

Network Working Group
Internet Draft
Intended status: Informational
Expires: March 2015

M. Georgescu
NAIST
September 24, 2014

IPv6 Transition Technologies Benchmarking Methodology
draft-georgescu-ipv6-transition-tech-benchmarking-00.txt

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

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

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

This Internet-Draft will expire on March 24, 2015.

Copyright Notice

Copyright (c) 2014 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.

Abstract

There are benchmarking methodologies addressing the performance of network interconnect devices which are IPv4 or IPv6-capable. However, the IPv6 transition technologies are outside of their scope. This document provides complementary guidelines for evaluating the performance of IPv6 transition technologies. The methodology also includes a tentative metric for benchmarking scalability.

Table of Contents

1. Introduction.....	3
1.1. IPv6 transition technologies.....	3
2. Conventions used in this document.....	4
3. Test environment setup.....	4
3.1. Single-stack transition technologies.....	4
3.2. Encapsulation/Translation based transition technologies...	5
4. Test traffic.....	5
4.1. Frame formats and sizes.....	5
4.1.1. Frame sizes to be used over Ethernet.....	6
4.1.2. Frame sizes to be used over SONET.....	6
4.2. Protocol addresses.....	6
4.3. Traffic setup.....	6
5. Modifiers.....	7
6. Benchmarking tests.....	7
6.1. Throughput.....	7
6.2. Latency.....	7
6.3. Frame loss rate.....	7
6.4. Back-to-back frames.....	7
6.5. System recovery.....	8
6.6. Reset.....	8
7. Scalability.....	8
7.1. Test setup.....	8
7.1.1. Single-stack transition technologies.....	8
7.1.2. Encapsulation/Translation transition technologies...	9
7.2. Benchmarking performance degradation.....	9
8. Security Considerations.....	10
9. IANA Considerations.....	10
10. Conclusions.....	10
11. References.....	11
11.1. Normative References.....	11
11.2. Informative References.....	11
12. Acknowledgments.....	11
Appendix A. Theoretical maximum frame rates.....	12
A.1. Ethernet.....	12

1. Introduction

The methodologies described in [RFC2544] and [RFC5180] help vendors and network operators alike analyze the performance of IPv4 and IPv6-capable network devices. The methodology presented in [RFC2544] is mostly IP version independent, while [RFC5180] contains complementary recommendations which are specific to the latest IP version, IPv6. However, [RFC5180] does not cover IPv6 transition technologies.

IPv6 is not backwards compatible, which means that IPv4-only nodes cannot directly communicate with IPv6-only nodes. To solve this issue, IPv6 transition technologies have been proposed and implemented, many of which are still in development.

This document presents benchmarking guidelines dedicated to IPv6 transition technologies. The benchmarking tests can provide insights about the performance of these technologies, which can act as useful feedback for developers, as well as for network operators going through the IPv6 transition process.

1.1. IPv6 transition technologies

Two of the basic transition technologies dual IP layer (also known as dual stack) and encapsulation are presented in [RFC4213]. IPv4/IPv6 Translation is presented in [RFC6144]. Most of the transition technologies employ at least one variation of these mechanisms. Some of the more complex ones (e.g. DSLite [RFC6333]) are using all three. In this context, a generic classification of the transition technologies can prove useful.

Tentatively, we can consider a basic production IP-based network as being constructed using the following components:

- o a Customer Edge (CE) segment
- o a Core network segment
- o a Provider Edge (PE) segment

According to the technology used for the core network traversal the transition technologies can be categorized as follows:

1. Single-stack: either IPv4 or IPv6 is used to traverse the core network and translation is used at one of the edges
2. Dual-stack: the core network devices implement both IP protocols

3. Encapsulation-based: an encapsulation mechanism is used to traverse the core network; CE nodes encapsulate the IPvX packets in IPvY packets, while PE nodes are responsible for the decapsulation process.
4. Translation-based: a translation mechanism is employed for the traversal of the network core; CE nodes translate IPvX packets to IPvY packets and PE nodes translate the packets back to IPvX.

The performance of Dual-stack transition technologies can be very well evaluated using the benchmarking methodology presented by [RFC2544] and [RFC5180]. Consequently the focus of this document is represented by the other 3 categories: Single-stack, Encapsulation-based and Translation-based transition technologies.

2. Conventions used in this document

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 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [RFC2119] significance.

