Review of BGP BCP in 2014
Seen from RIS collectors

Guillaume Valadon

Agence nationale de la sécurité des systèmes d'information

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

RIPE 69 - November, 3rd 2014

The observatory in a nutshell

The observatory is under the supervision of the ANSSI, the French national cyberdefence agency. French operators and Afnic are also involved in the project.

Some of our objectives

- Study the Internet in France in details:
  - presented during RIPE 67 plenary.
- Develop technical interactions with the networking community;
- Publish anonymized results;
  - see `http://www.ssi.gouv.fr/observatoire/`
- Publish recommendations and best practices:
  - BGP BCP presented during RIPE 68 BCOP WG.

ANSSI BGP Best Current Practices guide

About the guide

- available at: `http://www.ssi.gouv.fr/en/the-anssi/events/new-publication-bgp-configuration-best-practices.html`
- written in collaboration with 7 French operators
- configuration examples for: IOS, Junos, SR-OS, OpenBGPD
  - contributions are welcome !

Recommendations examples

- authenticate BGP sessions with TCP-MD5
- filter the default route
- filter special AS numbers (private, documentation, ...)
- filter too specific prefixes: IPv4 > /24, IPv6 > /48
- limit the number of prefixes received from a peer

ANSSI BGP Best Current Practices guide

About the guide

- available at: `http://www.ssi.gouv.fr/en`
- written in collaboration with 7 French operators
- configuration examples for: IOS, Junos, SR-OS, OpenBGPD
  - contributions are welcome !

Recommendations examples

- authenticate BGP sessions with TCP-MD5
- filter the default route
- filter special AS numbers (private, documentation, ...)
- filter too specific prefixes: IPv4 > /24, IPv6 > /48
- limit the number of prefixes received from a peer

ANSSI BGP Best Current Practices guide

About the guide

Some BCP can be observed in routing tables !

Recommendations examples

- authenticate BGP sessions with TCP-MD5
- filter the default route
- filter special AS numbers (private, documentation, ...)
- filter too specific prefixes: IPv4 > /24, IPv6 > /48
- limit the number of prefixes received from a peer

Default routes seen by the RIS collectors

Default routes seen by RIS

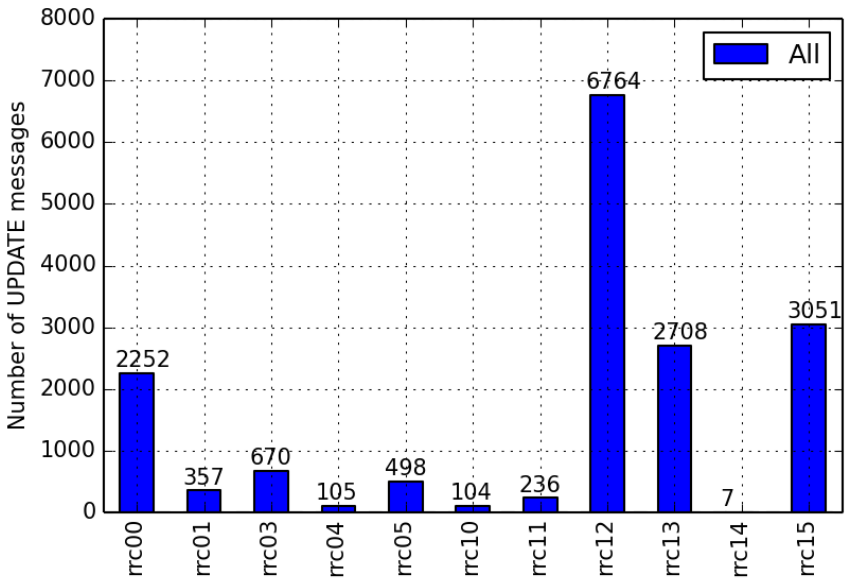

- ≈ 17000 UPDATEs received from January to September
- 11/13 active collectors received defaults

Some UPDATEs could be legitimate.

