

Network Working Group

Internet Draft

S. Dry  
F. Calabria  
I.Y Fung  
Cisco  
M. Napierala  
AT&T  
Y. Kamite  
NTT Communications

Expires: May 2008

November 9, 2007

Multicast VPN Scalability Benchmarking  
draft-sdry-bmwg-mvpnscale-03.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on May 9, 2008.

Abstract

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

Multicast VPN (MVPN) is a service deployed by VPN service providers to enable their customers to use IP multicast applications over VPNs. With the increased popularity the scalability of deploying such a service is becoming of a great interest. This document defines standard metric and test methodology for characterizing and comparing control plane MVPN scalability of Provider Edge (PE) devices that implement MVPN service.

## Table of Contents

<a href="#">1</a>	<a href="#">Introduction.....</a>	<a href="#">3</a>
<a href="#">2</a>	<a href="#">Document Scope.....</a>	<a href="#">4</a>
<a href="#">3</a>	<a href="#">Terminology.....</a>	<a href="#">5</a>
<a href="#">4</a>	<a href="#">Key Words to Reflect Requirements.....</a>	<a href="#">7</a>
<a href="#">5</a>	<a href="#">MVPN Metric Definition.....</a>	<a href="#">8</a>
<a href="#">6</a>	<a href="#">Test Environment.....</a>	<a href="#">8</a>
<a href="#">6.1</a>	<a href="#">Test Topologies.....</a>	<a href="#">8</a>
<a href="#">6.2</a>	<a href="#">Unicast Control Plane Setup.....</a>	<a href="#">9</a>
<a href="#">6.3</a>	<a href="#">Multicast Control Plane Setup.....</a>	<a href="#">10</a>
<a href="#">6.4</a>	<a href="#">Considerations for Number of Receivers behind remote PEs</a>	<a href="#">11</a>
<a href="#">6.5</a>	<a href="#">Data Traffic Characteristics.....</a>	<a href="#">11</a>
<a href="#">6.6</a>	<a href="#">Test Apparatus Considerations.....</a>	<a href="#">12</a>
<a href="#">7</a>	<a href="#">Test Categories, Stimulus and Execution Methodology.....</a>	<a href="#">12</a>
<a href="#">7.1</a>	<a href="#">Steady State Testing Execution Methodology.....</a>	<a href="#">14</a>
<a href="#">7.2</a>	<a href="#">Failure Recovery Testing Execution Methodology.....</a>	<a href="#">15</a>
<a href="#">8</a>	<a href="#">Results Content and Reporting Format.....</a>	<a href="#">16</a>
<a href="#">8.1</a>	<a href="#">Steady State Testing.....</a>	<a href="#">16</a>
<a href="#">8.2</a>	<a href="#">Failure Recovery Testing.....</a>	<a href="#">17</a>
<a href="#">9</a>	<a href="#">Test Cases Common to all MVPN Architectures.....</a>	<a href="#">18</a>
<a href="#">9.1</a>	<a href="#">"Empty" MVPNs Scale.....</a>	<a href="#">18</a>
<a href="#">9.2</a>	<a href="#">PIM Enabled VPN C-Interfaces Scale.....</a>	<a href="#">19</a>
<a href="#">9.3</a>	<a href="#">PIM C-instances Mroutes Scale.....</a>	<a href="#">20</a>
<a href="#">9.4</a>	<a href="#">PIM C-Instances OIF Scale.....</a>	<a href="#">22</a>
<a href="#">9.5</a>	<a href="#">Joined S-PMSI (Data MDT) Scale.....</a>	<a href="#">24</a>
<a href="#">9.6</a>	<a href="#">Sourced S-PMSI (Data MDT) Scale.....</a>	<a href="#">25</a>
<a href="#">9.7</a>	<a href="#">S-PMSI (Data MDT) Reuse.....</a>	<a href="#">26</a>
<a href="#">9.8</a>	<a href="#">Scale of mVPNs spanning large number of PEs.....</a>	<a href="#">28</a>
<a href="#">9.9</a>	<a href="#">Scale of mVPNs with larger amount of state.....</a>	<a href="#">30</a>
<a href="#">9.10</a>	<a href="#">Scale of "average" size mVPNs.....</a>	<a href="#">31</a>
<a href="#">9.11</a>	<a href="#">S-PMSI Switching Delay.....</a>	<a href="#">32</a>
<a href="#">9.12</a>	<a href="#">Convergence of C-Instance PIM Joins.....</a>	<a href="#">33</a>
<a href="#">9.13</a>	<a href="#">Effect of Co-locating C-RPs on a PE.....</a>	<a href="#">34</a>

<a href="#">10</a>	Test Cases Specific to PIM PE-PE signaling.....	<a href="#">35</a>
<a href="#">10.1</a>	PIM Neighborships Scale.....	<a href="#">35</a>
<a href="#">10.2</a>	PIM C-instances J/P Suppression Effectiveness.....	<a href="#">37</a>
<a href="#">11</a>	Test Cases Specific to PIM MI-PMSI trees.....	<a href="#">39</a>

<a href="#">11.1</a>	Default MDT's (MI-PMSI's) PIM P-Instance Mroutes Scale..	<a href="#">39</a>
<a href="#">12</a>	Security Considerations.....	<a href="#">41</a>
<a href="#">13</a>	IANA Considerations.....	<a href="#">41</a>
<a href="#">14</a>	Acknowledgments.....	<a href="#">41</a>
<a href="#">15</a>	References.....	<a href="#">41</a>
<a href="#">15.1</a>	Normative References.....	<a href="#">41</a>
<a href="#">15.2</a>	Informative References.....	<a href="#">42</a>
	Author's Addresses.....	<a href="#">43</a>
	Intellectual Property Statement.....	<a href="#">43</a>
	Disclaimer of Validity.....	<a href="#">44</a>
	Copyright Statement.....	<a href="#">44</a>
	Acknowledgment.....	<a href="#">44</a>

## [1](#) Introduction

Multicast Virtual Private Network (MVPN) is a service offered by BGP/MPLS VPN service providers, that provides a way for IP multicast traffic to travel from one customer site to another. [[L3VPN-MCAST](#)] is the framework describing how various protocols fit together to enable such functionality.

With the increased popularity, the scalability of deploying MVPN is becoming of a great interest. There is, however, no standard method defined to measure and compare different implementations. This document proposes a MVPN scalability metric and methodology for testing and comparing control plane MVPN scalability of (Provider Edge) PE devices in a standardized way.

Before describing the detailed test methodology, it is important to review the key factors that impact the scalability of MVPN deployments:

- o The MVPN Metric is 9-tuple comprised of a set of variables that indicate the overall scalability capabilities of a PE device implementation. MVPN scalability is multi-dimensional and can not be quantified with single parameter, thus defining such a metric set is necessary. MVPN Metric will be defined in [section 5](#). The remaining of this document focuses on a methodology that

characterizes different dimensions of MVPN Metric.

- o Choice of MVPN architecture or MVPN design and operational choices within specific architecture (such as selection of PIM protocol variant or extent of S-PMSI (data MDT) usage), impact the overall MVPN scalability of a PE device. Typically there is a tradeoff between optimality and scalability. More details on these choices with their tradeoffs are discussed in [MVPN-DEPLOY] and [L3VPN-

MCAST]. In this document design choices most suitable for a goal of any given test case will be used which may not necessarily be the same as recommended design choice for a realistic deployment.

- o MVPN is a service that is never deployed in isolation as it requires underlying unicast VPN offering. Typically SPs add MVPN service on PE devices that are already deployed and are providing a large number of other services such as unicast L3VPNs, L2VPNs, internet access, etc. Therefore, when considering MVPN scalability in realistic deployments one needs to take into consideration the level to which PE resources are already utilized and the available headroom amount remaining. In this document it will be assumed that MVPN service is deployed as an addition of a "minimized" unicast control plane.
- o MVPN Scalability of a PE device is different when the system is subjected to different stimuli. For example overall scalability achieved in steady state is typically higher than when the system is subjected to a network and/or device specific failures. In this document a limited set of mandatory test stimuli will be defined.

## [2](#) Document Scope

In IETF currently there are multiple proposals on architectures and protocols for implementing MVPN service, as documented in [L3VPN-MCAST]. The scope of this document is to provide MVPN scalability metric and benchmarking methodology set common to all architectures from [[L3VPN-MCAST](#)]. However, some architectures may require additional architecture specific test cases and considerations. This document provides such additional test cases for the stable and widely deployed MVPN architecture described in [[L3VPN-MCAST](#)] which uses PIM protocol for both PE-PE transmission of C-Multicast routing information (test cases in [section 10](#)) and to create 'tunnels' that instantiate Multidirectional Inclusive P-Multicast Service Interfaces (MI-PMSIs) and Selective P-Multicast Service Interfaces (S-PMSIs) (test cases in [section 11](#)). The same architecture is also described in [[ROSEN-8](#)] which is obsoleted by [[L3VPN-MCAST](#)]. In the rest of the

document this architecture will be referred to as "ROSEN-8" architecture.

