



Internet Resilience in France

2015



Document produced by the ANSSI with the participation of the Afnic.

Researched and written by: François Contat, Pierre Lorinquer, Florian Maury, Julie Rossi, Maxence Tury, Guillaume Valadon and Nicolas Vivet.

The editorial team would like to thank the observatory members and the reviewers for enriching this report with their comments and remarks.

Document formatted using \LaTeX . Figures produced with TikZ and PGFPlots.

Please send any comments or remarks to the following address:

`rapport.observatoire@ssi.gouv.fr`

Table of contents

Executive Summary	5
Presentation of the Observatory	7
Introduction	9
1 Resilience in Terms of the BGP Protocol	11
1.1 Introduction	11
1.2 Prefixes Hijacks	15
1.3 The Use of Route Objects	19
1.4 Declarations in the RPKI	22
2 Resilience in terms of the DNS protocol	25
2.1 Introduction	25
2.2 Dispersion of authoritative DNS servers	30
2.3 Implementing DNSSEC	34
2.4 Dispersion of inbound e-mail relays	36
3 Resilience in terms of the TLS protocol	43
3.1 Introduction	43
3.2 Session Negotiation	46
3.3 Robustness of Certificate Signatures	50
General Conclusion	53
Bibliography	55
Acronyms	59

Executive Summary

Since 2011, the Internet Resilience Observatory in France has studied the technologies that are critical for the proper operation of the Internet in France. In order to understand the dependence of the French economy and society on those of other countries, the Observatory focuses on the French Internet, a subset of the Internet in France that does not take into account foreign players.

Resilience is defined as the ability to operate during an incident and return to the nominal state. It can be characterized by measurable indicators, some of which come directly from engineering rules called best practices.

Written by ANSSI¹, this report analyzes resilience through the BGP², DNS³ and TLS⁴ protocols. The first two are used respectively to carry data using routing advertisements, and provide the mapping between a domain name and an IP address. The third is in particular used to encrypt communications between a web server and its clients.

In 2015, the Observatory identified that the recommended TLS version 1.2 is supported by 75% of web servers for delegated zones under the .fr TLD⁵. Regarding IPv6, the trends initiated in previous years are continuing, indicating that best practices in operating this protocol seem to be followed by few. To allow the reproduction of a part of the results, the main tools used for the BGP analyses have been published [1, 2].

The Observatory encourages all Internet players to appropriate the best engineering practices accepted for BGP [3], DNS [4], and TLS, and to anticipate the threat of DDoS [5]. In addition, the Observatory makes the following recommendations:

- **monitor prefix advertisements**, and be prepared to react in case of BGP hijacks;
- **use algorithms supporting forward secrecy** and **abandon SSLv2 and SHA-1** in favor of more robust mechanisms;
- **diversify the number of SMTP and DNS servers** in order to improve the robustness of the infrastructure;
- **apply best practices** including those contained in this document, to limit the effects of failures and operational errors;
- **pursue the deployments** of IPv6, DNSSEC, and RPKI to develop skills and to anticipate possible operational problems.

¹Agence nationale de la sécurité des systèmes d'information / The French national authority for the defence and the security of information systems.

²Border Gateway Protocol.

³Domain Name System.

⁴Transport Layer Security.

⁵Top Level Domain.

Presentation of the Observatory

The Internet is an essential infrastructure for economic and social activities at the global, national and local levels. A major outage would significantly affect the smooth running of France and its economy. In addition, the operation of the Internet as a whole is often misunderstood and can be perceived as an opaque system, managed by players whose roles are poorly identified. Despite the importance of that issue, no organization in France had been entrusted with studying the risks of a malfunction of the Internet at the national level.

Set up under the aegis of the ANSSI⁶ in 2011, the purpose of the Internet Resilience Observatory is to improve knowledge about the Internet by studying the technologies critical for its proper operation. One of its objectives is to increase the collective understanding of the French Internet in order to have a coherent vision as complete as possible. In particular this helps to identify the interactions between the various players concerned.

By nature, the Internet is international and has no borders. It is possible, however, to define the Internet in France as all of the French and international players engaged in an activity related to Internet technologies within the country. As part of its studies, the Observatory focuses on the French Internet, a subset of the Internet in France, which does not include foreign players. Studying the French Internet helps better understand the interdependencies of the French economy and society on foreign companies or organizations.

Resilience, in turn, is defined as the ability to operate during an incident and return to the nominal state. A natural extension of resilience is robustness, i.e. the capacity, beforehand, to minimize the impact of an incident on the status of the system as a whole. On the technical level, the resilience and robustness of the Internet can be characterized by a set of measurable technical indicators. Some are directly based on engineering rules, referred to as best practices, defined by the technical and scientific community.

The Internet Resilience Observatory in France is also tasked with defining and measuring representative resilience indicators, and to make its findings public. Stakeholders in the French Internet are involved in the initiative in order to increase the efficiency of the response and encourage the widest possible adoption of best practices.

⁶Agence nationale de la sécurité des systèmes d'information.

Introduction

Backed by their experience on BGP and DNS protocols, the Observatory team wanted to extend its analyses to other protocols to better understand the Internet in France. Their choice fell on the TLS protocol that in particular enables encryption of the communications between a web server and its clients. The scope of this new analysis consists of all the websites corresponding to delegated zones under the `.fr` TLD implementing HTTPS⁷. This new report presents various indicators and measurement methodologies, and the results associated with the observations made on TLS.

Regarding the BGP protocol, the main tools used for the analysis have been published [1, 2]. They can be used both to analyze the BGP archives of the RIS⁸ project, and to detect conflicts between prefix advertisements [6]. In 2015, a new methodology of detecting routing leaks was also developed. It is designed to facilitate manual analyses while reducing the number of false positives. This is an important step in the automation of repetitive classification tasks previously performed manually.

The study of mail relays via the DNS protocol, introduced in the 2014 report, has been improved, especially with respect to the dependencies on third-party domain names, and dispersions by country and by operator. This development helps better understand the phenomena of concentration on hosting platforms that can affect service availability. To echo the new indicator on the TLS protocol, observations on DNSSEC⁹ now incorporate an analysis of the cryptographic algorithms used.

For the sake of brevity, this new report provides a summary of the analyses and details the highlights of 2015. Previous reports are thus the reference in terms of the descriptions and methodologies for recurring technical indicators.

In a nutshell

French operators wishing to obtain detailed information about BGP indicators may request individual reports.

⁷The version of the HTTP protocol protected by TLS.

⁸Routing Information Service.

⁹Domain Name System Security Extensions.

Chapter 1

Resilience in Terms of the BGP Protocol

1.1 Introduction

1.1.1 Operation of the BGP Protocol

Each Internet operator manages sets of contiguous IP¹ addresses, called prefixes, which it can divide for its own needs or those of its customers. To form the infrastructure of the Internet, the operators connect to each other using the BGP protocol [7]. The aim of this protocol is to exchange reachability information about the prefixes between two operators which are then called AS² and identified by a unique number.

Each of the ASes informs its peer that that it can route traffic to its prefixes. Interconnections are divided into two categories:

- **peering:** an agreement in which each peer advertises to the other the prefixes that it manages. For example, if an ISP and a content broadcaster come to a peering agreement, they will exchange their traffic directly;
- **transit:** a commercial agreement between a customer and its transit operator. In practice, the customer advertises its prefixes to its operator so that the latter can propagate them. In return, the latter advertises the rest of the prefixes constituting the Internet.

In a BGP interconnection, each peer associates an AS_PATH with the prefixes it advertises. In 1.1 the router of AS65540 has learnt the AS_PATH 64510 64500 for prefix 192.0.2.0/24. To reach IP address 192.0.2.1, a packet from the AS65540 will cross AS64510 before arriving at the AS64500. The AS managing the prefix is located to the right in the list constituting an AS path.

In practice, a BGP message of the UPDATE type is used to indicate the AS path associated with a prefix. This BGP message is responsible for advertising the routes. In Figure 1.1 the router of AS65550 has two routes to reach prefix 192.0.2.0/24. One has been learnt via a *peering* interconnection (blue), and the other via a transit interconnection (purple). In the absence of any other information, the shortest AS path determines the route used. In this example, it is the *peering* link.

There is no robust authentication method for prefix advertisements. For this reason,

¹Internet Protocol.

²Autonomous System.

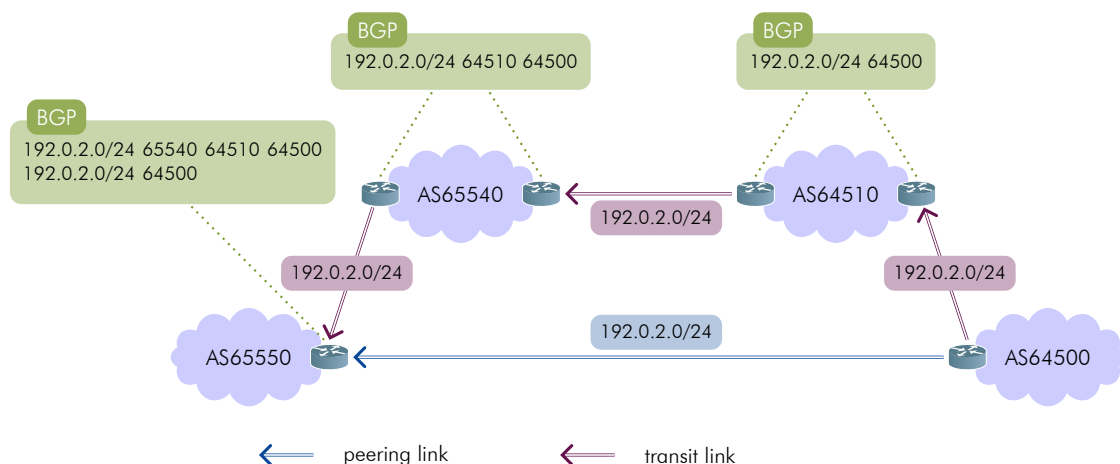


Figure 1.1: Example of AS paths on transit and peering links

a malicious AS can advertise a prefix belonging to another AS. This is referred to as prefixes hijacks. The consequences can be more or less serious according to the advertisement being made. For example, the victim network can become unreachable for all or part of the Internet. This type of incident can also result in the redirection of traffic destined for the victim network to the network having hijacked the prefixes.

1.1.2 Route Objects

According to best practices [3], an organization declares in the `whois` database the prefixes that it advertises by BGP. These declarations must be made via route objects and are stored in the database of an IRR³. This service is operated by each RIR⁴, that for Europe being the RIPE-NCC⁵. A route object clearly identifies the ASes liable to advertise the prefixes of the organization.

```

route:          198.18.7.0/24
descr:         Example prefix
origin:        AS64496
mnt-by:        MNTNER-RO-EXEMPLE

```


Figure 1.2: Example of a route object

The route object of of figure 1.2 indicates that prefix 198.18.7.0/24 is advertised by AS64496. The organization could delegate the use of the prefix to a client or a partner. In this case, the `origin` attribute would concern an AS number other than 64496.

³Internet Routing Registry.

⁴Regional Internet Registry.

⁵RIPE Network Coordination Centre.



To allow certain types of deployments, and some forms of anti-DDoS protection, it is legitimate to declare various `route` objects with identical `route` attributes but different origin attributes. The `mnt-by` attribute indicates the people in charge of the declaration and the maintenance of this `route` object.

In particular, `route` objects allow a transit provider to filter the advertisements of its clients. These filters allow it, for example, to guard against configuration errors resulting in advertisements of prefixes that do not belong to them.

1.1.3 The RPKI

A secure version of BGP, called BGPsec⁶ [8], is still under development at the IETF⁷. In this model, each AS has a certificate linking a public key to an AS number. When a prefix is advertised, the router includes a signature with the prefix, its AS number and that of its neighbor. Each AS propagating the advertisement adds a similar signature to the BGP message. In this way, the integrity of the AS path can be verified.

The RPKI⁸ [9] is a preliminary step to the implementation of BGPsec in particular by introducing a mechanism to verify the origin of an advertisement. Each RIR administers a PKI⁹ dedicated to the certification of IP resources (IP prefixes or AS number) under its management. For example, the RIPE-NCC is at the root of the chain of trust upon which European operators depend, and may issue a certificate to each of them.

