WITECH: Jurnal Teknologi Rekayasa Komputer dan Jaringan
Vol. 2 (No. 1), 2026: 57-71
https://journal.uwais.ac.id/index.php/witech/article/view/80

MAC Address Randomisation and Its Implications for Service Management in Multi-Access Point Community Networks

Evolusi Sistem Terdistribusi: Perkembangan Arsitektur, Teknologi, dan Tantangan Komputasi Modern

Ayu Angreini[1], Syahril Syam[2], Irhamullah[3], Iskandar Kato[4], Jumriati[5], Andi Yusril Cibu[6]
[1]Computer and Network Engineering Technology Study Program, Politeknik Wahdah Islamiyah Makassar, Makassar, Indonesia; ayuangreini28@gmail.com
[2]Multimedia Engineering Technology Study Program, Politeknik Negeri Ujung Pandang, Makassar, Indonesia; syahrilsyam@poliupg.ac.id
[3]Computer and Network Engineering Technology Study Program, Politeknik Wahdah Islamiyah Makassar, Makassar, Indonesia; irhamullah01@gmail.com
[4]Comparative Islamic Schools of Thought Study Program, Institut Agama Islam STIBA Makassar, Makassar, Indonesia; iskandarkato@stiba.ac.id
[5]Computer Science Study Program, Universitas Teknologi Akba Makassar, Makassar, Indonesia; perdana.jum@gmail.com
[6]Computer and Network Engineering Technology Study Program, Politeknik Wahdah Islamiyah Makassar, Makassar, Indonesia; andiyusrilc@gmail.com

Article History
Received: 2026-04-04
Revised: 2026-04-15
Accepted: 2026-04-20
Published: 2026-04-30

Keywords:
MAC address randomisation, community network, service identity, hotspot authentication, network service management

Corresponding author:
ayuangreini28@gmail.com

Paper type:
Research paper

Program Studi Teknologi Rekayasa Komputer dan Jaringan, Politeknik Wahdah Islamiyah Makassar, Indonesia

Abstract

Community networks built on centralised hotspot control commonly rely on the MAC address as a device marker to sustain authentication continuity. That assumption of address stability is now challenged by MAC address randomisation, a privacy feature enabled by default in mobile operating systems. This study measures the level of MAC address randomisation observed in authenticated sessions, analyses its persistence characteristics, and evaluates the implications for the choice of identification basis in service management architecture. A descriptive quantitative design was applied to ex post facto operational data from a small paid community network serving 31 service identities through approximately 30 access points. A total of 46 service activation records covering 26 July to 22 August 2026 were extracted from a configuration export of the controlling device. Randomisation status was classified using the locally administered bit test. Results show that 80.4 per cent of MAC addresses were randomised, with a 95 per cent confidence interval of 69.0 to 91.9 per cent, and the proportion remained stable across observation periods. Five addresses recurred, four of them randomised yet unchanged for three to eleven days, indicating that randomisation persists per network. This characteristic causes device-based continuity to function within the same network name but fail across different network names, whereas speed class and validity policies remain consistent because they are bound to the service identity. These findings provide empirical support for placing service identity at the centre of network service management.

Abstrak

Jaringan komunitas berbasis hotspot terpusat umumnya mengandalkan alamat MAC sebagai penanda perangkat untuk menjaga kesinambungan autentikasi. Asumsi kestabilan alamat tersebut kini tertantang oleh penerapan pengacakan alamat MAC sebagai fitur privasi bawaan pada sistem operasi perangkat bergerak. Penelitian ini bertujuan mengukur tingkat pengacakan alamat MAC pada sesi terautentikasi, menganalisis karakteristik persistensinya, serta mengevaluasi implikasinya terhadap pilihan dasar pengenalan dalam arsitektur pengelolaan layanan. Penelitian menggunakan desain deskriptif-kuantitatif dengan data operasional purna-fakta dari sebuah jaringan komunitas berbayar berskala kecil yang melayani 31 identitas layanan melalui sekitar 30 titik akses. Data berupa 46 catatan aktivasi layanan periode 26 Juli hingga 22 Agustus 2026 diekstraksi dari berkas ekspor konfigurasi perangkat pengendali. Klasifikasi status pengacakan dilakukan melalui uji bit alamat yang dikelola secara lokal. Hasil menunjukkan 80,4% alamat MAC bersifat teracak dengan interval kepercayaan 95% pada rentang 69,0% hingga 91,9%, dan proporsi tersebut stabil antarperiode pengamatan. Lima alamat teridentifikasi muncul berulang, empat di antaranya teracak namun tetap sama selama 3 hingga 11 hari, menunjukkan pengacakan bersifat persisten per jaringan. Karakteristik ini menyebabkan mekanisme kesinambungan berbasis perangkat berfungsi pada nama jaringan yang sama tetapi gagal pada nama jaringan berbeda. Sebaliknya, kebijakan kelas kecepatan dan masa aktif tetap konsisten karena melekat pada identitas layanan.

Cite this article:
Angreini, A., Irhamullah, Kato, I., Jumriati, & Cibu, A.Y. (2026). MAC Address Randomisation and Its Implications for Service Management in Multi-Access Point Community Networks. *WITECH: Jurnal Teknologi Rekayasa Komputer dan Jaringan*, 2(1), 57-71. https://journal.uwais.ac.id/index.php/witech/article/view/80.

