

Network Working Group

Internet-Draft

Updates: ???? (if approved)

Intended status: Informational

Expires: January 3, 2020

A. Morton

J. Uttaro

AT&T Labs

July 2, 2019

Benchmarks and Methods for Multihomed EVPN

draft-morton-bmwg-multihome-evpn-02

Abstract

Fundamental Benchmarking Methodologies for Network Interconnect Devices of interest to the IETF are defined in RFC 2544. Key benchmarks applicable to restoration and multi-homed sites are in RFC 6894. This memo applies these methods to Multihomed nodes implemented on Ethernet Virtual Private Networks (EVPN).

Comments from Sudhin Jacob

- Thanks! Detailed Reply on the BMWG- List.
- Summary of Changes in 02 follows.

Clarifications on Connections between SUT and Tester

The tester SHALL be connected to all CE and every PE, and be capable of simultaneously sending and receiving frames on all ports with connectivity. The tester SHALL be capable of generating multiple flows (according to a 5-tuple definition, or any sub-set of the 5-tuple). The tester SHALL be able to control the IP capacity of sets of individual flows, and the presence of sets of flows on specific interface ports.

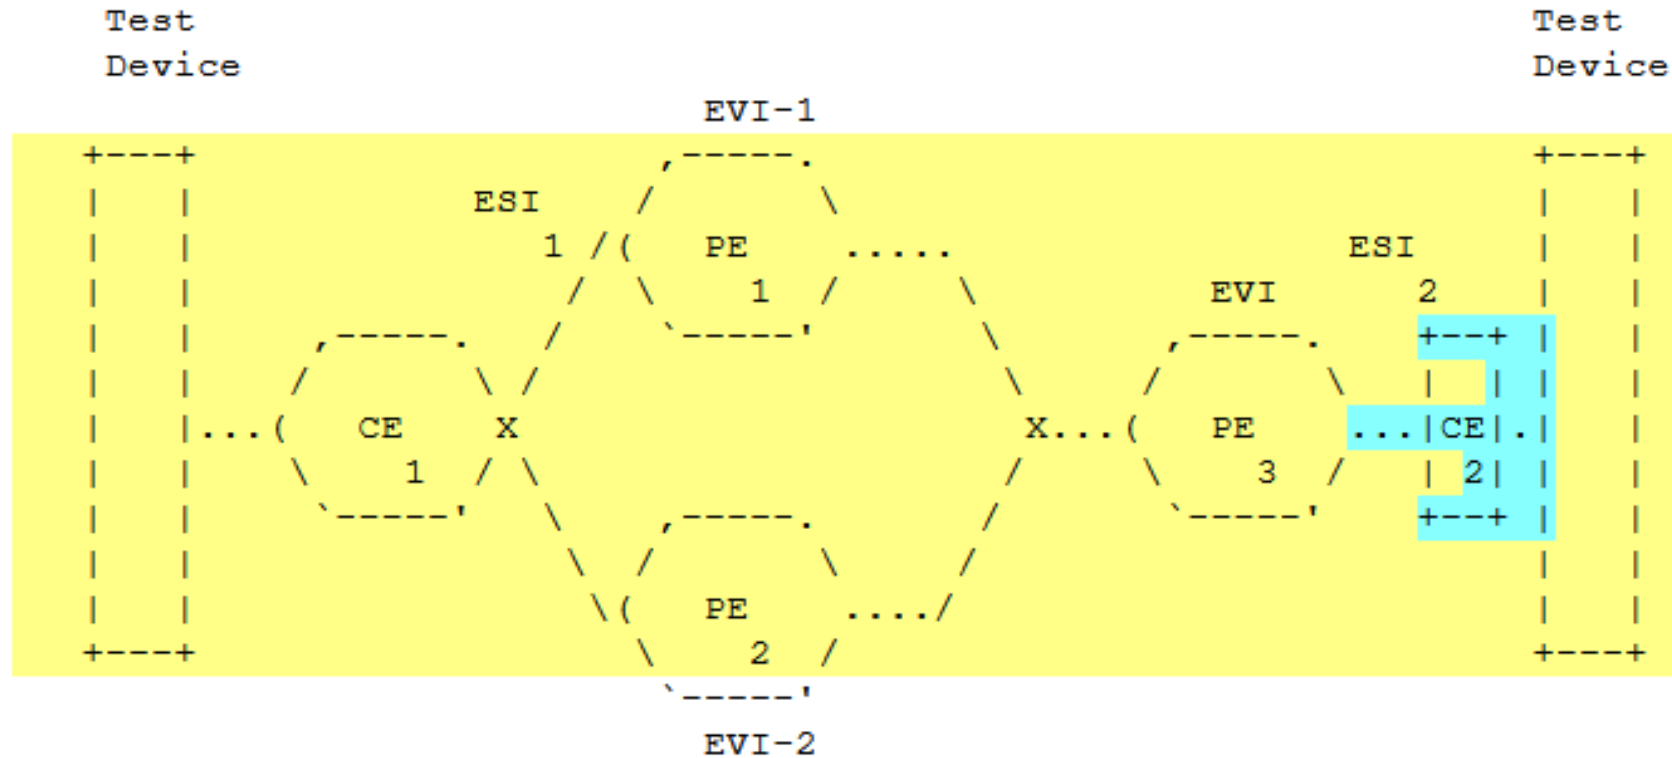
The tester SHALL be capable of generating and receiving a full mesh of Unicast flows, as described in section 3.0 of [RFC2889]:

