April 2022

LAYER 2 COLLISION AVOIDANCE USING COMPOSED MAC ADDRESS

Thomas Vegas

Anirban Karmakar

Vincent Cuissard

Loris Gazzarrini

LAYER 2 COLLISION AVOIDANCE USING COMPOSED MAC ADDRESS

AUTHORS:
Thomas Vegas
Anirban Karmakar
Vincent Cuissard
Loris Gazzarrini

ABSTRACT

Techniques are presented herein that leverage the ability of wireless clients to employ a non-constant media access control (MAC) address (to, for example, avoid long-term tracking and identification) and which support a method for segmenting a MAC address field while making sure that wireless clients can still use a per-association non-constant or randomized MAC address (thus solving the tracking and radio collision issue). Aspects of the presented techniques ensure that collisions are prevented on the Distribution System Medium (DSM)/Distribution System Services (DSS) and ensure that a wireless infrastructure can still identify a wireless client across associations and during roaming. Use of such (composed) client MAC addresses in a Layer 2 (L2) infrastructure avoids full randomization and allows for decentralized client lookup and access point (AP) access identification.

DETAILED DESCRIPTION

Currently, wireless clients now have the ability to employ a non-constant MAC address to avoid long-term tracking and identification. For example, on a wide wireless network (e.g., in a metropolitan area), with clients using Randomized and Changing media access control (MAC) Addresses (RCM), collisions may occur on the radio medium but can also be seen on the Distribution System Medium (DSM)/Distribution System Services (DSS) which is even more problematic.

Presented herein are techniques that provide a method for segmenting a MAC address field while ensuring that wireless clients can still use a per-association non-constant or randomized MAC address (thus solving the tracking and radio collision issue), such that collisions can be prevented on the DSS/DSM a wireless infrastructure can still identify a wireless client across associations and during roaming.

Figure 1, below, illustrates the six byte format of a general MAC address.

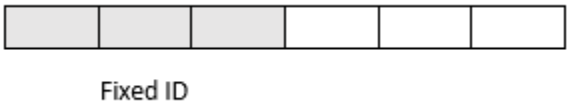

Figure 1: MAC Address Structure

Aspects of the techniques presented herein support an augmented MAC address field wherein the six bytes of a MAC address (as illustrated in Figure 1, above) are split into two parts. A first part of a MAC address, which may be referred to herein as MAC0, encompasses the first three bytes and includes a client or device identifier. Through such an organization a MAC0 address space may support more than 16 million devices. A second part of a MAC address, which may be referred to herein as MAC1, encompasses the last three bytes and includes either a random value or a client or device identifier.

Devices that understand an augmented MAC address (according to the techniques presented herein) may lookup devices and random parts as needed. Other non-supporting devices (such as, for example, older devices) may still employ the standard Layer 2 (L2) MAC address semantics.

Aspects of the techniques presented herein encompass a workflow that leverages the new MAC semantic that was described and illustrated above. For example, in connection with wireless client on-boarding, an access point (AP) may advertise, in its beacon, its support for an augmented MAC address. A wireless client may leverage such support during an association process as depicted in Figure 2, below.

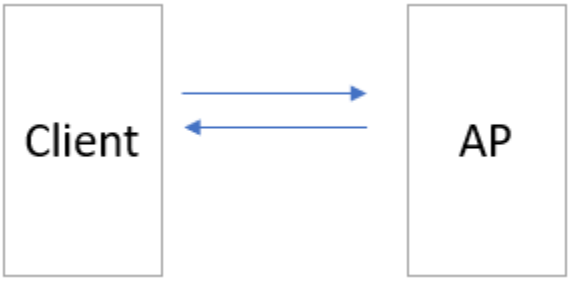

Figure 2: Exemplary Association Phase

The wireless client association as illustrated in Figure 2, above, encompasses two steps. First, the wireless client associates to the AP using a constant predefined MAC0 value. While MAC0 is a constant, MAC1 may be freely used by the client to avoid or resolve any detected radio collisions that could occur. In other words, an associating MAC address is equal to MAC0_assoc || MAC1_random where MAC0_assoc contains, for instance, a hexadecimal value such as `0xFF 0xFF 0xFF` and MAC1_random contains a, for example, random value. Figure 3, below, illustrate elements of the above-described step.