1. Introduction

Internet access has become a basic necessity underpinning education, employment, and public services. In areas not adequately reached by telecommunications infrastructure, community networks have emerged as a form of access provision that enables shared use of infrastructure within a locality (Braem et al., 2013). This provision model has developed across many countries with varying characteristics, ranging from volunteer-based wireless networks to small paid hotspot ecosystems operated by individuals (Astillo et al., 2026; Maccari & Lo Cigno, 2015).

Small-scale community networks typically consist of several access points connected to a single centralised controlling device. That controlling device performs authentication, policy enforcement, and session management, while the access points serve merely as connecting media. Such an architecture places all operational decisions at a single point of control, allowing management to be carried out without repeated configuration on every access device (Strassner, 2003).

Within this architecture, the MAC address plays an important role. Beyond serving as a link-layer identifier, the MAC address is commonly used as a device marker to sustain authentication continuity. Device-marker mechanisms allow previously recognised users to gain access without re-entering service credentials, thereby reducing friction in use. Such mechanisms rest on one fundamental assumption, namely that the MAC address is a stable identifier bound to the physical device.

That assumption no longer holds in full. To limit the privacy problems arising from the linkage between a device, its traffic, its location, and its user, device manufacturers and operating system developers have adopted MAC address randomisation (Henry & Lee, 2025). Randomised addresses are distinguished from manufacturer-assigned addresses by a marker bit in the most significant octet (Zuniga et al., 2025). When randomisation occurs, some state stored within the network may break, thereby affecting connectivity and user experience. Most importantly for the present study, commonly deployed randomisation schemes allow a device to present a different address on a different network name (Henry & Lee, 2025).

Prior research has mapped the adoption of MAC address randomisation in considerable detail. Martin et al. (2017) conducted the first large-scale study of randomisation in the field, covering techniques by operating system, manufacturer, and device model. Fenske et al. (2021) revisited the problem through a cross-sectional study of 160 phone models to determine the conditions under which randomisation occurs and the extent to which it closes tracking gaps. Both studies, like most work of this kind, measured at the level of probe request frames, that is, before a device associates with a network.

With respect to its impact on network services, the problem has been documented normatively. Standards documents note that many networks and applications depend on the MAC address to recognise devices, including authentication portals, so that frequent address changes may disrupt device recognition and service delivery. Another line of work proposes a persistent identity layer above the MAC address, distributed through a centralised authentication protocol, although that approach acknowledges that open hotspot scenarios without authentication offer no reliable correlation path (Seralathan, 2026).

Nevertheless, a gap remains at the level of empirical measurement. Available measurements have almost all been taken at the pre-association level and in public, retail,

or enterprise Wi-Fi environments. Measurement of MAC address randomisation in already-authenticated sessions within small paid community networks is rarely reported, as is analysis of its implications for the choice of service management architecture in such environments. Recent work on paid hotspot ecosystems in rural areas has indeed been carried out, but it focuses on security aspects and vulnerability to attack rather than on device identity or service management architecture (Astillo et al., 2026).

On this basis, the present study formulates three questions. First, how high is the level of MAC address randomisation in authenticated sessions within a voucher-based community network. Second, what are the persistence characteristics of randomised MAC addresses and their impact on the reliability of device-based continuity mechanisms. Third, how is the consistency of service policy maintained when device identity is unstable.

In line with these questions, this study aims to measure the level and distribution of MAC address randomisation in authenticated sessions, to analyse the persistence characteristics of randomised MAC addresses together with their implications for device-based continuity mechanisms, and to evaluate the resilience of service-identity-based management against instability of device identity.

This study is expected to offer two benefits. Academically, it supplies measured evidence on the adoption of MAC address randomisation from a post-authentication perspective in a community network context that has been under-represented in the measurement literature. Practically, its findings provide a basis on which community network operators may decide the identification basis used in their service management architecture.

2. Literature Review

2.1 Community Networks and Their Service Management

A community network is a form of network access provision that enables shared use of infrastructure within a locality, with management typically carried out by the community or by individuals within that locality (Braem et al., 2013). Its operational characteristics differ from those of conventional household networks because it involves several access points serving many users within a single service environment (Avonts et al., 2013; Maccari & Lo Cigno, 2015).

In their development, some community networks operate under a small paid model. Astillo et al. (2026) documented pay-per-use hotspot ecosystems widespread in rural areas of the Philippines as well as comparable initiatives in India, showing that this provision model has grown as an affordable solution for areas with limited telecommunications services. Their study focuses on the security properties of these ecosystems, which had previously received little research attention.

2.2 AAA and Policy-Based Service Management

Service management in networks serving many users generally rests on the authentication, authorisation, and accounting framework. The RADIUS protocol provides a centralised authentication mechanism so that identity verification can be performed consistently without depending on the configuration of individual access points (Rigney et al., 2000), accompanied by recording of service usage (Rigney, 2000) and the capability to modify authorisation during an active session (Chiba et al., 2008).

Building on that framework, policy-based management approaches place a set of rules as the basis for operational decision-making, so that the enforcement of access rights

and the allocation of resources can be determined centrally (Strassner, 2003). The separation of responsibilities among functions in such an architecture applies a common principle of software architecture design (Bass et al., 2012).

2.3 MAC Address Randomisation

The MAC address is a 48-bit identifier traditionally assigned by the manufacturer and globally unique. Locally administered addresses are distinguished from universally administered addresses by the second least significant bit of the most significant octet (Zuniga et al., 2025). This distinction forms the basis for classifying randomisation status in the present study.

