

# BGP Bestpath Selection Additional Criteria

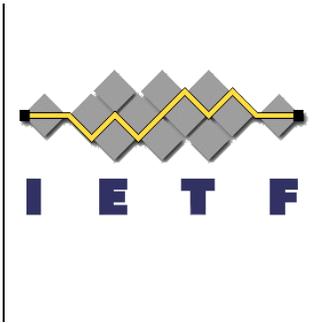
draft-ietf-idr-bgp-bestpath-selection-criteria-00

---

Rajiv Asati

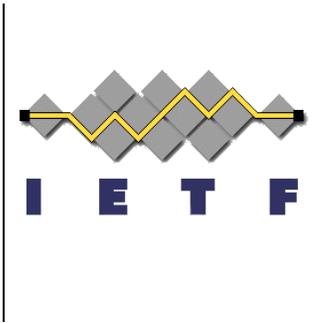


IETF 76, Hiroshima



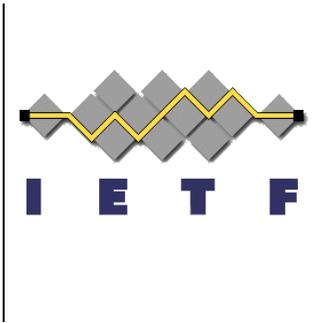
## Agenda

- Background / Motivation
- Problem Statement
- Solution
- Next Steps



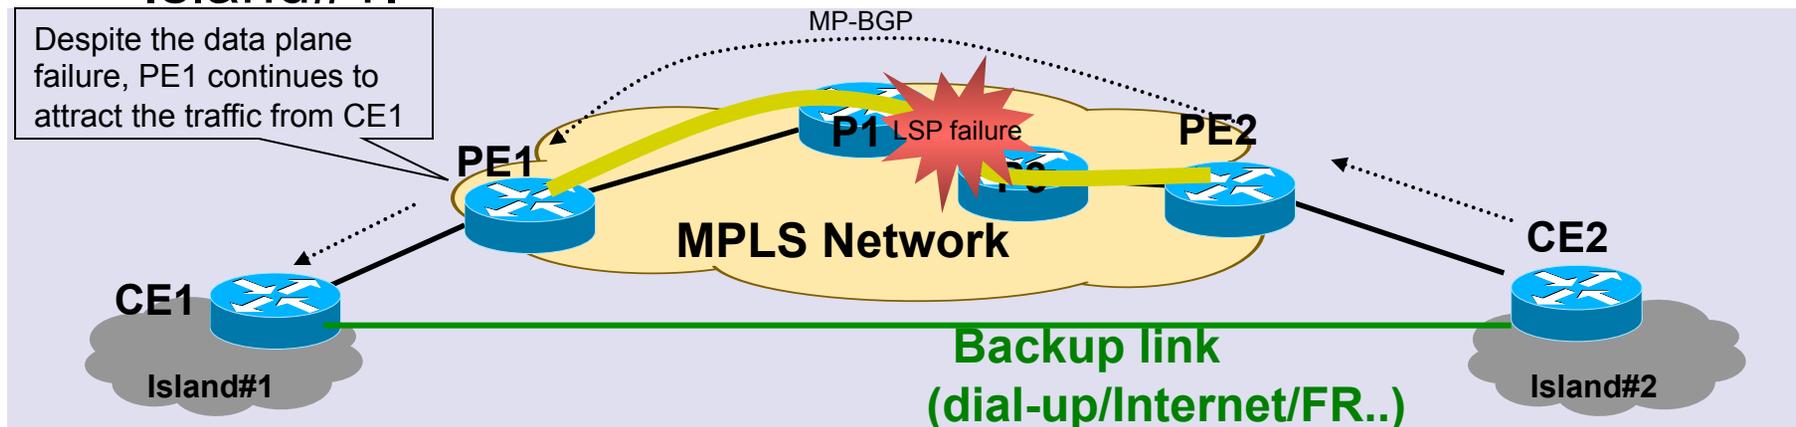
# Background / Motivation

- This draft is a result of WG feedback to ‘draft-asati-bgp-mpls-blackhole-avoidance-00’ that was presented at IETF68 last yr.
  - The feedback was that the problem was a generic BGP issue, not MPLS specific.
  - The feedback was that short/concise draft was needed.
- Ilya Varlashkin (Easynet), Chandra Appanna and John Scudder made the case for this draft.
- Yakov Rekhter helped out with extensive critical review.

