Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 12, 2013

W.J. Cerveny
Arbor Networks
March 11, 2013

Benchmarking Neighbor Discovery Problems
draft-cerveny-bmwg-ipv6-nd-00

Abstract

This document is a benchmarking instantiation of RFC 6583: "Operational Neighbor Discovery Problems". It describes a general testing procedure and measurements that can be performed to evaluate how the problems described in RFC 6583 may impact the functionality or performance of intermediate nodes.

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].

Table of Contents

1. Introduction

This document is a benchmarking instantiation of RFC 6583: "Operational Neighbor Discovery Problems" [RFC6583]. It describes a general testing procedure and measurements that can be performed to evaluate how the problems described in RFC 6583 may impact the functionality or performance of intermediate nodes.

2. Terminology

Neighbor Discovery See Section 1 of RFC 4861 [RFC4861]

NDP Triggering Event An event which forces the DUT (Device Under Test) to perform a neighbor solicitation. A triggering event could be an ICMPv6 echo request, but could also be any other packets which require discovering the MAC address of existing and non-existing nodes on an IPv6 subnet.

Scanner Network The network from which the scanning device is connected.

Target Network The network for which the scanner is targeting its scans.

Scanning Node The node which is conducting the scanning activity.

Target Network Measurement Node A node that resides on the target network, which is primarily used to measure DUT performance while the scanning activity is occurring.

Non-participating Measurement Node A node on a network directly connected to the DUT, but this node is not in the target network nor the scanner network.

3. Test Set-up

3.1. Device Under Test (DUT)

For purposes of this document, the intermediate node will be referred to as the device under test (DUT). The DUT may be any intermediate node which retains a neighbor cache. The tests in this document could also be completed with any intermediate node which maintains a list of addresses that traverse the intermediate node, although not all measurements and performance characteristics may apply.

3.2. Test Network

The test network design is fairly simple. The network needs to minimally have two subnets: one from which the scanner(s) source their scanning activity and the other which is the target network of the address scans.

It is assumed that the latency for all network segments is neglible.

At least one node should reside on the target network to confirm some of the performance characteristics.

Basic format of test network. Note that optional "non-participating node" is illustrated connected via a third network not related to the scanner or target network.

```
+---------------+            +-----------+              +--------------+
|               |   Scanner  |           |   Target     |              |
|   Scanning    |-------------|    DUT    |--------------|Target Network|
|     node      |   Network  |           |   Network    |     node     |
|               |            |           |              |              |
+---------------+            +-----------+              +--------------+
                                   |
                                   |
                         +-----------------+
                         |                 |
                         |Non-participating|
                         |      node       |
                         |                 |
                         +-----------------+
```

4. Modifiers (variables)

4.1. Frequency of NDP triggering events

The frequency of NDP triggering events could be as high as the maximum packet per second rate that the iscanner network will support (or is rated for). However, it may not be necessary to send packets at a particularly high rate and in fact a goal of testing could be to identify if the DUT is able to withstand scans at rates which otherwise would not impact the performance of the DUT.

Optimistically, the scanning rate should be incremented until the DUT’s performance begins deteriorating. Depending on the software and system being used to implement the scanning, it may be challenging to achieve a sufficient rate.

The lowest frequency is the lowest rate for which packets could be expected to have an impact on the DUT -\u002D this value is of course, subjective.

4.2. Prefix Length

The target network's subnet shall be 64-bits in length. It may be interesting to gauge performance when the subnet length is varied from 64-bits.

4.3. Duration of Test

The duration of the test needs to be evaluated

4.4. Packet Size

Although packet size shouldn't have a direct impact, packet per second (pps) rates will have an impact and smaller packet sizes should be utilized to facilitate higher packet per second rates.

4.5. Packet Type

For purposes of this test, the packet type being sent by the scanning device isn't important, although most scanning applications might want to send packets that would elicit responses from nodes within a subnet. Since it is not intended that responses be evoked from the target network node, such packets aren't necessary.

The hop limit for the scanning packets should be set to 2, to reduce the likelihood that scanning packets would escape the test network.

4.6. Packet Addressing

The destination address for the packet should be an address within the target network. While each packet sent should have a unique destination address in the destination network, it isn't clear if it matters what the sequence of addresses is. For purposes of thoroughness, it may be desirable to send each packet with a random address within the target network's address space.

The source address for the packet may be the same for all scanning packets. However, it may be interesting to vary the source address during the scanning activity

4.7. Testing of Mitigating Options

It may be desirable to perform some tests in the presence of mitigating techniques described in RFC 6583 [RFC6583]

4.8. Attack where node in target network responds to all neighbor solicitations

[Open Question: Is this an interesting condition, where a device on the network responds affirmatively to all incoming NDP requests?? Are there any non-malicious cases where this could happen?]

5. Exclusions

This benchmarking test is not intended to test DUT behavior in the presence of malformed packets, such as packets which do not confirm to designs consistent with IETF standards.

