

Network Working Group
Request for Comments: 5180
Category: Informational

C. Popoviciu
A. Hamza
G. Van de Velde
Cisco Systems
D. Dugatkin
FastSoft Inc.
May 2008

IPv6 Benchmarking Methodology for Network Interconnect Devices

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

The benchmarking methodologies defined in RFC 2544 are IP version independent. However, RFC 2544 does not address some of the specificities of IPv6. This document provides additional benchmarking guidelines, which in conjunction with RFC 2544, lead to a more complete and realistic evaluation of the IPv6 performance of network interconnect devices. IPv6 transition mechanisms are outside the scope of this document.

Table of Contents

1. Introduction	2
2. Existing Definitions	3
3. Tests and Results Evaluation	3
4. Test Environment Setup	3
5. Test Traffic	4
5.1. Frame Formats and Sizes	4
5.1.1. Frame Sizes to Be Used on Ethernet	5
5.1.2. Frame Sizes to Be Used on SONET	5
5.2. Protocol Addresses Selection	6
5.2.1. DUT Protocol Addresses	6
5.2.2. Test Traffic Protocol Addresses	7
5.3. Traffic with Extension Headers	7
5.4. Traffic Setup	9
6. Modifiers	9
6.1. Management and Routing Traffic	9
6.2. Filters	10
6.2.1. Filter Format	10
6.2.2. Filter Types	11
7. Benchmarking Tests	12
7.1. Throughput	13
7.2. Latency	13
7.3. Frame Loss	13
7.4. Back-to-Back Frames	13
7.5. System Recovery	14
7.6. Reset	14
8. IANA Considerations	14
9. Security Considerations	14
10. Conclusions	15
11. Acknowledgements	15
12. References	15
12.1. Normative References	15
12.2. Informative References	16
Appendix A. Theoretical Maximum Frame Rates Reference	17
A.1. Ethernet	17
A.2. Packet over SONET	18

1. Introduction

The benchmarking methodologies defined by RFC 2544 [9] are proving to be useful in consistently evaluating IPv4 forwarding performance of network elements. Adherence to these testing and result analysis procedures facilitates objective comparison of IPv4 performance data measured on various products and by various individuals. While the principles behind the methodologies introduced in RFC 2544 are largely IP version independent, the protocol has continued to evolve, particularly in its version 6 (IPv6).

This document provides benchmarking methodology recommendations that address IPv6-specific aspects, such as evaluating the forwarding performance of traffic containing extension headers, as defined in RFC 2460 [2]. These recommendations are defined within the RFC 2544 framework, and they complement the test and result analysis procedures as described in RFC 2544.

The terms used in this document remain consistent with those defined in "Benchmarking Terminology for Network Interconnect Devices", RFC 1242 [7]. This terminology SHOULD be consulted before using or applying the recommendations of this document.

Any methodology aspects not covered in this document SHOULD be assumed to be treated based on the RFC 2544 recommendations.

2. Existing Definitions

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 BCP 14, RFC 2119 [1]. RFC 2119 defines the use of these key words to help make the intent of standards track documents as clear as possible. While this document uses these key words, this document is not a standards track document.

3. Tests and Results Evaluation

The recommended performance evaluation tests are described in Section 7 of this document. Not all of these tests are applicable to all network element types. Nevertheless, for each evaluated device, it is recommended to perform as many of the applicable tests described in Section 6 as possible.

Test execution and results analysis MUST be performed while observing generally accepted testing practices regarding repeatability, variance, and statistical significance of small numbers of trials.

4. Test Environment Setup

The test environment setup options recommended for the IPv6 performance evaluation are the same as those described in Section 6 of RFC 2544, in both single-port and multi-port scenarios. Single-port testing measures per-interface forwarding performance, while multi-port testing measures the scalability of forwarding performance across the entire platform.

