

Network Working Group
Internet-Draft
Updates: 2544 (if approved)
Intended status: Informational
Expires: April 25, 2013

S. Bradner
Harvard University
K. Dubray
Juniper Networks
J. McQuaid
Turnip Video
A. Morton
AT&T Labs
October 22, 2012

RFC 2544 Applicability Statement:
Use on Production Networks Considered Harmful
draft-ietf-bmwg-2544-as-08

Abstract

Benchmarking Methodology Working Group (BMWG) has been developing key performance metrics and laboratory test methods since 1990, and continues this work at present. The methods described in RFC 2544 are intended to generate traffic that overloads network device resources in order to assess their capacity. Overload of shared resources would likely be harmful to user traffic performance on a production network, and there are further negative consequences identified with production application of the methods. This memo clarifies the scope of RFC 2544 and other IETF BMWG benchmarking work for isolated test environments only, and encourages new standards activity for measurement methods applicable outside that scope.

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 <http://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 April 25, 2013.

Copyright Notice

Copyright (c) 2012 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. Scope and Goals	4
3. The Concept of an Isolated Test Environment	4
4. Why RFC 2544 Methods are intended only for ITE	4
4.1. Experimental Control and Accuracy	4
4.2. Containing Damage	5
5. Advisory on RFC 2544 Methods in Production Networks	5
6. Considering Performance Testing in Production Networks	6
7. Security Considerations	7
8. IANA Considerations	7
9. Acknowledgements	8
10. Appendix - Example of RFC 2544 method failure in production network measurement	8
11. References	9
11.1. Normative References	9
11.2. Informative References	10
Authors' Addresses	10

1. Introduction

This memo clarifies the scope and use of IETF Benchmarking Methodology Working Group (BMWG) tests including [RFC2544], which discusses and defines several tests that may be used to characterize the performance of a network interconnecting device. All readers of this memo must read and fully understand [RFC2544].

Benchmarking methodologies (beginning with [RFC2544]) have always relied on test conditions that can only be produced and replicated reliably in the laboratory. These methodologies are not appropriate for inclusion in wider specifications such as:

1. Validation of telecommunication service configuration, such as the Committed Information Rate (CIR).
2. Validation of performance metrics in a telecommunication Service Level Agreement (SLA), such as frame loss and latency.
3. Telecommunication service activation testing, where traffic that shares network resources with the test might be adversely affected.

Above, we distinguish "telecommunication service" (where a network service provider contracts with a customer to transfer information between specified interfaces at different geographic locations) from the generic term "service". Below, we use the adjective "production" to refer to networks carrying live user traffic. [RFC2544] used the term "real-world" to refer to production networks and to differentiate them from test networks.

Although RFC 2544 has been held up as the standard reference for such testing, we believe that the actual methods used vary from [RFC2544] in significant ways. Since the only citation is to [RFC2544], the modifications are opaque to the standards community and to users in general.

Since applying the test traffic and methods described in [RFC2544] on a production network risks causing overload in shared resources there is direct risk of harming user traffic if the methods are misused in this way. Therefore, IETF BMWG developed this Applicability Statement for [RFC2544] to directly address the situation.

1.1. Requirements Language

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 [RFC2119].

2. Scope and Goals

This memo clarifies the scope of [RFC2544] with the goal to provide guidance to the industry on its applicability, which is limited to laboratory testing.

3. The Concept of an Isolated Test Environment

An Isolated Test Environment (ITE) used with [RFC2544] methods (as illustrated in Figures 1 through 3 of [RFC2544]) has the ability to:

- o contain the test streams to paths within the desired set-up
- o prevent non-test traffic from traversing the test set-up

These features allow unfettered experimentation, while at the same time protecting lab equipment management/control LANs and other production networks from the unwanted effects of the test traffic.

4. Why RFC 2544 Methods are intended only for ITE

The following sections discuss some of the reasons why [RFC2544] methods are applicable only for isolated laboratory use, and the consequences of applying these methods outside the lab environment.

4.1. Experimental Control and Accuracy

All of the tests described in RFC 2544 require that the tester and device under test are the only devices on the networks that are transmitting data. The presence of other traffic (unwanted on the ITE network) would mean that the specified test conditions have not been achieved and flawed results are a likely consequence.

If any other traffic appears and the amount varies over time, the repeatability of any test result will likely depend to some degree on the amount and variation of the other traffic.

The presence of other traffic makes accurate, repeatable, and consistent measurements of the performance of the device under test very unlikely, since the complete details of test conditions will not be reported.

For example, the RFC 2544 Throughput Test attempts to characterize a maximum reliable load, thus there will be testing above the maximum that causes packet/frame loss. Any other sources of traffic on the network will cause packet loss to occur at a tester data rate lower

than the rate that would be achieved without the extra traffic.

4.2. Containing Damage

[RFC2544] methods, specifically to determine Throughput as defined in [RFC1242] and other benchmarks, may overload the resources of the device under test, and may cause failure modes in the device under test. Since failures can become the root cause of more wide-spread failure, it is clearly desirable to contain all test traffic within the ITE.

In addition, such testing can have a negative effect on any traffic that shares resources with the test stream(s) since, in most cases, the traffic load will be close to the capacity of the network links.

Appendix C.2.2 of [RFC2544] (as adjusted by errata) gives the private IPv4 address range for testing:

"...The network addresses 198.18.0.0 through 198.19.255.255 have been assigned to the BMWG by the IANA for this purpose. This assignment was made to minimize the chance of conflict in case a testing device were to be accidentally connected to part of the Internet. The specific use of the addresses is detailed below."

In other words, devices operating on the Internet may be configured to discard any traffic they observe in this address range, as it is intended for laboratory ITE use only. Thus, if testers using the assigned testing address ranges are connected to the Internet and test packets are forwarded across the Internet, it is likely that the packets will be discarded and the test will not work.

We note that a range of IPv6 addresses has been assigned to BMWG for laboratory test purposes, in [RFC5180] (as amended by errata).

See the Security Considerations Section below for further considerations on containing damage.

5. Advisory on RFC 2544 Methods in Production Networks

The tests in [RFC2544] were designed to measure the performance of network devices, not of networks, and certainly not production networks carrying user traffic on shared resources. There will be undesirable consequences when applying these methods outside the isolated test environment.

One negative consequence stems from reliance on frame loss as an indicator of resource exhaustion in [RFC2544] methods. In practice,

link-layer and physical-layer errors prevent production networks from operating loss-free. The [RFC2544] methods will not correctly assess Throughput when loss from uncontrolled sources is present. Frame loss occurring at the SLA levels of some networks could affect every iteration of Throughput testing (when each step includes sufficient packets to experience facility-related loss). Flawed results waste the time and resources of the testing service user and of the service provider when called to dispute the measurement. These are additional examples of harm that compliance with this advisory should help to avoid. See the Appendix for an example.

The methods described in [RFC2544] are intended to generate traffic that overloads network device resources in order to assess their capacity. Overload of shared resources would likely be harmful to user traffic performance on a production network. These tests **MUST NOT** be used on production networks and as discussed above. The tests will not produce a reliable or accurate benchmarking result on a production network.

[RFC2544] methods have never been validated on a network path, even when that path is not part of a production network and carrying no other traffic. It is unknown whether the tests can be used to measure valid and reliable performance of a multi-device, multi-network path. It is possible that some of the tests may prove valid in some path scenarios, but that work has not been done or has not been shared with the IETF community. Thus, such testing is contra-indicated by the BMWG.

6. Considering Performance Testing in Production Networks

The IETF has addressed the problem of production network performance measurement by chartering a different working group: IP Performance Metrics (IPPM). This working group has developed a set of standard metrics to assess the quality, performance, and reliability of Internet packet transfer services. These metrics can be measured by network operators, end users, or independent testing groups. We note that some IPPM metrics differ from RFC 2544 metrics with similar names, and there is likely to be confusion if the details are ignored.

IPPM has not yet standardized methods for raw capacity measurement of Internet paths. Such testing needs to adequately consider the strong possibility for degradation to any other traffic that may be present due to congestion. There are no specific methods proposed for activation of a packet transfer service in IPPM at this time. Thus, individuals who need to conduct capacity tests on production networks should actively participate in standards development to ensure their

methods receive appropriate industry review and agreement, in the IETF or in alternate standards development organizations.

Other standards may help to fill gaps in telecommunication service testing. For example, the IETF has many standards intended to assist with network operation, administration and maintenance (OAM), and ITU-T Study Group 12 has a Recommendation on service activation test methodology [Y.1564].

The world will not spin off axis while waiting for appropriate and standardized methods to emerge from the consensus process.

7. Security Considerations

This Applicability Statement intends to help preserve the security of the Internet by clarifying that the scope of [RFC2544] and other BMWG memos are all limited to testing in a laboratory ITE, thus avoiding accidental Denial of Service attacks or congestion due to high traffic volume test streams.

All Benchmarking activities are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the other constraints [RFC2544].

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 device under test/ system under test (DUT/SUT).

Special capabilities SHOULD NOT exist in the DUT/SUT specifically for benchmarking purposes. Any implications for network security arising from the DUT/SUT SHOULD be identical in the lab and in production networks.

8. IANA Considerations

This memo makes no requests of IANA.

9. Acknowledgements

Thanks to Matt Zekauskas, Bill Cerveney, Barry Constantine, Curtis Villamizar, David Newman, and Adrian Farrel for suggesting improvements to this memo.

Specifically, Al Morton would like to thank his co-authors, who constitute the complete set of Chairmen-Emeritus of the BMWG, for returning from other pursuits to develop this statement and see it through to approval. This has been a rare privilege; one that likely will not be matched in the IETF again:

Scott Bradner	served as Chairman from 1990 to 1993
Jim McQuaid	served as Chairman from 1993 to 1995
Kevin Dubray	served as Chairman from 1995 to 2006

It's all about the band.

10. Appendix - Example of RFC 2544 method failure in production network measurement

This Appendix provides an example illustrating how [RFC2544] methods applied on production networks can easily produce a form of harm from flawed and misleading results.

The [RFC2544] Throughput benchmarking method usually includes the following steps:

- a. Set the offered traffic level, less than max of the ingress link(s).
- b. Send the test traffic through the device under test (DUT) and count all frames successfully transferred.
- c. If all frames are received, increment traffic level and repeat step b.
- d. If one or more frames are lost, the level is in the DUT-overload region (Step b may be repeated at a reduced traffic level to more exactly determine the maximum rate at which none of the frames are dropped by the DUT, defined as the Throughput [RFC1242]).
- e. Report the Throughput values, the x-y of graph of frame size and Throughput, and other information in accordance with [RFC2544].

In this method, frame loss is the sole indicator of overload and therefore the determining factor in the measurement of Throughput

using the [RFC2544] methodology (even though the results may not report frame loss per se).

Frame loss is subject to many factors in addition to operating above the Throughput traffic level. These factors include optical interference (which may be due to dirty interfaces, cross-over from other signals, fiber bend and temperature, etc.) and electrical interference (caused by local sources of radio signals, electrical spikes, solar particles, etc.). In the laboratory environment many of these issues can be carefully controlled through cleaning and isolation. Since [RFC2544] methodologies are primarily intended to test devices and not paths, the total length of path, the number of interfaces, and compound risk of random frame loss can be kept to a minimum.

In a production network, however, there will be many interfaces and many kilometres of path under test. This considerably increases the risk of random frame loss.

The risk of frame loss caused by outside effects is significantly higher in production networks, and significantly higher with long paths (both those with long physical path lengths, and those with large numbers of interfaces in the path). Thus, the risk of falsely low reported Throughput using an [RFC2544] methodology test is considerably increased in a production network.

Therefore, to successfully conduct tests with similar objectives to those in [RFC2544] in a production network, it will be necessary to develop modifications to the methodologies defined in [RFC2544] and standards to describe them. See [Bryant] for an in-progress effort and [Y.1564] for an approved method adapted to production service activation.

11. References

11.1. Normative References

- [RFC1242] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [RFC5180] Popoviciu, C., Hamza, A., Van de Velde, G., and D.

Dugatkin, "IPv6 Benchmarking Methodology for Network Interconnect Devices", RFC 5180, May 2008.

11.2. Informative References

- [Bryant] Bonica, R. and S. Bryant, "Work in-progress, "RFC2544 Testing in Production Network", (draft-bb-2544like-production-tests-00)", October 2012.
- [Y.1564] ITU-T Recommendation Y.1564, "Ethernet Service Activation Test Methodology", March 2011.

Authors' Addresses

Scott Bradner
Harvard University
29 Oxford St.
Cambridge, MA 02138
USA

Phone: +1 617 495 3864
Fax:
Email: sob@harvard.edu
URI: <http://www.sobco.com>

Kevin Dubray
Juniper Networks

Phone:
Fax:
Email: kdubray@juniper.net
URI:

Jim McQuaid
Turnip Video
6 Cobbleridge Court
Durham, North Carolina 27713
USA

Phone: +1 919-619-3220
Fax:
Email: jim@turnipvideo.com
URI: www.turnipvideo.com

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown,, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acmorton@att.com
URI: <http://home.comcast.net/~acmacm/>

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 5, 2013

M. Hamilton
Ixia
S. Banks
Aerohive Networks
Feb 2013

Benchmarking Methodology for Content-Aware Network Devices
draft-ietf-bmwg-ca-bench-meth-04

Abstract

This document defines a set of test scenarios and metrics that can be used to benchmark content-aware network devices. The scenarios in the following document are intended to more accurately predict the performance of these devices when subjected to dynamic traffic patterns. This document will operate within the constraints of the Benchmarking Working Group charter, namely black box characterization in a laboratory environment.

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 <http://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 August 5, 2013.

Copyright Notice

Copyright (c) 2013 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Requirements Language	5
2. Scope	5
3. Test Setup	5
3.1. Test Considerations	6
3.2. Clients and Servers	6
3.3. Traffic Generation Requirements	6
3.4. Discussion of Network Limitations	6
3.5. Framework for Traffic Specification	8
3.6. Multiple Client/Server Testing	8
3.7. Device Configuration Considerations	8
3.7.1. Network Addressing	9
3.7.2. Network Address Translation	9
3.7.3. TCP Stack Considerations	9
3.7.4. Other Considerations	9
4. Benchmarking Tests	9
4.1. Maximum Application Session Establishment Rate	10
4.1.1. Objective	10
4.1.2. Setup Parameters	10
4.1.3. Procedure	10
4.1.4. Measurement	10
4.1.4.1. Maximum Application Flow Rate	10
4.1.4.2. Application Flow Duration	11
4.1.4.3. Application Efficiency	11
4.1.4.4. Application Flow Latency	11
4.2. Application Throughput	11
4.2.1. Objective	11
4.2.2. Setup Parameters	11
4.2.3. Procedure	12
4.2.4. Measurement	12
4.2.4.1. Maximum Throughput	12
4.2.4.2. Maximum Application Flow Rate	12
4.2.4.3. Application Flow Duration	12
4.2.4.4. Application Efficiency	12
4.2.4.5. Packet Loss	12
4.2.4.6. Application Flow Latency	12
4.3. Malformed Traffic Handling	13
4.3.1. Objective	13
4.3.2. Setup Parameters	13
4.3.3. Procedure	13
4.3.4. Measurement	13

5. IANA Considerations	13
6. Security Considerations	13
7. References	14
7.1. Normative References	14
7.2. Informative References	15
7.3. URL References	15
Appendix A. Example Traffic Mix	15
Appendix B. Malformed Traffic Algorithm	17
Authors' Addresses	19

1. Introduction

Content-aware and deep packet inspection (DPI) device deployments have grown significantly in recent years. No longer are devices simply using Ethernet and IP headers to make forwarding decisions. This class of device now uses application-specific data to make these decisions. For example, a web-application firewall (WAF) may use search criteria upon the HTTP uniform resource indicator (URI) [1] to decide whether a HTTP GET method may traverse the network. In the case of lawful/legal intercept technology, a device could use the phone number within the Session Description Protocol [14] to determine whether a voice-over-IP phone may be allowed to connect. In addition to the development of entirely new classes of devices, devices that could historically be classified as 'stateless' or raw forwarding devices are now performing DPI functionality. Devices such as core and edge routers are now being developed with DPI functionality to make more intelligent routing and forwarding decisions.

The Benchmarking Working Group (BMWG) has historically produced Internet Drafts and Requests for Comment that are focused specifically on creating output metrics that are derived from a very specific and well-defined set of input parameters that are completely and unequivocally reproducible from test bed to test bed. The end goal of such methodologies is to, in the words of the RFC 2544 [2], reduce "specsmanship" in the industry and hold vendors accountable for performance claims.

The end goal of this methodology is to generate performance metrics in a lab environment that will closely relate to actual observed performance on production networks. By utilizing dynamic traffic patterns relevant to modern networks, this methodology should be able to closely tie laboratory and production metrics. It should be further noted that any metrics acquired from production networks SHOULD be captured according to the policies and procedures of the IPPM or PMOL working groups.

An explicit non-goal of this document is to replace existing methodology/terminology pairs such as RFC 2544 [2]/RFC 1242 [3] or RFC 3511 [4]/RFC 2647 [5]. The explicit goal of this document is to create a methodology more suited for modern devices while complementing the data acquired using existing BMWG methodologies. This document does not assume completely repeatable input stimulus. The nature of application-driven networks is such that a single dropped packet inherently changes the input stimulus from a network perspective. While application flows will be specified in great detail, it simply is not practical to require totally repeatable input stimulus.

1.1. Requirements Language

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 [6].

2. Scope

Content-aware devices take many forms, shapes and architectures. These devices are advanced network interconnect devices that inspect deep into the application payload of network data packets to do classification. They may be as simple as a firewall that uses application data inspection for rule set enforcement, or they may have advanced functionality such as performing protocol decoding and validation, anti-virus, anti-spam and even application exploit filtering. The document will universally call these devices middleboxes, as defined by RFC 3234 [7].

This document is strictly focused on examining performance and robustness across a focused set of metrics: throughput (min/max/avg/sample std dev), transaction rates (successful/failed), application response times, concurrent flows, and unidirectional packet latency. None of the metrics captured through this methodology are specific to a device and the results are DUT implementation independent. Functional testing of the DUT is outside the scope of this methodology.

Devices such as firewalls, intrusion detection and prevention devices, wireless LAN controllers, application delivery controllers, deep packet inspection devices, wide-area network (WAN) optimization devices, and unified threat management systems generally fall into the content-aware category. While this list may become obsolete, these are a subset of devices that fall under this scope of testing.

3. Test Setup

This document will be applicable to most test configurations and will not be confined to a discussion on specific test configurations. Since each DUT/SUT will have their own unique configuration, users SHOULD configure their device with the same parameters that would be used in the actual deployment of the device or a typical deployment, if the actual deployment is unknown. A summary of the DUT configuration MUST be published with the final benchmarking results. In order to improve repeatability, the published configuration information SHOULD include command-line scripts used to configure the DUT, if any, and SHOULD also include any configuration information

for the test equipment used."

3.1. Test Considerations

3.2. Clients and Servers

Content-aware device testing SHOULD involve multiple clients and multiple servers. As with RFC 3511 [4], this methodology will use the terms virtual clients/servers because both the client and server will be represented by the tester and not actual clients/servers. Similarly defined in RFC 3511 [4], a data source may emulate multiple clients and/or servers within the context of the same test scenario. The test report SHOULD indicate the number of virtual clients/servers used during the test. IANA has reserved address ranges for laboratory characterization. These are defined for IPv4 and IPv6 by RFC 2544 Appendix C [2] and RFC 5180 Section 5.2 [8] respectively and SHOULD be consulted prior to testing.

3.3. Traffic Generation Requirements

The explicit purposes of content-aware devices vary widely, but these devices use information deeper inside the application flow to make decisions and classify traffic. This methodology will utilize traffic flows that resemble real application traffic without utilizing captures from live production networks. Application Flows, as defined in Section 1.1 RFC 2724 [9] are able to be well-defined without simply referring to a network capture. An example traffic template is defined and listed in Appendix A of this document. A user of this methodology is free to utilize the example mix as provided in the appendix. If a user of this methodology understands the traffic patterns in their production network, that user MAY use the template provided in Appendix A to describe a traffic mix appropriate for their environment. In all cases, users MUST report the traffic mix used in the test, and SHOULD report this using a template similar to that in Appendix A.


The test tool SHOULD be able to create application flows between every client and server, regardless of direction. The tester SHOULD be able to open TCP connections on multiple destination ports and SHOULD be able to direct UDP traffic to multiple destination ports.

3.4. Discussion of Network Limitations

Prior to executing the methodology as outlined in the following sections, it is imperative to understand the implications of utilizing representative application flows for the traffic content of the benchmarking effort. One interesting aspect of utilizing application flows is that each flow is inherently different from

every other application flow. The content of each flow will vary from application to application, and in most cases, even varies within the same type of application flow. The following description of the methodology will individually benchmark every individual type and subset of application flow, prior to performing similar tests with a traffic mix as specified either by the example mix in Appendix A, or as defined by the user of this methodology.

The purpose of this process is to ensure that any performance implications that are discovered during the mixed testing aren't due to the inherent physical network limitations. As an example of this phenomena, it is useful to examine a network device inserted into a single path, as illustrated in the following diagram.



Simple Inline DUT Configuration

Figure 1: Simple Middle-box Example

For the purpose of this discussion, let's take a hypothetical application flow that utilizes UDP for the transport layer. Assume that the sample transaction we will be using to model this particular flow requires 10 UDP datagrams to complete the transaction. For simplicity, each datagram within the flow is exactly 64 bytes, including associated Ethernet, IP, and UDP overhead. With any network device, there are always three metrics which interact with each other: number of concurrent application flows, number of application flows per second, and layer-7 throughput.

Our example test bed is a single-path device connected by 1 gigabit Ethernet links. The purpose of this benchmark effort is to quantify the number of application flows per second that may be processed through our device under test. Let's assume that the result from our scenario is that the DUT is able to process 10,000 application flows per second. The question is whether that ceiling is the actual ceiling of the device, or if it is actually being limited by one of the other metrics. If we do the appropriate math, 10000 flows per second, with each flow at 640 total bytes means that we are achieving an aggregate bitrate of roughly 49 Mbps. This is dramatically less than the 1 gigabit physical link we are using. We can conclude that 10,000 flows per second is in fact the performance limit of the device.

If we change the example slightly and increase the size of each

datagram to 1312 bytes, then it becomes necessary to recompute the load. Assuming the same observed DUT limitation of 10,000 flows per second, it must be ensured that this is an artifact of the DUT, and not of physical limitations. For each flow, we'll require 104,960 bits. 10,000 flows per second implies a throughput of roughly 1 Gbps. At this point, we cannot definitively answer whether the DUT is actually limited to 10,000 flows per second. If we are able to modify the scenario, and utilize 10 Gigabit interfaces, then perhaps the flow per second ceiling will be reached at a higher number than 10,000.

This example illustrates why a user of this methodology SHOULD benchmark each application variant individually to ensure that the cause of a measured limit is fully understood

3.5. Framework for Traffic Specification

The following table SHOULD be specified for each application flow variant.

- o Data Exchanged By Flow, Bits
- o Offered Percentage of Total Flows
- o Transport Protocol(s)
- o Destination Port(s)

3.6. Multiple Client/Server Testing

In actual network deployments, connections are being established between multiple clients and multiple servers simultaneously. Device vendors have been known to optimize the operation of their devices for easily defined patterns. The connection sequence ordering scenarios a device will see on a network will likely be much less deterministic. In fact, many application flows have multiple layer 4 connections within a single flow, with client and server reversing roles. Flow initiation SHOULD be in a pseudo-random manner across ingress ports.

3.7. Device Configuration Considerations

The configuration of the DUT may have an effect on the observed results of the following methodology. A comprehensive, but certainly not exhaustive, list of potential considerations is listed below.

3.7.1. Network Addressing

The IANA has issued a range of IP addresses to the BMWG for purposes of benchmarking. Please refer to RFC 2544 [2] and RFC 5180 [8] for more details. If more IPv4 addresses are required than the RFC 2544 allotment provides, then allocations from the private address space as defined in RFC 1918 [10] may be used.

3.7.2. Network Address Translation

Many content-aware devices are capable of performing Network Address Translation (NAT) [5]. If the final deployment of the DUT will have this functionality enabled, then the DUT SHOULD also have it enabled during the execution of this methodology. It MAY be beneficial to perform the test series in both modes in order to determine the performance differential when using NAT. The test report SHOULD indicate whether NAT was enabled during the testing process.

3.7.3. TCP Stack Considerations

The IETF has historically provided guidance and information on TCP stack considerations. This methodology is strictly focused on performance metrics at layers above 4, thus does not specifically define any TCP stack configuration parameters of either the tester or the DUTs. The TCP configuration of the tester MUST remain constant across all DUTs in order to ensure comparable results. While the following list of references is not exhaustive, each document contains a relevant discussion on TCP stack considerations.

The general IETF TCP roadmap is defined in RFC 4614 [11] and congestion control algorithms are discussed in Section 2 of RFC 3148 [12] with even more detailed references. TCP receive and congestion window sizes are discussed in detail in RFC 6349 [13].

3.7.4. Other Considerations

Various content-aware devices will have widely varying feature sets. In the interest of representative test results, the DUT features that will likely be enabled in the final deployment SHOULD be used. This methodology is not intended to advise on which features should be enabled, but to suggest using actual deployment configurations.

4. Benchmarking Tests

Each of the following benchmark scenarios SHOULD be run with each of the single application flow templates. Upon completion of all iterations, the mixed test SHOULD be completed, subject to the

traffic mix as defined by the user.

4.1. Maximum Application Session Establishment Rate

4.1.1. Objective

To determine the maximum rate through which a device is able to establish and complete application flows as defined by draft-ietf-bmwg-ca-bench-term-00.

4.1.2. Setup Parameters

The following parameters SHOULD be used and reported for all tests:

For each application protocol in use during the test run, the table provided in Section 3.5 SHOULD be published.

4.1.3. Procedure

The test SHOULD generate application network traffic that meets the conditions of Section 3.3. The traffic pattern SHOULD begin with an application flow rate of 10% of expected maximum. The test SHOULD be configured to increase the attempt rate in units of 10% up through 110% of expected maximum. In the case where expected maximum is limited by physical link rate as discovered through Appendix A, the maximum rate will attempted will be 100% of expected maximum, or "wire-speed performance". The duration of each loading phase SHOULD be at least 30 seconds. This test MAY be repeated, each subsequent iteration beginning at 5% of expected maximum and increasing session establishment rate to 110% of the maximum observed from the previous test run.

