

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 18 February 2024

S. Jacob, Ed.
Independent
K. Tiruveedhula
Juniper Networks
17 August 2023

Benchmarking Methodology for EVPN and PBB-EVPN
draft-ietf-bmwg-evpntest-11

Abstract

This document defines methodologies for benchmarking EVPN and PBB-EVPN performance. EVPN is defined in RFC 7432, and is being deployed in Service Provider networks. Specifically, this document defines the methodologies for benchmarking EVPN/PBB-EVPN convergence, data plane performance, and control plane performance.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 18 February 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
1.2. Terminologies	3
2. Test Topology	4
3. Test Cases for EVPN Benchmarking	7
3.1. Data Plane MAC Learning	7
3.2. Control Plane MAC Learning	8
3.3. MAC Flush-Local Link Failure and Relearning	9
3.4. MAC Flush-Remote Link Failure and Relearning	10
3.5. MAC Aging	11
3.6. Remote MAC Aging	11
3.7. Control and Data Plane MAC Learning	12
3.8. High Availability	13
3.9. MAC+IP/MAC+IPv6 Learning Rate	14
3.10. Scaling of Services	15
3.11. Scale Convergence	15
3.12. SOAK Test.	16
4. Test Cases for PBB-EVPN Benchmarking	17
4.1. Data Plane Local MAC Learning	17
4.2. Data Plane Remote MAC Learning	18
4.3. MAC Flush-Local Link Failure	18
4.4. MAC Flush-Remote Link Failure	19
4.5. MAC Aging	20
4.6. Remote MAC Aging.	21
4.7. Local and Remote MAC Learning	21
4.8. High Availability	22
4.9. Scale	23
4.10. Scale Convergence	23
4.11. Soak Test	24
5. Acknowledgments	25
6. IANA Considerations	25
7. Security Considerations	26
8. References	26
8.1. Normative References	26
8.2. Informative References	26
Appendix A. Appendix	26
Authors' Addresses	26

1. Introduction

EVPN is defined in RFC 7432, and describes BGP MPLS based Ethernet VPNs (EVPN). PBB-EVPN is defined in RFC 7623, discusses how Ethernet Provider backbone Bridging can be combined with EVPNs to provide a new/combined solution. This draft defines methodologies that can be used to benchmark both RFC 7432 and RFC 7623 solutions. Further, this draft provides methodologies for benchmarking the performance of

EVPN data and control planes, MAC learning, MAC flushing, MAC aging, convergence, high availability, and scale.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminologies

Most of the terminology used in this documents comes from [RFC7432] and [RFC7632].

CE: Customer Router/Devices/Switch.

DF: Designated Forwarder.

DUT: Device under test.

Ethernet Segment (ES): When a customer site (device or network) is connected to one or more PEs via a set of Ethernet links, then that set of links is referred to as an 'Ethernet segment'.

EVI: An EVPN instance spanning the Provider Edge (PE) devices participating in that EVPN.

Ethernet Segment Identifier (ESI): A unique non-zero identifier that identifies an Ethernet segment is called an 'Ethernet Segment Identifier'.

Ethernet Tag: An Ethernet tag identifies a particular broadcast domain, e.g., a VLAN. An EVPN instance consists of one or more broadcast domains.

Interface: Physical interface of a router/switch.

IRB: Integrated routing and bridging.

MAC: Media Access Control addresses on a PE.

MHPE2: Multi homed Provider Edge router 2.

MHPE1: Multi homed Provider Edge router 1.

SHPE3: Single homed Provider Edge Router 3.

PE: Provider Edge device.

P: Provider Router.

RR: Route Reflector.

pps: Packets per sec.

