Benchmarking Methodology for IPv6 Transition Technologies

draft-georgescu-bmwg-ipv6-tran-tech-benchmarking-01

Marius Georgescu

Nara Institute of Science and Technology
Internet Engineering Laboratory

23 Jul. 2015

Draft overview

- RFC2544[1] and RFC5180[2] address both IPv4 and IPv6 performance benchmarking, but IPv6 transition technologies are outside their scope
- This draft provides complementary guidelines for evaluating the performance of IPv6 transition technologies
  - generic classification on IPv6 transition technologies → associated test setups
  - calculation formula for the maximum frame rate according to the *frame size overhead*
- Includes a tentative metric for benchmarking scalability
  - scalability as *performance degradation* under the stress of *multiple network flows*
- Proposes supplementary benchmarking tests for *stateful* IPv6 transition technologies in accordance with RFC3511[3]

---

[1] S. Bradner and J. McQuaid. *Benchmarking Methodology for Network Interconnect Devices*. United States, 1999.

[2] A. Hamza C. Popoviciu, G. Van de Velde, and D. Dugatkin. *IPv6 Benchmarking Methodology for Network Interconnect Devices*. RFC 5180. Internet Engineering Task Force, 2008.

[3] B. Hickman et al. *Benchmarking Methodology for Firewall Performance*. RFC 3511 (Informational). Internet Engineering Task Force, Apr. 2003. URL: http://www.ietf.org/rfc/rfc3511.txt.

Production networks generic design

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

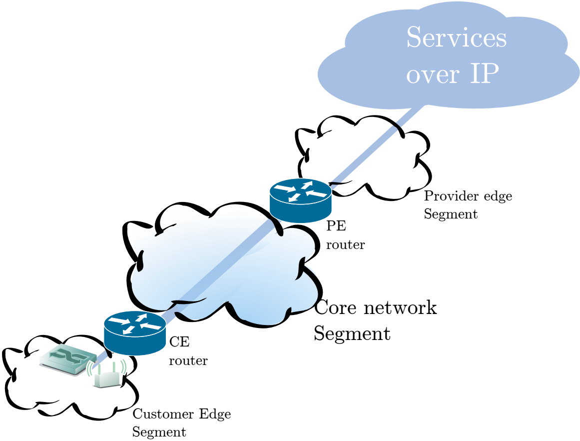

IPv6 transition technologies generic categories

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.

Test environment setup

Single-stack transition technologies

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

Test environment setup (contd.)

Encapsulation/Translation-based transition technologies

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

EMPIRICAL RESULTS

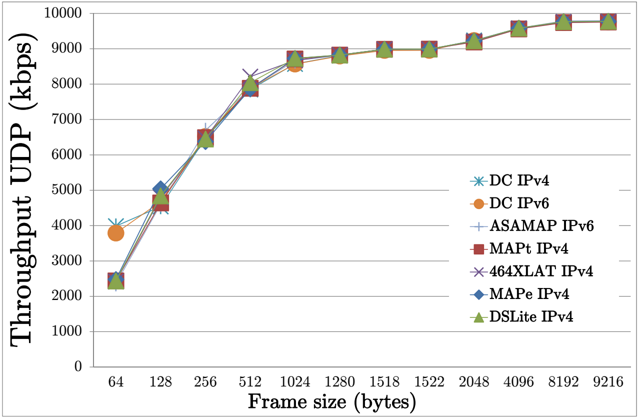

[4]

[4] M. Georgescu et al. "Empirical analysis of IPv6 transition technologies using the IPv6 Network Evaluation Testbed". In: *9th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities*. Guangzhou, China, 2014.

Scalability benchmarking

Benchmarking Scalability through **performance degradation**

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

Procedure: First the benchmarking tests 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:

Reporting format: relative performance change between the 1-flow results x_1 and the n-flow results x_n

$$X_{pd} = \frac{x_n - x_1}{x_1} \times 100 \quad (1)$$

Scalability test setup

Single-stack transition technologies

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

SCALABILITY TEST SETUP (CONTD.)

