French Internet Resilience Observatory

François Contat, Guillaume Valadon

Agence nationale de la sécurité des systèmes d'information

http://www.ssi.gouv.fr/en

RIPE 67 - October 15th, 2013

The observatory in a nutshell

Prior issues

- the Internet is misunderstood;
- network incidents analysis are rarely France-oriented;
- the usage of best current practices is unknown.

Some of our objectives

- study the French Internet in details;
- develop technical interactions with the networking community;
- publish anonymized results;
- publish recommendations and best practices.

Internet resilience?

«Resilience is the ability to respond to a major crisis and to quickly restore a normal service.»

The French White Paper on defence and national security, 2008

The Internet is often considered as a regular *industry*. Its resilience is mainly studied through:

- the dependency on electricity;
- the location of physical infrastructures.

The observatory aims to study the Internet resilience from a technical point of view.

Who?

The observatory is under the supervision of the ANSSI.

Created on July 7th 2009, the ANSSI is the national authority for the defence and the security of information systems:

- in French, ANSSI, Agence nationale de la sécurité des systèmes d'information;
- in English, French Network and Information Security Agency.

Main missions are:

- prevention;
- defence of information systems.

One of its priorities is the Internet resilience.

http://www.ssi.gouv.fr/en/

Who else?

Afnic

The French Registry for the `.fr` zone as well as overseas territories.

`http://www.afnic.fr/en/`

Afnic has been co-leading the project since the beginning.

French network actors

ISPs, IXP, transit providers...

What can be observed?

Two main possible directions:

- services (HTTPS usage, mail...);
- Internet structure (routing, name services).

Today, the observatory is focusing solely on the Internet structure through BGP and DNS.

How to observe?

Several technical indicators were defined:

- 7 indicators for BGP (route objects, hijacks, RPKI…);
- 5 indicators for DNS (topological distribution, DNSSEC…).

In the report, each indicator contains:

1. a description;
2. a methodology and its limitations;
3. an analysis.

Border Gateway Protocol

Data and indicators

RIS project - BGP updates

Data: AS origin, prefix, AS_PATH...

Indicators: hijacks classification, connectivity, IPv6, BCP...

RIPE-NCC Whois database

Data: route, route6, aut-num...

Indicators: hijacks classification, connectivity, IPv6, BCP...

Identifying the French Internet

Exisiting databases are not adequate: some ASes are missing.

Finding French AS

- more than 40,000 ASes in the Internet;
- automatically identify French ASes using an unsupervised learning algorithm.

Results

- 1270 French ASes;
- compared to existing public databases (Cymru, RIPE):
  - 9 ASes missing in our database;
  - 40 and 70 more ASes.

Connectivity

Motivations

- are French ASes well connected to each other?
- are there Single Point Of Failure (SPOF)?

Methodology

- build a representative graph of the French Internet:
  - use `AS_PATH` seen by the RIS collectors;
  - extract the subgraph of French ASes.
- identify the critical ASes (SPOF) for the French Internet:
  - highlight ASes whose loss can lead to a loss of connectivity.

Connectivity

IPv4

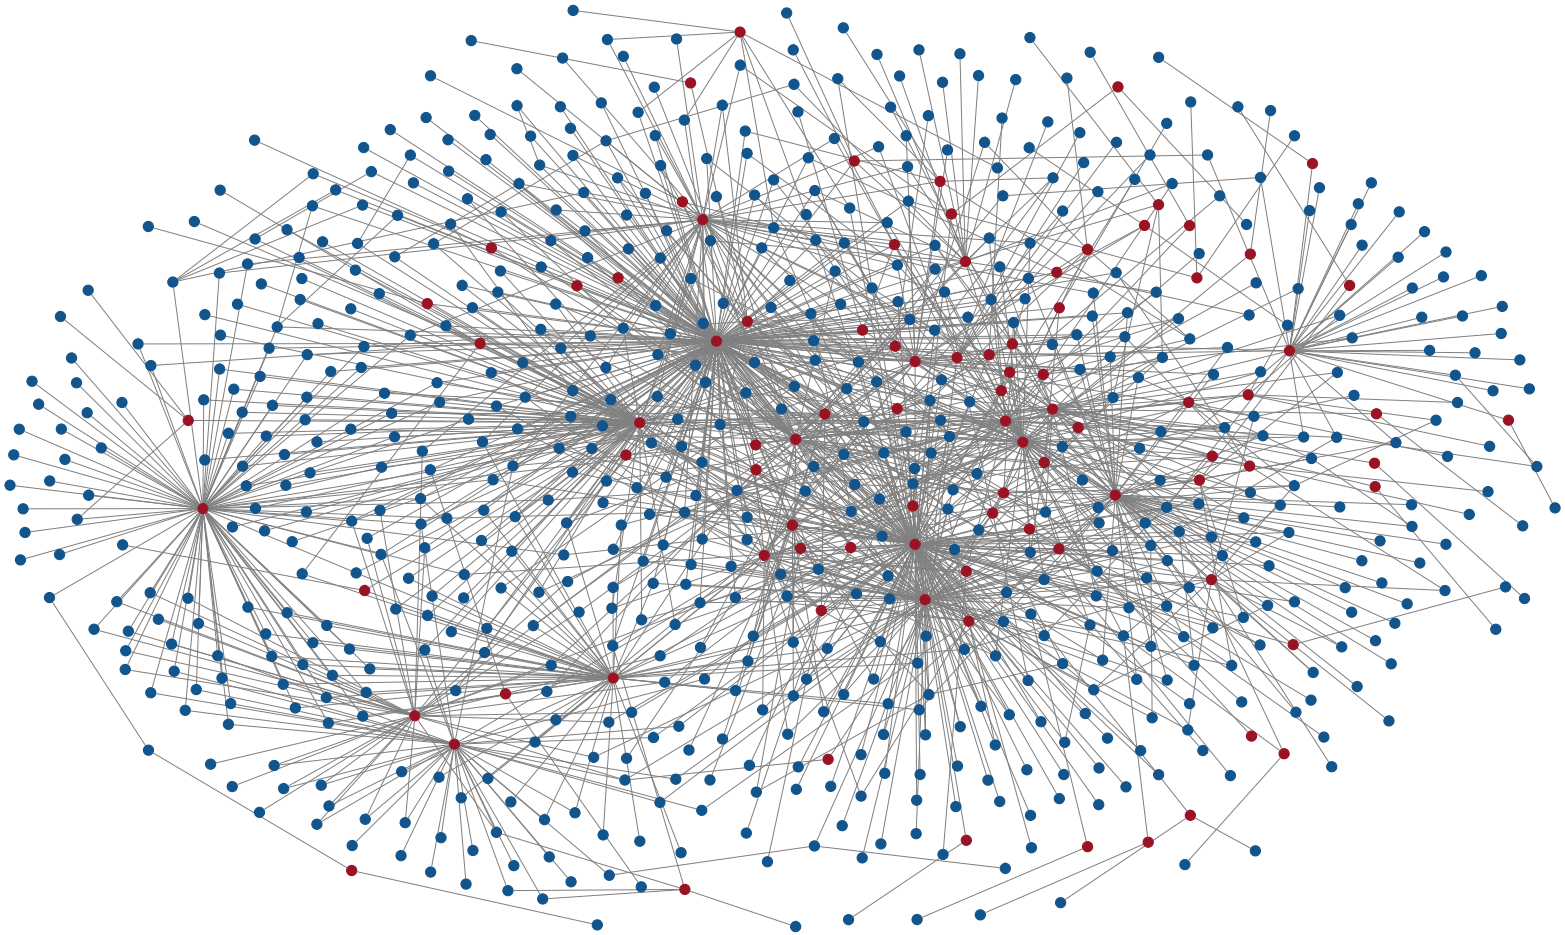

Blue: French ASes.

Red: ASes whose loss leads to a loss of connectivity.

Connectivity

IPv4

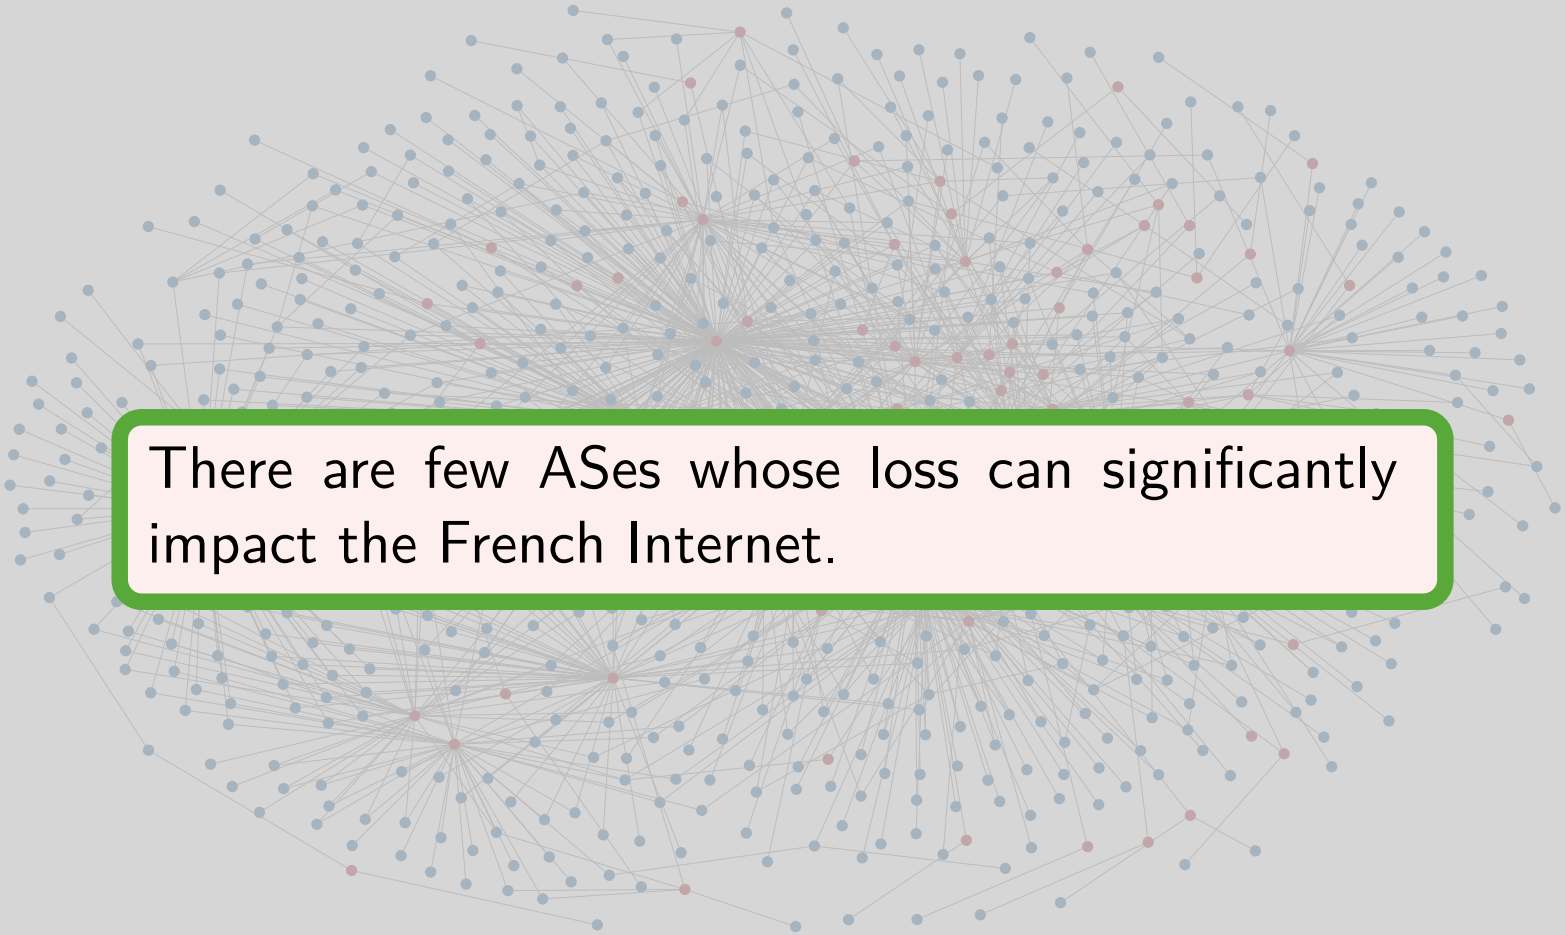

There are few ASes whose loss can significantly impact the French Internet.

