

Internet Engineering Task Force (IETF)
Request for Comments: 8161
Category: Informational
ISSN: 2070-1721

W. Cervený
Arbor Networks
R. Bonica
R. Thomas
Juniper Networks
May 2017

Benchmarking the Neighbor Discovery Protocol

Abstract

This document provides benchmarking procedures for the Neighbor Discovery Protocol (NDP). It also proposes metrics by which an NDP implementation's scaling capabilities can be measured.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8161>.

Copyright Notice

Copyright (c) 2017 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | |
|---|----|
| 1. Introduction | 2 |
| 1.1. Requirements Language | 4 |
| 2. Test Setup | 4 |
| 2.1. Device Under Test (DUT) | 4 |
| 2.1.1. Interfaces | 4 |
| 2.1.2. Neighbor Discovery Protocol (NDP) | 5 |
| 2.1.3. Routing | 5 |
| 2.2. Tester | 6 |
| 2.2.1. Interfaces | 6 |
| 2.2.2. Neighbor Discovery Protocol (NDP) | 6 |
| 2.2.3. Routing | 6 |
| 2.2.4. Test Traffic | 7 |
| 2.2.5. Counters | 8 |
| 3. Tests | 8 |
| 3.1. Baseline Test | 8 |
| 3.1.1. Procedure | 9 |
| 3.1.2. Baseline Test Procedure Flow Chart | 9 |
| 3.1.3. Results | 11 |
| 3.2. Scaling Test | 11 |
| 3.2.1. Procedure | 11 |
| 3.2.2. Scaling Test Procedure Flow Chart | 13 |
| 3.2.3. Results | 15 |
| 4. Measurements Explicitly Excluded | 15 |
| 4.1. DUT CPU Utilization | 15 |
| 4.2. Malformed Packets | 15 |
| 5. IANA Considerations | 16 |
| 6. Security Considerations | 16 |
| 7. Normative References | 16 |
| Acknowledgments | 16 |
| Authors' Addresses | 17 |

1. Introduction

When an IPv6 node forwards a packet, it executes the following procedure:

- o Identifies the outbound interface and IPv6 next hop.
- o Queries a local Neighbor Cache (NC) to determine the IPv6 next hop's link-layer address.
- o Encapsulates the packet in a link-layer header. The link-layer header includes the IPv6 next hop's link-layer address.
- o Forwards the packet to the IPv6 next hop.

IPv6 nodes use the Neighbor Discovery Protocol (NDP) [RFC4861] to maintain the NC. Operational experience [RFC6583] shows that when an implementation cannot maintain a sufficiently complete NC, its ability to forward packets is impaired.

NDP, like any other protocol, consumes processing, memory, and bandwidth resources. Its ability to maintain a sufficiently complete NC depends upon the availability of the above-mentioned resources.

This document provides benchmarking procedures for NDP. Benchmarking procedures include a Baseline Test and an NDP Scaling Test. In both tests, the Device Under Test (DUT) is an IPv6 router. Two physical links (A and B) connect the DUT to a Tester. The Tester sends traffic through Link A to the DUT. The DUT forwards that traffic, through Link B, back to the Tester.

The above-mentioned traffic stream contains one or more interleaved flows. An IPv6 Destination Address uniquely identifies each flow. Or, said another way, every packet within a flow has the same IPv6 Destination Address.

In the Baseline Test, the traffic stream contains exactly one flow. Because every packet in the stream has the same IPv6 Destination Address, the DUT can forward the entire stream using exactly one NC entry. NDP is exercised minimally, and no packet loss should be observed.

The NDP Scaling Test is identical to the Baseline Test, except that the traffic stream contains many flows. In order to forward the stream without loss, the DUT must maintain one NC entry for each flow. If the DUT cannot maintain one NC entry for each flow, packet loss will be observed and attributed to NDP scaling limitations.

