Privacy at the Link Layer

Piers O'Hanlon Joss Wright Ian Brown

Oxford Internet Institute
University of Oxford
forename.surname@oii.ox.ac.uk

Abstract

Many people now carry at least one device that routinely uses globally-unique link layer identifiers to locate and attach to local networks. Such link layer, or MAC, addresses are frequently emitted wirelessly in a variety of scenarios. This indiscriminate broadcast of a unique and effectively personal identifier allows for unregulated and broad-scale tracking of individuals via their personal devices, whether or not those devices have made use of a particular service or not. These addresses typically remain unchanged for the lifetime of a device, creating a persistent, lifelong tracking capability. These interface identifiers are now increasingly being monitored, collated, and analysed by a number of organisations, without meaningful regulation, for a variety of purposes without explicit permission from, or notification to, the individual.

At higher levels in the stack, on mobile devices, endpoint identifiers are generally ephemeral and change when attaching to different networks. Several efforts have been made to provide privacy or unlinkability at the IP level: private addressing, privacy extensions for IPv6 addresses, and Network Addresses Translators. Once a device seeks to connect to a network, however, there are a number of behaviours that can further weaken or compromise the privacy of devices and their owners, a notable example being the aggressive use of Detection of Network Attachment (DNA) services. Whilst, at a higher level, there is some level of user choice in access to communications services such as social media or cloud services, link layer identification is far less visible and accessible to normal users.

Given the ease with which link layer identifiers can be abused it would be advisable for for end systems to connect to networks without utilising an immutable, unique identifier where possible. We contend that, in many cases, the existence of unique *global* identifiers at the link layer is largely unjustified, and is increasingly a source of serious potential harm. End systems should therefore have the capability to employ an ephemeral address at the link layer to prevent long-term tracking and correlation of devices and individuals. We consider that the deployment of such systems should be both possible and practical within the constraints of many existing networks.

1 Introduction

Virtually all devices that connect to modern networks make use of globally unique link layer identifiers to locate and attach to local networks. Such link layer, or MAC, addresses are often emitted wirelessly providing the opportunity for unregulated wide scale track-

ing of specific devices whether or not they have connected a particular service or not. As many of these devices, particularly smartphones, are strongly bound to a single individual and accompany that individual at almost all times, these provide an effectively personal and unique identifier that is, in many cases, broadcast to anyone who cares to listen. In addition, as these addresses remain unchanged for the lifetime of the device they provide for a tracking capability for the entire active lifetime of the device.

The main wireless technologies in mobile devices, aside from the cellular radios, are WiFi (or 802.11), Bluetooth, and Near Field Communication (NFC), all of which typically operate in the industrial, scientific and medical (ISM) radio bands. WiFi utilises higher power levels, typically has a longer reach, and has been widely used as an approximate location service through global mapping of access points by corporations such as Google. Bluetooth provides for more persistent low level operation, and developments such as iBeacon, or Bluetooth Low Energy (BLE), promise to allow organisations to link device identifiers to their user's identity and location on a wide scale. NFC devices operate on very low power over short range, and have not seen the same scale of deployment in mobile devices, although they are widely and increasingly used for mobile payment and transport services.

At higher levels in the stack, on mobile devices, endpoint identifiers are generally ephemeral and usually change when attaching to different networks. Furthermore efforts have been made to provide for privacy at the IP level, such as the use of private addressing, privacy extensions for IPv6 addresses [7], and Network Addresses Translators. Once a device connects to a network there many opportunities to compromise the privacy of individuals further, for example in the aggressive use of proactive network attachment techniques such as Detection of Network Attachment (DNA) [1, 6] services. Clearly at the higher level communications can provide for further and more detailed tracking opportunities although there is a little more choice in the use of services such as social networks, web mail, and similar.

Thus far a number of companies, such as Google and Apple, and other organisations, such as wigle.net, have mapped a large proportion of the world's WiFi Access Point link layer addresses with which, amongst other things, they can provide approximate localisation services to mobile devices. A number of access point operators have also begun installing systems to monitor the link layer activity of any mobile devices, whether associated or not, that they observe in their vicinity. These sources of data, taken across a service provider, could enable detailed location and interaction traces of individuals. Furthermore, the major mobile handset OS providers maintain profiles pertaining to the owners of each handset, including its link layer addresses. Any of these information sources can potentially lead to significant losses in privacy.

This paper is structured as follows; in Section 2 we cover related work in this field. In Section 3 we detail the constraints of link layer addressing schemes. Section 4 outlines potential approaches to mitigating privacy problems at the link layer. Finally, Sections 5 and 6 propose future work and conclusions.

2 Related work

There have been various studies in the area of privacy and link layer addresses, with a number of potential solutions proposed and analysed.

Grutesser and Grunwald investigated the use of disposable interface identifiers for protection of location privacy [4], and found it yields significant privacy improvements. Re-

lated approaches have been taken by Jiang et al [5].

Pang et al. showed that whilst the use of disposable identifiers provided some additional privacy for a mobile device, the majority of users may be identified through the use 802.11 fingerprinting techniques [8].

In later work Greenstein et al. devised and implemented a new privacy enhanced link layer protocol, known as SlyFi [3], which included a number of security features including obfuscation of all transmitted bits including link identifiers. The protocol provided for improved performance over WPA-PSK for discovery and association.

Other work has shown that the use of aggressive proactive network attachment techniques in mobile OSes can already be exploited to uncover social relationships [2].

There have been various other examples of abuse of link layer identifiers in the press. Some notable examples include a UK based waste management company that was found to have deployed WiFi MAC address collection in recycling bins [1], whilst some commercial companies have already been brandishing potential solutions to such tracking by switching off WiFi when not in 'trusted' locations such as home or office [2].