Blue: French ASes.

Red: ASes whose loss leads to a loss of connectivity.

Prefix conflicts

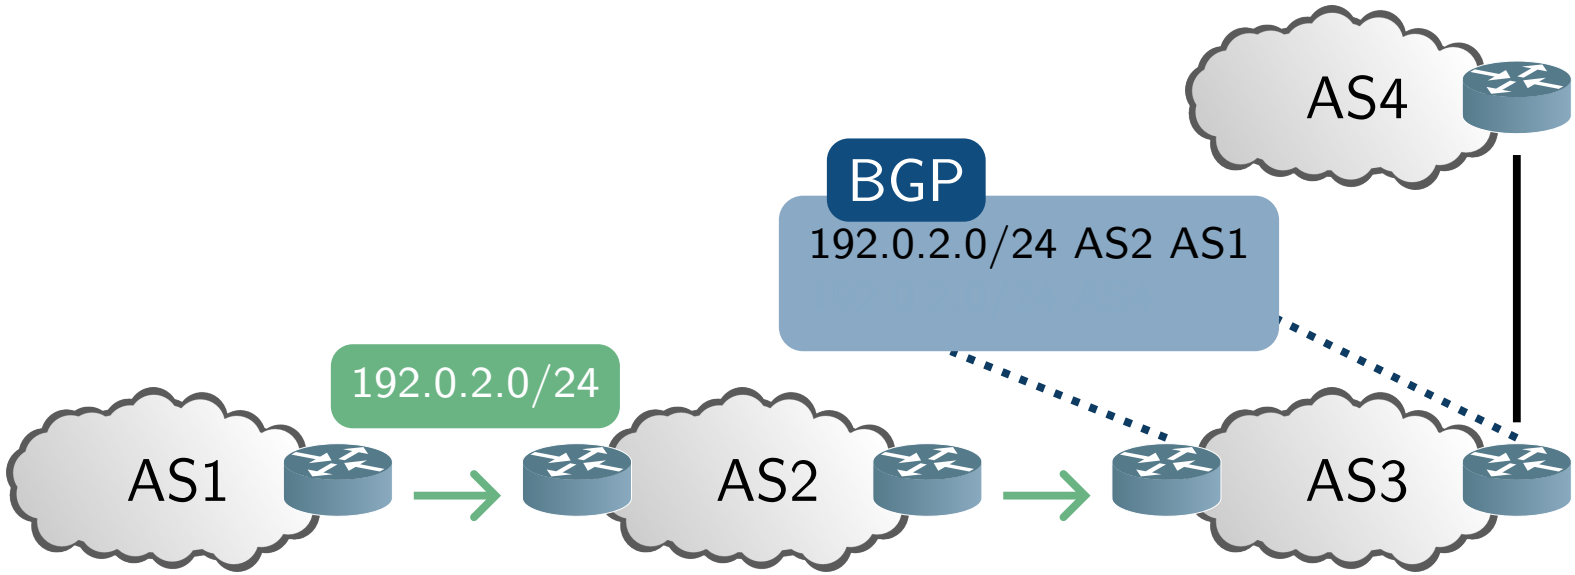

- prefix announcements between ASes: routes are exchanged;

Prefix conflicts

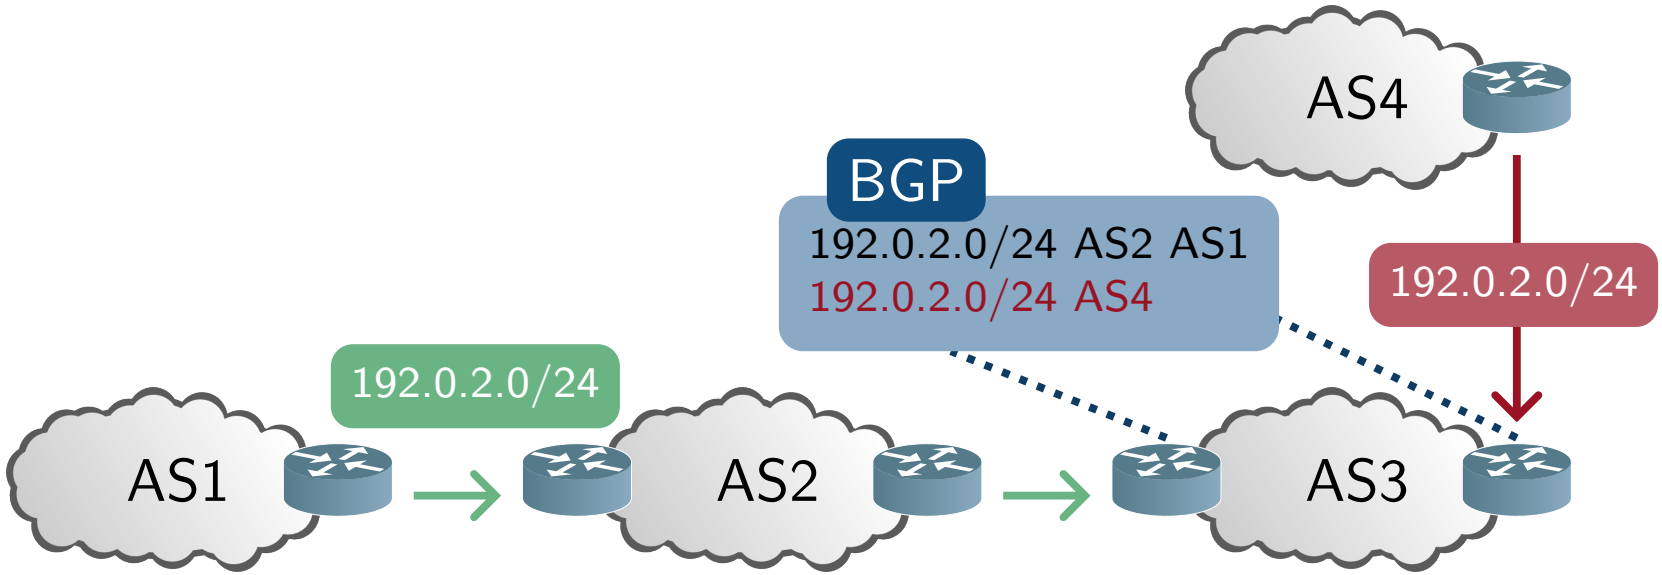

- prefix announcements between ASes: routes are exchanged;
- both ASes announce the same prefix: hijack?

Prefix conflicts

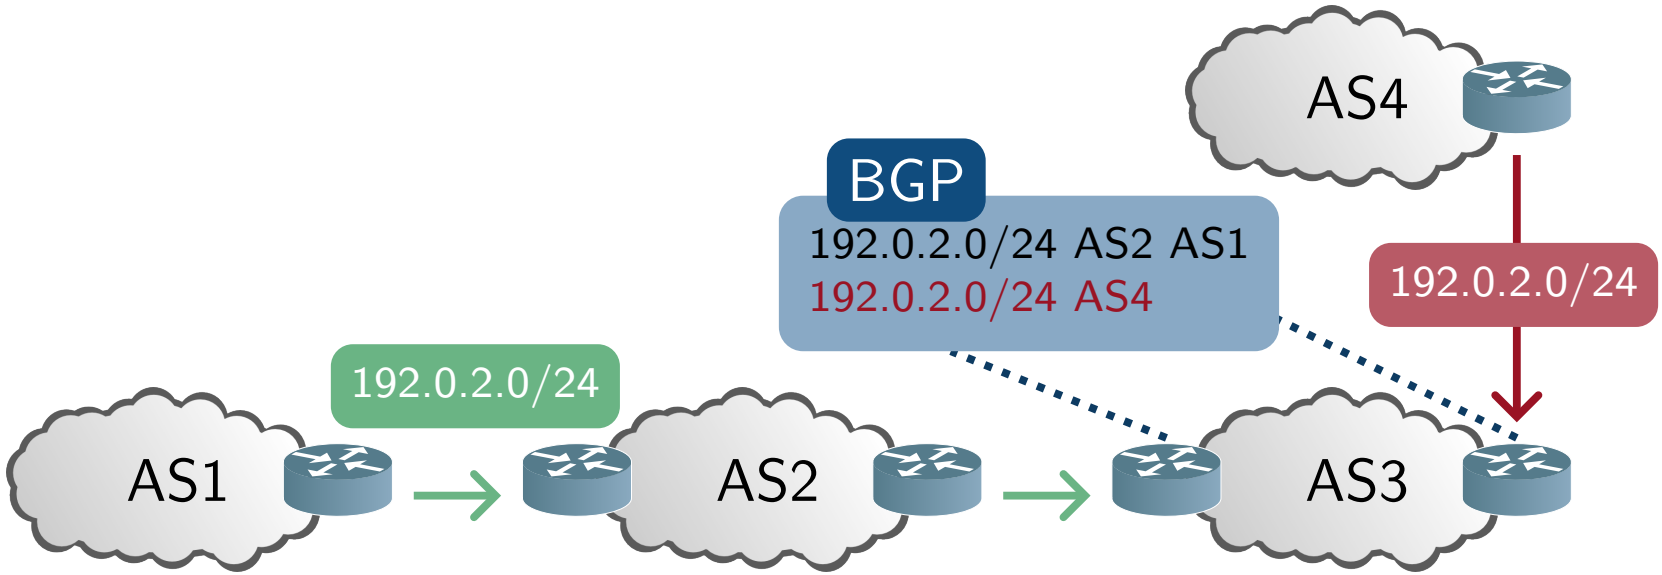

- prefix announcements between ASes: routes are exchanged;
- both ASes announce the same prefix: hijack?
- could be anycast, DDoS protection, customer...

Prefix conflicts

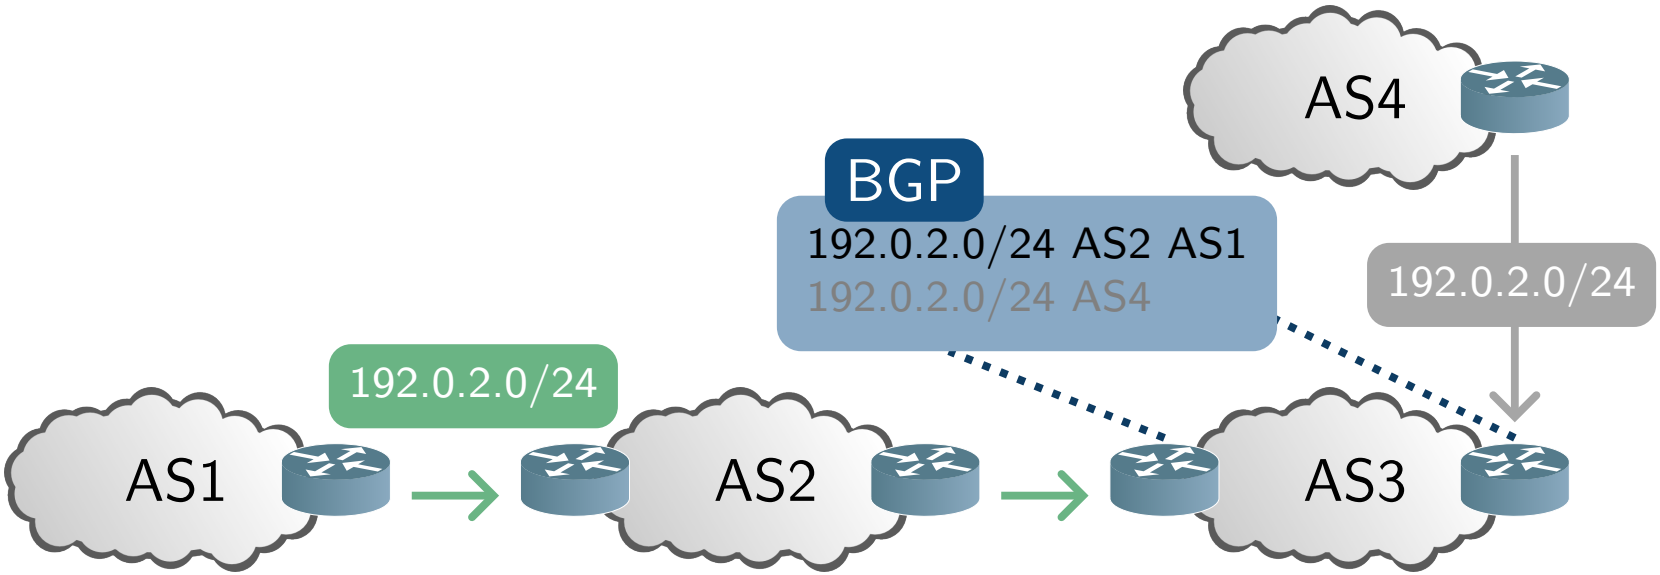

- prefix announcements between ASes: routes are exchanged;
- both ASes announce the same prefix: hijack?
- could be anycast, DDoS protection, customer...

