Network Working Group   M. Georgescu
Internet Draft   NAIST
Intended status: Informational   July 2, 2015
Expires: January 2016

Benchmarking Methodology for IPv6 Transition Technologies
draft-georgescu-bmwg-ipv6-tran-tech-benchmarking-01.txt

Abstract

There are benchmarking methodologies addressing the performance of network interconnect devices that are IPv4- or IPv6-capable, but 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.

Status of this Memo

Copyright Notice

Table of Contents

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:

- a Customer Edge (CE) segment
- a Core network segment
- 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 core network; 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 fully evaluated using the benchmarking methodologies presented by [RFC2544] and [RFC5180]. Consequently, this document focuses on the other 3 categories: Single-stack, Encapsulation-based, and Translation-based transition technologies.

Another important aspect by which the IPv6 transition technologies can be categorized is their use of stateful or stateless mapping algorithms. The technologies that use stateful mapping algorithms (e.g. Stateful NAT64 [RFC6146]) create dynamic correlations between IP addresses or {IP address, transport protocol, transport port number} tuples, which are stored in a state table. For ease of reference, the IPv6 transition technologies which employ stateful mapping algorithms will be called stateful IPv6 transition technologies. The efficiency with which the state table is managed can be an important performance indicator for these technologies. Hence, for the stateful IPv6 transition technologies additional benchmarking tests are RECOMMENDED.

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 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.

For the test setups presented in this memo dynamic routing SHOULD be employed. However, the presence of routing and management frames can represent unwanted background data that can affect the benchmarking result. To that end, the procedures defined in [RFC2544] (Sections 11.2 and 11.3) related to routing and management frames SHOULD be used here as well. Moreover, the "Trial description" recommendations presented in [RFC2544] (Section 23) are valid for this memo as well.

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.

```
                         +--------------------+
                         |                    |
            +------------|IPvX   tester   IPvY|<-------------+
            |            |                    |              |
            |            +--------------------+              |
            |                                                |
            |            +--------------------+              |
            |            |                    |              |
            +----------->|IPvX     DUT    IPvY|--------------+
                         |     (translator)   |
                         +--------------------+
```

Figure 1. Test setup 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.

```
                            +--------------------+
                            |                    |
  +---------------------|IPvX   tester   IPvX|<------------------+
  |                     |                    |                   |
  |                     +--------------------+                   |
  |                                                              |
  |      +--------------------+      +--------------------+      |
  |      |                    |      |                    |      |
  +----->|IPvX    DUT CE  IPvY|----->|IPvY   DUT PE   IPvX|------+
         |    trans/encaps    |      |    trans/decaps    |
         +--------------------+      +--------------------+
```

Figure 2. Test setup 2

In the case of translation based transition technology, the DUT CE and DUT PE machines MAY be tested separately as well. These tests can represent a fine grain performance analysis of the IPvX to IPvY translation direction versus the IPvY to IPvX translation direction. The tests SHOULD follow the test setup presented in Figure 1.

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. Multicast IP traffic is outside of the scope of this document.

4.1. Frame Formats and Sizes

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

In the context of frame size overhead MTU recommendations are needed in order to avoid frame loss due to MTU mismatch between the virtual encapsulation/translation interfaces and the physical network interface controllers (NICs). To avoid this situation, the larger MTU between the physical NICs and virtual encapsulation/translation interfaces SHOULD be set for all interfaces of the DUT and tester.

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.

Because of the simplicity of UDP, UDP measurements offer a more reliable basis for comparison than other transport layer protocols. Consequently, for the benchmarking tests described in Section 6 of this document UDP traffic SHOULD be employed.

Considering that the stateful transition technologies need to manage the state table for each connection, a connection-oriented transport layer protocol needs to be used with the test traffic. Consequently, TCP test traffic SHOULD be employed for the tests described in Section 7 of this document.

5. Modifiers

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 test conditions 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. Packet Delay Variation

Considering two of the metrics presented in [RFC5481], Packet Delay Variation (PDV) and Inter Packet Delay Variation (IPDV), it is RECOMMENDED to measure PDV. For a fine grain analysis of delay variation, IPDV measurements MAY be performed as well.

6.3.1. PDV

Objective: To determine the Packet Delay Variation as defined in [RFC5481].