AS PATH length

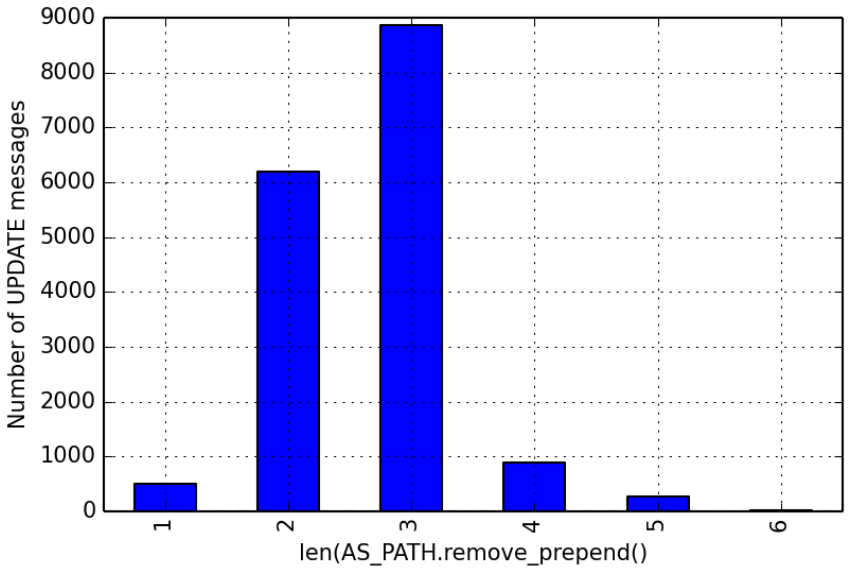

- $len() <= 2$: default announced by a RIS peer, or a transit provider of a RIS peer
- $len() > 2$: should not be seen
- 40% of the UPDATES have an AS PATH length strictly smaller than 3

Short AS PATH ($<= 2$) could identify legitimate announces.

Default routes seen by RIS - no short AS PATH

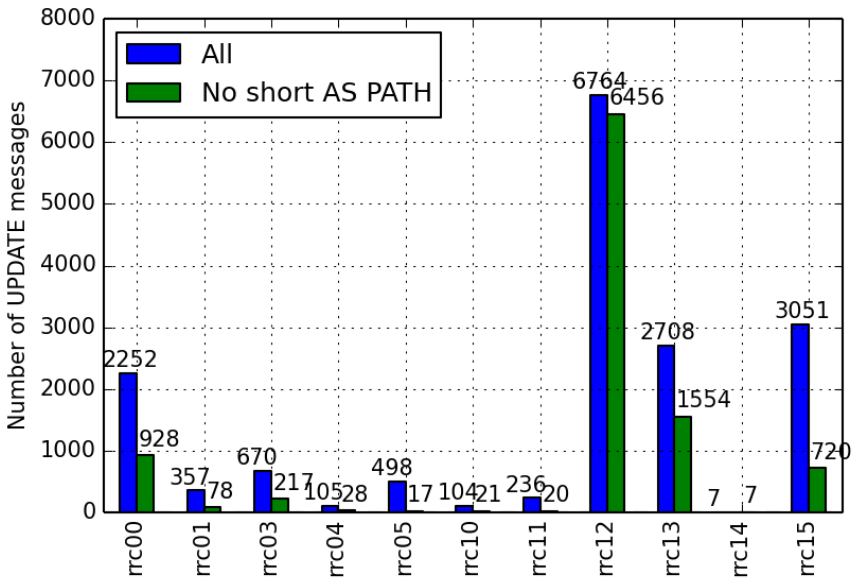

- ≈ 10000 UPDATEs received from January to September
  - IPv4: 12%
  - IPv6: 88%

Some collectors still received much more messages than the others.