Martin et al. (2017) conducted the first large-scale study of MAC address randomisation in the field, detailing techniques by operating system, manufacturer, and device model. That study also found that some devices erroneously transmitted their genuine addresses under conditions in which a randomised address should have been used, so that randomisation could be defeated on most Android devices at the time. Fenske et al. (2021) revisited the problem three years later through a cross-sectional study of 160 phone models, to determine when and under what conditions a device randomises its address and the extent to which previously identified weaknesses had been addressed.

Both studies measured at the level of management frames prior to association. This approach is consistent with their shared objective, namely to assess the effectiveness of randomisation as a privacy protection against passive tracking.

2.4 Impact on Network Services

The impact of MAC address randomisation on network operation has been documented normatively. Henry and Lee (2025) note that when randomisation occurs, some state stored within the network may break, thereby affecting connectivity and user experience, while devices may continue to use other stable identifiers. That document also explains that commonly deployed randomisation schemes allow a device to present a different address on a different network name.

In response to this problem, opt-in identification mechanisms have been developed at the link layer, allowing devices and networks to recognise one another without relying on a fixed address mapping. At the authentication level, Seralathan (2026) proposes a persistent device identity framework distributed through a centralised authentication protocol. That framework acknowledges that in open hotspot scenarios without authentication, the MAC address is the only available identifier and therefore no reliable correlation path exists.

2.5 Position of the Present Study

The existing literature is well established in three respects: the mechanism of MAC address randomisation and how to classify it, the level of its adoption on mobile devices, and its impact on network services that depend on device identifiers. This study claims to have discovered none of these three.

The contribution of this study lies in three points of difference. First, in terms of the measurement point, this study measures at the level of authenticated service activation events rather than at the level of probe request frames as in earlier measurements. Second, in terms of environment, the study was conducted in a small paid community network that is under-represented in the measurement literature, whereas

recent work in a comparable environment focuses on security aspects (Astillo et al., 2026). Third, in terms of perspective, the study connects measurement findings to the choice of identification basis in service management architecture rather than to the effectiveness of privacy protection.

Accordingly, this study complements the findings of Martin et al. (2017) and Fenske et al. (2021) with measurement from a different observation unit and context, and supplies field evidence for a problem that has been formulated normatively by Henry and Lee (2025).

3. Research Method

3.1 Research Design

This study employs a descriptive quantitative design with a case study approach in a single implementation environment. The data analysed are ex post facto operational data, that is, records produced by the system during normal operation without researcher intervention. The researchers performed no manipulation of variables and no control of conditions, so the study is not experimental in nature.

This approach was chosen because the aim of the study is to measure and explain a phenomenon occurring under real operational conditions rather than to test the effect of a treatment. Consequently, the findings are descriptive and contextual, and are not intended for generalisation to the whole population of community networks. The research flow is shown in Figure 1.

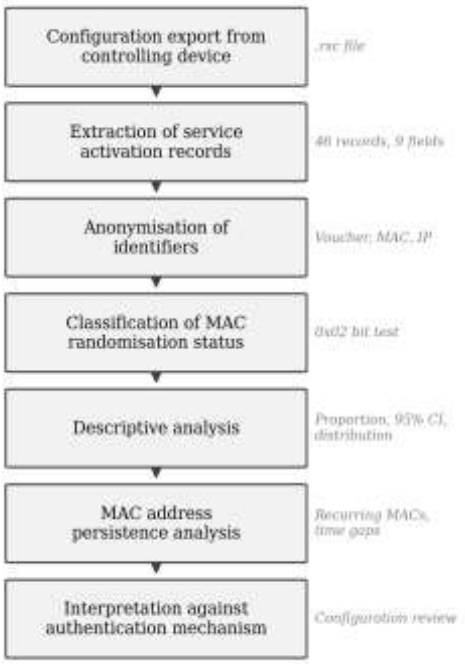

Figure 1. Research Flow
Source: Authors' design, 2026.

3.2 Research Object and Environment

The research object is a small paid community network operating in a residential area in Indonesia. The network serves 31 active service identities through approximately 30 wireless access points distributed across subscriber homes, with fibre optic links to a

single centralised controlling device. The architectural characteristics of the research environment are summarised in Table 1.

Table 1. Characteristics of the Research Environment

Parameter	Configuration
Controlling device	MikroTik RB1100AHx4, RouterOS 7.18.2
Service mechanism	Hotspot with voucher-based authentication
Login method	cookie, http-chap, trial, mac-cookie
Number of hotspot servers	1 (centralised)
Network segmentation	Single bridge, no VLAN per access point
Address pool	A single /23 block
Access points	Approximately 30 modems operating as bridges
SSID naming	Distinct per access point and per frequency band
Active service identities	31
Simultaneous device limit	1 (lowest class); 32 (other classes)
Aggregate hotspot capacity	200 Mbps

Source: Review of the controlling device configuration, 2026.

This configuration places all authentication, policy evaluation, and session management functions on the centralised controlling device, while the access points serve only as connecting media without any decision-making function.

3.3 Data Sources and Collection Technique

Data were obtained from a configuration export of the controlling device produced through a built-in command of the device operating system. The export contains the system configuration together with service activation transaction records stored by the hotspot management application as system script entries.

Each activation record contains nine fields, namely date, time, voucher identifier, transaction value, allocated IP address, device MAC address, service duration, service class, and voucher batch marker.