This procedure MAY be repeated any reasonable number of times with the results being averaged together.

4.1.4. Measurement

The following metrics MAY be determined from this test, and SHOULD be observed for each application protocol within the traffic mix:

4.1.4.1. Maximum Application Flow Rate

The test tool SHOULD report the maximum rate at which application flows were completed, as defined by RFC 2647 [5], Section 3.7. This rate SHOULD be reported individually for each application protocol present within the traffic mix.

4.1.4.2. Application Flow Duration

The test tool SHOULD report the minimum, maximum and average application duration, as defined by RFC 2647 [5], Section 3.9. This duration SHOULD be reported individually for each application protocol present within the traffic mix.

4.1.4.3. Application Efficiency

The test tool SHOULD report the application efficiency, similarly defined for TCP by RFC 6349 [13].

$$\text{App Efficiency \%} = \frac{\text{Transmitted Bytes} - \text{Retransmitted Bytes}}{\text{Transmitted Bytes}} \times 100$$

Figure 2: Application Efficiency Percent Calculation

Note that a calculation less than 100% does not necessarily imply noticeably degraded performance since certain applications utilize algorithms to maintain a quality user experience in the face of data loss.

4.1.4.4. Application Flow Latency

The test tool SHOULD report the minimum, maximum and average amount of time an application flow member takes to traverse the DUT, as defined by RFC 1242 [3], Section 3.8. This value SHOULD be reported individually for each application protocol present within the traffic mix.

4.2. Application Throughput

4.2.1. Objective

To determine the maximum rate through which a device is able to forward bits when using application flows as defined in the previous sections.

4.2.2. Setup Parameters

The same parameter reporting procedure as described in Section 4.1.2 SHOULD be used for all tests.

4.2.3. Procedure

This test will attempt to send application flows through the device at a flow rate of 30% of the maximum, as observed in Section 4.1. This procedure MAY be repeated with the results from each iteration averaged together.

4.2.4. Measurement

The following metrics MAY be determined from this test, and SHOULD be observed for each application protocol within the traffic mix:

4.2.4.1. Maximum Throughput

The test tool SHOULD report the minimum, maximum and average application throughput.

4.2.4.2. Maximum Application Flow Rate

The test tool SHOULD report the maximum rate at which application flows were completed, as defined by RFC 2647 [5], Section 3.7. This rate SHOULD be reported individually for each application protocol present within the traffic mix.

4.2.4.3. Application Flow Duration

The test tool SHOULD report the minimum, maximum and average application duration, as defined by RFC 2647 [5], Section 3.9. This duration SHOULD be reported individually for each application protocol present within the traffic mix.

4.2.4.4. Application Efficiency

The test tool SHOULD report the application efficiency as defined in Section 4.1.4.3.

4.2.4.5. Packet Loss

The test tool SHOULD report the number of packets lost or dropped from source to destination.

4.2.4.6. Application Flow Latency

The test tool SHOULD report the minimum, maximum and average amount of time an application flow member takes to traverse the DUT, as defined by RFC 1242 [3], Section 3.13. This value SHOULD be reported individually for each application protocol present within the traffic mix.

4.3. Malformed Traffic Handling

4.3.1. Objective

To determine the effects on performance and stability that malformed traffic may have on the DUT.

4.3.2. Setup Parameters

The same parameters SHOULD be used for Transport-Layer and Application Layer Parameters previously specified in Section 4.1.2 and Section 4.2.2.

4.3.3. Procedure

This test will utilize the procedures specified previously in Section 4.1.3 and Section 4.2.3. When performing the procedures listed previously, the tester should generate malformed traffic at all protocol layers. This is commonly known as fuzzed traffic. Fuzzing techniques generally modify portions of packets, including checksum errors, invalid protocol options, and improper protocol conformance.

The process by which the tester SHOULD generate the malformed traffic is outlined in detail in Appendix B.

4.3.4. Measurement

For each protocol present in the traffic mix, the metrics specified by Section 4.1.4 and Section 4.2.4 MAY be determined. This data may be used to ascertain the effects of fuzzed traffic on the DUT.

5. IANA Considerations

This memo includes no request to IANA.

All drafts are required to have an IANA considerations section (see the update of RFC 2434 [15] for a guide). If the draft does not require IANA to do anything, the section contains an explicit statement that this is the case (as above). If there are no requirements for IANA, the section will be removed during conversion into an RFC by the RFC Editor.

6. Security Considerations

Benchmarking activities as described in this memo are limited to

technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the other constraints RFC 2544 [2].

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 mis-route traffic to the test management network

7. References

7.1. Normative References

- [1] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [2] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [3] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [4] Hickman, B., Newman, D., Tadjudin, S., and T. Martin, "Benchmarking Methodology for Firewall Performance", RFC 3511, April 2003.
- [5] Newman, D., "Benchmarking Terminology for Firewall Performance", RFC 2647, August 1999.
- [6] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [7] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.
- [8] Popoviciu, C., Hamza, A., Van de Velde, G., and D. Dugatkin, "IPv6 Benchmarking Methodology for Network Interconnect Devices", RFC 5180, May 2008.
- [9] Handelman, S., Stibler, S., Brownlee, N., and G. Ruth, "RTFM: New Attributes for Traffic Flow Measurement", RFC 2724, October 1999.
- [10] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

- [11] Duke, M., Braden, R., Eddy, W., and E. Blanton, "A Roadmap for Transmission Control Protocol (TCP) Specification Documents", RFC 4614, September 2006.
- [12] Mathis, M. and M. Allman, "A Framework for Defining Empirical Bulk Transfer Capacity Metrics", RFC 3148, July 2001.
- [13] Constantine, B., Forget, G., Geib, R., and R. Schrage, "Framework for TCP Throughput Testing", RFC 6349, August 2011.

7.2. Informative References

- [14] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [15] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

7.3. URL References

- [16] Sandvine Corporation, "<http://www.sandvine.com/general/document.download.asp?docID=58&sourceID=0>", 2012.

Appendix A. Example Traffic Mix

This appendix shows an example case of a protocol mix that may be used with this methodology. This mix closely represents the research published by Sandvine [16] in their biannual report for the first half of 2012 on North American fixed access service provider networks.

Direction	Application Flow	Options	Value
Upstream	BitTorrent	Avg Flow Size (L7) Flow Percentage	512 MB 44.4%
	HTTP	Avg Flow Size (L7) Flow Percentage	128 kB 7.3%
	Skype	Avg Flow Size (L7) Flow Percentage	8 MB 4.9%
	SSL/TLS	Avg Flow Size (L7) Flow Percentage	128 kB 3.2%
	Netflix		

Downstream	PPStream	Avg Flow Size (L7)	500 kB
		Flow Percentage	3.1%
	YouTube	Avg Flow Size (L7)	500 MB
		Flow Percentage	2.2%
	Facebook	Avg Flow Size (L7)	4 MB
		Flow Percentage	1.9%
	Teredo	Avg Flow Size (L7)	2 MB
		Flow Percentage	1.9%
	Apple iMessage	Avg Flow Size (L7)	500 MB
		Flow Percentage	1.2%
	Bulk TCP	Avg Flow Size (L7)	40 kB
		Flow Percentage	1.1%
	Netflix	Avg Flow Size (L7)	128 kB
		Flow Percentage	28.8%
	YouTube	Avg Flow Size (L7)	512 MB
		Flow Percentage	32.9%
	HTTP	Avg Flow Size (L7)	5 MB
		Flow Percentage	13.8%
	BitTorrent	Avg Flow Size (L7)	1 MB
		Flow Percentage	12.1%
iTunes	Avg Flow Size (L7)	500 MB	
	Flow Percentage	6.3%	
Flash Video	Avg Flow Size (L7)	32 MB	
	Flow Percentage	3.8%	
MPEG	Avg Flow Size (L7)	100 MB	
	Flow Percentage	2.6%	
RTMP	Avg Flow Size (L7)	100 MB	
	Flow Percentage	2.0%	
Hulu	Avg Flow Size (L7)	50 MB	
	Flow Percentage	2.0%	
	SSL/TLS	Avg Flow Size (L7)	300 MB
		Flow Percentage	1.8%

		Avg Flow Size (L7)	256 kB
		Flow Percentage	1.6%
	Bulk TCP	Avg Flow Size (L7)	500 kB
		Flow Percentage	21.1%

Table 1: Example Traffic Pattern

Appendix B. Malformed Traffic Algorithm

Each application flow will be broken into multiple transport segments, IP packets, and Ethernet frames. The malformed traffic algorithm looks very similar to the IP Stack Integrity Checker project at <http://isic.sourceforge.net>.

The algorithm is very simple and starts by defining each of the fields within the TCP/IP stack that will be malformed during transmission. The following table illustrates the Ethernet, IPv4, IPv6, TCP, and UDP fields which are able to be malformed by the algorithm. The first column lists the protocol, the second column shows the actual header field name, with the third column showing the percentage of packets that should have the field modified by the malformation algorithm.

Protocol	Header Field	Malformed %
Total Frames		1%
Ethernet	Destination MAC	0%
	Source MAC	1%
	Ethertype	1%
	CRC	1%
IP Version 4	Version	1%
	IHL	1%
	Type of Service	1%
	Total Length	1%
	Identification	1%
	Flags	1%
	Fragment Offset	1%
	Time to Live	1%
	Protocol	1%
	Header Checksum	1%
	Source Address	1%
	Destination Address	1%
	Options	1%
	Padding	1%
UDP	Source Port	1%
	Destination Port	1%
	Length	1%
	Checksum	1%
TCP	Source Port	1%
	Destination Port	1%
	Sequence Number	1%
	Acknowledgement Number	1%
	Data Offset	1%
	Reserved(3 bit)	1%
	Flags(9 bit)	1%
	Window Size	1%
	Checksum	1%
	Urgent Pointer	1%
	Options(Variable Length)	1%

Table 2: Malformed Header Values

This algorithm is to be used across the regular application flows used throughout the rest of the methodology. As each frame is emitted from the test tool, a pseudo-random number generator will

indicate whether the frame is to be malformed by creating a number between 0 and 100. If the number is less than the percentage defined in the table, then that frame will be malformed. If the frame is to be malformed, then each of the headers in the table present within the frame will follow the same process. If it is determined that a header field should be malformed, the same pseudo-random number generator will be used to create a random number for the specified header field.

Authors' Addresses

Mike Hamilton
Ixia
Austin, TX 78730
US

Phone: +1 512 636 2303
Email: mhamilton@ixiacom.com

Sarah Banks
Aerohive Networks
San Jose, CA 95134
US

Email: sbanks@aerohive.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 5, 2013

A. Morton
AT&T Labs
June 3, 2013

IMIX Genome: Specification of variable packet sizes for additional
testing
draft-ietf-bmwg-imix-genome-05

Abstract

Benchmarking Methodologies have always relied on test conditions with constant packet sizes, with the goal of understanding what network device capability has been tested. Tests with constant packet size reveal device capabilities but differ significantly from the conditions encountered in operational deployment, and so additional tests are sometimes conducted with a mixture of packet sizes, or "IMIX". The mixture of sizes a networking device will encounter is highly variable and depends on many factors. An IMIX suited for one networking device and deployment will not be appropriate for another. However, the mix of sizes may be known and the tester may be asked to augment the fixed size tests. To address this need, and the perpetual goal of specifying repeatable test conditions, this draft defines a way to specify the exact repeating sequence of packet sizes from the usual set of fixed sizes, and other forms of mixed size specification.

Requirements Language

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 [RFC2119].

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 <http://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 December 5, 2013.

Copyright Notice

Copyright (c) 2013 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Scope and Goals	4
3. Specification of the IMIX Genome	5
4. Specification of a Custom IMIX	7
5. Reporting Long or Pseudo-Random Packet Sequences	8
6. Security Considerations	9
7. IANA Considerations	9
8. Acknowledgements	9
9. References	9
9.1. Normative References	9
9.2. Informative References	10
Author's Address	10

1. Introduction

This memo defines a method to unambiguously specify the sequence of packet sizes used in a load test.

Benchmarking Methodologies [RFC2544] have always relied on test conditions with constant packet sizes, with the goal of understanding what network device capability has been tested. Tests with the smallest size stress the header processing capacity, and tests with the largest size stress the overall bit processing capacity. Tests with sizes in-between may determine the transition between these two capacities.

Streams of constant packet size differ significantly from the conditions encountered in operational deployment, and so additional tests are sometimes conducted with a mixture of packet sizes. The set of sizes used is often called an Internet Mix, or "IMIX" [Spirent], [IXIA], [Agilent].

The mixture of sizes a networking device will encounter is highly variable and depends on many factors. An IMIX suited for one networking device and deployment will not be appropriate for another. However, the mix of sizes may be known and the tester may be asked to augment the fixed size tests. The references above cite the original studies and their methodologies. Similar methods can be used to determine new size mixes present on a link or network. We note that the architecture for IP Flow Information Export [RFC5470] provides one method to gather packet size information on private networks.

To address this need, and the perpetual goal of specifying repeatable test conditions, this memo proposes a way to specify the exact repeating sequence of packet sizes from the usual set of fixed sizes: the IMIX Genome. Other, less exact forms of size specification are also recommended for extremely complicated or customized size mixes. We apply the term "genome" to infer that the entire test packet size sequence can be replicated if this information is known, a parallel to the information needed for biological replication.

This memo takes the position that it cannot be proven for all circumstances that the sequence of packet sizes does not affect the test result, thus a standardized specification of sequence is valuable.

2. Scope and Goals

This memo defines a method to unambiguously specify the sequence of packet sizes that have been used in a load test, assuming that a

relevant mix of sizes is known to the tester and the length of the repeating sequence is not very long (<100 packets).

The IMIX Genome will allow an exact sequence of packet sizes to be communicated as a single-line name, resolving the current ambiguity with results that simply refer to "IMIX". This aspect is critical because no ability has been demonstrated to extrapolate results from one IMIX to another IMIX, even when the mix varies only slightly from another IMIX, and certainly no ability to extrapolate results to other circumstances.

While documentation of the exact sequence is ideal, the memo also covers the case where the sequence of sizes is very long or may be generated by a pseudo-random process.

It is a colossal non-goal to standardize one or more versions of the IMIX. This topic has been discussed on many occasions on the `bmwg-list` [`IMIXonList`]. The goal is to enable customization with minimal constraints while fostering repeatable testing once the fixed size testing is complete. Thus, the requirements presented in this specification, expressed in [`RFC2119`] terms, are intended for those performing/reporting laboratory tests to improve clarity and repeatability.

3. Specification of the IMIX Genome

The IMIX Genome is specified in the following format:

IMIX - 123456...x

where each number is replaced by the letter corresponding to the size of the packet at that position in the sequence. The following table gives the letter encoding for the [`RFC2544`] standard sizes (64, 128, 256, 512, 1024, 1280, and 1518 bytes) and "jumbo" sizes (2112, 9000, 16000). Note that the 4 octet Ethernet frame check sequence may fail to detect bit errors in the larger jumbo frames, see [`jumbo`].

Size, bytes	Genome Code Letter
64	a
128	b
256	c
512	d
1024	e
1280	f
1518	g
2112	h
9000	i
16000	j
MTU	z

For example: a five packet sequence with sizes 64,64,64,1280,1518 would be designated:

IMIX - aaafg

If z (MTU) is used, the tester MUST specify the length of the MTU in the report.

While this approach allows some flexibility, there are also constraints.

- o Non-RFC2544 packet sizes would need to be approximated by those available in the table.
- o The Genome for very long sequences can become undecipherable by humans.

Some questions testers must ask and answer when using the IMIX Genome are:

1. Multiple Source-Destination Address Pairs: is the IMIX sequence applicable to each pair, across multiple pairs in sets, or across all pairs?
2. Multiple Tester Ports: is the IMIX sequence applicable to each port, across multiple ports in sets, or across all ports?

The chosen configuration would be expressed in the following general form:

Source Address + Port AND/OR Blade	Destination Address + Port AND/OR Blade	Corresponding IMIX
x.x.x.x Blade2	y.y.y.y Blade3	IMIX - aaafg

where testers can specify the IMIX used between any two entities in the test architecture (and Blade is a component in a multi-component device chassis).

4. Specification of a Custom IMIX

This section describes how to specify an IMIX with locally-selected packet sizes

The Custom IMIX is specified in the following format:

CUSTOM IMIX - 123456...x

where each number is replaced by the letter corresponding to the size of the packet at that position in the sequence. The tester MUST complete the following table, giving the letter encoding for each size used, where each set of three lower-case letters would be replaced by the integer size in octets.

Size, bytes	Custom Code Letter
aaa	A
bbb	B
ccc	C
ddd	D
eee	E
fff	F
ggg	G
etc.	up to Z

For example: a five packet sequence with sizes
aaa=64,aaa=64,aaa=64,ggg=1020,ggg=1020 would be designated:

CUSTOM IMIX - AAAGG

5. Reporting Long or Pseudo-Random Packet Sequences

When the IMIX-Genome cannot be used (when the sheer length of the sequence would make the Genome unmanageable), two options are possible. When a sequence can be decomposed into a series of short repeating sequences, then a run-length encoding approach MAY be specified as shown in the table below (using the single lower-case letter Genome Codes from section 3):

Count of Repeating Sequences	Packet Size Sequence
20	abcd
5	ggga
10	dcb

The run-length encoding approach is also applicable to custom IMIX described in section 4 (where the single upper-case letter Genome Codes would be used instead).

When the sequence is designed to vary within some proportional constraints, a table simply giving the proportions of each size MAY be used instead.

IP Length	Percentage of Total	Length(s) at other layers
64	23	82
128	67	146
1000	10	1018

Note that the table of proportions also allows non-standard packet sizes, but trades the short Genome specification and ability to specify the exact sequence for other flexibilities.

If a deterministic packet size generation method is used (such as monotonic increase by one octet from start value to MTU), then the generation algorithm SHOULD be reported.

If a pseudo-random length generation capability is used, then the generation algorithm SHOULD be reported with the results along with the seed value used. We also recognize the opportunity to randomize inter-packet spacing from a test sender as well as the size, and both spacing and length pseudo-random generation algorithms and seeds SHOULD be reported when used.

Finally, we note another possibility: a pseudo-random sequence generates an index to the table of packet lengths, and the generation algorithm SHOULD be reported with the results along with the seed value if used.

6. Security Considerations

Benchmarking activities as described in this memo are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the other constraints [RFC2544].

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 DUT/SUT.

Special capabilities SHOULD NOT exist in the DUT/SUT specifically for benchmarking purposes. Any implications for network security arising from the DUT/SUT SHOULD be identical in the lab and in production networks.

7. IANA Considerations

This memo makes no requests of IANA, and hopes that IANA will leave it alone as well.

8. Acknowledgements

Thanks to Sarah Banks, Aamer Akhter, Steve Maxwell, and Scott Bradner for their reviews and comments. Ilya Varlashkin suggested the run-length coding approach in Section 5.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for

Network Interconnect Devices", RFC 2544, March 1999.

9.2. Informative References

- [Agilent] http://www.ixiacom.com/pdfs/test_plans/agilent_journal_of_internet_test_methodologies.pdf, "The Journal of Internet Test Methodologies", 2007.
- [IMIXonList] <http://www.ietf.org/mail-archive/web/bmwg/current/msg00691.html>, "Discussion on IMIX", 2003.
- [IXIA] http://www.ixiacom.com/library/test_plans/display?skey=testing_pppox, "Library: Test Plans", 2010.
- [RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, March 2009.
- [Spirent] <http://gospirent.com/whitepaper/IMIX%20Test%20Methodolgy%20Journal.pdf>, "Test Methodology Journal: IMIX (Internet Mix) Journal", 2006.
- [jumbo] <http://sd.wareonearth.com/~phil/jumbo.html> and <http://staff.psc.edu/mathis/MTU/arguments.html#crc>, "Discussion of Jumbo Packets and FCS Failure".

Author's Address

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown,, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acmorton@att.com
URI: <http://home.comcast.net/~acmacm/>

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 1, 2013

R. Papneja
Huawei Technologies
S. Vapiwala
J. Karthik
Cisco Systems
S. Poretsky
Allot Communications
S. Rao
Qwest Communications
JL. Le Roux
France Telecom
September 28, 2012

Methodology for Benchmarking MPLS-TE Fast Reroute Protection
draft-ietf-bmwg-protection-meth-11.txt

Abstract

This document describes the methodology for benchmarking MPLS Protection mechanisms for link and node protection as defined in [RFC 4090]. This document provides test methodologies and testbed setup for measuring failover times of Fast Reroute techniques while considering all other factors (such as underlying links) that might impact recovery times for real-time applications bound to MPLS traffic engineered (MPLS-TE) tunnels.

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 <http://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 April 1, 2013.

Copyright Notice

Copyright (c) 2012 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	5
2. Document Scope	6
3. Existing Definitions and Requirements	6
4. General Reference Topology	7
5. Test Considerations	8
5.1. Failover Events [RFC 6414]	8
5.2. Failure Detection [RFC 6414]	9
5.3. Use of Data Traffic for MPLS Protection benchmarking	10
5.4. LSP and Route Scaling	10
5.5. Selection of IGP	10
5.6. Restoration and Reversion [RFC 6414]	10
5.7. Offered Load	11
5.8. Tester Capabilities	11
5.9. Failover Time Measurement Methods	12
6. Reference Test Setup	12
6.1. Link Protection	13
6.1.1. Link Protection - 1 hop primary (from PLR) and 1 hop backup TE tunnels	13
6.1.2. Link Protection - 1 hop primary (from PLR) and 2 hop backup TE tunnels	14
6.1.3. Link Protection - 2+ hop (from PLR) primary and 1 hop backup TE tunnels	14
6.1.4. Link Protection - 2+ hop (from PLR) primary and 2 hop backup TE tunnels	15
6.2. Node Protection	16
6.2.1. Node Protection - 2 hop primary (from PLR) and 1 hop backup TE tunnels	16
6.2.2. Node Protection - 2 hop primary (from PLR) and 2 hop backup TE tunnels	17
6.2.3. Node Protection - 3+ hop primary (from PLR) and 1 hop backup TE tunnels	18
6.2.4. Node Protection - 3+ hop primary (from PLR) and 2 hop backup TE tunnels	19
7. Test Methodology	20
7.1. MPLS FRR Forwarding Performance	20
7.1.1. Headend PLR Forwarding Performance	20
7.1.2. Mid-Point PLR Forwarding Performance	21
7.2. Headend PLR with Link Failure	23
7.3. Mid-Point PLR with Link Failure	24
7.4. Headend PLR with Node Failure	26
7.5. Mid-Point PLR with Node Failure	27
8. Reporting Format	28
9. Security Considerations	30
10. IANA Considerations	30
11. Acknowledgements	30
12. References	30

12.1. Informative References	30
12.2. Normative References	31
Appendix A. Fast Reroute Scalability Table	31
Appendix B. Abbreviations	34
Authors' Addresses	35

1. Introduction

This document describes the methodology for benchmarking MPLS Fast Reroute (FRR) protection mechanisms. This document uses much of the terminology defined in [RFC 6414]. For any conflicting content, this document supersedes [RFC 6414]

Protection mechanisms provide recovery of client services from a planned or an unplanned link or node failures. MPLS FRR protection mechanisms are generally deployed in a network infrastructure where MPLS is used for provisioning of point-to-point traffic engineered tunnels (tunnel). MPLS FRR protection mechanisms aim to reduce service disruption period by minimizing recovery time from most common failures.

Network elements from different manufacturers behave differently to network failures, which impacts the network's ability and performance for failure recovery. It therefore becomes imperative for service providers to have a common benchmark to understand the performance behaviors of network elements.

There are two factors impacting service availability: frequency of failures and duration for which the failures persist. Failures can be classified further into two types: correlated and uncorrelated. Correlated and uncorrelated failures may be planned or unplanned.

Planned failures are predictable. Network implementations should be able to handle both planned and unplanned failures and recover gracefully within a time frame to maintain service assurance. Hence, failover recovery time is one of the most important benchmark that a service provider considers in choosing the building blocks for their network infrastructure.

A correlated failure is the simultaneous occurrence of two or more failures. A typical example is failure of a logical resource (e.g. layer-2 links) due to a dependency on a common physical resource (e.g. common conduit) that fails. Within the context of MPLS protection mechanisms, failures that arise due to Shared Risk Link Groups (SRLG) [RFC 4090] can be considered as correlated failures.

MPLS FRR [RFC 4090] allows for the possibility that the Label Switched Paths can be re-optimized in the minutes following Failover. IP Traffic would be re-routed according to the preferred path for the post-failure topology. Thus, MPLS-FRR may include additional steps following the occurrence of the failure detection [RFC 6414] and failover event [RFC 6414].

- (1) Failover Event - Primary Path (Working Path) fails
- (2) Failure Detection- Failover Event is detected
- (3)
 - a. Failover - Working Path switched to Backup path
 - b. Re-Optimization of Working Path (possible change from Backup Path)
- (4) Restoration [RFC 6414]
- (5) Reversion [RFC 6414]

2. Document Scope

This document provides detailed test cases along with different topologies and scenarios that should be considered to effectively benchmark MPLS FRR protection mechanisms and failover times on the Data Plane. Different Failover Events and scaling considerations are also provided in this document.

All benchmarking test-cases defined in this document apply to Facility backup [RFC 4090]. The test cases cover all possible failure scenarios and the associated procedures benchmark the performance of the Device Under Test (DUT) to recover from failures. Data plane traffic is used to benchmark failover times.

Benchmarking of correlated failures is out of scope of this document. Detection using Bi-directional Forwarding Detection (BFD) is outside the scope of this document, but mentioned in discussion sections.

The Performance of control plane is outside the scope of this benchmarking.

As described above, MPLS-FRR may include a Re-optimization of the Working Path, with possible packet transfer impairments. Characterization of Re-optimization is beyond the scope of this memo.

3. Existing Definitions and 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 BCP 14, RFC 2119 [Br97]. 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 document.

The reader is assumed to be familiar with the commonly used MPLS terminology, some of which is defined in [MPLS-FRR-EXT].

This document uses much of the terminology defined in [RFC 6414]. This document also uses existing terminology defined in other BMWG Work [Br91], [Ma98], [Po06].