3 Link Layer addressing

The link layer addressing used on Ethernet, WiFi and Bluetooth is specified by the IEEE. The address space is administered by the IEEE Registration Authority which handles the assignment of the Organisationally Unique Identifiers (OUI). The OUI is contained within the first 3 octets of an address, though actually utilises 22 bits. There are a number of address lengths available but the most often used is the OUI-48 address which is detailed in Figure 1.[3] There are two flag bits in the first octet which designate firstly whether the destination is a group (e.g. Broadcast or multicast) or individual address. The second bit indicates whether the address is universally administered by the IEEE RA or locally administered. This second flag bit of particular interest as it means that addresses with this bit set may be chosen without reference to the IEEE, though one should respect certain guidelines laid out by the IEEE in doing so.

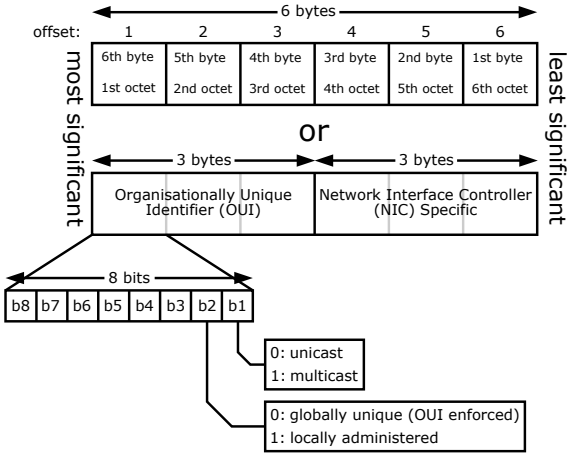

Figure 1: EUI-48 Addressing model

4 Ephemeral Addressing

Whilst it is possible to change the link layer addresses on most devices, in practice this rarely occurs. We performed some investigations on mobile OSes and noted that whilst older versions of Apple's iOS (iOS4-6) allowed for changing of an iPhone's MAC address (when rooted), later versions (iOS7) no longer appear to allow it.

[1] http://www.bbc.co.uk/news/technology-23665490

[2] http://www.economicvoice.com/avg-delivers-shopping-privacy-with-new-smartphone-wi-fi-do-not-track-feature/

[3] From the Wikipedia page on MAC Addresses

In certain systems, notably virtual machines, MAC addresses are routinely automatically generated and used on the LAN.

There are number of issues with the current use of MAC addresses: their global uniqueness and length is largely unnecessary in most LAN deployments where often only a few hundred or a few thousand devices exist on a single LAN segment. Although there are exceptions to this, such as large data centre deployments, the 48 bit address space is still unnecessarily large in the general case. Given the increasing deployment of virtual machines where MAC addresses are managed it seems that there are few compelling reasons for MAC addresses of other devices to be similarly managed.

In unmanaged environments devices could choose a random MAC address and optionally perform duplicate address detection [4] in a similar manner to the Address Resolution Protocol (ARP). In addition to providing a simple means to improve the privacy of individuals, this approach could also potentially free up the link layer address space for partitioning and other uses.

5 Conclusions

We have outlined the case for further investigations into privacy enhanced link layer interactions. We argue that whilst the link layer initially appears to be a concern only for interactions on a local network segment, in fact the high volume of deployed mobile devices today, the strong association between a user and their devices, and the current approaches to network access make link layer identifiers a serious candidate for mass tracking and surveillance of individuals. We therefore suggest that the link layer is a crucial, and under-studied, avenue to consider when attempting to address pervasive monitoring.

6 Future Work

We plan to investigate a range of approaches that can be used to minimise the leakage of identifying and long-term linkable information from mobile devices that could be used for uncontrolled monitoring.

7 Acknowledgments

We would like acknowledge funding from the UK Engineering and Physical Sciences Research Council for the Being There project, grant EP/L00416X/1.

References

[1] B. Aboba, J. Carlson, and S. Cheshire. Detecting Network Attachment in IPv4 (DNAv4). RFC 4436 (Proposed Standard), Mar. 2006.

[2] M. V. Barbera, A. Epasto, A. Mei, V. C. Perta, and J. Stefa. Signals from the crowd: Uncovering social relationships through smartphone probes. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 265–276, New York, NY, USA, 2013. ACM.

[3] B. Greenstein, D. McCoy, J. Pang, T. Kohno, S. Seshan, and D. Wetherall. Improving wireless privacy with an identifier-free link layer protocol. In *Proceedings of the 6th International Conference on Mobile Systems, Applications, and Services*, MobiSys '08, pages 40–53, New York, NY, USA, 2008. ACM.

[4] M. Gruteser and D. Grunwald. Enhancing location privacy in wireless lan through disposable interface identifiers: A quantitative analysis. In *Proceedings of the 1st ACM*

International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots, WMASH '03, pages 46–55, New York, NY, USA, 2003. ACM.

[5] T. Jiang, H. J. Wang, and Y.-C. Hu. Preserving location privacy in wireless lans. In *Proceedings of the 5th International Conference on Mobile Systems, Applications and Services*, MobiSys '07, pages 246–257, New York, NY, USA, 2007. ACM.

[6] S. Krishnan and G. Daley. Simple Procedures for Detecting Network Attachment in IPv6. RFC 6059 (Proposed Standard), Nov. 2010.

[7] T. Narten, R. Draves, and S. Krishnan. Privacy Extensions for Stateless Address Autoconfiguration in IPv6. RFC 4941 (Draft Standard), Sept. 2007.

[8] J. Pang, B. Greenstein, R. Gummadi, S. Seshan, and D. Wetherall. 802.11 user fingerprinting. In *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking*, MobiCom '07, pages 99–110, New York, NY, USA, 2007. ACM.