As other architectures from [[L3VPN-MCAST](#)] become stable and widely deployed, amendments to this document can be made to address any specific scalability considerations of such architectures.

Scope of this document is to address comparison between different implementations of the same MVPN architecture, and not between different MVPN architectures defined in [[L3VPN-MCAST](#)].

This document proposes a MVPN metric and a test methodology to compare the MVPN control plane scalability of PE devices in a standardized way. In contrast, forwarding performance benchmarking is not within the scope of this document.

Test methodology defines a standard set of test cases, their test execution procedures, results content and reporting format. Standard test environment is also defined for each test case.

Test cases 9.1-9.7,10.1,11.1 focuses on determining implementation limits individually for each of the key MVPN variables in a standard way. Test cases (9.8-9.13) focus on determining implementation limits for combination of all MVPN variables and will be helpful to operators with network engineering for their deployments. Choices of values of variables in test cases 9.8-9.13 were made using information from the MVPN requirements survey conducted as part of [[MVPN-REQ](#)].

Each test case addresses following two major testing types:

- . Steady State Testing: Device Under Test (DUT) and network as a whole are not subject to any failure stimulus/control plane events.
- . Failure Recovery Testing: DUT and or network components are subject to different failure stimulus that introduces one or more control plane instability events.

In this document limited set of mandatory test stimuli is also defined.

Note that, depending on MVPN architecture, the deployment of MVPN also consumes resources on network elements other than PE router such as P devices, route reflectors and CE devices. Their scalability is beyond the scope of this document.

### [3](#) Terminology

DUT (Device Under Test) term will be used interchangeably with MVPN PE device. All other PE devices in the test topology will be referred to as "remote PEs".

We will use term "MVPN architecture" to describe any specific subset of protocols and procedures from [[L3VPN-MCAST](#)] that can enable MVPN functionality on PE device. In contrast, we will use term "MVPN implementations" to describe practical implementations of such "MVPN architectures".

VPN related terms used in this document are defined in [RFC4364](#) and RFC2547bis. MVPN related terms used in this document are defined in [[L3VPN-MCAST](#)].

PIM (Protocol Independent Multicast) related terms are defined in [RFC4601](#).

For the reader's convenience, here is review of some key terms used in this document:

MVPN (Multicast Virtual Private Network): VPN that supports transport of IP multicast traffic from one site to another.

PMSI (P-Multicast Service Interface): A PMSI is a conceptual "overlay" on the P network with the following property: a PE in a given MVPN can give a packet to the PMSI, and the packet will be delivered to some or all of the other PEs in the MVPN, such that any PE receiving such a packet will be able to tell which MVPN the packet belongs to.

MI-PMSI (Multidirectional Inclusive PMSI): PMSI which enables ANY PE attaching to a particular MVPN to transmit a message such that it will be received by EVERY other PE attaching to that MVPN.

S-PMSI (Selective PMSI): PMSI which enables PE attaching to a MVPN to transmit a message such that it will be received by subset of other PEs attaching to that same MVPN.

Default MDT (Default Multicast Distribution Tree): Multicast distribution tree through the SP core that connects ALL PEs which belong to given MVPN. This is [[ROSEN-8](#)] terminology for transport service of MI-PMSIs.

Data MDT (Data Multicast Distribution Tree): Multicast distribution tree through the SP core that delivers VPN data traffic for a particular multicast group only to those PE routers which are on the path to receivers of that group. This is [[ROSEN-8](#)] terminology for transport service of S-PMSIs.

ASM (Any Source Multicast): Multicast service model in which a receiver subscribes to a multicast group to receive traffic sent to the group by any source.

SSM (Source Specific Multicast): Multicast service model in which a receiver subscribes to a multicast group to receive traffic sent to the group by the specific source.

Mroute: Multicast route. Term "state" will used interchangeable with "mroute" and "multicast route".

"ROSEN-8" architecture: architecture described in [[L3VPN-MCAST](#)] which uses PIM protocol for both PE-PE Transmission of C-Multicast Routing and to create 'tunnels' that instantiate Multidirectional Inclusive P-Multicast Service Interfaces (MI-PMSIs) and Selective P-Multicast Service Interfaces (S-PMSIs). This is a same as architecture described in [[ROSEN-8](#)].

Ingress C-instance multicast group or mroute: C-instance multicast group or mroute which has source behind VPN C-interface of DUT PE and at least one receiver behind one of remote PEs.

Egress C-instance multicast group or mroute: C-instance multicast group or mroute which has source behind one of remote PEs and at least one receiver behind VPN C-interface of DUT PE.

Ingress traffic flow, packet, traffic: respectively multicast traffic flow, packet, traffic forwarded by DUT using ingress C-instance mroute.

Egress traffic flow, packet, traffic: respectively multicast traffic flow, packet, traffic forwarded by DUT using egress C-instance mroute.

#### [4](#) Key Words to Reflect Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#).

[RFC 2119](#) defines the use of these key words to help make the intent of standards track documents as clear as possible. While this document uses these keywords, this document is not a standards track

#### [5](#) MVPN Metric Definition

MVPN control plane scalability of PE device can not be described as a single parameter but it requires a set of variables. We call such a set "MVPN Metric" and define it further in this section.

When providing scalability capabilities of a PE device one MUST provide values for all of the MVPN metric variables that were used during the test. For example, one should never claim that a PE device supports X number of MVPNs without disclosing the values of other MVPN Metric variables.

The MVPN Metric is defined as a tuple of the following 9 variables:

1. Num\_mVPN: Number of multicast VPN routing instances configured on DUT that have MI-PMSI (default MDT) active and forwarding
2. Num\_PE: Number of PE routers per multicast VPN
3. Num\_MC\_C\_ints: Number of PIM C-interfaces on DUT
4. Num\_PIM\_C\_neigh: Total number of PIM neighbors in PIM C-instances



- across all mVPNs on DUT
5. Num\_\*G\_C: Total number of (\*,G) multicast routes across all MVPNs on DUT capable of forwarding and created by PIM C-instances.
  6. Num\_SG\_C: Total number of (S,G) multicast routes across all MVPNs on DUT capable of forwarding and created by PIM C-instances.
  7. Num\_OIF\_C: Total number of OIFs on DUT across all multicast routes created by PIM C-instances.
  8. Num\_SPMSI\_Src: Total number of S-PMSIs (data MDTs) across all mVPNs on DUT that are sourced by DUT.
  9. Num\_SPMSI\_Rx: Total number of S-PMSIs (data MDTs) across all mVPNs on DUT for which DUT is a receiver.

## 6 Test Environment

All protocols involved MUST be deployed with default timers as specified by their corresponding RFC / standards.

### 6.1 Test Topologies

Following test topology is used by all test cases that don't explicitly specify different topology.




Figure 1. Test Topology 1

Legend:

D1 (DUT's C-facing interface): DUT's interface that connects to customer premise router (C-router).

D2 (DUT's P-facing interface): DUT's interface that connects to SP's core router (P-router).

RR/P (Route Reflector/P-router) - single router that will be performing roles of both P-router and route reflector

PE2 - Will also be referred to as "Remote PE" and is the router performing PE functionality to assist with evaluation of DUT PE router.

## [6.2](#) Unicast Control Plane Setup

All P facing interfaces MUST use OSPF as IGP. This requirement is made to provide a standard way to compare end to end convergence times which depend on the underlying unicast protocol. Only a minimum number of IGP routes required to establish connectivity should be seen on the DUT.

All PE routers in the topology including the DUT and emulated PE's MUST have one iBGP peer to the Route Reflector. DUT SHOULD NOT have any additional iBGP peering. Only the minimum number of VPNv4 iBGP routes required to establish site to site VPN connectivity should be imported on the DUT. There SHOULD NOT be any internet/ipv4 routes seen on the DUT.

Dry, et al.

Expires May 9, 2008

[Page 9]

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

A DUT MUST use static unicast routing on all C facing VPN interfaces. Only the minimum number of static routes required to establish end to end connectivity should be seen on the DUT. No dynamic unicast routing protocol is used in order to minimize processing overhead.