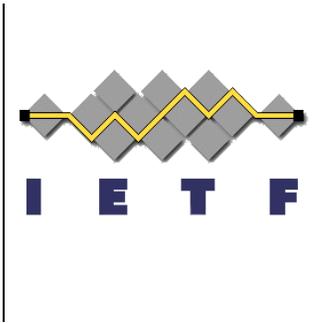


# Background / Motivation

- Deployment Scenario – Island#1 and #2 are connected via the MPLS network.
  - PE1->PE2 data plane failure (LSP failure, say) may result in blackholing of the island#1 to island#2 traffic.
- PE1 continues to advertise the Island#2 reachability to Island#1.

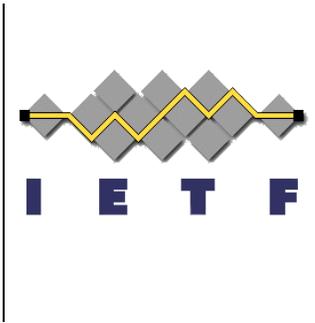


— PE1->PE2 LSP  
.....> Routing Protocol



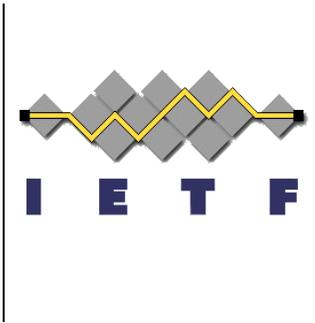
# Background / Motivation

- The BGP network pretends to have the reachability to the remote BGP prefixes.
- It is not fair to attract the (customer) traffic and **blackhole** it inside the BGP network.
  - Not fair or plain WRONG!!



# Problem Statement

- RFC4271 section 9.1.2 defines the BGP decision process aka bestpath selection, in which the 'Route Resolvability Condition' specifies checking for NEXT\_HOP route availability in the IP routing table.
- **This condition is not granular enough**, particularly for BGP networks that utilize data plane protocol other than IP.
- **The draft specifies two criteria for bestpath selection for further granularity.**



# Solution

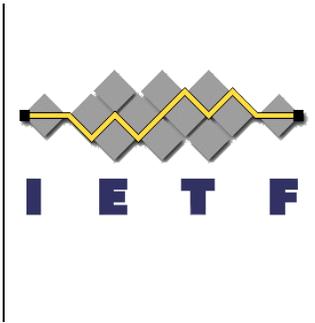
- Expand the route resolvability condition with the following two criteria –
  1. The reachability for the BGP next-hop SHOULD be resolved in a particular data plane protocol.
  2. The path availability check for the BGP next-hop MAY be performed.
- Draft has more text on this in the section 2.



# Solution – Criterion 1

The reachability for the BGP next-hop SHOULD be resolved in a particular data plane protocol.

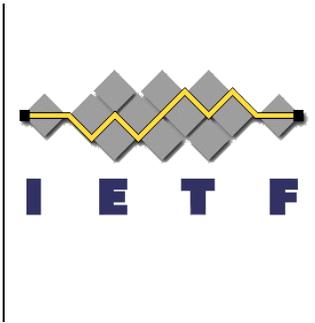
- The selection of particular data plane is a matter of a policy; Outside the scope of this document.
- Example :: if a BGP IPv4/v6 or VPNv4/v6 path wants to use MPLS data plane to the next-hop, as determined by the policy, then the BGP 'next-hop reachability' should be resolved using the MPLS data plane.



## Solution – Criterion 2

The path availability check for the BGP next-hop MAY be performed.

- This checks for the functioning path to the next-hop in a particular data plane protocol.
- The selection of particular data plane and the means to perform the path availability check are a matter of a policy; Outside the scope of this document.
- Example :: if a BGP VPNv4 path wants to use the MPLS as the data plane protocol to the next-hop, then liveness of MPLS LSP to the next-hop should be validated.



# Solution

- With these criteria, a BGP path may become the bestpath only if its NEXT\_HOP is reachable (and available) in the chosen data plane.
- Hence, the BGP speaker can either advertise or withdraw a BGP path.

# Next Step

- Get more WG feedback
- Assess WGLC.