The RIRs maintain the repositories containing RPKI objects that have been cryptographically signed. Among these objects, ROAs¹⁰ are `route` objects containing greater amounts of information, because they can be used to specify the maximum length of prefixes advertised by an AS. For example, a ROA can specify that AS64500 is entitled to advertise prefixes ranging from 198.18.0.0/15 to 198.18.0.0/17. Unlike `route` objects, ROAs can expire, since a validity period is associated with them.

1.1.4 Data and Tools

To study the resilience in terms of BGP, the Observatory uses BGP data archived by the RIS project [10]. Thirteen specific routers, called collectors, record in real time all of the BGP messages received from their peers. The geographical distribution of these collectors can be used to obtain the local vision of the Internet for a hundred ASes worldwide, mainly in North America and Europe.

The routing information is analyzed by the Observatory with dedicated tools, some of which have been published in open source. For example, the transformation of binary


⁶Border Gateway Protocol Security.

⁷Internet Engineering Task Force.

⁸Resource Public Key Infrastructure.

⁹Public Key Infrastructure.

¹⁰Route Origin Authorizations.



BGP messages into an intermediate textual format is ensured by the *MaBo* [1] tool. The detection of prefix hijacking is performed by *TaBi* [2] and the study of AS connectivity by the *AS Rank* tool [11].

The industrialization of these tools has also been the subject of large-scale work by the team of the Observatory, in order to produce the indicators more often and without manual intervention. For example, the performance of regular tasks, such as the recovery of BGP archives or *whois* and RPKI repositories [12], is performed using the *luigi* library [13]. In addition, certain tasks for which the processing time is too long are performed on a distributed computing platform by implementing *disco* software [14].

Between 2013 and 2014, the Observatory experimented the use of active measurements to try to better qualify prefix hijacking. The correlation of the data plan information with routing information from the control plane gives interesting results. However, this experiment was not renewed in 2015 because the *RIS* [10] and the array of *Atlas* probes induce a latency of a few minutes. The Observatory is continuing to work with the RIPE-NCC to make measurements in real time possible.

1.1.5 Development of the French ASes

In 2015, using the method defined in the previous reports, the Observatory identified 1588 French ASes. Of these, 1,001 are visible in the BGP archive, i.e. they advertised at least one prefix during the year.

A hard core of 869 active ASes advertised at least one prefix per day throughout 2015, representing approximately 87 % of the total number of distinct ASes visible during the year. Of the 13 % of the ASes remaining, about 50 % of them were visible for half of the year. Finally, 587 listed ASes did not advertise a prefix in 2015.

In a nutshell

In 2015, the Observatory identified 1588 French ASes. Among them, 1001 ASes were visible at the end of December 2015 against 880 at month-end December 2014.

1.2 Prefixes Hijacks

Overall results

In 2015, the Observatory detected 6,392 advertisement conflicts. They targeted 344 distinct French ASes and 1,350 prefixes. Their classification is provided in Figure 1.3. Nearly 50% of them were legitimate advertisements validated by route objects or ROAs. About 2% of the disputes were only validated by ROA.

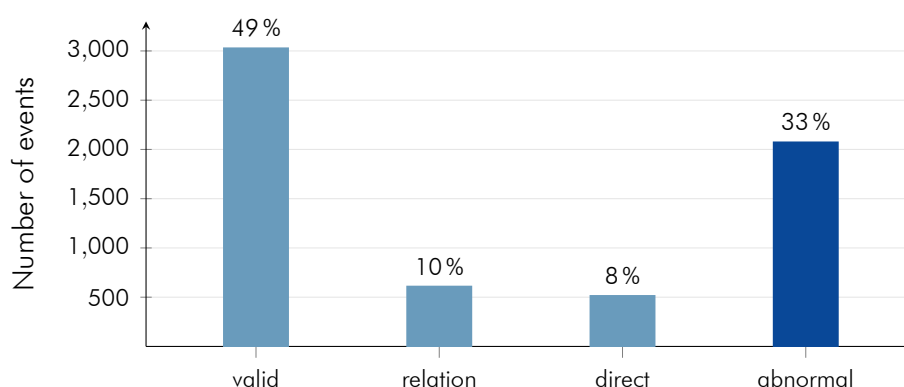


Figure 1.3: Types of conflicts detected in 2015


The "relation" and "direct" categories enabled the elimination of 18% of the conflicts. These simple validations, based on technical and commercial links between ASes, are nevertheless effective in limiting the number of conflicts that are not hijacks. It means that only 2,070 abnormal conflicts will be taken into consideration when identifying instances of prefixes hijacks.

Route leaks

Route leaks appear after BGP router configuration errors. They are characterized by a large number of conflicts originating from the same AS over a short period.

The Observatory has developed an algorithm to automatically detect these global leaks. It is designed to model a leak as an event in which there simultaneously occurs over a short period of time a significant increase in the number of prefixes advertised by an AS and the number of ASes in conflict with this AS. The correlation of these two criteria makes it possible to work with low values to detect peaks, thereby avoiding wrongly excluding certain cases.

During 2015, the algorithm detected several thousands of ASes with peaks in prefix advertisements and hundreds of ASes with peaks in the numbers of ASes in conflict. After completing the correlation, it appears that only 35 ASes were the cause of route leaks.



Some ASes having made route leaks several times in 2015, these results match the 38 route leaks this year, i.e. more than 3 per month. Of these 35 ASes, only 14 were in conflict with French ASes on the dates when the route leaks were detected, impacting a total of 175 French ASes.

Some ASes having made leaks several times in 2015, these results match the 38 leaks this year, i.e. more than 3 per month. Of these 35 ASes, only 14 were in conflict with French ASes on the dates when the route leaks were detected, impacting a total of 175 French ASes.

In particular, on October 27, from 10:00 to 11:00, an AS in Hong Kong advertised more than 15,000 additional prefixes, conflicting with over 6,000 ASes including 93 which were French. The second was an Indian AS [15] which advertised on November 6, nearly 30,000 additional prefixes, conflicting with over 3,000 ASes including 36 French ASes between 06:00 and 16:00. The third was a Greek AS that on October 9 at 13:00 advertised nearly 30,000 additional prefixes, conflicting with over 3,000 ASes including 33 which were French. The other 11 were in conflict with less than 20 French ASes.

The algorithm used is restrictive and tends to minimize the number of false positives at the expense of the number of false negatives. A manual analysis confirmed that the total number of route leaks exceeded the 38 detected. Certain events [16] were not detected because of their low impact on the French Internet or their appearance was not correlated with the leaks detected [15]. Out of the 2,070 abnormal conflicts, 1,480 correspond to global table re-advertisements.

In a nutshell


In 2015, the Observatory identified 35 ASes causing route leaks.

Protection against DDoS

In 2015, the threat of DDoS attacks remained high for French ASes. There are various techniques [5] for limiting the impact. One of them is based on BGP and is characterized by prefix advertisements made by a specialized operator, instead of the attacked AS. In the BGP data, this technique is seen as a conflict of advertisements.

The objective is to divert traffic to a specialized provider which has a large throughput capacity, and equipment both to clean up the traffic and protect the destination IP addresses. In practice, the specialized provider advertises more specific prefixes ¹¹ than those advertised by the client it protects, in order to recover all of the traffic. Once

¹¹ Usually /24 prefixes.



the clean-up has been done, legitimate traffic may, for example, be sent to the client in a tunnel.

In 2015, the Observatory showed 149 abnormal conflicts that corresponded to protection against DDoS, ranging from a few hours to several months. The French ASes protected were of different types, such as hosting, insurance, or online betting sites. It is interesting to note that none of the specialized operators used was French.

In a nutshell

Nearly 150 abnormal conflicts corresponding to protective mechanisms against DDoS were highlighted in 2015.

Automatic filtering of abnormal conflicts

To facilitate the manual analyses, the Observatory improved its automatic detection capabilities of prefixes hijacks. For example, a new filter automatically identifies relationships between ASes by studying the proximity of their names. This makes it possible, for example, to highlight conflicts between a foreign AS and its French subsidiaries. Likewise, a similar filter is used to identify typos in the AS numbers in BGP interconnection configurations.

The abnormal conflicts involving reserved prefixes¹² or which are too specific are filtered because it is difficult to determine their exact origin. Those from special AS numbers¹³ are filtered as well. This step helps to remove approximately 250 abnormal conflicts.


The following filter is to identify abnormal conflicts between two ASes for which there are conflicts of another class. If conflicts are validated by route objects for some prefixes, but not for others, there is probably a strong relationship between the two ASes. 80 abnormal conflicts are filtered in this way.

Finally, filters on the duration of conflicts, their visibility by RIS collectors, and the country of the hijacker AS are applied. Those that last more than two months, and are visible by less than 10 RIS peers out of 120, are filtered. Conflicts arising from French ASes are also filtered. Nearly 250 additional conflicts are thus excluded.

Filtering based on the similarity of AS names identified four conflicts between an operator and its French subsidiary. That involving AS numbers highlighted a configuration error for a BGP router that generated five separate conflicts. These automatic filters

¹²Such as prefix 6to4 2002::/16.

¹³These are private ASes, documentation, and AS_TRANS.



effectively limit the manual analyses that need to be performed. Only 89 abnormal conflicts had to be studied carefully.

Prefixes Hijacks

Prefixes hijacks has very specific characteristics. The malicious AS usually advertises a /24 prefix, more specific than the legitimate advertisement, for a short time. Its aim is to retrieve traffic, while limiting the counter-measures the hijacked AS can set up.

Following the automatic and manual analyses, there remained 40 abnormal conflicts that could be cases of prefixes hijacks. To reduce their number, only the abnormal conflicts with more specific advertisements were kept. The result was that 26 abnormal conflicts corresponded to cases of hijacking with terms ranging from five minutes to four days.

In early July, a Russian AS was responsible for five separate conflicts targeting a French AS. This behavior is suspect, and similar to observations made in 2014 related to spam campaigns [17]. During the year, it advertised nearly 150 different prefixes that it did not own. Of the other abnormal conflicts, ten corresponded most likely to other spam campaigns using BGP. A Romanian AS identified in the previous report conducted this type of hijacking in 2015.

Finally, 15 abnormal conflicts had characteristics that seemed to indicate they were most likely instances of prefixes hijacks affecting French ASes.

1.3 The Use of Route Objects

Best practices emphasize that a route object must be declared by an AS for each prefix that it advertises on the Internet. This indicator focuses on the analysis of the two sets illustrated in Figure 1.4: in blue, the route objects declared and in red, the prefixes advertised in BGP. By comparing them, it is possible to highlight the following three sub-indicators:

1. the route objects for which no prefix is advertised;
2. the prefixes having at least one associated route object;
3. the prefixes not covered by a route object.

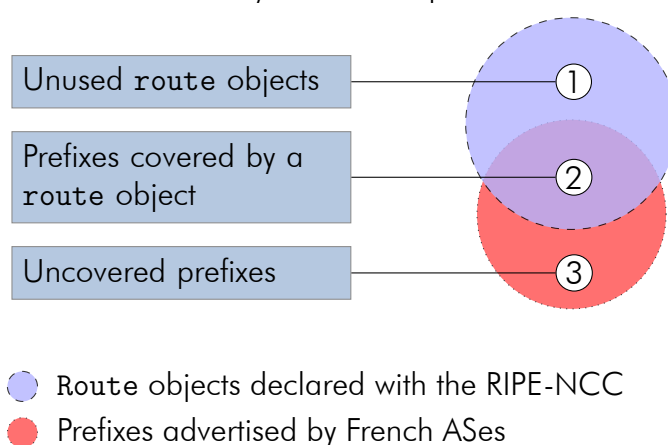


Figure 1.4: Representation of the sub-indicators of the use of route objects

Unused Route Objects

Declared route objects must match the prefixes advertised by an AS. The analysis here concerns the remaining orphan route objects, i.e. for which no prefix has been advertised during the year. The trend remained constant in comparison with previous years. During 2015, 137 orphan route objects were added to the 1,556 present on 1 January 2015. For IPv6, the quantity of orphan route6 objects increased from 182 to 232.

In a nutshell

The removal of unused route objects and route6 objects is not systematic and remains very marginal compared with the new declarations.