## [6.3](#) Multicast Control Plane Setup

In any given test, all MI-PMSIs (default MDTs) MUST be established using the same protocol. In cases where PIM is used to establish MI-PMSIs, different tests may require different PIM protocol variants,

so refer to individual test cases for the appropriate PIM P-instance configuration. In any test case where ASM (Any Source Multicast) mode of PIM-SM (Protocol Independent Multicast - Sparse mode) is the multicast routing protocol used for MI-PMSI (default MDT), a DUT SHOULD NOT be the RP (Rendezvous Point). Also dynamic RP discovery protocols SHOULD NOT be used for default MDT groups.

For S-PMSIs (data MDT) point to multipoint trees MUST be used. For example if PIM is used to build S-PMSI groups PIM-SSM (Source Specific Multicast) routing protocol MUST be used.

Following C-instance multicast configuration MUST be used in all test cases:

- a. Protocol for PIM C-instances: PIM-SM (ASM)
- b. RP Location for PIM C-instances: Test apparatus port closest to the source.
- c. SPT threshold for PIM C-instances: zero

Following per MVPN scale of common variables MUST be used in all test cases except where test case setup explicitly asks to test with different values:

- a. Number of PIM VPN C-interfaces: 1
- b. Number of sources per C-instance group: 2
- c. Ratio of ingress to egress C-instance groups: 10%:90%

Behind each VPN C-interface there MUST be one receiver for each C-instance group whose source doesn't reside behind the same C-interface.

Sources emulated by test apparatus ports that are physically directly connected to DUT (port A1 in Figures 1 and 2) not have IP address from DUT's connected subnets, i.e. the DUT MUST NOT be the first hop router.

For consistency, it is recommended for test apparatus ports that are physically directly connected to DUT (port A1 in Figures 1 and 2) MUST not to use IGMP protocol to emulate multicast receivers. Instead PIM protocol should be used, i.e. the DUT should not be the last hop router.

As an exception to previous paragraph it may exist specific network

design requirement to deploy IGMP receivers connected directly to the DUT in which case test results MUST specify number of C-interfaces with IGMP receivers. Regardless the IGMP protocol variant to be deployed (IGMPV2 / V3); receivers MUST be emulated by the test apparatus and NOT defined on the DUT in the form of static group reports. Test apparatus MUST be capable to emulate an IGMP Host or Querier and set a maximum Timer Interval between messages of 1/10th of a second.

#### [6.4](#) Considerations for Number of Receivers behind remote PEs

Number of C-instance receivers behind remote PEs that MUST be emulated with test tools for all ingress multicast groups depends on protocol used for PE-PE signaling.

If PIM is used as PE-PE signaling and test case 10.2 revealed that PE device does perform J/P suppression all test cases MUST be executed with receiver behind only one of remote PEs for any given ingress C-instance group. This is required to properly emulate remote PEs that support PIM J/P suppression with existing test tools which do not support PIM J/P suppression.

If PIM is used as PE-PE signaling and test case 10.2 revealed that PE device does not perform efficient J/P suppression, then for each test case that is requesting in Test Setup section for Number of remote PEs > 1, testing MUST be executed with receivers behind 50% of remote PEs per ingress C-instance group. It is DESIRABLE to execute all such test cases with 2 additional values of number receivers behind remote PEs.

If BGP is used as PE-PE signaling, then for each test case that is requesting in Test Setup section for Number of remote PEs > 1, testing MUST be executed with receivers behind 50% of remote PEs per ingress C-instance group. It is DESIRABLE to execute all such test cases with 2 additional values of number receivers behind remote PEs.

#### [6.5](#) Data Traffic Characteristics

For every C-instance multicast route there MUST be traffic flow associated with it and forwarded by DUT.

All C-instance flows SHOULD be transmitted with the same traffic rate and packet size.

As the focus of this document is on the control plane scalability and not on forwarding performance the data rate and packet size of traffic flows can be chosen by user but it MUST be reported in the test results. However it is suggested to use 10% of "idle system" throughput [[RFC1242](#)] so that it can be easily detected if hardware forwarding platforms start forwarding in software and at the same time in case of software forwarding platforms there will be enough processor headroom left for control plane scaling. By "idle system" we refer to system with all of MVPN metric variables minimized and single VPN traffic flow in each direction.

As an additional requirement, the reader of this document may also be interested in analyzing the "impact" that high traffic rate may have on the control plane. This would be of interest mostly for software forwarding platforms. For this specific requirement additional test cases SHOULD be performed increasing the rate of multicast traffic to 20%, 50% and 90% of "idle system" throughput [[RFC1242](#)].

## [6.6](#) Test Apparatus Considerations

Different test tools must generate C-instance PIM protocol control messages in a consistent way since they are directly connected to the DUT.

The following MUST be implemented on all PIM sessions on the test apparatus:

- 1) PIM Join/Prune aggregation MUST be utilized and set such that 80 PIM J/P messages are aggregated in each PDU
- 2) PIM Join/Prune aggregated PDUs MUST be sent at 10 PDUs/sec rate per PIM session, i.e. this translates to maximum of  $80 \times 10 \times 60 = 48,000$  state per minute.
- 3) PIM Register messages MUST be sent at 100 PDUs/sec rate.

## [7](#) Test Categories, Stimulus and Execution Methodology

Each test case specified in [section 9](#) MUST be executed for steady state and for each of six mandatory failure stimulus listed below. Optionally one can use methodology defined in this document for additional stimulus.

Mandatory failure stimulus:

- 1) DUT Power Cycle: Physical power cycle of DUT. All convergence times MUST be measured from the time DUT's power is turned back on. This time instance will be referred to as Tf (the time of failure recovery action) for this failure stimulus.
- 2) Main Processor Card Switchover: Physical removal of the active main processor card in the redundant system. All convergence times should be measured from the time active processor card is physically disconnected from the chassis (Tf). This stimulus can be omitted only for platforms that do not support redundant main processor cards.
- 3) P-facing Line Card OIR (online insertion and removal): Physical removal and insertion of P-facing line card. Time between removal and insertion SHOULD be at least 300 seconds. All convergence times should be measured from the time line card is physically inserted into chassis (Tf).
- 4) C-facing Line Card OIR: Physical removal and insertion of C-facing line card. Time between removal and insertion SHOULD be at least 300 seconds. All convergence times should be measured from the time line card is physically re-inserted into chassis (Tf).
- 5) P-facing Link Flap: Physical removal and insertion of the cable from P router side that is connected to P-facing interface of DUT. Time between removal and insertion SHOULD be at least 300 seconds. All convergence times should be measured from the time cable is physically re-inserted (Tf).
- 6) C-facing Link Flap: Physical removal and insertion of the cable from test apparatus side that is connected to C-facing interface of DUT. Time between removal and insertion SHOULD be at least 300 seconds. All convergence times should be measured from the time cable is physically re-inserted (Tf).

Since the test execution methodology is similar for all test cases we will describe it here for both steady state and failure recovery testing. Any deviation from this will be specified per test case in sections [9](#), [10](#) and [11](#).

Multiple iterations of each test are required to determine maximum value for certain set of variables. A single iteration will be referred to as a "Test Case Instance".

### [7.1](#) Steady State Testing Execution Methodology

The following test execution procedure MUST be used for all Test Case Instances during steady state testing of each test case defined in sections [9](#), [10](#) and [11](#) of this document:

- 1) Ensure the testbed is setup according to Test Setup instructions of individual test case
- 2) All interfaces MUST be operational and MI-PMSIs (default MDTs) required by the test case MUST be built as expected. Verification can be done by DUT internal tools.
- 3) All real and emulated PE devices required by test cases MUST have all C-instance PIM neighborships (including over MI-PMSIs) operational in both directions. Verification MUST be done by both external test apparatus and DUT internal tools.
- 4) All destination test apparatus ports configured to receive multicast traffic should join all configured multicast groups.
- 5) All source test apparatus ports configured to transmit multicast traffic should start transmitting to all multicast groups.
- 6) All multicast traffic MUST be received at all expected destination test ports without any packet drops. This MUST be verified using external test apparatus. If this state can not be reached within 10 minutes of execution of step 5, continue to next test case instance with reduced value of scaled variable/s.
- 7) After state in previous step is reached wait 10 minutes and start collecting data for this test instance required by individual test case. This time instance is considered steady state.
- 8) If any one of following conditions are reached continue to next test case instance with reduced value of scaled variable/s:
  - o 100% utilization of system resources (memory, processor, etc.)

- o Failing of any of test case specific criteria or criteria in steps 1-6 above

The number of Test Case Instances per test case is left to tester's discretion. However, it is DESIRABLE to have results for

Dry, et al.

Expires May 9, 2008

[Page 14]

---

Internet-Draft Multicast VPN Scalability Benchmarking November 2007

at least 5 test case instances. Having a range of values will help in variable's characterization. The characterization of a variable cannot be achieved with only one test case instance result.