3. Test environment setup

The test environment setup options recommended for IPv6 transition technologies benchmarking are very similar to the ones presented in Section 6 of [RFC2544]. In the case of the tester setup, the options presented in [RFC2544] can be applied here as well. However, the Device under test (DUT) setup options should be explained in the context of the 3 targeted categories of IPv6 transition technologies: Single-stack, Encapsulation-based and Translation-based transition technologies.

Although both single tester and sender/receiver setups are applicable to this methodology, the single tester setup will be used to describe the DUT setup options.

3.1. Single-stack transition technologies

For the evaluation of Single-stack transition technologies a single DUT setup (see Figure 1) SHOULD be used. The DUT is responsible for translating the IPvX packets into IPvY packets. In this context, the tester device should be configured to support both IPvX and IPvY.

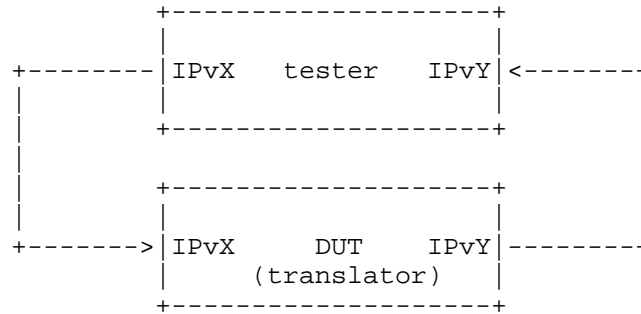


Figure 1

3.2. Encapsulation/Translation based transition technologies

For evaluating the performance of Encapsulation-based and Translation-based transition technologies a dual DUT setup (see Figure 2) SHOULD be employed. The tester creates a network flow of IPvX packets. The DUT CE is responsible for the encapsulation or translation of IPvX packets into IPvY packets. The IPvY packets are decapsulated/translated back to IPvX packets by the DUT PE and forwarded to the tester.

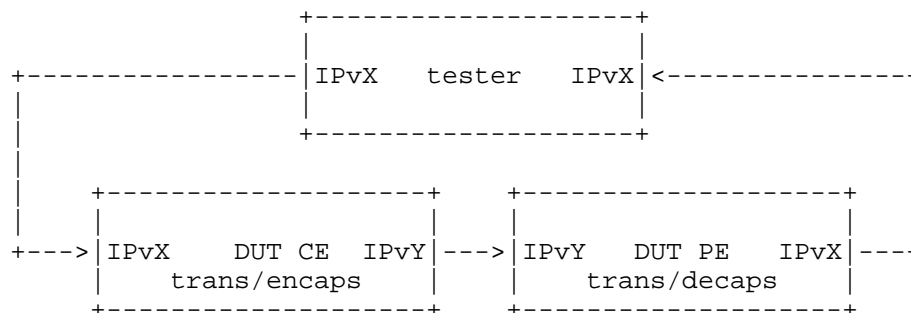


Figure 2

4. Test traffic

The test traffic represents the experimental workload and SHOULD meet the requirements specified in this section. The requirements are dedicated to unicast IP traffic.

4.1. Frame formats and sizes

[RFC5180] describes the frame size requirements for two commonly used media types: Ethernet and SONET (Synchronous Optical Network).

Internet-Draft IPv6 transition tech benchmarking September 2014
[RFC2544] covers also other media types, such as token ring and FDDI. The two documents can be referred for the dual-stack transition technologies. For the rest of the transition technologies the frame overhead introduced by translation or encapsulation MUST be considered.

The encapsulation/translation process generates different size frames on different segments of the test setup. For example, the single-stack transition technologies will create different frame sizes on the receiving segment of the test setup, as IPvX packets are translated to IPvY. This is not a problem if the bandwidth of the employed media is not exceeded. To prevent exceeding the limitations imposed by the media, the frame size overhead needs to be taken into account when calculating the maximum theoretical frame rates. The calculation methods for the two media types, Ethernet and SONET, as well as a calculation example are detailed in Appendix A.

4.1.1. Frame sizes to be used over Ethernet

Based on the recommendations of [RFC5180], the following frame sizes SHOULD be used for benchmarking Ethernet traffic: 64, 128, 256, 512, 1024, 1280, 1518, 1522, 2048, 4096, 8192 and 9216.

The theoretical maximum frame rates considering an example of frame overhead are presented in Appendix A1.

4.1.2. Frame sizes to be used over SONET

Based on the recommendations of [RFC5180], the frame sizes for SONET traffic SHOULD be: 47, 64, 128, 256, 512, 1024, 1280, 1518, 2048, 4096 bytes.

