

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 31, 2014

Y. Desmouceaux
Cisco Systems
July 2014

**Power consumption due to IPv6 multicast on WiFi devices
draft-desmouceaux-ipv6-mcast-wifi-power-usage-00**

Abstract

IPv6 networks make a consequent use of multicast for several purposes, including mandatory functions such as Neighbor Discovery. Although this use of multicast does not create real difficulties on wired networks, it can become painful on wireless ones, notably in terms of power consumption. There might be little effect on home networks, however, such effects become more important on large-scale networks. This memo provides statistics about the multicast traffic rate in a large IPv6 wireless network and the induced device power consumption, in response to a call emitted at IETF 89.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 31, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights

and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	IPv6 typical multicast traffic	3
2.1.	Behavior when joining a network	3
2.2.	Behavior once connected	3
3.	Power consumption induced by multicast traffic	4
4.	Large-scale wireless networks	5
4.1.	Typical orders of magnitude	5
4.1.1.	Arrival rates	5
4.1.2.	Connection durations	6
4.2.	Influence of such networks over devices	7
4.3.	Roaming	7
5.	Possible solutions	8
5.1.	Optimizing Router Solicitations retransmissions	8
5.2.	Proxying multicast traffic	8
6.	Acknowledgements	9
7.	IANA Considerations	9
8.	Security Considerations	9
9.	References	9
9.1.	Normative References	9
9.2.	Informative References	10
	Author's Address	10

[1.](#) Introduction

Multicast is widely used in IPv6 for configuration purposes through the Neighbor Discovery protocol [[RFC4861](#)]. On IEEE 802.11 wireless links, this can have a negative impact on low-energy devices such as smartphones, as stated in [[I-D.vyncke-6man-mcast-not-efficient](#)]. The 802.11 protocol [[IEEE80211](#)] includes a power-save mode for such devices, allowing them to be in a sleep state in which they will be notified when packets addressed to them are pending. However, there is no known powerful equivalent of wired-links Multicast Listener Discovery (MLD) snooping [[RFC4541](#)], hence all devices will be woken up when multicast traffic is pending.

In this document, we provide data illustrating the issue at 3 levels:

- o typical multicast traffic emitted by a device when joining an IPv6 network, and 'background multicast noise' generated once connected;
- o power consumption statistics for wireless devices when confronted

to multicast traffic;

o orders of magnitude of connections frequencies and durations in

large-scale wireless networks.

Confronting these 3 levels makes it possible to model the power consumption overhead of multicast traffic on typical large wireless networks.

We finally discuss possible solutions at the IP and the 802.11 level.

2. IPv6 typical multicast traffic

2.1. Behavior when joining a network

When joining an IPv6 network, devices usually emit the following multicast traffic:

1. duplicate address detection for the link-local address;
2. router solicitation;
3. duplicate address detection for the SLAAC address;
4. duplicate address detection for the SLAAC privacy address.

In addition, MLDv2 [[RFC3810](#)] frames are emitted for registration to the solicited-node multicast addresses needed for Duplicate Address Detection. Their number is not deterministic since other multicast protocols may interfere, but is typically at least 4.

Depending on the configuration of the router, the Router Advertisement sent by the router in response to the node's Router Solicitation may also be multicast.

On Linux and Apple MacOS X clients that were tested, joining the network also induces Multicast Domain Name System (mDNS) [[RFC6762](#)] traffic. Once more, the number of packets emitted depends on the network environment, but is around 20. Likewise, Microsoft Windows clients generate Link-local Multicast Name Resolution (LLMNR) [[RFC4795](#)] multicast traffic when they connect to the network.

2.2. Behavior once connected

Once connected, some devices keep on sending IPv6 multicast frames. Protocols inducing such traffic include Neighbor Discovery, MLD, and local discovery or name-resolution protocols, such as mDNS, LLMNR, Simple Service Discovery Protocol (SSDP), Web Services Dynamic Discovery (WS-D).