Throughout the test, the Device Under Test (DUT) can be monitored for relevant resource (processor, memory, etc.) utilization. This data could be beneficial in understanding traffic processing by each DUT and the resources that must be allocated for IPv6. It could reveal if the IPv6 traffic is processed in hardware, by applicable devices, under all test conditions, or if it is punted in the software-switched path. If such data is considered of interest, it **MUST** be collected out of band and be independent of any management data collected through the interfaces forwarding the test traffic.

Note: During testing, either static or dynamic options for neighbor discovery can be used. In the static case, the IPv6 neighbor information for the test tool is manually configured on the DUT, and the IPv6 neighbor information for the DUT is manually configured on the test tool. In the dynamic case, the IPv6 neighbor information is dynamically discovered by each device through the neighbor discovery protocol. The static option can be used as long as it is supported by the test tool. The dynamic option is preferred wherein the test tool interacts with the DUT for the duration of the test to maintain the respective neighbor caches in an active state. To avoid neighbor solicitation (NS) and neighbor advertisement (NA) storms due to the neighbor unreachability detection (NUD) mechanism [6], the test scenarios assume test traffic simulates endpoints and IPv6 source and destination addresses that are one hop beyond the DUT.

5. Test Traffic

Traffic used for all tests described in this document **SHOULD** meet the requirements described in this section. These requirements are designed to reflect the characteristics of IPv6 unicast traffic. Using the recommended IPv6 traffic profile leads to a complete evaluation of the network element performance.

5.1. Frame Formats and Sizes

Two types of media are commonly deployed, and each **SHOULD** be tested if the network element supports that type of media: Ethernet and SONET (Synchronous Optical Network). This section identifies the frame sizes that **SHOULD** be used for each media type. Refer to recommendations in RFC 2544 for all other media types.

Similar to IPv4, small frame sizes help characterize the per-frame processing overhead of the DUT. Note that the minimum IPv6 packet size (40 bytes) is larger than that of an IPv4 packet (20 bytes). Tests should compensate for this difference.

The frame sizes listed for IPv6 include the extension headers used in testing (see Section 5.3). By definition, extension headers are part of the IPv6 packet payload. Depending on the total length of the extension headers, their use might not be possible at the smallest frame sizes.

Note: Test tools commonly use signatures to identify test traffic packets to verify that there are no packet drops or out-of-order packets, or to calculate various statistics such as delay and jitter. This could be the reason why the minimum frame size selectable through the test tool might not be as low as the theoretical one presented in this document.

5.1.1. Frame Sizes to Be Used on Ethernet

Ethernet, in all its types, has become the most commonly deployed media in today's networks. The following frame sizes SHOULD be used for benchmarking over this media type: 64, 128, 256, 512, 1024, 1280, and 1518 bytes.

Note: The recommended 1518-byte frame size represents the maximum size of an untagged Ethernet frame. According to the IEEE 802.3as standard [13], the frame size can be increased up to 2048 bytes to accommodate frame tags and other encapsulation required by the IEEE 802.1AE MAC [14] security protocol. A frame size commonly used in operational environments is 1522 bytes, the max length for a VLAN-tagged frame, as per 802.1D [15].

Note: While jumbo frames are outside the scope of the 802.3 IEEE standard, tests SHOULD be executed with frame sizes selected based on the values supported by the device under test. Examples of commonly used jumbo frame sizes are: 4096, 8192, and 9216 bytes.

The maximum frame rates for each frame size and the various Ethernet interface types are provided in Appendix A.

5.1.2. Frame Sizes to Be Used on SONET

Packet over SONET (PoS) interfaces are commonly used for edge uplinks and high-bandwidth core links. Evaluating the forwarding performance of PoS interfaces supported by the DUT is recommended. The following frame sizes SHOULD be used for this media type: 47, 64, 128, 256, 512, 1024, 1280, 1518, 2048, 4096 bytes.

The theoretical maximum frame rates for each frame size and the various PoS interface types are provided in Appendix A.

5.2. Protocol Addresses Selection

There are two aspects of IPv6 benchmarking testing where IP address selection considerations **MUST** be analyzed: the selection of IP addresses for the DUT and the selection of addresses for test traffic.

