



Benchmarking Methodology for IPv6 Transition Technologies

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

Marius Georgescu

Nara Institute of Science and Technology
Internet Engineering Laboratory

7 Apr. 2016

DRAFT MOTIVATION: IPv6 TRANSITION

- ▶ IPv6 is not backwards compatible
- ▶ The Internet will undergo a period through which both protocols will coexist
- ▶ Currently only 5% of worldwide Internet users have IPv6 connectivity¹




¹APNIC. *IPv6 measurements for The World*. Asia-Pacific Network Information Centre, Apr. 2016. URL: <http://labs.apnic.net/ipv6-measurement/Regions/>.

²Original drawing by Andrew Bell @ www.creaturesinmyhead.com.

IPv6 TRANSITION TECHNOLOGIES EVOLUTION

- ▶ What benchmarks to use?
 - ▶ For Dual Stack RFC2544 or RFC5180 are enough
 - ▶ How about translation/encapsulation technologies?



3

³inspired by the APNIC35 presentation "The evolution of IPv6 transition technologies" by Jouni Korhonen.

DRAFT OVERVIEW

- ▶ This draft provides complementary guidelines to RFC2544⁴ and RFC5180⁵ 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 RFC3511⁶
- ▶ Proposes supplementary benchmarking tests for *DNS resolution performance*
 - ▶ contributed by Prof. Gábor Lencse [RG profile link]

⁴S. Bradner and J. McQuaid. *Benchmarking Methodology for Network Interconnect Devices*. United States, 1999.

⁵A. Hamza C. Popoviciu, G. Van de Velde, and D. Dugatkin. *IPv6 Benchmarking Methodology for Network Interconnect Devices*. RFC 5180. Internet Engineering Task Force, 2008.

⁶B. 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>.

UPDATE OVERVIEW

- ▶ New procedure for latency to report Typical Latency and Worst Latency
- ▶ Summarizing and variation functions
 - ▶ Mean + MoE → Median + 1st-99th Percentiles
- ▶ Simultaneous and Incremental benchmarks for network performance degradation
- ▶ Generic transition technologies association table
- ▶ DNS Resolution Performance
 - ▶ Tester configuration
 - ▶ Test duration
 - ▶ Requirements for the Tester and **dns64perf++** [link]⁷
- ▶ Various smaller editorial changes (detailed changelog [link])

⁷D. Bakai. *A C++11 DNS64 performance tester*. 2016. URL: <https://github.com/bakaid/dns64perfpp>.

UPDATE: TYPICAL & WORST CASE LATENCY

Text added to Section 7.2:

Identifying tags SHOULD be included in at least 500 frames after 60 seconds.

Typical Latency (TL) calculation formula:

$$TL = Median(L_i)$$

where L_i - the latency of frame i


(1)

Worst Case Latency (WCS) calculation formula:

$$WCS = 99.9thPercentile(L_i)$$
(2)


⁷ following the ML discussion: <http://www.ietf.org/mail-archive/web/bmwg/current/msg03371.html>.

UPDATE: SUMMARIZING & VARIATION



N = 600 Bandwidth = 0.01682

UPDATE: SUMMARIZING & VARIATION CONT'D



N = 20 Bandwidth = 17.8

NETWORK PERFORMANCE DEGRADATION WITH INCREMENTAL LOAD

Text added to Section 10.2.2:

The same tests have to be repeated for n network flows, where the network flows are started incrementally in succession, each after time T . In other words, if flow I is started at time x , flow $i+1$ will be started at time $x+T$. Considering the time T , the time duration of each iteration must be extended with the time necessary to start all the flows, namely $(n-1) \times T$.

⁷ following the suggestion from Fred Baker.

GENERIC TRANSITION TECHNOLOGIES ASSOCIATION TABLE

	Generic Category	IPv6Transition Technology
1	Dual-stack	Dual IP Layer Operations [RFC4213]
2	Single translation	NAT64 [RFC6146], IVI [RFC6219]
3	Double translation	464XLAT [RFC6877], MAP-T [RFC7599]
4	Encapsulation	DSLite[RFC6333], MAP-E [RFC7597] Lightweight 4over6 [RFC7596] 6RD [RFC 5569]

UPDATE: DNS RESOLUTION PERFORMANCE

- ▶ Tester configuration
 - ▶ Tester may be a single device or two physical devices
- ▶ Test duration
 - ▶ The duration should be at least 60 seconds and timeout should be no more than 1 second, otherwise “gaming” is possible
- ▶ Requirements for the Tester
 - ▶ For passing the self-test, the Tester SHOULD be able to answer AAAA record queries at $2 * (r + \delta)$ rate within $0.25 * t$ timeout, where the value of δ is at least 0.1.
- ▶ **dns64perf++** [\[link\]](#)
 - ▶ Developed by Dániel Bakai in compliance with the specifications of this draft

NEXT STEPS

- ▶ Comments not covered yet
 - ▶ Fred Baker's suggestion to use this methodology to benchmark NAT XX as well
- ▶ DNS Resolution Performance for **DNS46** ?

★ Questions for BMWG:

- ▶ Were the comments covered well enough?
- ▶ Is the 1st WGLC in IETF96 a realistic milestone?

CONTACT