An example of theoretical maximum frame rates calculation is shown in Appendix A2.

4.2. Protocol addresses

The selected protocol addresses should follow the recommendations of [RFC5180](Section 5) for IPv6 and [RFC2544](Section 12) for IPv4.

Note: testing traffic with extension headers might not be possible for the transition technologies which employ translation.

4.3. Traffic setup

Following the recommendations of [RFC5180], all tests described SHOULD be performed with bi-directional traffic. Uni-directional traffic tests MAY also be performed for a fine grained performance assessment.

The idea of testing under different operational conditions was first introduced in [RFC2544](Section 11) and represents an important aspect of benchmarking network elements, as it emulates to some extent the conditions of a production environment. [RFC5180] describes complementary testing conditions specific to IPv6. Their recommendations can be referred for IPv6 transition technologies testing as well.

6. Benchmarking tests

The benchmarking tests condition described in [RFC2544] (Sections 24, 25, 26) are also recommended here. The following sub-sections contain the list of all recommended benchmarking tests.

6.1. Throughput

Objective: To determine the DUT throughput as defined in [RFC1242].

Procedure: As described by [RFC2544].

Reporting Format: As described by [RFC2544].

6.2. Latency

Objective: To determine the latency as defined in [RFC1242].

Procedure: As described by [RFC2544].

Reporting Format: As described by [RFC2544].

6.3. Frame loss rate

Objective: To determine the frame loss rate, as defined in [RFC1242], of a DUT throughout the entire range of input data rates and frame sizes.

Procedure: As described by [RFC2544].

Reporting Format: As described by [RFC2544].

6.4. Back-to-back frames

Objective: To characterize the ability of a DUT to process back-to-back frames as defined in [RFC1242].

Procedure: As described by [RFC2544].

6.5. System recovery

Objective: To characterize the speed at which a DUT recovers from an overload condition.

Procedure: As described by [RFC2544].

Reporting Format: As described by [RFC2544].

6.6. Reset

Objective: To characterize the speed at which a DUT recovers from a device or software reset.

Procedure: As described by [RFC6201].

Reporting Format: As described by [RFC6201].

7. Scalability

Scalability has been often discussed, however, in the context of network devices, a formal definition or a measurement method have not been approached yet.

Scalability can be defined as the ability of each transition technology to accommodate network growth.

Poor scalability usually leads to poor performance. Considering this, scalability can be measured by quantifying the network performance degradation while the network grows.

7.1. Test setup

The test setups defined in Section 3 have to be modified to create network growth.

7.1.1. Single-stack transition technologies

In the case of single-stack transition technologies the network growth can be generated by increasing the number of network flows generated by the tester machine (see Figure 3).




Figure 3

7.1.2. Encapsulation/Translation transition technologies

Similarly, for the encapsulation/translation based technologies a multi-flow setup is recommended. As for most transition technologies the provider edge device is designed to support more than one customer edge network, the recommended test setup is a n:1 design, where n is the number of CE DUTs connected to the same PE DUT (See Figure 4).




Figure 4

7.2. Benchmarking performance degradation

Objective: To quantify the performance degradation introduced by n parallel network flows.

Internet-Draft IPv6 transition tech benchmarking September 2014
Procedure: First the benchmarking tests presented in Section 6 have to be performed for one network flow.

The same tests have to be repeated for n-network flows. The performance degradation of the X benchmarking dimension SHOULD be calculated as relative performance change between the 1-flow results and the n-flow results, using the following formula:

$$\text{Xpd} = \frac{\text{Xn} - \text{X1}}{\text{X1}} \times 100, \text{ where: } \text{X1} - \text{result for 1-flow} \\ \text{Xn} - \text{result for n-flows}$$

Reporting Format: The performance degradation SHOULD be expressed as a percentage. The number of tested parallel flows n MUST be clearly specified. For each of the performed benchmarking tests there SHOULD be a table containing a column for each frame size, stating also the applied frame rate.

8. Security Considerations

The benchmarking methodology described in this document MUST be used in conjunction with a controlled experimental environment.

The benchmarking environment MUST be isolated and the generated traffic MUST NOT be forwarded into production networks.

Given the isolated nature of the experimental environment, no other security considerations are required.

9. IANA Considerations

The IANA has allocated the prefix 2001:0002::/48 [RFC5180] for IPv6 benchmarking. For IPv4 benchmarking, the 198.18.0.0/15 prefix was reserved, as described in [RFC6890]. The two ranges are sufficient for benchmarking IPv6 transition technologies.