4. General Reference Topology

Figure 1 illustrates the basic reference testbed and is applicable to all the test cases defined in this document. The Tester is comprised of a Traffic Generator (TG) & Test Analyzer (TA) and Emulator. A Tester is connected to the test network and depending upon the test case, the DUT could vary. The Tester sends and receives IP traffic to the tunnel ingress and performs signaling protocol emulation to simulate real network scenarios in a lab environment. The Tester may also support MPLS-TE signaling to act as the ingress node to the MPLS tunnel.




Fig. 1 Fast Reroute Topology

The tester MUST record the number of lost, duplicate, and reordered packets. It should further record arrival and departure times so that Failover Time, Additive Latency, and Reversion Time can be measured. The tester may be a single device or a test system emulating all the different roles along a primary or backup path.

The label stack is dependent of the following 3 entities:

- (1) Type of protection (Link Vs Node)
- (2) # of remaining hops of the primary tunnel from the PLR[RFC 6414]
- (3) # of remaining hops of the backup tunnel from the PLR

Due to this dependency, it is RECOMMENDED that the benchmarking of failover times be performed on all the topologies provided in section 6.

5. Test Considerations

This section discusses the fundamentals of MPLS Protection testing:

- (1) The types of network events that causes failover
- (2) Indications for failover
- (3) the use of data traffic
- (4) Traffic generation
- (5) LSP Scaling
- (6) Reversion of LSP
- (7) IGP Selection

5.1. Failover Events [RFC 6414]

The failover to the backup tunnel is primarily triggered by either link or node failures observed downstream of the Point of Local repair (PLR). Some of these failure events are listed below.

Link Failure Events

- Interface Shutdown on PLR side with POS Alarm
- Interface Shutdown on remote side with POS Alarm
- Interface Shutdown on PLR side with RSVP hello enabled
- Interface Shutdown on remote side with RSVP hello enabled
- Interface Shutdown on PLR side with BFD
- Interface Shutdown on remote side with BFD
- Fiber Pull on the PLR side (Both TX & RX or just the TX)
- Fiber Pull on the remote side (Both TX & RX or just the RX)
- Online insertion and removal (OIR) on PLR side
- OIR on remote side
- Sub-interface failure on PLR side (e.g. shutting down of a VLAN)
- Sub-interface failure on remote side
- Parent interface shutdown on PLR side (an interface bearing multiple sub-interfaces)
- Parent interface shutdown on remote side

Node Failure Events

- A System reload initiated either by a graceful shutdown or by a power failure.
- A system crash due to a software failure or an assert.

5.2. Failure Detection [RFC 6414]

Link failure detection time depends on the link type and failure detection protocols running. For SONET/SDH, the alarm type (such as LOS, AIS, or RDI) can be used. Other link types have layer-two alarms, but they may not provide a short enough failure detection time. Ethernet based links do not have layer 2 failure indicators, and therefore relies on layer 3 signaling for failure detection. However for directly connected devices, remote fault indication in the ethernet auto-negotiation scheme could be considered as a type of layer 2 link failure indicator.

MPLS has different failure detection techniques such as BFD, or use of RSVP hellos. These methods can be used for the layer 3 failure indicators required by Ethernet based links, or for some other non-Ethernet based links to help improve failure detection time. However, these fast failure detection mechanisms are out of scope

The test procedures in this document can be used for a local failure or remote failure scenarios for comprehensive benchmarking and to evaluate failover performance independent of the failure detection techniques.

5.3. Use of Data Traffic for MPLS Protection benchmarking

Currently end customers use packet loss as a key metric for Failover Time [RFC 6414]. Failover Packet Loss [RFC 6414] is an externally observable event and has direct impact on application performance. MPLS protection is expected to minimize the packet loss in the event of a failure. For this reason it is important to develop a standard router benchmarking methodology for measuring MPLS protection that uses packet loss as a metric. At a known rate of forwarding, packet loss can be measured and the failover time can be determined. Measurement of control plane signaling to establish backup paths is not enough to verify failover. Failover is best determined when packets are actually traversing the backup path.

An additional benefit of using packet loss for calculation of failover time is that it allows use of a black-box test environment. Data traffic is offered at line-rate to the device under test (DUT) an emulated network failure event is forced to occur, and packet loss is externally measured to calculate the convergence time. This setup is independent of the DUT architecture.

In addition, this methodology considers the packets in error and duplicate packets [Po06] that could have been generated during the failover process. The methodologies consider lost, out-of-order [Po06] and duplicate packets to be impaired packets that contribute to the Failover Time.

5.4. LSP and Route Scaling

Failover time performance may vary with the number of established primary and backup tunnel label switched paths (LSP) and installed routes. However the procedure outlined here should be used for any number of LSPs (L) and number of routes protected by PLR(R). The amount of L and R must be recorded.

5.5. Selection of IGP

The underlying IGP could be ISIS-TE or OSPF-TE for the methodology proposed here. See [RFC 6412] for IGP options to consider and report.

5.6. Restoration and Reversion [RFC 6414]

Path restoration provides a method to restore an alternate primary LSP upon failure and to switch traffic from the Backup Path to the restored Primary Path (Reversion). In MPLS-FRR, Reversion can be implemented as Global Reversion or Local Reversion. It is important to include Restoration and Reversion as a step in each test case to

measure the amount of packet loss, out of order packets, or duplicate packets that is produced.

Note: In addition to restoration and reversion, re-optimization can take place while the failure is still not recovered but it depends on the user configuration, and re-optimization timers.

5.7. Offered Load

It is suggested that there be three or more traffic streams as long as there is a steady and constant rate of flow for all the streams. In order to monitor the DUT performance for recovery times, a set of route prefixes should be advertised before traffic is sent. The traffic should be configured towards these routes.

At least 16 flows should be used, and more if possible. Prefix-dependency behaviors are key in IP and tests with route-specific flows spread across the routing table will reveal this dependency. Generating traffic to all of the prefixes reachable by the protected tunnel (probably in a Round-Robin fashion, where the traffic is destined to all the prefixes but one prefix at a time in a cyclic manner) is not recommended. The reason why traffic generation is not recommended in a Round-Robin fashion to all the prefixes, one at a time is that if there are many prefixes reachable through the LSP the time interval between 2 packets destined to one prefix may be significantly high and may be comparable with the failover time being measured which does not aid in getting an accurate failover measurement.

5.8. Tester Capabilities

It is RECOMMENDED that the Tester used to execute each test case have the following capabilities:

- 1.Ability to establish MPLS-TE tunnels and push/pop labels.
- 2.Ability to produce Failover Event [RFC 6414].
- 3.Ability to insert a timestamp in each data packet's IP payload.
- 4.An internal time clock to control timestamping, time measurements, and time calculations.
- 5.Ability to disable or tune specific Layer-2 and Layer-3 protocol functions on any interface(s).

6.Ability to react upon the receipt of path error from the PLR

The Tester MAY be capable to make non-data plane convergence observations and use those observations for measurements.

5.9. Failover Time Measurement Methods

Failover Time is calculated using one of the following three methods

1. Packet-Loss Based method (PLBM): (Number of packets dropped/ packets per second * 1000) milliseconds. This method could also be referred as Loss-Derived method.
2. Time-Based Loss Method (TBLM): This method relies on the ability of the Traffic generators to provide statistics which reveal the duration of failure in milliseconds based on when the packet loss occurred (interval between non-zero packet loss and zero loss).
3. Timestamp Based Method (TBM): This method of failover calculation is based on the timestamp that gets transmitted as payload in the packets originated by the generator. The Traffic Analyzer records the timestamp of the last packet received before the failover event and the first packet after the failover and derives the time based on the difference between these 2 timestamps. Note: The payload could also contain sequence numbers for out-of-order packet calculation and duplicate packets.

The timestamp based method method would be able to detect Reversion impairments beyond loss, thus it is RECOMMENDED method as a Failover Time method.

6. Reference Test Setup

In addition to the general reference topology shown in figure 1, this section provides detailed insight into various proposed test setups that should be considered for comprehensively benchmarking the failover time in different roles along the primary tunnel

This section proposes a set of topologies that covers all the scenarios for local protection. All of these topologies can be mapped to the reference topology shown in Figure 1. Topologies provided in this section refer to the testbed required to benchmark failover time when the DUT is configured as a PLR in either Headend or midpoint role. Provided with each topology below is the label stack at the PLR. Penultimate Hop Popping (PHP) MAY be used and must be reported when used.

Figures 2 thru 9 use the following convention and are subset of figure 1:

- a) HE is Headend
- b) TE is Tail-End
- c) MID is Mid point
- d) MP is Merge Point
- e) PLR is Point of Local Repair
- f) PRI is Primary Path
- g) BKP denotes Backup Path and Nodes
- h) UR is Upstream Router

6.1. Link Protection

6.1.1. Link Protection - 1 hop primary (from PLR) and 1 hop backup TE tunnels




Figure 2.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	0	0
Layer3 VPN (PE-PE)	1	1
Layer3 VPN (PE-P)	2	2
Layer2 VC (PE-PE)	1	1
Layer2 VC (PE-P)	2	2
Mid-point LSPs	0	0

Note: Please note the following:

- a) For P-P case, R2 and R3 acts as P routers
- b) For PE-PE case, R2 acts as PE and R3 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R3 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2 and R3 act as shown in above figure HE, Mid point/PLR and TE respectively

6.1.2. Link Protection - 1 hop primary (from PLR) and 2 hop backup TE tunnels




Figure 3.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	0	1
Layer3 VPN (PE-PE)	1	2
Layer3 VPN (PE-P)	2	3
Layer2 VC (PE-PE)	1	2
Layer2 VC (PE-P)	2	3
Mid-point LSPs	0	1

Note: Please note the following:

- a) For P-P case, R2 and R3 acts as P routers
- b) For PE-PE case, R2 acts as PE and R3 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R3 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2 and R3 act as shown in above figure HE, Mid point/PLR and TE respectively

6.1.3. Link Protection - 2+ hop (from PLR) primary and 1 hop backup TE tunnels




Figure 4.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	1	1
Layer3 VPN (PE-PE)	2	2
Layer3 VPN (PE-P)	3	3
Layer2 VC (PE-PE)	2	2
Layer2 VC (PE-P)	3	3
Mid-point LSPs	1	1

Note: Please note the following:

- a) For P-P case, R2, R3 and R4 acts as P routers
- b) For PE-PE case, R2 acts as PE and R4 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R3 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2, R3 and R4 act as shown in above figure HE, Midpoint/PLR and TE respectively

6.1.4. Link Protection - 2+ hop (from PLR) primary and 2 hop backup TE tunnels

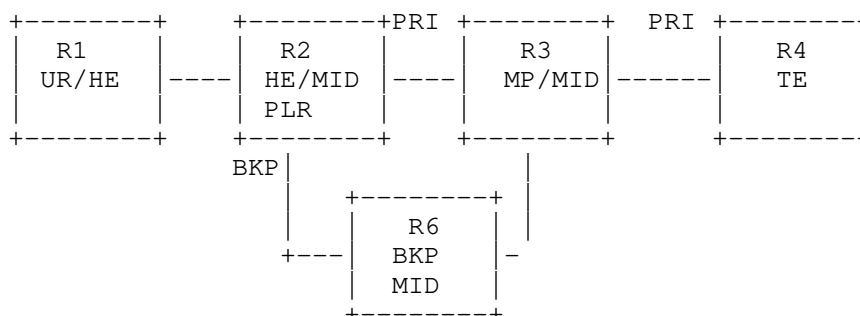


Figure 5.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	1	2
Layer3 VPN (PE-PE)	2	3
Layer3 VPN (PE-P)	3	4
Layer2 VC (PE-PE)	2	3
Layer2 VC (PE-P)	3	4
Mid-point LSPs	1	2

Note: Please note the following:

- a) For P-P case, R2, R3 and R4 acts as P routers
- b) For PE-PE case, R2 acts as PE and R4 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R3 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2, R3 and R4 act as shown in above figure HE, Midpoint/PLR and TE respectively

6.2. Node Protection

6.2.1. Node Protection - 2 hop primary (from PLR) and 1 hop backup TE tunnels




Figure 6.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	1	0
Layer3 VPN (PE-PE)	2	1
Layer3 VPN (PE-P)	3	2
Layer2 VC (PE-PE)	2	1
Layer2 VC (PE-P)	3	2
Mid-point LSPs	1	0

Note: Please note the following:

- a) For P-P case, R2, R3 and R3 acts as P routers
- b) For PE-PE case, R2 acts as PE and R4 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R4 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2, R3 and R4 act as shown in above figure HE, Midpoint/PLR and TE respectively

6.2.2. Node Protection - 2 hop primary (from PLR) and 2 hop backup TE tunnels

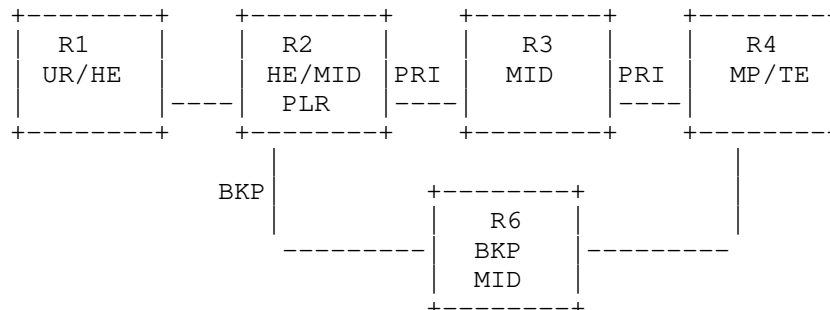


Figure 7.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	1	1
Layer3 VPN (PE-PE)	2	2
Layer3 VPN (PE-P)	3	3
Layer2 VC (PE-PE)	2	2
Layer2 VC (PE-P)	3	3
Mid-point LSPs	1	1

Note: Please note the following:

- a) For P-P case, R2, R3 and R4 acts as P routers
- b) For PE-PE case, R2 acts as PE and R4 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R4 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2, R3 and R4 act as shown in above figure HE, Midpoint/PLR and TE respectively

6.2.3. Node Protection - 3+ hop primary (from PLR) and 1 hop backup TE tunnels




Figure 8.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	1	1
Layer3 VPN (PE-PE)	2	2
Layer3 VPN (PE-P)	3	3
Layer2 VC (PE-PE)	2	2
Layer2 VC (PE-P)	3	3
Mid-point LSPs	1	1

Note: Please note the following:

- a) For P-P case, R2, R3, R4 and R5 acts as P routers
- b) For PE-PE case, R2 acts as PE and R5 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R4 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2, R3, R4 and R5 act as shown in above figure HE, Midpoint/PLR and TE respectively

6.2.4. Node Protection - 3+ hop primary (from PLR) and 2 hop backup TE tunnels

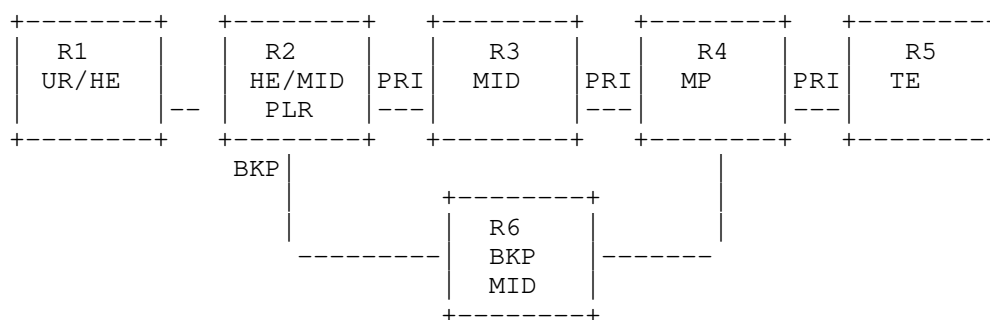


Figure 9.

Traffic	Num of Labels before failure	Num of labels after failure
IP TRAFFIC (P-P)	1	2
Layer3 VPN (PE-PE)	2	3
Layer3 VPN (PE-P)	3	4
Layer2 VC (PE-PE)	2	3
Layer2 VC (PE-P)	3	4
Mid-point LSPs	1	2

Note: Please note the following:

- a) For P-P case, R2, R3, R4 and R5 acts as P routers
- b) For PE-PE case, R2 acts as PE and R5 acts as a remote PE
- c) For PE-P case, R2 acts as a PE router, R4 acts as a P router and R5 acts as remote PE router (Please refer to figure 1 for complete setup)
- d) For Mid-point case, R1, R2, R3, R4 and R5 act as shown in above figure HE, Midpoint/PLR and TE respectively

7. Test Methodology

The procedure described in this section can be applied to all the 8 base test cases and the associated topologies. The backup as well as the primary tunnels are configured to be alike in terms of bandwidth usage. In order to benchmark failover with all possible label stack depth applicable as seen with current deployments, it is RECOMMENDED to perform all of the test cases provided in this section. The forwarding performance test cases in section 7.1 MUST be performed prior to performing the failover test cases.

The considerations of Section 4 of [RFC 2544] are applicable when evaluating the results obtained using these methodologies as well.

7.1. MPLS FRR Forwarding Performance

Benchmarking Failover Time [RFC 6414] for MPLS protection first requires baseline measurement of the forwarding performance of the test topology including the DUT. Forwarding performance is benchmarked by the Throughput as defined in [MPLS-FWD] and measured in units pps. This section provides two test cases to benchmark forwarding performance. These are with the DUT configured as a Headend PLR, Mid-Point PLR, and Egress PLR.

7.1.1. Headend PLR Forwarding Performance

Objective:

To benchmark the maximum rate (pps) on the PLR (as headend) over primary LSP and backup LSP.

Test Setup:

- A. Select any one topology out of the 8 from section 6.
- B. Select or enable IP, Layer 3 VPN or Layer 2 VPN services with DUT as Headend PLR.
- C. The DUT will also have 2 interfaces connected to the traffic Generator/analyzer. (If the node downstream of the PLR is not a simulated node, then the Ingress of the tunnel should have one link connected to the traffic generator and the node downstream to the PLR or the egress of the tunnel should have a link connected to the traffic analyzer).

Procedure:

1. Establish the primary LSP on R2 required by the topology selected.
2. Establish the backup LSP on R2 required by the selected topology.
3. Verify primary and backup LSPs are up and that primary is protected.
4. Verify Fast Reroute protection is enabled and ready.
5. Setup traffic streams as described in section 5.7.
6. Send MPLS traffic over the primary LSP at the Throughput supported by the DUT (section 6, RFC 2544).
7. Record the Throughput over the primary LSP.
8. Trigger a link failure as described in section 5.1.
9. Verify that the offered load gets mapped to the backup tunnel and measure the Additive Backup Delay (RFC 6414).
10. 30 seconds after Failover, stop the offered load and measure the Throughput, Packet Loss, Out-of-Order Packets, and Duplicate Packets over the Backup LSP.
11. Adjust the offered load and repeat steps 6 through 10 until the Throughput values for the primary and backup LSPs are equal.
12. Record the final Throughput, which corresponds to the offered load that will be used for the Headend PLR failover test cases.

7.1.2. Mid-Point PLR Forwarding Performance

Objective:

To benchmark the maximum rate (pps) on the PLR (as mid-point) over primary LSP and backup LSP.

Test Setup:

- A. Select any one topology out of the 8 from section 6.
- B. The DUT will also have 2 interfaces connected to the traffic generator.

Procedure:

1. Establish the primary LSP on R1 required by the topology selected.
2. Establish the backup LSP on R2 required by the selected topology.
3. Verify primary and backup LSPs are up and that primary is protected.
4. Verify Fast Reroute protection is enabled and ready.
5. Setup traffic streams as described in section 5.7.
6. Send MPLS traffic over the primary LSP at the Throughput supported by the DUT (section 6, RFC 2544).
7. Record the Throughput over the primary LSP.
8. Trigger a link failure as described in section 5.1.
9. Verify that the offered load gets mapped to the backup tunnel and measure the Additive Backup Delay (RFC 6414).
10. 30 seconds after Failover, stop the offered load and measure the Throughput, Packet Loss, Out-of-Order Packets, and Duplicate Packets over the Backup LSP.
11. Adjust the offered load and repeat steps 6 through 10 until the Throughput values for the primary and backup LSPs are equal.
12. Record the final Throughput which corresponds to the offered load that will be used for the Mid-Point PLR failover test cases.

7.2. Headend PLR with Link Failure

Objective:

To benchmark the MPLS failover time due to link failure events described in section 5.1 experienced by the DUT which is the Headend PLR.

Test Setup:

- A. Select any one topology out of the 8 from section 6.
- B. Select or enable IP, Layer 3 VPN or Layer 2 VPN services with DUT as Headend PLR.
- C. The DUT will also have 2 interfaces connected to the traffic Generator/analyzer. (If the node downstream of the PLR is not a simulated node, then the Ingress of the tunnel should have one link connected to the traffic generator and the node downstream to the PLR or the egress of the tunnel should have a link connected to the traffic analyzer).

Test Configuration:

1. Configure the number of primaries on R2 and the backups on R2 as required by the topology selected.
2. Configure the test setup to support Reversion.
3. Advertise prefixes (as per FRR Scalability Table described in Appendix A) by the tail end.

Procedure:

Test Case "7.1.1. Headend PLR Forwarding Performance" MUST be completed first to obtain the Throughput to use as the offered load.

1. Establish the primary LSP on R2 required by the topology selected.

2. Establish the backup LSP on R2 required by the selected topology.
3. Verify primary and backup LSPs are up and that primary is protected.
4. Verify Fast Reroute protection is enabled and ready.
5. Setup traffic streams for the offered load as described in section 5.7.
6. Provide the offered load from the tester at the Throughput [Br91] level obtained from test case 7.1.1.
7. Verify traffic is switched over Primary LSP without packet loss.
8. Trigger a link failure as described in section 5.1.
9. Verify that the offered load gets mapped to the backup tunnel and measure the Additive Backup Delay.
10. 30 seconds after Failover [RFC 6414], stop the offered load and measure the total Failover Packet Loss [RFC 6414].
11. Calculate the Failover Time [RFC 6414] benchmark using the selected Failover Time Calculation Method (TBLM, PLBM, or TBM) [RFC 6414].
12. Restart the offered load and restore the primary LSP to verify Reversion [RFC 6414] occurs and measure the Reversion Packet Loss [RFC 6414].
13. Calculate the Reversion Time [RFC 6414] benchmark using the selected Failover Time Calculation Method (TBLM, PLBM, or TBM) [RFC 6414].
14. Verify Headend signals new LSP and protection should be in place again.

IT is RECOMMENDED that this procedure be repeated for each of the link failure triggers defined in section 5.1.

7.3. Mid-Point PLR with Link Failure

Objective:

To benchmark the MPLS failover time due to link failure events described in section 5.1 experienced by the DUT which is the Mid-Point PLR.

Test Setup:

- A. Select any one topology out of the 8 from section 6.
- B. The DUT will also have 2 interfaces connected to the traffic generator.

Test Configuration:

1. Configure the number of primaries on R1 and the backups on R2 as required by the topology selected.
2. Configure the test setup to support Reversion.
3. Advertise prefixes (as per FRR Scalability Table described in Appendix A) by the tail end.

Procedure:

Test Case "7.1.2. Mid-Point PLR Forwarding Performance" MUST be completed first to obtain the Throughput to use as the offered load.

1. Establish the primary LSP on R1 required by the topology selected.
2. Establish the backup LSP on R2 required by the selected topology.
3. Perform steps 3 through 14 from section 7.2 Headend PLR with Link Failure.

IT is RECOMMENDED that this procedure be repeated for each of the link failure triggers defined in section 5.1.

7.4. Headend PLR with Node Failure

Objective:

To benchmark the MPLS failover time due to Node failure events described in section 5.1 experienced by the DUT which is the Headend PLR.

Test Setup:

- A. Select any one topology out of the 8 from section 6.
- B. Select or enable IP, Layer 3 VPN or Layer 2 VPN services with DUT as Headend PLR.
- C. The DUT will also have 2 interfaces connected to the traffic generator/analyzer.

Test Configuration:

1. Configure the number of primaries on R2 and the backups on R2 as required by the topology selected.
2. Configure the test setup to support Reversion.
3. Advertise prefixes (as per FRR Scalability Table described in Appendix A) by the tail end.

Procedure:

Test Case "7.1.1. Headend PLR Forwarding Performance" MUST be completed first to obtain the Throughput to use as the offered load.

1. Establish the primary LSP on R2 required by the topology selected.
2. Establish the backup LSP on R2 required by the selected topology.
3. Verify primary and backup LSPs are up and that primary is protected.

4. Verify Fast Reroute protection is enabled and ready.
5. Setup traffic streams for the offered load as described in section 5.7.
6. Provide the offered load from the tester at the Throughput [Br91] level obtained from test case 7.1.1.
7. Verify traffic is switched over Primary LSP without packet loss.
8. Trigger a node failure as described in section 5.1.
9. Perform steps 9 through 14 in 7.2 Headend PLR with Link Failure.

IT is RECOMMENDED that this procedure be repeated for each of the node failure triggers defined in section 5.1.

7.5. Mid-Point PLR with Node Failure

Objective:

To benchmark the MPLS failover time due to Node failure events described in section 5.1 experienced by the DUT which is the Mid-Point PLR.

Test Setup:

- A. Select any one topology from section 6.1 to 6.2.
- B. The DUT will also have 2 interfaces connected to the traffic generator.

Test Configuration:

1. Configure the number of primaries on R1 and the backups on R2 as required by the topology selected.
2. Configure the test setup to support Reversion.
3. Advertise prefixes (as per FRR Scalability Table described in Appendix A) by the tail end.

Procedure:

Test Case "7.1.1. Mid-Point PLR Forwarding Performance" MUST be completed first to obtain the Throughput to use as the offered load.

1. Establish the primary LSP on R1 required by the topology selected.
2. Establish the backup LSP on R2 required by the selected topology.
3. Verify primary and backup LSPs are up and that primary is protected.
4. Verify Fast Reroute protection is enabled and ready.
5. Setup traffic streams for the offered load as described in section 5.7.
6. Provide the offered load from the tester at the Throughput [Br91] level obtained from test case 7.1.1.
7. Verify traffic is switched over Primary LSP without packet loss.
8. Trigger a node failure as described in section 5.1.
9. Perform steps 9 through 14 in 7.2 Headend PLR with Link Failure.

IT is RECOMMENDED that this procedure be repeated for each of the node failure triggers defined in section 5.1.

8. Reporting Format

For each test, it is RECOMMENDED that the results be reported in the following format.

