ND on wireless links and/or with sleeping nodes

Problems

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

draft-yourtchenko-colitti-nd-reduce-multicast

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

Context

Original assumptions in RFC 1970 [1996]

- Based on shared medium (10BASE5, 10BASE-T with hubs)
  - Multicast as reliable as unicast
  - The network cost of multicast is the same as unicast (*total receiver cost is lower with filtering NIC*)
- Nodes are always on
  - The effort to power on a host was not optimized
  - Simple for multicast DAD probes to find duplicates

Sleeping nodes [RFC 6574 from IAB workshop]

- Small, low-cost, battery-powered nodes often "sleep" to extend battery life
  - Keep *just enough* of the system on to wake up on schedule
  - Radio **receiver** and transmitter are completely off
- Wake up periodically to perform functions
- Before transmitting to the network
  - Detect Network Attachment (DNA) by ucast NS to routers
  - Then DAD for LLA & global, then MLD reports for all solicited node mcast (hence at least 4 mcast packets if EUI-64 more if RFC 4941)

Wi-Fi Multicast Background

- Shared media: ucast, mcast and bcast are the same “thing” regarding radio spectrum use
- Additional 802.11 headers + management overhead
- Only ucast frames are acknowledged and retransmitted:
  - 10% packet loss appears to be common
  - Hence, 10% of multicast frames are lost
- Wi-Fi client mcast frames are first transmitted to the AP then retransmitted by the AP to all clients
- Depending on radio conditions, each Wi-Fi client has its own radio rate
  - => AP must transmit bcast/mcast frames at the lowest possible rate to ensure good reception
  - Makes bcast/mcast up to 10x more expensive than ucast

Radio Efficient Nodes
[IEEE 802.11 Low Power Wi-Fi clients]

- Even if main processor CPU is sleeping, try to keep radio going
- Radio is shut between AP beacons (100 msec)
- WiFi AP stores:
  - unicast frames destined to LP nodes
  - All multicast frames
- LP clients wake-up to listen to on AP beacons
  - Traffic Indicator Map (TIM) indicates whether to poll the AP to collect the frames
  - 1 bit states whether one or more multicast stored frames follow the beacon frame

Problems

How ND uses multicast

- DAD
  - 1 packet per IP address configured on the network
- RS
  - 1 packet per host that joins the network
- RA
  - Periodic: 1 packet every X seconds
  - Solicited: 1 packet for every host that joins the network
- NS
  - 1 packet for every new host/host pair

Problem areas

- Multicast, particularly on WiFi *(See subsequent slides)*
  - Wasting bandwidth
  - Waking up hosts unnecessarily
  - Looking at ND (RS, RA, DAD, address resolution)
- Duplicate address detection
  - Currently requires always-on to defend address
  - Requires waiting for 1 second for response (no “OK” response)
- Related DAD issues
  - Not robust against packet loss
  - Has loopback issues (see enhanced-dad draft)
  - Deployment issues for N:1 VLAN model in DSL (see dad-proxy draft)

Some data from IETF-hotel Wi-Fi

- Collected by a mostly silent node in promiscuous mode, 75% of IPv6 traffic was multicast