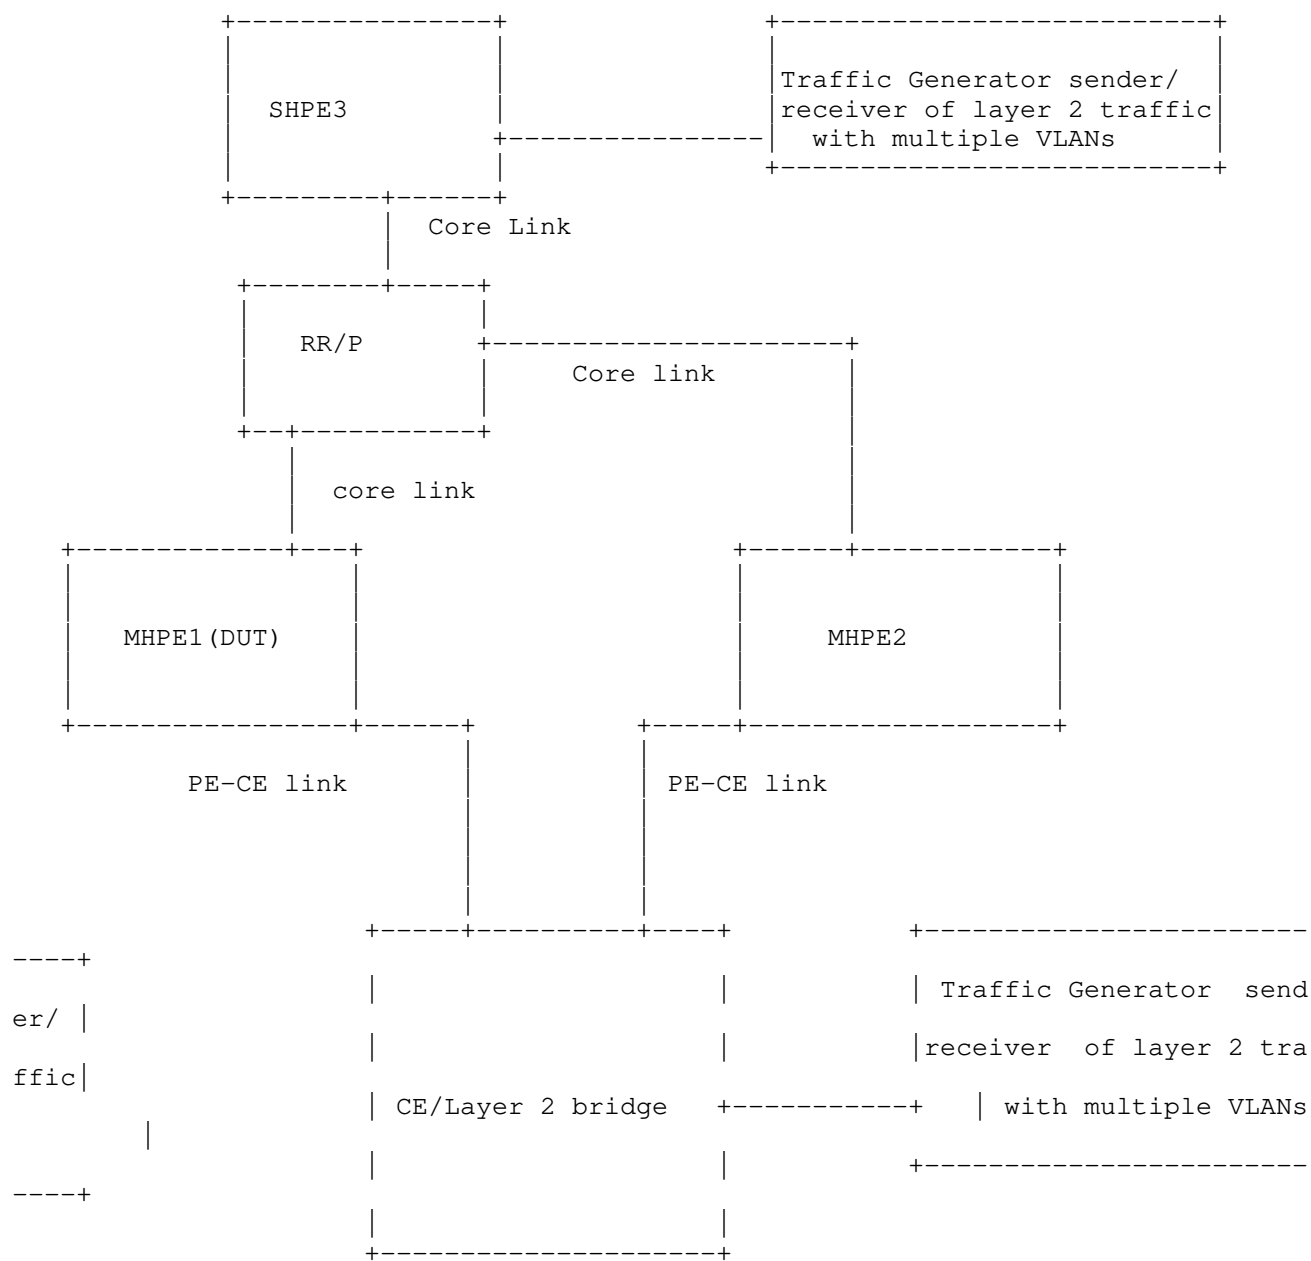
RT/Traffic Generator: Traffic Generator to generate traffic.