Procedure: As described by [RFC2544], first determine the throughput for the DUT at each of the listed frame sizes. Send a stream of frames at a particular frame size through the DUT at the determined throughput rate to a specific destination. The stream SHOULD be at least 60 seconds in duration. Measure the One-way delay as described by [RFC3393] for all frames in the stream. Calculate the PDV of the stream using the formula:

```
PDV=Avg(D(i) - Dmin)

Where: D(i) - the One-way delay of the i-th frame in the stream

       Dmin - the minimum One-way delay in the stream
```

As recommended in RFC 2544, the test MUST be repeated at least 20 times with the reported value being the average of the recorded values. Moreover, the margin of error from the average MAY be evaluated following the formula:

```
               StDev
MoE= alpha * ----------
              sqrt(N)

Where: alpha   - critical value; the recommended value is 2.576 for
a 99% level of confidence
      StDev    - standard deviation
      N        - number of repetitions
```

Reporting Format: The PDV results SHOULD be reported in a table with a row for each of the tested frame sizes and columns for the frame size and the applied frame rate for the tested media types. A column for the margin of error values MAY as well be displayed.

6.3.2. IPDV

Objective: To determine the Inter Packet Delay Variation as defined in [RFC5481].

Procedure: As described by [RFC2544], first determine the throughput for the DUT at each of the listed frame sizes. Send a stream of frames at a particular frame size through the DUT at the determined throughput rate to a specific destination. The stream SHOULD be at least 60 seconds in duration. Measure the One-way delay as described by [RFC3393] for all frames in the stream. Calculate the IPDV for each of the frames using the formula:

```
IPDV(i)=D(i) - D(i-1)

Where: D(i)   - the One-way delay of the i th frame in the stream

       D(i-1) - the One-way delay of i-1 th frame in the stream
```

Given the nature of IPDV, reporting a single number might lead to over-summarization. In this context, the report for each measurement SHOULD include 3 values: Dmin, Davg, and Dmax

Where: Dmin - the minimum One-way delay in the stream

Davg - the average One-way delay of the stream

Dmax - the maximum One-way delay in the stream

As recommended in RFC 2544, the test MUST be repeated at least 20 times. The average of the 3 proposed values SHOULD be reported. The IPDV results SHOULD be reported in a table with a row for each of the tested frame sizes. The columns SHOULD include the frame size and associated frame rate for the tested media types and sub-columns for the three proposed reported values.

6.4. 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].

6.5. 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].

Reporting Format: As described by [RFC2544].

6.6. 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].

6.7. Reset

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

Procedure: As described by [RFC2544].

Reporting Format: As described by [RFC2544].

7. Additional Benchmarking Tests for Stateful IPv6 Transition Technologies

This section describes additional tests dedicated to the stateful IPv6 transition technologies. For the tests described in this section the DUT devices SHOULD follow the test setup and test parameters recommendations presented in [RFC3511] (Sections 4, 5).

In addition to the IPv4/IPv6 transition function a network node can have a firewall function. This document is targeting only the network devices that do not have a firewall function, as this function can be benchmarked using the recommendations of [RFC3511]. Consequently, only the tests described in [RFC3511] (Sections 5.2, 5.3) are RECOMMENDED. Namely, the following additional tests SHOULD be performed:

7.1. Concurrent TCP Connection Capacity

Objective: To determine the maximum number of concurrent TCP connections supported through or with the DUT, as defined in [RFC 2647]. This test is supposed to find the maximum number of entries the DUT can store in its state table.

Procedure: As described by [RFC3511].

Reporting Format: As described by [RFC3511].

7.2. Maximum TCP Connection Establishment Rate

Objective: To determine the maximum TCP connection establishment rate through or with the DUT, as defined by RFC [2647]. This test is expected to find the maximum rate at which the DUT can update its connection table.

Procedure: As described by [RFC3511].

Reporting Format: As described by [RFC3511].

8. Scalability

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

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.

The following subsections describe how the test setups can be modified to create network growth and how the associated performance degradation can be quantified.