Default routes per day

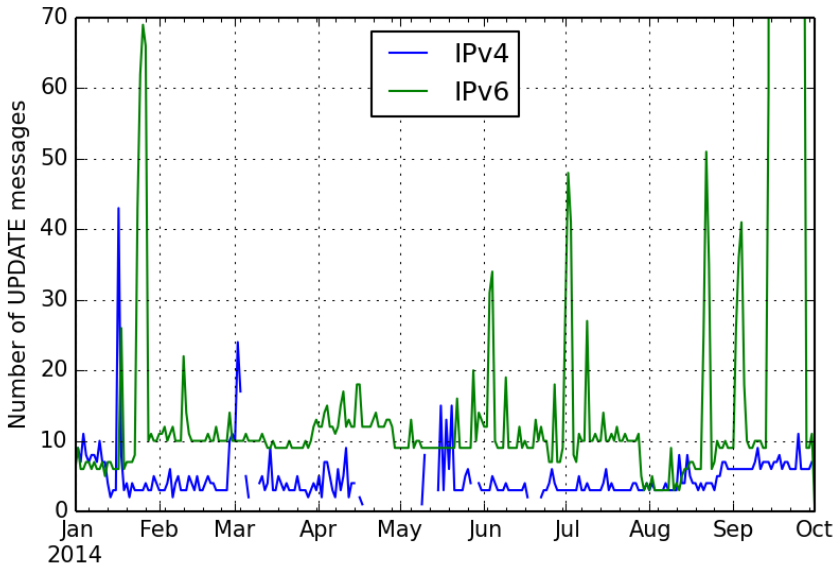

- IPv4: between 1 and 43 UPDATEs per day
  - some days no defaults are received
- IPv6: between 1 and 1436 UPDATEs per day
  - decrease at the end of September

Collectors see more IPv6 defaults than with IPv4.

Origin and transit AS

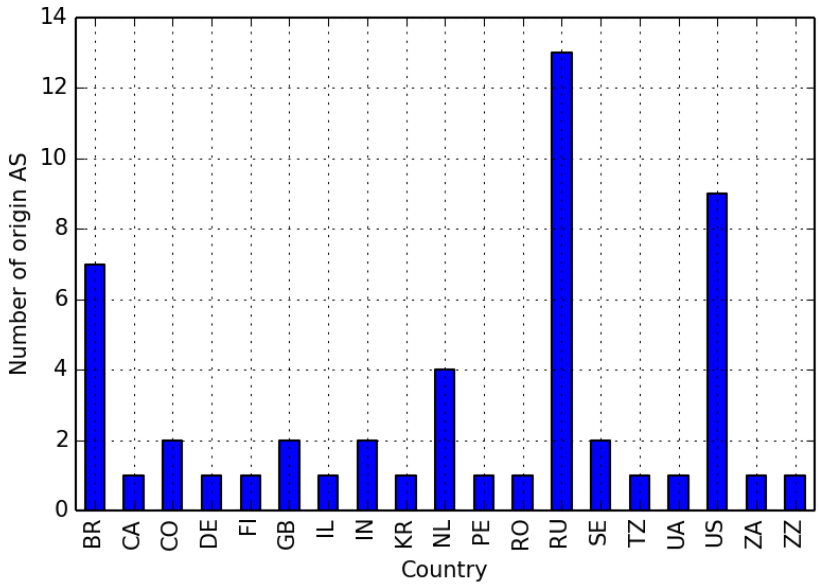

52 origin AS announced a default route

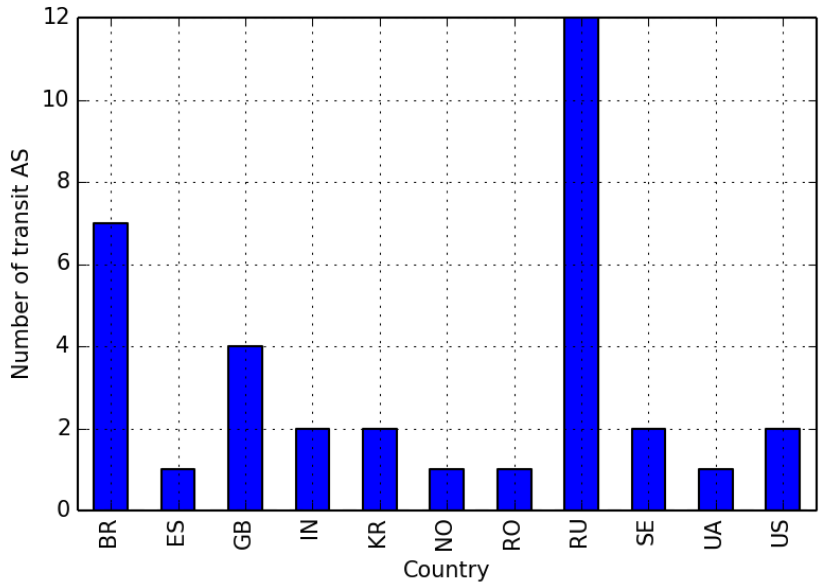

35 transit AS did not filter a default route

All of these transit providers should have filtered the default route.