SA: Single Active

Single-Active Redundancy Mode: When a single PE (among all the PEs attached to an Ethernet segment) is the only PE allowed to forward traffic to/from a given Ethernet segment for a given VLAN, then that Ethernet segment is defined to be operating in Single-Active redundancy mode.

2. Test Topology

There are five routers in the Test setup. SHPE3, RR/P, MHPE1 and MHPE2 emulating a service provider network. CE is a customer device connected to MHPE1 and MHPE2. it is configured with bridge domains in multiple VLANs. The traffic generator is connected to the CE and SHPE3. The MHPE1 acts as DUT. The traffic generator will be used as sender and receiver of traffic. The test measurements are taken from the DUT. MHPE1 and MHPE2 are multi-homed routers connected to CE running single active mode. The traffic generator will be generating traffic at 10% of the line rate.



Topology 1

Figure 1

Test Setup

Mode	Test	Traffic Direction	Sender	Receiver	
Single Active	Local MAC Learning	Uni	CE	SHPE3	Layer 2 traffic multiple M
Single Active	Remote MAC Learning	Uni	CE	SHPE3	Layer 2 traffic multiple M
Single Active	Scale Convergence	Bi	CE/SHPE3	CE/SHPE3	Layer 2 traffic multiple M VLANs

Figure 2

Table showing the traffic directions of various EVPN/PBB-EVPN benchmarking test cases. Depending on the test scenario, the traffic can be uni-directional or bi-directional (configured in the traffic generator).

Test Setup Configurations:

SHPE3 is configured with Interior Gateway protocols like OSPF or IS-

IS for the underlay, LDP for MPLS support, and Interior Border Gateway with EVPN address family for overlay support. This router must be configured with N EVPN/PBB-EVPN instances for testing. A traffic generator is connected to this router for sending and receiving traffic.

RR is configured with Interior Gateway protocols like OSPF or IS-IS for the underlay, LDP for MPLS support, and Interior Border Gateway with EVPN address family for overlay support. This router functions as both a provider router and a route reflector.

MHPE1 is configured with Interior Gateway protocols like OSPF or IS-IS for the underlay, LDP for MPLS support, and Interior Border Gateway with EVPN address family for overlay support. This router must be configured with N EVPN/PBB-EVPN instances for testing. It is also configured with ESI per VLAN or ESI per interface. MHPE1 functions as a multi-homing PE in Single Active EVPN mode. This router serves as the Device Under Test (DUT) and is connected to the CE.

MHPE2 is configured with Interior Gateway protocols like OSPF or IS-IS for the underlay, LDP for MPLS support, and Interior Border Gateway with EVPN address family for overlay support. This router must be configured with N EVPN/PBB-EVPN instances for testing. It is also configured with ESI per VLAN or ESI per interface. MHPE2 functions as a standby router during the tests and is connected to the CE.

CE is acting as a bridge configured with multiple VLANs. The same VLANs are configured on MHPE1, MHPE2, and SHPE3. A traffic generator is connected to the CE, which acts as both a sender and receiver of traffic.

Depending on the test scenarios, the traffic generators will be used to generate uni-directional or bi-directional flows.

The above configuration will serve as the base configuration for all test cases.

The 'X' below is used as a variable to denote the scale factor of the testing parameters. It must be a multiple of 100.

'N' is an arbitrary number used to collect statistically significant samples. The range can vary from 5 to 15.

3. Test Cases for EVPN Benchmarking

3.1. Data Plane MAC Learning

Objective: Measure the time taken to learn the Data Plane MAC in the DUT.

Topology: Topology 1

Procedure:

The data plane MAC learning can be measured using the parameters defined in RFC 2889 section 5.8.

Confirm that the DUT is up and running with EVPN.

The traffic generator connected to CE must send frames with X different source and destination MAC addresses for one VLAN. The same VLAN must be present in all the devices except RR.

Send X unicast frames from CE to MHPE1 (DUT) for one EVPN instance working in Single Active (SA) mode.

The DUT will learn these X MAC addresses in the data plane.

Measurement: Measure the time taken to learn X MAC addresses locally in the DUT's EVPN MAC table. The data plane measurement is taken by considering the DUT as a black box. The range of MAC addresses is known from the traffic generator; the same must be learned in the DUT. The time taken to learn X MAC addresses is measured. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp. Note that at present, there are no standard models to measure the counters.

The test is repeated N times, and the values are collected. The MAC learning rate is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn.

