

Internet Engineering Task Force (IETF)
Request for Comments: 6985
Category: Informational
ISSN: 2070-1721

A. Morton
AT&T Labs
July 2013

IMIX Genome: Specification of Variable Packet Sizes for Additional Testing

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 a constant packet size reveal device capabilities but differ significantly from the conditions encountered in operational deployment, so additional tests are sometimes conducted with a mixture of packet sizes, or "IMIX" ("Internet Mix"). 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 document defines a way to specify the exact repeating sequence of packet sizes from the usual set of fixed sizes and from other forms of mixed-size specification.

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/rfc6985>.

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	2
2. Requirements Language	3
3. Scope and Goals	3
4. Specification of the IMIX Genome	4
5. Specification of a Custom IMIX	6
6. Reporting Long or Pseudorandom Packet Sequences	7
6.1. Run-Length Encoding	7
6.2. Table of Proportions	7
6.3. Deterministic Algorithm	7
6.4. Pseudorandom Length Algorithm	8
6.5. Pseudorandom Sequence Algorithm	8
7. Security Considerations	8
8. Acknowledgements	8
9. References	9
9.1. Normative References	9
9.2. Informative References	9

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, 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. 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].

3. 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 -- and certainly no ability to extrapolate results to other circumstances -- even when the mix varies only slightly from another IMIX.

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 pseudorandom 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 mailing 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.

4. 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, and 16000 bytes). Note that the 4-octet Ethernet frame check sequence may fail to detect bit errors in the larger jumbo frames [Jumbo1] [Jumbo2].

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 Packet sizes not defined by RFC 2544 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).

5. 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

6. Reporting Long or Pseudorandom Packet Sequences

When the IMIX Genome cannot be used (when the sheer length of the sequence would make the genome unmanageable), five options are possible, as noted in the following subsections.

6.1. Run-Length Encoding

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 4):

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

The run-length encoding approach is also applicable to the custom IMIX as described in Section 5 (where the single upper-case letter Genome Codes would be used instead).

6.2. Table of Proportions

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.

6.3. Deterministic Algorithm

If a deterministic packet-size generation method is used (such as a monotonic increase by 1 octet from start value to MTU), then the generation algorithm SHOULD be reported.

6.4. Pseudorandom Length Algorithm

If a pseudorandom 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 pseudorandom generation algorithms and seeds SHOULD be reported when used.

6.5. Pseudorandom Sequence Algorithm

Finally, we note another possibility: a pseudorandom 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.

7. 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 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 (DUT) or System Under Test (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 Sarah Banks, Aamer Akhter, Steve Maxwell, and Scott Bradner for their reviews and comments. Ilya Varlashkin suggested the run-length encoding approach in Section 6.1.

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] Agilent, "The Journal of Internet Test Methodologies", September 2007, <http://www.ixiacom.com/pdfs/test_plans/agilent_journal_of_internet_test_methodologies.pdf>.
- [IMIXonList] IETF Benchmarking Methodology Working Group, "Discussion on IMIX", October 2003, <<http://www.ietf.org/mail-archive/web/bmwg/current/msg00691.html>>.
- [IXIA] IXIA, "Testing PPPoX and L2TP Broadband Access Devices", 2004, <http://www.ixiacom.com/library/test_plans/display?skey=testing_pppox>.
- [Jumbo1] Dykstra, P., "Gigabit Ethernet Jumbo Frames, and why you should care", WareOnEarth Communications, Inc., December 1999, <<http://sd.wareonearth.com/~phil/jumbo.html>>.
- [Jumbo2] Mathis, M., "The Ethernet CRC limits packets to about 12 kBytes. (NOT)", Pittsburgh Supercomputing Center, April 2003, <<http://staff.psc.edu/mathis/MTU/arguments.html#crc>>.
- [RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, March 2009.
- [Spirent] Spirent, "Test Methodology Journal: IMIX (Internet Mix) Journal", January 2006, <<http://gospirent.com/whitepaper/IMIX%20Test%20Methodolgy%20Journal.pdf>>.

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/>

