

Network Working Group	S. Bradner
Internet-Draft	Harvard University
Intended status: Informational	K. Dubray
Expires: April 22, 2012	Juniper Networks
	J. McQuaid
	Turnip Video
	A. Morton
	AT&T Labs
	October 20, 2011

RFC 2544 Applicability Statement: Use on Production Networks Considered Harmful

draft-ietf-bmwg-2544-as-01

Abstract

Benchmarking Methodology Working Group (BMWG) has been developing key performance metrics and laboratory test methods since 1990, and continues this work at present. Recent application of the methods beyond their intended scope is cause for concern. This memo clarifies the scope of RFC 2544 and other benchmarking work for the IETF community.

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.

Table of Contents

- *1. [Introduction](#)
- *1.1. [Requirements Language](#)
- *2. [Scope and Goals](#)
- *3. [The Concept of an Isolated Test Environment](#)
- *4. [Why RFC 2544 Methods are intended for ITE](#)
- *4.1. [Experimental Control, Repeatability, and Accuracy](#)
- *4.2. [Containment of Implementation Failure Impact](#)
- *5. [Advisory on RFC 2544 Methods in Real-world Networks](#)
- *6. [What to do without RFC 2544?](#)
- *7. [Security Considerations](#)
- *8. [IANA Considerations](#)
- *9. [Acknowledgements](#)
- *10. [References](#)
- *10.1. [Normative References](#)
- *10.2. [Informative References](#)
- *[Authors' Addresses](#)

1. Introduction

This memo clarifies the scope of RFC 2544 [\[RFC2544\]](#), and other benchmarking work for the IETF community.

Benchmarking Methodologies (beginning with [\[RFC2544\]](#)) have always relied on test conditions that can only be reliably produced in the laboratory. Thus it was surprising to find that this foundation methodology was being cited in several unintended applications, such as:

[\[RFC2544\]](#) used the term "real-world" to refer production networks and to differentiate them from test networks.

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. As an integral part of 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 in the real world) from the generic term "service". Also, we use the adjective "production" to refer to networks carrying live user traffic.

Although RFC 2544 is held up as the standard reference for such testing, we believe that the actual methods used vary from RFC 2544 in significant ways. Since the only citation is to RFC 2544, the modifications are opaque to the standards community and to users in general (an undesirable situation).

To directly address this situation, the past and present Chairs of the IETF Benchmarking Methodology Working Group (BMWG) have prepared this Applicability Statement for RFC 2544.

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 community 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:

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

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

4. Why RFC 2544 Methods are intended for ITE

The following sections discuss some of the reasons why RFC 2544 [\[RFC2544\]](#) methods were intended only for isolated laboratory use, and the difficulties of applying these methods outside the lab environment.

4.1. Experimental Control, Repeatability, and Accuracy

All of the tests described in RFC 2544 assume that the tester and device under test are the only devices on the networks that are transmitting data. The presence of other unwanted traffic on the network would mean that the specified test conditions have not been achieved.

Assuming that the unwanted traffic appears in variable amounts over time, the repeatability of any test result will likely depend to some degree on the unwanted traffic.

The presence of unwanted or unknown traffic makes accurate, repeatable, and consistent measurements of the performance of the device under test very unlikely, since the actual 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. Containment of Implementation Failure Impact

RFC 2544 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 DUT traffic within the ITE.

In addition, such testing can have a negative affect on any traffic which 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, testers using the assigned testing address ranges MUST NOT be connected to the Internet.

We note that a range of IPv6 addresses has been assigned to BMWG for laboratory test purposes, in [\[RFC5180\]](#). Also, the strong statements in the Security Considerations Section of this memo make the scope even more clear; this is now a standard fixture of all BMWG memos.

5. Advisory on RFC 2544 Methods in Real-world 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 unanticipated difficulties when applying these methods outside the lab environment. Operating test equipment on production networks according to the methods described in [\[RFC2544\]](#), where overload is a possible outcome, would no doubt be harmful to user traffic performance. These tests **MUST NOT** be used on active networks. And as discussed above, the tests will never produce a reliable or accurate benchmarking result.

[\[RFC2544\]](#) 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 to be 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. What to do without RFC 2544?

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 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.

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), the ITU-T Study Group 12 has a recommendation on service activation test methodology.

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 is also intended 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 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, and hopes that IANA will leave it alone as well.

9. Acknowledgements

Thanks to Matt Zekauskas, Bill Cervený, Barry Constantine, and Curtis Villamizar for reading and suggesting improvements for this memo.

10. References

10.1. Normative References

[RFC2544]	Bradner, S. and J. McQuaid , " Benchmarking Methodology for Network Interconnect Devices ", RFC 2544, March 1999.
[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.
[RFC2679]	Almes, G. , Kalidindi, S. and M. Zekauskas , " A One-way Delay Metric for IPPM ", RFC 2679, September 1999.
[RFC5180]	Popoviciu, C. , Hamza, A. , Van de Velde, G. and D. Dugatkin , " IPv6 Benchmarking Methodology for Network Interconnect Devices ", RFC 5180, May 2008.

10.2. Informative References

Authors' Addresses

Scott Bradner Bradner Harvard University 29 Oxford St.
Cambridge, MA 02138 USA Phone: +1 617 495 3864 EMail:
sob@harvard.edu URI: <http://www.sobco.com>

Kevin Dubray Dubray Juniper Networks EMail: kdubray@juniper.net

Jim McQuaid 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 Morton AT&T Labs 200 Laurel Avenue South Middletown,, NJ
07748 USA Phone: +1 732 420 1571 EMail: acmorton@att.com URI:
<http://home.comcast.net/~acmacm/>