names, against virus infection or other forms of technical infection of workstations, or against attacks on the websites themselves (SQL injection etc.).

In addition, each security level only makes sense within the chain of trust: for example, signed records cannot be authenticated as long as the key to the zone has not been published in the parent zone. Conversely, no specific security is provided for users if their resolver has not implemented DNSSEC.

Finally, DNSSEC does not protect the integrity of data that have been modified either accidentally or intentionally upstream of their publication in the DNS.

5 Using keys in DNSSEC

DNSSEC uses a mechanism based on a key pair with complementary roles. The first key, the private one, signs by encryption, while the second key, the public one, verifies the signatures by decrypting them.

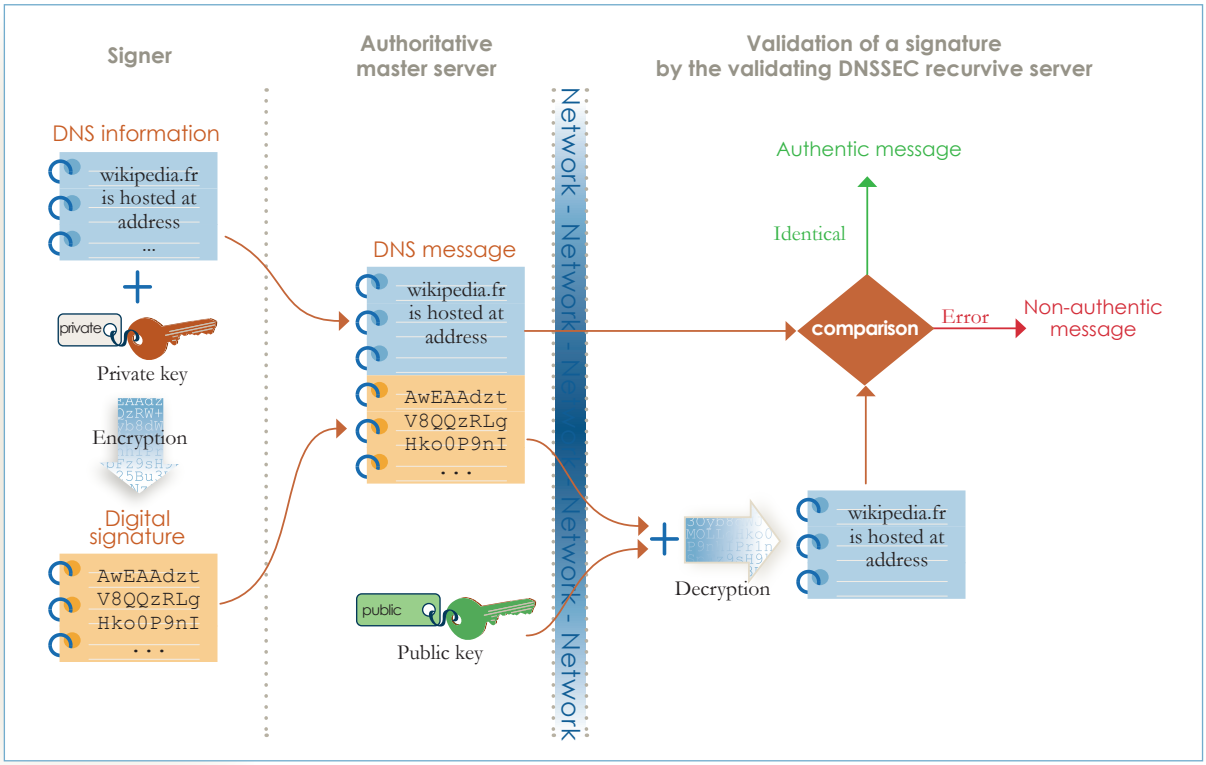

Signature and signature validation in the case of DNS

A message is encrypted using the private key, creating a signature accompanying the message. This signature can be authenticated by DNSSEC validating resolvers using the corresponding public key published in the zone.

To prevent an attacker from using its own set of keys, the DNSSEC designers have proceeded such that each parent zone, located on the upper level of the DNS tree, guarantees the authenticity of the keys of their child zones by signing them, the exchange of keys between parent and child zones exclusively taking place through secure channels. This system creates real chains of trust down to the root of the DNS.

DNSSEC also requires proper management of the life of signatures in relation to the publication of the keys, since expired signatures may block the resolution of the zone in question.

Like other DNS records, DNSSEC records have finite lifetimes in order to ensure the resilience of the system and to enable updates in due course. When a signature expires, it is considered invalid and is no longer possible to connect to the service. It is therefore necessary to regularly re-sign records in addition to new signatures produced due to the creation of, or changes in records.

The deployment of DNSSEC therefore induces extra work and potentially generates new risks of errors (incorrect configurations, expired signatures, etc.). For this reason, if the key inserted in the parent zone is not that used for the signature, the child zone will be considered invalid and there will be a resolution error.