Parameter	Units
IGP used for the test	ISIS-TE/ OSPF-TE

Interface types	Gige, POS, ATM, VLAN etc.
Packet Sizes offered to the DUT	Bytes (at layer 3)
Offered Load (Throughput)	packets per second
IGP routes advertised	Number of IGP routes
Penultimate Hop Popping	Used/Not Used
RSVP hello timers	Milliseconds
Number of Protected tunnels	Number of tunnels
Number of VPN routes installed on the Headend	Number of VPN routes
Number of VC tunnels	Number of VC tunnels
Number of mid-point tunnels	Number of tunnels
Number of Prefixes protected by Primary	Number of LSPs
Topology being used	Section number, and figure reference
Failover Event	Event type
Re-optimization	Yes/No

Benchmarks (to be recorded for each test case):

Failover-

Failover Time	seconds
Failover Packet Loss	packets
Additive Backup Delay	seconds
Out-of-Order Packets	packets
Duplicate Packets	packets
Failover Time Calculation Method	Method Used

Reversion-

Reversion Time	seconds
Reversion Packet Loss	packets
Additive Backup Delay	seconds
Out-of-Order Packets	packets
Duplicate Packets	packets
Failover Time Calculation Method	Method Used

9. Security Considerations

Benchmarking activities as described in this memo are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the constraints specified in the sections above.

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 DUT/SUT.

Special capabilities SHOULD NOT exist in the DUT/SUT specifically for benchmarking purposes. Any implications for network security arising from the DUT/SUT SHOULD be identical in the lab and in production networks.

10. IANA Considerations

This draft does not require any new allocations by IANA.

11. Acknowledgements

We would like to thank Jean Philip Vasseur for his invaluable input to the document, Curtis Villamizar for his contribution in suggesting text on definition and need for benchmarking Correlated failures and Bhavani Parise for his textual input and review. Additionally we would like to thank Al Morton, Arun Gandhi, Amrit Hanspal, Karu Ratnam, Raveesh Janardan, Andrey Kiselev, and Mohan Nanduri for their formal reviews of this document.

12. References

12.1. Informative References

- [RFC2285] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.
- [RFC4689] Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms", RFC 4689, October 2006.

12.2. Normative References

- [RFC1242] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [RFC5695] Akhter, A., Asati, R., and C. Pignataro, "MPLS Forwarding Benchmarking Methodology for IP Flows", RFC 5695, November 2009.

Appendix A. Fast Reroute Scalability Table

This section provides the recommended numbers for evaluating the scalability of fast reroute implementations. It also recommends the typical numbers for IGP/VPNv4 Prefixes, LSP Tunnels and VC entries. Based on the features supported by the device under test (DUT), appropriate scaling limits can be used for the test bed.

A1. FRR IGP Table

No. of Headend TE Tunnels	IGP Prefixes
1	100
1	500
1	1000
1	2000
1	5000
2 (Load Balance)	100
2 (Load Balance)	500
2 (Load Balance)	1000
2 (Load Balance)	2000
2 (Load Balance)	5000
100	100
500	500
1000	1000
2000	2000

A2. FRR VPN Table

No. of Headend TE Tunnels	VPNv4 Prefixes
1	100
1	500
1	1000
1	2000
1	5000
1	10000
1	20000
1	Max
2 (Load Balance)	100
2 (Load Balance)	500
2 (Load Balance)	1000
2 (Load Balance)	2000
2 (Load Balance)	5000
2 (Load Balance)	10000
2 (Load Balance)	20000
2 (Load Balance)	Max

A3. FRR Mid-Point LSP Table

No of Mid-point TE LSPs could be configured at recommended levels - 100, 500, 1000, 2000, or max supported number.

A2. FRR VC Table

No. of Headend TE Tunnels	VC entries
1	100
1	500
1	1000
1	2000
1	Max
100	100
500	500
1000	1000
2000	2000

Appendix B. Abbreviations

AIS	- Alarm Indication Signal
BFD	- Bidirectional Fault Detection
BGP	- Border Gateway protocol
CE	- Customer Edge
DUT	- Device Under Test
FRR	- Fast Reroute
IGP	- Interior Gateway Protocol
IP	- Internet Protocol
LOS	- Loss of Signal
LSP	- Label Switched Path
MP	- Merge Point
MPLS	- Multi Protocol Label Switching
N-Nhop	- Next - Next Hop
Nhop	- Next Hop
OIR	- Online Insertion and Removal
P	- Provider
PE	- Provider Edge
PHP	- Penultimate Hop Popping
PLR	- Point of Local Repair
RSVP	- Resource reSerVation Protocol
SRLG	- Shared Risk Link Group
TA	- Traffic Analyzer
TE	- Traffic Engineering
TG	- Traffic Generator
VC	- Virtual Circuit
VPN	- Virtual Private Network

Authors' Addresses

Rajiv Papneja
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: rajiv.papneja@huawei.com

Samir Vapiwala
Cisco Systems
300 Beaver Brook Road
Boxborough, MA 01719
USA

Email: svapiwal@cisco.com

Jay Karthik
Cisco Systems
300 Beaver Brook Road
Boxborough, MA 01719
USA

Email: jkarthik@cisco.com

Scott Poretsky
Allot Communications
USA

Email: sporetsky@allot.com

Shankar Rao
950 17th Street
Suite 1900
Denver, CO 80210
USA

Email: shankar.rao@du.edu

JL. Le Roux
France Telecom
2 av Pierre Marzin
22300 Lannion
France

Email: jeanlouis.leroux@orange.com

Benchmarking Methodology Working Group
Internet-Draft
Expires: April 25, 2013

C. Davids
Illinois Institute of Technology
V. Gurbani
Bell Laboratories,
Alcatel-Lucent
S. Poretsky
Allot Communications
October 22, 2012

Methodology for Benchmarking SIP Networking Devices
draft-ietf-bmwg-sip-bench-meth-05

Abstract

This document describes the methodology for benchmarking Session Initiation Protocol (SIP) performance as described in SIP benchmarking terminology document. The methodology and terminology are to be used for benchmarking signaling plane performance with varying signaling and media load. Both scale and establishment rate are measured by signaling plane performance. The SIP Devices to be benchmarked may be a single device under test (DUT) or a system under test (SUT). Benchmarks can be obtained and compared for different types of devices such as SIP Proxy Server, SBC, and server paired with a media relay or Firewall/NAT device.

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 <http://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 April 25, 2013.

Copyright Notice

Copyright (c) 2012 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Terminology	4
2. Introduction	4
3. Benchmarking Topologies	5
4. Test Setup Parameters	5
4.1. Selection of SIP Transport Protocol	5
4.2. Signaling Server	5
4.3. Associated Media	5
4.4. Selection of Associated Media Protocol	6
4.5. Number of Associated Media Streams per SIP Session	6
4.6. Session Duration	6
4.7. Attempted Sessions per Second	6
4.8. Stress Testing	6
4.9. Benchmarking algorithm	6
5. Reporting Format	9
5.1. Test Setup Report	9
5.2. Device Benchmarks for IS	10
5.3. Device Benchmarks for NS	10
6. Test Cases	10
6.1. Baseline Session Establishment Rate of the test bed	10
6.2. Session Establishment Rate without media	11
6.3. Session Establishment Rate with Media on DUT/SUT	11
6.4. Session Establishment Rate with Media not on DUT/SUT	12
6.5. Session Establishment Rate with Loop Detection Enabled	13
6.6. Session Establishment Rate with Forking	13
6.7. Session Establishment Rate with Forking and Loop Detection	14
6.8. Session Establishment Rate with TLS Encrypted SIP	14
6.9. Session Establishment Rate with IPsec Encrypted SIP	15
6.10. Session Establishment Rate with SIP Flooding	15
6.11. Maximum Registration Rate	16
6.12. Maximum Re-Registration Rate	16
6.13. Maximum IM Rate	17
6.14. Session Capacity without Media	17
6.15. Session Capacity with Media	18
6.16. Session Capacity with Media and a Media Relay/NAT and/or Firewall	18
7. IANA Considerations	19
8. Security Considerations	19
9. Acknowledgments	19
10. References	19
10.1. Normative References	19
10.2. Informative References	20
Authors' Addresses	20

1. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, conforming to [RFC2119] and indicate requirement levels for compliant implementations.

Terms specific to SIP [RFC3261] performance benchmarking are defined in [I-D.sip-bench-term].

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 document. The term Throughput is defined in [RFC2544].

2. Introduction

This document describes the methodology for benchmarking Session Initiation Protocol (SIP) performance as described in Terminology document [I-D.sip-bench-term]. The methodology and terminology are to be used for benchmarking signaling plane performance with varying signaling and media load. Both scale and establishment rate are measured by signaling plane performance.

The SIP Devices to be benchmarked may be a single device under test (DUT) or a system under test (SUT). The DUT is a SIP Server, which may be any [RFC3261] conforming device. The SUT can be any device or group of devices containing RFC 3261 conforming functionality along with Firewall and/or NAT functionality. This enables benchmarks to be obtained and compared for different types of devices such as SIP Proxy Server, SBC, SIP proxy server paired with a media relay or Firewall/NAT device. SIP Associated Media benchmarks can also be made when testing SUTs.

The test cases covered in this methodology document provide benchmarks metrics of Registration Rate, SIP Session Establishment Rate, Session Capacity, and IM Rate. These can be benchmarked with or without associated Media. Some cases are also included to cover Forking, Loop detection, Encrypted SIP, and SIP Flooding. The test topologies that can be used are described in the Test Setup section. Topologies are provided for benchmarking of a DUT or SUT. Benchmarking with Associated Media can be performed when using a SUT.

SIP permits a wide range of configuration options that are also explained in the Test Setup section. Benchmark metrics could possibly be impacted by Associated Media. The selected values for

Session Duration and Media Streams Per Session enable benchmark metrics to be benchmarked without Associated Media. Session Setup Rate could possibly be impacted by the selected value for Maximum Sessions Attempted. The benchmark for Session Establishment Rate is measured with a fixed value for maximum Session Attempts.

Finally, the overall value of these tests is to serve as a comparison function between multiple SIP implementations. One way to use these tests is to derive benchmarks with SIP devices from Vendor-A, derive a new set of benchmarks with similar SIP devices from Vendor-B and perform a comparison on the results of Vendor-A and Vendor-B. This document does not make any claims on the interpretation of such results.

3. Benchmarking Topologies

Familiarity with the benchmarking models in Section 2.2 of [I-D.sip-bench-term] is assumed. Figures 1 through 10 in [I-D.sip-bench-term] contain the canonical topologies that can be used to perform the benchmarking tests listed in this document.

4. Test Setup Parameters

4.1. Selection of SIP Transport Protocol

Test cases may be performed with any transport protocol supported by SIP. This includes, but is not limited to, SIP TCP, SIP UDP, and TLS. The protocol used for the SIP transport protocol must be reported with benchmarking results.

4.2. Signaling Server

The Signaling Server is defined in the companion terminology document, ([I-D.sip-bench-term], Section 3.2.2) It is a SIP-speaking device that complies with RFC 3261. Conformance to [RFC3261] is assumed for all tests. The Signaling Server may be the DUT or a component of a SUT. The Signaling Server may include Firewall and/or NAT functionality. The components of the SUT may be a single physical device or separate devices.

4.3. Associated Media

Some tests require Associated Media to be present for each SIP session. The test topologies to be used when benchmarking SUT performance for Associated Media are shown in [I-D.sip-bench-term], Figures 4 and 5.

4.4. Selection of Associated Media Protocol

The test cases specified in this document provide SIP performance independent of the protocol used for the media stream. Any media protocol supported by SIP may be used. This includes, but is not limited to, RTP, RTSP, and SRTP. The protocol used for Associated Media MUST be reported with benchmarking results.

4.5. Number of Associated Media Streams per SIP Session

Benchmarking results may vary with the number of media streams per SIP session. When benchmarking a SUT for voice, a single media stream is used. When benchmarking a SUT for voice and video, two media streams are used. The number of Associated Media Streams MUST be reported with benchmarking results.

4.6. Session Duration

SUT performance benchmarks may vary with the duration of SIP sessions. Session Duration MUST be reported with benchmarking results. A Session Duration of zero seconds indicates transmission of a BYE immediately following successful SIP establishment indicate by receipt of a 200 OK. An infinite Session Duration indicates that a BYE is never transmitted.

4.7. Attempted Sessions per Second

DUT and SUT performance benchmarks may vary with the the rate of attempted sessions offered by the Tester. Attempted Sessions per Second MUST be reported with benchmarking results.

4.8. Stress Testing

The purpose of this document is to benchmark SIP performance; this document does not benchmark stability of SIP systems under stressful conditions such as a high rate of Attempted Sessions per Second.

4.9. Benchmarking algorithm

In order to benchmark the test cases uniformly in Section 6, the algorithm described in this section should be used. Both, a prosaic description of the algorithm and a pseudo-code description are provided.

The goal is to find the largest value of a SIP session-request-rate, measured in sessions-per-second, which the DUT/SUT can process with zero errors. To discover that number, an iterative process (defined below) is used to find a candidate for this rate. Once the candidate

rate has been found, the DUT/SUT is subjected to an offered load whose arrival rate is set to that of the candidate rate. This test is run for an extended period of time, which is referred to as infinity, and which is, itself, a parameter of the test labeled T in the pseudo-code. This latter phase of testing is called the steady-state phase. If errors are encountered during this steady-state phase, then the candidate rate is reduced by a defined percent, also a parameter of test, and the steady-state phase is entered again until a final (new) steady-state rate is achieved.

The iterative process itself is defined as follows: a starting rate of 100 sessions per second (sps) is selected. The test is executed for the time period identified by t in the pseudo-code below. If no failures occur, the rate is increased to 150 sps and again tested for time period t. The attempt rate is continuously ramped up until a failure is encountered before the end of the test time t. Then an attempt rate is calculated that is higher than the last successful attempt rate by a quantity equal to half the difference between the rate at which failures occurred and the last successful rate. If this new attempt rate also results in errors, a new attempt rate is tried that is higher than the last successful attempt rate by a quantity equal to half the difference between the rate at which failures occurred and the last successful rate. Continuing in this way, an attempt rate without errors is found. The operator can specify margin of error using the parameter G, measured in units of sessions per second.

The pseudo-code corresponding to the description above follows.

```
; ---- Parameters of test, adjust as needed
t := 5000      ; local maximum; used to figure out largest
               ; value
T := 50000     ; global maximum; once largest value has been
               ; figured out, pump this many requests before calling
               ; the test a success
m := {...}    ; other attributes that affect testing, such
               ; as media streams, etc.
s := 100       ; Initial session attempt rate (in sessions/sec)
G := 5         ; granularity of results - the margin of error in sps
C := 0.05      ; calibration amount: How much to back down if we
               ; have found candidate s but cannot send at rate s for
               ; time T without failures

; ---- End of parameters of test
; ---- Initialization of flags, candidate values and upper bounds

f := false    ; indicates that you had a success after the upper limit
```

```

F := false ; indicates that test is done
c := 0      ; indicates that we have found an upper limit

proc main
  find_largest_value ; First, figure out the largest value.

  ; Now that the largest value (saved in s) has been figured out,
  ; use it for sending out s requests/s and send out T requests.

  do {
    send_traffic(s, m, T) ; send_traffic not shown
    if (all requests succeeded) {
      F := true ; test is done
    } else if (one or more requests fail) {
      s := s - (C * s) ; Reduce s by calibration amount
      steady_state
    }
  } while (F == false)
end proc

proc find_largest_value
  ; Iterative process to figure out the largest value we can
  ; handle with no failures
  do {
    send_traffic(s, m, t) ; Send s request/sec with m
                          ; characteristics until t requests have
                          ; been sent
    if (all requests succeeded) {
      s' := s ; save candidate value of metric

      if ( c == 0 ) {
        s := s + (0.5 * s)

      } else if ((c == 1) && (s''-s')) > 2*G ) {
        s := s + ( 0.5 * (s'' - s) );

      } else if ((c == 1) && ((s''-s') <= 2*G ) {
        f := true;

      }
      else if (one or more requests fail) {
        c := 1 ; we have found an upper bound for the metric
        s'' := s ; save new upper bound
        s := s - (0.5 * (s - s'))
      }
    } while (f == false)
  }
end proc

```


5. Reporting Format

5.1. Test Setup Report

SIP Transport Protocol = _____
(valid values: TCP|UDP|TLS|SCTP|specify-other)
Session Attempt Rate = _____
(session attempts/sec)
IS Media Attempt Rate = _____
(IS media attempts/sec)
Total Sessions Attempted = _____
(total sessions to be created over duration of test)
Media Streams Per Session = _____
(number of streams per session)
Associated Media Protocol = _____
(RTP|RTSP|specify-other)
Media Packet Size = _____
(bytes)
Media Offered Load = _____
(packets per second)
Media Session Hold Time = _____
(seconds)
Establishment Threshold time = _____
(seconds)
Loop Detecting Option = _____
(on|off)
Forking Option
 Number of endpoints request sent to = _____
 (1, means forking is not enabled)
 Type of forking = _____
 (serial|parallel)
Authentication option = _____
 (on|off; if on, please see Notes 2 and 3 below).

Note 1: Total Sessions Attempted is used in the calculation of the Session Establishment Performance ([I-D.sip-bench-term], Section 3.4.5). It is the number of session attempts ([I-D.sip-bench-term], Section 3.1.6) that will be made over the duration of the test.

Note 2: When the Authentication Option is "on" the test tool must be set to ignore 401 and 407 failure responses in any test described as a "test to failure." If this is not done, all such tests will yield trivial benchmarks, as all attempt rates will lead to a failure after the first attempt.

Note 3: When the Authentication Option is "on" the DUT/SUT uses two

transactions instead of one when it is establishing a session or accomplishing a registration. The first transaction ends with the 401 or 407. The second ends with the 200 OK or another failure message. The Test Organization interested in knowing how many times the EA was intended to send a REGISTER as distinct from how many times the EA wound up actually sending a REGISTER may wish to record the following data as well: Number of responses of the following type: 401: _____ (if authentication turned on; N/A otherwise) 407: _____ (if authentication turned on; N/A otherwise)

5.2. Device Benchmarks for IS

Registration Rate = _____
(registrations per second)
Re-registration Rate = _____
(registrations per second)
Session Capacity = _____
(sessions)
Session Overload Capacity = _____
(sessions)
Session Establishment Rate = _____
(sessions per second)
Session Establishment Performance = _____
(total established sessions/total sessions attempted) (no units)
Session Attempt Delay = _____
(seconds)

5.3. Device Benchmarks for NS

IM Rate = _____ (IM messages per second)

6. Test Cases

6.1. Baseline Session Establishment Rate of the test bed

Objective:

To benchmark the Session Establishment Rate of the Emulated Agent (EA) with zero failures.

Procedure:

1. Configure the DUT in the test topology shown in Figure 1 in [I-D.sip-bench-term].

2. Set media streams per session to 0.
3. Execute benchmarking algorithm as defined in Section 4.9 to get the baseline session establishment rate. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: This is the scenario to obtain the maximum Session Establishment Rate of the EA and the test bed when no DUT/SUT is present. The results of this test might be used to normalize test results performed on different test beds or simply to better understand the impact of the DUT/SUT on the test bed in question.

6.2. Session Establishment Rate without media

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with no associated media and zero failures.

Procedure:

1. If the DUT/SUT is being benchmarked as a user agent client or a user agent server, configure the DUT in the test topology shown in Figure 1 or Figure 2 in [I-D.sip-bench-term]. Alternatively, if the DUT is being benchmarked as a proxy or a B2BUA, configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 7 in [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: This is the scenario to obtain the maximum Session Establishment Rate of the DUT/SUT.

6.3. Session Establishment Rate with Media on DUT/SUT

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with zero failures when Associated Media is included in the benchmark test and the media is running through the DUT/SUT.

Procedure:

1. If the DUT is being benchmarked as a user agent client or a user agent server, configure the DUT in the test topology shown in Figure 3 or Figure 4 of [I-D.sip-bench-term]. Alternatively, if the DUT is being benchmarked as a B2BUA,

- configure the DUT in the test topology shown in Figure 6 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 9 in [I-D.sip-bench-term].
 3. Set media streams per session to 1.
 4. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with media. This rate **MUST** be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with Associated Media with any number of media streams per SIP session are expected to be identical to the Session Establishment Rate results obtained without media in the case where the server is running on a platform separate from the platform on which the Media Relay, NAT or Firewall is running. Session Establishment Rate results obtained with Associated Media may be lower than those obtained without media in the case where the server and the NAT, Firewall or Media Relay are running on the same platform.

6.4. Session Establishment Rate with Media not on DUT/SUT

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with zero failures when Associated Media is included in the benchmark test but the media is not running through the DUT/SUT.

Procedure:

1. If the DUT is being benchmarked as proxy or B2BUA, configure the DUT in the test topology shown in Figure 7 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 8 in [I-D.sip-bench-term].
3. Set media streams per session to 1.
4. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with media. This rate **MUST** be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with Associated Media with any number of media streams per SIP session are expected to be identical to the Session Establishment Rate results obtained without media in the case where the server is running on a platform separate from the platform on which the Media Relay, NAT or Firewall is running. Session Establishment Rate results obtained with Associated Media may be lower than those obtained without media in the case where the server and the NAT, Firewall or Media Relay are running on the same platform.

6.5. Session Establishment Rate with Loop Detection Enabled

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with zero failures when the Loop Detection option is enabled and no media streams are present.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, and loop detection is supported in the DUT, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term]. If the DUT does not support loop detection, then this step can be skipped.
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Turn on the Loop Detection option in the DUT or SUT.
5. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with loop detection enabled. This rate **MUST** be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with Loop Detection may be lower than those obtained without Loop Detection enabled.

6.6. Session Establishment Rate with Forking

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with zero failures when the Forking Option is enabled.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, and forking is supported in the DUT, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term]. If the DUT does not support forking, then this step can be skipped.
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Set the number of endpoints that will receive the forked invitation to a value of 2 or more (subsequent tests may increase this value at the discretion of the tester.)
5. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with forking. This rate **MUST** be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with Forking may be lower than those obtained without Forking enabled.

6.7. Session Establishment Rate with Forking and Loop Detection

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with zero failures when both the Forking and Loop Detection Options are enabled.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Enable the Loop Detection Options on the DUT.
5. Set the number of endpoints that will receive the forked invitation to a value of 2 or more (subsequent tests may increase this value at the discretion of the tester.)
6. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with forking and loop detection. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with Forking and Loop Detection may be lower than those obtained with only Forking or Loop Detection enabled.

6.8. Session Establishment Rate with TLS Encrypted SIP

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with zero failures when using TLS encrypted SIP.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Configure Tester to enable TLS over the transport being benchmarked. Make a note the transport when compiling results. May need to run for each transport of interest.

5. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with encryption. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with TLS Encrypted SIP may be lower than those obtained with plaintext SIP.

6.9. Session Establishment Rate with IPsec Encrypted SIP

Objective:

To benchmark the Session Establishment Rate of the DUT/SUT with zero failures when using IPsec Encrypted SIP.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Configure Tester for IPSec.
5. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with encryption. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with IPSec Encrypted SIP may be lower than those obtained with plaintext SIP.

6.10. Session Establishment Rate with SIP Flooding

Objective:

To benchmark the Session Establishment Rate of the SUT with zero failures when SIP Flooding is occurring.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Set s = 500 (c.f. Section 4.9).

5. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with flooding. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Establishment Rate results obtained with SIP Flooding may be degraded.

6.11. Maximum Registration Rate

Objective:

To benchmark the maximum registration rate of the DUT/SUT with zero failures.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. Set media streams per session to 0.
4. Set the registration timeout value to at least 3600 seconds.
5. Execute benchmarking algorithm as defined in Section 4.9 to get the maximum registration rate. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results:

6.12. Maximum Re-Registration Rate

Objective:

To benchmark the maximum re-registration rate of the DUT/SUT with zero failures.

Procedure:

1. If the DUT is being benchmarked as a proxy or B2BUA, then configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 8 of [I-D.sip-bench-term].
3. First, execute test detailed in Section 6.11 to register the endpoints with the registrar.
4. After at least 5 minutes of Step 2, but no more than 10 minutes after Step 2 has been performed, execute test detailed in Section 6.11 again (this will count as a re-registration).
5. Execute benchmarking algorithm as defined in Section 4.9 to get the maximum re-registration rate. This rate MUST be recorded using any pertinent parameters as shown in the

reporting format of Section 5.1.

Expected Results: The rate should be at least equal to but not more than the result of Section 6.11.

6.13. Maximum IM Rate

Objective:

To benchmark the maximum IM rate of the SUT with zero failures.

Procedure:

1. If the DUT/SUT is being benchmarked as a user agent client or a user agent server, configure the DUT in the test topology shown in Figure 1 or Figure 2 in [I-D.sip-bench-term]. Alternatively, if the DUT is being benchmarked as a proxy or a B2BUA, configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 5 in [I-D.sip-bench-term].
3. Execute benchmarking algorithm as defined in Section 4.9 to get the maximum IM rate. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results:

6.14. Session Capacity without Media

Objective:

To benchmark the Session Capacity of the SUT without Associated Media.

Procedure:

1. If the DUT/SUT is being benchmarked as a user agent client or a user agent server, configure the DUT in the test topology shown in Figure 1 or Figure 2 in [I-D.sip-bench-term]. Alternatively, if the DUT is being benchmarked as a proxy or a B2BUA, configure the DUT in the test topology shown in Figure 5 in [I-D.sip-bench-term].
2. Configure a SUT according to the test topology shown in Figure 7 in [I-D.sip-bench-term].
3. Set the media treams per session to be 0.
4. Set the Session Duration to be a value greater than T.
5. Execute benchmarking algorithm as defined in Section 4.9 to get the baseline session establishment rate. This rate MUST be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.
6. The Session Capacity is the product of T and the Session Establishment Rate.

Expected Results: The maximum rate at which the DUT/SUT can handle session establishment requests with no media for an infinitely long period with no errors. This is the SIP "throughput" of the system with no media.

6.15. Session Capacity with Media

Objective:

To benchmark the session capacity of the DUT/SUT with Associated Media.

