draft-gashinsky-6man-v6nd-enhance-01

Joel Jaeggli
Kiran Chittimaneni
And others

History

- Genesis of this document was the original
  - Draft-gashinsky-v6nd-enhancements
- Document split into:
    - Draft-gashinsky-v6ops-v6nd-problems (now published as rfc 6583)
    - draft-gashinsky-6man-v6nd-enhancements
- Part of v6nd-enhancements was supplemental to draft-nordmark-6man-impatient-nud (should have completed a WGLC about now)

History

- Genesis of this document was the original
  - Draft-gashinsky-v6nd-enhancements
- Document split into:
    - Draft-gashinsky-v6ops-v6nd-problems (now published as rfc 6583)
    - draft-gashinsky-6man-v6nd-enhancements
- Part of v6nd-enhancements was supplemental to draft-nordmark-6man-impatient-nud (should have completed a WGLC about now)

History – continued

- What does that leave us with?
  - Gratuitous neighbor discovery is the remaining component of the original draft still floating around.
  - In initial discussion it proved to be the most controversial element of the draft.
  - That said, the problem it addresses isn't hypothetical.
  - loss of neighbor cache entries due to cpu exhaustion in the degenerate case and in fact can be induced in the course of normal operation of large l2 ethernets.

Why discuss this here and not just in 6man?

- It's pretty clear that changing NDP in this fashion is within scope for 6man and not v6ops.
- But the operational problem is here. (we're not looking for adoption here)
- If operators aren't demanding a solution to this problem, then the incentive implement isn't there.
- If there isn't consensus that this is useful work among those affected there's little point in advancing the work elsewhere.

Problem

In some network environments, legitimate Neighbor Discovery traffic from a large number of connected hosts could induce a DoS condition even without the use of scanning tools or deliberate attack. For e.g., consider a campus network with a pair of core routers that aggregate traffic from a few thousand wifi clients. In this scenario, high volume of regular ND traffic from clients on 1 or 2 large subnets worth of hosts can easily overwhelm the routers such that they are no longer able to process regular traffic anymore. Perversely subnets with large numbers of mostly inactive hosts may generate more NDP traffic than with hosts that are in regular communication with each other.

Solution(s)

- Some of the implementation suggestions in suggested in RFC6583 as well as draft-nordmark-6man-impatient-nud if approved would contribute at least partially towards alleviating this issue.
  - But if your workload is to high based on outstanding requests in your NDP queue from hosts on the subnet eventually hosts that you should be able to reach are going to fall out of the cache or not be learned when in fact they should be.

Solution - continued

- This draft:
  - proposes alterations that allow the update orinstallation of neighbor entries without the instigation of a full neighbor solicitation.

Proposal in detail

- RFC 4861, section 7.2.5 and 7.2.6 [RFC4861] requires that unsolicited neighbor advertisements result in the receiver setting it's neighbor cache entry to STALE, kicking off the resolution of the neighbor using neighbor solicitation.
- What we propose is that, If the link layer address in an unsolicited neighbor advertisement matches that of the existing ND cache entry, routers SHOULD retain the existing entry updating it's status with regards to LRU retention policy.

Proposal Continued.

- I envisioned this originally largely being confined to a datacenter, so the proposal had several deliberately imposed limitations.
- Hosts MAY be configured to send unsolicited Neighbor advertisement at a rate set at the discretion of the operators.
- The rate SHOULD be appropriate to the sizing of ND cache parameters and the host count on the subnet.
- An unsolicited NA rate parameter MUST NOT be enabledby default.

Proposal Continued

- The unsolicited rate interval as interpreted by hosts must jitter the value for the interval between transmissions.
- Hosts receiving a neighbor solicitation requests from a router following each of three subsequent gratuitous NA intervals MUST revert to RFC 4861 behavior.

Caveats

- This may impact one-way reachability detection
- A configuration method is not currently specified
  - DHCPv6 option seems likely.
- Without a fallback to normal (4861) behavior this results in more NDP rather than less.

Discussion

- Is this a problem we can solve?
- DHCP option, is it a useful inclusion.