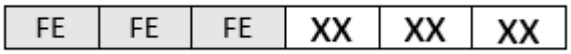

Figure 3: Exemplary MAC Address During Association Phase

The first three bytes of the exemplary MAC address that is depicted in Figure 3, above, contain a fixed value and a last three bytes containing a per-association random value.

Second, the AP, using the wireless infrastructure, allocates a three byte Client identifier (ID) that is to be used by the wireless client as its MAC0 field. The wireless client may use that field and then freely select a random value for the MAC1 field to resolve any radio collisions. In other words, under such an approach a MAC address is equal to MAC0_client || MAC1_randomX where MAC0_client contains the allocated Client ID and MAC1_randomX contains, for example, a random value.

Further aspects of the techniques presented herein encompass an AP MAC address translation process that supports different cases. For example, a first (preferred) case encompasses a hierarchical MAC address paradigm. Under this case, APs are themselves using the augmented MAC scheme where they are allocated, or negotiate with the wireless infrastructure, a three byte MAC0 identifier that is to be used when forwarding client traffic. That identifier may be referred to herein as either ap_MAC0 or apX_MAC0. Such a three byte MAC0 identifier is negotiated by the wireless infrastructure between one or many wireless and L2 infrastructures (using MDID, broadcast facilities, etc.). It is important to

note that many APs in a given wireless infrastructure can use and reuse the same ap_MAC0 value.

When an AP receives traffic from a wireless client, the AP discards the random value in MAC1 and forwards the packet using as a source MAC address (source_MAC) the value ap_MAC0 || client_MAC0. Such a source MAC address is collision free as the value of ap_MAC0 is hierarchically allocated and is used when the AP forwards client traffic. Importantly, there is also no translation table or any lookup to use or invalidate.

A wireless client can receive any packet containing its client_MAC0 in the destination MAC address field MAC0, where it may then discard the randomized MAC1 field.

Figure 4, below, illustrates elements of the process that was described above.

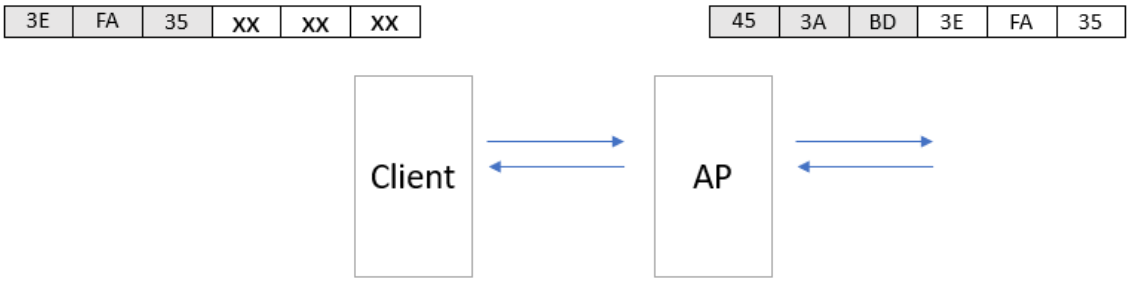

Figure 4: Exemplary MAC Addresses

The exemplary MAC address that is depicted on the left-hand side of Figure 4, above, represents an over-the-air (OTA) address where (according to aspects of the techniques presented herein and as described above) the first three bytes contain an allocated Client ID and the last three bytes contain a per-association random value (to, for example, prevent radio collisions).

The exemplary MAC address that is depicted on the right-hand side of the figure represents a MAC address that may be used to forward traffic on the infrastructure where (again, according to aspects of the techniques presented herein and as described above) the first three bytes contain an AP or AP pool identifier and the last three bytes contain an allocated Client ID.