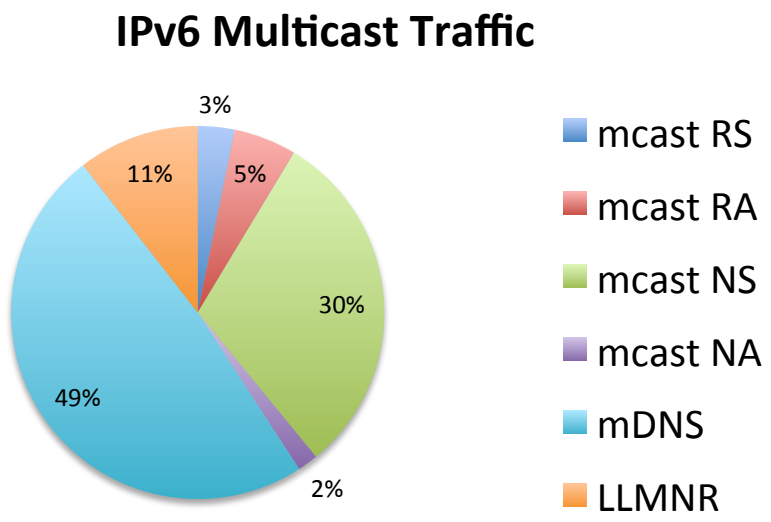

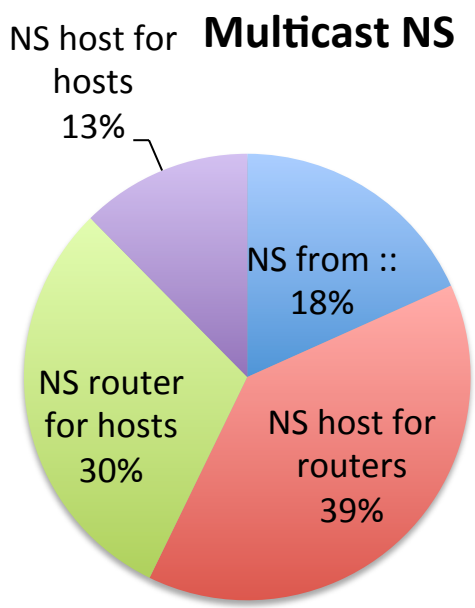

IPv6 is Chatty

Assumptions: RA interval 600 sec, 2 routers, NS lifetime 35 seconds, sleeping nodes wake-up cycle 1200 seconds, Wi-Fi max rate 54 Mbps, Wi-Fi lowest rate 1 Mbps

