

BGP Data-Plane Benchmarking Applicable to Modern Routers

Ilya Varlashkin

Rajiv Papneja

Bhavani Parise

presented by Grégory CAUCHIE

Background/ Current Status

- Terminology for Benchmarking BGP Device Convergence in the (RFC 4098) / June 2005
- Methodology draft focusing on the BGP Data Plane convergence was posted prior IETF 80

<http://tools.ietf.org/html/draft-papneja-bgp-basic-dp-convergence-01>

- Not addressed (Missing item): Data-plane convergence on modern routers, which have local restoration capability
- Probable solution to address missing item: Proposing additional draft (may merge later) to be submitted prior to IETF 82


Elaborating Missing Item

- Assumption was: packet flow restored – protocol converged
 - This may NO LONGER be true – modern routers can restore traffic before protocols converge
 - May not be applicable when mapped to modern routers

Proposed Solution!

- Measure data-plane convergence at DUT (not behind) and do NOT call it BGP convergence
 - Relevant to modern world
 - By definition can be measured at data-plane
- Measure BGP convergence only behind DUT, not at
 - Still relevant even today
 - Capture protocol packets rather than data-plane
 - Use same reference clock for all timestamps

Key points of methodology



- Eliminate or minimize dependence on non-DUT
- DUT Detection time
- DUT Reaction time
- DUT Propagation time

Two drafts

- DUT ability to tell routers, which depend on it, about new network topology
 - Revised draft-papneja-bgp-basic-dp-convergence
 - Either exact but in control plane,
 - Or leveraging results from the new draft
- DUT ability to send packets that already arrived to it
 - New draft (pending upload to ietf.org)
 - Strictly data-plane

Current Draft Summary

(draft-papneja-bgp-basic-dp-convergence-01)

BGP Data plane FIB convergence for both IPv4 and IPv6

Limited to Basic BGP convergence (RFC 4271 functionality with Multi-Protocol BGP (MP-BGP) for IPv6)

BGP Failure/Convergence Events

Considers dependencies on factors impacting convergence:

- Number of peers,

- Number of routes/peers

- Policy Processing/Reconfiguration

Data Traffic characterization – offered load

Various test cases that covers iBGP, eBGP and failure convergence events

Topologies – Several 3 node, and 4 node setups

New draft summary

- Metrics derived from packets, not from clock
- Single test procedure, 5 failure scenarios
 - Loss of Signal on the link attached to DUT
 - Attached link failure without LoS
 - Non-direct link failure
 - Best route withdrawal
 - BGP next-hop failure
- Background topological noise for realistic result
- Large topology scenario

Next Steps

- Current draft will be revised to focus only on “behind DUT”
- Probably NOT data-plane
 - Metrics obtained directly, not derived
- Still data-plane
 - Use results from new draft

Action Items

- Can we agree on approach?
- Post revised old draft and new draft before IETF82
- Aim for WG-item readiness by IETF82
- Inputs welcome!