6. Measurements

6.1. Round-trip time across DUT

This consists of pinging the target network measurement node from a non-participating measurement node and recording reported round-trip time. This measurement should be conducted with an address not yet present in the DUT's neighbor cache. This measurement is included because it is perhaps the easiest to conduct and capture.

6.2. Rate DUT adds a valid node in the target network to its neighbor cache

There are three distinct time elements associated with this measurement:

1. The difference in time for which the DUT receives the packet which must be forwarded to a node in the target network not yet listed in the neighbor cache and the time the DUT sends a neighbor solicitation.

2. The difference in time between when the target network measurement node receives the neighbor solicitation and the time the target network measurement node responds with a neighbor advertisement. This time is outside the control of the DUT and measurements should account for this time if it is significant.

3. The difference in time from which the DUT receives the packet to the time for which the DUT adds the neighbor in its neighbor cache.

The first time element may be measurable via a device which can observe packets on both the scanner network and the target network. The second time element may be measured by monitoring the target network and observing the specific neighbor solicitation for the node and the node’s solicited[Is this the right term?] neighbor advertisement.

Of the above time elements, the third is perhaps the hardest to measure for times smaller than a few seconds.

A challenge with this measurement is to conduct it where the target network node has an address that is not in the DUT’s neighbor cache in any state (such as "INCOMPLETE"). As tested with a router, the router’s "clear neighbor cache" command did not always flush the target network node’s neighbor entry. One method of implementing this may be to configure the target network node with sufficient addresses for a unique NDP request per test interval.

6.3. Adherence to prioritization of NDP activity prioritization

As discussed in RFC 6583 [RFC6583], this measurement would require confirming that a set prioritization is adhered to. [Insert more text here.]

6.4. DUT CPU utilization

Measured in percent utilization, captured via a non-intrusive query of the DUT.

6.5. Rate DUT forwards packets

This measures the impact that the scan may have on the DUT’s ability to forward packets. The measurement should be documented in packets per seconds (pps) or (bps). However, if the DUT handles NDP in the "management plane" and packets are forwarded in a separate "forwarding plane", the scanning tests described in this document may not have any impact on the DUT’s ability to forward packets.

It may be beneficial to conduct two RFC 5180 [RFC5180] style throughput tests even if it is assumed that scanning activity won’t have any bearing on the DUT’s packet forwarding capabilities:

1. Baseline test without any scanning activity.

2. Test while worse-case scanning activity is occuring.

6.6. Rate DUT responds to neighbor solicitations for its own address

This is the difference in time from when a node on the target network network sends a neighbor solicitation for the DUT’s MAC address and when the DUT responds with a neighbor advertisement in response to the neighbor solicitation. This can be determined by observing the target network and measuring the difference in time (in milliseconds) between when the neighbor solicitation leaves the target network measurement node and when the solicited neighbor advertisement is returned from the DUT.

6.7. Impact on unaffected interfaces/subnets

This measurement would require having a node on a network directly connected to the DUT, but not on either the scanner network or target network. Although not itemized, this measurement could consist of any combination of measurements which are conducted relating to the target network.

6.8. Maximum number of enteries in the DUT’s neighbor cache

This measurement confirms how many entries can effectively reside in the DUT’s neighbor cache. This measurement would support or refute any value documented by the DUT manufacturer. [Need to describe how this is done.]

7. Measurement Interval

To be determined.

8. DUT initialization

At the beginning of each test, the neighbor cache of the DUT should be initialized

9. General Test Procedure

This test can be completed with publicly available scanning software. The methodology to implement this scan is fairly straightforward and could be implemented using open-source network scripting tools.

The algorithm for such a scanner could be as simple as:

```
Dest_address = <ip prefix>::1000

While True:

Send(ICMPv6(dst=Dest_address)

Dest_address = Dest_address + 1
```

As described in [RFC6583], four instances of a scanner on a single computer was able to impact the performance of high-end routers. If multiple scanner instances are used, the starting address should be in different "regions" of the subnet.

Some existing software for completing network scans is discussed in [RFC6583], although other applications may exist.

Although not tested, commercial network testing solutions may be effectively implemented and may provide dedsired throughput.

10. Other Potential Testing Scenarios

10.1. Exhaustion of Address Tables (NCE) in Intermediate Nodes

[Question: Where a large number of addresses are being scanned for, would there be an impact on intermediate nodes, such as firewalls?]

10.2. Link-local network attack

In this attack, a node in the subnet simulates a condition where it is sending packets to every address in the subnet and where the destination MAC address is the DUT[Is this an allowed scenario?]. In this scenario, it "could" be possible to send neighbor solicition messages to every link local address via the default gateway.

11. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

12. 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.

13. Acknowledgements

14. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

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

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, March 2012.

Author's Address

Bill Cerveny
Arbor Networks