It is important to note that the approach that was described above in the first three bytes of a MAC address for forwarding traffic contains an AP or AP pool identifier is provided for illustrative purposes only. These bytes do not need to be constant and do not necessarily need to identify an AP or an AP group. For example, there may be a set of values of arbitrary size that may be used by one or many APs, leading to pseudo-random values being observed and thus enhancing privacy. The only constraint is that the values that are used do not result in collisions.

A second case for AP MAC address translation encompasses a sub-MAC address randomization process. Under this case, a wireless client itself uses the MAC1 part of the source MAC address (as described above) to randomize and avoid radio collisions. To avoid infrastructure collisions, an AP may also forward a packet that is received from the wireless client by only updating or re-randomizing the MAC1 field of the source MAC address. In such a way, the MAC0 client_MAC0 value is preserved and it may be used for identity management.

Aspects of the techniques presented herein also support the identification of a client on a wireless infrastructure for various different cases. For example, under an inline case, a controller, or any supporting device in between, may extract the client ID from a MAC1 value when identifying the MAC0 part apX_MAC0. That apX_MAC0 value, which is currently described as a three byte field, may itself be split into further parts to, for example, facilitate DSM/DSS detection. The first nibble or byte of such an approach could, for instance, be constant and identify a wireless controller where the AP belongs.

Importantly, the roaming case is preserved in accordance with techniques of this proposal. A wireless client that is roaming may still use the same client_MAC0 field, and a freely randomized MAC1 field, with a new AP. The new AP may extract the MAC0 client_MAC0 value and provide it to a controller. The controller can then identify the client as a roaming client.

A mesh case features a hierarchy of APs where, for example, any AP can join another parent. In that case, the connectivity is preserved as a child AP can inform its neighbors about its wireless clients and the neighbors may extract the Client ID from the augmented MAC address using the first AP MAC address translation case (encompassing a hierarchical MAC address paradigm) procedure that was described above.

It is important to note that existing MAC rotation and communication techniques may still be employed and are independent from the techniques presented herein. Further, the client address space can always be reduced to the MAC0 field size according to the presented techniques.

In connection with client identification on L2 services, the Institute of Electrical and Electronics Engineers (IEEE) standard 802.1X (Dot1X) and the Dynamic Host Configuration Protocol (DHCP), for instance, often bind the client states to the actual MAC address. Supporting Dot1X and DHCP servers, which are a part of the wireless infrastructure, may also perform client MAC address extraction as described above.

Additionally, edge nodes may perform a translation in front of those services, using randomization for the MAC1 field and/or translation tables, as needed, to, when necessary, maintain a consistent view of a client MAC address. Such edge nodes may also present a constant MAC address view, translating a source MAC address comprising the concatenation of MAC0:apX_MAC0 and MAC1:client_MAC0 into the concatenation of MAC0:edge_MAC0 and MAC1:client_MAC0.

As described above, according to the techniques presented herein a constant MAC0 value is used during an association phase. It is important to note that that fixed value is employed just during the association phase. Consequently, to obtain a collision it would be necessary for some four thousand clients to be associating at the exact same moment in time. However, any collision that did arise would lead to an association failure and a retry, with another MAC1 randomized part. Thus, such a collision would not lead to ongoing traffic disruption.

Aspects of the techniques presented herein may be employed in locally owned L2 networks, which can grow to be quite large, without implying full network equipment vendor support for aspects of the techniques presented herein. Further, aspects of the presented techniques can help remove the centralization step of client lookup and identification. Upon observing a MAC address (as described and illustrated above) the other APs and L2 services (such as, for example, Dot1X, etc.) could identify the AP that is providing access and the client by its association.

In summary, techniques have been presented herein that leverage the ability of wireless clients to employ a non-constant MAC address (to, for example, avoid long-term

tracking and identification) and which support a method for segmenting a MAC address field while making sure that wireless clients can still use a per-association non-constant or randomized MAC address (thus solving the tracking and radio collision issue). Aspects of the presented techniques ensure that collisions are prevented on the DSM/DSS and ensure that a wireless infrastructure can still identify a wireless client across associations and during roaming. Use of such (composed) client MAC addresses in a L2 infrastructure avoids full randomization and allows for decentralized client lookup and AP access identification.