Open questions

- do these UPDATEs are only seen by RIS collectors ?
- how many UPDATEs are seen by different RIS collectors ?
- ...

Too specific prefixes

Number of too specific prefixes

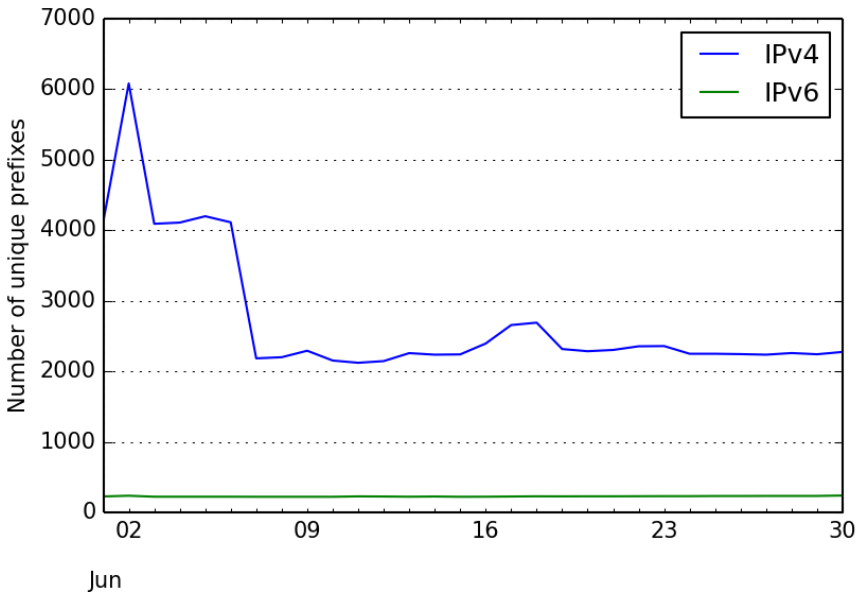

- IPv6: ≈ 200 distinct prefixes per day

≈ 2100 distinct prefixes seen every day.