Data collection was performed once, on 23 August 2026. The researchers installed no additional monitoring devices, performed no packet capture, and did not access the content of user traffic.

3.4 Population and Sample

The population of the study comprises all service activation events occurring on the network throughout its period of operation. The sample comprises the 46 activation records still stored on the controlling device at the time of data collection, covering the period from 26 July to 22 August 2026.

The sampling technique is total sampling of available data rather than random selection. Activation records from earlier periods were no longer stored at the time the export was taken, so the sample does not represent the entire operational period of the network. This limitation is stated explicitly and taken into account in drawing conclusions.

3.5 Variables and Operational Definitions

The research variables and their operational definitions are presented in Table 2.

Table 2. Variables and Operational Definitions

Variable	Operational definition	Scale
MAC randomisation status	Status of a MAC address as locally administered based on the marker bit in the first octet	Nominal
MAC persistence	Occurrence of the same MAC address in more than one activation event	Nominal
Persistence gap	Difference in days between the first and last occurrence of the same MAC address	Ratio
Service class	Speed and validity profile bound to the service identity	Nominal
Activation time	Date and time at which the first service session was established	Interval

Source: Authors' design, 2026.

3.6 Research Instruments

The data recording instrument comprises the network controlling device together with the hotspot management application. Both components record activation events automatically as part of normal system operation, so they introduce no observer effect on user behaviour.

The classification instrument determines MAC address randomisation status on the basis of the locally administered bit, that is, the second least significant bit of the most significant octet of the MAC address. Following the provision referenced in Zuniga et al. (2025), an address whose bit is set to one is locally administered, whereas an address whose bit is zero is universally administered by the manufacturer. Operationally, a MAC address is classified as randomised when the bitwise conjunction of its first octet with the hexadecimal value 0x02 yields a non-zero result.

The processing instrument comprises a script written by the researchers, covering parsing of the export file, classification of randomisation status, computation of frequencies, and identification of recurring MAC addresses.

3.7 Data Analysis Technique

Analysis proceeded through five sequential stages. The first stage was extraction, namely parsing the activation records from the export file into nine fields according to their storage structure.

The second stage was anonymisation. Voucher identifiers were replaced with sequential labels, while MAC addresses were retained only in their first octet because that is the portion required for classification. MAC addresses occurring more than once were assigned device labels so that persistence analysis remained possible without displaying the actual addresses.

The third stage was classification, namely determining the status of each MAC address as randomised or genuine on the basis of the classification instrument described in Subsection 3.6.

The fourth stage was descriptive analysis, covering computation of the frequency and proportion of randomised addresses together with the 95 per cent confidence interval for that proportion, the distribution of activations by service class, the

distribution of activation times, and the distribution of first octet values among randomised addresses.

The fifth stage was persistence analysis, namely identification of MAC addresses occurring in more than one activation event together with computation of the time gaps between occurrences. The results of this stage were related to the configuration parameters of the device-based continuity mechanism in order to explain the observed authentication behaviour. The analysis was supplemented by a review of the system configuration to identify parameters affecting the authentication mechanism and the enforcement of service policy.

3.8 Ethical Considerations and Data Protection

This study analyses operational metadata and does not access the content of user communications. All identifiers capable of identifying individuals or devices were masked before data processing. Voucher identifiers are not displayed in any form because they function as access credentials. MAC addresses are displayed only in the portion required for classification. The network name and the geographical location of the research environment are not disclosed.

3.9 Replicability

The research procedure can be replicated in other environments with comparable characteristics through the following steps: obtaining a configuration export from a hotspot-based network controlling device; extracting activation records together with the MAC address field; classifying randomisation status using the bit test described in Subsection 3.6; and computing proportions and conducting persistence analysis in accordance with Subsection 3.7. The classification instrument is deterministic and involves no subjective judgement by the researchers, so classification of the same data will always yield identical results.

4. Results and Discussion

4.1 Characteristics of Service Activation

During the observation period from 26 July to 22 August 2026, 46 service activation events were recorded on the network under study. These activations were distributed across six service classes, concentrated in two mid-range classes, namely the 7 Mbps class with a 31-day duration at 18 activations and the 10 Mbps class with a 31-day duration at 12 activations. Together these two classes account for 65.2 per cent of all activations. The full distribution is presented in Table 3.

Table 3. Distribution of Service Activations by Service Class

Service class	Activations	Randomised MAC	Randomised proportion (%)
7 Mbps, 31 days	18	13	72.2
10 Mbps, 31 days	12	9	75.0
7 Mbps, 7 days	6	6	100.0
7 Mbps, 30 days	6	5	83.3
10 Mbps, 30 days	3	3	100.0
5 Mbps, 30 days	1	1	100.0

Service class	Activations	Randomised MAC	Randomised proportion (%)
Total	**46**	**37**	**80.4**

Source: Processed research data, 2026.

The randomised proportion within each service class cannot be interpreted as a meaningful difference because of the very small cell sizes in several classes. A value of 100 per cent in a class containing a single activation, for instance, carries no statistical explanatory power. That column is presented for descriptive completeness only.

The concentration in two mid-range classes indicates that although the network offers eight service profiles spanning 5 to 30 Mbps, actual uptake is centred on the 7 to 10 Mbps range. The three highest-speed service classes recorded no activations at all during the observation period. This finding indicates a gap between the service capacity offered and the demand actually absorbed in a small-scale community environment.