5.2.1. DUT Protocol Addresses

IANA reserved an IPv6 address block for use with IPv6 benchmark testing (see Section 8). It **MAY** be assumed that addresses in this block are not globally routable, and they **MUST NOT** be used as Internet source or destination addresses.

Similar to Appendix C of RFC 2544, addresses from the first half of this range **SHOULD** be used for the ports viewed as input and addresses from the other half of the range for the output ports.

The prefix length of the IPv6 addresses configured on the DUT interfaces **MUST** fall into either of the following:

- o Prefix length is /126, which would simulate a point-to-point link for a core router.
- o Prefix length is smaller or equal to /64.

No prefix lengths **SHOULD** be selected in the range between 64 and 128 except the 126 value mentioned above.

Note that /126 prefixes might not always be the best choice for addressing point-to-point links such as back-to-back Ethernet unless the auto-provisioning mechanism is disabled. Also, not all network elements support addresses of this prefix length.

While with IPv6, the DUT interfaces can be configured with multiple global unicast addresses, the methodology described in this document does not require testing such a scenario. It is not expected that such an evaluation would bring additional data regarding the performance of the network element.

The Interface ID portion of global unicast IPv6 DUT addresses **SHOULD** be set to ::1. There are no requirements in the selection of the Interface ID portion of the link local IPv6 addresses.

It is recommended that multiple iterations of the benchmark tests be conducted using the following prefix lengths: 48, 64, 126, and 128 for the advertised traffic destination prefix. Other prefix lengths can be used. However, the indicated range reflects major prefix boundaries expected to be present in IPv6 routing tables, and they should be representative to establish baseline performance metrics.

5.2.2. Test Traffic Protocol Addresses

IPv6 source and destination addresses for the test streams SHOULD belong to the IPv6 range assigned by IANA, as defined in Section 8. The source addresses SHOULD belong to one half of the range and the destination addresses to the other, reflecting the DUT interface IPv6 address selection.

Tests SHOULD first be executed with a single stream leveraging a single source-destination address pair. The tests SHOULD then be repeated with traffic using a random destination address in the corresponding range. If the network element prefix lookup capabilities are evaluated, the tests SHOULD focus on the IPv6 relevant prefix boundaries: 0-64, 126, and 128.

Note: When statically defined neighbors are not used, the parameters affecting Neighbor Unreachability Detection should be consistently set. The IPv6 prefix-reachable time in the router advertisement SHOULD be set to 30 seconds.

5.3. Traffic with Extension Headers

Extension headers are an intrinsic part of the IPv6 architecture [2]. They are used with various types of practical traffic such as: fragmented traffic, mobile IP-based traffic, and authenticated and encrypted traffic. For these reasons, all tests described in this document SHOULD be performed with both traffic that has no extension headers and traffic that has a set of extension headers. Extension header types can be selected from the following list [2] that reflects the recommended order of multiple extension headers in a packet:

- o Hop-by-hop header
- o Destination options header
- o Routing header
- o Fragment header
- o Authentication header
- o Encapsulating security payload (ESP) header
- o Destination options header
- o Mobility header

Since extension headers are an intrinsic part of the protocol and they fulfill different roles, benchmarking of traffic containing each extension header SHOULD be executed individually.

The special processing rules for the hop-by-hop extension header require a specific benchmarking approach. Unlike other extension headers, this header must be processed by each node that forwards the traffic. Tests with traffic containing these extension header types will not measure the forwarding performance of the DUT, but rather its extension-header processing capability, which is dependent on the information contained in the extension headers. The concern is that this traffic, at high rates, could have a negative impact on the operational resources of the router, and it could be used as a security threat. When benchmarking with traffic that contains the hop-by-hop extension header, the goal is not to measure throughput [9] as in the case of the other extension headers, but rather to evaluate the impact of such traffic on the router. In this case, traffic with the hop-by-hop extension headers should be sent at 1%, 10%, and 50% of the interface total bandwidth. Device resources must be monitored at each traffic rate to determine the impact.