On a typical middle-scale enterprise network, IPv6 multicast traffic induced by these protocols has a noticeable impact. Our measures on a 160-hosts network show that the average rate of multicast traffic

transiting through the link is 4.5 frames per second. The following table shows the repartition of the multicast traffic between these

protocols, as it was measured:

Protocol	ICMPv6	mDNS	LLMNR	WS-D
Ratio	0.53	0.33	0.12	0.03

Based on the previous data, we can deduce how much multicast traffic is generated by a "representative" device. Doing a simple cross-multiplication, we obtain a rate of 0.028 multicast packets/s for a single device, or 101 frames/hour. (This result is confirmed on a small isolated test network of two hosts.) The following table shows the indicative number of multicast packets emitted by a such a representative device on our test network in an hour:

Protocol	ICMPv6	mDNS	LLMNR	WS-D
Frames per hour	53	33	12	3

3. Power consumption induced by multicast traffic

While multicast traffic has no significant influence on actively connected devices, it might have a poor impact on the ones that are in power-save mode.

In power-save mode, the hardware of such a device needs to regularly wake up in order to check whether pending frames are buffered for it. To do so, it wakes up every DTIM_PERIOD beacons (DTIM_PERIOD usually being set to 1 or 2) and retrieves the beacon emitted by the Access Point (AP). If multicast traffic is pending, this will be indicated by the AP by setting a specific bit in the beacons's Time Information Management (TIM) information element to 1.

If a device sees that this bit is set, it will stay awake in order to retrieve the frames. The device CPU will then be awakened and frames will be transmitted by the wireless firmware to the IP stack. (Note that this might not happen if the firmware implements a multicast filter, but even in this case, current will be drawn to retrieve the frames.)

Power consumption measurements on smartphone devices confirm the negative impact of multicast traffic on sleeping devices. The measurement was performed on a Samsung i9195 by attaching an ammeter between the battery and the phone. When idle (screen off, GSM off, WiFi on, 802.11 power-save enabled by the device), current drawn is 10 mA in average. When a multicast frame is received and the CPU awakened to process it, the current draw goes to between 100 and 150 mA during a small peak of time.

Desmouceaux

Expires December 31, 2014

[Page 4]

Assuming that the duration of the current peak induced by processing a multicast frame is 0.1 s, a device would hence use K times more energy when it receives RATE multicast packets per second, than when it receives none, where $K = (10 \cdot (1 - \text{RATE} \cdot 0.1) + 150 \cdot \text{RATE} \cdot 0.1) / 10$.

The following table gives K as a function of RATE.

RATE (pkts/s)	0.01	0.1	0.5	1	2	4	>10
K	1.014	1.14	1.7	2.4	3.8	6.6	15

A possible workaround to this issue is to implement a multicast filter at the firmware level, which will only forward multicast packets to the IP stack if their destination address matches one of those registered by the device. Although this would have a positive impact on battery consumption, the following measurements show that this is not entirely satisfying.

Indeed, current drawn to retrieve pending multicast frames at L2 without waking up the main CPU is around 40 mA, which is still 4 times more than idle consumption. In order to simulate this, one can tweak a router so that it always advertises the presence of multicast frames by setting the TIM multicast bit to 1. The device's radio will then always try to retrieve frames following this beacon.

4. Large-scale wireless networks

Previous sections provide data about the battery usage induced by individual multicast frames, and the number of multicast frames generated by a single host when connecting and once connected.

In order to model the battery usage of real-life networks, the following section provides data about connection arrival rates and connection durations in usual large-scale wireless networks.

4.1. Typical orders of magnitude

The following results are based on anonymized data from a dualstack WiFi network in a conference setup. These data provide interesting statistics about user habits in a network characterized by mobility, where users move much from one room to another. Plus, it is a good example of a large wireless network, since the user count during working hours is between 600 and 700.

4.1.1. Arrival rates

Arrival rates in this test network follow a probability distribution which is very close to an exponential law (the error rate being only 6%). The observed parameter for this law is such that $1/\lambda = 6$

seconds.