The distribution of activation times exhibits a pattern consistent with household usage characteristics. As shown in Figure 2, activations are concentrated in the late afternoon and evening, peaking between 16.00 and 19.59, which accounts for 37.0 per cent of all activations. By contrast, the midnight to early morning range recorded only one activation.

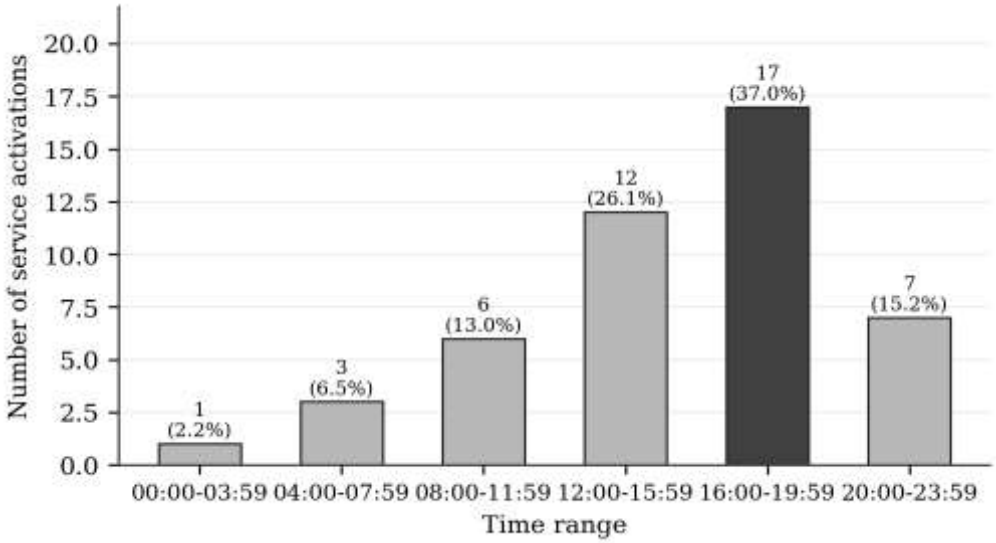

Figure 2. Distribution of Service Activation Times
Source: Processed research data, 2026.

4.2 Level of MAC Address Randomisation

Classification of the 46 MAC addresses recorded at activation events shows that 37 addresses, or 80.4 per cent, are locally administered, whereas 9 addresses, or 19.6 per cent, are universally administered by the manufacturer. With a sample size of 46, that proportion carries a 95 per cent confidence interval spanning 69.0 to 91.9 per cent. Details of the classification are presented in Table 4.

Table 4. Classification of MAC Address Randomisation Status

Description	Value
Randomised addresses (locally administered)	37 (80.4%)
Genuine addresses (universally administered)	9 (19.6%)
Total activation events	46 (100.0%)
95% confidence interval of randomised proportion	69.0% - 91.9%
Standard error of the proportion	5.8%
Unique MAC addresses	41
Distinct first octet values among randomised addresses	23

Source: Processed research data, 2026.

That proportion is relatively stable across periods. July 2026 recorded 81.8 per cent of 11 activations, while August 2026 recorded 80.0 per cent of 35 activations, as shown in Figure 3. This stability indicates that the observed proportion is not an artefact of any particular period. It should be emphasised that this cross-period comparison is descriptive; the July sample is too small to support meaningful statistical testing.

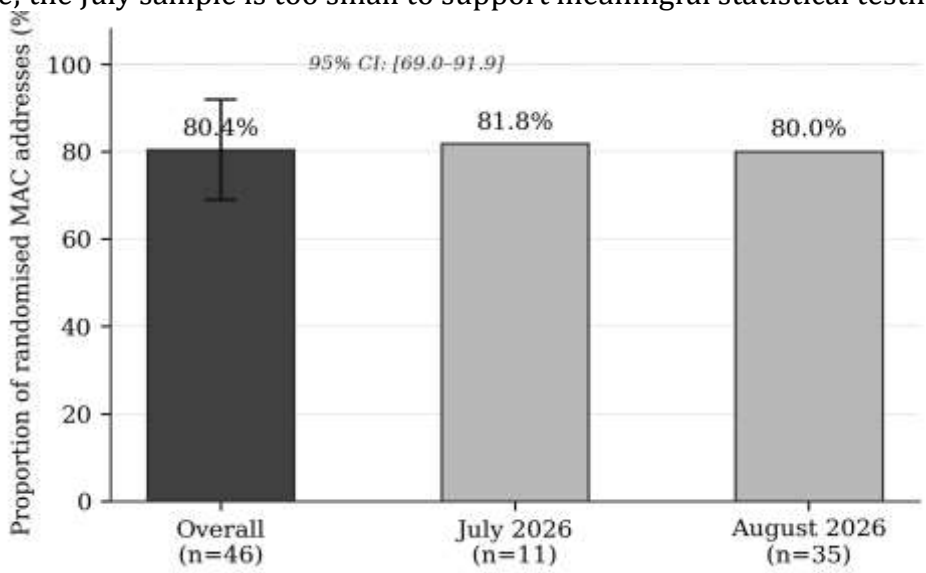

Figure 3. Proportion of Randomised MAC Addresses with Confidence Interval
Source: Processed research data, 2026.

The distribution of first octet values among randomised addresses strengthens the validity of the classification. As shown in Figure 4, the 37 randomised addresses contain 23 distinct first octet values, with the highest frequency being only four occurrences. This even distribution is consistent with a random address generation process and indicates that these addresses do not originate from a single manufacturer identifier block.