Tests with traffic containing each individual extension header MUST be complemented with tests containing a chain of two or more extension headers (the chain MUST NOT contain the hop-by-hop extension header). This chain should also exclude the ESP [5] extension header, since traffic with an encrypted payload cannot be used in tests with modifiers such as filters based on upper-layer information (see Section 5). Since the DUT is not analyzing the content of the extension headers, any combination of extension headers can be used in testing. The extension header chain recommended for testing is:

- o Routing header - 24-32 bytes
- o Destination options header - 8 bytes
- o Fragment header - 8 bytes

This is a real-life extension-header chain that would be found in an IPv6 packet between two mobile nodes exchanged over an optimized path that requires fragmentation. The listed extension headers' lengths represent test tool defaults. The total length of the extension header chain SHOULD be larger than 32 bytes.

Extension headers add extra bytes to the payload size of the IP packets, which MUST be factored in when used in testing at low frame sizes. Their presence will modify the minimum packet size used in

testing. For direct comparison between the data obtained with traffic that has extension headers and with traffic that doesn't have them at low frame size, a common value SHOULD be selected for the smallest frame size of both types of traffic.

For most cases, the network elements ignore the extension headers when forwarding IPv6 traffic. For these reasons, it is likely the performance impact related to extension headers will be observed only when testing the DUT with traffic filters that contain matching conditions for the upper-layer protocol information. In those cases, the DUT MUST traverse the chain of extension headers, a process that could impact performance.

5.4. Traffic Setup

All tests recommended in this document SHOULD be performed with bi-directional traffic. For asymmetric situations, tests MAY be performed with uni-directional traffic for a more granular characterization of the DUT performance. In these cases, the bi-directional traffic testing would reveal only the lowest performance between the two directions.

All other traffic profile characteristics described in RFC 2544 SHOULD be applied in this testing as well. IPv6 multicast benchmarking is outside the scope of this document.

6. Modifiers

RFC 2544 underlines the importance of evaluating the performance of network elements under certain operational conditions. The conditions defined in Section 11 of RFC 2544 are common to IPv4 and IPv6, except that IPv6 does not employ layer 2 or 3 broadcast frames. IPv6 does not use layer 2 or layer 3 broadcasts. This section provides additional conditions that are specific to IPv6. The suite of tests recommended in this document SHOULD be first executed in the absence of these conditions and then repeated under each of these conditions separately.

6.1. Management and Routing Traffic

The procedures defined in Sections 11.2 and 11.3 of RFC 2544 are applicable for IPv6 management and routing update frames as well.

6.2. Filters

The filters defined in Section 11.4 of RFC 2544 apply to IPv6 benchmarking as well. The filter definitions must be expanded to include upper-layer protocol information matching in order to analyze the handling of traffic with extension headers, which are an important architectural component of IPv6.

6.2.1. Filter Format

The filter format defined in RFC 2544 is applicable to IPv6 as well, except that the source addresses (SA) and destination addresses (DA) are IPv6 addresses. In addition to these basic filters, the evaluation of IPv6 performance SHOULD analyze the correct filtering and handling of traffic with extension headers.

While the intent is not to evaluate a platform's capability to process the various extension header types, the goal is to measure performance impact when the network element must parse through the extension headers to reach upper-layer information. In IPv6, routers do not have to parse through the extension headers (other than hop-by-hop) unless, for example, upper-layer information has to be analyzed due to filters.

To evaluate the network element handling of IPv6 traffic with extension headers, the definition of the filters must be extended to include conditions applied to upper-layer protocol information. The following filter format SHOULD be used for this type of evaluation:

```
[permit|deny] [protocol] [SA] [DA]
```

where permit or deny indicates the action to allow or deny a packet through the interface the filter is applied to. The protocol field is defined as:

- o ipv6: any IP Version 6 traffic
- o tcp: Transmission Control Protocol
- o udp: User Datagram Protocol

and SA stands for the source address and DA for the destination address.

