

Internet Engineering Task Force (IETF)
Request for Comments: 6815
Updates: 2544
Category: Informational
ISSN: 2070-1721

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

Applicability Statement for RFC 2544:
Use on Production Networks Considered Harmful

Abstract

The 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 it encourages new standards activity for measurement methods applicable outside that scope.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6815>.

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	4
2. Scope and Goals	4
3. The Concept of an Isolated Test Environment	4
4. Why the Methods of RFC 2544 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. Acknowledgements	7
9. References	8
9.1. Normative References	8
9.2. Informative References	8
Appendix A. Example of RFC 2544 Method Failure in Production Network Measurement	9

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 [RFC2544] has been held up as the standard reference for the testing listed above, 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, the 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 of providing 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 the methods of [RFC2544] (as illustrated in Figures 1 through 3 of [RFC2544]) has the ability to:

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

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 the Methods of RFC 2544 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 they may cause failure modes in the device under test. Since failures can become the root cause of more widespread 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 Appendix A 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 contraindicated 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 Operations, Administration, and Maintenance (OAM). 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. 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 coauthors, 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.

9. References

9.1. 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.

9.2. Informative References

- [Bryant] Bonica, R. and S. Bryant, "RFC2544 Testing in Production Network", Work in Progress, October 2012.
- [Y.1564] ITU-T Recommendation Y.1564, "Ethernet Service Activation Test Methodology", March 2011.

Appendix A. 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, crossover 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 kilometers 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.

Authors' Addresses

Scott Bradner
Harvard University
1350 Mass. Ave., Room 760
Cambridge, MA 02138
USA

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

Kevin Dubray
Juniper Networks

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

Phone: +1 919-619-3220
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/>