This document proposes an NDP scaling metric, called NDP-MAX-NEIGHBORS. NDP-MAX-NEIGHBORS is the maximum number of neighbors to which an IPv6 node can send traffic during periods of high NDP activity.

The procedures described herein reveal how many IPv6 neighbors an NDP implementation can discover. They also provide a rough estimate of the time required to discover those neighbors. However, that estimate does not reflect the maximum rate at which the implementation can discover neighbors. Maximum rate discovery is a topic for further exploration.

The test procedures described herein assume that NDP does not compete with other applications for resources on the DUT. When NDP competes for resources, its scaling characteristics may differ from those reported by the benchmarks described and may vary over time.

1.1. Requirements Language

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 RFC 2119 [RFC2119].

2. Test Setup




Figure 1: Test Setup

The DUT is an IPv6 router. Two links (A and B) connect the DUT to the Tester. Link A capabilities must be identical to Link B capabilities. For example, if the interface to Link A is a 10 Gigabit Ethernet port, the interface to Link B must also be a 10 Gigabit Ethernet port.

2.1. Device Under Test (DUT)

2.1.1. Interfaces

DUT interfaces are numbered as follows:

- o Link A - 2001:2:0:0::2/64
- o Link B - 2001:2:0:1::1/64

Both DUT interfaces should be configured with a 1500-byte MTU. However, if they cannot support a 1500-byte MTU, they may be configured with a 1280-byte MTU.

2.1.2. Neighbor Discovery Protocol (NDP)

NDP is enabled on both DUT interfaces. Therefore, the DUT emits both solicited and unsolicited Router Advertisement (RA) messages. The DUT emits an RA message at least once every 600 seconds and no more frequently than once every 200 seconds.

When the DUT sends an RA message, it includes the following information:

- o Router Lifetime - 1800 seconds
- o Reachable Time - 0 seconds
- o Retrans Time - 0 seconds
- o Source Link-Layer Address - link-layer address of DUT interface
- o M-bit is clear (0)
- o O-bit is clear (0)

The above-mentioned values are chosen because they are the default values specified in RFC 4861.

NDP manages the NC. Each NC entry represents an on-link neighbor and is identified by the neighbor's on-link unicast IP address. As per RFC 4861, each NC entry needs to be refreshed periodically. NDP refreshes NC entries by exchanging Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages.

No static NC entries are configured on the DUT.

2.1.3. Routing

The DUT maintains a direct route to 2001:2:0:0/64 through Link A. It also maintains a direct route to 2001:2:0:1/64 through Link B. No static routes or dynamic routing protocols are configured on the DUT.

2.2. Tester

2.2.1. Interfaces

Interfaces are numbered as follows:

- o Link A - 2001:2:0:0::1/64
- o Link B - Multiple addresses are configured on Link B. These addresses are drawn sequentially from the 2001:2:0:1::/64 address block. The first address is 2001:2:0:1::2/64. Subsequent addresses are 2001:2:0:1::3/64, 2001:2:0:1::4/64, 2001:2:0:1::5/64, etc. The number of configured addresses should be the expected value of NDP-MAX-NEIGHBORS times 1.1.

Both Tester interfaces should be configured with a 1500-byte MTU. However, if they cannot support a 1500-byte MTU, they may be configured with a 1280-byte MTU.

2.2.2. Neighbor Discovery Protocol (NDP)

NDP is enabled on both Tester interfaces. Therefore, upon initiation, the Tester sends Router Solicitation (RS) messages and waits for Router Advertisement (RA) messages. The Tester also exchanges Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages with the DUT.

No static NC entries are configured on the Tester.

2.2.3. Routing

The Tester maintains a direct route to 2001:2:0:0/64 through Link A. It also maintains a direct route to 2001:2:0:1/64 through Link B. No static routes or dynamic routing protocols are configured on the Tester.

2.2.4. Test Traffic

The Tester sends a stream of test traffic through Link A to the DUT. The test traffic stream contains one or more interleaved flows. Flows are numbered 1 through N, sequentially.