The upper-layer protocols listed above are a recommended selection. However, they do not represent an all-inclusive list of upper-layer protocols that could be used in defining filters. The filters described in these benchmarking recommendations apply to native IPv6 traffic and upper-layer protocols (tcp, udp) transported in native IPv6 packets.

6.2.2. Filter Types

Based on RFC 2544 recommendations, two types of tests are executed when evaluating performance in the presence of modifiers: one with a single filter and another with 25 filters. Examples of recommended filters are illustrated using the IPv6 documentation prefix [11] 2001:DB8::.

Examples of single filters are:

```
Filter for TCP traffic - permit tcp 2001:DB8::1 2001:DB8::2
Filter for UDP traffic - permit udp 2001:DB8::1 2001:DB8::2
Filter for IPv6 traffic - permit ipv6 2001:DB8::1 2001:DB8::2
```

The single line filter case SHOULD verify that the network element permits all TCP/UDP/IPv6 traffic with or without any number of extension headers from IPv6 SA 2001:DB8::1 to IPv6 DA 2001:DB8::2 and deny all other traffic.

Example of 25 filters:

```
deny tcp 2001:DB8:1::1 2001:DB8:1::2
deny tcp 2001:DB8:2::1 2001:DB8:2::2
deny tcp 2001:DB8:3::1 2001:DB8:3::2
...
deny tcp 2001:DB8:C::1 2001:DB8:C::2
permit tcp 2001:DB8:99::1 2001:DB8:99::2
deny tcp 2001:DB8:D::1 2001:DB8:D::2
deny tcp 2001:DB8:E::1 2001:DB8:E::2
...
deny tcp 2001:DB8:19::1 2001:DB8:19::2
deny ipv6 any any
```

The router SHOULD deny all traffic with or without extension headers except TCP traffic with SA 2001:DB8:99::1 and DA 2001:DB8:99::2.

7. Benchmarking Tests

This document recommends the same benchmarking tests described in RFC 2544 while observing the DUT setup and the traffic setup considerations described above. The following sections state the test types explicitly, and they highlight only the methodology differences that might exist with respect to those described in Section 26 of RFC 2544.

The specificities of IPv6, particularly the definition of extension header processing, require additional benchmarking steps. The tests recommended by RFC 2544 MUST be repeated for IPv6 traffic without extension headers and for IPv6 traffic with one or multiple extension headers.

IPv6's deployment in existing IPv4 environments and the expected long coexistence of the two protocols leads network operators to place great emphasis on understanding the performance of platforms processing both types of traffic. While device resources are shared between the two protocols, it is important that IPv6-enabled platforms not experience degraded IPv4 performance. Thus, IPv6 benchmarking SHOULD be performed in the context of a stand-alone protocol as well as in the context of its coexistence with IPv4.

The modifiers defined are independent of the extension header type, so they can be applied equally to each one of the above tests.

The benchmarking tests described in this section SHOULD be performed under each of the following conditions:

Extension header specific conditions:

- i) IPv6 traffic with no extension headers
- ii) IPv6 traffic with one extension header from the list in Section 5.3
- iii) IPv6 traffic with the chain of extension headers described in Section 5.3

Coexistence-specific conditions:

- iv) IPv4 ONLY traffic benchmarking
- v) IPv6 ONLY traffic benchmarking
- vi) IPv4-IPv6 traffic mix with the ratio 90% vs 10%

vii) IPv4-IPv6 traffic mix with the ratio 50% vs 50%

viii) IPv4-IPv6 traffic mix with the ratio 10% vs 90%

Combining the test conditions listed for benchmarking IPv6 as a stand-alone protocol and the coexistence tests leads to a large-coverage matrix. At a minimum requirement, the coexistence tests should use IPv6 traffic with no extension headers and the 10%-90%, 90%-10%, or IPv4/IPv6 traffic mix.

