

Benchmarking Neighbor Discovery (draft-cerveney-ippm-nd- benchmarking)

March 7, 2014

History

- Suggested by Ron Bonica at IETF 85 BMWG meeting
- Draft v00 presented and discussed at IETF 86 (Orlando)
- Draft v02 discussed at IETF 88 (Vancouver)
- Draft v04 completed / submitted February 2014

Neighbor Discovery (ND) Problem Background

- The problem is described and documented in RFC 6583, “Operational Neighbor Discovery Problems.”
- An IPv4 subnet is “typically” no larger than 510 addresses and scanning is relatively quick.
- Since the default size of any IPv6 user subnet is 2^{64} , there can be a lot of addresses
- Scanning the IPv6 subnet takes a really long time, but one can still start scanning it.

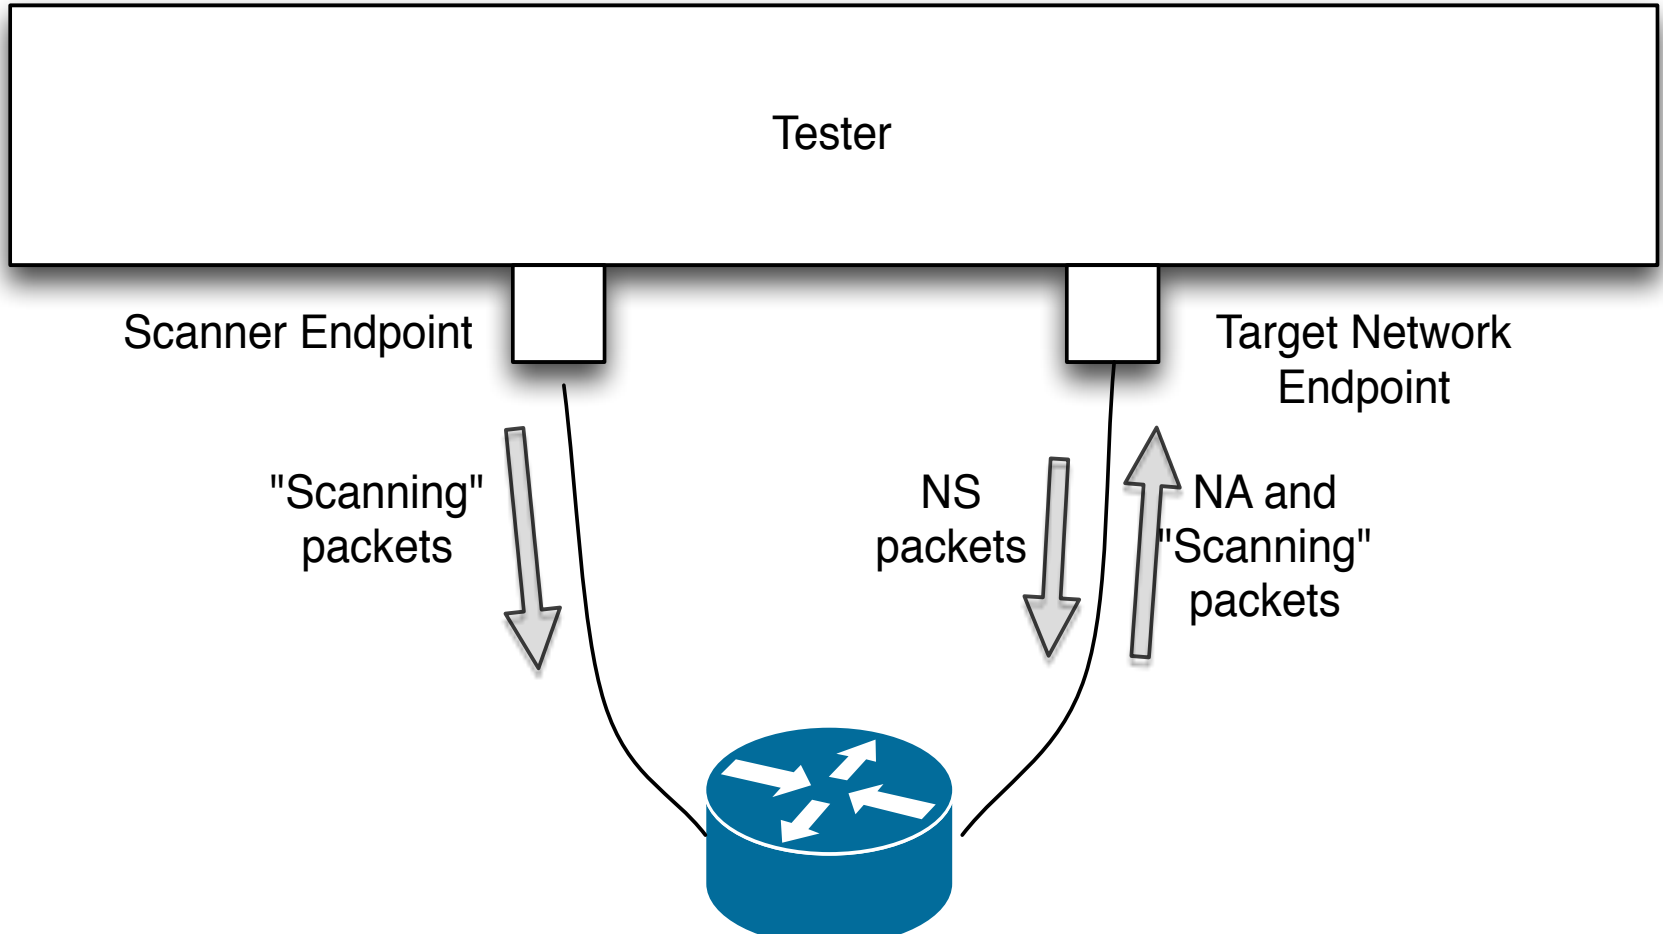
ND Problem con't

- The number of addresses one can scan for is limited only by the available bandwidth.
- The DUT (router or intermediate node) needs to perform ND for the addresses being scanned, even if the addresses aren't "live" in the subnet
- This can create a lot of state in the DUT, so much so that the DUT may be unable to complete ND for real, valid nodes in subnet.

Benchmarking Neighbor Discovery

- Build a network and a set of tests which illustrate the neighbor discovery problems.
- Create measurements which characterize how a node behaves under stress due to heavy neighbor discovery activity.

Basic Test Network and Methodology



What's New in -04

- Concept of non-participating network was removed
- Reworked and simplified tests
- Grammatical corrections (thanks to Ann Cervený)
- “Problems” removed from draft name
- Ron Bonica added as co-author

Tests Defined in v04

- Stale Entry Time Determination
- Neighbor Cache Exhaustion Determination
- Determine Neighbor Discovery Behavior During Address Scans
- Pre-established Flow Treatment
- Stopped Flow Recovery Behavior

Proposed Charter Text

IPv6 Neighbor Discovery-related benchmarking

Large address space in IPv6 subnets presents several networking challenges, as described in RFC 6583. Indexes to describe the performance of network devices, such as the number of reachable devices on a sub-network, are useful benchmarks to the operations community. The working group will develop the necessary terminology and methodologies to measure such benchmarks.