## 7.2 Failure Recovery Testing Execution Methodology

The following test execution procedure MUST be used for all Test Case Instances during failure recovery testing of each test case defined in [section 9](#) of this document:

- 1) Execute steps 1-7 from [section 7.1](#)
- 2) After steady state in previous step is reached wait 10 minutes to initiate one of mandatory failure stimulus listed in [section 7](#). Note the time of failure recovery action (Tf) as displayed by the external test apparatus that is measuring the received multicast traffic.
- 3) Using the external test apparatus note the time when THE FIRST multicast packet has been received on at least ONE of expected ports. Refer to this time instance as Tre for the first such ingress packet and Trd for the first such egress packet.
- 4) Using the external test apparatus note the time when ALL multicast traffic has been received on ALL expected ports, i.e. it has returned to the same initial rate (in pps). Refer to this time instance Trall.
- 5) After state in previous step is reached execute steps 2-3 from 7.1.
- 6) If all data verified in step 5 is the same as before failure wait 10 minutes and start collecting data for this test instance required by each individual test case
- 7) If any one of following conditions are reached continue to next test case instance with reduced value of scaled variable/s:



- a. Value of MVPN metric in steady state reached after failure stimuli (step 6 above) is not the same as in original steady state.
- b. Failing of any of test case specific criteria or criteria in steps 1-6 above

The number of Test Case Instances per test case is left to tester's discretion. However, it is DESIRABLE to have results for

at least 5 test case instances. Having a range of values will help in variable's characterization. The characterization of a variable cannot be achieved with only one test case instance result.

## [8](#) Results Content and Reporting Format

### [8.1](#) Steady State Testing

For steady state portion of testing for each test case the following results MUST be included in the test case report:

1. Protocol used for establishment of MI-PMSIs and S-PMSIs
2. Protocol used for PE-PE signaling
3. Maximum value achieved for variables requested to be varied in individual test case
4. Values of MVPN Metric variables in the test instance in which item 1 of this report was achieved. The MVPN Metric as defined in [section 5](#) of this document MUST be used
5. Num\_\*G\_P: Total number of (\*,G) multicast routes on DUT created by PIM P-instance on DUT.
6. Num\_SG\_P: Total number of (S,G) multicast routes on DUT created by PIM P-instance.
7. Num\_OIF\_P: Total number of OIFs (outgoing interfaces) on DUT across all multicast routes created by PIM P-instance.
8. Forwarding rate(in pps)[\[RFC2285\]](#) and packet sizes (in bytes) of all flows in ingress direction at DUT
9. Forwarding rate(in pps)[\[RFC2285\]](#) and packet sizes (in bytes) of all flows in egress direction at DUT
10. Multicast Latency [\[RFC2432\]](#) averaged over all C-instance multicast flows in ingress direction
11. Multicast Latency [\[RFC2432\]](#) averaged over all C-instance multicast flows in egress direction

12. Utilization of all processors in the system including the main processor and line card processors where applicable. A description of the way processor utilization is measured SHOULD be included in the report.
13. Utilization of all relevant DUT memory components including the main route processor memory and line cards where applicable.
14. Utilization of any relevant hardware components where applicable
15. Any deviations in DUT configuration from the configuration defined in this document.

16. Any deviations in test execution procedure

It is DESIRABLE to include in the report items 1-16 above for all optional test case instances executed, where instead of maximum value achieved one would report tested value for each test case instance.

## [8.2](#) Failure Recovery Testing

In addition to data included in steady state reports defined in the previous section the following MUST be included in the result report of each failure recovery test case:

1. The worst case end to end traffic convergence time ( $T_{\text{rall-Tf}}$ )
2. The best case end to end traffic convergence time ( $T_{\text{re-Tf}}$ ) for ingress traffic and ( $T_{\text{rd-Tf}}$ ) for egress traffic)

Note: Determination of whether all multicast flows had recovered to the original traffic rate MUST be made by external test tools and not by any available tools internal to the DUT or other routers in the test topology.

It is DESIRABLE to include:

1. A graph from all test tool ports showing transmitted and received packet rate starting from 60 seconds prior to failure action to 60 seconds after all multicast flows had recovered to the traffic rate they had prior to the failure.
2. The best case MI-PMSI PIM neighborhood convergence time: time interval from instance  $T_f$  to instance when the first C-instance PIM

neighbor across one of MI-PMSIs comes up on both DUT and neighboring device (i.e. "bi-directional" neighborships are established).

3. The worst case MI-PMSI PIM neighborship convergence time: time interval from instance T<sub>f</sub> to instance when all expected C-instance PIM neighbors across one of MI-PMSIs comes up on both DUT and neighboring device (i.e. "bi-directional" neighborships are established).

## [9](#) Test Cases Common to all MVPN Architectures

Test cases in this section are applicable to all MVPN architectures described in [[L3VPN-MCAST](#)]. As [[L3VPN-MCAST](#)] specifies use of S-PMSIs as optional, test cases 9.5-9.7,9.11 can be omitted for implementers that don't support S-PMSIs. For such implementers test cases 9.8-9.10,9.12-9.13 SHOULD still be executed but without use of S-PMSIs and the exception MUST be documented in the test report.

### [9.1](#) "Empty" MVPNs Scale

#### Test Objective:

To determine maximum number of MVPN instances that can be configured and operational on the MVPN PE router. Note that we refer here to mVPNs as "empty" as amount of PIM neighborships, interfaces, C-instances multicast routes and SI-PMSIs associated with given mVPN is negligible or zero in this test case.

#### Metric Variables Relationships:

Num\_mVPN=Num\_MC\_C\_ints

Num\_\*G\_C=Num\_SG\_C=2\*Num\_mVPN

Num\_OIF\_C=4\*Num\_mVPN

#### Test Setup:

Following test setup MUST be performed prior to executing this test

case

1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT) used: NO
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs configured on DUT(Num\_mVPN): varies
  - b. Number of remote PEs: 1
  - c. Number of C-instance multicast groups : 2
  - d. Ratio of ingress vs. egress C-instance groups: 50%:50%

Test Execution Procedure:

Dry, et al.

Expires May 9, 2008

[Page 18]

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

Execute number of test case instances where in each test case instance number of configured mVPNs is varied with the goal of finding maximum number of mVPNs that can be configured and operational on DUT. Configured mVPN will be considered operational if all traffic flows are being received on ALL expected ports without any drops.

For each test case instance perform steps 1-8 from [section 7.1](#). and 1-7 from [section 7.2](#) for all mandatory stimuli in [section 7](#).

Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of number of mVPNs achieved. It is DESIRABLE to include the same data for at least 5 different values of number of mVPNs (i.e. for at least 5 test case instances).

## [9.2](#) PIM Enabled VPN C-Interfaces Scale

Test Objective:

To determine maximum number of PIM enabled VPN C-interfaces that can be operational on the MVPN PE router for couple of fixed values of number of mVPNs. Amount of all other MVPN Metric is minimized in this test case.

Metric Variables Relationships:

Num\_MC\_C\_ints >= Num\_mVPN

Num\_\*G\_C=Num\_SG\_C=Num\_OIF\_C=0

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT) used: NO
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs configured on DUT (Num\_mVPN): varies
  - b. Number of PIM VPN C-interfaces: varies
  - c. Number of remote PEs: 1

- d. Number of C-instance multicast groups : 0

#### Test Execution Procedure:

Following are steps to execute this test case:

1. Configure 100 mVPNs on DUT and PE2. Execute number of test case instances where in each test case instance number of PIM enabled VPN C-interfaces per mVPN is varied with the goal of finding maximum number of PIM enabled VPN C-interfaces that can be configured and operational on DUT. Configured VPN C-interface will be considered operational if there is at least one PIM neighbor in VPN C-instance on each configured C-interface.
2. Repeat step 1 for 100\*I mVPNs where "i=2<sup>00</sup>N" where N is integer value for which either maximum number of PIM enabled VPN C-interfaces per mVPN becomes smaller than one or maximum number of mVPNs found in test case 9.1 is reached.

Note that in this test case there SHOULD NOT be any multicast C-instance traffic sources or receivers thus one MUST modify test execution procedure from 7.1 and 7.2. For each test case instance perform steps 1-3,7 from [section 7.1](#). and 1-2,5-7 from [section 7.2](#)

for all mandatory stimuli in [section 7](#).

#### Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum achieved value of number of PIM enabled VPN C-interfaces. It is DESIRED to include the same data for at least 5 different values of PIM enabled VPN C-interfaces (i.e. for at least 5 test case instances).

### [9.3](#) PIM C-instances Mroutes Scale

#### Test Objective:

To determine the maximum amount of PIM C-instance mroutes that a PE router can create, maintain and forward on. Amount of most of other MVPN Metric is minimized in this test case.