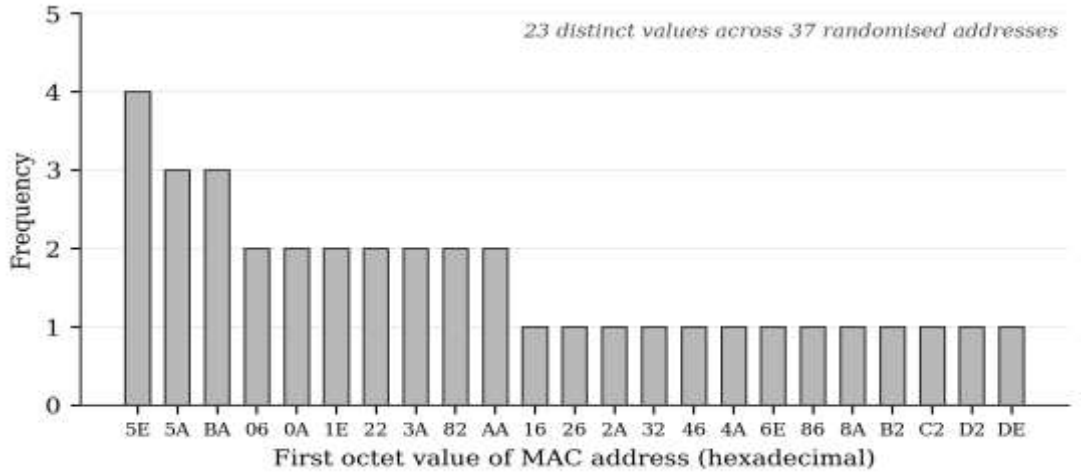

Figure 4. Distribution of First Octet Values among Randomised MAC Addresses
Source: Processed research data, 2026.

This finding is consistent in direction with earlier research reporting increased adoption of MAC address randomisation on mobile devices (Fenske et al., 2021; Martin et al., 2017). Direct comparison with the figures in those studies must nonetheless be made with care, because the observation units differ. Earlier research generally measured at the level of probe request frames, that is, before a device associates with a network, whereas the present study measures at the level of service activation events that have already passed through authentication. Agreement in the direction of findings across two different observation units in fact reinforces the general picture that MAC address randomisation has become default behaviour on the majority of user devices.

4.3 MAC Address Persistence

The 46 activation events yielded 41 unique MAC addresses. The difference arises from five addresses appearing in two separate activation events with different voucher identifiers, as presented in Table 5. This pattern indicates that the same device renewed its service using a new voucher.

Table 5. MAC Address Persistence across Repeated Activations

Device code	MAC status	Number of activations	Gap between activations (days)
P-03	Randomised	2	11
P-01	Randomised	2	6
P-05	Randomised	2	4
P-02	Randomised	2	3
P-04	Genuine	2	1

Source: Processed research data, 2026.

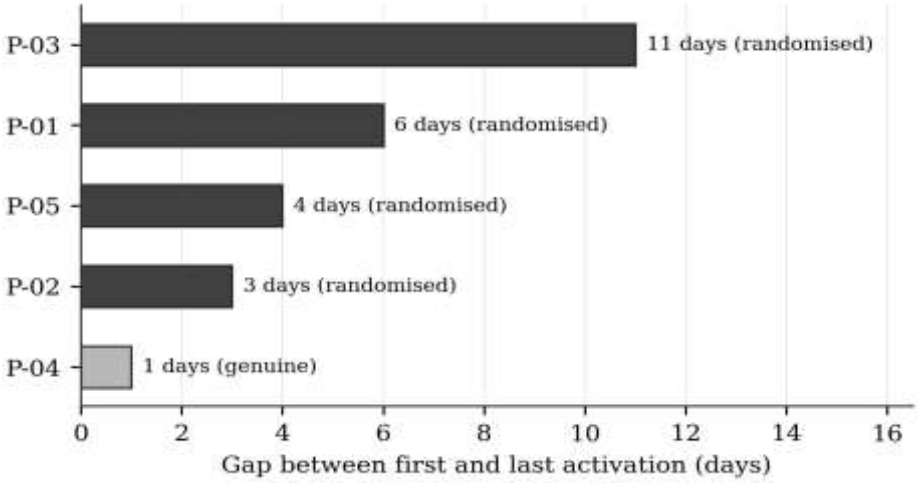

Figure 5. MAC Address Persistence across Repeated Activations
Source: Processed research data, 2026.

Four of the five addresses are randomised, with gaps between activations ranging from 3 to 11 days. This finding indicates that randomised addresses in the environment under study are persistent, that is, they remain unchanged over a span of several days rather than changing at every connection.

This persistence characteristic is important for understanding the observed authentication behaviour. Randomisation schemes commonly implemented by mobile operating systems produce an address that is stable for a given network, yet allow a device to present a different address on a different network name (Henry & Lee, 2025). The configuration of the system under study enables a device-marker continuity mechanism, so that a previously recognised device can obtain access without re-entering service credentials.

The combination of these two factors produces two distinct behaviours, as depicted in Figure 6. When a device reconnects to the same network name, the address it presents remains unchanged, so the device marker is recognised and access is granted without re-authentication. Conversely, when a device moves to a different network name, the address it presents changes, so the device marker is not recognised and the system requests service credentials again.