Prefixes lengths

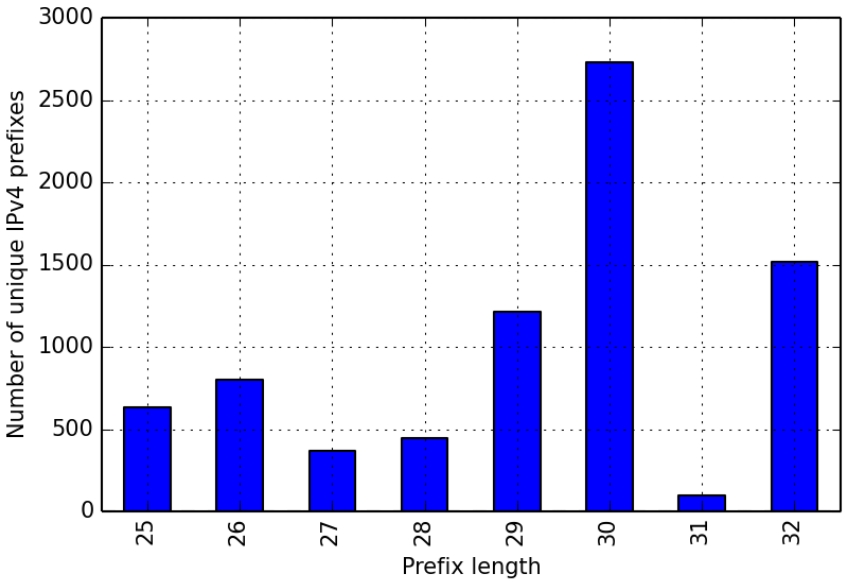

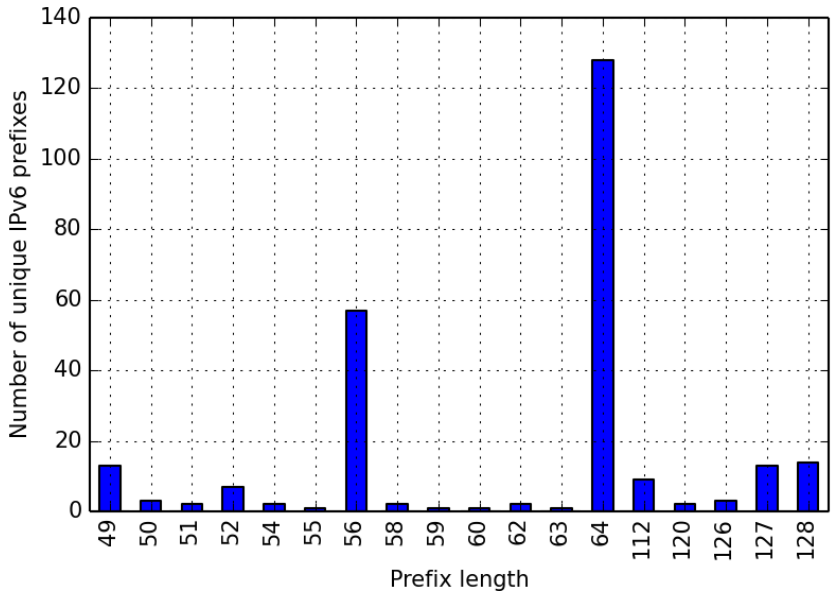

Unique IPv4 prefixes: 7797
Unique IPv6 prefixes: 261

Unique AS PATH length

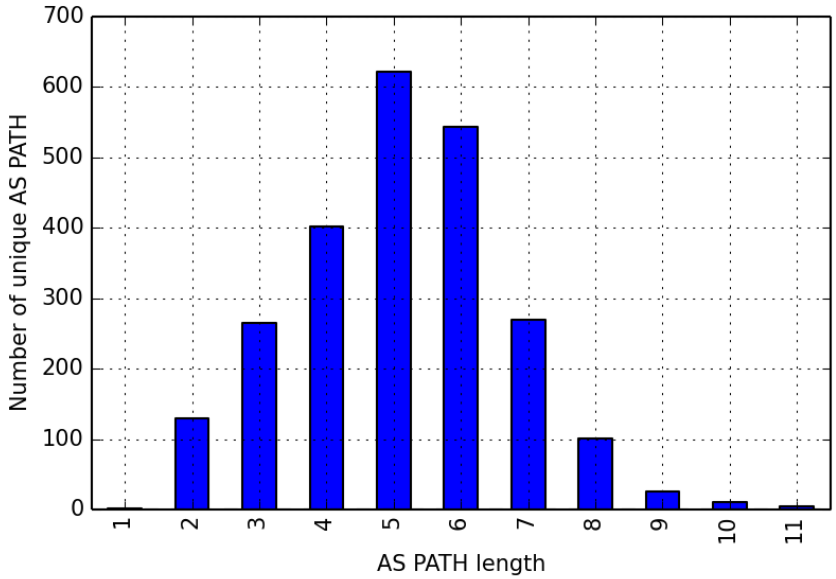

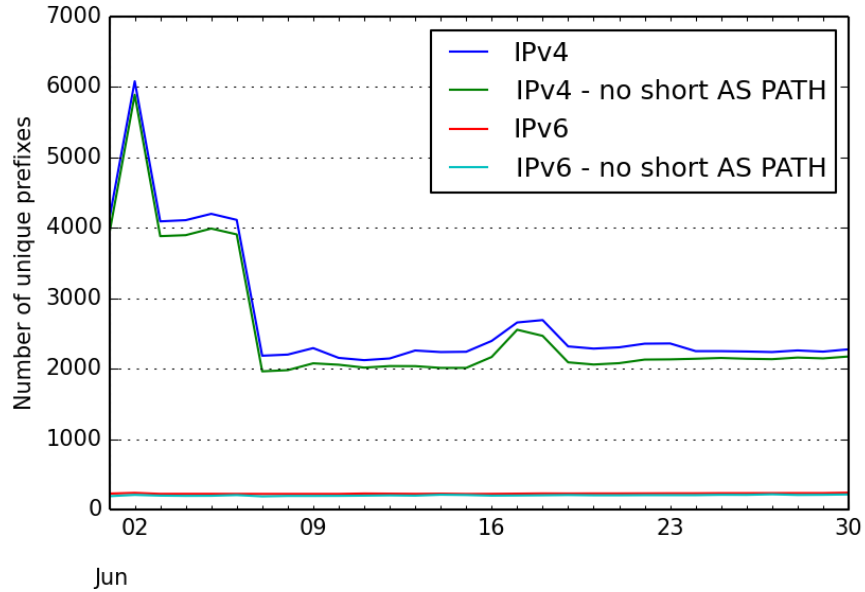

Most of the too specific prefixes cross the Internet.