#### Metric Variables Relationships:

$$(\text{Num\_G\_C} + \text{Num\_SG\_C}) \gg \text{Num\_mVPN}$$
$$\text{Num\_OIF\_C} \gg \text{Num\_mVPN}$$
$$\text{Num\_MC\_C\_ints} = \text{Num\_mVPN}$$

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups(if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT) used: NO
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs (Num\_mVPN) configured on DUT: varies
  - b. Number of remote PEs: 1
  - c. Ratio of ingress vs. egress C-instance groups: 100%:0%, 0%:100% and 10%:90%

- d. Number of C-instance sources per group: 2 and 50

#### Test Execution Procedure:

Total number of C-instance PIM mroutes is proportional to product of number of mVPNs DUT belongs to and average number of C-instance PIM mroutes per mVPN. Total number of C-instance mroutes can vary depending on number of mVPNs, number of sources per group and location of the sources/receivers. Thus this test should be executed for all specified values of parameters in Test Setup.

Each test will consist of finding maximum number of C-instance PIM mroutes by varying average number of C-instance PIM mroutes per mVPN for set of fixed values of number of mVPNs. Procedure is as follows:

1. On DUT and PE2 configure 100 mVPNs. Setup environment such that all PIM C-instance sources are behind DUT. Execute number of test case instances using steps 1-7 in [section 7.1](#) where in each test case instance number of C-instance PIM groups is varied until maximum number of C-instance PIM mroutes is found.
2. Repeat step 1 for  $100 \times I$  mVPNs where " $i=2^{00N}$ " where N is integer value for which either maximum number of C-instance PIM mroutes per

mVPN becomes smaller than one or maximum number of mVPNs found in test case 9.1 is reached.

3. Repeat steps 1 and 2 for two more values of ingress vs. egress C-instance groups ratio: 0%:100% and 10%:90%.
4. Repeat steps 1 and 2 for one more value of number of sources per group

For each test case instance perform steps 1-8 from [section 7.1](#). and 1-7 from [section 7.2](#) for all mandatory stimuli in [section 7](#).

#### Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of average number of PIM C-instance mroutes for every tested value of number of mVPNs per PE. It is DESIRED to

include the same data for at least 5 different values of number of PIM C-instance mroutes per mVPN for each of tested values of number of mVPNs per PE(i.e. for at least 5 test case instances per each tested value of number of mVPNs).

#### 9.4 PIM C-Instances OIF Scale

##### Test Objective:

To determine the maximum amount of PIM C-instance OIFs that a PE router can create and maintain. Amount of other MVPN Metric is minimized in this test case.

##### Metric Variables Relationships:

$\text{Num\_MC\_C\_ints} > \text{Num\_mVPN}$

$(\text{Num\_*G\_C} + \text{Num\_SG\_C}) \gg \text{Num\_mVPN}$

$\text{Num\_OIF\_C} \gg \text{Num\_mVPN}$

##### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir

2. S-PMSI (Data MDT)used: NO
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs(Num\_mVPN)configured on DUT: 100 and maximum value tested in 9.3
  - b. Number of PIM VPN C-interfaces: maximum found in 9.2
  - c. Number of remote PEs: 1
  - d. Ratio of ingress vs. egress C-instance groups: 0%:100%

##### Test Execution Procedure:

Test will consist of finding the maximum number of C-instance PIM OIFs by varying the average number OIFs per PIM C-instance mroute.



Maximum number will be found for couple of values of number of C-instance PIM mroutes. Test will be executed for two values of number of mVPNs: 100 and maximum value tested in 9.3. All C-instance PIM mroutes will be in egress direction. Procedure is as follows:

1. For the first test iteration number of C-instance groups should be set to 25% of maximum value achieved in test case instance of 9.3 where C-instance group ingress vs. egress ratio was 0:100% and 100 mVPNs was used. Execute number of test case instances using steps 1-8 in [section 7.1](#) where in each test case instance average number of C-instance OIFs per mroute is varied in increments of 2 until maximum number of OIFs is reached.
2. Repeat step 1 for 50%, 75% and 100% of egress C-instance groups achieved in test case 9.3.
3. Repeat steps 1 and 2 for the case of maximum number of mVPNs tested in 9.3.

For each test case instance perform steps 1-8 from [section 7.1](#). and 1-7 from [section 7.2](#) for all mandatory stimuli in [section 7](#).

#### Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of OIFs per C-instance mroute for every tested value of number of egress groups per PE. It is DESIRED to include the same data for at least 5 different values of number of OIFs for each of tested values of number of egress groups (i.e. for at least 5 test case instances per each tested value of number of egress groups).

#### [9.5](#) Joined S-PMSI (Data MDT) Scale

##### Test Objective:

To determine the maximum number of S-PMSIs (data MDTs) that a PE can join. In order to assess maximum number of S-PMSI (data MDTs) joined, we minimize resources taken by C-instance mroutes by requiring that no data MDT (S-PMSI) reuse is utilized in this test case. Note that depending on specific deployment context data MDT

reuse might or might not be desirable.

#### Metric Variables Relationships:

Num\_SPMSI\_Rx > Num\_mVPN

Num\_SPMSI\_Src = 0

Num\_MC\_C\_ints > Num\_mVPN

(Num\_\*G\_C + Num\_SG\_C)>> Num\_mVPN

Num\_OIF\_C >> Num\_mVPN

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT)used: YES
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs (Num\_mVPN) configured on DUT: maximum number of MVPNs obtained in test case 9.3 (refer to it as Vmax)
  - b. Number of PIM VPN C-interfaces: max found in 9.2 for Vmax mVPNs
  - c. Number of remote PEs: 1
  - d. Number of C-instance multicast groups: :[Smax/4] where Smax is maximum number of C-instance groups obtained in test case 9.3 for Vmax number of mVPNs
  - e. Ratio of ingress vs. egress C-instance groups: 0%:100%

#### Test Execution Procedure:

Test will consist of varying number of S-PMSIs utilized by flows that have receivers behind DUT (refer to those S-PMSI as "joined S-PMSIs"). During all test case instances total number of C-instance PIM mroutes MUST remain constant and will be [Smax/4] rounded to

the first lower integer. We will vary total number of joined S-PMSIs by varying number of mVPNs configured to use S-PMSIs at the remote PE that has sources behind it, while number of data S-PMSIs per mVPNs will be same for all mVPNs that use them. If given mVPN is using S-PMSIs in particular test case instance number of them should be  $D_{vpn} = \lceil S_{max} / (4 * V_{max}) \rceil$  rounded to first lower value that can be represented as  $2^I$  where I is an integer. Note that number of S-PMSIs configured and sourced by DUT MUST be zero in this test case. Procedure is as follows:

1. Configure one mVPN on the remote PE and all flows so that this mVPN has  $D_{vpn}$  S-PMSIs (data MDTs) utilized and all are sourced at the remote PE and received at DUT. Execute steps 1-7 in [section 7.1](#)
2. Repeat step 1 where number of mVPNs that utilize S-PMSIs/data MDTs (in the exact same way as 1 mVPN described in step1) takes following values 25%,50%,75%,and 100% of  $V_{max}$ .
3. If platform limit is not reached during execution of step 2, increase number of joined S-PMSIs (data MDTs) per mVPN and repeat steps 1 and 2.

For each test case instance perform steps 1-8 from [section 7.1](#). and 1-7 from [section 7.2](#) for all mandatory stimuli in [section 7](#).

Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for all test case instances executed.

## [9.6](#) Sourced S-PMSI (Data MDT) Scale

Test Objective:

To determine maximum number of S-PMSIs (data MDTs) that PE can source. In order to assess maximum number of S-PMSIs sourced, we minimize resources taken by C-instance mroutes by requiring that no S-PMSI (data MDT) reuse is utilized in this test case. Note that depending on specific deployment context data MDT reuse might or might not be desirable.

Metric Variables Relationships:

Num\_SPMSI\_Src > Num\_mVPN

Num\_SPMSI\_Rx = 0

Num\_MC\_C\_ints = Num\_mVPN

(Num\_\*G\_C + Num\_SG\_C)>> Num\_mVPN

Num\_OIF\_C >> Num\_mVPN

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT)used: YES
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs (Num\_mVPN) configured on DUT: maximum number of mVPNs obtained in test case 9.3 (refer to it as Vmax)
  - b. Number of PIM VPN C-interfaces: max found in 9.2 for Vmax mVPNs
  - c. Number of remote PEs: 1
  - d. Number of C-instance multicast groups: :[Smax/4] where Smax is maximum number of C-instance groups obtained in test case 9.5 for Vmax number of mVPNs
  - e. Ratio of ingress vs. egress C-instance groups: 100%:0%

#### Test Execution Procedure:

Reverse role of DUT and remote PE from test case 9.5, where now DUT is sourcing all data MDTs (S-PMSIs) while remote PE is on the receiving end of them. Repeat test case 9.5 for this reversed role scenario.