	Type	1 January	31 December
IPv4	no declared route object	653	575
	no route object used	91	103
IPv6	no declared route6 object	1,256	1,195
	no route6 object used	108	132

Table 1.1: Distribution of ASes according to the use of route objects in 2015

In order to analyze the ASes in terms of unused route objects, the Observatory classifies them into two categories, represented by 1.1:

1. ASes with no declared route object;
2. ASes using no declared route object.

The quantities of ASes with no declared route object in 2015 decreased by 78 and 61 ASes for IPv4 and IPv6 respectively. 575 ASes under IPv4 and 1,195 under IPv6 were in this category on 31 December. Despite this encouraging number, the number of ASes using none of the route object they declared increased by 12 under IPv4 and 24 under IPv6, ending at 103 under IPv4 and 132 under IPv6. It is important to note that most of these ASes have stopped advertising prefixes on the Internet: it is therefore not a question of new operators not respecting best practices with respect to the RIPE-NCC.

Prefixes Covered by Route Objects

A BGP interconnection can be filtered. Most of the time, this filtering is done via a list based on declared route objects. In this case we focused on highlighting the prefixes and ASes that would be covered by filters such as these.

The coverage of prefixes by route objects has been improving since 2011. IPv4 prefixes covered by route objects increased from 4,211 to 4,709 in 2016. For IPv6, 69 prefixes were added to the 358 present at the beginning of the year.

To get an idea of the ASes reachability, those in which all of the prefixes were covered by route objects were studied. For IPv4, 726 ASes fell into this category at 1 January, and reached 805 on 31 December, i.e. an improvement in accessibility for 79 ASes. Note that among the ASes, 77 were ASes created during the year. The new ASes strongly tend to apply best practices. For IPv6, the situation went from 212 to 249 ASes for which all of their prefixes were covered by route6 objects. In this set of ASes, 15 were created in 2015.



Prefixes not Covered by Route Objects

As with the previous sub-indicator, the situation is continuing to improve. From 841 IPv4 prefixes not covered at the start of 2015, the year ended with 785 prefixes not covered. For IPv6, the number of prefixes not covered increased from 121 to 125 in 2015. This means the situation is worsening for IPv6, and shows the need to strengthen efforts on this protocol.

Finally, we consider here the ASes for which at least one of the advertised prefixes was not covered by a route object. Under IPv4, 171 ASes were affected at the beginning of the year. On 31 December, there were 163 ASes in this situation. For IPv6, the overall situation has changed very little: the number of ASes missing at least one route6 object decreased from 51 to 48 in 2015.

In a nutshell

Unlike previous years, the number of IPv6 prefixes not covered by route6 objects increased.

1.4 Declarations in the RPKI

Changes in the Coverage of the Address Space

The study of the declarations made in the RIPE-NCC repository of the RPKI shows that the number of ASes participating grew in 2015. At the start of January, 198 French ASes had ROAs in the RPKI. On 31 December, the RPKI contained declarations from 237 French ASes, which represents an increase of almost 20%. Note that this increase is much lower than that observed during the previous year. In 2014, the number of ASes participating in the RPKI increased by nearly 80% between the beginning of January and the end of December.

In order to characterize the effects of this increase, the development of the coverage of the IPv4 address space managed by the French ASes during 2015 was studied. The percentage of valid address space changed little during the year. At December 31, 2015, approximately 65% of the address space was valid according to the RPKI.

In parallel, the percentage of uncovered address space and the invalid address space remained stable during the year. At December 31, 2015, 34.4% of the address space was not covered. The percentage of invalid address space remained relatively low during the year. At the end of 2015, the percentage was 0.4%.

As for IPv6, the address space managed by the French ASes was covered very little by the declarations of the RPKI. At the end of 2015, the ROAs covered less than 1% of the address space. There has therefore been no significant change in this coverage since 2014.

Validity of Advertisements Made by French ASes

In order to have an overview of the potential impact on the connectivity of strict filtering based on data from the RPKI, the study also focused on the number of ASes only issuing valid or invalid advertisements.

The number of ASes performing only valid prefix advertisements increased during 2015. From 102 in January 2015, the number rose to 118 at the end of the year. Furthermore, only one AS made invalid prefix advertisements during the month of December 2015.

These results show that in case of strict filtering based on the RPKI, nearly 12% of the active ASes at the end of 2015 would see all of their prefixes propagated over the Internet. This is comparable to the value observed in 2014. Furthermore, in the case of filtering only invalid advertisements, the address space managed by a single AS was longer reachable during the month of December.

Potential Use of the RPKI by French ASes

The analyses performed on the data in the RIPE-NCC repository do not measure the actual use by French ASes of the RPKI. For example, these data do not provide information on the use of ROAs for filtering purposes. However, a study of the change in the consistency of declarations compared with actual advertisements provides clues as to the maintenance over time of the ROAs in the RIPE-NCC repository.

	Type	january	december
Number of ASes (IPv4)	No ROA used	7	8
	Some ROAs used	20	33
	All ROAs used	161	181
Number of ASes (IPv6)	No ROA used	6	11
	Some ROAs used	2	4
	All ROAs used	58	69

Table 1.2: Change in the use of ROAs

Table 1.2 shows the results of the study of the potential use of declarations made by French ASes in the RPKI. Under IPv4 as IPv6, note that a large part of the ASes used all of their ROAs. However, it should also be noted that despite the low and recent adoption of the RPKI, there already are ASes that do not use any of their ROAs.

In a nutshell

At the end of 2015, the general conclusion remains the same as at the end of 2014: the declarations made in the RPKI are far from being exhaustive. As a result, about a third of the IPv4 address space is not covered. For IPv6, the coverage is still very low.

Chapter 2

Resilience in terms of the DNS protocol

2.1 Introduction

The Domain Name System, managed by the DNS protocol [18, 19], is a distributed, hierarchical naming system whose main objective is to associate an IP address with a name readable by users. For example, the name `www.afnic.fr` retrieves the IP address `192.134.5.5`. In the case of a change of hosting service provider, only the domain manager needs to change the IP address pointed to by the name. Thanks to the DNS, the change is therefore transparent for users.

The structure of the DNS is illustrated in Figure 2.1. At the top of the hierarchy is the root, represented by a dot “.”. This is the dot that can be found in domain names such as “`www.afnic.fr.`”. The names just below the root, like `.fr`, are called TLD.

At each level of the hierarchy there are one or more nodes in the DNS tree. The tree structure under a given node is called a domain. In turn, it may have sub-domains, and so on. This report does not take into account the subtle difference between a domain and a zone. In it, the two terms are therefore used interchangeably.

A zone can be “delegated” to entrust the management of its data to an organization different from that which administers the parent zone. For example, the `.fr` zone has been delegated to Afnic, which stipulates the rules for assigning domain names under

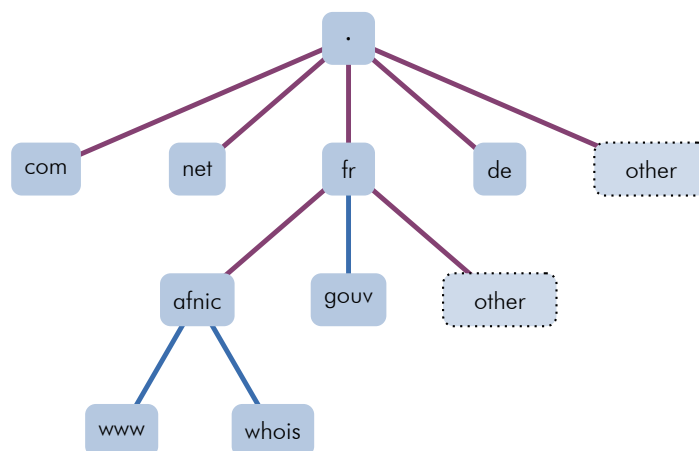



Figure 2.1: DNS structure



the `.fr` zone, regardless of its parent zone, the root, managed by ICANN¹. Delegations are illustrated by the purple links in Figure 2.1.

The information stored in the DNS

The resources attached to a zone are described by DNS records. Each DNS record has a domain name that stems from that of the zone (e.g. the name `www.afnic.fr` under the zone `afnic.fr`), a type, and data which depend on the type in question.

The different types of DNS records are published and maintained by IANA² in a registry dedicated to the DNS parameters [20]. The following record types are studied in this report:

- **A**: an IPv4 address;
- **AAAA**: an IPv6 address;
- **MX**: the name of an inbound mail relay;
- **NS**: the name of a DNS server;
- **Delegation Signer** and **DNSKEY**: useful cryptographic information for DNSSEC.

Querying the DNS

DNS resolution is the mechanism that retrieves records associated with a given domain name and type. This resolution mechanism involves two types of DNS servers, as shown in Figure 2.2, which highlights numbered interactions:

- **a recursive server** (also called a cache server or resolver). The user's machine knows it and submits its DNS queries to it (interaction 1). This server, usually managed by an ISP³, queries the DNS tree starting from the root (interaction 2) and following the delegation points one after another to the authoritative server for the domain name concerned by the query (interactions 3-4). Finally, the recursive server responds to the user's computer (interaction 5) and memorizes (cache function) the information received;
- **authoritative servers** for given zones, which reply to the recursive server. Either they are authoritative for the domain name requested by the recursive server, and return the answer to it; or they refer the query to other servers that are more likely to be authoritative for the domain name in question.

Domain name bailiwick

The NS and MX types of DNS records contain server names, as shown in Figure 2.3. The name on the left of the record type, in these cases `ssi.gouv.fr`, is the location in the DNS tree for this record. The name on the right of the type, for instance in this

¹Internet Corporation for Assigned Names and Numbers.

²Internet Assigned Numbers Authority.

³Internet Service Provider.

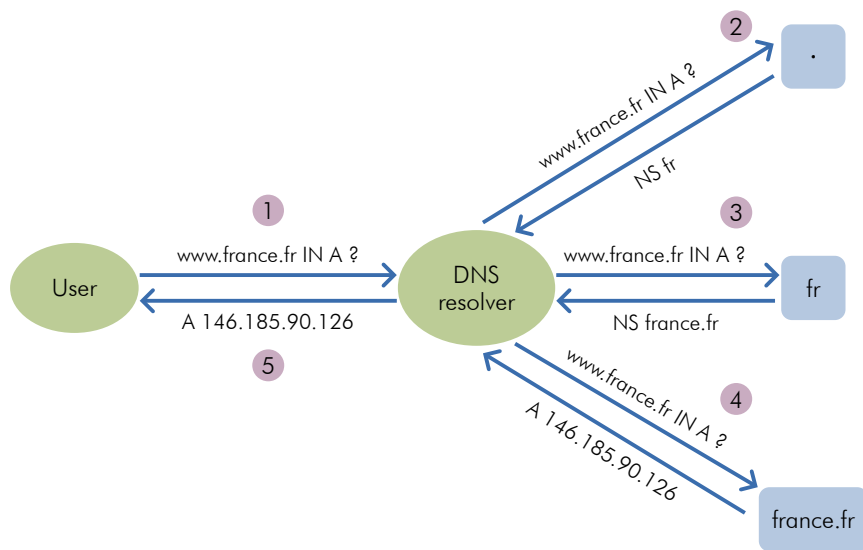


Figure 2.2: Example of DNS resolution

case `dns1.ssi.gouv.fr`, is the data. Generally, the names on the right must first be resolved into IP addresses by a recursive server to meet the expectations of a DNS user.

<code>ssi.gouv.fr.</code>	NS	<code>ns6.gandi.net.</code>
<code>ssi.gouv.fr.</code>	NS	<code>dns1.ssi.gouv.fr.</code>
<code>ssi.gouv.fr.</code>	MX	<code>10 smtp.ssi.gouv.fr.</code>


Figure 2.3: Examples of NS and MX records

In a DNS record when the name on the left of the type is included in the name on the right, the name on the right is said to be “in-bailiwick”. For example, `dns1.ssi.gouv.fr` is in-bailiwick of `ssi.gouv.fr`. The server sending such an NS or MX record may also, in addition, respond with the IP address corresponding to the name. Sometimes this is even required by the protocol.

Conversely, server names can be located in a third-party domain. They are then out-of-bailiwick. This is the case of the NS record using `ns6.gandi.net` to delegate the domain name `ssi.gouv.fr` in Figure 2.3.

