

BGP Data-Plane Benchmarking

Rajiv Papneja
Bhavani Parise
Sue Hares
Dean Lee
Ilya Varlashkin

Current Status

- New version of BGP Data Plane convergence after WG adoption was posted 2 weeks ago
 - draft-ietf-bmwg-bgp-basic-convergence-00
 - Author team: Rajiv Papneja, Bhavani Parise, Sue Hares, Dean Lee, Ilya Varlashkin
 - Presented and solicited feedback from IDR at IETF-83
 - Performed Benchmarking tests based on the proposed methodology at Ixia labs. Presented the findings to BMWG group and also at MPLS Ethernet World Congress 2013, Paris
 - Current update include:
 - Published as a worked group item
 - Addressed comments from BMWG and IDR group members
 - Addressed all the comments received from IDR chair
 - Comments/extensions received by SP
 - Test vendors have implemented this draft/methodology as automated protocol suite – used the test suite to perform tests on different implementations
 - Presented results in MPLS & Ethernet World Congress 2013, Paris: highlights presented as part of this presentation
 - Applicable to VPNv4 route convergence benchmarking

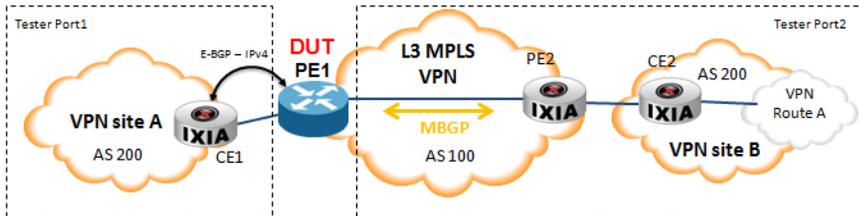
Highlights from MPLS World Congress Presentation

- Test Environment
 - Involved testing in 1-2 node Black Box settings
 - Used tester to generate traffic, routes and routers emulation, packet sniffing, throughput and timestamp measurement
- Target Test Cases
 - Address family convergence
 - RIB-IN
 - RIB-OUT
 - Attribute Change
 - Comparing the impact of label stack
- Results – Analysis & Comparison

Test Setup used

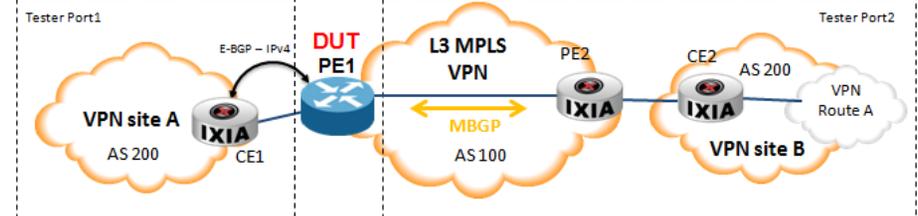
RIB-IN Convergence

RIB-OUT Convergence



Objective - To measure the convergence time required to receive and install a VPN route in RIB using MBGP

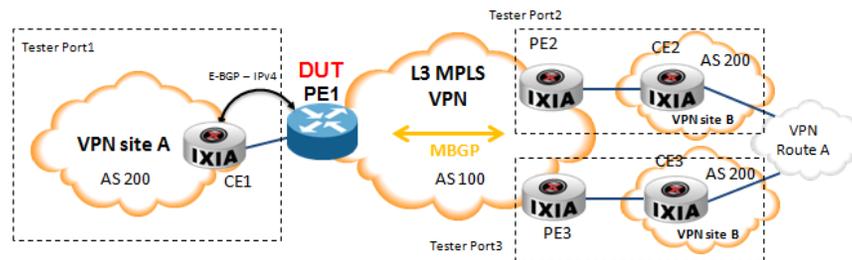
- Test steps –
 - CE1 sends traffic toward VPN route A
 - PE2 advertises VPN route A at RCV-Rt-time
 - PE2 receives traffic for VPN route A at DUT-XMT-Data-Time
 - RIB-IN Convergence time = DUT-XMT-Data-Time - RCV-Rt-time



Objective - This test measures the convergence time taken to receive, install and advertise a VPN route using eBGP

- Test steps –
 - PE2 advertises VPN route A at RCV-Rt-time
 - DUT-PE1 forwards VPN route A toward CE1 at DUT-XMT-Rt-Time
 - RIB-Out Convergence = (DUT-XMT-Rt-Time – DUT-RCV-Rt-Time)