#### [9.7](#) S-PMSI (Data MDT) Reuse

##### Test Objective:

To determine maximum number of C-instance flows that can utilize data MDTs (S-PMSIs) and assess impact data MDT reuse has. Note that depending on specific deployment context data MDT reuse might or might not be desirable. This is optional test case for implementations that do support S-PMSI reuse.

#### Metric Variables Relationships:

$$(\text{Num\_*G\_C} + \text{Num\_SG\_C}) > \text{Num\_SPMSI\_Rx} \geq \text{Num\_mVPN}$$
$$(\text{Num\_*G\_C} + \text{Num\_SG\_C}) > \text{Num\_SPMSI\_Src} \geq \text{Num\_mVPN}$$
$$\text{Num\_MC\_C\_ints} = \text{Num\_PIM\_C\_neigh} > \text{Num\_mVPN}$$
$$(\text{Num\_*G\_C} + \text{Num\_SG\_C}) \gg \text{Num\_mVPN}$$
$$\text{Num\_OIF\_C} \gg \text{Num\_mVPN}$$

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT)used: YES
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs configured on DUT (Num\_mVPN): maximum number of mVPNs obtained in test case 9.5 (refer to it as Vmax)
  - b. Number of PIM VPN C-interfaces: max found in 9.2 for Vmax mVPNs
  - c. Number of remote PEs: 1
  - d. Number of C-instance multicast groups:  $\lceil S_{\text{max}}/2 \rceil$  where Smax is maximum number of C-instance groups obtained in test case 9.3 for Vmax number of mVPNs
  - e. Ratio of ingress vs. egress C-instance groups: 10%:90%
  - f. Number of sourced data MDTs (S-PMSIs): 2
  - g. Number of received data MDTs (S-PMSIs): 8

#### Test Execution Procedure:

Test will consist of varying number of C-instance flows that will utilize S-PMSIs, while keeping number of C-instance mroutes and data MDTs (S-PMSIs) constant. By doing this one can assess impact of data MDT reuse.

Procedure is as follows:

1. Configure test apparatus such that number of flows using S-PMSIs is the same as number of S-PMSIs, i.e. there is no S-PMSI reuse by multiple traffic flows. Execute steps 1-7 in [section 7.1](#) and 1-8 in [section 7.2](#)
2. Repeat step 1 for  $10 \times I$  flows where " $i=2^N$ " where N is integer value for which either maximum number of flows mapped to data MDT (S-PMSI) is reached or number of flows becomes equal to number of (S,G) C-instance mroutes.

For each test case instance perform steps 1-8 from [section 7.1](#). and 1-7 from [section 7.2](#) for all mandatory stimuli in [section 7](#).

Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for all test case instances executed.

## [9.8](#) Scale of mVPNs spanning large number of PEs

Test Objective:

As we noted mVPN scale is multidimensional and depends on number of variables. While test cases 9.1-9.7 focused on only one or two variables at the time while minimizing impact of all others, they don't give good representation of platform capabilities in more realistic deployment scenarios where none of variables are minimized. Objective of this test case (and test cases 9.8-9.13) is to assess capabilities of platform in more realistic deployment scenario. In particular this test case will focus on finding maximum number of mVPN instances that span large number of PE routers while they have values for other MVPN variables chosen to be on the order of magnitude used by MVPN deployments at the time this draft was written. Specific values are defined in Test Setup section.

Metric Variables Relationships:

$$\text{Num\_MC\_C\_ints} = 2 \times \text{Num\_mVPN}$$

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

$$(\text{Num\_G\_C} + \text{Num\_SG\_C}) = 30 * \text{Num\_mVPN}$$
$$\text{Num\_OIF\_C} = \text{Num\_mVPN} = 60 * \text{Num\_mVPN}$$
$$\text{Num\_SPMSI\_Rx} = 8 * \text{Num\_mVPN}$$
$$\text{Num\_SPMSI\_Src} = 2 * \text{Num\_mVPN}$$

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. S-PMSI (Data MDT)used: YES
2. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs configured on DUT: varies
  - b. Number of PIM VPN C-interfaces: 2
  - c. Number of remote PEs: 500
  - d. Number of C-instance multicast groups:10
  - e. Number of data MDTs (S-PMSIs) sourced from DUT:2
  - f. Number of data MDTs (S-PMSIs) with receivers behind DUT:8

#### Test Execution Procedure:

This test case SHOULD be repeated for all mechanisms of instantiating MI-PMSIs (default MDTs) supported by implementer. If PIM protocol is used all PIM variants(PIM ASM, SSM and Bi-dir) supported by implementer should be tested. At minimum PIM ASM MUST be tested.

Refer to 6.4 for guidelines on how many PE routers should be sending J/P messages for any given ingress C-instance mroute.

Execute number of test case instances where in each test case instance number of configured mVPNs is varied with the goal of finding maximum number of mVPNs that platform can support. mVPN instance here includes C-instance state, OIFs, PIM neighborships and data MDTs (S-PMSIs) as specified by Test Setup.

For each test case instance perform steps 1-8 from [section 7.1](#). and

1-7 from [section 7.2](#) for all mandatory stimuli in [section 7](#).

## Test Result Report:

Dry, et al.

Expires May 9, 2008

[Page 29]

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of number of mVPNs achieved. It is DESIRED to include the same data for at least 5 different values of number of mVPNs (i.e. for at least 5 test case instances).

### [9.9](#) Scale of mVPNs with larger amount of state

#### Test Objective:

Objective of this test case is on finding maximum number of mVPN instances that contain large number of C-instance PIM state while they have values for other MVPN variables chosen to be on the order of magnitude used by MVPN deployments at the time this draft was written. Specific values are defined in Test Setup section.

#### Metric Variables Relationships:

$$\text{Num\_MC\_C\_ints} = 2 * \text{Num\_mVPN}$$
$$(\text{Num\_*G\_C} + \text{Num\_SG\_C}) = 300 * \text{Num\_mVPN}$$
$$\text{Num\_OIF\_C} = \text{Num\_mVPN} = 600 * \text{Num\_mVPN}$$
$$\text{Num\_SPMSI\_Rx} = 8 * \text{Num\_mVPN}$$
$$\text{Num\_SPMSI\_Src} = 2 * \text{Num\_mVPN}$$

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. S-PMSI (Data MDT)used: YES
- 3.Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs configured on DUT (Num\_mVPN) : varies
  - b. Number of PIM VPN C-interfaces: 2
  - c. Number of remote PEs: 50
  - d. Number of C-instance multicast groups:100



- e. Number of data MDTs (S-PMSIs) sourced from DUT:2
- f. Number of data MDTs (S-PMSIs) with receivers behind DUT:8

#### Test Execution Procedure:

Same as Test Execution Procedure in 9.8

Dry, et al.

Expires May 9, 2008

[Page 30]

---

Internet-Draft Multicast VPN Scalability Benchmarking November 2007

#### Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of number of mVPNs achieved. It is DESIRED to include the same data for at least 5 different values of number of mVPNs (i.e. for at least 5 test case instances).

#### [9.10](#) Scale of "average" size mVPNs

##### Test Objective:

While test cases 9.12 and 9.13 assess two more extreme deployment scenarios with respect to number of PE routers and mVPN routes, objective of this test case is to assess number of mVPNs for the case where each mVPN represents average size mVPN customer. Specific values are defined in Test Setup section.

##### Metric Variables Relationships:

$$\text{Num\_MC\_C\_ints} = 2 * \text{Num\_mVPN}$$

$$(\text{Num\_*G\_C} + \text{Num\_SG\_C}) = 100 * \text{Num\_mVPN}$$

$$\text{Num\_OIF\_C} = \text{Num\_mVPN} = 200 * \text{Num\_mVPN}$$

$$\text{Num\_SPMSI\_Rx} = 8 * \text{Num\_mVPN}$$

$$\text{Num\_SPMSI\_Src} = 2 * \text{Num\_mVPN}$$

##### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. S-PMSI (Data MDT)used: YES
3. Multicast Control Plane Profile (all per mVPN except a.; all

from DUT's perspective):

- a. Number of MVPNs configured on DUT (Num\_mVPN): varies
- b. Number of PIM VPN C-interfaces: 2
- c. Number of remote PEs: 100
- d. Number of C-instance multicast groups: 33
- e. Number of data MDTs (S-PMSIs) sourced from DUT: 2
- f. Number of data MDTs (S-PMSIs) with receivers behind DUT: 8

Test Execution Procedure:

Dry, et al.