This conflict must be named differently: event.

How can we classify an announcement as valid?

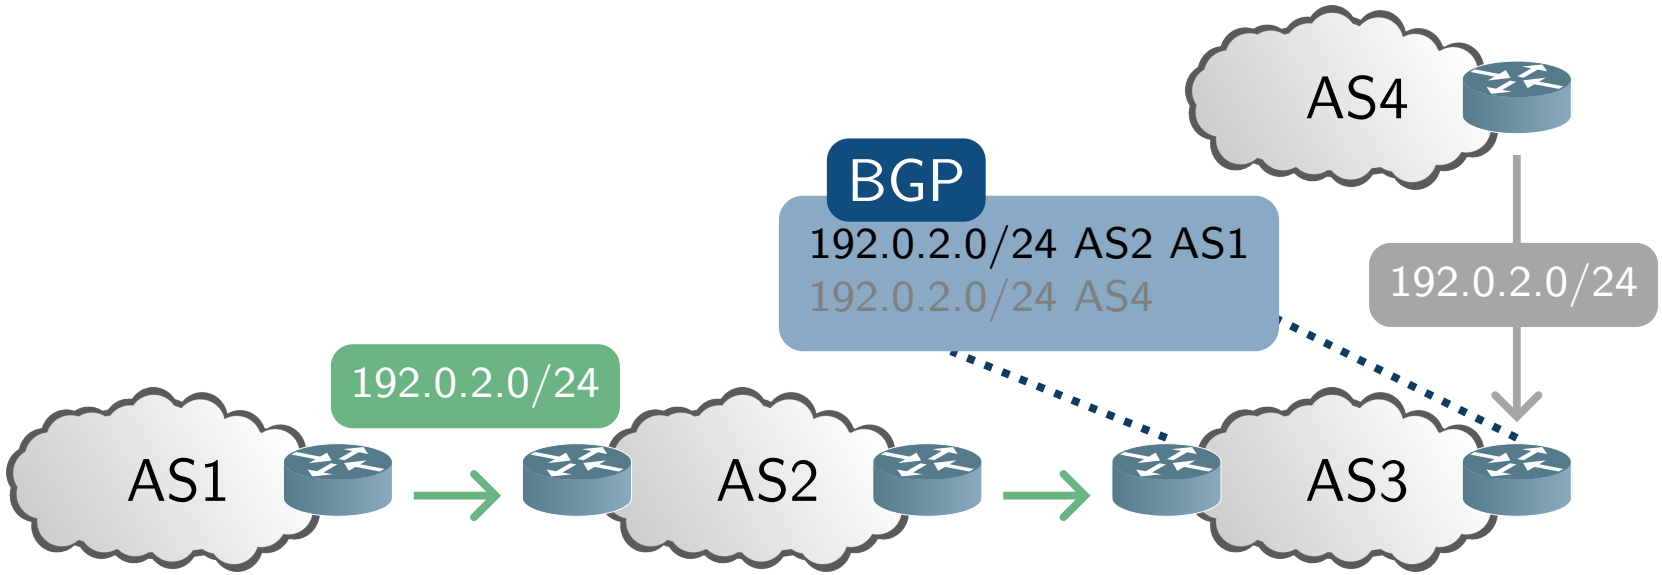

In this example, we look for the prefix 192.0.2.0/24 in whois database:

```
$ whois -T route 192.0.2.0/24
```

How can we classify an announcement as valid?

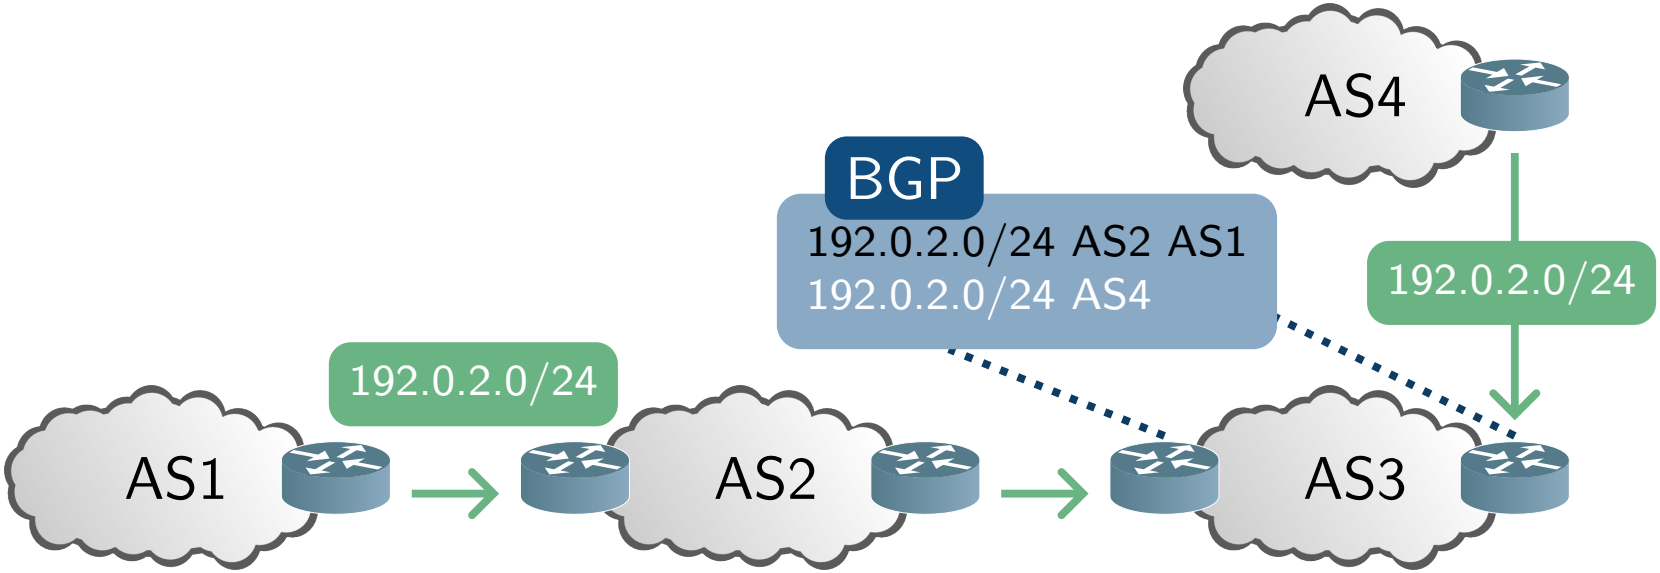

In this example, we look for the prefix 192.0.2.0/24 in whois database:

```
$ whois -T route 192.0.2.0/24
descr:          Route object example
route:          192.0.2.0/24
origin:         AS4
mnt-by:         AS1-MNT
```

Classifying events

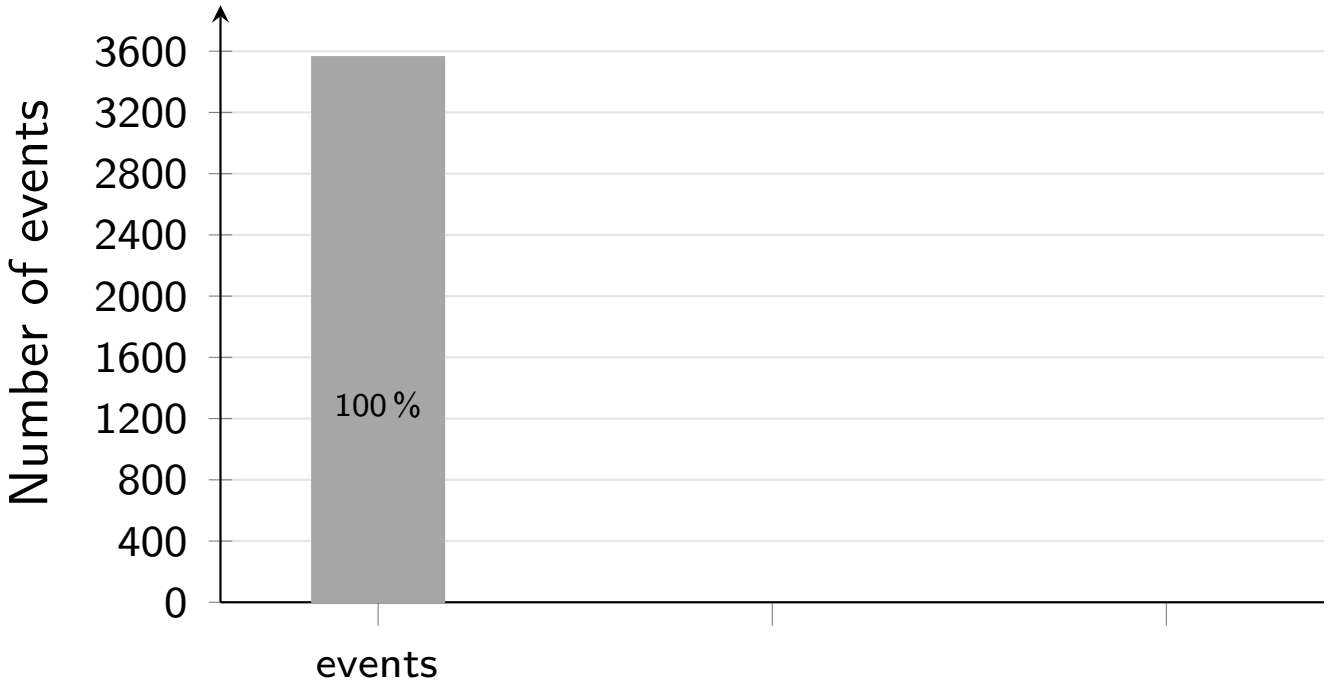

Classifying events

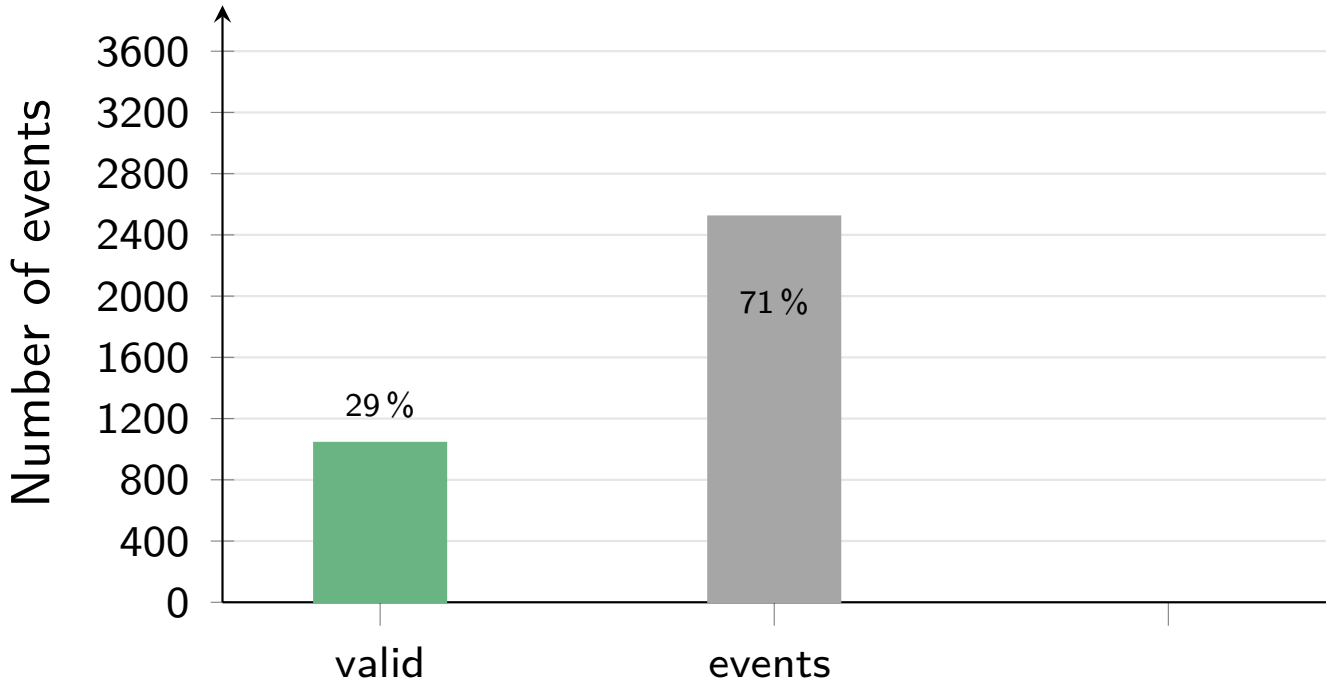

Valid: a route object exists for the AS including the prefix.