Within each flow, each packet contains an IPv6 header, and each IPv6 header contains the following information:

- o Version - 6
- o Traffic Class - 0
- o Flow Label - 0
- o Payload Length - 0
- o Next Header - IPv6-NoNxt (59)
- o Hop Limit - 255
- o Source Address - 2001:2:0:0::1
- o Destination Address - The first 64 bits of the Destination Address are 2001:2:0:1::. The next 64 are uniquely associated with the flow. Every packet in the first flow carries the Destination Address 2001:2:0:1::2. Every subsequent flow has an IP address one greater than the last (i.e., 2001:2:0:1::3, 2001:2:0:1::4, etc.).

In order to avoid link congestion, test traffic is offered at a rate not to exceed 50% of available link bandwidth. In order to avoid burstiness and buffer occupancy, every packet in the stream is exactly 40 bytes long (i.e., the length of an IPv6 header with no IPv6 payload). Furthermore, the gap between packets is identical.

During the course of a test, the number of flows that the test stream contains may increase. When this occurs, the rate at which test traffic is offered remains constant. For example, assume that a test stream is offered at a rate of 1,000 packets per second. This stream contains two flows, each contributing 500 packets per second to the 1,000 packet per second aggregate. When a third stream is added to the flow, all three streams must contribute 333 packets per second in order to maintain the 1,000 packet per second limit. (As in this example, rounding error is acceptable.)

The DUT attempts to forward every packet in the test stream through Link B to the Tester. It does this because:

- o Every packet in the test stream has a Destination Address drawn from the 2001:2:0:1::/64 address block.
- o The DUT has a direct route to 2001:2:0:1/64 through Link B.

2.2.5. Counters

On the Tester, two counters are configured for each flow. One counter, configured on Link A, increments when the Tester sends a packet belonging to the flow. The other counter, configured on Link B, increments when the Tester receives a packet from the flow. In order for a packet to be associated with a flow, the following conditions must all be true:

- o The IPv6 Destination Address must be that of the flow.
- o The IPv6 Next Header must be IPv6-NoNxt (59).

The following counters also are configured on both Tester Interfaces:

- o RS packets sent
- o RA packets received
- o NS packets sent
- o NS packets received
- o NA packets sent
- o NA packets received
- o Total packets sent
- o Total packets received

3. Tests

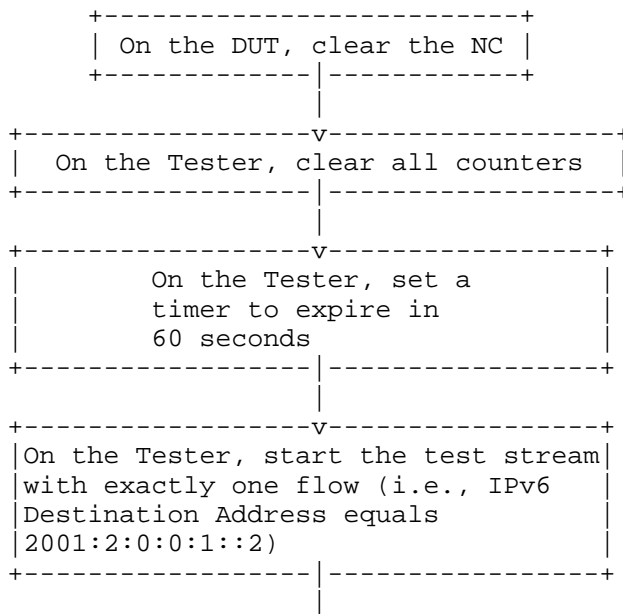
3.1. Baseline Test

The purpose of the Baseline Test is to ensure that the DUT can forward every packet in the test stream, without loss, when NDP is minimally exercised and not operating near its scaling limit.