Procedure:

1. Configure the DUT in the test topology shown in Figure 3 or Figure 4 of [I-D.sip-bench-term] depending on whether the DUT is being benchmarked as a user agent client or user agent server. Alternatively, configure the DUT in the test topology shown in Figure 6 or Figure 7 in [I-D.sip-bench-term] depending on whether the DUT is being benchmarked as a B2BUA or as a proxy. If a SUT is being benchmarked, configure the SUT as shown in Figure 9 of [I-D.sip-bench-term].
2. Set the media streams per session to 1.
3. Set the Session Duration to be a value greater than T.
4. Execute benchmarking algorithm as defined in Section 4.9 to get the baseline session establishment rate. This rate **MUST** be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.
5. The Session Capacity is the product of T and the Session Establishment Rate.

Expected Results: Session Capacity results obtained with Associated Media with any number of media streams per SIP session will be identical to the Session Capacity results obtained without media.

6.16. Session Capacity with Media and a Media Relay/NAT and/or Firewall

Objective:

To benchmark the Session Establishment Rate of the SUT with Associated Media.

Procedure:

1. Configure the SUT as shown in Figure 7 or Figure 10 in [I-D.sip-bench-term].
2. Set media streams per session to 1.
3. Execute benchmarking algorithm as defined in Section 4.9 to get the session establishment rate with media. This rate **MUST** be recorded using any pertinent parameters as shown in the reporting format of Section 5.1.

Expected Results: Session Capacity results obtained with Associated Media with any number of media streams per SIP session may be lower than the Session Capacity without Media result if the Media Relay, NAT or Firewall is sharing a platform with the server.

7. IANA Considerations

This document does not requires any IANA considerations.

8. Security Considerations

Documents of this type do not directly affect the security of Internet or corporate networks as long as benchmarking is not performed on devices or systems connected to production networks. Security threats and how to counter these in SIP and the media layer is discussed in RFC3261, RFC3550, and RFC3711 and various other drafts. This document attempts to formalize a set of common methodology for benchmarking performance of SIP devices in a lab environment.

9. Acknowledgments

The authors would like to thank Keith Drage and Daryl Malas for their contributions to this document.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [I-D.sip-bench-term]
Davids, C., Gurbani, V., and S. Poretsky, "SIP Performance Benchmarking Terminology",
draft-ietf-bmwg-sip-bench-term-04 (work in progress),
March 2012.

10.2. Informative References

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

Authors' Addresses

Carol Davids
Illinois Institute of Technology
201 East Loop Road
Wheaton, IL 60187
USA

Phone: +1 630 682 6024
Email: davids@iit.edu

Vijay K. Gurbani
Bell Laboratories, Alcatel-Lucent
1960 Lucent Lane
Rm 9C-533
Naperville, IL 60566
USA

Phone: +1 630 224 0216
Email: vkg@bell-labs.com

Scott Poretsky
Allot Communications
300 TradeCenter, Suite 4680
Woburn, MA 08101
USA

Phone: +1 508 309 2179
Email: sporetsky@allot.com

Benchmarking Methodology Working Group
Internet-Draft
Expires: April 25, 2013

C. Davids
Illinois Institute of Technology
V. Gurbani
Bell Laboratories,
Alcatel-Lucent
S. Poretsky
Allot Communications
October 22, 2012

Terminology for Benchmarking Session Initiation Protocol (SIP)
Networking Devices
draft-ietf-bmwg-sip-bench-term-05

Abstract

This document provides a terminology for benchmarking the SIP performance of networking devices. The term performance in this context means the capacity of the device- or system-under-test to process SIP messages. Terms are included for test components, test setup parameters, and performance benchmark metrics for black-box benchmarking of SIP networking devices. The performance benchmark metrics are obtained for the SIP signaling plane only. The terms are intended for use in a companion methodology document for characterizing the performance of a SIP networking device under a variety of conditions. The intent of the two documents is to enable a comparison of the capacity of SIP networking devices. Test setup parameters and a methodology document are necessary because SIP allows a wide range of configuration and operational conditions that can influence performance benchmark measurements. A standard terminology and methodology will ensure that benchmarks have consistent definition and were obtained following the same procedures.

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 <http://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 April 25, 2013.

Copyright Notice

Copyright (c) 2012 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Terminology	5
2. Introduction	6
2.1. Scope	7
2.2. Benchmarking Models	9
3. Term Definitions	14
3.1. Protocol Components	14
3.1.1. Session	14
3.1.2. Signaling Plane	17
3.1.3. Media Plane	18
3.1.4. Associated Media	18
3.1.5. Overload	19
3.1.6. Session Attempt	20
3.1.7. Established Session	20
3.1.8. Invite-initiated Session (IS)	21
3.1.9. Non-INVITE-initiated Session (NS)	22
3.1.10. Session Attempt Failure	22
3.1.11. Standing Sessions Count	23
3.2. Test Components	23
3.2.1. Emulated Agent	24
3.2.2. Signaling Server	24
3.2.3. SIP-Aware Stateful Firewall	24
3.2.4. SIP Transport Protocol	25
3.3. Test Setup Parameters	26
3.3.1. Session Attempt Rate	26
3.3.2. IS Media Attempt Rate	26
3.3.3. Establishment Threshold Time	27
3.3.4. Session Duration	27
3.3.5. Media Packet Size	28
3.3.6. Media Offered Load	28
3.3.7. Media Session Hold Time	29
3.3.8. Loop Detection Option	29
3.3.9. Forking Option	30
3.4. Benchmarks	31
3.4.1. Registration Rate	31
3.4.2. Session Establishment Rate	31
3.4.3. Session Capacity	32
3.4.4. Session Overload Capacity	33
3.4.5. Session Establishment Performance	33
3.4.6. Session Attempt Delay	34
3.4.7. IM Rate	34
4. IANA Considerations	35
5. Security Considerations	35
6. Acknowledgments	36
7. References	36
7.1. Normative References	36
7.2. Informational References	36

Appendix A. White Box Benchmarking Terminology	37
Authors' Addresses	37

1. Terminology

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 BCP 14, RFC2119 [RFC2119]. 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 document. The term Throughput is defined in RFC2544 [RFC2544].

For the sake of clarity and continuity, this document adopts the template for definitions set out in Section 2 of RFC 1242 [RFC1242].

The terms Device Under Test (DUT) and System Under Test (SUT) are defined in the following BMWG documents:

Device Under Test (DUT) (c.f., Section 3.1.1 RFC 2285 [RFC2285]).
System Under Test (SUT) (c.f., Section 3.1.2, RFC 2285 [RFC2285]).

Many commonly used SIP terms in this document are defined in RFC 3261 [RFC3261]. For convenience the most important of these are reproduced below. Use of these terms in this document is consistent with their corresponding definition in [RFC3261].

- o Call Stateful: A proxy is call stateful if it retains state for a dialog from the initiating INVITE to the terminating BYE request. A call stateful proxy is always transaction stateful, but the converse is not necessarily true.
- o Stateful Proxy: A logical entity that maintains the client and server transaction state machines defined by this specification during the processing of a request, also known as a transaction stateful proxy. The behavior of a stateful proxy is further defined in Section 16 of RFC 3261 [RFC3261]. A transaction stateful proxy is not the same as a call stateful proxy.
- o Stateless Proxy: A logical entity that does not maintain the client or server transaction state machines defined in this specification when it processes requests. A stateless proxy forwards every request it receives downstream and every response it receives upstream.
- o Back-to-back User Agent: A back-to-back user agent (B2BUA) is a logical entity that receives a request and processes it as a user agent server (UAS). In order to determine how the request should be answered, it acts as a user agent client (UAC) and generates requests. Unlike a proxy server, it maintains dialog state and must participate in all requests sent on the dialogs it has established. Since it is a concatenation of a UAC and a UAS, no explicit definitions are needed for its behavior.

- o Loop: A request that arrives at a proxy, is forwarded, and later arrives back at the same proxy. When it arrives the second time, its Request-URI is identical to the first time, and other header fields that affect proxy operation are unchanged, so that the proxy will make the same processing decision on the request it made the first time. Looped requests are errors, and the procedures for detecting them and handling them are described by the SIP protocol[RFC3261] and also by RFC 5393

2. Introduction

Service Providers and IT Organizations deliver Voice Over IP (VoIP) and Multimedia network services based on the IETF Session Initiation Protocol (SIP) [RFC3261]. SIP is a signaling protocol originally intended to be used to dynamically establish, disconnect and modify streams of media between end users. As it has evolved it has been adopted for use in a growing number of services and applications. Many of these result in the creation of a media session, but some do not. Examples of this latter group include text messaging and subscription services. The set of benchmarking terms provided in this document is intended for use with any SIP-enabled device performing SIP functions in the interior of the network, whether or not these result in the creation of media sessions. The performance of end-user devices is outside the scope of this document.

A number of networking devices have been developed to support SIP-based VoIP services. These include SIP Servers, Session Border Controllers (SBC), Back-to-back User Agents (B2BUA), and SIP-Aware Stateful Firewalls. These devices contain a mix of voice and IP functions whose performance may be reported using metrics defined by the equipment manufacturer or vendor. The Service Provider or IT Organization seeking to compare the performance of such devices will not be able to do so using these vendor-specific metrics, whose conditions of test and algorithms for collection are often unspecified. SIP functional elements and the devices that include them can be configured many different ways and can be organized into various topologies. These configuration and topological choices impact the value of any chosen signaling benchmark. Unless these conditions-of-test are defined, a true comparison of performance metrics will not be possible. Some SIP-enabled network devices terminate or relay media as well as signaling. The processing of media by the device impacts the signaling performance. As a result, the conditions-of-test must include information as to whether or not the device under test processes media and if the device does process media, a description of the media handled and the manner in which it is handled. This document and its companion methodology document [I-D.ietf-bmwg-sip-bench-meth] provide a set of black-box benchmarks

for describing and comparing the performance of devices that incorporate the SIP User Agent Client and Server functions and that operate in the network's core.

The definition of SIP performance benchmarks necessarily includes definitions of Test Setup Parameters and a test methodology. These enable the Tester to perform benchmarking tests on different devices and to achieve comparable results. This document provides a common set of definitions for Test Components, Test Setup Parameters, and Benchmarks. All the benchmarks defined are black-box measurements of the SIP signaling plane. The Test Setup Parameters and Benchmarks defined in this document are intended for use with the companion Methodology document. Benchmarks of internal DUT characteristics (also known as white-box benchmarks) such as Session Attempt Arrival Rate, which is measured at the DUT, are described in Appendix A to allow additional characterization of DUT behavior with different distribution models.

2.1. Scope

The scope of this work item is summarized as follows:

- o This terminology document describes SIP signaling performance benchmarks for black-box measurements of SIP networking devices. Stress and debug scenarios are not addressed in this work item.
- o The DUT must be an RFC 3261 capable network equipment. This may be a Registrar, Redirect Server, Stateless Proxy or Stateful Proxy. A DUT MAY also include a B2BUA, SBC functionality. The DUT MAY be a multi-port SIP-to-switched network gateway implemented as a SIP UAC or UAS.
- o The DUT MAY include an internal SIP Application Level Gateway (ALG), firewall, and/or a Network Address Translator (NAT). This is referred to as the "SIP Aware Stateful Firewall."
- o The DUT or SUT MUST NOT be end user equipment, such as personal digital assistant, a computer-based client, or a user terminal.
- o The Tester acts as multiple "Emulated Agents" (EA) that initiate (or respond to) SIP messages as session endpoints and source (or receive) associated media for established connections.
- o SIP Signaling in presence of Media
 - * The media performance is not benchmarked in this work item.
 - * It is RECOMMENDED that SIP signaling plane benchmarks be performed with media present, but this is optional.
 - * The SIP INVITE requests MUST include the SDP body.
 - * The type of DUT dictates whether the associated media streams traverse the DUT or SUT. Both scenarios are within the scope of this work item.
 - * SIP is frequently used to create media streams; the signaling plane and media plane are treated as orthogonal to each other in this document. While many devices support the creation of

media streams, benchmarks that measure the performance of these streams are outside the scope of this document and its companion methodology document [I-D.ietf-bmwg-sip-bench-meth]. Tests may be performed with or without the creation of media streams. The presence or absence of media streams MUST be noted as a condition of the test as the performance of SIP devices may vary accordingly. Even if the media is used during benchmarking, only the SIP performance will be benchmarked, not the media performance or quality.

- o Both INVITE and non-INVITE scenarios (such as Instant Messages or IM) are addressed in this document. However, benchmarking SIP presence is not a part of this work item.
- o Different transport mechanisms -- such as UDP, TCP, SCTP, or TLS -- may be used. The specific transport mechanism MUST be noted as a condition of the test as the performance of SIP devices may vary accordingly.
- o Looping and forking options are also considered since they impact processing at SIP proxies.
- o REGISTER and INVITE requests may be challenged or remain unchallenged for authentication purpose. Whether or not the REGISTER and INVITE requests are challenged is a condition of test which will be recorded along with other such parameters which may impact the SIP performance of the device or system under test.
- o Re-INVITE requests are not considered in scope of this work item since the benchmarks for INVITEs are based on the dialog created by the INVITE and not on the transactions that take place within that dialog.
- o Only session establishment is considered for the performance benchmarks. Session disconnect is not considered in the scope of this work item. This is because our goal is to determine the maximum throughput of the device or system under test, that is the number of simultaneous SIP sessions that the device or system can support. It is true that there are BYE requests being created during the test process. These transactions do contribute to the load on the device or system under test and thus are accounted for in the metric we derive. We do not seek a separate metric for the number of BYE transactions a device or system can support.
- o SIP Overload [I-D.ietf-soc-overload-design] is within the scope of this work item. We test to failure and then can continue to observe and record the behavior of the system after failures are recorded. The cause of failure is not within the scope of this work. We note the failure and may continue to test until a different failure or condition is encountered. Considerations on how to handle overload are deferred to work progressing in the SOC working group [I-D.ietf-soc-overload-control]. Vendors are, of course, free to implement their specific overload control behavior as the expected test outcome if it is different from the IETF recommendations. However, such behavior MUST be documented and

interpreted appropriately across multiple vendor implementations. This will make it more meaningful to compare the performance of different SIP overload implementations.

- o IMS-specific scenarios are not considered, but test cases can be applied with 3GPP-specific SIP signaling and the P-CSCF as a DUT.

2.2. Benchmarking Models

This section shows ten models to be used when benchmarking SIP performance of a networking device. Figure 1 shows the configuration needed to benchmark the tester itself. This model will be used to establish the limitations of the test apparatus.

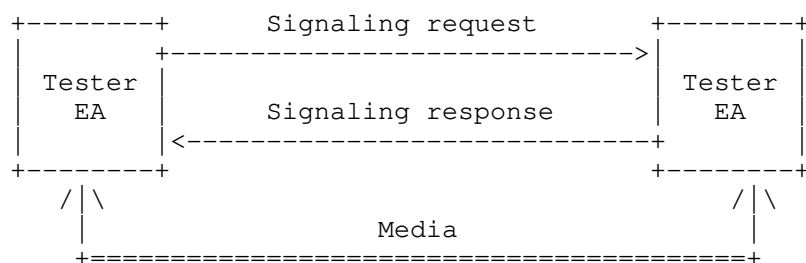


Figure 1: Baseline performance of the Emulated Agent without a DUT present

Figure 2 shows the DUT playing the role of a user agent client (UAC), initiating requests and absorbing responses. This model can be used to baseline the performance of the DUT acting as an UAC without associated media.

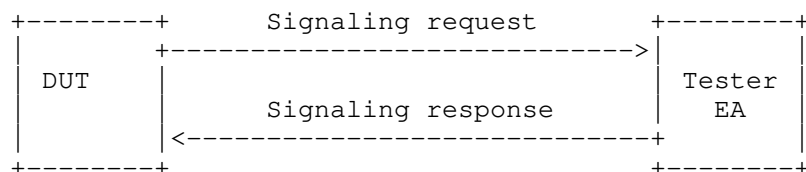


Figure 2: Baseline performance for DUT acting as a user agent client without associated media

Figure 3 shows the DUT plays the role of a user agent server (UAS), absorbing the requests and sending responses. This model can be used as a baseline performance for the DUT acting as a UAS without

associated media.




Figure 3: Baseline performance for DUT acting as a user agent server without associated media

Figure 4 shows the DUT plays the role of a user agent client (UAC), initiating requests and absorbing responses. This model can be used as a baseline performance for the DUT acting as a UAC with associated media.

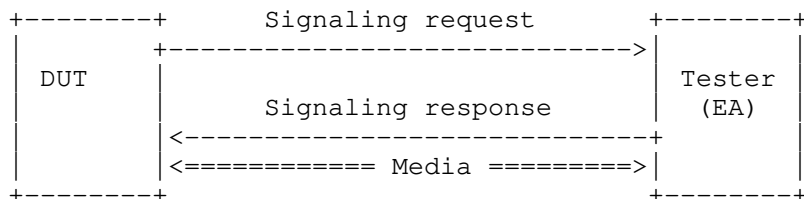


Figure 4: Baseline performance for DUT acting as a user agent client with associated media

Figure 5 shows the DUT plays the role of a user agent server (UAS), absorbing the requests and sending responses. This model can be used as a baseline performance for the DUT acting as a UAS with associated media.

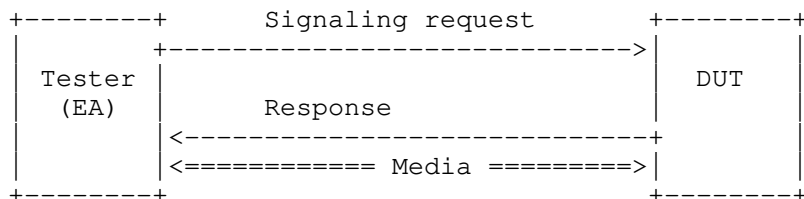


Figure 5: Baseline performance for DUT acting as a user agent server

with associated media

Figure 6 shows that the Tester acts as the initiating and responding EA as the DUT/SUT forwards Session Attempts.




Figure 6: DUT/SUT performance benchmark for session establishment without media

Figure 7 is used when performing those same benchmarks with Associated Media traversing the DUT/SUT.

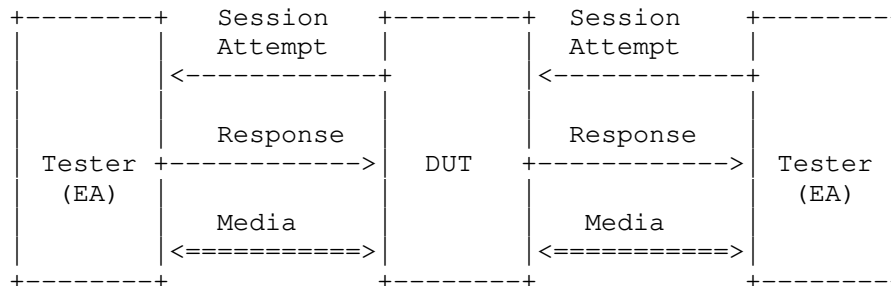


Figure 7: DUT/SUT performance benchmark for session establishment with media traversing the DUT

Figure 8 is to be used when performing those same benchmarks with Associated Media, but the media does not traverse the DUT/SUT. Again, the benchmarking of the media is not within the scope of this work item. The SIP control signaling is benchmarked in the presence of Associated Media to determine if the SDP body of the signaling and the handling of media impacts the performance of the DUT/SUT.

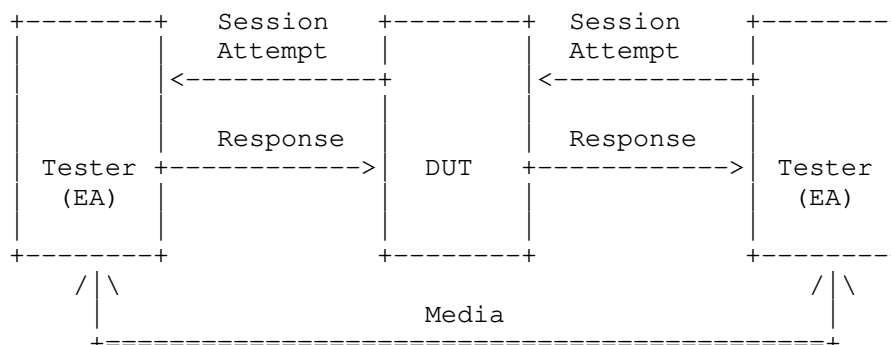


Figure 8: DUT/SUT performance benchmark for session establishment with media external to the DUT

Figure 9 is used when performing benchmarks that require one or more intermediaries to be in the signaling path. The intent is to gather benchmarking statistics with a series of DUTs in place. In this topology, the media is delivered end-to-end and does not traverse the DUT.

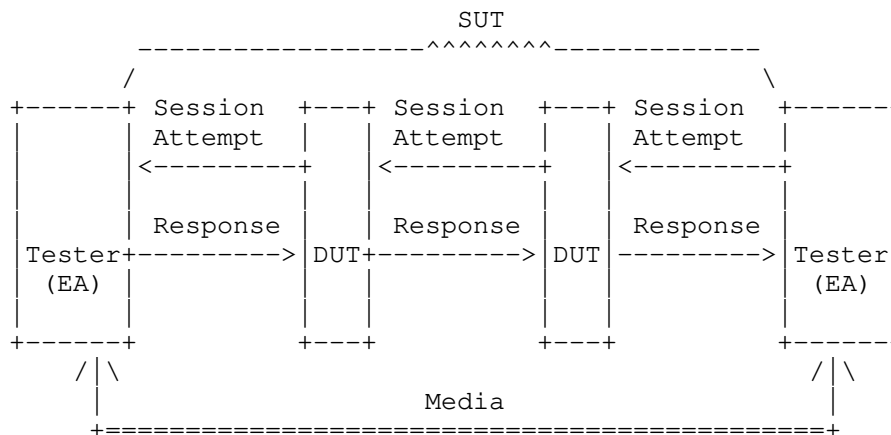


Figure 9: DUT/SUT performance benchmark for session establishment with multiple DUTs and end-to-end media

Figure 10 is used when performing benchmarks that require one or more intermediaries to be in the signaling path. The intent is to gather benchmarking statistics with a series of DUTs in place. In this topology, the media is delivered hop-by-hop through each DUT.

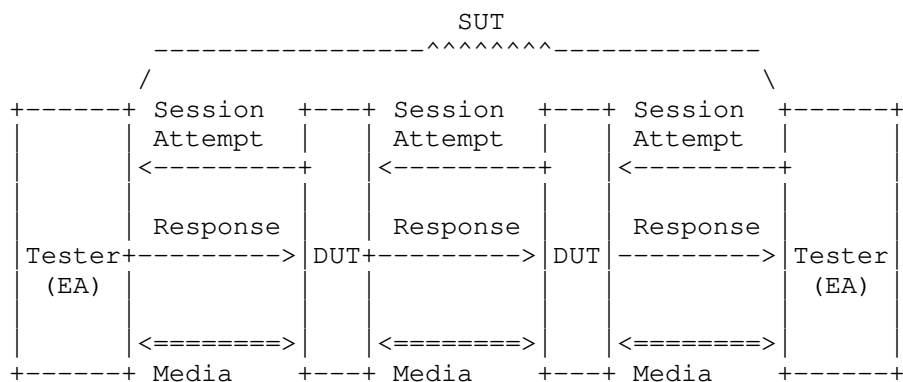


Figure 10: DUT/SUT performance benchmark for session establishment with multiple DUTs and hop-by-hop media

Figure 11 illustrates the SIP signaling for an Established Session. The Tester acts as the EAs and initiates a Session Attempt with the DUT/SUT. When the Emulated Agent (EA) receives a 200 OK from the DUT/SUT that session is considered to be an Established Session. The illustration indicates three states of the session bring created by the EA - Attempting, Established, and Disconnecting. Sessions can be one of two type: Invite-Initiated Session (IS) or Non-Invite Initiated Session (NS). Failure for the DUT/SUT to successfully respond within the Establishment Threshold Time is considered a Session Attempt Failure. SIP Invite messages MUST include the SDP body to specify the Associated Media. Use of Associated Media, to be sourced from the EA, is optional. When Associated Media is used, it may traverse the DUT/SUT depending upon the type of DUT/SUT. The Associated Media is shown in Figure 11 as "Media" connected to media ports M1 and M2 on the EA. After the EA sends a BYE, the session disconnects. Performance test cases for session disconnects are not considered in this work item (the BYE request is shown for completeness.)

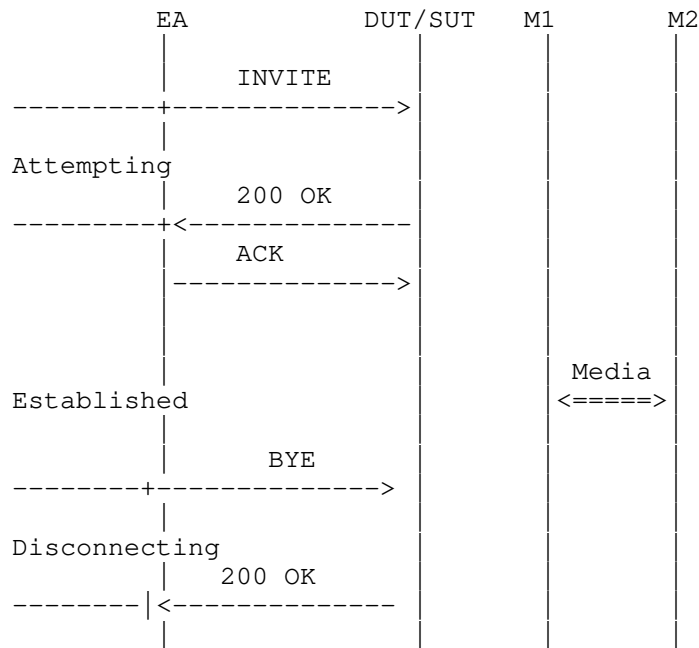


Figure 11: Invite-initiated Session States

3. Term Definitions

3.1. Protocol Components

3.1.1. Session

Definition:

The combination of signaling and media messages and processes that support a SIP-based service.