The subsequent sections each describe specific tests that **MUST** be executed under the conditions listed above for a complete benchmarking of IPv6-forwarding performance.

7.1. Throughput

Objective: To determine the DUT throughput as defined in RFC 1242.

Procedure: Same as RFC 2544.

Reporting Format: Same as RFC 2544.

7.2. Latency

Objective: To determine the latency as defined in RFC 1242.

Procedure: Same as RFC 2544.

Reporting Format: Same as RFC 2544.

7.3. Frame Loss

Objective: To determine the frame-loss rate (as defined in RFC 1242) of a DUT throughout the entire range of input data rates and frame sizes.

Procedure: Same as RFC 2544.

Reporting Format: Same as RFC 2544.

7.4. Back-to-Back Frames

Based on the IPv4 experience, the back-to-back frames test is characterized by significant variance due to short-term variations in the processing flow. For these reasons, this test is no longer recommended for IPv6 benchmarking.

7.5. System Recovery

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

Procedure: Same as RFC 2544.

Reporting Format: Same as RFC 2544.

7.6. Reset

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

Procedure: Same as RFC 2544.

Reporting Format: Same as RFC 2544.

8. IANA Considerations

The IANA has allocated 2001:0200::/48 for IPv6 benchmarking, which is a 48-bit prefix from the RFC 4773 pool. This allocation is similar to 198.18.0.0/15, defined in RFC 3330 [10]. This prefix length (48) provides similar flexibility as the range allocated for IPv4 benchmarking, and it takes into consideration address conservation and simplicity of usage concerns. The requested size meets the requirements for testing large network elements and large emulated networks.

Note: Similar to RFC 2544 avoiding the use of RFC 1918 address space for benchmarking tests, this document does not recommend the use of RFC 4193 [4] (Unique Local Addresses) in order to minimize the possibility of conflicts with operational traffic.

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 (System Under Test).

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.

The isolated nature of the benchmarking environments and the fact that no special features or capabilities, other than those used in operational networks, are enabled on the DUT/SUT requires no security considerations specific to the benchmarking process.

10. Conclusions

The Benchmarking Methodology for Network Interconnect Devices document, RFC 2544 [9], is for the most part applicable to evaluating the IPv6 performance of network elements. This document addresses the IPv6-specific requirements that MUST be observed when applying the recommendations of RFC 2544. These additional requirements stem from the architecture characteristics of IPv6. This document is not a replacement for, but a complement to, RFC 2544.

11. Acknowledgements

Scott Bradner provided valuable guidance and recommendations for this document. The authors acknowledge the work done by Cynthia Martin and Jeff Dunn with respect to defining the terminology for IPv6 benchmarking. The authors would like to thank Bill Kine for his contribution to the initial document and to Tom Alexander, Bill Cervený, Silvija Dry, Sven Lanckmans, Dean Lee, Athanassios Liakopoulos, Benoit Lourdelet, Al Morton, David Newman, Rajiv Papejna, Dan Romascanu, and Pekka Savola for their very helpful feedback. Maryam Hamza inspired the authors to complete this document.

12. References

12.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

- [3] Malis, A. and W. Simpson, "PPP over SONET/SDH", RFC 2615, June 1999.
- [4] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [5] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [6] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

12.2. Informative References

- [7] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.
- [8] Simpson, W., Ed., "PPP in HDLC-like Framing", STD 51, RFC 1662, July 1994.
- [9] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [10] IANA, "Special-Use IPv4 Addresses", RFC 3330, September 2002.
- [11] Huston, G., Lord, A., and P. Smith, "IPv6 Address Prefix Reserved for Documentation", RFC 3849, July 2004.
- [12] Newman, D. and T. Player, "Hash and Stuffing: Overlooked Factors in Network Device Benchmarking", RFC 4814, March 2007.
- [13] LAN/MAN Standards Committee of the IEEE Computer Society, "IEEE Std 802.3as-2006, Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, Amendment 3: Frame format extensions", November 2006.
- [14] Allyn Romanow (editor), "IEEE Std 802.3ae, Media Access Control (MAC) Security", June 2006.
- [15] Mick Seaman (editor), "IEEE Std 802.1D-2004, MAC Bridges", February 2004.