3.1.1.1. Procedure

- o On the DUT, clear the NC.
- o On the Tester, clear all counters.
- o On the Tester, set a timer to expire in 60 seconds.
- o On the Tester, start the test stream with exactly one flow (i.e., IPv6 Destination Address equals 2001:2:0:1::2).
- o Wait for either the timer to expire or the packets-received counter associated with the flow to increment.
- o If the timer expires, stop the test stream and end the test.
- o If the packets-received counter increments, pause the traffic stream, log the initial counter values, clear the counters, reset the timer to expire in 1800 seconds, and restart the traffic stream.
- o When the timer expires, stop the test stream, wait sufficient time for any queued packets to exit, log the final counter values, and end the test.

3.1.1.2. Baseline Test Procedure Flow Chart



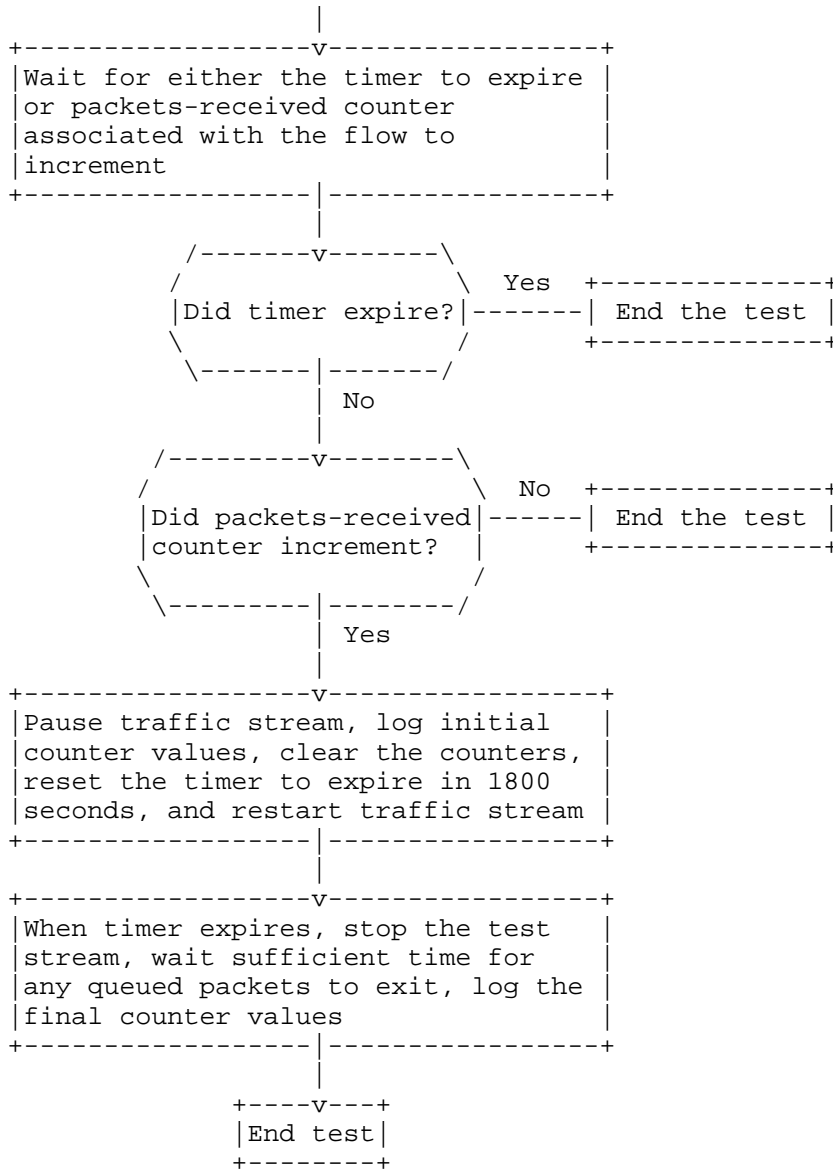


Figure 2: Baseline Test Procedure Flow Chart

3.1.3. Results

The log contains initial and final values for the following counters:

- o packets-sent
- o packets-received

The initial values of packets-sent and packets-received may be equal to one another. If these values are identical, none of the initial packets belonging to the flow were lost. However, if the initial value of packets-sent is greater than the initial value of packets-received, initial packets were lost. This loss of initial packets is acceptable.