Classifying events

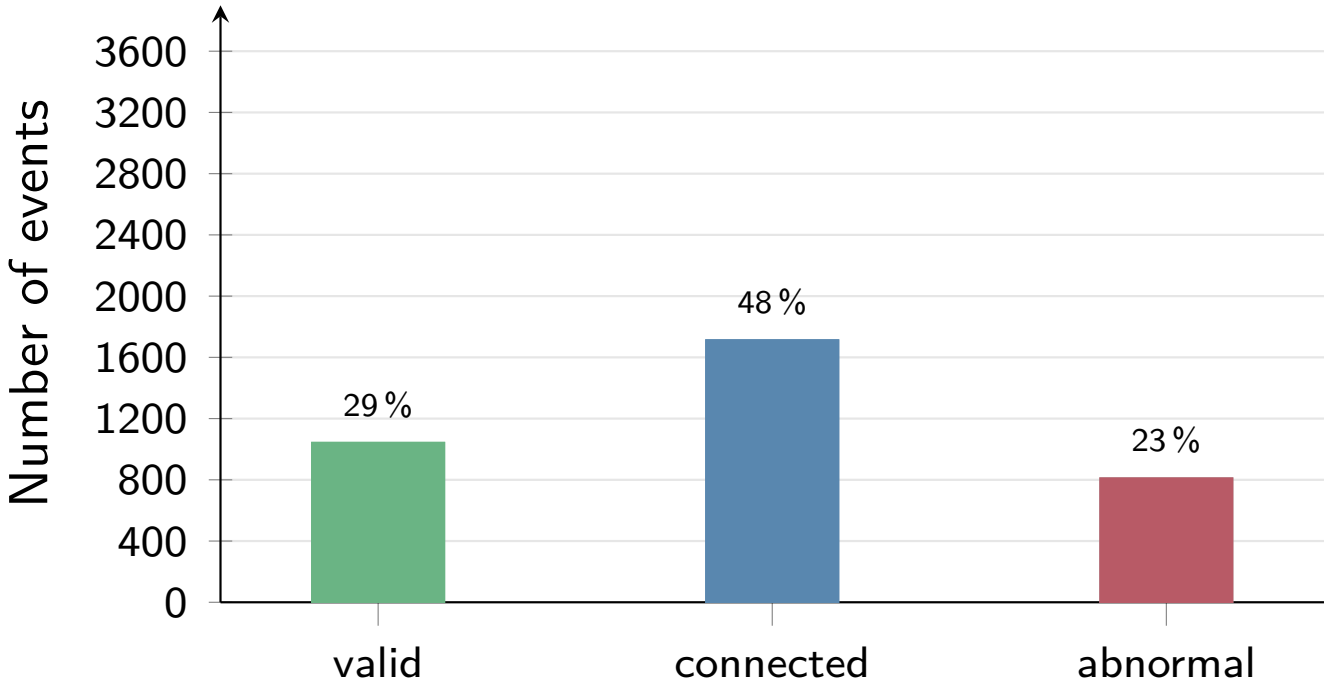

Valid: a route object exists for the AS including the prefix.
Connected: one of the ASes provides transit to the other.
Abnormal: it might be a prefix hijack.

Classifying events

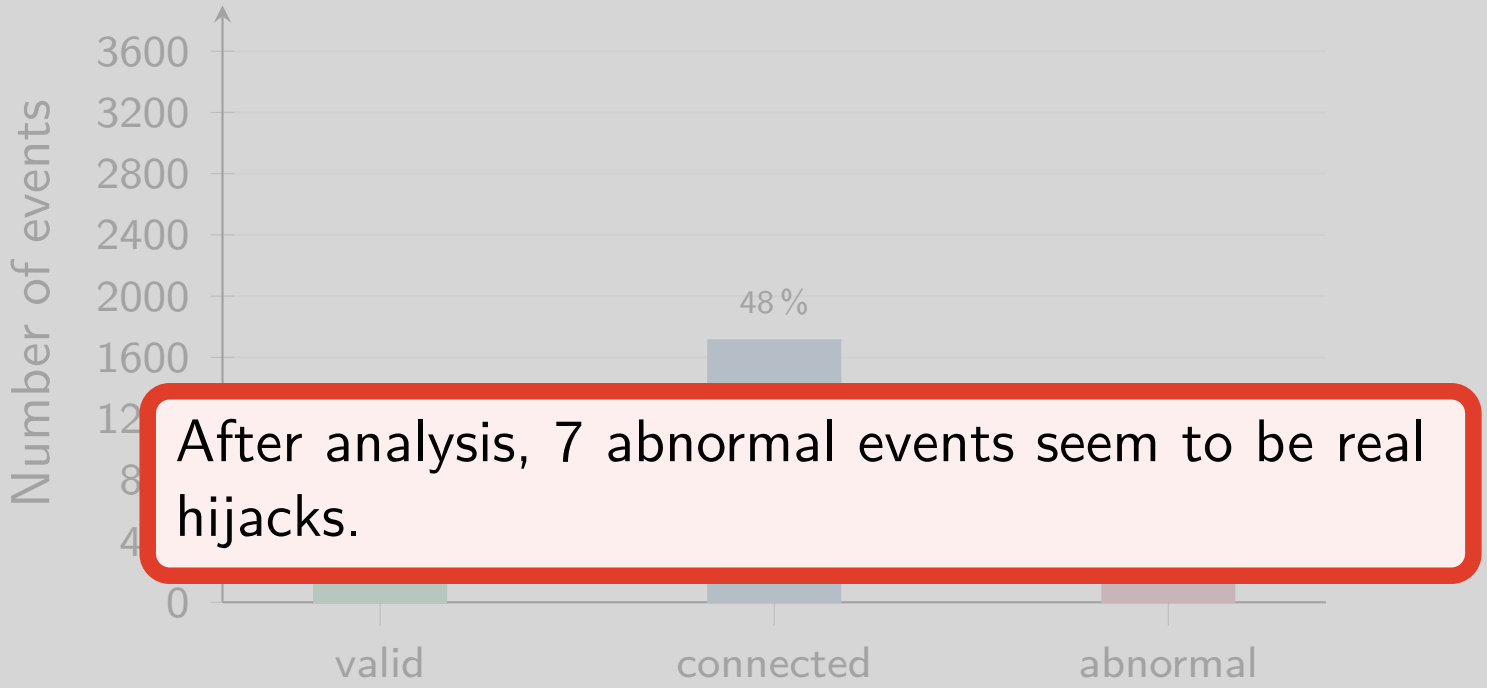

After analysis, 7 abnormal events seem to be real hijacks.

Valid: a route object exists for the AS including the prefix.
Connected: one of the ASes provides transit to the other.
Abnormal: it might be a prefix hijack.