Discussion:

SIP messages are used to create and manage services for end users. Often, these services include the creation of media streams that are defined in the SDP body of a SIP message and carried in RTP protocol data units. However, SIP messages can also be used to create Instant Message services and subscription services, and such services are not associated with media streams. SIP reserves the term "session" to describe services that are analogous to telephone calls on a circuit switched network. SIP reserves the term "dialog" to refer to a signaling-only relationship between User Agent peers. SIP reserves the term "transaction" to refer to

the brief communication between a client and a server that lasts only until the final response to the SIP request. None of these terms describes the entity whose performance we want to benchmark. For example, the MESSAGE request does not create a dialog and can be sent either within or outside of a dialog. It is not associated with media, but it resembles a phone call in its dependence on human rather than machine initiated responses. The SUBSCRIBE method does create a dialog between the originating end-user and the subscription service. It too is not associated with a media session. In light of these observations we have extended the term "session" to include SIP-based services that are not initiated by INVITE requests and that do not have associated media. In this extended definition, a session always has a signaling component and may also have a media component. Thus, a session can be defined as signaling-only or a combination of signaling and media. We define the term "Associated Media", see Section 3.1.4, to describe the situation in which media is associated with a SIP dialog. The terminology "Invite-initiated Session" (IS) Section 3.1.8 and "Non-invite-Initiated Session" (NS) (add xref target="NS") are used to distinguish between these two types of session. An Invite-initiated Session is a session as defined in SIP. The performance of a device or system that supports Invite-initiated Sessions that do not create media sessions, "Invite-initiated Sessions without Associated Media", can be measured and is of interest for comparison and as a limiting case. The REGISTER request can be considered to be a "Non-invite-initiated Session without Associated Media." A separate set of benchmarks is provided for REGISTER requests since most implementations of SIP-based services require this request and since a registrar may be a device under test.

A Session in the context of this document, can be considered to be a vector with three components:

1. A component in the signaling plane (SIP messages), sess.sig;
2. A media component in the media plane (RTP and SRTP streams for example), sess.med (which may be null);
3. A control component in the media plane (RTCP messages for example), sess.medc (which may be null).

An IS is expected to have non-null sess.sig and sess.med components. The use of control protocols in the media component is media dependent, thus the expected presence or absence of sess.medc is media dependent and test-case dependent. An NS is expected to have a non-null sess.sig component, but null sess.med and sess.medc components.

Packets in the Signaling Plane and Media Plane will be handled by different processes within the DUT. They will take different paths within a SUT. These different processes and paths may produce variations in performance. The terminology and benchmarks defined in this document and the methodology for their use are designed to enable us to compare performance of the DUT/SUT with reference to the type of SIP-supported application it is handling.

Note that one or more sessions can simultaneously exist between any participants. This can be the case, for example, when the EA sets up both an IM and a voice call through the DUT/SUT. These sessions are represented as an array session[x].

Sessions will be represented as a vector array with three components, as follows:

session->

session[x].sig, the signaling component

session[x].medc[y], the media control component (e.g. RTCP)

session[x].med[y], an array of associated media streams (e.g. RTP, SRTP, RTSP, MSRP). This media component may consist of zero or more media streams.

Figure 12 models the vectors of the session.

Measurement Units:

N/A.

Issues:

None.

See Also:

Media Plane

Signaling Plane

Associated Media

Invite-initiated Session (IS)

Non-invite-initiated Session (NS)




Figure 12: Session components

3.1.2. Signaling Plane

Definition:

The plane in which SIP messages [RFC3261] are exchanged between SIP Agents [RFC3261].

Discussion:

SIP messages are used to establish sessions in several ways: directly between two User Agents [RFC3261], through a Proxy Server [RFC3261], or through a series of Proxy Servers. The Session Description Protocol is included in the Signaling Plane. (SDP). The Signaling Plane for a single Session is represented by session.sig.

Measurement Units:

N/A.

Issues:

None.

See Also:

Media Plane

EAs

3.1.3. Media Plane

Definition:

The data plane in which one or more media streams and their associated media control protocols are exchanged between User Agents after a media connection has been created by the exchange of signaling messages in the Signaling Plane.

Discussion:

Media may also be known as the "bearer channel". The Media Plane MUST include the media control protocol, if one is used, and the media stream(s). Examples of media are audio and video. The media streams are described in the SDP of the Signaling Plane. The media for a single Session is represented by session.med. The media control protocol for a single media description is represented by session.medc.

Measurement Units:

N/A.

Issues:

None.

See Also:

Signaling Plane

3.1.4. Associated Media

Definition:

Media that corresponds to an 'm' line in the SDP payload of the Signaling Plane.

Discussion:

Any media protocol MAY be used.

For any session's signaling component, `session.sig`, there may be zero, one, or multiple associated media streams. When there are multiple media streams, these are represented by a vector array `session.med[y]`. When there are multiple media streams there will be multiple media control protocol descriptions as well. They are represented by a vector array `session.medc[y]`.

Measurement Units:

N/A.

Issues:

None.

3.1.5. Overload

Definition:

Overload is defined as the state where a SIP server does not have sufficient resources to process all incoming SIP messages [I-D.ietf-soc-overload-design].

Discussion:

The distinction between an overload condition and other failure scenarios is outside the scope of black box testing and of this document. Under overload conditions, all or a percentage of Session Attempts will fail due to lack of resources. In black box testing the cause of the failure is not explored. The fact that a failure occurred for whatever reason, will trigger the tester to reduce the offered load, as described in the companion methodology document, [I-D.ietf-bmwg-sip-bench-meth]. SIP server resources may include CPU processing capacity, network bandwidth, input/output queues, or disk resources. Any combination of resources may be fully utilized when a SIP server (the DUT/SUT) is in the overload condition. For proxy-only type of devices, it is expected that the proxy will be driven into overload based on the delivery rate of signaling requests.

For UA-type of network devices such as gateways, it is expected that the UA will be driven into overload based on the volume of media streams it is processing.

Measurement Units:

N/A.

Issues:

The issue of overload in SIP networks is currently a topic of discussion in the SIPPING WG. The normal response to an overload stimulus -- sending a 503 response -- is considered inadequate and new response codes and behaviors may be specified in the future. From the perspective of this document, all these responses will be considered to be failures. There is thus no dependency between this document and the ongoing work on the treatment of overload failure.

3.1.6. Session Attempt**Definition:**

A SIP request sent by the EA that has not received a final response.

Discussion:

The attempted session may be Invite Initiated or Non-Invite Initiated. When counting the number of session attempts we include all INVITEs that are rejected for lack of authentication information. The EA needs to record the total number of session attempts including those attempts that are routinely rejected by a proxy that requires the UA to authenticate itself. The EA is provisioned to deliver a specific number of session attempts per second. But the EA must also count the actual number of session attempts per given tie interval.

Measurement Units:

N/A.

Issues:

None.

See Also:

Session
Session Attempt Rate
Invite-initiated Session
Non-Invite initiated Session

3.1.7. Established Session**Definition:**

A SIP session for which the EA acting as the UE/UA has received a 200 OK message.

Discussion:

An Established Session MAY be Invite Initiated or Non-Invite Initiated.

Measurement Units:

N/A.

Issues:

None.

See Also:

Invite-initiated Session
Session Attempting State
Session Disconnecting State

3.1.8. Invite-initiated Session (IS)**Definition:**

A Session that is created by an exchange of messages in the Signaling Plane, the first of which is a SIP INVITE request.

Discussion:

When an IS becomes an Established Session its signaling component is identified by the SIP dialog parameter values, Call-ID, To-tag, and From-tag (RFC3261 [RFC3261]). An IS may have zero, one or multiple Associated Media descriptions in the SDP body. The inclusion of media is test case dependent. An IS is successfully established if the following two conditions are met:

1. Sess.sig is established by the end of Establishment Threshold Time (c.f. Section 3.3.3), and
2. If a media session is described in the SDP body of the signaling message, then the media session is established by the end of Establishment Threshold Time (c.f. Section 3.3.3). An SBC or B2BUA may receive media from a calling or called party before a signaling dialog is established and certainly before a confirmed dialog is established. The EA can be built in such a way that it does not send early media or it needs to include a parameter that indicates when it will send media. This parameter must be included in the list of test setup parameters in Section 5.1 of [I-D.ietf-bmwg-sip-bench-meth]

Measurement Units:

N/A.

Issues:

None.

See Also:

Session

Non-Invite initiated Session

Associated Media

3.1.9. Non-INVITE-initiated Session (NS)

Definition:

A session that is created by an exchange of SIP messages in the Signaling Plane the first of which is not a SIP INVITE message.

Discussion:

An NS is successfully established if the Session Attempt via a non- INVITE request results in the EA receiving a 2xx reply before the expiration of the Establishment Threshold timer (c.f., Section 3.3.3). An example of a NS is a session created by the SUBSCRIBE request.

Measurement Units:

N/A.

Issues:

None.

See Also:

Session

Invite-initiated Session

3.1.10. Session Attempt Failure

Definition:

A session attempt that does not result in an Established Session.

Discussion:

The session attempt failure may be indicated by the following observations at the EA:

1. Receipt of a SIP 4xx, 5xx, or 6xx class response to a Session Attempt.
2. The lack of any received SIP response to a Session Attempt within the Establishment Threshold Time (c.f. Section 3.3.3).

Measurement Units:

N/A.

Issues:

None.

See Also:

Session Attempt

3.1.11. Standing Sessions Count

Definition:

The number of Sessions currently established on the DUT/SUT at any instant.

Discussion:

The number of Standing Sessions is influenced by the Session Duration and the Session Attempt Rate. Benchmarks MUST be reported with the maximum and average Standing Sessions for the DUT/SUT for the duration of the test. In order to determine the maximum and average Standing Sessions on the DUT/SUT for the duration of the test it is necessary to make periodic measurements of the number of Standing Sessions on the DUT/SUT. The recommended value for the measurement period is 1 second. Since we cannot directly poll the DUT/SUT, we take the number of standing sessions on the DUT/SUT to be the number of distinct calls as measured by the number of distinct Call-IDs that the EA is processing at the time of measurement. The EA must make that count available for viewing and recording.

Measurement Units:

Number of sessions

Issues:

None.

See Also:

Session Duration
Session Attempt Rate
Session Attempt Rate
Emulated Agent

3.2. Test Components

3.2.1. Emulated Agent

Definition:

A device in the test topology that initiates/responds to SIP messages as one or more session endpoints and, wherever applicable, sources/receives Associated Media for Established Sessions.

Discussion:

The EA functions in the Signaling and Media Planes. The Tester may act as multiple EAs.

Measurement Units:

N/A

Issues:

None.

See Also:

Media Plane
Signaling Plane
Established Session
Associated Media

3.2.2. Signaling Server

Definition:

Device in the test topology that acts to create sessions between EAs. This device is either a DUT or a component of a SUT.

Discussion:

The DUT MUST be an RFC 3261 capable network equipment such as a Registrar, Redirect Server, User Agent Server, Stateless Proxy, or Stateful Proxy. A DUT MAY also include B2BUA or SBC.

Measurement Units:

NA

Issues:

None.

See Also:

Signaling Plane

3.2.3. SIP-Aware Stateful Firewall

Definition:

Device in the test topology that provides protection against various types of security threats to which the Signaling and Media Planes of the EAs and Signaling Server are vulnerable.

Discussion:

Threats may include Denial-of-Service, theft of service and misuse of service. The SIP-Aware Stateful Firewall MAY be an internal component or function of the Session Server. The SIP-Aware Stateful Firewall MAY be a standalone device. If it is a standalone device it MUST be paired with a Signaling Server. If it is a standalone device it MUST be benchmarked as part of a SUT. SIP-Aware Stateful Firewalls MAY include Network Address Translation (NAT) functionality. Ideally, the inclusion of the SIP-Aware Stateful Firewall in the SUT does not lower the measured values of the performance benchmarks.

Measurement Units:

N/A

Issues:

None.

See Also:

3.2.4. SIP Transport Protocol

Definition:

The protocol used for transport of the Signaling Plane messages.

Discussion:

Performance benchmarks may vary for the same SIP networking device depending upon whether TCP, UDP, TLS, SCTP, or another transport layer protocol is used. For this reason it MAY be necessary to measure the SIP Performance Benchmarks using these various transport protocols. Performance Benchmarks MUST report the SIP Transport Protocol used to obtain the benchmark results.

Measurement Units:

TCP,UDP, SCTP, TLS over TCP, TLS over UDP, or TLS over SCTP

Issues:

None.

See Also:

3.3. Test Setup Parameters

3.3.1. Session Attempt Rate

Definition:

Configuration of the EA for the number of sessions per second that the EA attempts to establish using the services of the DUT/SUT.

Discussion:

The Session Attempt Rate is the number of sessions per second that the EA sends toward the DUT/SUT. Some of the sessions attempted may not result in a session being established. A session in this case may be either an IS or an NS.

Measurement Units:

Session attempts per second

Issues:

None.

See Also:

Session

Session Attempt

3.3.2. IS Media Attempt Rate

Definition:

Configuration on the EA for the rate, measured in sessions per second, at which the EA attempts to establish INVITE-initiated sessions with Associated Media, using the services of the DUT/SUT.

Discussion:

An IS is not required to include a media description. The IS Media Attempt Rate defines the number of media sessions we are trying to create, not the number of media sessions that are actually created. Some attempts might not result in successful sessions established on the DUT.

Measurement Units:

session attempts per second (saps)

Issues:

None.

See Also:
IS

3.3.3. Establishment Threshold Time

Definition:

Configuration of the EA for representing the amount of time that an EA will wait before declaring a Session Attempt Failure.

Discussion:

This time duration is test dependent.

It is RECOMMENDED that the Establishment Threshold Time value be set to Timer B (for ISs) or Timer F (for NSs) as specified in RFC 3261, Table 4 [RFC3261]. Following the default value of T1 (500ms) specified in the table and a constant multiplier of 64 gives a value of 32 seconds for this timer (i.e., $500\text{ms} * 64 = 32\text{s}$).

Measurement Units:
seconds

Issues:
None.

See Also:
session establishment failure

3.3.4. Session Duration

Definition:

Configuration of the EA that represents the amount of time that the SIP dialog is intended to exist between the two EAs associated with the test.

Discussion:

The time at which the BYE is sent will control the Session Duration

Normally the Session Duration will be the same as the Media Session Hold Time. However, it is possible that the dialog established between the two EAs can support different media sessions at different points in time. Providing both parameters allows the testing agency to explore this possibility.

Measurement Units:
seconds

Issues:
None.

See Also:
Media Session Hold Time

3.3.5. Media Packet Size

Definition:
Configuration on the EA for a fixed size of packets used for media streams.

Discussion:
For a single benchmark test, all sessions use the same size packet for media streams. The size of packets can cause variation in performance benchmark measurements.

Measurement Units:
bytes

Issues:
None.

See Also:

3.3.6. Media Offered Load

Definition:
Configuration of the EA for the constant rate of Associated Media traffic offered by the EA to the DUT/SUT for one or more Established Sessions of type IS.

Discussion:
The Media Offered Load to be used for a test MUST be reported with three components:
1. per Associated Media stream;
2. per IS;
3. aggregate.
For a single benchmark test, all sessions use the same Media Offered Load per Media Stream. There may be multiple Associated Media streams per IS. The aggregate is the sum of all Associated Media for all IS.

Measurement Units:
packets per second (pps)

Issues:
None.

See Also:
Established Session
Invite Initiated Session
Associated Media

3.3.7. Media Session Hold Time

Definition:
Parameter configured at the EA, that represents the amount of time that the Associated Media for an Established Session of type IS will last.

Discussion:
The Associated Media streams may be bi-directional or uni-directional as indicated in the test methodology. Normally the Media Session Hold Time will be the same as the Session Duration. However, it is possible that the dialog established between the two EAs can support different media sessions at different points in time. Providing both parameters allows the testing agency to explore this possibility.

Measurement Units:
seconds

Issues:
None.

See Also:
Associated Media
Established Session
Invite-initiated Session (IS)

3.3.8. Loop Detection Option

Definition:
An option that causes a Proxy to check for loops in the routing of a SIP request before forwarding the request.

Discussion:

This is an optional process that a SIP proxy may employ; the process is described under Proxy Behavior in RFC 3261 [RFC3261] in Section 16.3 Request Validation and that section also contains suggestions as to how the option could be implemented. Any procedure to detect loops will use processor cycles and hence could impact the performance of a proxy.

Measurement Units:

NA

Issues:

None.

See Also:**3.3.9. Forking Option****Definition:**

An option that enables a Proxy to fork requests to more than one destination.

Discussion:

This is an process that a SIP proxy may employ to find the UAS. The option is described under Proxy Behavior in RFC 3261 in Section 16.1. A proxy that uses forking must maintain state information and this will use processor cycles and memory. Thus the use of this option could impact the performance of a proxy and different implementations could produce different impacts. SIP supports serial or parallel forking. When performing a test, the type of forking mode MUST be indicated.

Measurement Units:

The number of endpoints that will receive the forked invitation. A value of 1 indicates that the request is destined to only one endpoint, a value of 2 indicates that the request is forked to two endpoints, and so on. This is an integer value ranging between 1 and N inclusive, where N is the maximum number of endpoints to which the invitation is sent.
Type of forking used, namely parallel or serial.

Issues:

None.

See Also:

3.4. Benchmarks

3.4.1. Registration Rate

Definition:

The maximum number of registrations that can be successfully completed by the DUT/SUT in a given time period without registration failures in that time period.

Discussion:

This benchmark is obtained with zero failure in which 100% of the registrations attempted by the EA are successfully completed by the DUT/SUT. The registration rate provisioned on the Emulated Agent is raised and lowered as described in the algorithm in the companion methodology draft [I-D.ietf-bmwg-sip-bench-meth] until a traffic load consisting of registrations at the given attempt rate over the sustained period of time identified by T in the algorithm completes without failure.

Measurement Units:

registrations per second (rps)

Issues:

None.

See Also:

3.4.2. Session Establishment Rate

Definition:

The maximum number of sessions that can be successfully completed by the DUT/SUT in a given time period without session establishment failures in that time period.

Discussion:

This benchmark is obtained with zero failure in which 100% of the sessions attempted by the Emulated Agent are successfully completed by the DUT/SUT. The session attempt rate provisioned on the EA is raised and lowered as described in the algorithm in the accompanying methodology document, until a traffic load at the given attempt rate over the sustained period of time identified by T in the algorithm completes without any failed session attempts. Sessions may be IS or NS or a mix of both and will be defined in the particular test.

Measurement Units:
sessions per second (sps)

Issues:
None.

See Also:
Invite-initiated Sessions
Non-INVITE initiated Sessions
Session Attempt Rate

3.4.3. Session Capacity

Definition:
The maximum value of Standing Sessions Count achieved by the DUT/SUT during a time period T in which the EA is sending session establishment messages at the Session Establishment Rate.

Discussion:
Sessions may be IS or NS. If they are IS they can be with or without media. When benchmarking Session Capacity for sessions with media it is required that these sessions be permanently established (i.e., they remain active for the duration of the test.) This can be achieved by causing the EA not to send a BYE for the duration of the testing. In the signaling plane, this requirement means that the dialog lasts as long as the test lasts. When media is present, the Media Session Hold Time MUST be set to infinity so that sessions remain established for the duration of the test. If the DUT/SUT is dialog-stateful, then we expect its performance will be impacted by setting Media Session Hold Time to infinity, since the DUT/SUT will need to allocate resources to process and store the state information. The report of the Session Capacity must include the Session Establishment Rate at which it was measured.

Measurement Units:
sessions

Issues:
None.

See Also:
Established Session

Session Attempt Rate
Session Attempt Failure

3.4.4. Session Overload Capacity

Definition:

The maximum number of Established Sessions that can exist simultaneously on the DUT/SUT until it stops responding to Session Attempts.

Discussion:

Session Overload Capacity is measured after the Session Capacity is measured. The Session Overload Capacity is greater than or equal to the Session Capacity. When benchmarking Session Overload Capacity, continue to offer Session Attempts to the DUT/SUT after the first Session Attempt Failure occurs and measure Established Sessions until no there is no SIP message response for the duration of the Establishment Threshold. Note that the Session Establishment Performance is expected to decrease after the first Session Attempt Failure occurs.

Units:

Sessions

Issues:

None.

See Also:

Overload
Session Capacity
Session Attempt Failure

3.4.5. Session Establishment Performance

Definition:

The percent of Session Attempts that become Established Sessions over the duration of a benchmarking test.

Discussion:

Session Establishment Performance is a benchmark to indicate session establishment success for the duration of a test. The duration for measuring this benchmark is to be specified in the Methodology. The Session Duration SHOULD be configured to infinity so that sessions remain established for the entire test duration.

Session Establishment Performance is calculated as shown in the following equation:

$$\text{Session Establishment Performance} = \frac{\text{Total Established Sessions}}{\text{Total Session Attempts}}$$

Session Establishment Performance may be monitored real-time during a benchmarking test. However, the reporting benchmark MUST be based on the total measurements for the test duration.

Measurement Units:
Percent (%)

Issues:
None.

See Also:
Established Session
Session Attempt

3.4.6. Session Attempt Delay

Definition:

The average time measured at the EA for a Session Attempt to result in an Established Session.

Discussion:

Time is measured from when the EA sends the first INVITE for the call-ID in the case of an IS. Time is measured from when the EA sends the first non-INVITE message in the case of an NS. Session Attempt Delay MUST be measured for every established session to calculate the average. Session Attempt Delay MUST be measured at the Session Establishment Rate.

Measurement Units:
Seconds

Issues:
None.

See Also:
Session Establishment Rate

3.4.7. IM Rate

Definition:

Maximum number of IM messages completed by the DUT/SUT.

Discussion:

For a UAS, the definition of success is the receipt of an IM request and the subsequent sending of a final response.

For a UAC, the definition of success is the sending of an IM request and the receipt of a final response to it. For a proxy, the definition of success is as follows:

- A. the number of IM requests it receives from the upstream client MUST be equal to the number of IM requests it sent to the downstream server; and
- B. the number of IM responses it receives from the downstream server MUST be equal to the number of IM requests sent to the downstream server; and
- C. the number of IM responses it sends to the upstream client MUST be equal to the number of IM requests it received from the upstream client.

Measurement Units:

IM messages per second

Issues:

None.

See Also:

4. IANA Considerations

This document requires no IANA considerations.

5. Security Considerations

Documents of this type do not directly affect the security of Internet or corporate networks as long as benchmarking is not performed on devices or systems connected to production networks. Security threats and how to counter these in SIP and the media layer is discussed in RFC3261 [RFC3261], RFC 3550 [RFC3550], RFC3711 [RFC3711] and various other drafts. This document attempts to formalize a set of common terminology for benchmarking SIP networks. Packets with unintended and/or unauthorized DSCP or IP precedence values may present security issues. Determining the security consequences of such packets is out of scope for this document.

6. Acknowledgments

The authors would like to thank Keith Drage, Cullen Jennings, Daryl Malas, Al Morton, and Henning Schulzrinne for invaluable contributions to this document.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [I-D.ietf-bmwg-sip-bench-meth] Davids, C., Gurbani, V., and S. Poretsky, "Methodology for Benchmarking SIP Networking Devices", draft-ietf-bmwg-sip-bench-meth-04 (work in progress), March 2012.

7.2. Informational References

- [RFC2285] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.
- [RFC1242] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [I-D.ietf-soc-overload-design] Hilt, V., Noel, E., Shen, C., and A. Abdelal, "Design Considerations for Session Initiation Protocol (SIP) Overload Control", draft-ietf-soc-overload-design-08 (work

in progress), July 2011.

[I-D.ietf-soc-overload-control]

Gurbani, V., Hilt, V., and H. Schulzrinne, "Session Initiation Protocol (SIP) Overload Control", draft-ietf-soc-overload-control-10 (work in progress), October 2012.

Appendix A. White Box Benchmarking Terminology

Session Attempt Arrival Rate

Definition:

The number of Session Attempts received at the DUT/SUT over a specified time period.

Discussion:

Sessions Attempts are indicated by the arrival of SIP INVITES OR SUBSCRIBE NOTIFY messages. Session Attempts Arrival Rate distribution can be any model selected by the user of this document. It is important when comparing benchmarks of different devices that same distribution model was used. Common distributions are expected to be Uniform and Poisson.

Measurement Units:

Session attempts/sec

Issues:

None.

See Also:

Session Attempt

Authors' Addresses

Carol Davids
Illinois Institute of Technology
201 East Loop Road
Wheaton, IL 60187
USA

Phone: +1 630 682 6024
Email: davids@iit.edu

Vijay K. Gurbani
Bell Laboratories, Alcatel-Lucent
1960 Lucent Lane
Rm 9C-533
Naperville, IL 60566
USA

Phone: +1 630 224 0216
Email: vkg@bell-labs.com

Scott Poretsky
Allot Communications
300 TradeCenter, Suite 4680
Woburn, MA 08101
USA

Phone: +1 508 309 2179
Email: sporetsky@allot.com

Benchmarking Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 27, 2012

R. Papneja
Huawei Technologies
B. Parise
Cisco Systems
S. Hares
Huawei Technologies
I. Varlashkin
Easynet Global Services
March 26, 2012

Basic BGP Convergence Benchmarking Methodology for Data Plane
Convergence
draft-papneja-bgp-basic-dp-convergence-03.txt

Abstract

BGP is widely deployed and used by several service providers as the default Inter AS routing protocol. It is of utmost importance to ensure that when a BGP peer or a downstream link of a BGP peer fails, the alternate paths are rapidly used and routes via these alternate paths are installed. This document provides the basic BGP Benchmarking Methodology using existing BGP Convergence Terminology, RFC 4098.

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 <http://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 September 27, 2012.

Copyright Notice

Copyright (c) 2012 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

(<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	4
1.1. Precise Benchmarking Definition	4
1.2. Purpose of BGP FIB (Data Plane) Convergence	4
1.3. Control Plane Convergence	5
1.4. Benchmarking Testing	5
2. Existing Definitions and Requirements	5
3. Test Topologies	6
3.1. General Reference Topologies	6
4. Test Considerations	8
4.1. Number of Peers	9
4.2. Number of Routes per Peer	9
4.3. Policy Processing/Reconfiguration	9
4.4. Configured Parameters (Timers, etc..)	9
4.5. Interface Types	11
4.6. Measurement Accuracy	11
4.7. Measurement Statistics	11
4.8. Authentication	12
4.9. Convergence Events	12
4.10. High Availability	12
5. Test Cases	12
5.1. Basic Convergence Tests	12
5.1.1. RIB-IN Convergence	13
5.1.2. RIB-OUT Convergence	14
5.1.3. eBGP Convergence	16
5.1.4. iBGP Convergence	16
5.1.5. eBGP Multihop Convergence	16
5.2. BGP Failure/Convergence Events	18
5.2.1. Physical Link Failure on DUT End	18
5.2.2. Physical Link Failure on Remote/Emulator End	19
5.2.3. ECMP Link Failure on DUT End	19
5.3. BGP Adjacency Failure (Non-Physical Link Failure) on Emulator	19
5.4. BGP Hard Reset Test Cases	21
5.4.1. BGP Non-Recovering Hard Reset Event on DUT	21
5.5. BGP Soft Reset	22
5.6. BGP Route Withdrawal Convergence Time	23
5.7. BGP Path Attribute Change Convergence Time	25
5.8. BGP Graceful Restart Convergence Time	26
6. Reporting Format	28
7. IANA Considerations	31
8. Security Considerations	31
9. References	31
9.1. Normative References	31
9.2. Informative References	32
Authors' Addresses	32

1. Introduction

This document defines the methodology for benchmarking data plane FIB convergence performance of BGP in router and switches for simple topologies of 3 or 4 nodes. The methodology proposed in this document applies to both IPv4 and IPv6 and if a particular test is unique to one version, it is marked accordingly. For IPv6 benchmarking the device under test will require the support of Multi-Protocol BGP (MP-BGP) [RFC4760, RFC2545].