8.1. Test Setup

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

8.1.1. 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).

```
             +-------------------------+
+------------|NF1                   NF1|<-------------+
|  +---------|NF2      tester       NF2|<----------+  |
|  |      ...|                         |           |  |
|  |   +-----|NFn                   NFn|<------+    |  |
|  |   |     +-------------------------+       |    |  |
|  |   |                                       |    |  |
|  |   |     +-------------------------+       |    |  |
|  |   +---->|NFn                   NFn|-------+    |  |
|  |      ...|           DUT           |           |  |
|  +-------->|NF2    (translator)   NF2|-----------+  |
+----------->|NF1                   NF1|--------------+
             +-------------------------+
```

Figure 3. Test setup 3

8.1.2. Encapsulation/Translation Transition Technologies

Similarly, for the encapsulation/translation based technologies a multi-flow setup is recommended. For most transition technologies, the provider edge device is designed to support more than one customer edge network. Hence, 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).

```
                          +-------------------------+
     +--------------------|NF1                   NF1|<--------------+
     |  +-----------------|NF2      tester       NF2|<-----------+  |
     |  |              ...|                         |            |  |
     |  |   +-------------|NFn                   NFn|<-------+   |  |
     |  |   |             +-------------------------+        |   |  |
     |  |   |                                                |   |  |
     |  |   |    +-----------------+    +---------------+    |   |  |
     |  |   +--->|NFn  DUT CEn  NFn|--->|NFn         NFn| ---+   |  |
     |  |        +-----------------+    |               |        |  |
     |  |     ...                       |               |        |  |
     |  |        +-----------------+    |    DUT PE     |        |  |
     |  +------->|NF2  DUT CE2  NF2|--->|NF2         NF2|--------+  |
     |           +-----------------+    |               |           |
     |           +-----------------+    |               |           |
     +---------->|NF1  DUT CE1  NF1|--->|NF1         NF1|-----------+
                 +-----------------+    +---------------+
```

Figure 4. Test setup 4

This test setup can help to quantify the scalability of the PE device. However, for testing the scalability of the DUT CEs additional recommendations are needed.
For encapsulation based transition technologies a m:n setup can be created, where m is the number of flows applied to the same CE device and n the number of CE devices connected to the same PE device.
For the translation based transition technologies the CE devices can be separately tested with n network flows using the test setup presented in Figure 3.

8.2. Benchmarking Performance Degradation

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

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:

```
              Xn - X1
       Xpd= ----------- * 100, where: X1 - result for 1-flow
                 X1                    Xn - 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. The table SHOULD also state the applied frame rate.

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

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.

11. 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.

12. References

12.1. 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.

[RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.

[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.

12.2. Informative References

[RFC1242] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", [RFC1242], July 1991.

[RFC2544] Bradner, S., McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", [RFC2544], March 1999.

[RFC2647] Newman, D., "Benchmarking Terminology for Firewall Devices", [RFC2647], August 1999.

[RFC3511] Hickman, B., Newman, D., Tadjudin, S., Martin, T., "Benchmarking Methodology for Firewall Performance", [RFC3511], April 2003.

[RFC5180] Popoviciu, C., Hamza, A., Van de Velde, G., and D. Dugatkin, "IPv6 Benchmarking Methodology for Network Interconnect Devices", RFC 5180, May 2008.

[RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.

13. Acknowledgments

The author would like to thank Professor Youki Kadobayashi for his constant feedback and support. The thanks should be extended to the NECOMA project members for their continuous support. Helpful comments and suggestions were offered by Scott Bradner, Al Morton, Bhuvaneswaran Vengainathan, Andrew McGregor, Nalini Elkins, Kaname Nishizuka and Yasuhiro Ohara. A special thank you to the RFC Editor Team for their thorough editorial review and helpful suggestions. This document was prepared using 2-Word-v2.0.template.dot.

Appendix A. Theoretical Maximum Frame Rates

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:

```
              Line Rate (bps)
       ------------------------------
       (8bits/byte)*(X+O+20)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:

```
              10,000,000(bps)
       ------------------------------      = 12,019 fps
       (8bits/byte)*(64+20+20)bytes/frame
```

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:

```
              Line Rate (bps)
      ------------------------------
      (8bits/byte)*(X+O+1)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:

```
              10,000,000(bps)
      ------------------------------      = 14,706 fps
      (8bits/byte)*(64+20+1)bytes/frame
```

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

Authors' Addresses

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