Cross-check routing table and whois database

RIS LINX

Whois database

Cross-check routing table and whois database

RIS LINX

Whois database

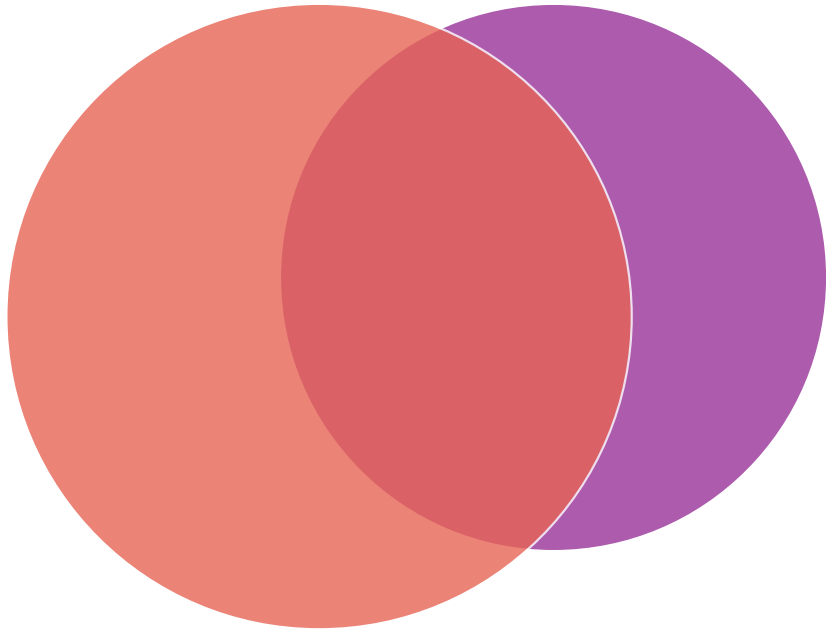

Prefix filtering

Whois consistency

Cross-check routing table and whois database

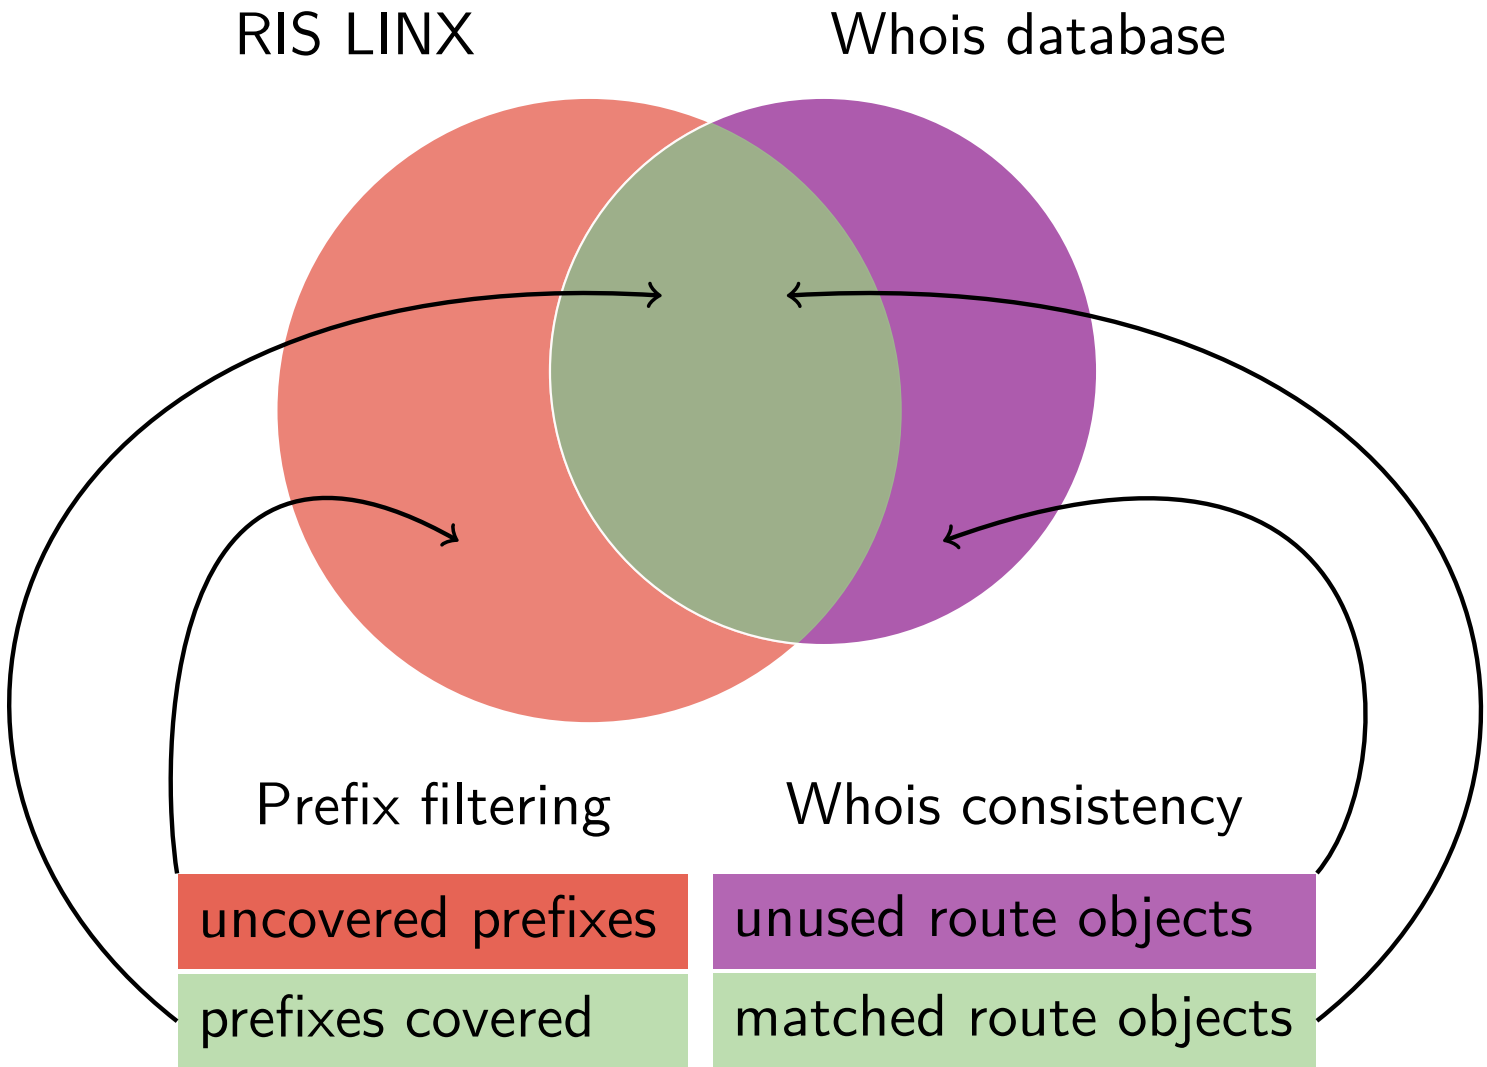

Whois database consistency

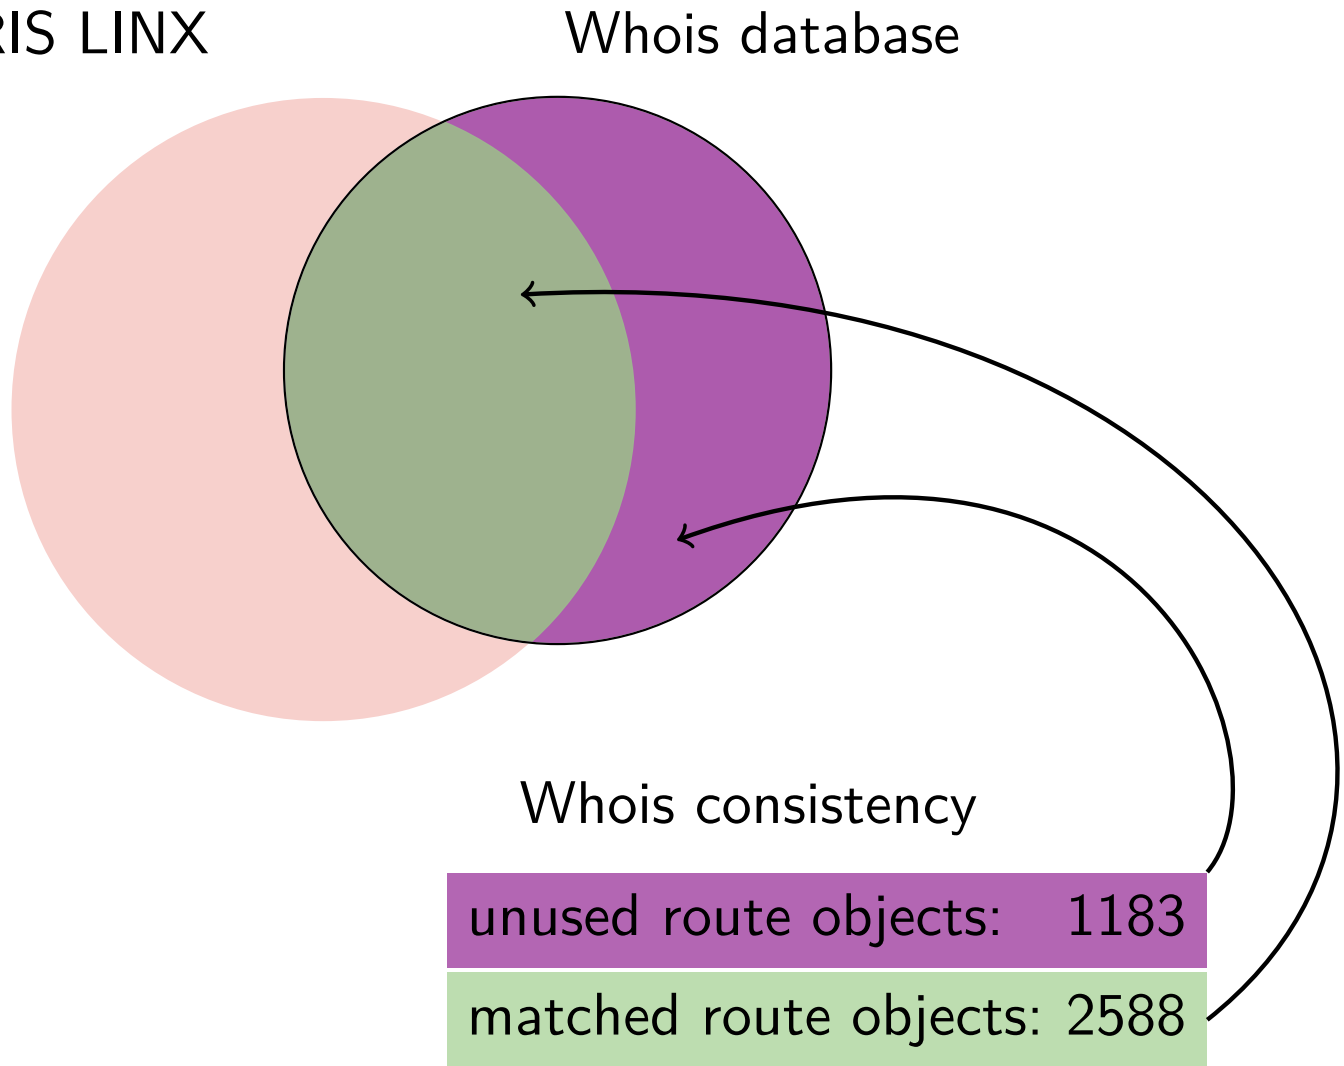

31% of route objects are unused in 2012

Prefix filtering

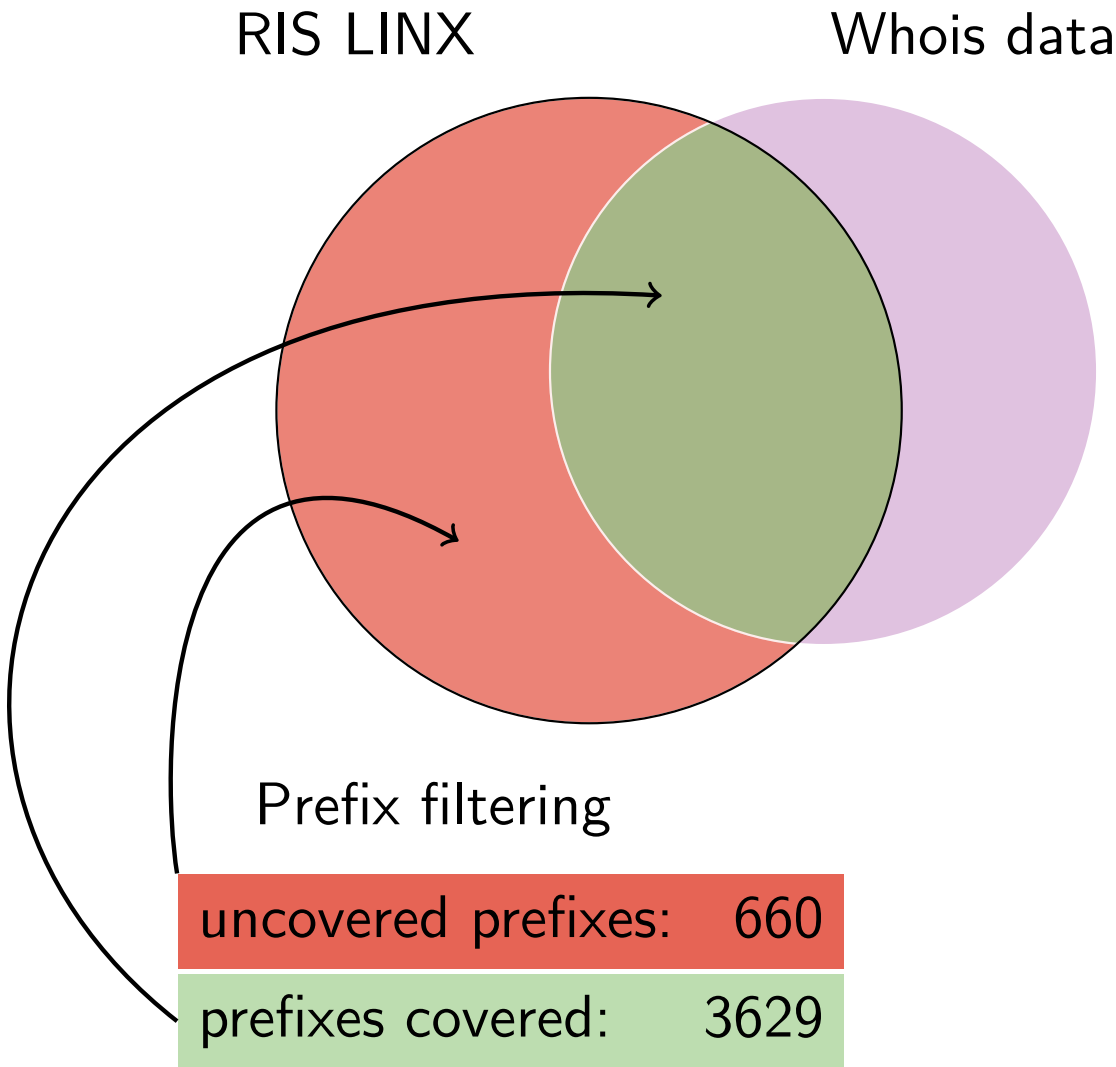

15% of French prefixes could be blackholed

Prefix filtering

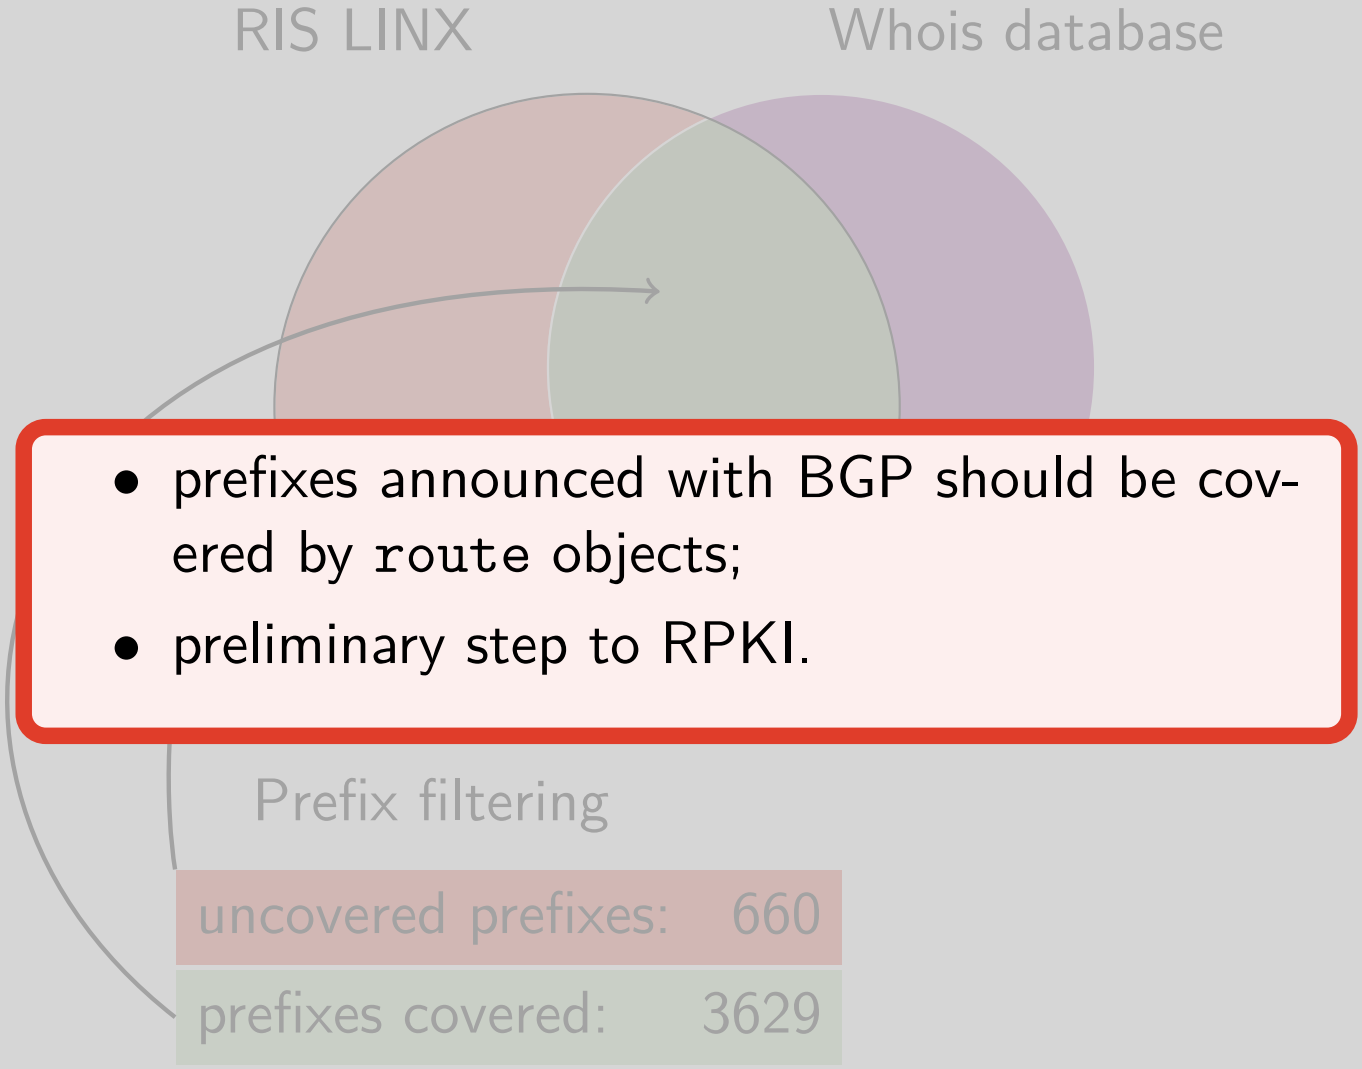

- prefixes announced with BGP should be covered by `route` objects;
- preliminary step to RPKI.

15% of French prefixes could be blackholed