Desmouceaux

Expires December 31, 2014

[Page 5]

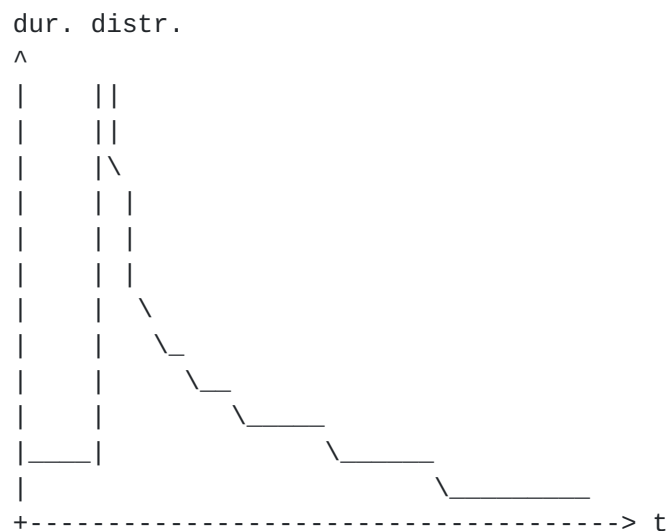
We are therefore in a scheme expected by the classical queuing theory, where arrivals are memoryless and have the Markov property of independence of the past. The smallness of the parameter is also an indicator of the great mobility of such networks: in average, a client joined the network every 6 seconds.

There is a small bias in previous measures due to people arriving en masse at the beginning of the day and going to lunch at noon. This bias does not affect the exponential nature of the law. However, it has a non-negligible effect over the parameter of the distribution. Computing the probability distribution over a shorter unbiased amount of time leads to a greater arrival rate: $1/\lambda$ becomes 4.5 seconds.

From these data it appears that arrival rates in wireless networks characterized by mobility are high. Since arrivals mean multicast traffic issued by the client's IPv6 stack, this already provides a rough intuition that large wireless networks feature high rates of multicast frames.

4.1.2. Connection durations

Once clients are connected, duration of their connections follow a less typical probability distribution, which form is represented below:



From these data it appears that there is a peak of connections which last 5 minutes, whereas almost no connection lasts less than this. A plausible explanation for this is that there are two kinds of behaviors: devices which automatically join the network without human intervention and disconnect after an idle timeout, and users who

consciously connect to the network. After this peak, the distribution is roughly an exponential law, which correspond to the latter case.

An average connection time of 55 minutes has been measured: at least, clients do not seem to be eager to leave the network too soon once they have joined it.

In our test network, arrivals follow an exponential law, and service times do not follow a close-form probability distribution. We can therefore model it as an M/G/oo queue. Once stabilized, the number of hosts in such a system follows a Poisson law of parameter S/T , where S is the expected service time and $T=1/\lambda$ the expected time between two arrivals.

[4.2.](#) Influence of such networks over devices

In this section, we model the influence of a large wireless network over a device, in terms of power consumption.

Let us assume that in average, N devices are present in the network, and that the arrival rate is λ . (See [Section 4.1](#) for typical numerical values). Then, [Section 2](#) implies that the rate of multicast frames in the network is $RATE = 0.025*N + 4*\lambda$. The following table gives RATE for a few values of N and λ :

N	5	10	50	100	500	500
1/lambda (s)	600	600	60	60	60	5
RATE (pkts/s)	0.13	0.26	1.32	2.57	12.6	13.3

Finally, using [Section 3](#), we can compute what is the energy overhead induced by the multicast traffic on a device. This overhead is $K = (10*(1 - (0.025*N + 4*\lambda)*10 + 150*(0.025*N + 4*\lambda)*10)/10$. The following table gives K as a function of N and λ :

N	5	10	50	100	500	500
1/lambda (s)	600	600	60	60	60	5
K	1.18	1.36	2.84	4.59	15	15

Thus, battery consumption induced by multicast traffic will be doubled in a network of 30 nodes where the arrival rate is 10 minutes.

[4.3.](#) Roaming

Depending on the configuration of the wireless network, roaming might have different consequences on multicast traffic emission.

Desmouceaux

Expires December 31, 2014

[Page 7]

When a host moves from one access point to another, roaming is supposed to be seamless. If a device detects a drop in signal quality, it will probe for a nearer access point with the same SSID. If it finds one, it will send an Authentication frame and a Reassociation Request. This will seem transparent to L3, except that some timers may time out, causing retransmissions.

However, if access points are not properly geographically distributed within the network, a device might lose connection to one before it can reconnect to another. In such a case, seamless roaming cannot happen and the device will have to go through the whole process of connection at the IPv6 level. This will generate multicast traffic as discussed in [Section 2.1](#).