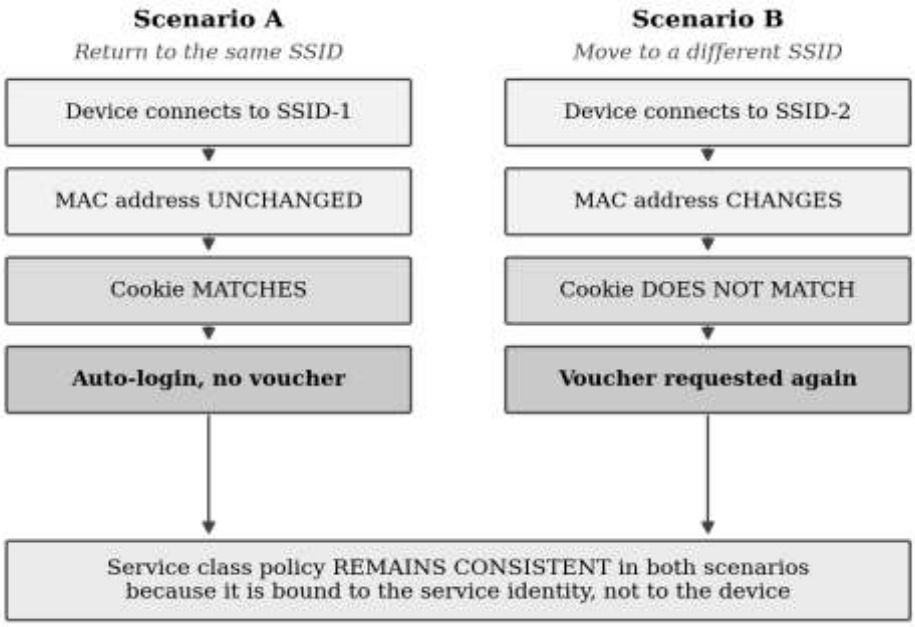

Figure 6. Authentication Mechanism under Same and Different Network Name Scenarios
Source: Authors' analysis, 2026.

Because the environment under study applies distinct network naming for every access point and every frequency band, the practical consequence is that movement between access points and movement between frequency bands on the same access point alike trigger a request for re-authentication.

4.4 Implications for the Choice of Service Management Architecture

The findings on the level and characteristics of MAC address randomisation carry direct implications for the choice of identification basis in service management architecture. A comparison of the resilience of the two mechanisms operating within the system under study is presented in Table 6.

Table 6. Comparison of Management Mechanism Resilience against Device Identity Instability

Aspect	Device-based	Service-identity-based
Identification basis	MAC address	Service credential
Behaviour on the same network name	Functions	Functions
Behaviour on a different network name	Does not function	Functions
Dependence on MAC stability	High	None
Consistency of speed class	Not applicable	Maintained
Consistency of validity period	Not applicable	Maintained

Source: Authors' analysis based on configuration review and activation data, 2026.

The device-based continuity mechanism relies on the assumption that the MAC address is a stable identifier. In the environment under study, that assumption holds only

in part: the address is stable within the context of a single network name, but unstable when the device moves between network names. As a result, the mechanism functions in some usage scenarios and fails in others.

By contrast, a management mechanism that places the service credential as the identification basis is unaffected by that instability. Speed class, validity period, and the simultaneous device limit continue to be enforced consistently because these parameters are bound to the service identity rather than to the device using it. A user moving between access points still receives the same service class, even though re-authentication is required.

Instability of device identity therefore causes degradation at the level of usage convenience but does not cause failure at the level of service policy enforcement. This separation of consequences indicates that placing service identity as the basis of management affords better resilience against changes in user device behaviour than placing device identity as the basis of management.

These findings supply empirical support, from a small-scale operational environment, for the direction already formulated in standards documents and discussed in the literature, namely a shift in the identification basis from hardware identifiers towards credential-based identifiers (Henry & Lee, 2025; Seralathan, 2026). The contribution of this study lies in providing measured evidence from a paid community network context that has hitherto been under-represented in empirical measurement.

4.5 Practical Implications

For operators of community networks with comparable characteristics, these findings carry several practical implications. First, counting devices on the basis of MAC addresses cannot be relied upon. With 80.4 per cent of addresses randomised and changing between network names, a single physical device may be recorded as several distinct devices. Policies limiting the number of simultaneous devices that rest on MAC address counting are therefore liable to be inaccurate.

Second, the network naming strategy directly affects the frequency of re-authentication experienced by users. Distinct naming for every access point and every frequency band multiplies the number of network names, and each movement between those names triggers a renewed request for credentials.

Third, centralised service-identity-based management is shown to preserve policy consistency under conditions of unstable device identity. For environments with limited management resources, this characteristic affords reliability without requiring additional mechanisms.

4.6 Limitations of the Study

Several limitations should be borne in mind when interpreting these results. The study was conducted in a single implementation environment with 31 active service identities, so the findings are contextual and not intended for generalisation.

The sample comprises 46 activation events still stored on the controlling device at the time of data collection, covering a period of 28 days. Records from earlier periods were no longer available. This sample size yields a relatively wide confidence interval, and the persistence analysis rests on only five cases of recurring MAC addresses, making it indicative rather than conclusive.

The operating systems of user devices were not recorded in the available data, so randomisation behaviour could not be disaggregated by platform. The originating access

point was likewise not recorded because all access points reside on a single network segment without logical separation, so movement between access points could not be counted directly.