Domain Name System

Data & tools

The `DNSWitness` platform is used to collect data.

Active measurements

Data: `.fr` domains retrieved from the whole `.fr` zone.

Tool: `DNSdelve`.

Indicators: number of DNS servers, IPv6 services, DNSSEC...

Passive measurements

Data: requests received by Afnic authoritative servers.

Tool: `DNSmezzo`.

Indicators: *Kaminsky* attack, IPv6 queries...

Distribution of authoritative DNS servers

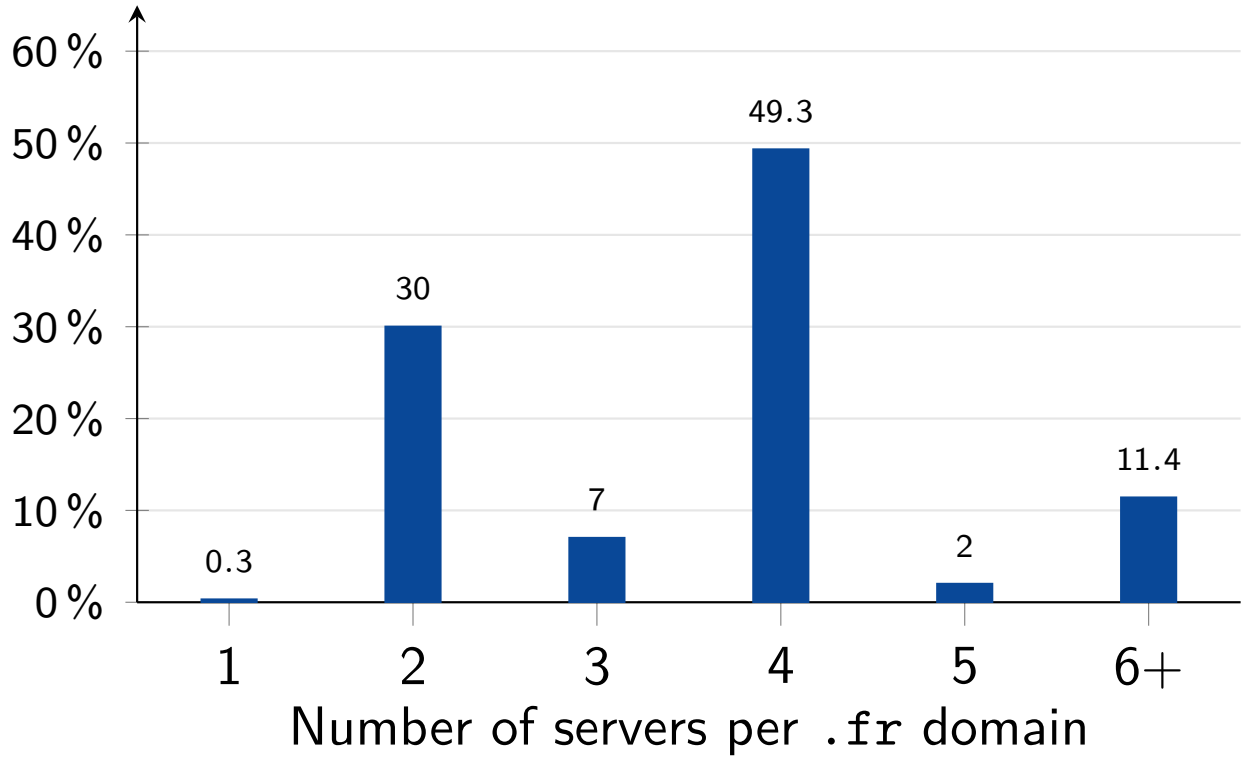

A high number of servers per zone increases the resilience.

Distribution of authoritative DNS servers

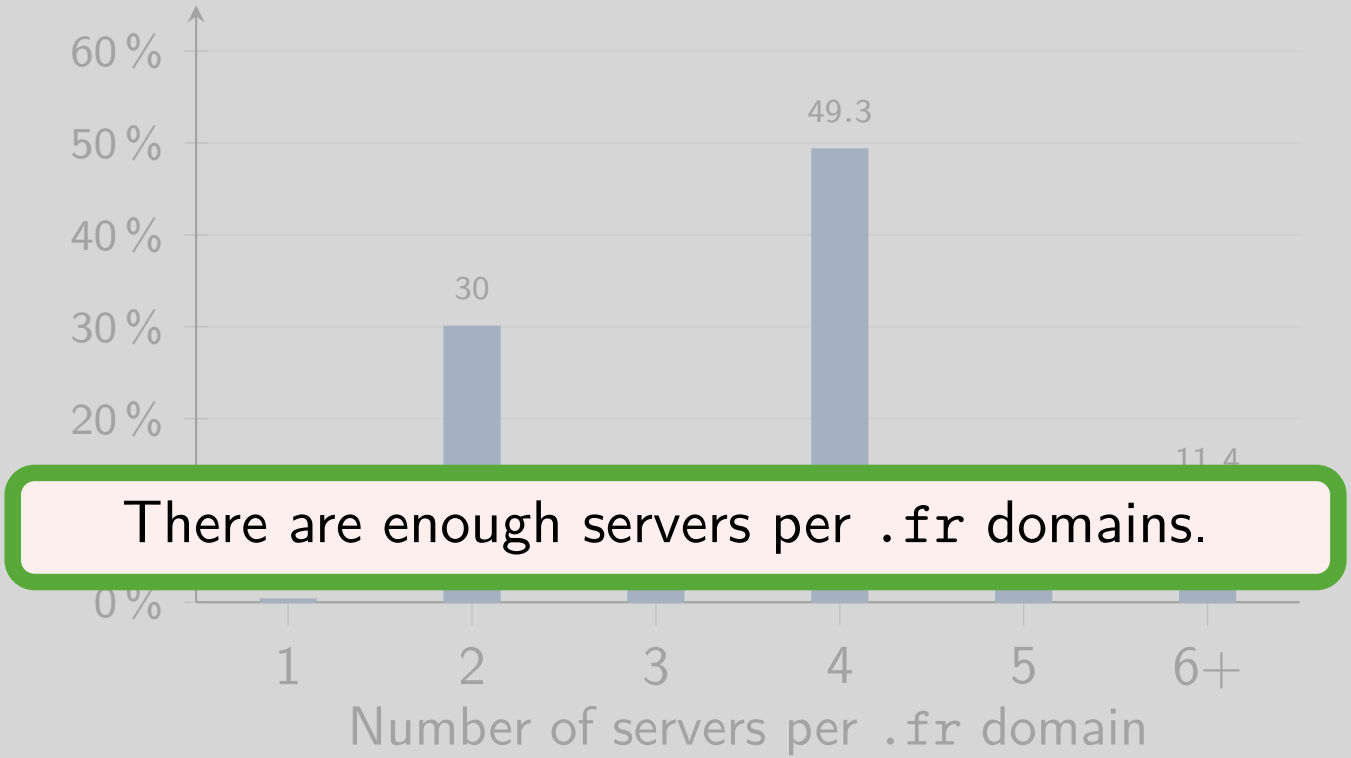

There are enough servers per `.fr` domains.

A high number of servers per zone increases the resilience.