The scope of this companion document is limited to basic BGP protocol FIB convergence measurements. BGP extensions outside of carrying IPv6 in (MP-BGP) [RFC4760, RFC2545] are outside the scope of this document. Interaction with IGPs (IGP interworking) is outside the scope of this document.

1.1. Precise Benchmarking Definition

Since benchmarking is science of precision, let us restate the purpose of this document in benchmarking terms. This document defines methodology to test

- data plane convergence on a single BGP device that supports the BGP [RFC4271] functionality
- in test topology of 3 or 4 nodes
- using Basic BGP

Data plane convergence is defined as the completion of all FIB changes so that all forwarded traffic now takes the new proposed route. RFC 4098 defines the terms BGP device, FIB and the forwarded traffic. Data plane convergence is different than control plane convergence within a node.

Basic BGP is defined as RFC 4271 functional with Multi-Protocol BGP (MP-BGP) [RFC4760, RFC2545] for IPv6. The use of other extensions of BGP to support layer-2, layer-3 virtual private networks (VPN) are out of scope of this document.

The terminology used in this document is defined in [RFC4098]. One additional term is defined in this draft: FIB (Data plane) BGP Convergence.

1.2. Purpose of BGP FIB (Data Plane) Convergence

In the current Internet architecture the Inter-Autonomous System (inter-AS) transit is primarily available through BGP. To maintain a

reliable connectivity within intra-domains or across inter-domains, fast recovery from failures remains most critical. To ensure minimal traffic losses, many service providers are requiring BGP implementations to converge the entire Internet routing table within sub-seconds at FIB level.

Furthermore, to compare these numbers amongst various devices, service providers are also looking at ways to standardize the convergence measurement methods. This document offers test methods for simple topologies. These simple tests will provide a quick high-level check, of the BGP data plane convergence across multiple implementations.

1.3. Control Plane Convergence

The convergence of BGP occurs at two levels: RIB and FIB convergence. RFC 4098 defines terms for BGP control plane convergence. Methodologies which test control plane convergence are out of scope for this draft.

1.4. Benchmarking Testing

In order to ensure that the results obtained in tests are repeatable, careful setup of initial conditions and exact steps are required.

This document proposes these initial conditions, test steps, and result checking. To ensure uniformity of the results all optional parameters SHOULD be disabled and all settings SHOULD be changed to default, these may include BGP timers as well.

2. Existing Definitions and Requirements

RFC 1242, "Benchmarking Terminology for Network Interconnect Devices" [RFC1242] and RFC 2285, "Benchmarking Terminology for LAN Switching Devices" [RFC2285] SHOULD be reviewed in conjunction with this document. WLAN-specific terms and definitions are also provided in Clauses 3 and 4 of the IEEE 802.11 standard [802.11]. Commonly used terms may also be found in RFC 1983 [RFC1983].

For the sake of clarity and continuity, this document adopts the general template for benchmarking terminology set out in Section 2 of RFC 1242. Definitions are organized in alphabetical order, and grouped into sections for ease of reference. The following terms are assumed to be taken as defined in RFC 1242 [RFC1242]: Throughput, Latency, Constant Load, Frame Loss Rate, and Overhead Behavior. In addition, the following terms are taken as defined in [RFC2285]: Forwarding Rates, Maximum Forwarding Rate, Loads, Device Under Test

(DUT), and System Under Test (SUT).

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 [RFC2119].

3. Test Topologies

This section describes simple test setups for use in BGP benchmarking tests measuring convergence of the FIB (data plane) after the BGP updates has been received.

These simple test nodes have 3 or 4 nodes with the following configuration:

1. Basic Test Setup
2. Three node setup for iBGP or eBGP convergence
3. Setup for eBGP multihop test scenario
4. Four node setup for iBGP or eBGP convergence

Individual tests refer to these topologies.

Figures 1-4 use the following conventions

- o AS-X: Autonomous System X
- o Loopback Int: Loopback interface on the BGP enabled device
- o R2: Helper router

3.1. General Reference Topologies

Emulator acts as 1 or more BGP peers for different testcases.




Figure 1 Basic Test Setup




Figure 2 Three Node Setup for eBGP and iBGP Convergence

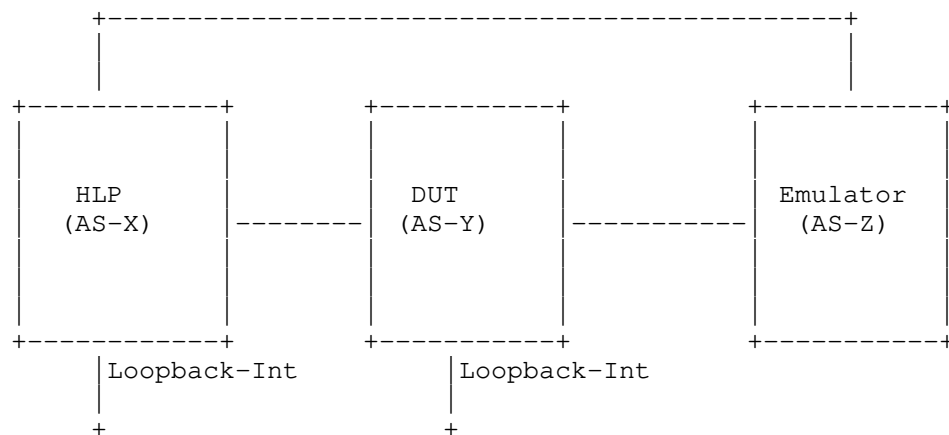


Figure 3 BGP Convergence for eBGP Multihop Scenario

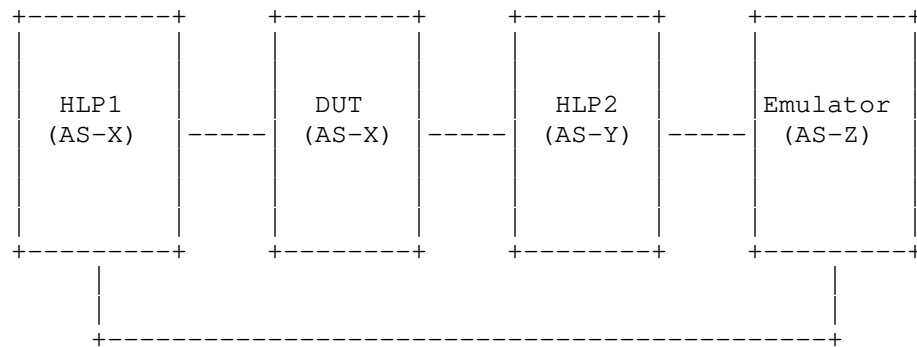


Figure 4 Four Node Setup for EBGP and IBGP Convergence

4. Test Considerations

The test cases for measuring convergence for iBGP and eBGP are different. Both iBGP and eBGP use different mechanisms to advertise, install and learn the routes. Typically, an iBGP route on the DUT is installed and exported when the next-hop is valid. For eBGP the

route is installed on the DUT with the remote interface address as the next-hop with the exception of the multihop case.

4.1. Number of Peers

Number of Peers is defined as the number of BGP neighbors or sessions the DUT has at the beginning of the test. The peers are established before the tests begin. The relationship could be either, iBGP or eBGP peering depending upon the test case requirement.

The DUT establishes one or more BGP sessions with one more emulated routers or helper nodes. Additional peers can be added based on the testing requirements. The number of peers enabled during the testing should be well documented in the report matrix.

4.2. Number of Routes per Peer

Number of Routes per Peer is defined as the number of routes advertized or learnt by the DUT per session or through neighbor relationship with an emulator or helper node. The tester, emulating as neighbor MUST advertise at least one route per peer.

Each test run must identify the route stream in terms of route packing, route mixture, and number of routes. This route stream must be well documented in the reporting stream. RFC 4098 defines these terms.

It is RECOMMENDED that the user may consider advertizing the entire current Internet routing table per peering session using an Internet route mixture with unique or non-unique routes. If multiple peers are used, it is important to precisely document the timing sequence between the peer sending routes (as defined in RFC 4098).

4.3. Policy Processing/Reconfiguration

The DUT MUST run one baseline test where policy is Minimal policy as defined in RFC 4098. Additional runs may be done with policy set-up before the tests begin. Exact policy settings should be documented as part of the test.

4.4. Configured Parameters (Timers, etc..)

There are configured parameters and timers that may impact the measured BGP convergence times.

The benchmark metrics MAY be measured at any fixed values for these configured parameters.

It is RECOMMENDED these configure parameters have the following settings: a) default values specified by the respective RFC b) platform-specific default parameters and c) values as expected in the operational network. All optional BGP settings MUST be kept consistent across iterations of any specific tests

Examples of the configured parameters that may impact measured BGP convergence time include, but are not limited to:

1. Interface failure detection timer
2. BGP Keepalive timer
3. BGP Holdtime
4. BGP update delay timer
5. ConnectRetry timer
6. TCP Segment Size
7. Minimum Route Advertisement Interval (MRAI)
8. MinASOriginationInterval (MAOI)
9. Route Flap Dampening parameters
10. TCP MD5
11. Maximum TCP Window Size
12. MTU

The basic-test settings for the parameters should be:

1. Interface failure detection timer (0 ms)
2. BGP Keepalive timer (1 min)
3. BGP Holdtime (3 min)
4. BGP update delay timer (0 s)

5. ConnectRetry timer (1 s)
6. TCP Segment Size (4096)
7. Minimum Route Advertisement Interval (MRAI) (0 s)
8. MinASOriginationInterval (MAOI) (0 s)
9. Route Flap Dampening parameters (off)
10. TCP MD5 (off)

4.5. Interface Types

The type of media dictate which test cases may be executed, each interface type has unique mechanism for detecting link failures and the speed at which that mechanism operates will influence the measurement results. All interfaces MUST be of the same media and throughput for each test case.

4.6. Measurement Accuracy

Since observed packet loss is used to measure the route convergence time, the time between two successive packets offered to each individual route is the highest possible accuracy of any packet-loss based measurement. When packet jitter is much less than the convergence time, it is a negligible source of error and hence it will be treated as within tolerance.

Other options to measure convergence are the Time-Based Loss Method (TBLM) and Timestamp Based Method (TBM) [MPLSProt]

An exterior measurement on the input media (such Ethernet) is defined by this specification.

4.7. Measurement Statistics

The benchmark measurements may vary for each trial, due to the statistical nature of timer expirations, CPU scheduling, etc. It is recommended to repeat the test multiple times. Evaluation of the test data must be done with an understanding of generally accepted testing practices regarding repeatability, variance and statistical significance of a small number of trials.

For any repeated tests that are averaged to remove variance, all parameters MUST remain the same.

4.8. Authentication

Authentication in BGP is done using the TCP MD5 Signature Option [RFC5925]. The processing of the MD5 hash, particularly in devices with a large number of BGP peers and a large amount of update traffic, can have an impact on the control plane of the device. If authentication is enabled, it SHOULD be documented correctly in the reporting format

4.9. Convergence Events

Convergence events or triggers are defined as abnormal occurrences in the network, which initiate route flapping in the network, and hence forces the re-convergence of a steady state network. In a real network, a series of convergence events may cause convergence latency operators desire to test.

These convergence events must be defined in terms of the sequences defined in RFC 4098. This basic document begins all tests with a router initial set-up. Additional documents will define BGP data plane convergence based on peer initialization.

The convergence events may or may not be tied to the actual failure A Soft Reset (RFC 4098) does not clear the RIB or FIB tables. A Hard reset clears the BGP peer sessions, the RIB tables, and FIB tables.

4.10. High Availability

Due to the different Non-Stop-Routing (sometimes referred to High-Availability) solutions available from different vendors, it is RECOMMENDED that any redundancy available in the routing processors should be disabled during the convergence measurements.

5. Test Cases

All tests defined under this section assume the following:

- a. BGP peers should be brought to BGP Peer established state
- b. Furthermore the traffic generation and routing should be verified in the topology

5.1. Basic Convergence Tests

These test cases measure characteristics of a BGP implementation in non-failure scenarios like:

1. RIB-IN Convergence
2. RIB-OUT Convergence
3. eBGP Convergence
4. iBGP Convergence

5.1.1. RIB-IN Convergence

Objective:

This test measures the convergence time taken to receive and install a route in RIB using BGP

Reference Test Setup:

This test uses the setup as shown in figure 1

Procedure:

- A. All variables affecting Convergence should be set to a basic test state (as defined in section 4-4).
- B. Establish BGP adjacency between DUT and peer x of Emulator.
- C. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test.
- D. Start the traffic from the Emulator peer-x towards the DUT targeted at a routes specified in route mixture (ex. route A) Initially no traffic SHOULD be observed on the egress interface as the route A is not installed in the forwarding database of the DUT.
- E. Advertise route A from the Peer-x to the DUT and record the time.

This is $Tup(EMx, Rt-A)$ also named 'XMT-Rt-time'.

- F. Record the time when the route-A from Peer-x is received at the DUT.

This $Tup(DUT, Rt-A)$ also named 'RCV-Rt-time'.

- G. Record the time when the traffic targeted towards route A is received by Emulator on appropriate traffic egress interface.

This is $TR(TDr, Rt-A)$. This is also named DUT-XMT-Data-Time.

- H. The difference between the $Tup(DUT, RT-A)$ and traffic received time ($TR(TDr, Rt-A)$) is the FIB Convergence Time for route-A in the route mixture. A full convergence for the route update is the measurement between the 1st route (Route-A) and the last route ($Rt-last$)

Route update convergence is

$TR(TDr, RT-last) - Tup(DUT, Rt-A)$ or

$(DUT-XMT-Data-Time - RCV-Rt-Time)(Rt-A)$

Note: It is recommended that a single test with the same route mixture be repeated several times. A report should provide the Standard Deviation of all tests and the Average.

Running tests with a varying number of routes and route mixtures is important to get a full characterization of a single peer.

5.1.2. RIB-OUT Convergence

Objective:

This test measures the convergence time taken by an implementation to receive, install and advertise a route using BGP

Reference Test Setup:

This test uses the setup as shown in figure 2

Procedure:

- A. The Helper node (HLP) run same version of BGP as DUT.

- B. All devices MUST be synchronized using NTP or some local reference clock.
- C. All configuration variables for HLP, DUT, and Emulator SHOULD be set to the same values. These values MAY be basic-test or a unique set completely described in the test set-up.
- D. Establish BGP adjacency between DUT and Emulator.
- E. Establish BGP adjacency between DUT and Helper Node.
- F. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- G. Start the traffic from the Emulator towards the Helper Node targeted at a specific route say route A. Initially no traffic SHOULD be observed on the egress interface as the route-A is not installed in the forwarding database of the DUT.
- H. Advertise routeA from the Emulator to the DUT and note the time.

This is $Tup(EMx, Route-A)$. (also named EM-XMT-Rt-Time)

- I. Record when Route-A is received by DUT.

This is $Tup(DUTr, Route-A)$. (also named DUT-RCV-Rt-Time)

- J. Record the time when the ROUTE is forwarded by DUT towards the Helper node.

This is $Tup(DUTx, Route-A)$. (also named DUT-XMT-Rt-Time)

- K. Record the time when the traffic targeted towards route-A is received on the Route Egress Interface. This is $TR(EMr, Route-A)$. (also named DUT-XMT-Data Time).

$FIB\ convergence = (DUT-RCV-Rt-Time - DUT-XMT-Data-Time)$

$RIB\ convergence = (DUT-RCV-Rt-Time - DUT-XMT-Rt-Time)$

Convergence for a route stream is characterized by

a) Individual route convergence for FIB, RIB

b) All route convergence of

FIB-convergence =DUT-RCV-Rt-Time(A) -DUT-XMT-Data-Time(last)

RIB-convergence =DUT-RCV-Rt-Time(A) -DUT-XMT-Rt-Time(last)

5.1.3. eBGP Convergence

Objective:

This test measures the convergence time taken by an implementation to receive, install and advertise a route in an eBGP Scenario

Reference Test Setup:

This test uses the setup as shown in figure 2 and the scenarios described in RIB-IN and RIB-OUT are applicable to this test case.

5.1.4. iBGP Convergence

Objective:

This test measures the convergence time taken by an implementation to receive, install and advertise a route in an iBGP Scenario

Reference Test Setup:

This test uses the setup as shown in figure 2 and the scenarios described in RIB-IN and RIB-OUT are applicable to this test case.

5.1.5. eBGP Multihop Convergence

Objective:

This test measures the convergence time taken by an implementation to receive, install and advertise a route in an eBGP Multihop Scenario

Reference Test Setup:

This test uses the setup as shown in figure 3. DUT is used along with a helper node.

Procedure:

- A. The Helper Node (HLP) runs the same BGP version as DUT
- B. All devices to be synchronized using NTP
- C. All variables affecting Convergence like authentication, policies, timers should be set to basic-settings
- D. All 3 devices, DUT, Emulator and Helper Node are configured with different Autonomous Systems
- E. Loopback Interfaces are configured on DUT and Helper Node and connectivity is established between them using any config options available on the DUT
- F. Establish BGP adjacency between DUT and Emulator
- G. Establish BGP adjacency between DUT and Helper Node
- H. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- I. Start the traffic from the Emulator towards the DUT targeted at a specific route say routeA
- J. Initially no traffic SHOULD be observed on the egress interface as the routeA is not installed in the forwarding database of the DUT
- K. Advertise routeA from the Emulator to the DUT and note the time (Tup(EMx,RouteA) also named (Route-Tx-time)
- L. Record the time when the route is received by the DUT. This is Tup(EMr,DUT) named (Route-Rcv-time)
- M. Record the time when the traffic targeted towards routeA is received from Egress Interface of DUT on emulator. This is Tup(EMd,DUT) named (Data-Rcv-time)
- N. Record the time when the routeA is forwarded by DUT towards the Helper node. This is Tup(EMf,DUT) also named (Route-Fwd-time)

FIB Convergence = (Data-Rcv-time - Route-Rcv-time)

RIB Convergence = (Route-Fwd-time - Route-Rcv-time)

Note: It is recommended that the test be repeated with varying number

of routes and route mixtures. With each set route mixture, the test should be repeated multiple times. The results should record average, mean, Standard Deviation

5.2. BGP Failure/Convergence Events

5.2.1. Physical Link Failure on DUT End

Objective:

This test measures the route convergence time due to local link failure event at DUT's Local Interface

Reference Test Setup:

This test uses the setup as shown in figure 1. Shutdown event is defined as an administrative shutdown event on the DUT

Procedure:

- A. All variables affecting Convergence like authentication, policies, timers should be set to basic-test policy
- B. Establish 2 BGP adjacencies from DUT to Emulator, one over the peer interface and the other using a second peer interface
- C. Advertise the same route, route A over both the adjacencies and (Tx1)Interface to be the preferred next hop
- D. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- E. Start the traffic from the Emulator towards the DUT targeted at a specific route say route A. Initially traffic would be observed on the best egress route (Emp1) instead of Trr2
- F. Trigger the shutdown event of Best Egress Interface on DUT (Drr1)
- G. Measure the Convergence Time for the event to be detected and traffic to be forwarded to Next-Best Egress Interface (rr2)

Time = Data-detect(rr2) - Shutdown time

H. Stop the offered load and wait for the queues to drain and Restart

I. Bring up the link on DUT Best Egress Interface

J. Measure the convergence time taken for the traffic to be rerouted from (rr2) to Best Interface (rr1)

Time = Data-detect(rr1) - Bring Up time

K. It is recommended that the test be repeated with varying number of routes and route mixtures or with number of routes & route mixtures closer to what is deployed in operational networks

5.2.2. Physical Link Failure on Remote/Emulator End

Objective:

This test measures the route convergence time due to local link failure event at Tester's Local Interface

Reference Test Setup:

This test uses the setup as shown in figure 1. Shutdown event is defined as shutdown of the local interface of Tester via logical shutdown event. The procedure used in 5.2.1 is used for the termination

5.2.3. ECMP Link Failure on DUT End

Objective:

This test measures the route convergence time due to local link failure event at ECMP Member. The FIB configuration and BGP is set to allow two ECMP routes to be installed. However, policy directs the routes to be sent only over one of the paths

Reference Test Setup:

This test uses the setup as shown in figure 1 and the procedure uses 5.2.1

5.3. BGP Adjacency Failure (Non-Physical Link Failure) on Emulator

Objective:

This test measures the route convergence time due to BGP Adjacency Failure on Emulator

Reference Test Setup:

This test uses the setup as shown in figure 1

Procedure:

- A. All variables affecting Convergence like authentication, policies, timers should be basic-policy set
- B. Establish 2 BGP adjacencies from DUT to Emulator, one over the Best Egress Interface and the other using the Next-Best Egress Interface
- C. Advertise the same route, routeA over both the adjacencies and make Best Egress Interface to be the preferred next hop
- D. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- E. Start the traffic from the Emulator towards the DUT targeted at a specific route say routeA. Initially traffic would be observed on the Best Egress interface
- F. Remove BGP adjacency via a software adjacency down on the Emulator on the Best Egress Interface. This time is called BGPadj-down-time also termed BGPpeer-down
- G. Measure the Convergence Time for the event to be detected and traffic to be forwarded to Next-Best Egress Interface. This time is Tr-rr2 also called TR2-traffic-on

$$\text{Convergence} = \text{TR2-traffic-on} - \text{BGPpeer-down}$$

- H. Stop the offered load and wait for the queues to drain and Restart
- I. Bring up BGP adjacency on the Emulator over the Best Egress Interface. This time is BGP-adj-up also called BGPpeer-up
- J. Measure the convergence time taken for the traffic to be rerouted to Best Interface. This time is BGP-adj-up also called BGPpeer-up

5.4. BGP Hard Reset Test Cases

5.4.1. BGP Non-Recovering Hard Reset Event on DUT

Objective:

This test measures the route convergence time due to Hard Reset on the DUT

Reference Test Setup:

This test uses the setup as shown in figure 1

Procedure:

- A. The requirement for this test case is that the Hard Reset Event should be non-recovering and should affect only the adjacency between DUT and Emulator on the Best Egress Interface
- B. All variables affecting SHOULD be set to basic-test values
- C. Establish 2 BGP adjacencies from DUT to Emulator, one over the Best Egress Interface and the other using the Next-Best Egress Interface
- D. Advertise the same route, routeA over both the adjacencies and make Best Egress Interface to be the preferred next hop
- E. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- F. Start the traffic from the Emulator towards the DUT targeted at a specific route say routeA. Initially traffic would be observed on the Best Egress interface
- G. Trigger the Hard Reset event of Best Egress Interface on DUT
- H. Measure the Convergence Time for the event to be detected and traffic to be forwarded to Next-Best Egress Interface

Time of convergence = time-traffic flow - time-reset

- I. Stop the offered load and wait for the queues to drain and Restart
- J. It is recommended that the test be repeated with varying number of routes and route mixtures or with number of routes & route mixtures closer to what is deployed in operational networks
- K. When varying number of routes are used, convergence Time is measured using the Loss Derived method [IGPData]
- L. Convergence Time in this scenario is influenced by Failure detection time on Tester, BGP Keep Alive Time and routing, forwarding table update time

5.5. BGP Soft Reset

Objective:

This test measures the route convergence time taken by an implementation to service a BGP Route Refresh message and advertise a route

Reference Test Setup:

This test uses the setup as shown in figure 2

Procedure:

- A. The BGP implementation on DUT & Helper Node needs to support BGP Route Refresh Capability [RFC2918]
- B. All devices to be synchronized using NTP
- C. All variables affecting Convergence like authentication, policies, timers should be set to basic-test defaults
- D. DUT and Helper Node are configured in the same Autonomous System whereas Emulator is configured under a different Autonomous System
- E. Establish BGP adjacency between DUT and Emulator
- F. Establish BGP adjacency between DUT and Helper Node

- G. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- H. Configure a policy under BGP on Helper Node to deny routes received from DUT
- I. Advertise routeA from the Emulator to the DUT
- J. The DUT will try to advertise the route to Helper Node will be denied
- K. Wait for 3 KeepAlives
- L. Start the traffic from the Emulator towards the Helper Node targeted at a specific route say routeA. Initially no traffic would be observed on the Egress interface, as routeA is not present
- M. Remove the policy on Helper Node and issue a Route Refresh request towards DUT. Note the timestamp of this event. This is the RefreshTime
- N. Record the time when the traffic targeted towards routeA is received on the Egress Interface. This is RecTime
- O. The following equation represents the Route Refresh Convergence Time per route

$$\text{Route Refresh Convergence Time} = (\text{RecTime} - \text{RefreshTime})$$

5.6. BGP Route Withdrawal Convergence Time

Objective:

This test measures the route convergence time taken by an implementation to service a BGP Withdraw message and advertise the withdraw

Reference Test Setup:

This test uses the setup as shown in figure 2

Procedure:

- A. This test consists of 2 steps to determine the Total Withdraw Processing Time
- B. Step 1:
- (1) All devices to be synchronized using NTP
 - (2) All variables should be set to basic-test parameters
 - (3) DUT and Helper Node are configured in the same Autonomous System whereas Emulator is configured under a different Autonomous System
 - (4) Establish BGP adjacency between DUT and Emulator
 - (5) To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
 - (6) Start the traffic from the Emulator towards the DUT targeted at a specific route say routeA. Initially no traffic would be observed on the Egress interface as the routeA is not present on DUT
 - (7) Advertise routeA from the Emulator to the DUT
 - (8) The traffic targeted towards routeA is received on the Egress Interface
 - (9) Now the Tester sends request to withdraw routeA to DUT, TRx(Awith) also called WdrawTime1
 - (10) Record the time when no traffic is observed on the Egress Interface. This is the RouteRemoveTime1(A)
$$\text{WdrawConvTime1} = \text{RouteRemoveTime1(A)}$$
 - (11) The difference between the RouteRemoveTime1 and WdrawTime1 is the WdrawConvTime1
- C. Step 2:
- (1) Continuing from Step 1, re-advertise routeA back to DUT from Tester