The final values of packets-sent and packets-received should be equal to one another. If they are not, an error has occurred. Because this error is likely to affect Scaling Test results, the error must be corrected before the Scaling Test is executed.

3.2. Scaling Test

The purpose of the Scaling Test is to discover the number of neighbors to which an IPv6 node can send traffic during periods of high NDP activity. We call this number NDP-MAX-NEIGHBORS.

3.2.1. Procedure

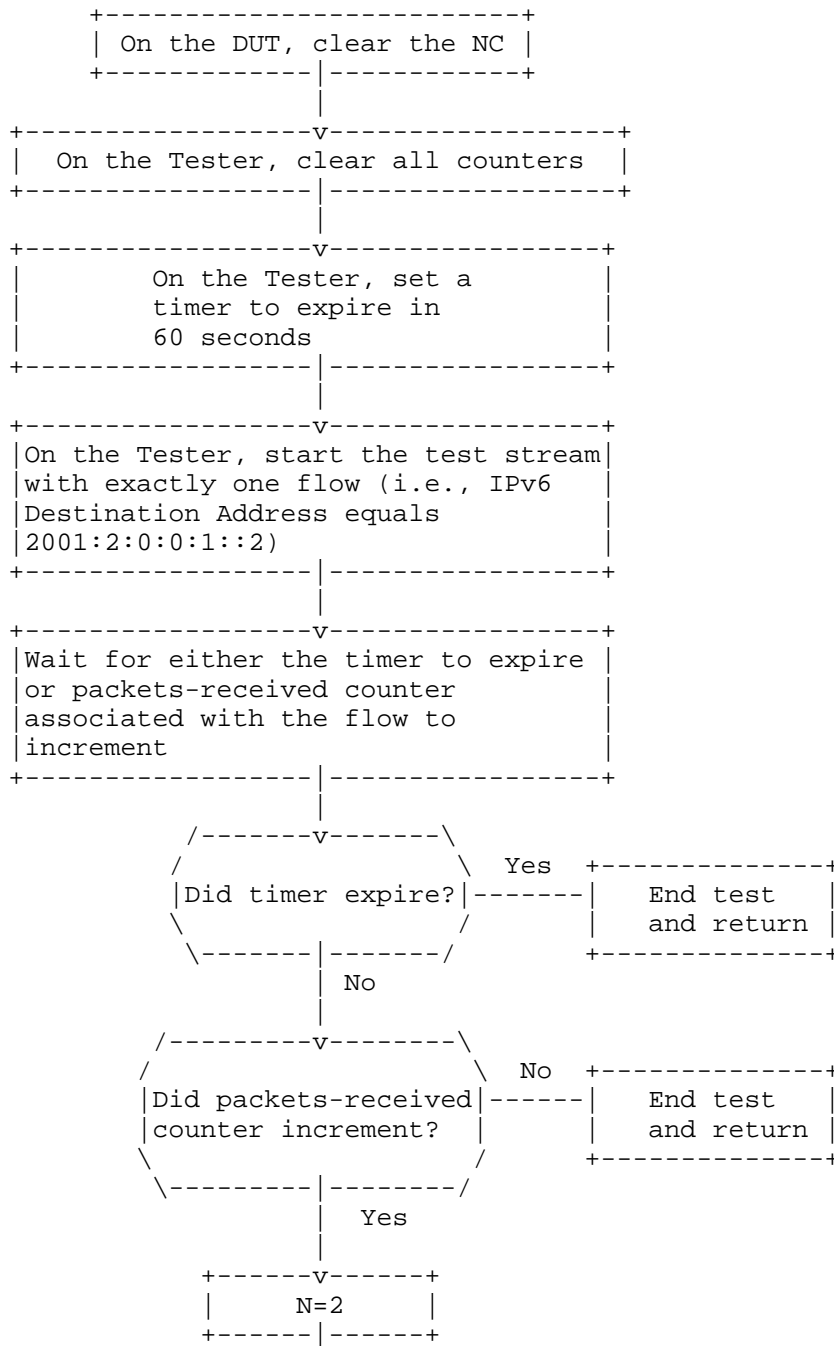
Execute the following procedure:

- o On the DUT, clear the NC.
- o On the Tester, clear all counters.
- o On the Tester, set a timer to expire in 60 seconds.
- o On the Tester, start the test stream with exactly one flow (i.e., IPv6 Destination Address equals 2001:2:0:1::2).
- o Wait for either the timer to expire or the packets-received counter associated with the flow to increment.
- o If the timer expires, stop the test stream and end the test.

- o If the packets-received counter increments, execute the following procedure N times, starting at 2 and ending at the expected value of NDP-MAX-NEIGHBORS times 1.1.
 - * Pause the test stream.
 - * Log the time and the value of N minus one.
 - * Clear the packets-sent and packets-received counters associated with the previous flow (i.e., N minus one).
 - * Reset the timer to expire in 60 seconds.
 - * Add the next flow to the test stream (i.e., IPv6 Destination Address is a function of N).
 - * Restart the test stream.
 - * Wait for either the timer to expire or the packets-received counter associated with the new flow to increment.

After the procedure described above has been executed N times, clear the timer and reset it to expire in 1800 seconds. When the timer expires, stop the stream, log all counters, and end the test (after waiting sufficient time for any queued packets to exit).

3.2.2. Scaling Test Procedure Flow Chart



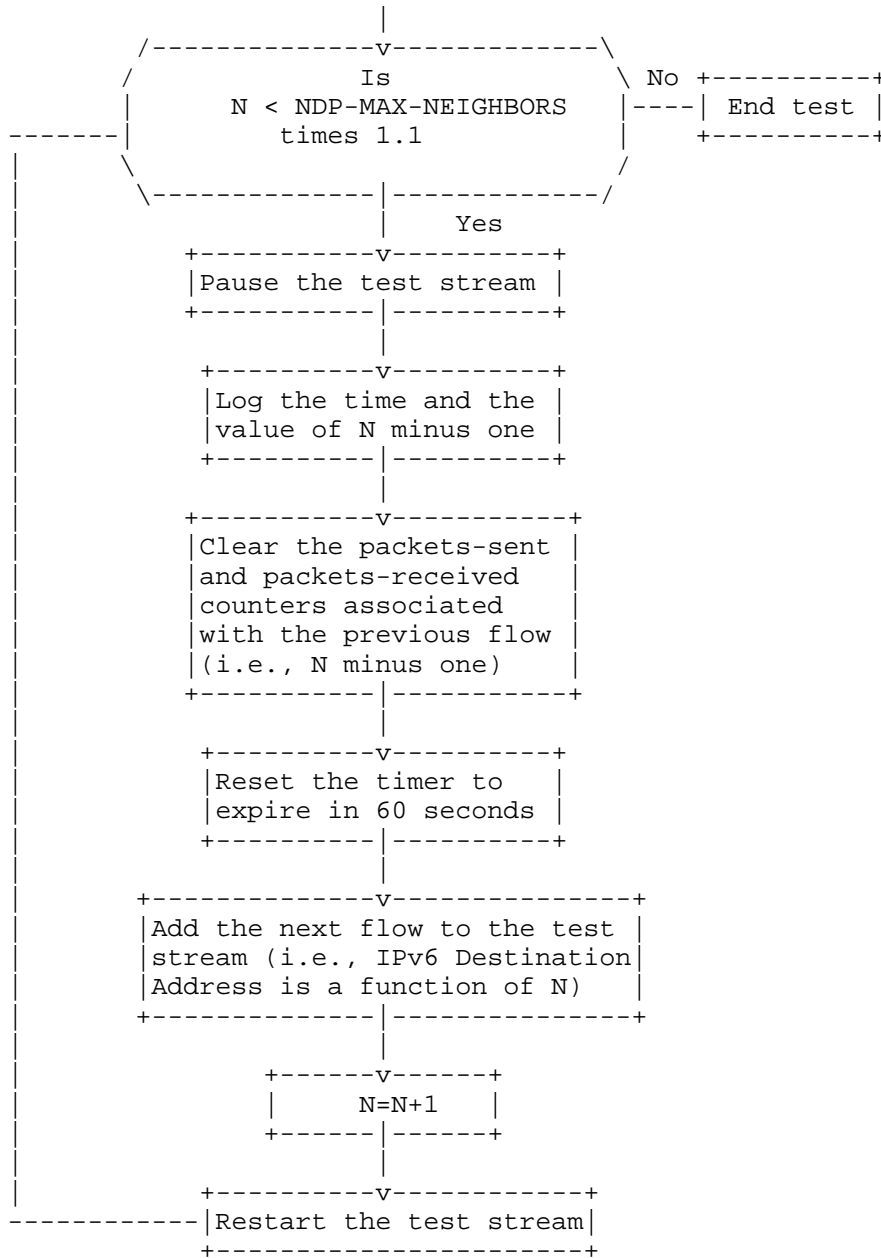


Figure 3: Scaling Test Procedure Flow Chart

3.2.3. Results

The test report includes the following:

- o A description of the DUT (make, model, processor, memory, and interfaces)
- o Rate at which the Tester offers test traffic to the DUT (measured in packets per second)
- o A log that records the time at which each flow was introduced to the test stream and the final value of all counters
- o The expected value of NDP-MAX-NEIGHBORS
- o The actual value of NDP-MAX-NEIGHBORS