Expires May 9, 2008

[Page 31]

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

Same as Test Execution Procedure in 9.8

Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of number of mVPNs achieved. It is DESIRED to include the same data for at least 5 different values of number of mVPNs (i.e. for at least 5 test case instances).

### [9.11](#) S-PMSI Switching Delay

Test Objective:

The test objective is to measure the elapsed time for traffic to start flowing on S-PMSI, i.e., the time from the moment of signaling of an S-PMSI to the moment when traffic starts flowing on the S-PMSI. We will refer to this measure as S-PMSI switching delay [S-PMSI\_DELAY].

Metric Variables Relationships:

Same as for test case 9.10.

Test Setup:

Same as for test case 9.10.

Test Execution Procedure:

Test will measure the time for traffic to start flowing on S-PMSI, i.e., the time from the moment of signaling of an S-PMSI to the moment when traffic starts flowing on the S-PMSI on the DUT. This test MUST be repeated multiple (at least 20) times (for each value

of number of mVPNs) across multi-second intervals in order to isolate any timing issues. This test MUST be performed with varying number of average size MVPNs on DUT (up to the maximum) as defined in the test case 9.10. With a given number of MVPNs on DUT, the switching delay of several S-PMSIs sourced at the DUT in different MVPNs will be measured. In case S-PSMI creation is triggered by rate of C-instance traffic flow, the S-PMSI threshold should be set to min possible value, depending on the implementation. Such threshold value used MUST be documented in the test report.

#### Test Result Report:

The test results MUST include the range of S-PMSI switching delays: minimum, average and maximum [S-PMSI\_DELAY]. Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of number of mVPNs achieved. It is DESIRED to include the same data for at least 5 different values of number of mVPNs (i.e. for at least 5 test case instances).

#### [9.12](#) Convergence of C-Instance PIM Joins

##### Test Objective:

The test objective is to measure how long it takes for the first receiver on the DUT, issuing C-instance PIM (\*,G) Join, to receive the traffic from already active C-instance source (S,G).

##### Metric Variables Relationships:

Same as for test case 9.14.

##### Test Setup:

Same as in test case 9.14

##### Test Execution Procedure:

Test will consist of an active C-instance source (S,G) attached to/behind a remote PE (PE2 in Topology #1). There will be at least

one active C-instance receiver of (\*,G) behind this remote PE (PE2). This setup ensures that C-instance PIM Register procedure for (S,G) has been completed and that there are no receivers across the core network. With this initial setup, the test will add one C-instance (\*,G) receiver attached to DUT and issuing PIM (\*,G) Join towards the DUT. The test will measure the time for traffic from (S,G) behind the remote PE to be received by (\*,G) receiver behind the DUT. This test MUST be repeated multiple (at least 10) times in order to isolate any timing issues. This test MUST be performed with varying number of average size MVPNs (up to the maximum) as defined in the test case 9.14.

#### Test Result Report:

The test results MUST include the range of C-instance PIM (\*,G) Join convergence times: minimum, average, maximum. Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of number of mVPNs achieved. It is DESIRED to include the same data for at least 5 different values of number of mVPNs (i.e. for at least 5 test case instances).

#### [9.13](#) Effect of Co-locating C-RPs on a PE

##### Test Objective:

The test objective is to assess scaling impact of SP hosting customer's RPs (C-RPs) on PE router. This is the optional deployment model as stated in [[MVPN-REQ](#)] and is deployed today by number of service providers. Only difference between test cases 9.10 and 9.13 is location of C-RP; thus by comparing test results of 9.10 and 9.13 one will be able to determine additional impact co-locating RP has on PE scale compared to deployment model of customers hosting their own RPs.

##### Metric Variables Relationships:

Same as in 9.10.

##### Test Setup:

Same as in 9.10, except for placing C-instance RPs on DUT.

##### Test Execution Procedure:

This test case SHOULD be repeated for all mechanisms of instantiating MI-PMSIs (default MDTs) supported by implementer. If PIM protocol is used all PIM variants(PIM ASM, SSM and Bi-dir) supported by implementer should be tested. At minimum PIM ASM MUST be tested.

Refer to 6.4 for guidelines on how many PE routers should be sending J/P messages for any given ingress C-instance mroute.

Execute number of test case instances where in each test case instance number of configured mVPNs is varied with the goal of finding maximum number of mVPNs that platform can support in this environment. mVPN instance here includes C-instance state, OIFs,

Dry, et al.

Expires May 9, 2008

[Page 34]

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

PIM neighborships and data MDTs (S-PMSIs) as specified by Test Setup.

For each test case instance perform following steps:

- a. Steps 1-4 from [section 7.1](#).
- b. Send PIM Register messages from all "source" test apparatus ports
- c. Test apparatus should verify that correct (S,G) PIM Join messages had been received by "source" test apparatus port. In reaction to receipt of (S,G) joins, source test apparatus ports should start transmitting multicast traffic to appropriate multicast groups and start sending Data-header Registers [[RFC4601](#)].
- d. Steps 6-8 from [section 7.1](#).

For all mandatory stimuli defined in [section 7](#) perform following steps:

- a. Steps a-d from paragraph above
- b. Steps 2-7 from [section 7.2](#)

Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of number of mVPNs achieved. It is DESIRED to include the same data for at least 5 different values of number of mVPNs (i.e. for at least 5 test case instances).

## [10](#) Test Cases Specific to PIM PE-PE signaling

Test cases in this section are applicable only to MVPN architectures which use PIM protocol for PE-PE transmission of C-Multicast routing information, such as "ROSEN-8" architecture.

### [10.1](#) PIM Neighborships Scale

#### Test Objective:

To determine maximum number of PIM C-instance neighborships across MI-PMSIs that PE router can create and maintain. Amount of most of other MVPN Metric such as C-instance multicast routes is minimized in this test case.

Dry, et al.

Expires May 9, 2008

[Page 35]

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

#### Metric Variables Relationships:

$$\text{Num\_MC\_C\_ints} = \text{Num\_PIM\_C\_neigh} = \text{Num\_mVPN}$$
$$\text{Num\_*G\_C} = \text{Num\_SG\_C} = 2 * \text{Num\_mVPN}$$
$$\text{Num\_OIF\_C} = 4 * \text{Num\_mVPN}$$

#### Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT) used: NO
4. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs configured on DUT (Num\_mVPN): varies
  - b. Number of remote PEs: varies
  - c. Number of C-instance multicast groups : 2
  - d. Ratio of ingress vs. egress C-instance groups: 50%:50%

## Test Execution Procedure:

Number of C-instance PIM neighborships across MI-PMSIs is proportional to product of number of mVPNs DUT belongs to and average number of PEs belonging to the same mVPNs.

Test will consist of finding maximum number of C-instance PIM neighborships across MI-PMSIs by varying average number of PEs per mVPN for set of fixed values of number of mVPNs. Procedure is as follows:

1. Configure 100 mVPNs on DUT. Execute number of test case instances where in each test case instance number of PE routers belonging to each mVPN is varied until maximum number of such PEs is found. All mVPNs should have same number of PE routers.
2. Repeat step 1 for  $100 \times i$  mVPNs where " $i=2^N$ " where N is integer value for which either maximum number of PEs per mVPN becomes smaller than one or maximum number of mVPNs found in test case 9.1 is reached.

For each test case instance perform steps 1-8 from [section 7.1](#). and 1-7 from [section 7.2](#) for all mandatory stimuli in [section 7](#).

## Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of average number of PEs for every tested value of number of mVPNs per PE (Num\_mVPN). It is DESIRED to include the same data for at least 5 different values of number of PEs for each of tested values of number of mVPNs per PE (i.e. for at least 5 test case instances per each tested value of number of mVPNs).

## [10.2](#) PIM C-instances J/P Suppression Effectiveness

### Test Objective:

If PIM is used as PE-PE signaling, optional feature of PIM J/P Suppression is performed by all PE devices that have receivers behind

them (refer to such PE as egress PE). Goal of this test case is to assess capability of egress PE to perform C-instance J/P suppression for large amount of C-instance multicast routes.