Out-of-bailiwick names can introduce dependence to the proper functioning of a third-party domain. This is because each one is a SPOF⁴, if its unavailability may make it impossible to resolve a name into an IP address. The notion of dependency degrees quantifies the number of SPOF. Some third-parties are essential and are therefore excluded from this count. These include the parent zones of a name, such as, for example, the root or `.fr` for the name `france.fr`.

⁴Single Point of Failure.



This means `ssi.gouv.fr` would have a single degree of dependence if the domain was delegated to a DNS server in the `example.com` domain and another server in the `example2.com` domain. This is because the failure of `example.com` could be offset by the availability of `example2.com` and vice versa. The only SPOF would be `.com`. Similarly, `ssi.gouv.fr` would have a degree of dependence of two, if it was only delegated with server names in the `example.com` domain. This is because the proper functioning of the servers of both actors would be required: those of `.com` and those of `example.com`.

The risk of an unavailable TLD is not theoretical. For example, in December 2015, the `.tr` TLD suffered a DDoS⁵ attack for three weeks [21], with periods during which all of its DNS servers were unreachable.

Public domain names

In this report, the term “public domain name” means domain names delegated from one of the domains in a list, called PSL⁶ [22]. This comes from Mozilla’s initiative to strengthen the isolation of websites in browsers. This list, while not directly related to the DNS, is used to reference the domains managed by entities acting like registries. This information is not systematically visible in the DNS. This is because some registries do not always operate domain names consisting of a single label. This is the case of the Nominet registry, responsible for `.co.uk`.

Record security

Designed at a time when the threats were less pervasive, the initial DNS did not integrate advanced security mechanisms. The DNSSEC protocol has been designed to remedy this shortcoming [23]. It is used to ensure the authenticity and integrity of the data, based on asymmetric cryptographic mechanisms. The public keys and signatures are respectively stored in `DNSKEY` and `RRSIG` records. The DNSSEC chain of trust is established and maintained through `DS` records. This mechanism prevents attacks referred to as “cache poisoning” that attempt to inject fraudulent records into a recursive server.


2.1.1 Data and tools

The observatory uses *ad hoc* scripts to carry out active measurements of DNS zones. The `dnspython` library [24] is used in particular for this purpose.

The authoritative servers for the delegated zones of the `.fr` TLD are directly queried by the scripts. When multiple authoritative servers exist for the same domain name,

⁵Distributed Denial of Service.

⁶Public Suffix List.



the queried server is chosen randomly in order to limit the load imposed by the measurement campaign. The distribution on which the random draw is performed is not uniform, however. Servers have a probability of being selected that is inversely proportional to the number of zones studied that they host.

To summarize, the more zones a server hosts, the more its probability of being chosen to resolve a domain is low. If it returns an error, or does not respond within the time limit, the scripts operate a fallback action. They then use a recursive server that applies its usual resolution algorithm.

Domain names that have not responded to our queries are not counted in the statistics. They constituted a statistical bias of almost 3% of the delegated domain names in the .fr zone in December 2015, or about 77,500 domain names. In December 2014, this bias was 2.5%, or 66,000 domain names.

Data used

All of the active measurements were made using the .fr zone, which varies with the creations, deletions and modifications of delegated zones. During the analysis, only the zones for which all of the measurements have been carried out without fail are taken into account. They are called the “studied zones”. For example, from 2014 to 2015, the number of these zones increased by 7% to around 2,810,000 on December 7, 2015, against approximately 2,630,000 on December 31, 2014. This is due to the creation of 617,000 new zones, the removal of 438,000 zones and an increase in the number of failed zones during the measurements.

The decision to use the .fr zone as a data source introduced a bias with respect to the representativeness of the Internet in France. This is because domain name registration in the .fr zone is not limited only to people registered within French territory. Furthermore, other TLDs exist in France, including geographical names, such as .re or generic names such as .paris. Finally, in the domain name market in France, foreign TLDs are often used such as .com or .net.

It should be noted that the data used are authoritative. This means that the DNS records used come from the authoritative servers that host them. For example, the DS⁷ records used for the DNSSEC indicator are extracted from the .fr zone, while the NS, MX, A, and AAAA records are resolved in accordance with the method described on page 28.

The list of public suffixes used in this report was downloaded on December 2, 2015 [22]. The notion of second-level domain names, used in the 2014 report, was replaced by that of public domain names. The creation dates of the zones studied were obtained by the analysis of Afnic public data, published as part of the Open Data initiative [25].

⁷Delegation Signer.

2.2 Dispersion of authoritative DNS servers

Number of servers per delegated zone

The number of NS records per zone studied remained roughly equivalent to that reported in 2014. Thus, around 70% of the zones are hosted on two servers against three servers for about 18% of the zones. The minor fluctuations observed are mainly due to market dynamics and the creation of new names hosted on service platforms with various numbers of DNS servers.

The number of NS records per delegated zone remains sufficient to allow good resilience, from the perspective of this indicator. In fact, less than 1% of the zones use a single NS record, which could therefore be a SPOF.

The same study can be performed once the hostnames contained in the NS records have been resolved. For IPv4 as IPv6, the distribution is substantially equivalent to that observed when only taking NS records into account. Consequently, for the zones studied, the use of multiple IP addresses from the same version of the IP protocol for the same NS record, can be regarded as an anecdotal practice.

It should however be noted that almost 41% of zones have no server with an IPv6 address. The zones in this case are based solely on the availability of IPv4 servers, even if the recursive servers querying them simultaneously had an IPv4 and IPv6 connectivity. The proportion in 2015 is similar to that reported by the observatory in 2014.

In a nutshell

The number of DNS servers per zone seems sufficient to ensure good resilience. IPv6 deployment on authoritative DNS servers stagnated in 2015. Thus, about 41% of the zones are reachable only over IPv4.

Topological dispersion of delegated zones

The topological dispersion of DNS servers is a requirement resulting from resilience engineering [26, 27]. The dispersion of name servers in separate ASes may, in some cases, help to prevent downtime in case of incidents affecting the whole of an operator network. In 2015, the average number of ASes per zone remained the same as in 2014, and stagnated at 1.2. Furthermore, the quantity of zones hosted by a single AS has remained stable since 2011, peaking at 83% of the studied zones.

The dispersion of the DNS servers in different ASes is therefore still very low. Based on this fact, however, it is difficult to draw a direct conclusion on the potential impact

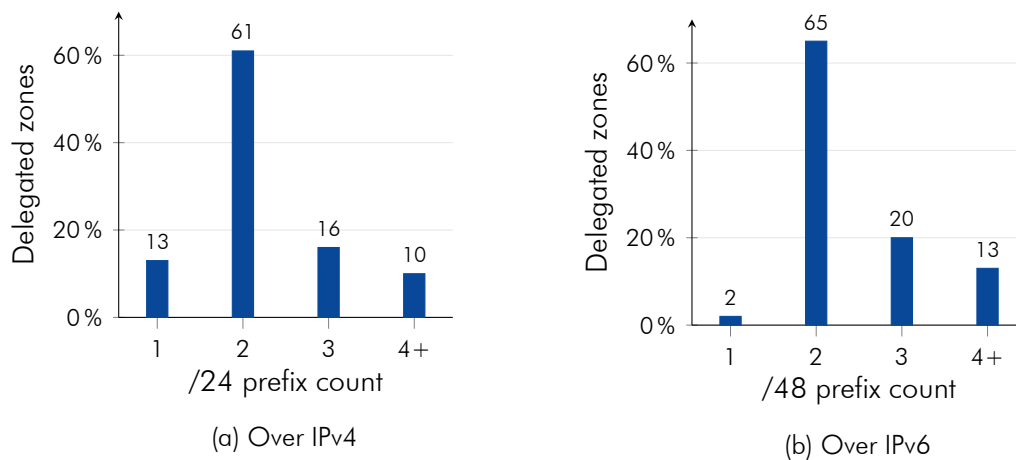


Figure 2.4: Dispersion of zones by number of prefixes in December 2015

on availability. A more detailed analysis of operator networks would be required to estimate the likelihood of a single incident affecting their entire network.

In a nutshell

Virtually all of the zones have at least two DNS servers, but they are generally located in a single AS.

Distributing or being able to distribute name servers within separate /24 IPv4 prefixes and /48 IPv6 prefixes can also be a good resilience practice⁸. In particular, this makes it possible to limit BGP pollution zones and to increase agility in the case of distributed denials of service.

The study of the distribution of name servers in the prefixes is performed only on the zones hosted on at least two IP addresses. Over IPv4, Figure 2.4a shows that 87 % of the zones can be or are advertised in distinct /24 prefixes. Over IPv6, this figure rises to 98 % for separate /48 prefixes, as detailed in Figure 2.4b.

In a nutshell

Best current resilience practices with respect to the diversity of prefixes and BGP hijacking resistance are applied to virtually all of the zones studied.

⁸The sizes of /24 prefixes under IPv4 [28] and /48 prefixes under IPv6 [29] are the longest that can be advertised on the Internet, according to BGP best current practices.

Dispersion of authoritative DNS servers per country

The geographic dispersion of authoritative DNS servers can also have an impact on the availability of French Internet services. For example, this may be the case following the rupture of submarine cables isolating users from the DNS servers indicating the IP address to connect to, in order to access a service.

Using the Maxmind GeoLite database [30] downloaded in December 2015, it is possible to estimate the geolocation of the authoritative servers in the zones studied. This practice is however of limited value if the anycast routing technique, presented in the 2013 report, is used. This is because in this case, the IP geolocated in a country will be advertised with BGP for several locations worldwide. This sub-indicator nevertheless provides a first approximation of the location of these servers.

As in previous years, in 2015, both over IPv4 and IPv6, nearly 82 % of the zones studied are served exclusively by authoritative servers located in the same country.

Over IPv4, the zones for which all of the authoritative DNS servers are located in the same foreign country accounted for 27 % of the zones studied. It should nevertheless be noted that 75 % of them are hosted in a country sharing a terrestrial border with mainland France. Therefore, this situation may not represent a significant risk. For nearly 20 % of these zones, the situation is more mixed, the latter being hosted in North America.

For IPv6, the zones hosted in a foreign country only represent 30 % of zones with DNS servers accessible over IPv6. Of these, nearly 85 % are hosted in a country sharing a terrestrial border with mainland France.


In a nutshell

The geographical dispersion of the zones studied is satisfactory, both over IPv4 and IPv6. Nevertheless, almost 20 % of zones are served over IPv4, exclusively from North America.

Dependence on third-party names

The analysis of NS records reveals that 99 % of the zones are delegated exclusively using out-of-bailiwick hostnames.

The risk of downtime caused by the dependence degree is not a theoretical risk. For example, in 2015, the `tools.ietf.org` site was down for several hours. All of the authoritative DNS servers on this zone were designated by hostnames in the third-party



domain `levkowetz.com`. When this third-party domain suffered an availability incident, `tools.ietf.org` became unreachable in turn.

Observatory data shows that 89% of the delegated zones studied had at least one degree of dependence: all of the NS records use a name located in the same TLD distinct from the `.fr`, such as `.net`.

Similarly, 75% of the zones studied have two degrees of dependence, due to the use of a single public out-of-bailiwick domain name in a third-party TLD, such as `tools.ietf.org` that was dependent on `levkowetz.com`.

In a nutshell

In 75% of the zones studied, there is an increased risk of downtime due to the names of the servers chosen to delegate those zones.

2.3 Implementing DNSSEC

Analysis of DS records

For this report, the enumeration of DS records is based on the data contained in the `.fr` zone after filtering to keep only the zones where the measurements for the rest of the report were carried out without failure. This new methodology is expected to improve the reproducibility of results by only using public information. In previous years, certain DS records corresponded in fact to tests, or to zones not published by Afnic.

The evolution of zone count with at least one DS record under the new methodology was observed using the zones studied on December 14, 2014 and December 6, 2015.

With these new data, between December 2014 and December 2015, the percentage of zones studied with DS records rose from 6.4 % to 8.8 %. For simplicity, the term DNSSEC zone is used to designate these zones. In December 2015 it was possible to count about 248,000 zones in this case, against 180,000 in December 2014.