NDP-MAX-NEIGHBORS is equal to the number of counter pairs where packets-sent is equal to packets-received. Two counters are members of a pair if they are both associated with the same flow. If packets-sent is equal to packets-received for every counter pair, the test should be repeated with a larger expected value of NDP-MAX-NEIGHBORS.

If an implementation abides by the recommendation of Section 7.1 of RFC 6583, for any given counter pair, packets-received will either be equal to zero or packets-sent.

The log documents the time at which each flow was introduced to the test stream. This log reveals the effect of NC size to the time required to discover a new IPv6 neighbor.

4. Measurements Explicitly Excluded

These measurements aren't recommended because of the itemized reasons below:

4.1. DUT CPU Utilization

This measurement relies on the DUT to provide utilization information, which is not externally observable (not black-box). However, some testing organizations may find the CPU utilization is useful auxiliary information specific to the DUT model, etc.

4.2. Malformed Packets

This benchmarking test is not intended to test DUT behavior in the presence of malformed packets.

5. IANA Considerations

This document does not require any IANA actions.

6. 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 or System Under Test (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.

7. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, DOI 10.17487/RFC6583, March 2012, <<http://www.rfc-editor.org/info/rfc6583>>.

Acknowledgments

Helpful comments and suggestions were offered by Al Morton, Joel Jaeggli, Nalini Elkins, Scott Bradner, and Ram Krishnan on the BMWG email list and at BMWG meetings. Precise grammatical corrections and suggestions were offered by Ann Cerveny.

Authors' Addresses

Bill Cervený
Arbor Networks
2727 South State Street
Ann Arbor, MI 48104
United States of America

Email: wcervený@arbor.net

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20170
United States of America

Email: rbonica@juniper.net

Reji Thomas
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore, KA 560103
India

Email: rejithomas@juniper.net