MAC learning rate = (T1 + T2 + ... + Tn) / N

3.2. Control Plane MAC Learning

Objective: Measure the time taken to learn the control plane MAC.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

The traffic generator connected to SHPE3 must send frames with X different source and destination MAC addresses for one VLAN. The same VLAN must be present in all devices except RR.

Ensure the frames are destined for one EVPN instance.

The DUT will learn these X MAC addresses in the control plane.

Measurement: Measure the time taken by the DUT to learn the X MAC addresses in the control plane. The test is repeated N times, and the values are collected. The remote MAC learning rate is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp. Note that there are currently no standard models to measure the counters.

MAC learning rate = (T1 + T2 + ... + Tn) / N

3.3. MAC Flush-Local Link Failure and Relearning

Objective: Measure the time taken to flush the Data Plane MAC and the time taken to relearn the same amount of MAC.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

Send X frames with X different source and destination MAC addresses to the DUT from CE using the traffic generator for one VLAN.

Ensure the DUT learns all X MAC addresses in the data plane.

Fail the link between DUT and CE, and measure the time taken to flush these X MAC addresses from the EVPN MAC table.

Bring up the previously failed DUT-CE link and measure the time taken by the DUT to relearn these X MAC addresses.

The DUT and MHPE2 are both running in Single Active (SA) mode.

Measurement: Measure the time taken for flushing these X MAC addresses. Measure the time taken to relearn these X MAC addresses in the DUT. The test is repeated N times, and the values are collected. The flush and relearning times are calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn, R1, R2, ..., Rn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp.

$$\text{Flush rate} = (T1 + T2 + \dots + Tn) / N$$
$$\text{Relearning rate} = (R1 + R2 + \dots + Rn) / N$$

3.4. MAC Flush-Remote Link Failure and Relearning

Objective: Measure the time taken to flush the Control plane MAC learned in DUT during remote link failure and the time taken to relearn.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

Send X frames with X different source and destination MAC addresses to the DUT from SHPE3 using the traffic generator for one VLAN.

Bring down the link between SHPE3 and the traffic generator.

SHPE3 will withdraw the routes from DUT due to the link failure.

Measure the time taken to flush the DUT's EVPN MAC table. The DUT and MHPE2 are both running in Single Active (SA) mode.

Bring up the previously failed link between SHPE3 and the traffic generator.

Measure the time taken by the DUT to relearn these X MAC addresses from the control plane.

Measurement: Measure the time taken to flush X remote MAC addresses from the EVPN MAC table of the DUT. Measure the time taken to relearn these X MAC addresses in the DUT. The test is repeated N times, and the values are collected. The flush rate is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn, R1, R2, ..., Rn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp.

$$\text{Flush rate} = (T1 + T2 + \dots + Tn) / N$$
$$\text{Relearning rate} = (R1 + R2 + \dots + Rn) / N$$

3.5. MAC Aging

Objective: To measure the MAC Aging time.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

Send X frames with X different source and destination MAC addresses to the DUT from CE using the traffic generator for one VLAN.

Ensure that these X MAC addresses are learned in the DUT.

Then stop the traffic.

Ensure that the DUT and other devices in the test are using the default timers for aging.

Measure the time taken to flush X MAC addresses from the DUT's EVPN MAC table due to aging.

The DUT and MHPE2 are both running in Single Active (SA) mode.

Measurement: Measure the time taken to flush X MAC addresses due to aging. The test is repeated N times, and the values are collected. The aging time is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp.

Aging time for X MAC addresses in seconds = $(T1 + T2 + \dots + Tn) / N$

3.6. Remote MAC Aging

Objective: Measure the control plane learned MAC Aging time.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

Send X frames with X different source and destination MAC addresses to the DUT from SHPE3 using the traffic generator for one VLAN.

Ensure that these X MAC addresses are learned in the DUT via control plane.

Then stop the traffic.

Ensure that the DUT and other devices in the test are using the default timers for aging.

Measure the time taken to flush X MAC addresses from the DUT's EVPN MAC table due to aging.

The DUT and MHPE2 are both running in Single Active (SA) mode.

Measurement: Measure the time taken to flush X remote MAC addresses learned in the DUT's EVPN MAC table due to aging. The test is repeated N times, and the values are collected. The aging is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp.

Aging time for X MAC addresses in seconds = $(T1 + T2 + \dots + Tn) / N$

3.7. Control and Data Plane MAC Learning

Objective: To record the time taken to learn both local and remote MAC.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

Send X frames with X different source and destination MAC addresses to the DUT from SHPE3 using the traffic generator for one VLAN.

Send X frames with different source and destination MAC addresses from the traffic generator connected to CE for one VLAN.