Appendix A. Theoretical Maximum Frame Rates Reference

This appendix provides the formulas to calculate and the values for the theoretical maximum frame rates for two media types: Ethernet and SONET.

A.1. Ethernet

The throughput in frames per second (fps) for various Ethernet interface types and for a frame size X can be calculated with the following formula:

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

The 20 bytes in the formula is the sum of the preamble (8 bytes) and the inter-frame gap (12 bytes). The throughput for various Ethernet interface types and frame sizes:

Size Bytes	10Mb/s pps	100Mb/s pps	1000Mb/s pps	10000Mb/s pps
64	14,880	148,809	1,488,095	14,880,952
128	8,445	84,459	844,594	8,445,945
256	4,528	45,289	452,898	4,528,985
512	2,349	23,496	234,962	2,349,624
1024	1,197	11,973	119,731	1,197,318
1280	961	9,615	96,153	961,538
1518	812	8,127	81,274	812,743
1522	810	8,106	81,063	810,635
2048	604	6,044	60,444	604,448
4096	303	3,036	30,396	303,691
8192	152	1,522	15,221	152,216
9216	135	1,353	13,534	135,339

Note: Ethernet's maximum frame rates are subject to variances due to clock slop. The listed rates are theoretical maximums, and actual tests should account for a +/- 100 ppm tolerance.

A.2. Packet over SONET

ANSI T1.105 SONET provides the formula for calculating the maximum available bandwidth for the various Packet over SONET (PoS) interface types:

STS-Nc (N = 3Y, where Y=1,2,3,etc)

$$[(N*87) - N/3]*(9 \text{ rows})*(8 \text{ bit/byte})*(8000 \text{ frames/sec})$$

Packets over SONET can use various encapsulations: PPP [3], High-Level Data Link Control (HDLC) [8], and Frame Relay. All these encapsulations use a 4-byte header, a 2- or 4-byte Frame Check Sequence (FCS) field, and a 1-byte Flag that are all accounted for in the overall frame size. The maximum frame rate for various interface types can be calculated with the formula (where X represents the frame size in bytes):

Line Rate (bps)

(8bits/byte)*(X+1)bytes/frame

The theoretical maximum frame rates for various PoS interface types and frame sizes:

Size Bytes	OC-3c fps	OC-12c fps	OC-48c fps	OC-192c fps	OC-768c fps
47	390,000	1,560,000	6,240,000	24,960,000	99,840,000
64	288,000	1,152,000	4,608,000	18,432,000	73,728,000
128	145,116	580,465	2,321,860	9,287,441	37,149,767
256	72,840	291,361	1,165,447	4,661,789	18,647,159
512	36,491	145,964	583,859	2,335,438	9,341,754
1024	18,263	73,053	292,214	1,168,858	4,675,434
2048	9,136	36,544	146,178	584,714	2,338,857
4096	4,569	18,276	73,107	292,428	1,169,714

It is important to note that throughput test results may vary from the values presented in Appendices A.1 and A.2 due to bit stuffing performed by various media types [12]. The theoretical throughput numbers were rounded down.

Authors' Addresses

Ciprian Popoviciu
Cisco Systems
Kit Creek Road
RTP, North Carolina 27709
USA

Phone: 919 787 8162
EMail: cpopovic@cisco.com

Ahmed Hamza
Cisco Systems
3000 Innovation Drive
Kanata K2K 3E8
Canada

Phone: 613 254 3656
EMail: ahamza@cisco.com

Gunter Van de Velde
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2704 5473
EMail: gunter@cisco.com

Diego Dugatkin
FastSoft, Inc.
150 S. Los Robles Ave.
Pasadena, CA 91101
USA

Phone: +1-626-357-7012
EMail: diego@fastsoft.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