Distribution of ASes per DNS zone

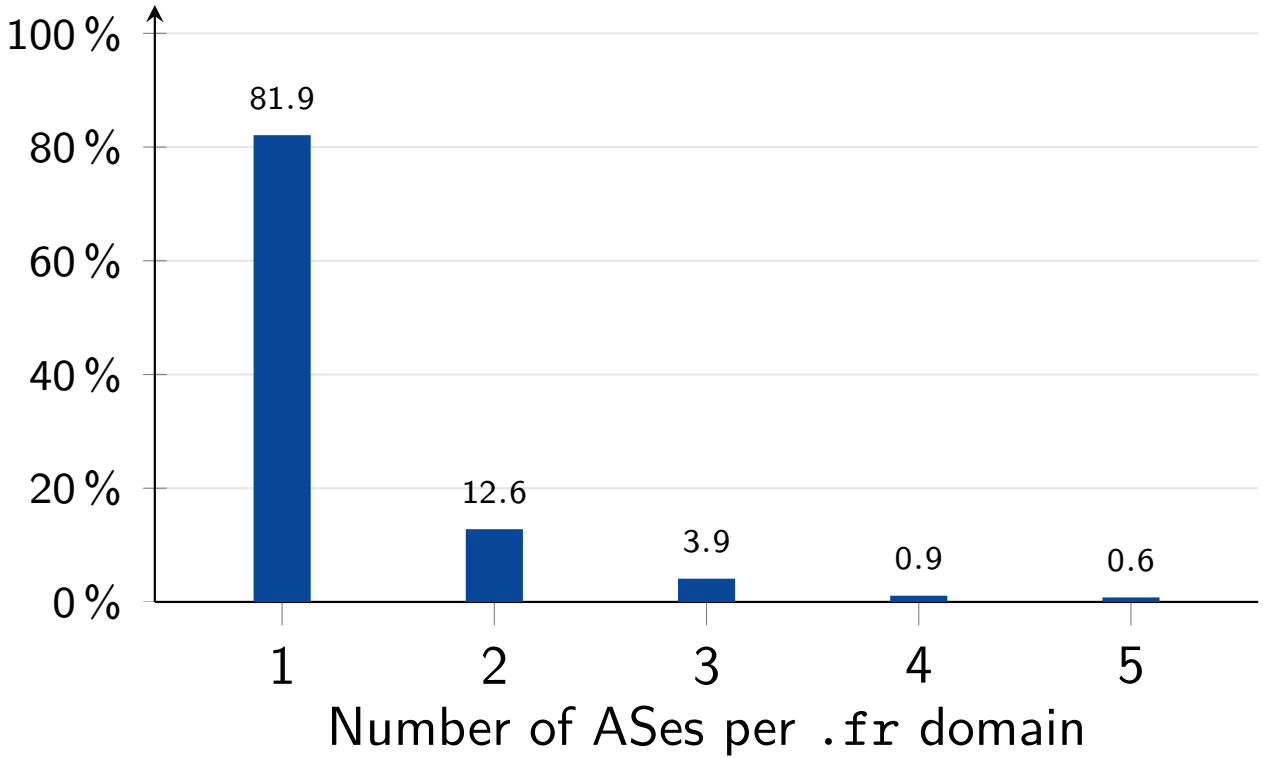

A high number of AS per zone increases resilience.

Distribution of ASes per DNS zone

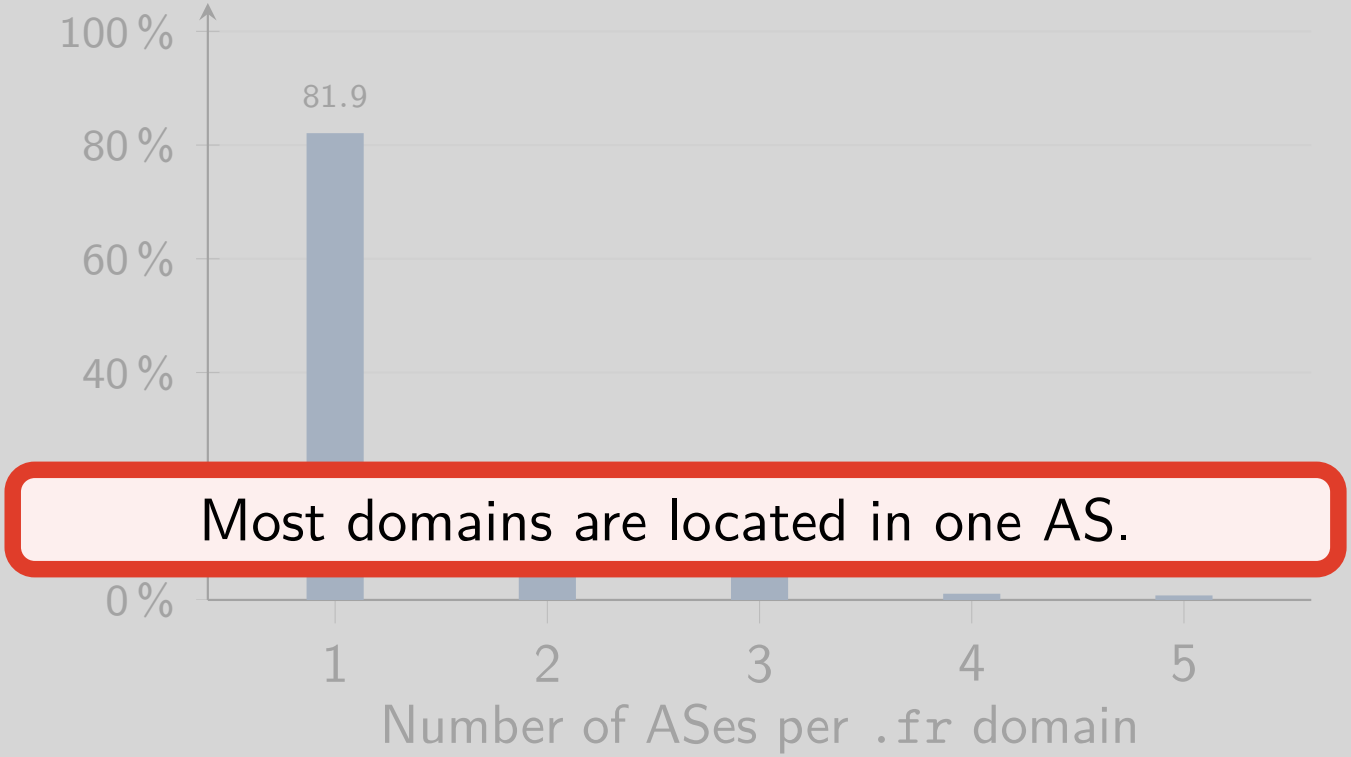

Most domains are located in one AS.

A high number of AS per zone increases resilience.

DNSSEC deployment in .fr domains

DNSSEC prevents DNS cache poisoning.

Deployment history

- .fr zone signed and published: September 14th, 2010;
- .fr zone accepts signed delegation since April, 2011.

All .fr domains could be signed.

Deployment in practice

- only 1.5% of the whole .fr zone is signed;
- thanks to a single French DNS registrar.

DNSSEC deployment in .fr domains

DNSSEC prevents DNS cache poisoning.

Deployment history

- .fr zone signed and published: September 14th, 2010;
- .fr zone accepts signed delegation since April, 2011.

DNSSEC is not widely deployed.

All .fr domains could be signed.

Deployment in practice

- only 1.5% of the whole .fr zone is signed;
- thanks to a single French DNS registrar.

IPv6 deployment of servers within `.fr` domains

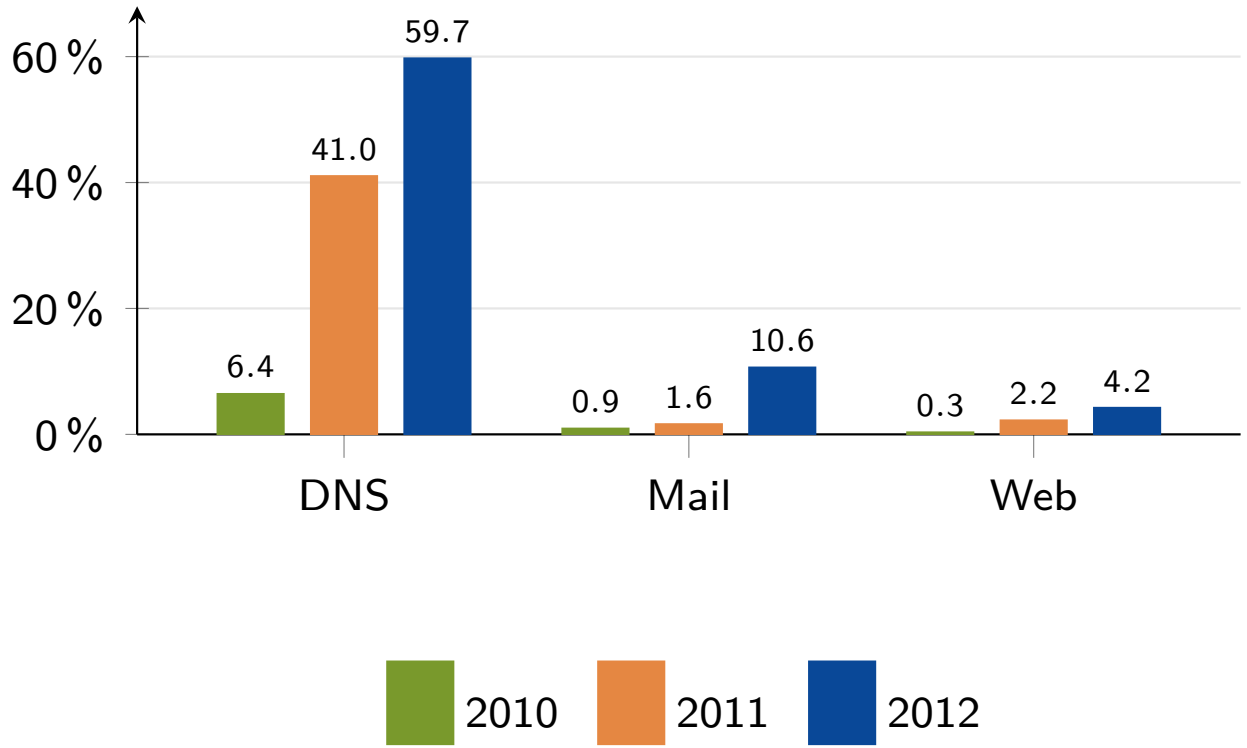

DNS: `NS` record points to a name with a `AAAA` record;

mail: `MX` record points to a name with a `AAAA` record;

web: `www.zone.fr` has a `AAAA` record.

IPv6 deployment of servers within .fr domains

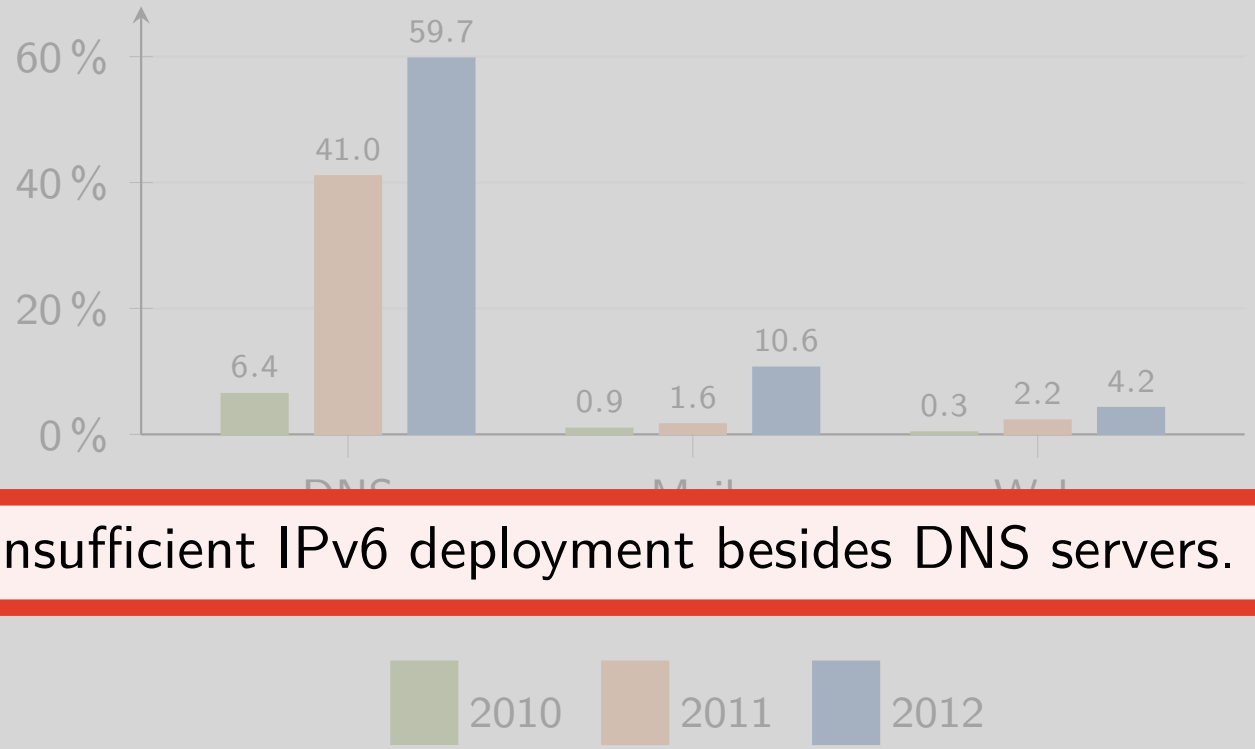

Insufficient IPv6 deployment besides DNS servers.