- (2) The DUT will try to advertise the routeA to Helper Node (assumption there exists a session between DUT and helper node)
- (3) Start the traffic from the Emulator towards the Helper Node targeted at a specific route say routeA. Traffic would be observed on the Egress interface after routeA is received by the Helper Node

WATime=time traffic first flows

- (4) Now the Tester sends a request to withdraw routeA to DUT. This is the WdrawTime2

WAWtime-TRx(RouteA) = WdrawTime2

- (5) DUT processes the withdraw and sends it to Helper Node
- (6) Record the time when no traffic is observed on the Egress Interface of Helper Node. This is

TR-WAW(DUT,RouteA) = RouteRemoveTime2

- (7) Total withdraw processing time is

TotalWdrawTime = ((RouteRemoveTime2 - WdrawTime2) - WdrawConvTime1)

5.7. BGP Path Attribute Change Convergence Time

Objective:

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

Reference Test Setup:

This test uses the setup as shown in figure 1

Procedure:

- A. This test only applies to Well-Known Mandatory Attributes like Origin, AS Path, Next Hop
- B. In each iteration of test only one of these mandatory attributes need to be varied whereas the others remain the

same

- C. All devices to be synchronized using NTP
- D. All variables should be set to basic-test parameters
- E. Advertise the route, routeA over the Best Egress Interface only, making it the preferred named Tbest
- F. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- G. Start the traffic from the Emulator towards the DUT targeted at the specific route say routeA. Initially traffic would be observed on the Best Egress interface
- H. Now advertise the same route routeA on the Next-Best Egress Interface but by varying one of the well-known mandatory attributes to have a preferred value over that interface. We call this Tbetter. The other values need to be same as what was advertised on the Best-Egress adjacency

$$TRx(\text{Path-Change}) = \text{Path Change Event Time}$$

- I. Measure the Convergence Time for the event to be detected and traffic to be forwarded to Next-Best Egress Interface

$$DUT(\text{Path-Change}, \text{RouteA}) = \text{Path-switch time}$$
$$\text{Convergence} = \text{Path-switch time} - \text{Path Change Event Time}$$

- J. Stop the offered load and wait for the queues to drain and Restart
- K. Repeat the test for various attributes

5.8. BGP Graceful Restart Convergence Time

Objective:

This test measures the route convergence time taken by an implementation during a Graceful Restart Event

Reference Test Setup:

This test uses the setup as shown in figure 4

Procedure:

- A. It measures the time taken by an implementation to service a BGP Graceful Restart Event and advertise a route
- B. The Helper Nodes are the same model as DUT and run the same BGP implementation as DUT
- C. The BGP implementation on DUT & Helper Node needs to support BGP Graceful Restart Mechanism [RFC4724]
- D. All devices to be synchronized using NTP
- E. All variables are set to basic-test values
- F. DUT and Helper Node-1(HLP1) are configured in the same Autonomous System whereas Emulator and Helper Node-2(HLP2) are configured under different Autonomous Systems
- G. Establish BGP adjacency between DUT and Helper Nodes
- H. Establish BGP adjacency between Helper Node-2 and Emulator
- I. To ensure adjacency establishment, wait for 3 KeepAlives from the DUT or a configurable delay before proceeding with the rest of the test
- J. Configure a policy under BGP on Helper Node-1 to deny routes received from DUT
- K. Advertise routeA from the Emulator to Helper Node-2
- L. Helper Node-2 advertises the route to DUT and DUT will try to advertise the route to Helper Node-1 which will be denied
- M. Wait for 3 KeepAlives
- N. Start the traffic from the Emulator towards the Helper Node-1 targeted at the specific route say routeA. Initially no traffic would be observed on the Egress interface as the routeA is not present
- O. Perform a Graceful Restart Trigger Event on DUT and note the time. This is the GREventTime

- P. Remove the policy on Helper Node-1
- Q. Record the time when the traffic targeted towards routeA is received on the Egress Interface

TRr(DUT, routeA). This is also called RecTime
- R. The following equation represents the Graceful Restart Convergence Time

$$\text{Graceful Restart Convergence Time} = (\text{RecTime} - \text{GREventTime}) - \text{RIB-IN}$$
- S. It is assumed in this test case that after a Switchover is triggered on the DUT, it will not have any cycles to process BGP Refresh messages. The reason for this assumption is that there is a narrow window of time where after switchover when we remove the policy from Helper Node -1, implementations might generate Route-Refresh automatically and this request might be serviced before the DUT actually switches over and reestablishes BGP adjacencies with the peers

6. Reporting Format

For each test case, it is recommended that the reporting tables below are completed and all time values SHOULD be reported with resolution as specified in [RFC4098]

Parameter	Units
Test case	Test case number
Test topology	1,2,3 or 4
Parallel links	Number of parallel links
Interface type	GigE, POS, ATM, other
Convergence Event	Hard reset, Soft reset, link failure, or other defined
eBGP sessions	Number of eBGP sessions
iBGP sessions	Number of iBGP sessions
eBGP neighbor	Number of eBGP neighbors
iBGP neighbor	Number of iBGP neighbors
Routes per peer	Number of routes
Total unique routes	Number of routes
Total non-unique routes	Number of routes
IGP configured	ISIS, OSPF, static, or other
Route Mixture	Description of Route mixture
Route Packing	Number of routes in an update
Policy configured	Yes, No
Packet size offered to the DUT	Bytes
Offered load	Packets per second
Packet sampling interval on tester	Seconds
Forwarding delay threshold	Seconds
Timer Values configured on DUT	
Interface failure indication delay	Seconds
Hold time	Seconds
MinRouteAdvertisementInterval (MRAI)	Seconds
MinASOriginationInterval (MAOI)	Seconds
Keepalive Time	Seconds
ConnectRetry	Seconds
TCP Parameters for DUT and tester	
MSS	Bytes
Slow start threshold	Bytes
Maximum window size	Bytes

Test Details:

- a. If the Offered Load matches a subset of routes, describe how this subset is selected
- b. Describe how the Convergence Event is applied; does it cause instantaneous traffic loss or not

c. If there is any policy configured, describe the configured policy

Complete the table below for the initial Convergence Event and the reversion Convergence Event

Parameter	Unit
Convergence Event	Initial or reversion
Traffic Forwarding Metrics	
Total number of packets offered to DUT	Number of packets
Total number of packets forwarded by DUT	Number of packets
Connectivity Packet Loss	Number of packets
Convergence Packet Loss	Number of packets
Out-of-order packets	Number of packets
Duplicate packets	Number of packets
Convergence Benchmarks	
Rate-derived Method [IGP-Data]:	
First route convergence time	Seconds
Full convergence time	Seconds
Loss-derived Method [IGP-Data]:	
Loss-derived convergence time	Seconds
Route-Specific Loss-Derived Method:	
Minimum R-S convergence time	Seconds
Maximum R-S convergence time	Seconds
Median R-S convergence time	Seconds
Average R-S convergence time	Seconds
Loss of Connectivity Benchmarks	
Loss-derived Method:	
Loss-derived loss of connectivity period	Seconds
Route-Specific loss-derived Method:	
Minimum LoC period [n]	Array of seconds
Minimum Route LoC period	Seconds
Maximum Route LoC period	Seconds
Median Route LoC period	Seconds

Average Route LoC period Seconds

7. IANA Considerations

This draft does not require any new allocations by IANA.

8. Security Considerations

Benchmarking activities as described in this memo are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the constraints specified in the sections above.

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 DUT/SUT.

Special capabilities SHOULD NOT exist in the DUT/SUT specifically for benchmarking purposes. Any implications for network security arising from the DUT/SUT SHOULD be identical in the lab and in production networks.

9. References

9.1. Normative References

- [I-D.ietf-bmwg-igp-dataplane-conv-term]
Poretsky, S., Imhoff, B., and K. Michielsen, "Terminology for Benchmarking Link-State IGP Data Plane Route Convergence", draft-ietf-bmwg-igp-dataplane-conv-term-23 (work in progress), February 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2918] Chen, E., "Route Refresh Capability for BGP-4", RFC 2918, September 2000.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

9.2. Informative References

- [RFC1242] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [RFC1983] Malkin, G., "Internet Users' Glossary", RFC 1983, August 1996.
- [RFC2285] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.
- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, March 1999.
- [RFC4098] Berkowitz, H., Davies, E., Hares, S., Krishnaswamy, P., and M. Lepp, "Terminology for Benchmarking BGP Device Convergence in the Control Plane", RFC 4098, June 2005.
- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", RFC 4724, January 2007.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

Authors' Addresses

Rajiv Papneja
Huawei Technologies

Email: rajiv.papneja@huawei.com

Bhavani Parise
Cisco Systems

Email: bhavani@cisco.com

Susan Hares
Huawei Technologies (USA)

Email: shares@huawei.com

Ilya Varlashkin
Easynet Global Services

Email: ilya.varlashkin@easynet.com

Dean Lee
Ixia

Email: dlee@ixiacom.com

Eric Brendel
Independent Consultant

Email: brendel@pektel.com

Mohan Nanduri
Microsoft

Email: mnanduri@microsoft.com

Jay Karthik
Cisco Systems

Email: jkarthik@cisco.com

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 22, 2012

I. Varlashkin
Easynet Global Services
R. Papneja
Huawei Technologies (USA)
B. Parise
Cisco
T. Van Unen
Ixia
October 20, 2011

Convergence benchmarking on contemporary routers
draft-varlashkin-router-conv-bench-00

Abstract

This document specifies methodology for benchmarking convergence of routers without making assumptions about relation and dependencies between data- and control-planes. Provided methodology is primary intended for testing routers running BGP and some form of link-state IGP with or without MPLS. It may also be applicable for environments using MPLS-TE or GRE, however they're beyond scope of this document and such application is left for further study.

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 <http://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 April 22, 2012.

Copyright Notice

Copyright (c) 2011 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

(<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	4
2. Test topology	5
3. TEST PARAMETERS	6
3.1. Packing ratios	7
3.2. Test traffic	7
3.3. IGP metrics	7
3.4. Internal routers matrix	7
3.5. Number of next-hops	8
3.6. 'e' - Failure and Restoration start entropy	8
4. TEST PROCEDURES	8
4.1. Initialisation time	8
4.2. Generic data-plane failure test	9
4.3. Generic test procedure for	9
5. Failure and restoration scenarios	10
5.1. Loss of Signal on the link attached to DUT	10
5.2. Link failure without LoS	10
5.3. Non-direct link failure	11
5.4. Best route withdrawal	11
5.5. iBGP next-hop failure	12
6. Test report	12
7. Link bundling and Equal Cost Multi-Path	13
8. Graceful Restart and Non-Stop Forwarding	13
9. Security considerations	13
10. IANA Considerations	14
11. Acknowledgments	14
12. Normative References	14
Authors' Addresses	14

1. Introduction

Ability of the network to restore traffic flow when primary path fails has always been important subject for network engineers, researchers and equipment manufacturers. Time to recover from a link or node failure has often been linked to routing protocols convergence; and benchmarking of a routing protocol convergence has often been considered sufficient for quantifying recovery performance. As long as routers could obtain new best path only after relevant routing protocols perform their calculations such methodology was reasonable. However continuous improvements in hardware and software result in more and more routers being able to restore traffic flow even before routing protocols converge. Methodology described in this document takes such fact into account.

When a failure occurs on the network a router needs to:

1. select new best path so that the packets, which already arrived to the router, can be forwarded
2. let other routers know about new network state so they can find new best path from their perspective

How fast a router can perform these two functions characterise router's performance with regards to convergence. Note that in general case each of these characteristics may or may not be related to the other. For example, some platform may need to perform calculations to find new best path and only then update local FIB and send relevant protocol updates to other routers, another platform can update local FIB without waiting for calculations to complete but still needs to wait for calculations before sending routing protocol updates, third platform can use different optimisation for both FIB changes and routing protocol updates without waiting for completion of the calculations. Other variations are also possible. This document makes no assumption about whether local FIB changes and routing protocol updates dependencies on each other or on routing protocol calculations.


Since it is not known whether local FIB is updated before or after routing protocol calculations, forwarding-plane method is proposed to benchmark local convergence. And because it is not known whether routing protocol updates are linked to FIB modification or not the control-plane approach is used to benchmark how fast updates are propagated. However both characteristics are benchmarked using very similar test topologies and procedures. Also, an attempt is made to to minimise dependency on performance on non-DUT elements involved in the tests.

At the time of writing of this document it is not known whether existing network testers and protocol emulators are able to execute described tests out of the box. Nevertheless the authors believe that required functionality can be added with reasonable effort. Alternatively the tests can be performed with help of physical routers to create necessary test topology, which may have impact on time required to perform the test but expected to provide same degree of the test results accuracy. This also means that tests performed using a protocol simulator can be repeated using physical routers and results expected to be comparable.

This document complements draft-papneja-bgp-basic-dp-convergence.

2. Test topology

Unless specified otherwise all tests use same basic test topology outlined below:



S is source of test traffic for data-plane tests, while for control-plane tests S is an emulated or physical router with packet capturing (sniffing) capability.

Unidirectional test traffic goes from Source to NetA.


IGP between DUT and R1-R4; BGP between DUT and R3, R4; no BGP between R3 and R4 (important). If tunnelling (e.g. MPLS or GRE) is used then R1 and R2 do not need to run BGP, otherwise they MUST run BGP. Source has static default to DUT; R3 and R4 have static to NetA. NetA is in BGP but not in IGP. M1 is K*M matrix of internal routers. Metrics C1 is used to control whether R2 is LFA for DUT to NetA. Metric C2 is used to control whether R3 or R4 are best exit towards NetA. All other metrics are fixed for all tests and MUST be set to

exact values provided in the above diagram. IGP metrics from M1 to ER1 throughout ERn can be set arbitrarily, their exact values are irrelevant to this test as long as they're valid for given IGP.

Routers ER1 throughout ERn together with prefixes NetB-1 throughout NetB-N are presented to create realistic environment but not used directly in measurements. NetB-1 throughout NetB-N are distinct single-prefix sets.

Traffic restoration depends on ability of R2 and M1 to forward traffic after failure. To eliminate this dependency R2 is set to always forward traffic to R3 and NetA via M1 which in turn always forwards traffic directly via R3 or R2 depending on the test. One possibility to achieve this is to use static routes. Another alternative is to use different IGP between R2 and R3 from the one used by DUT and make routes learned via this IGP preferred on R2. E.g. DUT uses OSPF, then in addition to it R2&R3 also run ISIS and prefer ISIS routes over OSPF ones. A protocol simulator can have internal mechanism to provide required behaviour. There are no other dependencies on non-DUT devices in this tests.

For evaluating eBGP performance following topology is used:



Test topology for eBGP

In "Link failure without LoS" test direct cable between DUT and R1 is replaced with connection over an L2 switch as follow:

[DUT]---[SW1]---[R1]

3. TEST PARAMETERS

3.1. Packing ratios

Routes with different prefixes but same attributes can potentially be packed into single update message. Since both number of update messages and number of prefixes per update can affect convergence time, the tests SHOULD be performed with various prefix packing ratios. This document does not specify values of individual BGP attributes used to control packing ratio.

3.2. Test traffic

Traffic is sent from single source address located at the Source port of the tester to one address in each prefix in NetA set. Packets are sent at rate 1000 per second, which provides 1ms resolution of the convergence time as measured by tests in this document. All packets SHOULD be 64 bytes at IP layer, that is IP header plus IP payload.

3.3. IGP metrics

Basic test topology specifies fixed IGP metrics for some links. These metrics SHOULD be used verbatim. There are also two variable metrics - C1 and C2 - intended for controlling whether R2 is Loop-Free-Alternate (LFA) for DUT towards NetA, and whether R3 remains best exit towards NetA after path failure between DUT and R3. Following values SHOULD be used for C1 and C2 depending on required behaviour:

R2 is LFA?	R3 best?	C1	C2
yes	yes	1	1
yes	no	1	3
no	yes	5	1
no	no	5	3

3.4. Internal routers matrix

Basic test topology has N*K grid of internal routers denoted as M1. When N>1 or K>1 the cost of all links within grid MUST be set to 1 (one). This matrix is intended for controlling topology size, which has affect on particularly SPF run-time.

If traffic is forwarded using a tunneling mechanism, such as MPLS or GRE, the internal routers only need to have reachability information about tunnel end-points. However if traditional hop-by-hop forwarding is used, then internal routers MUST have routes to each and every prefix within NetA set.

This document does not specify how internal routers should obtain necessary reachability information. The only requirement is that after primary DUT-NetA path failure internal routers are able to forward traffic to NetA instantly. Using values of IGP metrics as described earlier addresses this requirement. Also, protocol simulator may have built-in mechanism to achieve desired behaviour.

3.5. Number of next-hops

Basic test topology has set of N edge routers ER1 throughout ERn, each advertising unique prefix. Some BGP implementations may exhibit different performance depending on number of next-hops for which IGP cost has changed after failure. By varying overall number of next-hops such dependency can be detected.

Note that prefixes NetB-1 throughout NetB-n are not used as destinations for test traffic, they're only present for creating "background environment".

3.6. 'e' - Failure and Restoration start entropy

Tests described in this document use fixed time T2 and variable offset 'e' as starting point for simulating failure or restoration event.

Fixing time T2 is necessary as reference point to which variable offset e is added for each iteration of the test. Introduction of such variable offset allows better analysis of the test results. For example, DUT may run FIB changes at certain intervals. If failure introduced close to the end of such interval, shorter outage will be observed, and if introduced close to the beginning of such interval longer outage will be observed. Running test multiple times each time using different offset will help to profile DUT better.

Test report must contain value of T2 (same for all iterations) and values of e for each iterations. This document recommends to use $T2=T1+8s$ and e from 0 to 1s in 0.01s (10ms) increments.

4. TEST PROCEDURES

This section provides generic steps that are used in all tests.

4.1. Initialisation time

The objective of this test is to measure time that must elapse between starting protocols and ability of the test topology to forward traffic. This test is not intended to reflect DUT

performance but used only as a way to find time T1 that is used in all subsequent tests.

To execute test perform following steps:

1. Configure DUT and protocol simulator (or auxiliary nodes)
2. At T0 start traffic and then immediately start routing protocols
3. When traffic starts arriving Sink Port 1 stop test.

The time of arrival of the first packet is T1.

4.2. Generic data-plane failure test

The purpose of failure test is to measure time required by DUT to resume traffic flow after best path to destination fails. Following steps are common for all failure tests:

1. Start protocols and mark time as T0
2. At time T1 start traffic to each prefix in set NetA
3. At T2+e simulate failure or restoration event (see Section 5)
4. From T2+e until T3 packets do not arrive to NetA
5. After packets are seen again at NetA (T3) wait until time T4
6. Stop traffic
7. Measure total number of lost packets and calculate outage knowing packet-per-second

4.3. Generic test procedure for

1. At T0 bring up all interfaces and protocols, and start capturing BGP packets at RS1
2. At T1+e simulate failure/restoration event (see Section 5)
3. At T2-d1 first UPDATE message is sent by DUT and at T2 it will be observed at RS1
4. At T3-d2 last UPDATE message is sent by DUT and at T3 it will be observed at RS1

d1 and d2 represent serialisation and propagation delay and can be

disregarded unless DUT-RS1 link has large delay. With this in mind, T2-(T1+e) and T3-(T1+e) represent convergence time for the first and last prefix respectively.

5. Failure and restoration scenarios

This section defines set of various failure and restoration scenarios used in step 3 of the generic test procedures described in previous section. Unless otherwise specified all scenarios are applicable to both data- and control-plane test procedures.

5.1. Loss of Signal on the link attached to DUT

This scenario simulates situation where link attached to DUT fails and Loss of Signal (LoS) can be observed by DUT. In other words link fails and results in interface on the DUT going down.

To simulate LoS failure at the time defined by the test procedure shut down R1 side of the link to DUT.

To simulate LoS restoration at the time defined by the test procedure re-activate R1 side of the link to DUT.

5.2. Link failure without LoS

This scenario simulates situation where link between DUT and adjacent node fails but DUT does not observe LoS. In practice such failure can occur when, for example, link between DUT and adjacent node is implemented via carrier equipment that does not shut link down when remote side of the link fails.

DUT can use various methods to detect such failures, including but not limited to protocol HELLO or Keep-alive packets, BFD, OAM. This document does not restrict methods which DUT can use, but requires use of particular method to be recorded in the test report.

Basic network topology is modified for the purpose of this test only as follow: rather than using direct cabling between DUT and R1 the link is implemented via intermediate L2 switch that supports concept of VLAN's. Initially switch ports connected to DUT and R1 are placed into the same VLAN (same L2 broadcast domain).

To simulate failure at the time defined by the test procedure move switch port connected to R1 to a VLAN different from the one used for switch port connected to DUT.

To simulate restoration at the time defined by the test procedure

move switch port connected to R1 back to the same VLAN as the one used for switch port connected to DUT.

5.3. Non-direct link failure

This scenario simulates situation where a link not directly connected to DUT but located on the primary path to destination fails. Unmodified basic network topology is used.

Depending on technologies used in the setup different failure detection techniques can be employed by DUT. This document assumes that DUT relies exclusively on IGP information to learn about failure and that nodes adjacent to the failed link flood this information within D seconds since the event. If required exact value of D can be obtained through simple additional test, but in this document D is assumed to be 0 (zero).

It is possible, though undesirable, that some traffic and protocol simulators may continue accepting packets coming through the port that leads to simulated failed link. It is essential to assert such behaviour prior to the tests and if confirmed, exclude packets received after failure from calculations in step 7 of the test.

Failure event is triggered by simulating shutdown of R3 side of the link to R1 at the time defined by the test procedure. R1 MUST send IGP update (depending on which protocol is used) to DUT within D seconds.

Restoration event is triggered by simulating recovery of R3 side of the link to R1 at the time defined by the test procedure. R1 MUST send IGP update (depending on which protocol is used) to DUT within D seconds.

5.4. Best route withdrawal

This scenario simulates situation where best AS exit path to a destination is no longer valid and ASBR sends BGP UPDATE to its iBGP peers. Unmodified basic network topology is used.

Disconnecting R3 from NetA implies that R3 will send BGP WITHDRAW for this prefixes in its update to DUT. It is possible, though undesirable, that some protocol simulator and traffic generators will still count packets received at sink port 1 even after prefixes were withdrawn. To correctly execute this test it's mandatory that traffic received at sink port 1 after withdrawing prefixes is ignored and not counted as delivered. If traffic generator is not able to assure such functionality (should be asserted prior to the test), then packets received at the sink port 1 MUST be excluded from

calculation in step 7 of the test.

Failure event is triggered by simulating failure of the link between R3 and NetA and immediate withdrawal of all corresponding prefixes by R3.

Restoration event is triggered by simulating recovery of the link between R3 and NetA and immediate BGP UPDATE for all corresponding prefixes by R3.

5.5. iBGP next-hop failure

This scenario simulates situation where ASBR used as best exit to a destination unexpectedly fails both at control and forwarding plane. Both R1 and a router within M1 connected to R3 MUST send appropriate IGP update message to the rest of the network within D seconds. To detect failure DUT MAY rely on IGP information provided by rest of the network or it MAY employ additional techniques. This document does not restrict what detection mechanism should DUT use but requires that particular mechanism is recorded in the test report.

Failure event is triggered by simulating removal of R3 from the test topology at the time defined by the test procedure, followed by IGP update as described in previous paragraph.

Recovery event is triggered by re-introducing R3 into the test topology, followed by IGP update as described in first paragraph of this section and immediate re-activation of BGP session between R3 and DUT. Note that recovery time calculated by this method depends on DUT performance in respect to bringing up new BGP session. This is intentional. Control plane convergence benchmarking can be performed separately by a method that is outside of the scope of this document and two results can be correlated netto data-plane convergence value should that be necessary.

6. Test report

TODO: Report format is to be discussed.

Test report MUST contain following data for each test:

1. T1 and 'e'
2. Number of prefixes NetA and NetB
3. Size of M1 (recored as N*K)

4. Traffic rate, in packets per second, and packet size at IP layer in octets
5. Number of lost packets during failure, and number of lost packets during restoration

7. Link bundling and Equal Cost Multi-Path

Scenarios where DUT can balance traffic to NetA across multiple best paths is explicitly excluded from scope of this document. There are two reasons.

First, two different DUT may choose different path (out of all equal) to forward given packet, which makes it unreasonably difficult to define generic traffic that would produce comparable results when testing different platforms.

Second, mechanisms used to handle failures in ECMP (but not necessarily in link-bundling) environment are similar to those handling single-path failures. Therefore it's expected that convergence in ECMP scenario will be of the same order as in single-path scenario.

8. Graceful Restart and Non-Stop Forwarding

While Graceful Restart and Non-Stop Forwarding mechanisms are related to DUT ability to forward traffic under certain failure conditions, the test covering DUT own ability to restore or preserve traffic flow already covered in RFC6201.

9. Security considerations

The tests described in this document intended to be performed in isolated lab environment, which inherently has no security implication on the live network of the organisation or Internet as whole.

Authors foresee that some people or organisations might be interested to benchmark performance of the live networks. The tests described in this document are disruptive by their nature and will have impact at least on the network where they're executed, and depending on the role of that network effect can extend to other parts of the Internet. Such tests MUST NOT be attempted in live environment without careful consideration.

The fact of publishing this document does not increase potential negative consequences if tests are executed in live environment because information provided here is mere recording of widely known and used techniques.

10. IANA Considerations

None.

11. Acknowledgments

Authors would like to thank Gregory Cauchie, Rob Shakir, David Freedman, Anton Elita, Saku Ytti, Andrew Yourtchenko, for their valuable contribution and peer-review of this work.

12. Normative References

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter,
"Multiprotocol Extensions for BGP-4", RFC 4760,
January 2007.

Authors' Addresses

Ilya Varlashkin
Easynet Global Services

Email: ilya.varlashkin@easynet.com

Rajiv Papneja
Huawei Technologies (USA)

Email: rajiv.papneja@huawei.com

Bhavani Parise
Cisco

Email: bhavani@cisco.com

Tara Van Unen
Ixia

Email: TVanUnen@ixiacom.com