The source and destination addresses of flows must be complementary to have unicast flows.

Measure the time taken by the DUT to learn 2X MAC addresses in the EVPN MAC table.

The DUT and MHPE2 are both running in Single Active (SA) mode.

Measurement: Measure the time taken to learn 2X MAC addresses in the DUT's EVPN MAC table. The test is repeated N times, and the values are collected. The MAC learning time is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp.

MAC learning rate = $(T1 + T2 + \dots + Tn) / N$

3.8. High Availability

Objective: Measure traffic loss during routing engine failover.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

Send X frames from CE to DUT using the traffic generator with X different source and destination MAC addresses.

Send X frames from the traffic generator to SHPE3 with X different source and destination MAC addresses, resulting in 2X MAC addresses being learned in the DUT.

Set up bidirectional traffic flow with X packets per second (pps) in each direction.

Ensure that the DUT learns 2X MAC addresses.

Then perform a routing engine failover.

Measurement: The expectation of the test is to observe 0 traffic loss with no change in the Designated Forwarder (DF) role. The DUT should not withdraw any routes. However, if the DUT is not properly synchronized between the master and standby, resulting in packet loss, measure that packet loss. The test is repeated N times, and the values are collected. The packet loss is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn. This measurement is carried out using an external server that polls the DUT using automated scripts to ensure the DUT learned 2X MAC addresses. The packet drop is measured using the traffic generator.

Packet loss in seconds with 2X MAC addresses = $(T1 + T2 + \dots + Tn) / N$

3.9. MAC+IP/MAC+IPv6 Learning Rate

Measure the DUT learning rate of MAC+IP/MAC+IPv6.

Objective: Learning rate measurement.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with EVPN.

Send X ARP/Neighbor Discovery (ND) requests per second from the traffic generator to the DUT with different sender source MAC addresses to the target IRB address configured in the EVPN instance.

Measure the time taken by the DUT to learn X MAC+IP/MAC+IPv6 in the EVPN MAC table.

The DUT and MHPE2 are both running in Single Active (SA) mode.

Measurement: Measure the time taken to learn X MAC+IP/MAC+IPv6 addresses in the DUT's EVPN MAC table. The test is repeated N times, and the values are collected. The MAC learning time is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the counter value with a timestamp. The measurements are taken for MAC+IP and MAC+IPv6 separately.

MAC+IP Learning Rate = $(T1 + T2 + \dots + Tn) / N$

MAC+IPv6 Learning Rate = $(T1 + T2 + \dots + Tn) / N$

3.10. Scaling of Services

Objective: Measure the scale of EVPN instances that a DUT can hold.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with EVPN.

The DUT, MHPE2 and SHPE3 are scaled to N EVI.

Ensure routes received from MHPE2 and SHPE3 for N EVI in the DUT.

Then increment the scale of N by 5% of N till the limit is reached.

The limit is where the DUT cant learn any EVPN routes from its peers.

Measurement : There should not be any loss of route types 1,2,3 and 4 in DUT. DUT must relearn all type 1, 2, 3 and 4 from remote routers. The DUT must be subjected to various values of N to find the optimal scale limit. The scope of the test is find out the maximum EVPN instance that a DUT can hold. The measurement is carried out using external server which polls the DUT using automated scripts to find the scale limit of EVPN instances. Once the limit is reached the DUT will no longer learn any incremental type 1,2,3 routes.

3.11. Scale Convergence

Objective: Measure the convergence time of DUT when the DUT is scaled with EVPN instance along with traffic.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with EVPN.

Scale N EVIs in DUT, SHPE3 and MHPE2.

Send F frames to DUT from CE using traffic generator with X different source and destination MAC for N EVI's.

Send F frames from traffic generator to SHPE3 with X different source and destination MAC.

There will be 2X number of MAC will be learned in DUT EVPN MAC table.

There is a bi directional traffic flow with F pps in each direction.

Then clear the BGP neighbours in the DUT.

Once the BGP session is in established state in DUT.

Measure the time taken to learn 2X MAC in DUT MAC table.

Measurement : The DUT must learn 2X MAC. Measure the time taken to learn 2X MAC in DUT. The test is repeated for N times and the values are collected. The convergence time is calculated by averaging the values obtained by N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1, T2...Tn. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

Time taken to learn 2X MAC in DUT = $(T1+T2+...Tn)/N$

3.12. SOAK Test.

Objective: This test is carried out to measure the stability of the DUT in a scaled environment with traffic over a period of time "T' ". In each interval "t1" the DUT CPU usage, memory usage are measured. The DUT is checked for any crashes till the end of the cycle which is 24 hours.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with EVPN.