To determine the origin of this change, it is interesting to detail the growth of the `.fr` zone. For example, 617,000 zones were created between December 2014 and December 2015. This represents 22 % of the zones studied at year-end 2015. In addition, during the same period, about 438,000 zones were deleted, or 16 % of the zones studied in December 2014. Of these 438,000 zones, 173,000 were created for a period of one year and were not renewed.

Comparatively, the number of zones implementing DNSSEC consists of approximately 98,000 newly registered zones, or 40 % of the DNSSEC zones in December 2015. In addition, approximately 54,000 zones were deleted from the `.fr` zone, or 30 % of the DNSSEC zones in December 2014. Of the 54,000 zones deleted, about 36,000 had been created in 2014.

In addition to the growth dynamic of DNSSEC zones, it is interesting to note that approximately 29,000 zones already registered in 2014 implemented DNSSEC in 2015. In addition, about 6,000 zones signed in 2014 disabled DNSSEC during the year. These zones represented 3 % of DNSSEC zones in 2014.

DNSSEC growth is mainly due to the creation of zones. For example, only 1 % of the zones studied, which already existed in 2014, implemented DNSSEC in 2015.

In a nutshell

Just under 9 % of the zones studied implemented DNSSEC in 2015. The growth is mainly due to the creation of new zones in 2015.

Analysis of cryptographic algorithms

Almost 92 % of the zones studied and with a DS record indicated they use the cryptographic suite RSASHA1-NSEC3-SHA1. Almost all of the remaining 8 % use the RSASHA256 suite. It should be noted that the use of the SHA-1 hash algorithm as used in the cryptographic suite RSASHA1-NSEC3-SHA1, is contrary to cryptographic best current practices [31], and to the recommendations of the RGS⁹ [32].

The analysis of the algorithms used to hash the DNSSEC keys of the zones studied contrasts with the previous result. Indeed, SHA-256 is used by almost 98 % of zones with a DS record to create the DNSSEC chain of trust. The remaining 2 % are using SHA-1.

In a nutshell

Nearly 92 % of the zones studied implement the SHA-1 hash algorithm, considered insufficient with respect to the cryptographic best current practices. Conversely, nearly 98 % of the zones studied use the recommended SHA-256 hash algorithm to create the DNSSEC chain of trust.

⁹Référentiel Général de Sécurité.

2.4 Dispersion of inbound e-mail relays

Number of relays per delegated zone

The observatory changed its methodology for this indicator. In 2014, the percentages were based on the number of zones having at least one MX record. Yet, some zones have no inbound e-mail relay, which produced a statistical bias. In 2015, about 9 % of the zones studied had no MX record.

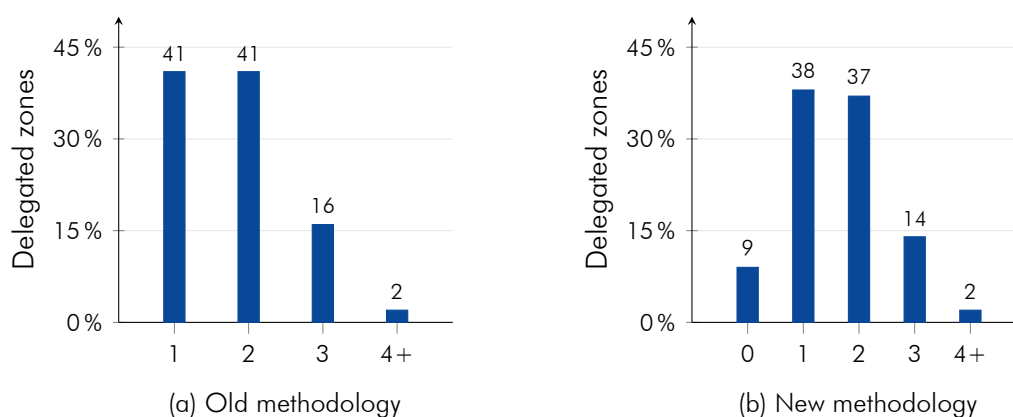


Figure 2.5: Number of inbound e-mail relays per domain in 2015


In 2015, the zones studied had a low number of inbound email relays, on average. Thus, a single relay was available for 38 % of the delegated zones, as shown in Figure 2.5b. In the case of an incident, the failure of this single relay then causes the total unavailability of the service. If it remains short, the impact is however limited. The delivery protocol for electronic mail (SMTP) is, indeed, itself designed to be resilient.

The 2014 figures are not included in Figure 2.5 as they were identical to those of 2015.

In a nutshell

A single mail relay is filled for 38 % of the delegated zones. The risk of being unable to receive new e-mail is therefore increased by the presence of a SPOF.

A single MX record, however, does not indicate the presence of a SPOF in the addressing plan. This is because several IP addresses can be specified for a single domain name or an IP address may be advertised on the Internet using the anycast routing technique. Similarly, an IP address can be virtual and distributed across multiple real servers using load balancers.



Consequently, only the case in which several IP addresses are associated with a name can be analyzed through the DNS. To do so, the hostnames listed in the MX records have been resolved into IPv4 and IPv6 addresses.

The number of IPv4 addresses generally proves to be identical to the number of MX records. There is nevertheless a special case for almost 14 % of the zones. These zones are wholly or partly hosted on a certain service platform. It provides a secondary e-mail relay¹⁰ designated by a single MX record. The hostname contained in this MX record is resolved in five separate IPv4 addresses. The zones using this service thus have a greater resilience since they have at least six IPv4 addresses in total.

In a nutshell

Analysis of the IPv4 addresses of inbound e-mail relays shows a slightly better situation than analyzing only the hostnames in the MX records. Regarding the e-mail relays, almost 38 % of the zones are still hosted on a single IPv4 address.

In contrast, on nearly 89 % of the zones, none of the MX record hostnames have an associated IPv6 address. Among the remaining 11 %, for 81 % of these zones, only one hostname can be resolved into a single IPv6.

It is worth noting that a service platform has an atypical behavior. It represents 10 % of the zones whose relays have IPv6 addresses. When the hostnames of this platform are resolved, a single IPv6 address is returned. It is distinct but constant depending on the authoritative server queried. The selection of the IPv6 address of the mail relay is therefore dependent on the recursive server and its selection algorithm of the authoritative server to query. The observatory has chosen to aggregate all of the IP addresses returned as if it were a single AAAA resource record set. The size of this resource record set varies from 14 to 18 IPv6 addresses per hostname.

In a nutshell

IPv6 is still little deployed. Indeed, for nearly 89 % of the zones, no hostnames contained in the MX records have an associated IPv6 address. Among the remaining 11 %, for 81 % of these zones, only one hostname can be resolved into a single IPv6.

¹⁰In particular, a secondary relay may store messages if the primary servers are unavailable. The messages are then sent to the primary servers, when they are reachable once again.

Analysis of hostnames for inbound relays

The study of the inbound relay hostnames found in MX records can be used to compute their degrees of dependency, using the methodology presented on page 27. A statistical bias also slipped into this sub-indicator in 2014. This was due to the percentages calculated from the entire zone instead of counting only the domains with MX records. The figures for 2014 and 2015 are nonetheless similar.

In 2015, 86 % of the zones studied had one or more degrees of dependence for their e-mail relays. The remaining 14 % had at least one e-mail relay designated by an in-bailiwick name.

Approximately 1,700,000 zones studied exclusively use hostnames located in another TLD than .fr. For 99 % of these zones, one degree of dependence is introduced because all of the relays are designated by hostnames under a single TLD. For 69 % of them, it is the .net TLD and for 23 %, the TLD is .com. Both TLDs are the responsibility of the same US organization and are hosted on the same set of DNS servers [33].

A degree of dependence is also introduced for the 19 % of zones using hostnames in a single domain out-of-bailiwick but delegated under the .fr TLD. Finally, 64 % of the zones studied are afflicted with two degrees of dependence, since their relays are exclusively designated with hostnames located in a unique public name in a third-party TLD.

In a nutshell

Up to 86 % of the zones studied introduce at least one SPOF due to the choice of hostnames for their e-mail relays. In particular, 64 % of the zones studied create two SPOFs using hostnames located in a public name under a single third-party TLD.

Concentration of inbound e-mail relays

The analysis of public names can also be used to measure the concentration of e-mail relays on some shared hosting platforms.

For the following results, only the label furthest to the left of the public name is considered, regardless of the name of the registry. It is thus possible to group together certain platforms that are diversified across multiple registries or TLDs. For example, e-mail relays located in the sub-domains `1and1.com` and `1and1.co.uk` will be considered to be hosted by `1and1`. To better assess the risks, only the zones for which all of the

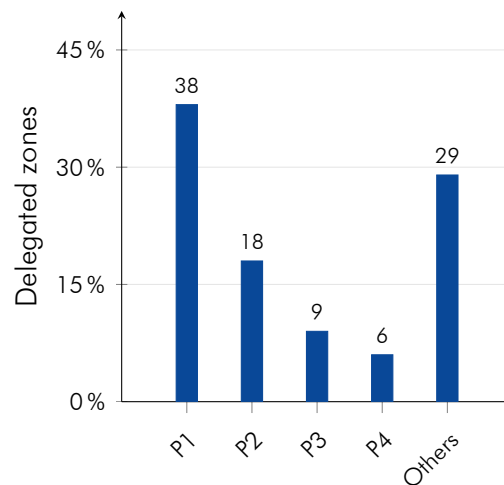


Figure 2.6: Concentration of e-mail relays on hosting platforms in 2015

relays are hosted by a single platform are counted. This set consists of approximately 1,800,000 zones.

As shown in Figure 2.6, e-mail relays are very highly concentrated on a handful of hosting platforms. In particular, 71 % of the e-mail relays are hosted by four operators.

This concentration can sometimes help the pooling of resources. This may be useful in defending against certain denial of service attacks. Spam filtering can also be shared. It should be noted, however, that there is a risk of collective failure in the case of incidents affecting the pooled components.

Dispersion of e-mail relays per country

The distribution of e-mail relays in several countries, in the same way as DNS servers, may be a factor affecting availability. In particular, ensuring connectivity with users who may send e-mail is necessary.

Over IPv4, for the zones studied that have MX records, 99 % of them all use e-mail relays located in the same country. This represents approximately 2,540,000 zones. Over IPv4, 69 % of these zones have their relays located in France. As illustrated in Figure 2.7, the only country outside Europe containing a significant number of e-mail relays is the United States of America, with 6 % of the zones concerned.

Over IPv6, it should be noted that almost 89 % of zones have no IPv6 e-mail relay. Only 248,000 zones are therefore concerned by this study. Thus, 85 % of these zones use e-mail relays all located in the same country. For 221,000 zones, this country is France. Almost all of the remaining zones are located in Europe.

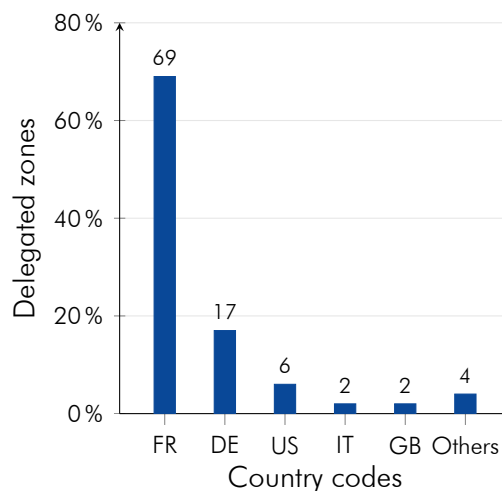


Figure 2.7: Geolocation of relays hosted in a single country, over IPv4, in 2015

In a nutshell

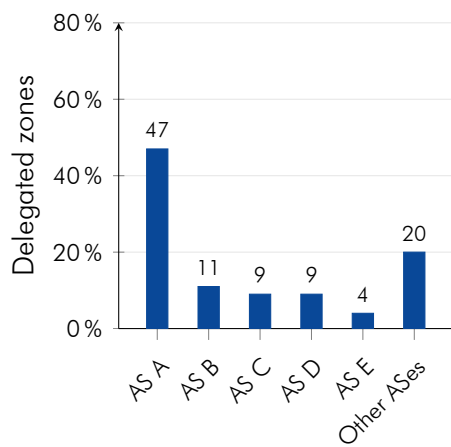
Most of the e-mail relays for the zones studied are located in France, both for IPv4 and IPv6.

