

# Requirements for planned maintenance of BGP sessions

draft-dubois-bgp-pm-reqs-02.txt

[Nicolas.dubois@francetelecom.com](mailto:Nicolas.dubois@francetelecom.com)

[Bruno.decraene@francetelecom.com](mailto:Bruno.decraene@francetelecom.com)

[Benoit.fondeviole@francetelecom.com](mailto:Benoit.fondeviole@francetelecom.com)

[Zubair.ahmad@equant.com](mailto:Zubair.ahmad@equant.com)

# Context and objective

## Context

- Q Maintenance operations impacting forwarding plane
- Q Back up paths are available

## Objective

- Q To ease the maintenance of BGP sessions and limit the amount of traffic that is lost during planned maintenance operations on routers, a solution is required in order to gracefully shutdown a router or a session.

No packet lost when doing BGP router maintenance

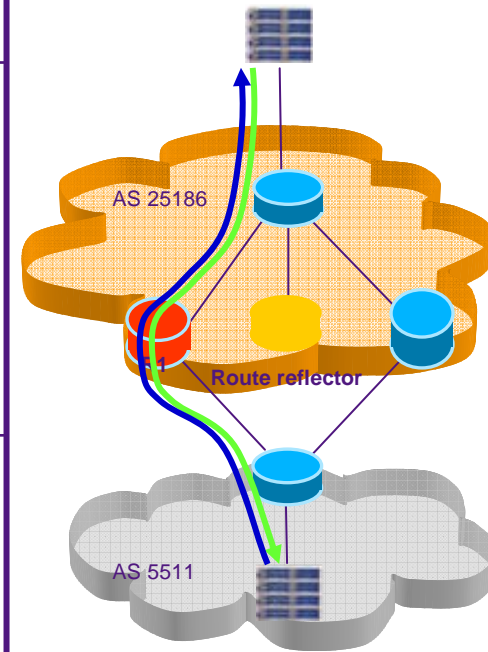
No operational burden in maintenance situations

# Current BGP Behavior

- S** When reloading, upgrading cards, shutting down a router, BGP withdraws existing routes and then informs its peers.
  - Q** Packets are lost.
- S** Current behavior is the following:
  1. Session is shut down.
  2. Forwarding is interrupted.
  3. Peers try to find an alternate route.
  4. As a result packets are lost during a few seconds (if not minutes)
- S** One would prefer to have a “make before break” behavior
  - Q** Peer would find an alternate route before the session, link or router disappears.

# Current BGP Behavior: Example

Test case	With 100 BGP routes	With 100 K BGP routes
Shutdown of R1	<p>Green: <b>4s</b> interruption Blue: <b>3s</b> interruption</p>	<p>Green: <b>20s</b> perturbation Blue: <b>0s</b> interruption</p>
Clear ip BGP R1	<p>Green: <b>1,1s</b> interruption Blue: <b>1s</b> interruption</p>	<p>Green: <b>40s</b> perturbation Blue: <b>20s</b> perturbation</p>



# Requirement Highlights

**S** The proposed solution SHOULD be applicable to all kinds of BGP sessions and any address family. The following cases are important:

- The maintenance of one particular e-BGP/MP-eBGP session.
- The reload/shutdown of one AS border router.
- The shutdown of PE <-> CE links (Static & eBGP) in a MPLS-VPN environment.

**S** Expected behavior

**Q**Step 1: A mechanism is implemented on the router under maintenance to gracefully reroute packets towards and from the BGP next-hop that is about to be unavailable.

**Q**Step 2: Once traffic is correctly rerouted, BGP sessions are shut down.

# Deployment Considerations

- S** The proposed solution **SHOULD** not change the BGP convergence behavior for the ASes exterior to the maintenance process
  
- S** An incremental deployment on a per AS basis **MUST** be possible.
  
- S** Transient routing loops
  - Q**Case of transient routing loops **MAY** be considered.

# Pending Issues

**S** Interaction with other (new) BGP mechanisms

**Q**BGP multi-path, communities, ORF?

**S** Does the requirement imply protocol changes ?

**Q**Can we find best practices that are applicable to most peering situations ?

**Q** Is it possible to use some existing extensions with a specific treatment in the router?

**S** Variety of situations to solve

**Q**In the draft we listed a set of use cases where hitless maintenance should be possible: Is that the way to go?

# Possible Next Steps

- S Interest for this work ?
- S Consensus for a new WG item ?
- S Please send your comments on the mailing list.



Thank you !