Origin and transit ASes

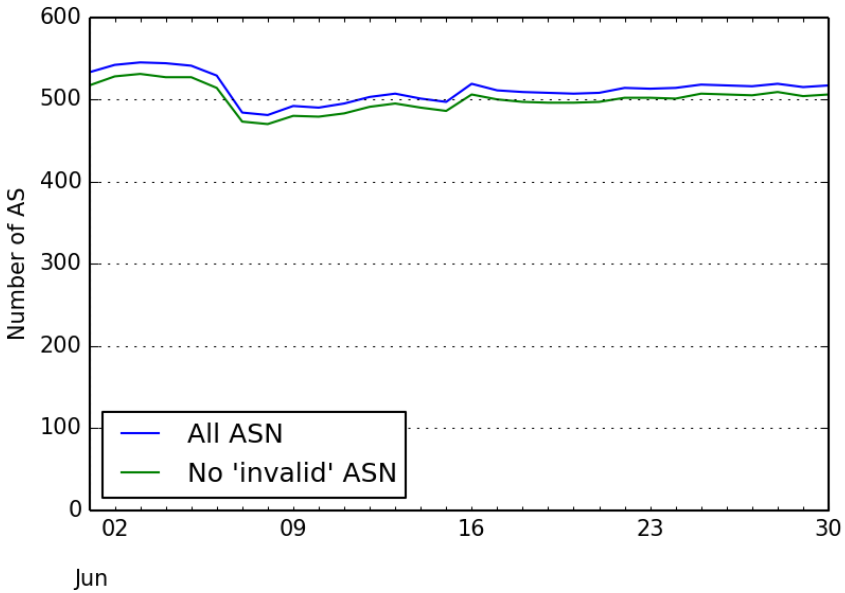

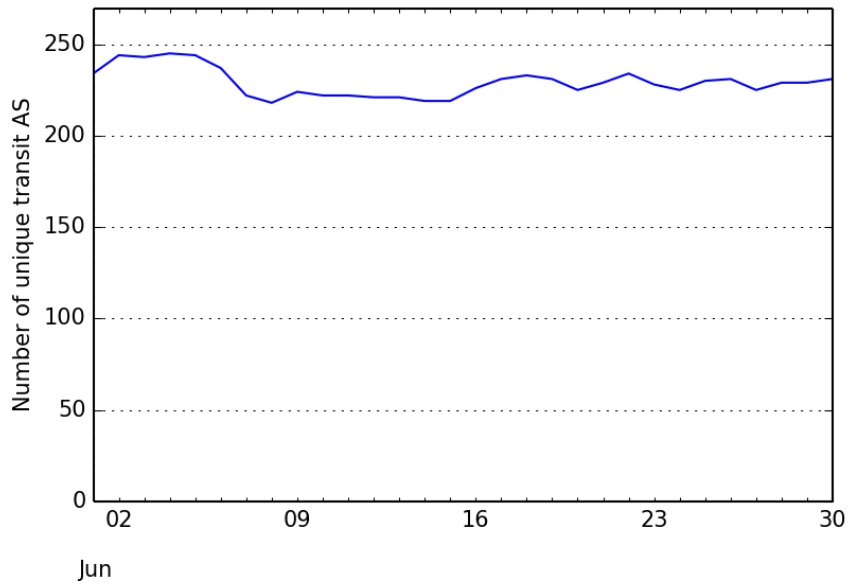

≈ 450 distinct origin AS seen every day.

≈ 200 transit AS seen every day.

Can these prefixes be reached otherwise ?

- on June 30th, there are 2089 unique too specific IP prefixes
- on July 1st: 125 prefixes can't be reached globally:
  - 46 are only reachable through the specific announce
  - 79 are not reachable at all

Most of the too specific prefixes can be reached by a less specific prefix.

Conclusion

Closing remarks

Still a work in progress !

- the observation of BCP adoption is a good awareness tool
- the same methodology can be applied to AS numbers, ...

  28220 3549 3356 8220 23456 198648

Will it be useful to contact operators ?

Questions?

Published material)

- 2011 report (French);
- 2012 report (French);
- 2013 report (French & English - soon);
- BGP configuration best practices (French & English).