Dispersion of the e-mail relay network

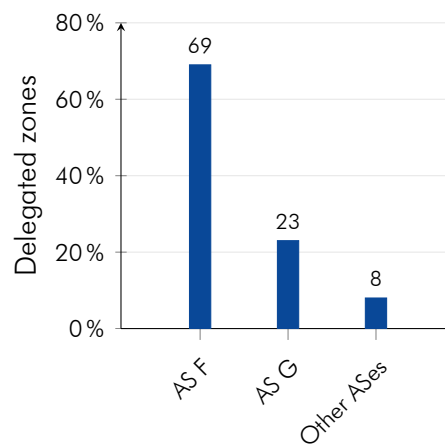
This sub-indicator examines the distribution of e-mail relays on one or more operators identified by their AS number.

Over IPv4, in 96 % of the cases, all of the e-mail relays in a zone are hosted in a single AS. The rest of the zones have their relays hosted almost entirely in two ASes. Over IPv6, the numbers are similar, with 98 % of the zones having all of their relays in the same AS. These numbers indicate a low diversity of hosting operators. These findings corroborate those of the sub-indicator using public domain names, presented on page 38.

By analyzing the number of e-mail relays by AS number, it is possible to draw a portrait of this concentration, from the routing point of view. Thus, over IPv4, 47 % of the e-mail relays are concentrated within a single AS. In addition, four hosting operators alone account for 80 % of the e-mail relays of the zones studied. This distribution is shown in Figure 2.8a. Over IPv6, the concentration is even higher, since 69 % of the relays are hosted by the same player. The second player hosts 23 % of them, as shown in Figure 2.8b.



(a) IPv4



(b) IPv6

Figure 2.8: Distribution by AS of relays hosted in a single country in 2015

In a nutshell

The concentration of inbound e-mail relays on a few network operators is significant. Over IPv4, one operator is responsible for the connectivity of 47 % of the relays of the zones studied. Over IPv6, this figure rises to 69 %.

Chapter 3

Resilience in terms of the TLS protocol

3.1 Introduction

Setting up a TLS session between a client and a server ensures the integrity and confidentiality of communications, regardless of the nature of the underlying applications. Among the most common uses of the protocol is HTTPS, which consists in the protection of the HTTP data flows inside TLS tunnels.

The development of the TLS protocol followed several iterations [34, 35, 36] since the design of the SSL¹ protocol, which is obsolete [37]. For the sake of interoperability, the specifications allow both parties to negotiate the protocol version they will commonly adopt.

This parameter is set during a *TLS handshake* phase which precedes the actual encryption of the data. Similarly, the specifications allow the use of different combinations of cryptographic algorithms. The cipher suite chosen for the session is determined by messages of type *handshake*.

Figure 3.1 illustrates the negotiation of these parameters in a generic case. It does not substitute for more precise references [38, 39].

1. The client initiates a request by sending a message of type `ClientHello` which contains the cipher suites it supports;
2. the server responds with a `ServerHello` containing the adopted suite;
3. the server sends a `Certificate` message, which contains its public key in a digital certificate;
4. the server transmits in a `ServerKeyExchange` an ephemeral value that it signs with the private key associated with the previous public key;
5. the server indicates it is now waiting with a `ServerHelloDone`;
6. after checking the certificate and authenticating the previous value, the client in turn chooses an ephemeral value that it encrypts using the public key of the certificate and the transmits it in a `ClientKeyExchange`;
7. the client signals the adoption of the negotiated suite with a `ChangeCipherSpec`;

¹Secure Sockets Layer.

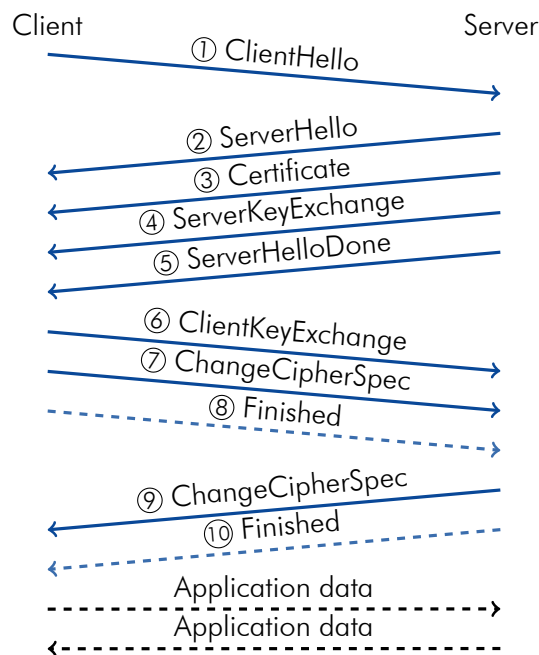


Figure 3.1: Generic initiation of a TLS session

8. the client sends a `Finished`, the first message protected by the cipher suite with secrets from the previous exchange of ephemeral keys;
9. the server signals the adoption of the same suite with a `ChangeCipherSpec`;
10. the server in turn sends a `Finished`, its first secure message.


The generic case described here implies the adoption of one of the cipher suites that verify perfect forward secrecy (PFS). This is to prevent the decryption of past sessions of messages even when the private key of the server is compromised, by negotiating an ephemeral secret using a Diffie–Hellman exchange [40].

The protocol specifications also define additional messages and extensions that make it possible to manage and enhance the protection of communications [41]. For instance, a client may advertise its elliptic curve cryptography capabilities through specific `ClientHello` extensions.

3.1.1 Public Key Infrastructures

The validity of the certificate sent by the server during the initiation of the TLS session is crucial to the security of the protocol. All the mechanisms and entities that formulate and maintain this validity represent a PKI.

Considering an X.509-compliant certificate, the assurance that the public key it contains actually belongs to the server advertised as a subject (usually in the form of a



domain name) is based on the transmission of trust from an authority already recognized through to the server in question. Successive links of trust established by CA² are materialized by cryptographic signatures backing the various certificates.

Thus, the `Certificate` message from which the key at the origin of the session secrets is extracted actually contains a chain of certificates. The client expects it to form a link from a trusted root to the polled server. In the case of web browsers, these roots usually correspond to a public registry such as the NSS certificate store maintained by Mozilla for its Firefox browser [42]. Some applications interact directly with the appropriate public registry in order to update their trusted roots, while others rely on maintenance performed by the host operating system.

The certificates contain several attributes, such as a public key and a validity period, usually supplemented by X.509v3 extensions. These include the ability to specify the context of use of the certificate in question and strengthen the assurances of the PKI. To that purpose, the ANSSI recommends following Annex A4 of the RGS [43].

3.1.2 Data and Tools

The measurements of the Observatory focused on the web resources accessible through the French Internet. They more specifically concerned the resources exposed through port 443, which is traditionally allocated by servers for HTTPS exchanges. The issues relating to online messaging resources differ in several aspects [44] and are not addressed in this report.

Domain names maintained by the Afnic³ were prefixed with `www.` before being resolved. For instance, the measurements on the `afnic.fr` domain match the IPv4 address resolved for `www.afnic.fr`. When the port 443 of the polled server was open, the Observatory sent various `ClientHello`, stimuli. The variations were designed to evaluate several capabilities of the server, such as its support for forward secrecy or its tolerance for outdated versions of the protocol.

The responses were dissected and inserted into the database using Parsifal [45]. The tools used were also available to check and if necessary reconstruct the chains of certificates observed. A closer examination of certain certificates made use of the X.509 support by Scapy [46].

Since the SNI⁴ [41] extension had not been used within the `ClientHello`s, the measurements did not meet all of the resources exposed via HTTPS on the `.fr` zone. Thus in polling one server per domain name resolved, a subset of 61,216 accessible servers were polled in July 2015, against 26,261 in February 2014.

²Certificate Authority.

³Association Française pour le Nommage Internet en Coopération.

⁴Server Name Indication.

3.2 Session Negotiation

Among the session parameters initially set by the specifications, some were recognized as being safe for use in 2015, while others have been declared obsolete for security reasons. The attributes established through the negotiation phase therefore hold a direct impact of the security of the subsequent exchanges. For this reason, the Observatory sought to establish a profile of the servers in the .fr zone exposing HTTPS resources.

As a single server subjected to the appropriate ClientHellos may accept the recommended version TLS 1.2 or the proscribed versions SSLv2 and SSLv3, the examination of a single parameter does not make it possible to judge the absolute security of a set of servers. By pooling these parameters and monitoring the evolution from 2014 to 2015, however, a qualitative declaration of compliance with best practices was possible.

General State of the TLS Servers

In 2015, considering all of the stimuli trying to negotiate a TLS 1.2 or TLS 1.0 session, with variations especially at the level of the suggested cipher suites, 80% of the servers with port 443 open enabled the initiation of a TLS session. Only 3.5% of servers sent no data whatever the stimulus used, and 0.4% responded with data never recognized as TLS messages.

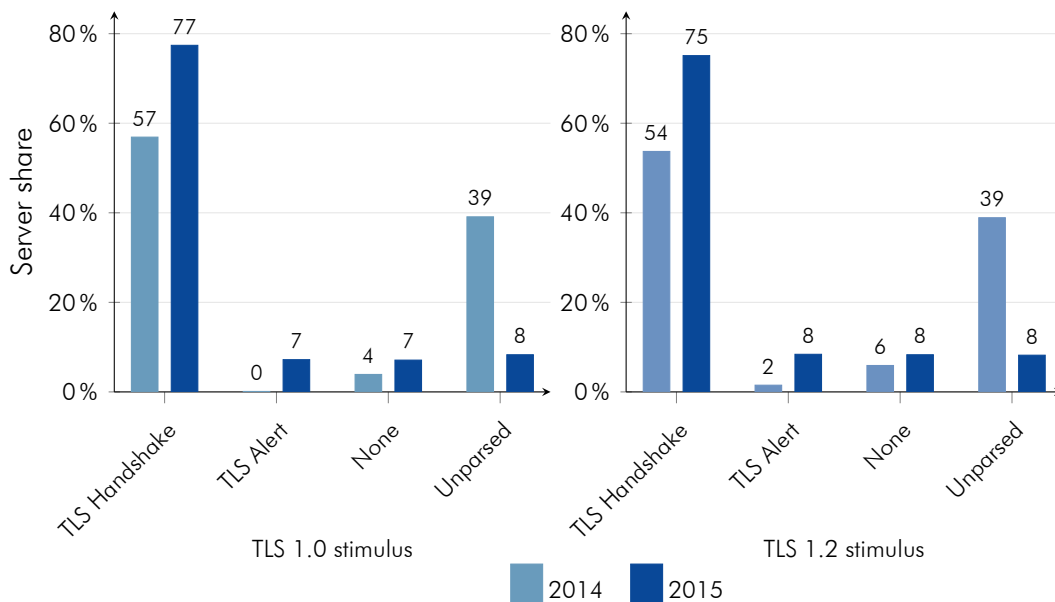



Figure 3.2: Developments in the responses of HTTPS servers to TLS 1.0 and 1.2 stimuli

Figure 3.2 specifies the success rate of session establishment during the main measurement campaign with TLS 1.2 in 2014 and then in 2015, in which the stimulus proposed



several cipher suites without any particular constraint. 75 % of the servers were able to negotiate a TLS 1.2 session, a significant increase from 54 % in 2014. The proportion of unacknowledged messages dropped from 39 % to 8 %. Manual review revealed several response profiles, including unencrypted HTML pages and SSH headers.

During the main campaign on TLS 1.0, the stimulus used still allowed the negotiation of a varied number of suites. The results obtained are also shown in Figure 3.2. The proportions and trends are similar to those observed during the previous campaign. For example, the adoption of version 1.2 is not synonymous with the disappearance of version 1.0. Although version 1.2 protects the communications against some attacks that affect version 1.0 [47], the servers generally maintain interoperability with dated clients.

In a nutshell

TLS 1.2 was frequently supported by servers in the .fr zone in 2015. TLS 1.0 was present to an equal degree, but clients and servers should favor the latest version of the protocol.