Scale N EVI's in DUT, SHPE3 and MHPE2. Send F frames to DUT from CE using traffic generator with different X source and destination MAC for N EVI's.

Send F frames from traffic generator to SHPE3 with X different source and destination MAC.

There will be 2X number of MAC will be learned in DUT EVPN MAC table.

There is a bi directional traffic flow with F pps in each direction.

The DUT must run with traffic for 24 hours.

Every hour check for memory leak in EVPN process, CPU usage and crashes in DUT, this must be repeated till the end of the cycle.

If the DUT crashed during the test and it is not able to recover, then the test will be terminated.

Measurement : Take the hourly reading of CPU, process memory. There should not be any memory leak, crashes, CPU spikes. The CPU spike is determined as the sudden increase of CPU usage to 100 percent compared to the average usage. The average value varies from device to device. Memory leak is determined as the increase of memory usage with respect to time. The expectation is under steady state the memory usage for EVPN process should not increase with respect to time. The measurement is carried out using external server which polls the DUT using automated scripts which captures the CPU usage and process memory. The data collected from the 24 hour duration must be evaluated to see any leaks and crashes during the entire period.

4. Test Cases for PBB-EVPN Benchmarking

4.1. Data Plane Local MAC Learning

Objective: Measure the time taken to learn the Data Plane MAC in DUT.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with PBB-EVPN.

Traffic generator connected to CE must send frames with X different source and destination MAC for one VLAN, the same VLAN must be present in all the devices except RR.

Send X unicast frames from CE to MHPE1(DUT) for one PBB-EVPN instance working in SA mode.

The DUT will learn these X MAC in data plane.

Measurement : Measure the time taken to learn X MAC locally in DUT PBB-EVPN MAC table. The data plane measurement is taken by considering DUT as black box. The range of MAC are known from traffic generator, the same must be learned in DUT, the time taken to learn X MAC is measured. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

The test is repeated for N times and the values are collected. The MAC learning rate is calculated by averaging the values obtained from N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1,T2...Tn.

MAC learning rate = $(T1+T2+..Tn)/N$

4.2. Data Plane Remote MAC Learning

Objective: To Record the time taken to learn the remote MAC.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with PBB-EVPN.

Traffic generator connected to SHPE3 must send frames with X different source and destination MAC for one VLAN, the same VLAN must be present in all the devices except RR.

Ensure the frames must be destined to one PBB-EVPN instance.

The DUT will learn these X MAC in data plane.

Measurement : Measure the time taken by the DUT to learn the X MAC in the data plane. The test is repeated for N times and the values are collected. The remote MAC learning rate is calculated by averaging the values obtained from N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1,T2...Tn. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

MAC learning rate = $(T1+T2+..Tn)/N$

4.3. MAC Flush-Local Link Failure

Objective: Measure the time taken to flush the locally learned MAC and the time taken to relearn the same amount of MAC.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with PBB-EVPN.

Send X frames with X different source and destination MAC to DUT from CE using traffic generator for one VLAN.

Ensure the DUT learns all X MAC in data plane.

Fail the DUT-CE link and measure the time taken to flush these X MAC from the PBB-EVPN MAC table.

Bring up the link which was made Down(the link between DUT and CE).Measure time taken by the DUT to relearn these X MAC.

The DUT and MHPE2 are running SA mode.

Measurement : Measure the time taken for flushing these X MAC. Measure the time taken to relearn these X MAC in DUT. The test is repeated for N times and the values are collected. The flush and the relearning time is calculated by averaging the values obtained by N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1,T2...Tn,R1,R2...Rn. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

Flush rate = $(T1+T2+..Tn)/N$

Relearning rate = $(R1+R2+..Rn)/N$

4.4. MAC Flush-Remote Link Failure

Objective: Measure the time taken to flush the remote MAC learned in DUT due to remote link failure and relearning it.

Topology : Topology 1

Procedure:

confirm the DUT is up and running with PBB-EVPN.

Send X frames with X different source and destination MAC to DUT from SHPE3 using traffic generator for one VLAN.

Bring down the link between SHPE3 and traffic generator.

Measure the time taken to flush the DUT PBB-EVPN MAC table. The DUT and MHPE2 are running SA mode.

Bring up the link which was made Down(the link between SHPE3 and traffic generator).