BGP Path Attribute Change Convergence

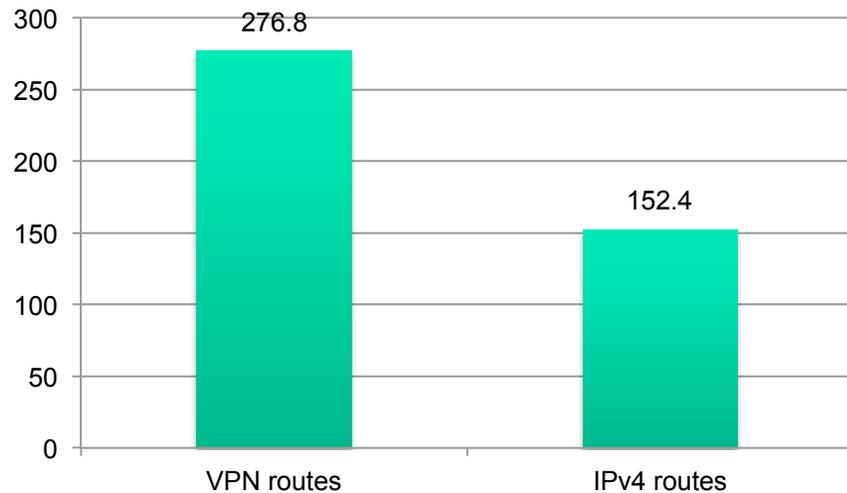


Objective - This test measures the convergence time taken by an implementation to service a MBGP Path Attribute Change

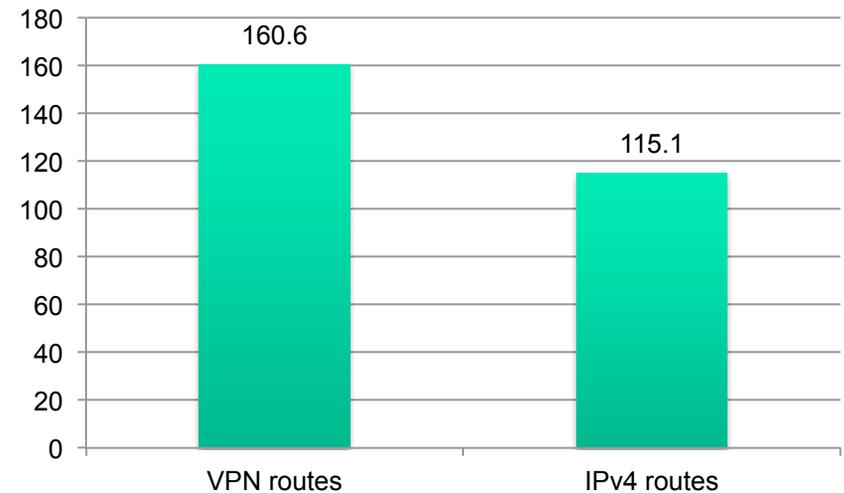
- Test steps –
 - PE3 advertises the VPN route A. PE3 is the optimal path
 - CE1 sends traffic toward VPN route A
 - PE2 advertises the same VPN route A with better attribute at "Path-Change-Event-Time"
 - PE2 receives traffic at "Path-Switch-Time"
 - BGP convergence time = "Path-Switch-Time" - "Path-Change-Event-Time"

Address Family Convergences

RIB-IN Convergence (ms)



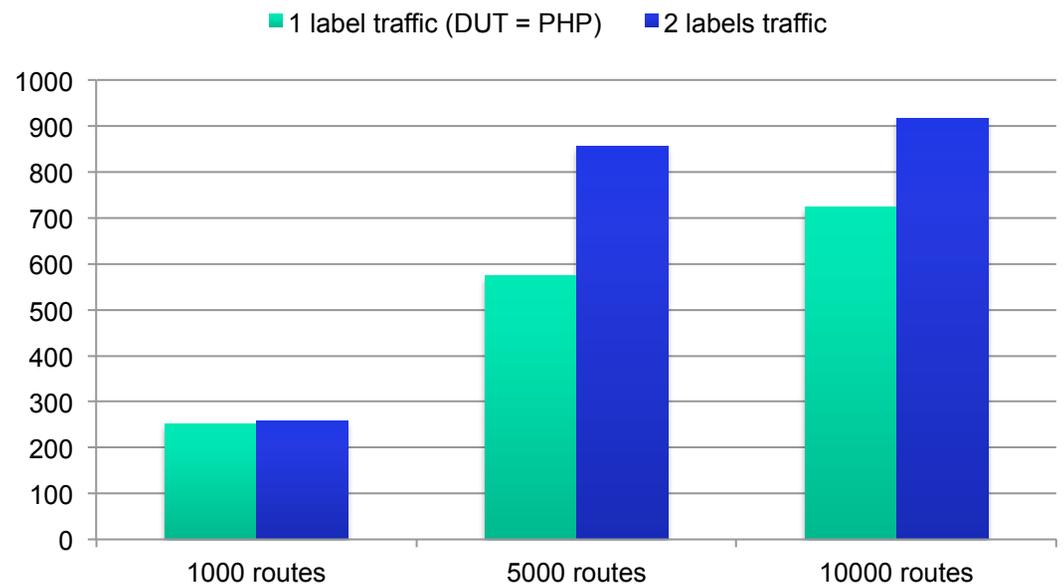
RIB-OUT Convergence (ms)



Compare convergence performance of different BGP address families with repeatable test results

BGP Path Attribute Change Convergence - Results

5.7 BGP Attribute Change Convergence	
Iteration	Time (ms)
1	257.3
2	256
3	263
4	253.2
5	257.3
Average	257.3
Deviation	2.2



Impact of attribute change on convergence performance with different label stack

Next Steps

- Solicit feedback from fellow BMWGers
- Progress the draft to informational RFC