10. Conclusions

The methodologies described in [RFC2544] and [RFC5180] can be used for benchmarking the performance of IPv4-only, IPv6-only and dual-stack supporting network devices. This document presents complementary recommendations dedicated to IPv6 transition technologies. Furthermore, the methodology includes a tentative approach for benchmarking scalability by quantifying the performance degradation associated with network growth.

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2234] Crocker, D. and Overell, P.(Editors), "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, Internet Mail Consortium and Demon Internet Ltd., November 1997.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.
- [RFC6333] Cotton, M., Vegoda, L., Bonica, R., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC6890, April 2013.

11.2. Informative References

- [RFC1242] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", [RFC1242], July 1991.
- [RFC2544] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", [RFC1242], July 1991.
- [RFC5180] Popoviciu, C., Hamza, A., Van de Velde, G., and D. Dugatkin, "IPv6 Benchmarking Methodology for Network Interconnect Devices", RFC 5180, May 2008.
- [RFC6201] Asati, R., Pignataro, C., Calabria, F., and C. Olvera, "Device Reset Characterization", RFC 6201, March 2011.

12. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

This appendix describes the recommended calculation formulas for the theoretical maximum frame rates to be employed over two types of commonly used media. The formulas take into account the frame size overhead created by the encapsulation or the translation process. For example, the 6in4 encapsulation described in [RFC4213] adds 20 bytes of overhead to each frame.

A.1. Ethernet

Considering X to be the frame size and O to be the frame size overhead created by the encapsulation on translation process, the maximum theoretical frame rate for Ethernet can be calculated using the following formula:

$$\frac{\text{Line Rate (bps)}}{(8\text{bits/byte}) * (X + O + 20)\text{bytes/frame}}$$

The calculation is based on the formula recommended by RFC5180 in Appendix A1. As an example, the frame rate recommended for testing a 6in4 implementation over 10Mb/s Ethernet with 64 bytes frames is:

$$\frac{10,000,000(\text{bps})}{(8\text{bits/byte}) * (64 + 20 + 20)\text{bytes/frame}} = 12,019 \text{ fps}$$

The complete list of recommended frame rates for 6in4 encapsulation can be found in the following table:

Frame size (bytes)	10 Mb/s (fps)	100 Mb/s (fps)	1000 Mb/s (fps)	10000 Mb/s (fps)
64	12,019	120,192	1,201,923	12,019,231
128	7,440	74,405	744,048	7,440,476
256	4,223	42,230	422,297	4,222,973
512	2,264	22,645	226,449	2,264,493
1024	1,175	11,748	117,481	1,174,812
1280	947	9,470	94,697	946,970
1518	802	8,023	80,231	802,311
1522	800	8,003	80,026	800,256
2048	599	5,987	59,866	598,659
4096	302	3,022	30,222	302,224
8192	152	1,518	15,185	151,846
9216	135	1,350	13,505	135,048

A.2. SONET

Similarly for SONET, if X is the target frame size and O the frame size overhead, the recommended formula for calculating the maximum theoretical frame rate is:

$$\frac{\text{Line Rate (bps)}}{(8\text{bits/byte}) * (X+O+1)\text{bytes/frame}}$$

The calculation formula is based on the recommendation of RFC5180 in Appendix A2.

As an example, the frame rate recommended for testing a 6in4 implementation over a 10Mb/s PoS interface with 64 bytes frames is:

$$\frac{10,000,000(\text{bps})}{(8\text{bits/byte}) * (64+20+1)\text{bytes/frame}} = 14,706 \text{ fps}$$

The complete list of recommended frame rates for 6in4 encapsulation can be found in the following table:

Frame size (bytes)	10 Mb/s (fps)	100 Mb/s (fps)	1000 Mb/s (fps)	10000 Mb/s (fps)
47	18,382	183,824	1,838,235	18,382,353
64	14,706	147,059	1,470,588	14,705,882
128	8,389	83,893	838,926	8,389,262
256	4,513	45,126	451,264	4,512,635
512	2,345	23,452	234,522	2,345,216
1024	1,196	11,962	119,617	1,196,172
2048	604	6,042	60,416	604,157
4096	304	3,036	30,362	303,619

Marius Georgescu
Nara Institute of Science and Technology (NAIST)
Takayama 8916-5
Nara
Japan

Phone: +81 743 72 5216
Email: liviumarius-g@is.naist.jp