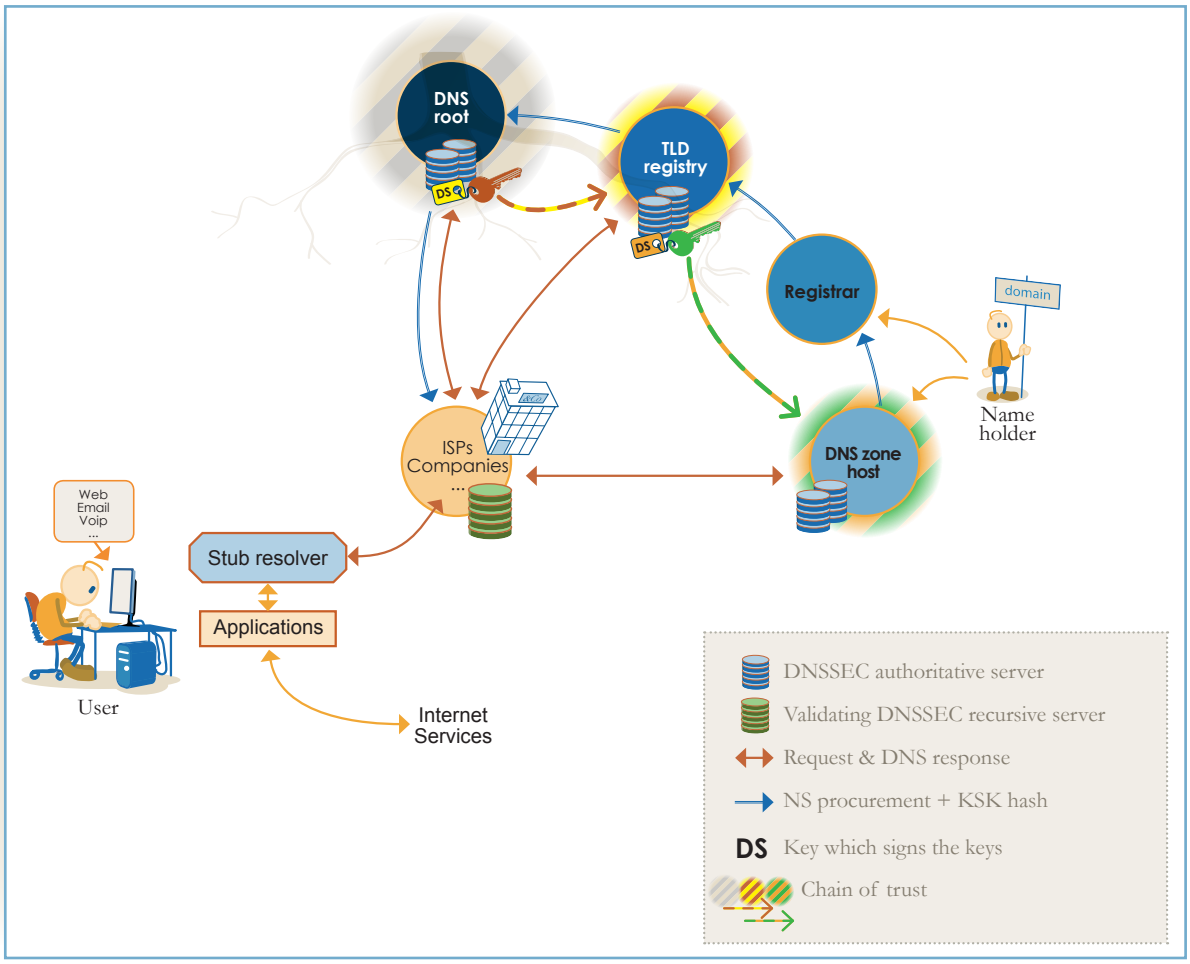

Components of the DNSSEC chain of trust

6 The deployment of DNSSEC

Although there has been a sharp build-up in the use of DNSSEC since the disclosure of the Kaminsky flaw, the security extensions are nothing new for DNS experts. Indeed, the IETF (Internet Engineering Task Force) started working on such a protocol for securing the DNS as early as 1995.

It was not until 2005, after many iterations in the technical community, that the registry for the *.se* Top Level Domain (Sweden) was the first to sign its zone. The country was also the first to open the service to registrants in 2007.

Viewed with interest but without any sense of urgency until 2008, DNSSEC became a top priority for all the TLD registries from the summer of 2008 onwards, when the Kaminsky flaw made everyone realise the magnitude of the potential problem. By mid-2010 fifteen registries had signed their TLD, and most of the others are working on the issue. The *.fr* TLD was signed on 14th September 2010.

Even though DNSSEC involves the registries, its deployment does not stop there: the managers of the resolvers (ISPs, registrars, companies, etc.) must in turn implement DNSSEC so that it is fully operational.

A number of solutions exist, from open-source software solutions such as open-dnssec, to proprietary boxes, including programming libraries that simplify developments.

The workload caused by the deployment of DNSSEC, both in costs and in terms of refresher training for the teams, can be seen as important, given the risks that it helps prevent. Although potentially justified in the short term, this approach does not take into account, however, the increase in network throughput and machine capacity, which mechanically advances the chances of successful cache poisoning attacks in the current format of DNS operation. The implementation of DNSSEC should be considered as an unavoidable necessity in the medium term.