DNS: `NS` record points to a name with a `AAAA` record;

mail: `MX` record points to a name with a `AAAA` record;

web: `www.zone.fr` has a `AAAA` record.

IPv6 deployment of DNS cache and clients

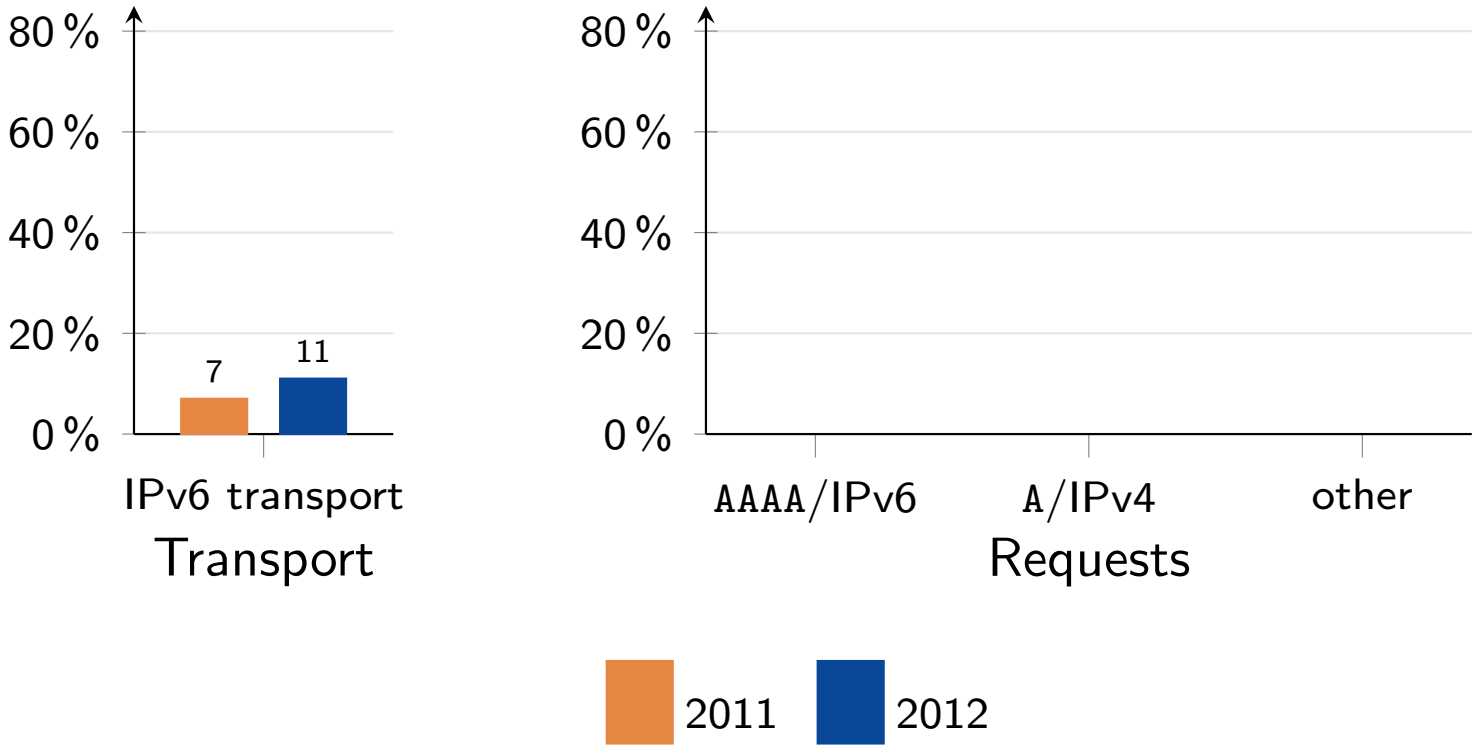

Data received by Afnic servers is now analyzed based on:

- transport: IP version preferred by DNS caches;

IPv6 deployment of DNS cache and clients

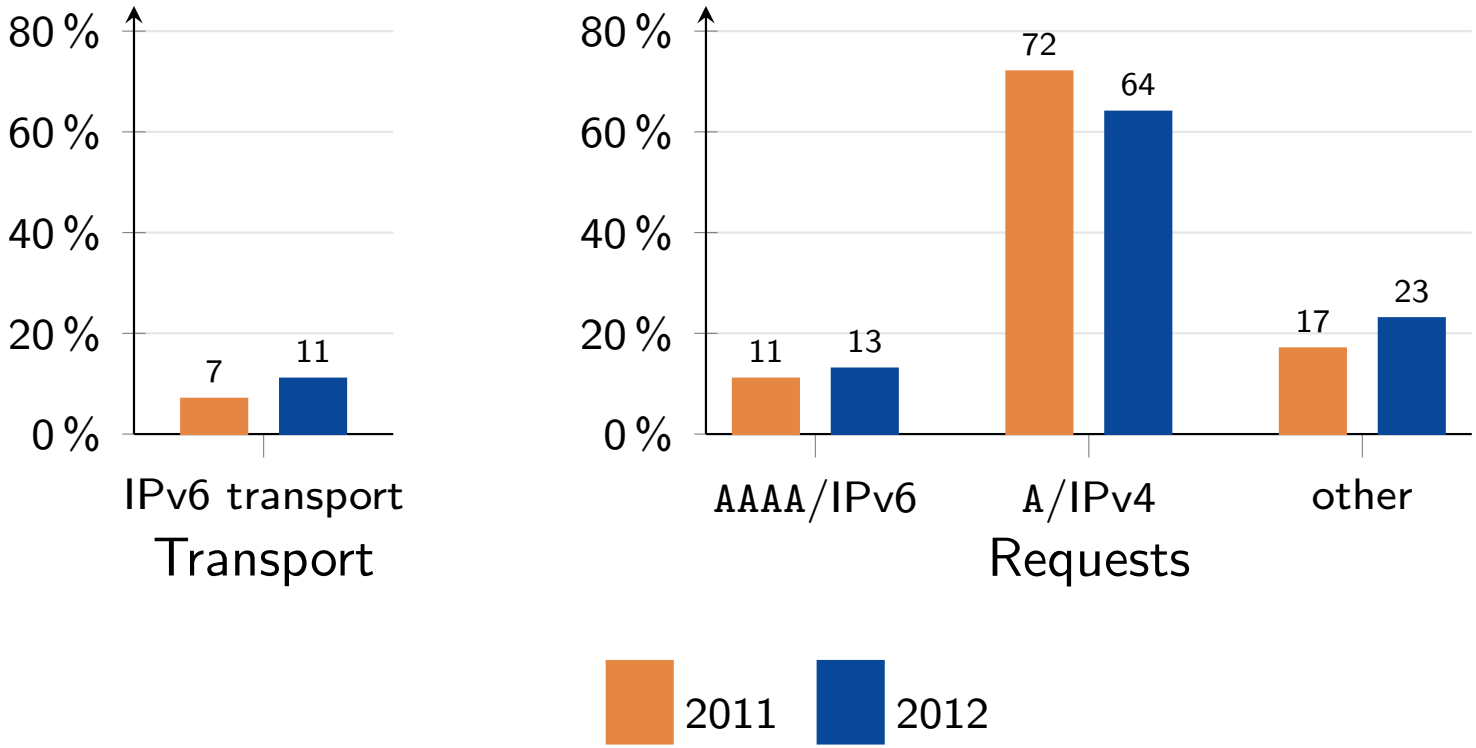

Data received by Afnic servers is now analyzed based on:

- transport: IP version preferred by DNS caches;
- requests: IP version preferred by clients.

IPv6 deployment of DNS cache and clients

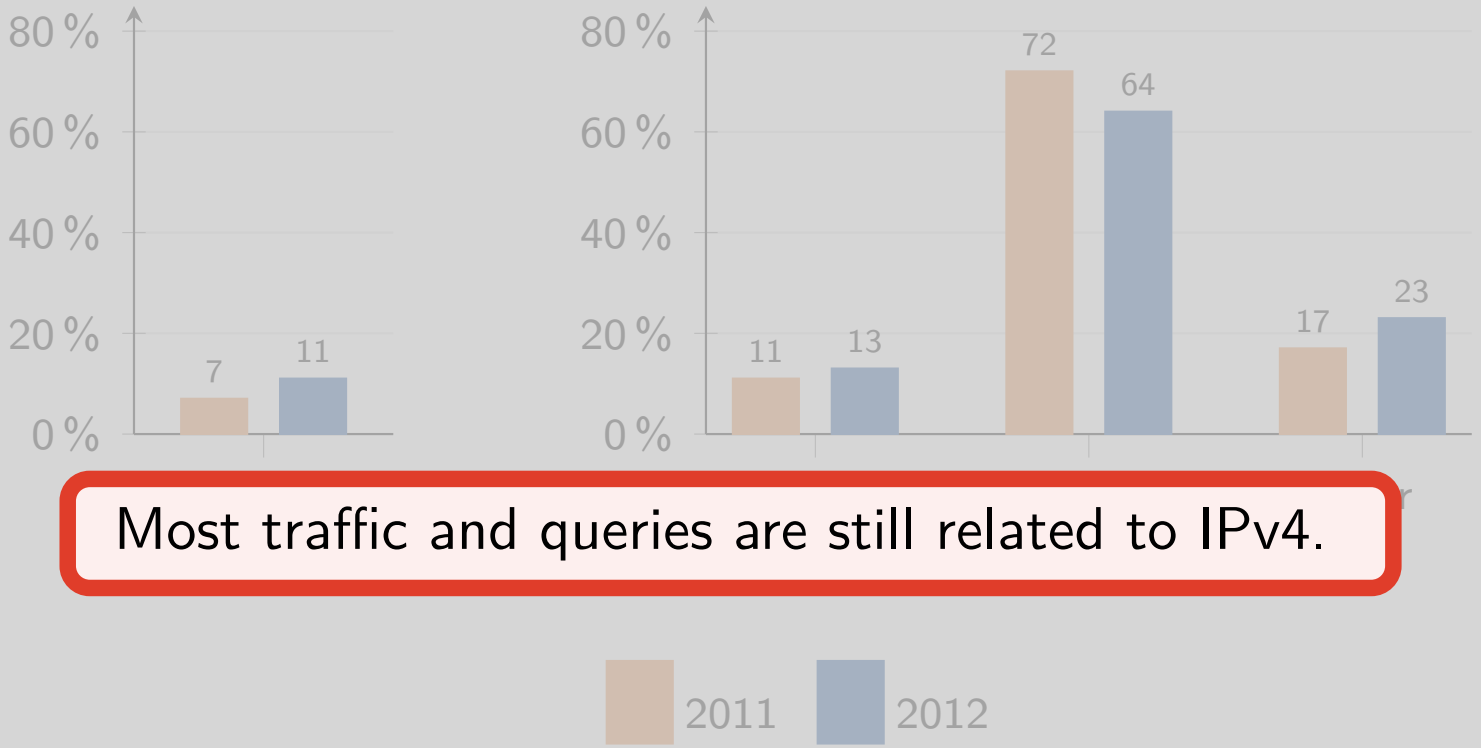

Most traffic and queries are still related to IPv4.

Data received by Afnic servers is now analyzed based on:

- transport: IP version preferred by DNS caches;
- requests: IP version preferred by clients.

Conclusion & recommendations

« For BGP & DNS, the French Internet status is acceptable. However, there is no evidence that it will be true in the future. »

2012 report

Recommendations

1. declare route objects, and keep declarations up-to-date, in order to ease filtering and hijack detection;
2. deploy IPv6 to anticipate problems;
3. apply BGP best current practices;
4. distribute authoritative DNS servers across several ASes.

Future work

Tools

- scale to handle 40k ASes;
- reduce indicator limitations;
- use more than one BGP collector from RIS.

The next report

- will be published mid-2014;
- indicators will be enhanced;
- items will be written in English.

Questions?

Published material (French only)

- 2011 report;
- 2012 report;
- BGP configuration best practices.