Re-authentication events were not counted directly because the available records capture only voucher activation. Conclusions regarding the impact on re-authentication are derived from the mechanism identified through configuration review and persistence analysis rather than from a count of events.

5. Conclusion

This study measured the level of MAC address randomisation in authenticated sessions within a small paid community network and analysed its implications for the choice of identification basis in service management architecture.

In line with the first objective, of the 46 service activation events observed over 28 days, 80.4 per cent presented locally administered MAC addresses, with a 95 per cent confidence interval spanning 69.0 to 91.9 per cent. That proportion was stable across observation periods, and the distribution of first octet values, comprising 23 distinct values among 37 randomised addresses, exhibits a pattern consistent with random address generation.

In line with the second objective, five MAC addresses were identified as appearing in two separate activation events. Four of these were randomised addresses that remained unchanged over spans of 3 to 11 days. This finding indicates that randomisation in the environment under study is persistent for a given network name. The combination of that persistence with distinct network naming for every access point and every frequency band causes the device-marker continuity mechanism to function when a device returns to the same network name but to fail when a device moves to a different network name.

In line with the third objective, policies governing speed class, validity period, and the simultaneous device limit continued to be enforced consistently under both scenarios, because these parameters are bound to the service identity rather than to the device. Instability of device identity therefore produces degradation at the level of usage convenience but not failure at the level of service policy enforcement.

Theoretically, these findings supply empirical support, from a small-scale operational environment, for the shift in identification basis from hardware identifiers towards credential-based identifiers, a shift hitherto formulated normatively and discussed chiefly in enterprise network contexts. Practically, the findings show that counting devices on the basis of MAC addresses cannot be relied upon in comparable environments, and that the network naming strategy directly affects the frequency of re-authentication experienced by users.

This study is limited to a single implementation environment with 46 activation events over 28 days, so its findings are contextual. Further research is recommended to extend the observation period and the number of environments studied, to record re-authentication events directly so that their impact can be counted rather than inferred from mechanism, and to design and test mechanisms that preserve authentication continuity across network names without depending on the stability of device addresses.

References

Astillo, P. V. B., Sengupta, J., & Vanhoef, M. (2026). Open challenges for secure and scalable Wi-Fi connectivity in rural areas. In *Proceedings of the 1st Workshop on Security and Privacy for Asian Internet Communities (SPAIC), AsiaCCS Workshops*. https://arxiv.org/abs/2604.02774

Avonts, J., Braem, B., & Blondia, C. (2013). A questionnaire based examination of community networks. In *2013 IEEE 9th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)* (pp. 8-15). IEEE.

Bass, L., Clements, P., & Kazman, R. (2012). *Software architecture in practice* (3rd ed.). Addison-Wesley.

Braem, B., Baig Vinas, R., Kaplan, A. L., Neumann, A., Vilata i Balaguer, I., Tatum, B., Matson, M., Blondia, C., Barz, C., Rogge, H., Freitag, F., Navarro, L., Bonicioli, J., Papathanasiou, S., & Escrich, P. (2013). A case for research with and on community networks. *ACM SIGCOMM Computer Communication Review, 43*(3), 68-73. https://doi.org/10.1145/2500098.2500108

Chiba, M., Dommety, G., Eklund, M., Mitton, D., & Aboba, B. (2008). *Dynamic authorization extensions to Remote Authentication Dial In User Service (RADIUS)* (RFC 5176). Internet Engineering Task Force.

Fenske, E., Brown, D., Martin, J., Mayberry, T., Ryan, P., & Rye, E. (2021). Three years later: A study of MAC address randomization in mobile devices and when it succeeds. *Proceedings on Privacy Enhancing Technologies, 2021*(3), 164-181. https://doi.org/10.2478/popets-2021-0042

Henry, J., & Lee, Y. (2025). *Randomized and changing Media Access Control (MAC) addresses: Context, network impacts, and use cases* (RFC 9797). Internet Engineering Task Force. https://www.rfc-editor.org/info/rfc9797

Maccari, L., & Lo Cigno, R. (2015). A week in the life of three large wireless community networks. *Ad Hoc Networks, 24*, 175-190.

Martin, J., Mayberry, T., Donahue, C., Foppe, L., Brown, L., Riggins, C., Rye, E. C., & Brown, D. (2017). A study of MAC address randomization in mobile devices and when it fails. *Proceedings on Privacy Enhancing Technologies, 2017*(4), 365-383. https://doi.org/10.1515/popets-2017-0054

Rigney, C. (2000). *RADIUS accounting* (RFC 2866). Internet Engineering Task Force.

Rigney, C., Willens, S., Rubens, A., & Simpson, W. (2000). *Remote Authentication Dial In User Service (RADIUS)* (RFC 2865). Internet Engineering Task Force.

Seralathan, P. (2026). *Persistent device identity for network access control in the era of MAC address randomization: A RADIUS-based framework* [Preprint]. arXiv. https://arxiv.org/abs/2603.16745

Strassner, J. (2003). *Policy-based network management: Solutions for the next generation*. Morgan Kaufmann.

Zuniga, J. C., Bernardos, C. J., & Andersdotter, A. (2025). *State of affairs for randomized and changing Media Access Control (MAC) addresses* (RFC 9724). Internet Engineering Task Force. https://www.rfc-editor.org/info/rfc9724

Acknowledgment

The authors thank the operator of the implementation environment for granting permission to use operational data for the purposes of this study.