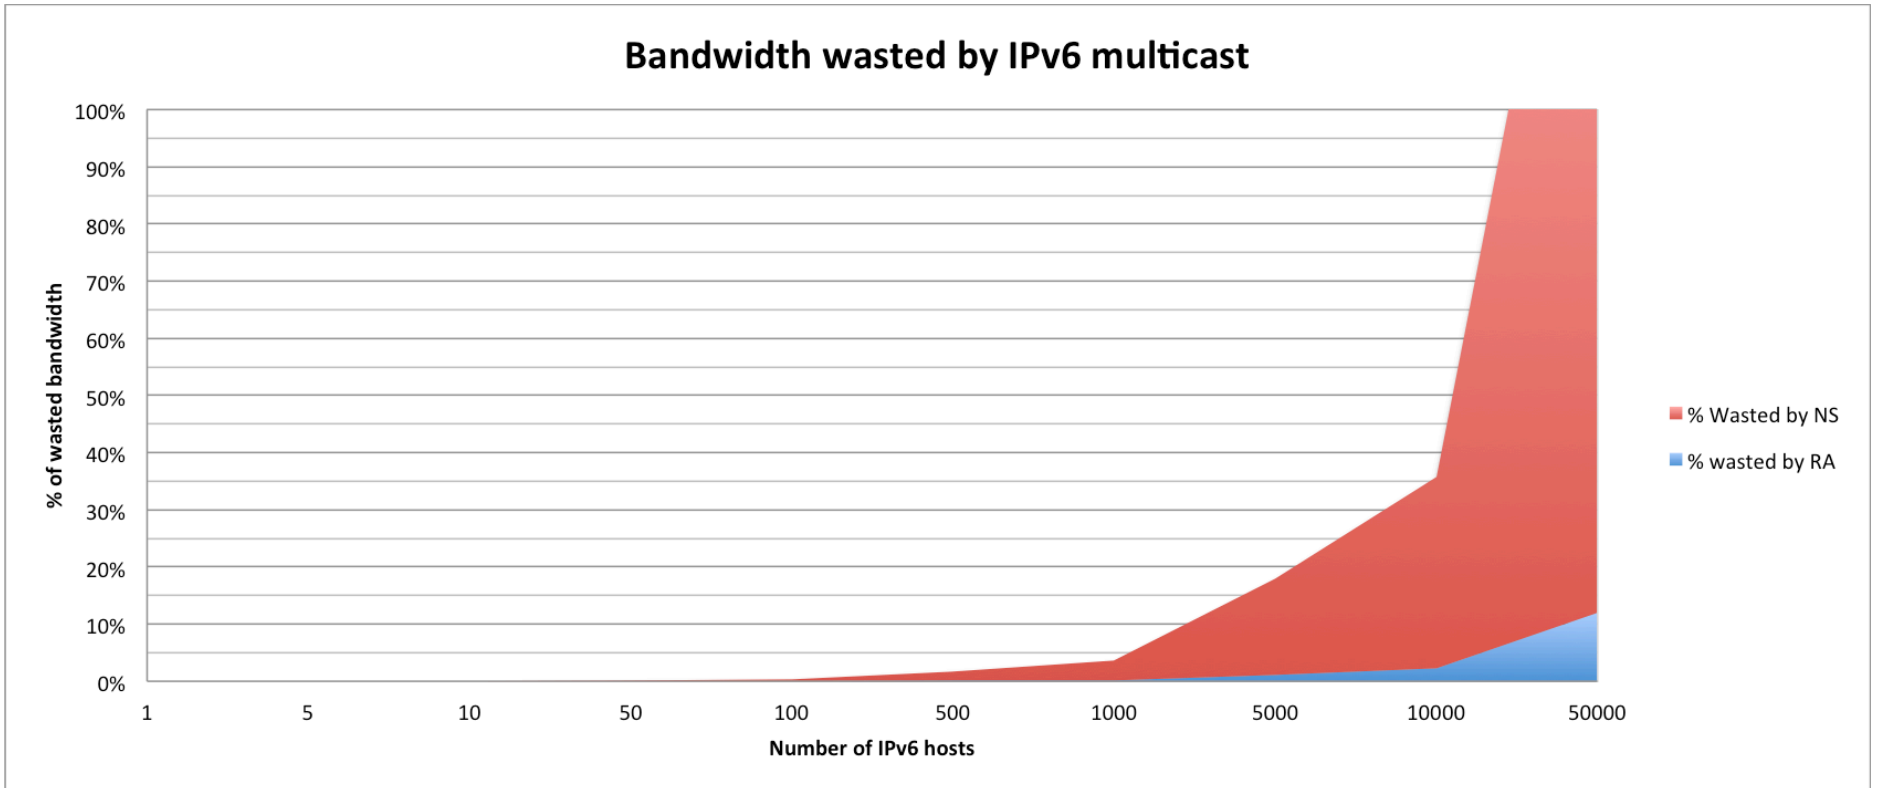

RFC 4541: MLD Snooping is not Enough

- Solicited-node multicast **was a good** idea:
  - MLD snooping could filter in switches
  - NICs could filter multicasts
  - Works best EUI-64-based link-local & global addresses (same group)
- RFC 4541 switches
  - Implement MLD snooping for global mcast
  - Cannot economically/physically implement RFC 4541 for solicited nodes mcast esp after RFC 4941 (privacy addresses)
  - Results in flooding the mcast NS
  - What if MLD report is lost?
- (v)NIC filters help with battery lifetime
  - Do not help with multicast bandwidth use
  - How many layer-2 mcast addresses can a (v)NIC support before interrupting the CPU?

Other problems in the neighborhood

- mDNS (Bonjour) is also quite chatty...

MITIGATION

Reducing Multicast in IPv6 Neighbor Discovery

draft-yourtchenko-colitti-nd-reduce-multicast

Andrew Yourtchenko

Lorenzo Colitti

On-device multicast filtering

- Have firmware drop packets that don't go to groups the host has joined
- Problems:
  - Still requires receiving the packet
  - The depth of sleep may vary
  - Airtime inefficiency remains
- Implementation effort: host-local

Unicast solicited RAs

- Respond to RS with unicast RA
  - Standard already allows this
- Substantial airtime and battery savings
- Should probably rate-limit and send multicast RA if too many RS/s

```
16:09:21.867637 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:23.599514 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:25.368509 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:27.039620 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:29.303073 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:30.899102 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:32.888595 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:35.186113 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:37.330651 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:39.781627 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
16:09:42.079485 IP6 fe80::20d:48ff:fe0a:69cb > ff02::1: ICMP6, router
```

Infrastructure multicast filtering

- Perform multicast snooping
  - e.g., like SAVI does - RFC6620
- Convert multicast → unicast
  - Pure 802.11
  - Unicast ethernet (RFC6085)

Proxy ND on the Access Points

- Selectively translate multicast NS to unicast
  - Based on snooping L3/L2 mappings
- AP has all information it needs to do this
  - In 802.11 infra mode, traffic always goes through AP

Maximize interval for periodic RAs

- Set all lifetimes to current maximum limits:
  - AdvDefaultLifetime: 9000 seconds
  - MaxRtrAdvInterval: 1800 seconds

- Ideas at the IETF to augment vastly those limits
  - Experiments show that host implementations are fine with this ;-)

Increase Reachable Interval

- Makes hosts more tolerant to short failures
  - Avoid multicast NS when NUD fails
- Hosts can't learn about unreachable routers
  - OK if there is only one router

Clear on-link bit in prefixes

- Reduces NS multicasts from hosts
  - Everything non link-local is off-link
- Caveats:
  - By itself doesn't fix router->host multicast NS
  - Need to disable redirects on first-hop router
  - All host-host traffic must go through router
    - If router is AP, or APs wired, not an issue

Client link shutdown when asleep

- Best savings (no power = no power drain!)

- Issues:
  - Additional delay on wake-up
  - No network-side wake-up
  - Breaks unicast, too

IETF works

- Several people working on this issue
  - Different layer-2 have different issues
  - 3GPP is different than WiFi