"In fully meshed traffic, each interface of a DUT/SUT is set up to both receive and transmit frames to all the other interfaces under test."

Other mandatory testing aspects described in [RFC2544] and [RFC2889] MUST be included, unless explicitly modified in the next section.

Clarifications on Connections between SUT and Tester

A second test case is where a BGP backbone implements MPLS-LDP to provide connectivity between multiple PE - ESI - CE locations.



5. Procedure for Full Mesh Throughput Char.

5.2. Test for a Single Frame Size and Number of Unicast Flows

Each trial in the test requires configuring a number of flows (from 100 to 100k) and a fixed frame size (64 octets to 128, 256, 512, 1024, 1280 and 1518 bytes, as per [RFC2544]). Frame formats MUST be specified, they are as described in section 4 of [RFC2889].

5.3. Detailed Procedure

The Procedure SHALL follow section 5.1 of [RFC2889].

Specifically, the Throughput measurement parameters found in section 5.1.2 of [RFC2889] SHALL be configured and reported with the results.

The procedure for transmitting Frames on each port is described in section 5.1.3 of [RFC2889] and SHALL be followed (adapting to the number of ports in the test setup).

Once the traffic is started, the procedure for Measurements described in section 5.1.4 of [RFC2889] SHALL be followed (adapting to the number of ports in the test setup). The section on Throughput measurement (5.1.4 of [RFC2889]) SHALL be followed.

In the case that one or more of the CE and PE are virtual implementations, then the search algorithm of [TST009] that provides consistent results when faced with host transient activity SHOULD be used (Binary Search with Loss Verification).

Other Comments

- Lets discuss (or draw the solution)

Suggestions:-

If the remote PE is Single home and there must be a router tester connected to it.

[acm]

Are you suggesting that we remove the single-homed CE_2 in Figure 2? We can do that, but then we lose a simple and symmetrical test case.

Another method is if we place a single home interface and multihome interface in the same routing instance which reduces the other routers/elements.

[acm]

So, a new setup like Figure 2 should include a CE which is single-homed to PE1 with ESI=0 ?? This new scenario could be used to reduce the test equipment needed, and it is supported by the EVPN technology. (We could do this in a new Figure 3)

Or, the CE would have two links to the same Router or PE?

We're not sure how that affects the performance we are trying to measure...

Thanks for your attention today!