5. Possible solutions

[I-D.yourtchenko-colitti-nd-reduce-multicast] already provides solutions at the IP layer. These include:

- o increasing intervals between Router Advertisements, and possibly remove the limit of 9000 s set by [\[RFC4861\]](#);
- o increasing the timer value for Neighbor Unreachability Detection;
- o unicasting Router Advertisements;
- o multicast filtering at the infrastructure level (MLD snooping).

At the L2 layer, it suggests using an efficient on-device multicast filter which would send frames to the IP layer only if their destination address is registered.

We explore other possible solutions in the next sections.

5.1. Optimizing Router Solicitations retransmissions

Some wireless systems retransmit all multicast packets received, so that hosts have a better chance to obtain them. This causes a problem when retransmitted packets are not desirable for hosts.

Without having to implementing a full multicast snooping mechanism, which can be costly in terms of resources and complexity for small systems like home boxes, a simple fix can be applied to APs in order to reduce the amount of multicast traffic retransmitted. APs should refrain from retransmitting packets destined to the all routers address (ff02::2), since routers should not be present on the wireless side. Essentially, Router Solicitations will be concerned.

5.2. Proxying multicast traffic

Desmouceaux

Expires December 31, 2014

[Page 8]

Some multicast traffic is only relevant to routers (for instance, MLD reports) or can be proxied by them (mDNS, ND). For instance, [\[RFC4389\]](#) specifies a means to proxy Neighbor Discovery traffic. Hosts are informed of the router proxying Neighbor Discovery traffic thanks to a bit in Router Advertisements.

On the same model, a more general option in Router Advertisements could indicate which multicast addresses a router is able to proxy. When a host receives such a Router Advertisement, it will record those addresses. It will then use the router L2 address instead of a 33:33:XX:XX:XX:XX multicast L2 address as destination for corresponding packets, thus reducing multicast traffic emission. The router will still be able to know that the final destination is multicast, since the destination IP address remains multicast (a behavior permitted by [\[RFC6085\]](#)). A disadvantage of this solution, though, is that it requires changes to the hosts.

6. Acknowledgements

The author would like to thank Andrew Yourtchenko, Eric Vyncke, Ole Troan, Brian Hart and Mark Townsley for their precious opinions on the subject.

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

None.

9. References

9.1. Normative References

- [IEEE80211] Institute of Electrical and Electronics Engineers, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Standard 802.11, 2012.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W. and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.
- [RFC6085] Gundavelli, S., Townsley, M., Troan, O. and W. Dec, "Address Mapping of IPv6 Multicast Packets on Ethernet",

[RFC 6085](#), January 2011.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", [RFC 6762](#),
February 2013.

Desmouceaux

Expires December 31, 2014

[Page 9]

9.2. Informative References

- [I-D.vyncke-6man-mcast-not-efficient]
Vyncke, E., Thubert, P., Levy-Abegnoli, E. and A.
Yourtchenko, "Why Network-Layer Multicast is Not Always
Efficient At Datalink Layer", Internet-Draft [draft-vyncke-6man-mcast-not-efficient-01](#), February 2014.
- [I-D.yourtchenko-colitti-nd-reduce-multicast]
Yourtchenko, A. and L. Colitti, "Reducing Multicast in
IPv6 Neighbor Discovery", Internet-Draft [draft-yourtchenko-colitti-nd-reduce-multicast-00](#), February 2014.
- [RFC4389] Thaler, D., Talwar, M. and C. Patel, "Neighbor Discovery
Proxies (ND Proxy)", [RFC 4389](#), April 2006.
- [RFC4541] Christensen, M., Kimball, K. and F. Solensky,
"Considerations for Internet Group Management Protocol
(IGMP) and Multicast Listener Discovery (MLD) Snooping
Switches", [RFC 4541](#), May 2006.
- [RFC4795] Aboba, B., Thaler, D. and L. Esibov, "Link-local Multicast
Name Resolution (LLMNR)", [RFC 4795](#), January 2007.

Author's Address

Yoann Desmouceaux
Cisco Systems
Issy Les Moulineaux, 92130
France

Email: ydesmouc@cisco.com

Desmouceaux

Expires December 31, 2014

[Page 10]