Encapsulation/Translation-based transition technologies

```
                  +-------------------------+
+-----------------|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|----------+
            +-----------------+    +--------------+
```

SCALABILITY EMPIRICAL RESULTS

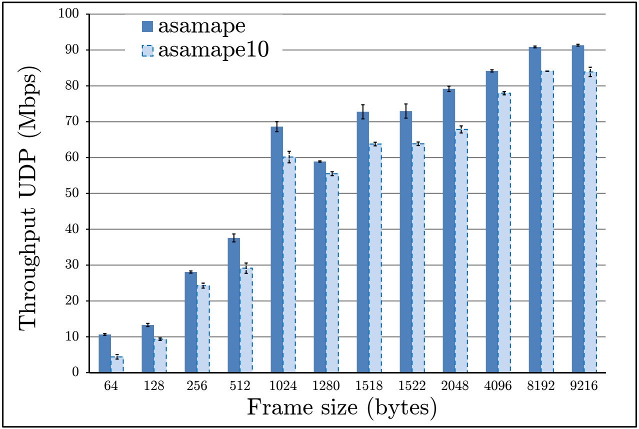

[5]

[5] Marius Georgescu et al. "Benchmarking the Load Scalability of IPv6 Transition Technologies: a Black-Box Analysis". In: *20th IEEE Symposium on Computers and Communication (ISCC) (ISCC2015).* Larnaca, Cyprus, July 2015, pp. 238–243.

Update 01 overview

- Recommended dynamic routing as the only option for the test setups
- Added benchmarking tests for Packet Delay variation
  - SHOULD measure PDV as defined by RFC5481[6]
  - MAY measure IPDV as defined by RFC5481

[6] A. Morton and B. Claise. *Packet Delay Variation Applicability Statement*. RFC 5481 (Informational). Internet Engineering Task Force, Mar. 2009. URL: http://www.ietf.org/rfc/rfc5481.txt.

Update: Routing recommendations

Text modified in Section 3:

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.

[6] **following the comment from Scott Bradner**.

UPDATE: PACKET DELAY VARIATION - CALCULATION FORMULA

Metric1: Packet Delay Variation (PDV) - Formula

$$PDV_i = D_i - D_{min}$$

D_i - One-Way-Delay of i-th frame (2)

D_{min} - The minimum One-Way-Delay of the frame stream

Metric2: Inter Packet Delay Variation (IPDV) - Formula

$$IPDV_i = D_i - D_{i-1}$$

D_i - One-Way-Delay of i th frame (3)

D_{i-1} - The One-Way-Delay of the i-1 th frame

[6] **following the comments from Al Morton and Bhuvan Vengainathan and the discussion in IETF92**.

Update: Packet Delay Variation - Summarization

Metric1: Packet Delay Variation (PDV) - Single number summarization

$$
\begin{aligned}
&PDV = Avg(D_i - D_{min}) \\
&D_i \text{ - One-Way-Delay of i-th frame} \\
&D_{min} \text{ - The minimum One-Way-Delay of the frame stream}
\end{aligned} \tag{4}
$$

Metric2: Inter Packet Delay Variation (IPDV) - 3 numbers summarization

$$
\begin{aligned}
IPDV_{min} &= Min(D_i - D_{i-1}) \\
IPDV_{avg} &= Avg(D_i - D_{i-1}) \\
IPDV_{max} &= Max(D_i - D_{i-1})
\end{aligned} \tag{5}
$$

Update: Error expression

Margin of Error(MoE) expression

$$MoE = z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \tag{6}$$

σ - standard deviation, n - sample size

$z_{\alpha/2}$ - critical value

$z_{\alpha/2}$=2.576 for 99%level of confidence

Comments not covered yet

- The comment from Nalini Elkins related to DNS resolution
  - **Still* Considering a DNS Resolution Performance metric: *Number of processed DNS requests/sec* [7]
  - Planning to include it in *version 02* feat. Prof. Gábor Lencse [RG profile link]
- Many comments received from Al Morton on the ML
  - Planning to address them in *version 02*

[7] Gabor Lencse and Sandor Repas. “Performance Analysis and Comparison of Different DNS64 Implementations for Linux, OpenBSD and FreeBSD”. . In: *Proceedings of the 2013 IEEE 27th International Conference on Advanced Information Networking and Applications*. AINA ’13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 877–884. ISBN: 978-0-7695-4953-8. DOI: 10.1109/AINA.2013.80. URL: http://dx.doi.org/10.1109/AINA.2013.80.

Next steps

- Propose solutions for *DNS Resolution Performance*
- Continue the revisions

⋆ Question for BMWG:

- Were the comments covered well enough?
- How adoption-ready is the draft?

Contact


Nara Institute of Science and Technology
NAIST
Marius Georgescu
liviumarius-g@is.naist.jp
www.ipv6net.ro
IPv6
IPv4