Forward Secrecy

The carrying out of DHE⁵ and ECDHE⁶ exchanges is conditioned by the choice of an appropriate cipher suite. Respect for forward secrecy can therefore be measured by counting the servers that accept the use of such suites.

Various measures were aggregated in order to identify the servers enabling forward secrecy. Figure 3.3 shows the proportion of separate IP addresses which enabled to negotiate suites with DHE or ECDHE, compared with all of the addresses in which an open port 443 was observed. Mid-2015, nearly three-quarters of the servers offered forward secrecy, an increase since early 2014.

This positive trend must however be tempered by the fact that tolerance to DHE or ECDHE does not guarantee that an ephemeral key exchange is preferred in every circumstance. In addition, for interoperability reasons, it is rare that such protection is demanded by a server.

⁵Diffie-Hellman Ephemeral.

⁶Elliptic Curve Diffie-Hellman Ephemeral.

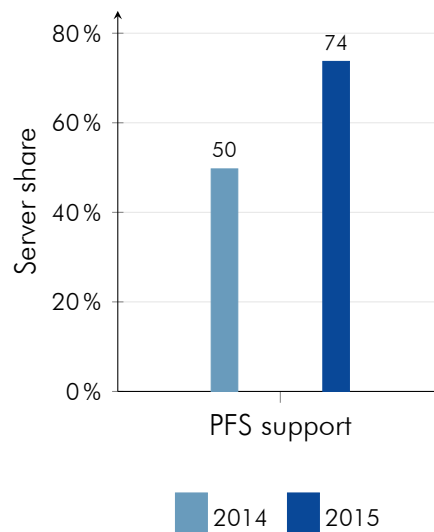


Figure 3.3: Changes in the support of PFS by HTTPS servers

In a nutshell

In 2015, the forward secrecy offer was widespread on the `.fr` zone. Servers should favor this approach with clients that support it.

Obsolescence of SSLv2

Since its initial publication in 1995 by Netscape, the security of SSLv2 has been subject to criticism which motivated the definition of SSLv3 one year later. Ultimately, a formal declaration of obsolescence was emitted by the IETF in 2011 [48]. In 2015, SSLv2 was disabled by default in all of modern web browsers, while its assessed danger continues to rise [49].

Figure 3.4 represents the responses of servers in the `.fr` zone that received an SSLv2 `ClientHello`. The message was ignored by 75% of the servers in early 2014, and 78% of the servers in 2015. Although there is an error code to reject the unsupported protocol versions, only 2% of the servers responded with a message of type `alert` in 2015.

Less than one server in five accepted to mount an SSLv2 session, although the decrease compared with 2014 remains low. Furthermore, adjusted to the larger number of servers in 2015, the development indicated that at least 5,000 new servers accepting SSLv2 were configured between the two measurements, despite its acknowledged dangerousness.

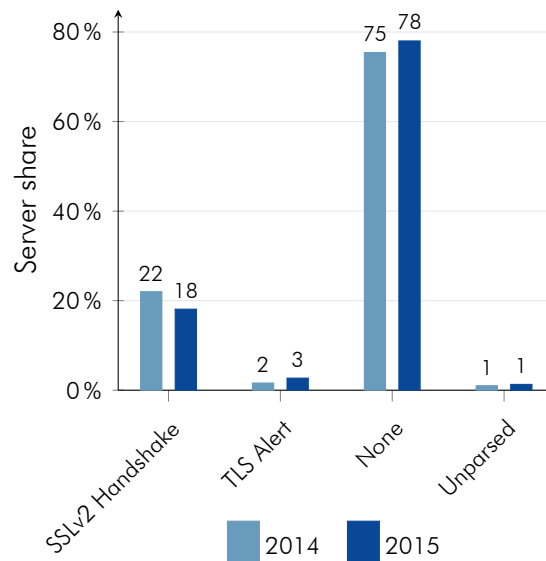


Figure 3.4: Responses from French servers to an SSLv2 stimulus

The Observatory did not measure the support of SSLv3. The use of this version is also to be avoided, as recent studies have reasserted [50].

In a nutshell

In 2015, many servers tolerating SSLv2 emerged. This protocol version should be abandoned as quickly as possible.

3.3 Robustness of Certificate Signatures

The use of robust signature algorithms for building certificate chains is essential for the security of exchanges with the servers involved in the PKI. These algorithms generally combine a method of asymmetric cryptography with a hash function such as SHA-1 or SHA-2. However several theoretical attacks against SHA-1 have been discovered since 2005 [51, 52]. Thus web browser editors chose to plan the obsolescence of this function within the PKI.

This way, in 2015, most browsers detecting the use of SHA-1 addressed a warning to the Internet users. Since early 2016, certificates using SHA-1 and valid at the earliest on January 1, 2016 are rejected. The rejection of all of the certificates signed using SHA-1, originally scheduled for January 1, 2017, could be advanced to early summer 2016 [53, 54, 55].

Changes In All of The Certificates

These considerations about the strength of the trust paradigm motivated the Observatory to study the profiles of certificates observable in the .fr zone. Because the SNI extension was not used, the measurements noted at most one terminal certificate per domain name resolved and per stimulus. Therefore, they are not exhaustive. However, they can be used to characterize significant trends between 2014 and 2015.

Certificates	February 2014	July 2015
Self-signed	5,324	15,712
Issued by a CA	7,554	22,759

Table 3.1: Number of separate certificates observed on the .fr zone

Table 3.1 reports the number of separate certificates noted during the two measurements campaigns launched on the .fr zone in 2014 and 2015. At nearly eighteen months apart, the proportion of self-signed certificates remained unchanged at 41 %. Out of the 15,712 representatives observed in 2015, 47 were recognized as trusted roots by the NSS certificate store.

Reliance on a self-signed certificate is arbitrary and not based on its signature. The quality of the PKI is therefore independent of the signature algorithm used in its trusted roots. For this reason, Figure 3.5 represents the change in the presence of different hashing algorithms between 2014 and 2015 while excluding the self-signed certificates, whether recognized or not by public trust stores.

In early 2014, the share of certificates signed using SHA-2 was only 3 %, while 96 % of them were signed with SHA-1. The residue of 1 % used MD5, an algorithm vulnerable

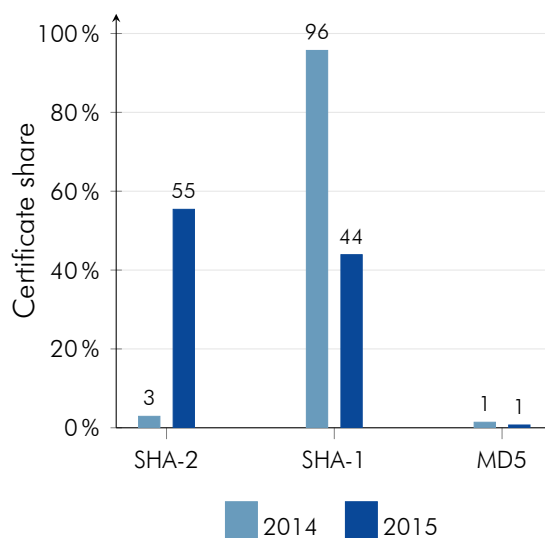


Figure 3.5: Change in certificate signatures

to collision attacks which might endanger the PKI [56].

Mid-2015, the presence of SHA-2 rose to 55 %, while that for SHA-1 fell accordingly to 44 %. This indicates that the announced gradual rejection of SHA-1 has clearly changed into a positive adoption of SHA-2, although a substantial proportion of certificates signed using SHA-1 is still in use. At 98 %, SHA-256 is the most commonly used function of the SHA-2 family.

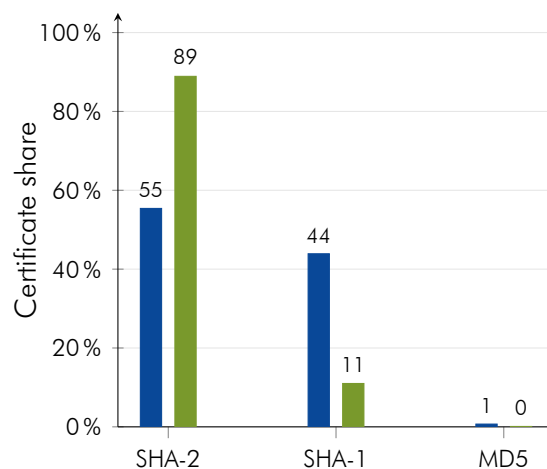
A residue of 1 % of certificates signed with MD5 remains. A manual examination of the relevant IP address shows they are essentially those of unused servers, sending empty or default HTML pages. The impact of this persistence of MD5 is therefore deemed minimal.

Emergence of New Certificates

Among the previously observed certificates, some were signed several years before the measurements. For this reason, the overall analysis of the certificates in the .fr zone only partially reflects the change in emission practices by the certification authorities. The latter can be established by observing the profiles of the most recent certificates.

However, apart from the Certificate Transparency initiative which is being standardized [57], there is no public registry of the times of issuance of X.509 certificates. The issuing date differ from that of the start of the validity period, but also from that of the first public appearance. Aware of the approximation made, the Observatory isolated the certificates observed in July 2015 and valid from January 1, 2015 at the earliest.

The comparison of the 22,759 non-self-signed certificates noted in July 2015 and the



■ retrieved in 2015 ■ retrieved in 2015, valid from January 1st 2015 or later

Figure 3.6: Change in signatures in issuing new certificates

subset of 9,475 valid certificates from January 1, 2015 at the earliest is shown in Figure 3.6. For the new certificates, the utilization rate of SHA-2 and SHA-1 are 89% and 11% respectively, characterizing an obvious transition to SHA-2. Furthermore, only 10 certificates signed with MD5 were observed.

In a nutshell

In July 2015, more than half of the certificates in the `.fr` zone were signed using the SHA-2 algorithm. In accordance with best practices, few of the newly issued certificates are based on SHA-1.

General Conclusion

Among the various improvements made to the methodologies of the Observatory, the TLS analyses offer a new vision of the Internet in France. It is now possible to understand the parameters affecting the security of exchanges carried out using HTTPS. For example, the recommended TLS version 1.2 is supported by 75 % of the servers in zones delegated under the `.fr` TLD. Similarly, many of them offer perfect forward secrecy (PFS) to their users.


Furthermore, with respect to the robustness of the certificates, the Observatory showed the virtual disappearance of signatures made with `SHA-1` in favor of `SHA-2`. In July 2015, more than half of the certificates were signed in accordance with best practices, with `SHA-2`. This is a particularly encouraging result for the study of the TLS protocol, confirming the relevance of observing and comparing the development of a protocol using stable technical indicators.

Regarding the IPv6 protocol, the trends initiated in previous years were confirmed in 2015. From the point of view of BGP, the number of French ASes implementing IPv6 has increased by 13 % during the year against 6 % in 2014, to finish at about 300 by year end. The situation is less satisfactory on other aspects. For example, the adoption of IPv6 for DNS servers and mail servers is stagnant, and changed little between 2014 and 2015. Unlike previous years, the number of IPv6 prefixes not covered by `route6` objects increased. The coverage also remains very low with RPKI: over 99 % of the IPv6 address space is not covered by Route Origin Authorization (ROA). While not alarming, these various observations suggest changes in the implementation of IPv6.

Among the observations made in 2015, some appear problematic. For example, the dispersion of inbound email relays remains low under IPv4, and worrying under IPv6. Similarly, these relays are highly concentrated on a small number of service platforms, which could affect their availability. For name servers, best practices in terms of resilience and dispersion have been implemented. However the dependencies of domain names outside the `.fr` zone introduce additional risks.

In addition, the Observatory noted a slowdown in the adoption of DNSSEC, and the massive use of the `SHA-1` hash algorithm, deemed insufficient by the currently accepted best practices in cryptography. It thus seems necessary to quickly begin the transition to a more robust algorithm such as `SHA-256`. Finally, with respect to the TLS protocol, the Observatory deplores the results for SSLv2. For example, between 2014 and 2015, nearly 5,000 new servers were configured with this protocol, despite its recognized dangerousness.

