

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

Authors:

Bill Cerveney, Ron Bonica

BMWG - IETF 92

March 24, 2015

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
- Draft v05 (minor tweaks to test cases) submitted June 2014
- Draft v06 (clarification on benchmarks in tests) submitted January 2015


Neighbor Discovery (ND) Problem Background

- The problem is described and documented in RFC 6583, “Operational Neighbor Discovery Problems.”
- Scan a typical IPv4 subnet ($2^{**}8$ addresses), no “harm” done.
- Start scanning a typical IPv6 subnet ($2^{**}64$ addresses), create excessive neighbor cache state on intermediate node, impacting node’s forwarding

Benchmarking Neighbor Discovery

- Creates a methodology, including measurements, which characterize how a node behaves under stress due to heavy neighbor discovery activity.

Basic Test Network and Methodology



Test Summary

Test	Measurement
Stale Entry Time Determination	Time interval when the intermediate node (DUT) identifies an address as stale
Neighbor Cache Exhaustion Determination	Interim Step: Isolate point where cache is exhausted and evaluate behavior when this occurs.
Dropped Flows Per Second	Number of flows which are dropped in a minute

Status

- Considered Ready For Implementation
- Adoption Call