Measure time taken by the DUT to relearn these X MAC

Measurement : Measure the time taken to flush X remote MAC from PBB-EVPN MAC table of the DUT. Measure the time taken to relearn these X MAC in DUT. The test is repeated for N times and the values are collected. The flush rate is calculated by averaging the values obtained by N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1,T2...Tn,R1,R2..Rn. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

Flush rate = $(T1+T2+...Tn)/N$

Relearning rate = $(R1+R2+...Rn)/N$

4.5. MAC Aging

Objective: Measure the MAC Aging time.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with PBB-EVPN.

Send X frames with X different source and destination MAC to DUT from CE using traffic generator for one VLAN.

Ensure these X MAC are learned in DUT.

Then stop the traffic.

Ensure the DUT and other devices in the test are using the default timers for aging.

Measure the time taken to flush X MAC from DUT PBB-EVPN MAC table due to aging.

The DUT and MHPE2 are running SA mode.

Measurement : Measure the time taken to flush X MAC due to aging. The test is repeated for N times and the values are collected. The aging is calculated averaging the values obtained by N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1,T2...Tn. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

Aging time for X MAC in sec = $(T1+T2+...Tn)/N$

4.6. Remote MAC Aging.

Objective: Measure the remote MAC Aging time.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with PBB-EVPN.

Send X frames with X different source and destination MAC to DUT from SHPE3 using traffic generator for one VLAN.

Ensure these X MAC are learned in DUT.

Then stop the traffic.

Ensure the DUT and other devices in the test are using the default timers for aging.

Measure the time taken to flush X MAC from DUT PBB-EVPN MAC table due to aging.

The DUT and MHPE2 are running SA mode.

Measurement : Measure the time taken to flush X remote MAC learned in DUT EVPN MAC table due to aging. The test is repeated for N times and the values are collected. The aging is calculated by averaging the values obtained by N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1,T2...Tn. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

Aging time for X MAC in sec = $(T1+T2+...Tn)/N$

4.7. Local and Remote MAC Learning

Objective: Measure the time taken to learn both local and remote MAC.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with PBB-EVPN.

Send X frames with X different source and destination MAC to DUT from SHPE3 using traffic generator for one VLAN.

Send X frames with different source and destination MAC from traffic generator connected to CE for one VLAN.

The source and destination addresses of flows must be complimentary to have unicast flows.

Measure the time taken by the DUT to learn 2X in PBB-EVPN MAC table.

DUT and MHPE2 are running in SA mode.

Measurement : Measure the time taken to learn 2X MAC in DUT PBB-EVPN MAC table. The test is repeated for N times and the values are collected. The MAC learning time is calculated by averaging the values obtained by N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1,T2...Tn. The measurement is carried out using external server which polls the DUT using automated scripts which measures the MAC counter value with timestamp.

MAC learning rate = $(T1+T2+..Tn)/N$

4.8. High Availability

Objective: Measure traffic loss during routing engine failover.

Topology : Topology 1

Procedure:

Confirm the DUT is up and running with PBB-EVPN.

Send X frames from CE to DUT from traffic generator with X different source and destination MAC.

Send X frames from traffic generator to SHPE3 with X different source and destination MAC, so that 2X MAC will be learned in the DUT.

There is a bi directional traffic flow with X pps in each direction.

Ensure the DUT learn 2X MAC.

Then do a routing engine fail-over.

Measurement : The expectation of the test is 0 traffic loss with no change in the DF role. DUT should not withdraw any routes. But in cases where the DUT is not properly synchronised between master and standby, due to that packet loss are observed. In that scenario the packet loss is measured. The test is repeated for N times and the values are collected. The packet loss is calculated by averaging the values obtained by N samples. N is an arbitrary number to get a sufficient sample. The time measured for each sample is denoted by T1, T2...Tn. The measurement is carried out using external server which polls the DUT using automated scripts to ensure the DUT learned 2X MAC. The packet drop is measured using traffic generator.

Packet loss in sec with 2X MAC = $(T1+T2+...Tn)/N$

4.9. Scale

Objective: Measure the scale limit of the DUT for PBB-EVPN.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with PBB-EVPN.

Scale the DUT, MHPE2, and SHPE3 to N PBB-EVI instances.

Ensure that routes are received from MHPE2 and SHPE3 for N PBB-EVI instances in the DUT.

Increment the scale by 5% of N until the limit is reached.

The limit is reached when the DUT can no longer learn any EVPN routes from its peers.

Measurement: There should be no loss of route types 2, 3, and 4 in the DUT. The DUT must relearn all type 2, 3, and 4 routes from remote routers. The DUT should be subjected to various values of N to find the optimal scale limit. The scope of the test is to determine the maximum number of EVPN instances that a DUT can handle. The measurement is carried out using an external server that polls the DUT using automated scripts to determine the scale limit of PBB-EVPN instances.