In light of the analyses for 2015, the Observatory reiterates its encouragement to the



Internet players concerning the appropriation of best engineering practices generally accepted for the BGP [3], DNS [4], and TLS protocols. The Observatory also encourages them to anticipate the threat of DDoS [5]. Regarding IPv6, 2015 seems to be a pivotal year. Despite the increase in the number of French ASes implementing IPv6, the best practices for operating this protocol, studied in this report, seem to be little followed. In addition, the Observatory makes the following recommendations:


- **monitor prefix advertisements**, and be prepared to react in case of BGP hijacks;
- **use algorithms supporting forward secrecy** and **abandon SSLv2 and SHA-1** in favor of more robust mechanisms;
- **diversify the number of SMTP and DNS servers** in order to improve the robustness of the infrastructure;
- **apply best practices** including those contained in this document, to limit the effects of failures and operational errors;
- **pursue the deployments** of IPv6, DNSSEC, and RPKI to develop skills and to anticipate possible operational problems.


In a nutshell


Organizations wishing to participate in the Observatory can contact ANSSI and Afnic.

Bibliography

- [1] ANSSI, "MaBo - MRT and BGP in OCaml." <<https://github.com/ANSSI-FR/mabo>>.
- [2] ANSSI, "TaBi - Track BGP Hijacks." <<https://github.com/ANSSI-FR/tabii>>.
- [3] ANSSI, "Bonnes pratiques de configuration de BGP," tech. rep., 2013.
- [4] ANSSI, "Bonnes pratiques pour l'acquisition et l'exploitation de noms de domaine," tech. rep., May 2014.
- [5] ANSSI, "Comprendre et anticiper les attaques DDoS," tech. rep., 2015.
- [6] Vivet, Nicolas and Valadon, Guillaume, "Tools to Detect Routing Anomalies," tech. rep., May 2016.
- [7] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)." RFC 4271 (Draft Standard), Jan. 2006. Updated by RFCs 6286, 6608, 6793, 7606, 7607, 7705.
- [8] M. Lepinski, Ed., "BGPSEC Protocol Specification - draft-ietf-sidr-bgpsec-protocol-15." <<http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-15>>, 2016.
- [9] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing." RFC 6480 (Informational), Feb. 2012.
- [10] RIPE-NCC, "Routing Information Service (RIS)." <<http://www.ripe.net/data-tools/stats/ris/>>.
- [11] "asrank - Implementation of CAIDA AS ranking algorithm," tech. rep., Feb. 2014.
- [12] RIPE-NCC, "Dépôt RPKI." <<rsync://rpki.ripe.net/>>.
- [13] Spotify, "Luigi - Build complex pipelines of batch jobs." <<https://github.com/spotify/luigi>>.
- [14] "Disco MapReduce," tech. rep.
- [15] BGPMON, "Large scale bgp hijack out of india." <<http://www.bgpmon.net/large-scale-bgp-hijack-out-of-india/>>, November 2015.


- 
- [16] BGPMON, "Bgp optimizer causes thousands of fake routes." <<http://www.bgpmon.net/bgp-optimizer-causes-thousands-of-fake-routes/>>, May 2015.
- [17] DynResearch, "The vast world of fraudulent routing." <<http://research.dyn.com/2015/01/vast-world-of-fraudulent-routing/>>, January 2015.
- [18] P. Mockapetris, "Domain names - concepts and facilities." RFC 1034 (INTERNET STANDARD), Nov. 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936.
- [19] P. Mockapetris, "Domain names - implementation and specification." RFC 1035 (INTERNET STANDARD), Nov. 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 1995, 1996, 2065, 2136, 2181, 2137, 2308, 2535, 2673, 2845, 3425, 3658, 4033, 4034, 4035, 4343, 5936, 5966, 6604, 7766.
- [20] IANA, "Domain Name System (DNS) Parameters." <<http://www.iana.org/assignments/dns-parameters/dns-parameters.xhtml>>.
- [21] Attila Özgit, ".tr DDoS Attack." <<https://meetings.icann.org/en/marrakech55/schedule/mon-tech/presentation-ddos-07mar16-en.pdf>>, Dec. 2015.
- [22] "Public Suffix List." <<https://publicsuffix.org/>>.
- [23] "Observatoire de la Résilience de l'Internet français - rapport 2013." <<http://www.ssi.gouv.fr/observatoire>>, Sept. 2013.
- [24] Bob Halley, "DNSPython Library." <<http://www.dnspython.org/>>.
- [25] Afnic, "Opendata .fr." <<https://opendata.afnic.fr/fr/produits-et-services/le-fr/opendata-fr.html>>.
- [26] R. Elz, R. Bush, S. Bradner, and M. Patton, "Selection and Operation of Secondary DNS Servers." RFC 2182 (Best Current Practice), July 1997.
- [27] R. Bush, D. Karrenberg, M. Koster, and R. Plzak, "Root Name Server Operational Requirements." RFC 2870 (Best Current Practice), June 2000. Obsoleted by RFC 7720.
- [28] RIPE-NCC, "RIPE Routing Working Group Recommendations on Route Aggregation." <<http://www.ripe.net/ripe/docs/ripe-399>>, Dec. 2006.
- [29] RIPE-NCC, "RIPE Routing Working Group Recommendations on IPv6 Route Aggregation." <<http://www.ripe.net/ripe/docs/ripe-532>>, Nov. 2011.
- [30] MaxMind, "GeoIP | IP Address Location Technology." <<http://www.maxmind.com/app/ip-location>>.

- 
- [31] Damien Giry, "Cryptographic Key Length Recommendation." <<http://www.keylength.com/fr>>, Sept. 2015.
- [32] ANSSI, "Référentiel Général de Sécurité," tech. rep., June 2014.
- [33] ICANN, "Root zone." <https://www.internic.net/domain/root.zone>.
- [34] T. Dierks and C. Allen, "The TLS Protocol Version 1.0." RFC 2246 (Proposed Standard), Jan. 1999. Obsoleted by RFC 4346, updated by RFCs 3546, 5746, 6176, 7465, 7507.
- [35] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1." RFC 4346 (Proposed Standard), Apr. 2006. Obsoleted by RFC 5246, updated by RFCs 4366, 4680, 4681, 5746, 6176, 7465, 7507.
- [36] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2." RFC 5246 (Proposed Standard), Aug. 2008. Updated by RFCs 5746, 5878, 6176, 7465, 7507, 7568, 7627, 7685.
- [37] R. Barnes, M. Thomson, A. Pironti, and A. Langley, "Deprecating Secure Sockets Layer Version 3.0." RFC 7568 (Proposed Standard), June 2015.
- [38] O. Levillain, "SSL/TLS, 3 ans plus tard." <http://www.ssi.gouv.fr/uploads/2015/06/SSTIC2015-Article-ssltls_soa_reloaded-levillain_c0bDbqp.pdf>, Juin 2015.
- [39] I. Ristić in *Bulletproof SSL and TLS*, Feisty Duck, August 2014.
- [40] E. Rescorla, "Diffie-Hellman Key Agreement Method." RFC 2631 (Proposed Standard), June 1999.
- [41] D. E. 3rd, "Transport Layer Security (TLS) Extensions: Extension Definitions." RFC 6066 (Proposed Standard), Jan. 2011.
- [42] Mozilla, "Network security services." <<https://developer.mozilla.org/en-US/docs/Mozilla/Projects/NSS>>.
- [43] Agence nationale de la sécurité des systèmes d'information (ANSSI), "Référentiel Général de Sécurité - Annexe A4." <https://references.modernisation.gouv.fr/sites/default/files/RGS_v-2-0_A4.pdf>.
- [44] Z. Durumeric, D. Adrian, A. Mirian, J. Kasten, E. Bursztein, N. Lidzborski, K. Thomas, V. Eranti, M. Bailey, and J. A. Halderman, "Neither Snow Nor Rain Nor MITM... An Empirical Analysis of Email Delivery Security," in *Internet Measurement Conference (IMC)*, October 2015.
- [45] "Parsifal: an OCaml-based parsing engine." <<https://github.com/ANSSI-FR/parsifal>>.

- 
- [46] “Scapy: the Python-based interactive packet manipulation program & library.” <<https://github.com/secdev/scapy>>.
- [47] J. Rizzo and T. Duong, “Browser Exploit Against SSL/TLS.” <<https://packetstormsecurity.com/files/105499/Browser-Exploit-Against-SSL-TLS.html>>, October 2011.
- [48] S. Turner and T. Polk, “Prohibiting Secure Sockets Layer (SSL) Version 2.0.” RFC 6176 (Proposed Standard), Mar. 2011.
- [49] N. Aviram, S. Schinzel, J. Somorovsky, N. Heninger, M. Dankel, J. Steube, L. Valenta, D. Adrian, J. A. Halderman, V. Dukhovni, E. Kasper, S. Cohney, S. Engels, C. Paar, and Y. Shavitt, “DROWN: Breaking TLS using SSLv2.” <<https://drownattack.com/drown-attack-paper.pdf>>, March 2016.
- [50] B. Möller, T. Duong, and K. Kotowicz, “This POODLE bites: Exploiting the SSL 3.0 Fallback,” tech. rep., September 2014.
- [51] H. Y. Xiaoyun Wang, “Advances in cryptology – crypto 2005: 25th annual international cryptology conference, santa barbara, california, usa, august 14-18, 2005. proceedings,” 2005.
- [52] M. Stevens, P. Karpman, and T. Peyrin, “Freestart collision for full SHA-1.” <<https://eprint.iacr.org/2015/967>>, 2016.
- [53] Microsoft, “SHA-1 Deprecation Update.” <<https://blogs.windows.com/msedgedev/2015/11/04/sha-1-deprecation-update/>>, November 2015.
- [54] Google, “An update on SHA-1 certificates in Chrome.” <<https://security.googleblog.com/2015/12/an-update-on-sha-1-certificates-in.html>>, November 2015.
- [55] Mozilla, “Continuing to Phase Out SHA-1 Certificates.” <<https://blog.mozilla.org/security/2015/10/20/continuing-to-phase-out-sha-1-certificates/>>, December 2015.
- [56] H. Y. Xiaoyun Wang, “How to break md5 and other hash functions,” in *EUROCRYPT’05*, pp. 19–35, 2005.
- [57] B. Laurie, A. Langley, and E. Kasper, “Certificate Transparency.” RFC 6962 (Experimental), June 2013.

Acronyms

Afnic	Association Française pour le Nommage Internet en Coopération
ANSSI	Agence nationale de la sécurité des systèmes d'information
AS	Autonomous System
BGP	Border Gateway Protocol
BGPsec	Border Gateway Protocol Security
CA	Certificate Authority
DDoS	Distributed Denial of Service
DHE	Diffie-Hellman Ephemeral
DNS	Domain Name System
DNSSEC	Domain Name System Security Extensions
DS	Delegation Signer
ECDHE	Elliptic Curve Diffie–Hellman Ephemeral
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IETF	Internet Engineering Task Force
IP	Internet Protocol
IRR	Internet Routing Registry
ISP	Internet Service Provider
PKI	Public Key Infrastructure
PSL	Public Suffix List
RGS	Référentiel Général de Sécurité
RIPE-NCC	RIPE Network Coordination Centre
RIR	Regional Internet Registry
RIS	Routing Information Service



ROAs	Route Origin Authorizations
RPKI	Resource Public Key Infrastructure
SNI	Server Name Indication
SPOF	Single Point of Failure
SSL	Secure Sockets Layer
TLD	Top Level Domain
TLS	Transport Layer Security

About ANSSI

The Agence nationale de la sécurité des systèmes d'information (ANSSI - French Network and Information Security Agency) was created on July 7 2009 as an agency with national jurisdiction.

By Decree No. 2009-834 of July 7 2009 as amended by Decree No. 2011-170 of February 11 2011, the agency has responsibility at national level concerning the defence and security of information systems. It is attached to the Secretariat-General for National Defence and Security (*Secrétaire général de la défense et de la sécurité nationale*) under the authority of the Prime Minister.

To learn more about ANSSI and its activities, please visit www.ssi.gouv.fr.

October 2016

Licence ouverte / Open Licence (Etalab v1)

Agence nationale de la sécurité des systèmes d'information
ANSSI - 51 boulevard de la Tour-Maubourg - 75700 PARIS 07 SP
Web site: www.ssi.gouv.fr
Email address: [communication \[at\] ssi.gouv.fr](mailto:communication@ssi.gouv.fr)