The main issue in the deployment of DNSSEC is that of setting up effective key management. Keys must be regularly changed in order to prevent their theft or recalculation, but they must also be protected by physical and digital means. Updates must also be carried out in order to take into account the propagation lead-time for information in the DNS, so that the resolvers do correlate a new signature with an old key or vice-versa. This imposes a certain period of latency between the declaration of the keys and the signature with them.

The implementation of DNSSEC means each zone must carry out the following operations:

- Create the keys;
- Sign its records;
- Publish the signed zone;
- Manage their validity periods;
- Manage the publications of the key summary in the parent zone with each KSK rollover;
- Check the publication of a new key before signing with it.

7 Questions to ask about implementing DNSSEC

The implementation of DNSSEC is not only a major technological development for the DNS: it also induces an adaptation of the organisation and management of domain names and may require changes in both, mainly because of the emergence of new issues such as key management.

It may therefore be necessary to foresee the need for a certain increase in the technical competence of the teams - the subject remains little known for the time being - but also a change in procedures while maintaining overall consistency with the various aspects of the company's security policy.

Here is a list of the various questions to ask, without presuming to be exhaustive:

- Who hosts the company's DNS zones?
- How skilled is the provider in DNSSEC management?
- What level of transparency does the provider offer its customers in terms of practices?
- What security systems does the provider use in addition to DNSSEC?
- Does the provider offer secure transaction channels?
- What type of organisation does provider recommend for key management, within its own teams and in conjunction with those of the company?
- Is there an impact in implementing DNSSEC in terms of costs?
- If the company calls on a provider to sign its zones, how does it manage transfers to another provider of its signed domain names?
- Is the resulting level of dependency on that provider acceptable, or is it liable to incite the company to internalise that function at least for its most strategic domain names?
- If such is the case, what resources and capabilities are available to the company to solve these issues, and what means would be required to support that build-up?

8 Find out more

- DNSSEC page on AFNIC's website:
  www.afnic.fr/dnssec
- Wikipedia page on DNSSEC in English:
  http://en.wikipedia.org/wiki/Domain_Name_System_Security_Extensions

For further information about DNSSEC at AFNIC :

solutions@afnic.fr

9 Glossary

DNS : Domain Name System.

DNSSEC : Domain Name System Security Extensions.

DS : Delegation Signer.

KSK : Key Signing Key.

NS : Name Server.

SSL : Secure Sockets Layer.

SQL : Structured Query Language.

ZSK : Zone Signing Key.

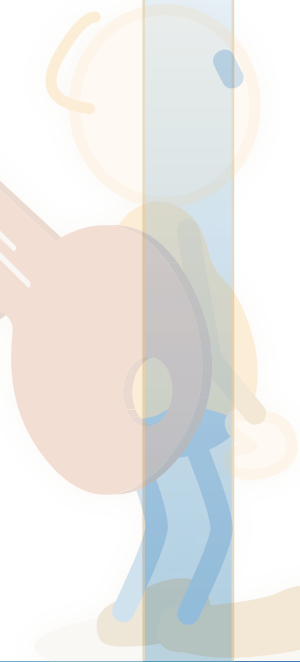

Read all of our issue papers:
http://www.afnic.fr/actu/presse/liens-utiles_en

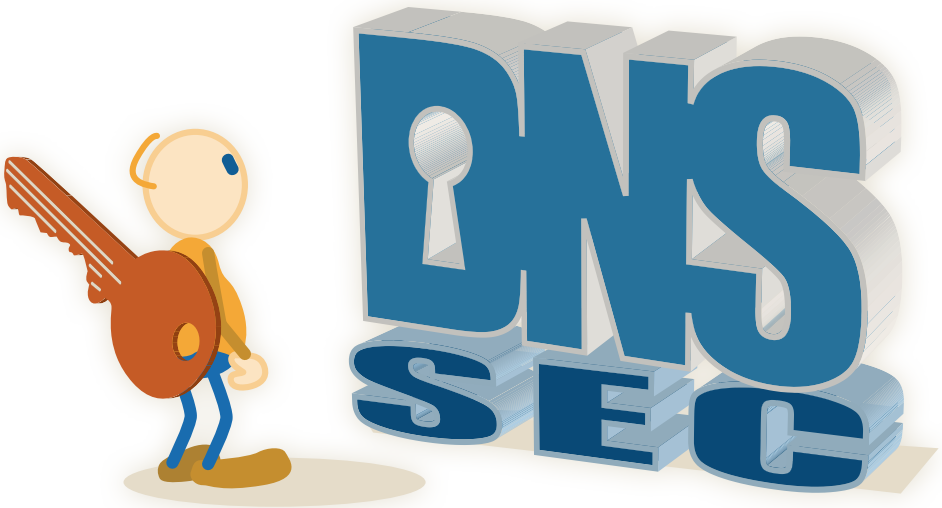

web **www.afnic.fr - afnic@afnic.fr**