Metric Variables Relationships:

$$\text{Num\_MC\_C\_ints} = \text{Num\_mVPN}$$

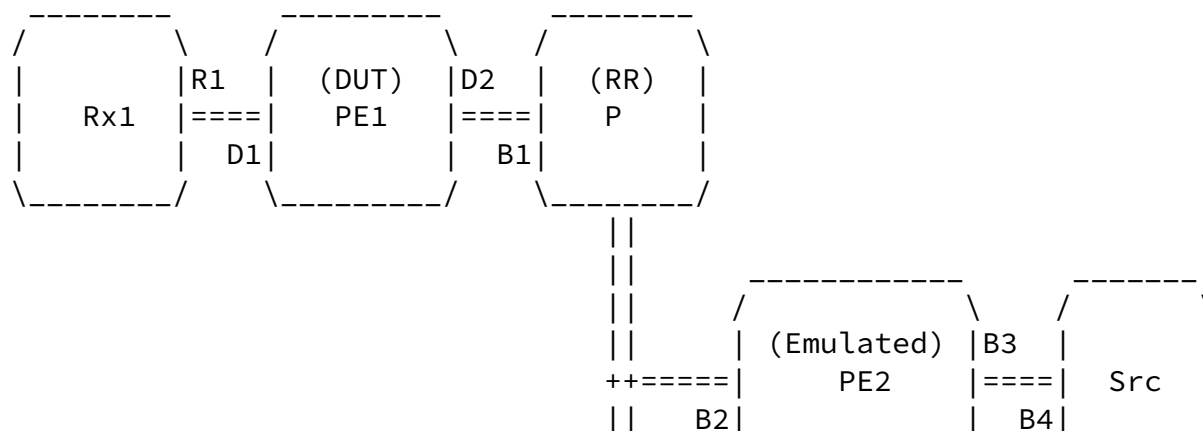
$$(\text{Num\_*G\_C} + \text{Num\_SG\_C}) \gg \text{Num\_mVPN}$$

$$\text{Num\_OIF\_C} \gg \text{Num\_mVPN}$$


Test Setup:

Following test setup MUST be performed prior to executing this test case

## 1. Topology:







1. Protocol for MI-PMSI P-instance PIM groups (if PIM used): ASM or Bi-dir
2. S-PMSI (Data MDT) used: NO
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs (Num\_mVPN) configured on DUT (Num\_mVPN): varies
  - b. Number of PIM VPN C-interfaces: 1
  - c. Number of remote PEs: 2
  - d. Ratio of ingress vs. egress C-instance groups: 0%:100%

#### Test Execution Procedure:

For maximum number of C-instances multicast routes obtained in test case 9.3 for 100 mVPNs and 100% C-instance mroutes in egress direction perform following:

- 1) Establish all PIM sessions required to emulate defined topology
- 2) Perform all C-instance PIM joins from "Rx2" (test apparatus port B5 in topology diagram)

- 3) Start all traffic from "Src" (test apparatus port R1 in topology diagram) and wait until steady state is achieved.
- 4) On C-instance PIM session used for PE-PE signaling of test apparatus port "Src" (B2) measure number of C-instance J/Ps received in 10 minute (J1) interval and calculate rate of J/Ps as  $JR1 = J1 / (60 \times 10)$ .

- 5) Perform all C-instance PIM joins from test apparatus port "Rx1" (R1) and wait until steady state is achieved on DUT.
- 6) On C-instance PIM session used for PE-PE signaling of test apparatus port "Src" (B2) measure number of C-instance J/Ps received in 10 minute (J2) interval and calculate rate of J/Ps as  $JR2 = J2 / (60 \times 10)$ .
- 7) If  $JR2 < 1.2 \times JR1$  we can conclude that DUT is suppressing J/P messages successfully.

Repeat steps 1-7, for maximum number of C-instances multicast routes obtained in test case 9.3 for maximum number of mVPN and 100% state in egress direction.

Note that no failure recovery testing is required in this test case.

#### Test Result Report:

Data listed in 8.1 MUST be reported in tabular format for all test case instances. In addition rates JR1 and JR2 MUST be reported. Optionally one can report absolute numbers or rates of number of PIM J/P PDUs transmitted by DUT and PE3 (test apparatus port B5).

## [11](#) Test Cases Specific to PIM MI-PMSI trees

Test cases in this section are applicable only to MVPN architectures which use PIM protocol to create 'tunnels' that instantiate MI-PMSIs and optional S-PMSIs, such as "ROSEN-8" architecture.

### 11.1 Default MDT's (MI-PMSI's) PIM P-Instance Mroutes Scale

#### Test Objective:

Dry, et al.

Expires May 9, 2008

[Page 39]

To determine maximum number of mVPNs and PE routers per mVPN when PIM P-instance is using protocol variant that generates maximum amount of PIM P-instance mroutes. Amount of most of other MVPN Metric such C-instance multicast mroutes is minimized in this test case.

## Metric Variables Relationships:

$\text{Num\_MC\_C\_ints} = \text{Num\_mVPN}$

$\text{Num\_*G\_C} = \text{Num\_SG\_C} = 2 * \text{Num\_mVPN}$

$\text{Num\_OIF\_C} = 4 * \text{Num\_mVPN}$

## Test Setup:

Following test setup MUST be performed prior to executing this test case

1. Protocol for MI-PMSI P-instance PIM groups: SSM
2. S-PMSI (Data MDT) used: NO
3. Multicast Control Plane Profile (all per mVPN except a.; all from DUT's perspective):
  - a. Number of MVPNs configured on DUT (Num\_mVPN): varies
  - b. Number of remote PEs: varies
  - c. Number of C-instance multicast groups : 2
  - d. Ratio of ingress vs. egress C-instance groups: 50%:50

## Test Execution Procedure:

Amount of PIM P-instance mroutes on PE router created by default MDTs (MI-PMSIs) depends in general on choice of PIM protocol variant, number of mVPNs and average number of PE routers per mVPN. In order to assess the impact of PIM P-instance mroutes created by MVPN default MDTs (MI-PMSIs) has on resources, test case 10.1 will be repeated with changing PIM P-instance protocol mode to SSM. Note that test cases 9.1-9.7 and 10.1-2 use Bi-dir PIM or PIM-SM (ASM) with SPT threshold of infinity in order to minimize impact PIM P-instance mroutes has on resources while focusing on characterizing other variables described in test cases 9.1-9.7 and 10.1-2.

## Test Result Report:

Data listed in 8.1 and 8.2 MUST be reported in tabular format for at least maximum value of average number of PEs for every tested value of number of mVPNs per PE. It is DESIRED to include the same

data for at least 5 different values of number of PEs for each of tested values of number of mVPNs per PE(i.e. for at least 5 test case instances per each tested value of number of mVPNs).

## 12 Security Considerations

Documents of this type do not directly affect the security of the Internet or of corporate networks as long as benchmarking is not performed on devices or systems connected to operating networks.

## 13 IANA Considerations

This document requires no IANA considerations.

## 14 Acknowledgments

We would like to thank Aamer Akhter, Arjen Boers, Yiqun Cai, Min Li, Amal Maalouf, Mike McBride, Ciprian Popoviciu, Dan Williston, Rajiv Asati and Thomas Morin for their valuable feedback on content of this draft. We would like to thank Nick Satsia for his support with test verification of this draft.

## 15 References

### 15.1 Normative References

[MVPN-REQ] T. Morin, Ed., "Requirements for Multicast in L3 Provider-Provisioned VPNs", [draft-ietf-l3vpn-ppvnp-mcast-reqts-09.txt](#)

[L3VPN-MCAST] E. Rosen, R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", [draft-ietf-l3vpn-2547bis-mcast-03.txt](#)

[RFC4364] E. Rosen, Y. Rekhter, "'BGP/MPLS IP Virtual Private Networks (VPNs)'"

[RFC4601] B. Fenner, M. Handley, H. Holbrook, I. Kouvelas, "'Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification'"

## [15.2](#) Informative References

[ROSEN-8] E. Rosen, Y. Cai, I. Wijnands, "Multicast in MPLS/BGP IP VPNs", [draft-rosen-vpn-mcast-08.txt](#)

[MVPN-BCP] Y. Cai, M. McBride, C. Hall, M. Napierala, "'Multicast VPN Deployment Recommendations'", [draft-ycai-mboned-mvpn-deploy-00.txt](#)

## Author's Addresses

Silvija A. Dry  
Cisco  
7025 Kit Creek Rd.  
Research Triangle Park, NC 27709  
sdry@cisco.com

Fernando Calabria  
Cisco  
7025 Kit Creek Rd.  
Research Triangle Park, NC 27709  
fcalabri@cisco.com

Maria Napierala  
AT&T  
200 Laurel Avenue  
Middletown, NJ 07748  
mnapierala@att.com

Yuji Kamite  
NTT Communications  
Tokyo Opera City Tower  
3-20-2 Nishi Shinjuku, Shinjuku-ku  
Tokyo 163-1421  
Japan  
y.kamite@ntt.com

Ian Yee Yan Fung  
Cisco Systems, Inc.  
ifung@cisco.com

## Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

---

Internet-Draft   Multicast VPN Scalability Benchmarking   November 2007

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

#### Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

#### Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

Dry, et al.

Expires May 9, 2008

[Page 44]