

batman-adv - Feature #206

Distributed IPv6-NDP cache to reduce overhead

03/12/2015 04:46 PM - Ruben Kelevra

Status:	New	Start date:	03/12/2015
Priority:	Normal	Due date:	
Assignee:	batman-adv developers	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
<p>Currently the Neighbor Discovery Protocol does takes much air-time and idle-bandwidth because of the broadcasts which are send thru the network.</p> <p>It would be nice if the queries could be stored on the nodes, distributed, to use some of ram of the nodes usefully and reduce network overhead.</p> <p>One possible solution would be:</p> <ul style="list-style-type: none">• If an IPv6 is queried by the local client, the node make three hashes and match them to the nearest mac-address of other nodes, and query them.• * If they all send NX do send the query as normal broadcast.• ** If the broadcast get an answer, send an update to the three nodes.• * If they does not return any answers for more than 20 seconds, do a normal broadcast. (redo queries for each Neighbor-Discovery-Query the node get)• If a node get no query for 2h, delete the entry.• If a node get more than \$StoreLimit entries, delete the oldest one.			

History

#1 - 03/12/2015 04:47 PM - Ruben Kelevra

Sorry, syntax was broken:

- If an IPv6 is queried by the local client, the node make three hashes and match them to the nearest mac-address of other nodes, and query them.
 - If they all send NX do send the query as normal broadcast.
 - If the broadcast get an answer, send an update to the three nodes.
 - If they does not return any answers for more than 20 seconds, do a normal broadcast. (redo queries for each Neighbor-Discovery-Query the node get)
- If a node get no query for 2h, delete the entry.
- If a node get more than \$StoreLimit entries, delete the oldest one.

#2 - 03/13/2015 04:15 PM - Antonio Quartulli

Ruben,
thanks for your interest about this problem. The mechanism you just described is almost the same we have implemented in DAT. Unfortunately DAT is currently working only for IPv4/ARP messages and needs to be extended in order to be able to work with IPv6/ND messages too. In the past we had somebody that started to work on this, but he had not enough time to continue working on this.

Maybe you can have a look and continue his work? :)

The last version of his code has been slightly re-arranged and pushed by me into the ordex/dat6 branch that you can find here:

<https://git.open-mesh.org/batman-adv.git/shortlog/refs/heads/ordex/dat6>

I hope this was helpful

#3 - 03/14/2015 12:42 PM - Ruben Kelevra

- File 00«.png added

Hey Antonio,

I'm sorry, but my C-skills are not good enough for a nice implementation.

I've added a graph which shows why we need a distributed cache for this info

#4 - 05/12/2015 03:12 PM - Linus Lüssing

Hi Ruben,

like the slides mention in the end too, the high IPv6 Neighbor Discovery overhead is supposed to get fixed soon with the multicast optimizations for bridges... please read and understand first :(. And as I said on the other ticket before, please use the mailinglist and IRC first.

<https://wiki.freifunk.net/images/e/e1/Batman-adv-scalability.pdf>

You can keep track of the current state of the batman-adv multicast optimizations for bridges on the mailinglist or here:

<https://git.open-mesh.org/batman-adv.git/shortlog/refs/heads/linus/multicast-bridge>

Hey, maybe you'd like to give me a hand with testing :)? There are never enough testers for new, upcoming features and will make features come sooner and safer. If everything works for you, you can write a few lines on what setup you tested to the according patchset on the mailinglist and could get a shiny "Tested-by: ..." reference in the Linux kernel hall of fame (git) :).

No C coding skills or even kernel coding knowledge required for that ;).

Edit: Oh, sorry, it actually does mention IPv6 DAT on these old slides of mine. Thought you had referenced the newer slides here: <https://www.metameute.de/~tux/Freifunk/batman-adv-assessment%2blookout-ffnordcon2014.pdf> (which was supposed to be hold at the FFNordCon but which I had to cancel). I can't read either, sorry again :D.

#5 - 03/10/2016 08:15 PM - Sven Eckelmann

The last version of the IPv6-DAT patchset (I know about) can be found at:

- <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/1373242365-763-1-git-send-email-mihail.costea2005@gmail.com/>
- <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/1373242365-763-2-git-send-email-mihail.costea2005@gmail.com/>
- <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/1373242365-763-3-git-send-email-mihail.costea2005@gmail.com/>
- <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/1373242365-763-4-git-send-email-mihail.costea2005@gmail.com/>
- <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/1373242365-763-5-git-send-email-mihail.costea2005@gmail.com/>
- <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/1373242365-763-6-git-send-email-mihail.costea2005@gmail.com/>

#6 - 02/11/2017 05:56 PM - Sven Eckelmann

- Assignee set to *batman-adv developers*

#7 - 04/24/2017 07:59 PM - Linus Lüssing

Just as an update: The batman-adv multicast optimization for bridges patches are part of batman-adv since v2016.3. Which should solve this ticket.

I'm expecting confirming results with the next, LEDE based Gluon release.

#8 - 04/25/2017 08:39 AM - Sven Eckelmann

@Linus We were starting to compile the multicast optimization as part of batman-adv in our gluon fork. (which uses batman-adv 2017.0). We were not actually enabling the multicast option at runtime but already noticed that the TT (full) sync stuff was going crazy and was not able to correctly sync anymore. The result were incorrect checksums and full TT syncs after each OGM. Please discuss this with Simon before starting a large scale test of the multicast optimization to find out whether the multicast optimization solves the IPv6-NDP overhead. He actually started to debug this problem - or actually the symptom of the high management overhead in our Freifunk mesh.

Files

00«.png	83.2 KB	03/14/2015	Ruben Kelevra
---------	---------	------------	---------------