4.10. Scale Convergence

Objective: To measure the convergence time of the DUT when it's scaled with EVPN instances along with traffic.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with PBB-EVPN.

Scale N PBB-EVIs in DUT, SHPE3, and MHPE2.

Send F frames to DUT from CE using the traffic generator with X different source and destination MAC addresses for N PBB-EVIs.

Send F frames from the traffic generator to SHPE3 with X different source and destination MAC addresses.

There will be 2X number of MAC addresses learned in the DUT's PBB-EVPN MAC table.

Set up bidirectional traffic flow with F packets per second (pps) in each direction.

Then clear the BGP neighbors in the DUT.

Once the BGP session is established in the DUT, measure the time taken to learn 2X MAC addresses in the DUT's MAC table.

Measurement: The DUT must learn 2X MAC addresses. Measure the time taken to learn 2X MAC addresses in the DUT. The test is repeated N times, and the values are collected. The convergence time is calculated by averaging the values obtained from N samples. N is an arbitrary number to obtain a sufficient sample size. The time measured for each sample is denoted by T1, T2, ..., Tn. This measurement is carried out using an external server that polls the DUT using automated scripts, which records the MAC counter value with a timestamp.

Time taken to learn 2X MAC addresses in DUT = $(T1 + T2 + \dots + Tn) / N$

4.11. Soak Test

Objective: To measure the stability of the DUT in a scaled environment with traffic.

Topology: Topology 1

Procedure:

Confirm that the DUT is up and running with PBB-EVPN.

Scale N PBB-EVIs in DUT, SHPE3, and MHPE2.

Send F frames to DUT from CE using the traffic generator with X different source and destination MAC addresses for N EVIs.

Send F frames from the traffic generator to SHPE3 with X different source and destination MAC addresses.

There will be 2X number of MAC addresses learned in the DUT's PBB-EVPN MAC table.

Set up bidirectional traffic flow with F packets per second (pps) in each direction.

The DUT must run with traffic for 24 hours.

Every hour, check for memory leaks in the PBB-EVPN process, CPU usage, and crashes in the DUT until the end of the 24-hour cycle.

If the DUT crashes during the test and is unable to recover, the test will be terminated.

Measurement: Take hourly readings of CPU usage and process memory. There should be no memory leaks, crashes, or CPU spikes. A CPU spike is defined as a sudden increase in CPU usage to 100 percent compared to the average usage. The average value may vary from device to device. A memory leak is defined as an increase in memory usage over time. Under steady-state conditions, the memory usage for the PBB-EVPN process should not increase significantly. The measurement is carried out using an external server that polls the DUT using automated scripts, capturing the CPU usage and process memory.

5. Acknowledgments

We would like to thank Fioccola Giuseppe of Telecom Italia reviewing our draft and commenting it. We would like to thank Sarah Banks for guiding and mentoring us. We take the opportunity to thank Al for reviewing our draft and gave us valuable comments.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

The benchmarking tests described in this document are limited to the performance characterisation of controllers in a lab environment with isolated networks. The benchmarking network topology will be an independent test setup and MUST NOT be connected to devices that may forward the test traffic into a production network or misroute traffic to the test management network. Further, benchmarking is performed on a "black-box" basis, relying solely on measurements observable external to the controller. Security features mentioned in the RFC 7432 will affect the test results. Special capabilities SHOULD NOT exist in the controller specifically for benchmarking purposes. Any implications for network security arising from the controller SHOULD be identical in the lab and in production networks.

8. References

8.1. Normative References

- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <<https://www.rfc-editor.org/info/rfc2544>>.
- [RFC2899] Ginoza, S., "Request for Comments Summary RFC Numbers 2800-2899", RFC 2899, DOI 10.17487/RFC2899, May 2001, <<https://www.rfc-editor.org/info/rfc2899>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

8.2. Informative References

- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7623] Sajassi, A., Ed., Salam, S., Bitar, N., Isaac, A., and W. Henderickx, "Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)", RFC 7623, DOI 10.17487/RFC7623, September 2015, <<https://www.rfc-editor.org/info/rfc7623>>.

Appendix A. Appendix

Authors' Addresses

Sudhin Jacob (editor)
Independent
Bangalore
India
Phone: +91 8061212543
Email: sudhinjacob@rediffmail.com

Kishore Tiruveedhula
Juniper Networks
10 Technology Park Dr
Westford, MA 01886
United States of America
Phone: +1 9785898861
Email: kishoret@juniper.net