Efficientnd Design Team Status Report

Presented by Erik Nordmark

Overview

- Design team was formed at London IETF-89
- Different discussions topics
  - Measurements of impact of ND traffic
  - Problems in different ND functionality
  - Operational techniques to reduce that impact
  - Any minor standards tweaks to reduce impact?
  - Any useful protocol options?

Measurement results

- draft-vyncke-6man-mcast-not-efficient
- draft-desmouceaux-ipv6-mcast-wifi-power-usage

Areas of potential concern

- Multicast RS and periodic RAs
- Multicast NS and NA for address resolution
- Duplicate Address Detection
- Look at WiFi and cellular/3GPP
- Requirement to not disrupt ND operation on links where it works fine today
- Consider hosts that:
  - sleep and wakeup based on packets
  - sleep based on a schedule
- Note that there is other traffic/applications which result in host-to-host multicast packets
  - mDNS has been observed to result in more multicast load than ND

RS/RA issues

- Multicast RS is not an issue
  - Can be filtered so that it is only sent towards the routers
- Multicast RA is unreliable on WiFi
  - And at least one every ~ 1800 seconds
- Periodic RA might be an issue for cellular
  - Each host needs to be sent an unsolicited RA each ~ 1800 seconds
  - If host is not active this results in paging to lots of base stations (1 million hosts; 1000 base stations => 1 billion paging messages for each RA)
  - See draft-garneij-6man-nd-m2m-issues

NS/NA address resolution

- Not an issue
  - Can configure RAs to have L=0 (prefix not on-link); only link-locals then subject to NS/NA
- In some cases sleep proxies are believed to cause unsolicited multicast NA
  - Don’t have much data on that

DAD issues

draft-yourtchenko-6man-dad-issues-00

1. Duplicate L2 address detection
2. Interaction with delay in forwarding on the link
3. **Behavior on links with unreliable multicast**
4. Interaction with looped interfaces [6man-enhanced-dad]
5. Delays before an address can be used [RFC4429]
6. Partition-join tolerance
7. Behavior on collision
8. **Energy efficiency**
9. **Wake-up and L2 events**
10. Usage of DAD to create state
11. Support of multi-link subnets [RFC6275 has a partial solution]
12. Anycast Addresses and Duplicate Address Detection
13. Implementations doing DAD once per IID
14. Backwards compatibility and presence of the DAD proxies [RFC6957]

Operational techniques

- Some already documented draft-yourtchenko-colitti-nd-reduce-multicast
  - Tuning timers etc
  - Having smaller L2 networks is an obvious one, but not always easy
- In a large conference or stadium on WiFi
  - Most users do not expect to do file sharing or other local protocols with others
  - Many such large networks filter multicast from host to host (Might break DAD)
  - Could also delegate a /64 to each host/subscriber
    - Makes it more clear that they do not share a

RS/RA potential improvements

- Use higher timers for periodic RA
  - Increase the max allowed from 9000 to 64k seconds
  - See draft-krishnan-6man-maxra-01
  - Seems to be already deployed
- Encourage router implementations to unicast solicited RAs
  - Configurable option so satellite links can multicast?
- Allow for hosts that optionally want to use RS to refresh the RA information
  - Configurable on router. Don’t change default behavior
  - See draft-nordmark-6man-rs-refresh-01

RS refresh draft

- RA includes a Refresh Time Option
  - If included, host will unicast RS before that time
  - Try ~ 3 RAs during the refresh time
  - Hosts randomize the timeout - avoid synchronization
- RS includes a flag (refresh capable)
  - Allows routers/admin to see if hosts are not capable
- Makes sense to implement together with resilient-RS
  - Resilient-RS handles the initial RS
- Updated RA information propagate slowly
  - A multicast RA would not reach hosts on sleep schedule

Address resolution - potential improvements?

- If NS/NA address resolution traffic is a problem, then L=0 will reduce it
  - Will still have NS/NA for link-local addresses
  - DAD will still be multicast
  - Already in RFC 4861
- Unsolicited NAs (sent to all-nodes) can happen with some proxies?
  - Guessing purpose is to update link-layer address
  - Need more data on this potential issue

DAD potential improvements?

- First we need to decide which of the 11+ DAD problems to address
  - Reliability of DAD? Efficiency and allowing sleep?
- Need WG input on draft-yourtchenko-6man-dad-issues-00
- FWIW DT has discussed
  - Implicit registration (some form of DAD proxy) to address reliability of DAD and flooding
  - Explicit registration (example is the ARO option in RFC6775/draft-chakrabarti-nordmark-6man-efficient-nd)

DAD choices

1. Deprecate DAD - it is expensive and duplicates are not common
2. Only send and receive DAD for manually configured addresses
3. Improve DAD
   a. Better at detecting duplicates (partition-join, etc)
   b. Less network and host impact (allow sleep schedule)
4. Do nothing
   a. aka go to the beach ;-)

Summary and Next Steps

- Please review the new
  - draft-yourtchenko-6man-dad-issues
  - draft-krishnan-6man-maxra-01
  - draft-nordmark-6man-rs-refresh-01
- We need WG feedback on what to approach in DAD issues
- We also need to look at 802 privacy - any implications on DAD?

Draft relevant to DT

draft-vyncke-6man-mcast-not-efficient

draft-desmouceaux-ipv6-mcast-wifi-power-usage

draft-yourtchenko-colitti-nd-reduce-multicast

draft-yourtchenko-6man-dad-issues

draft-krishnan-6man-maxra-01

draft-nordmark-6man-rs-refresh

draft-ietf-6man-resilient-rs

draft-chakrabarti-nordmark-6man-efficient-nd

draft-ietf-6man-enhanced-dad