Domain Segmentation, Trust Federation and Performance Trade-offs for the Future Railway Mobile Communication System

Martin Reed
University of Essex, UK
mjreed@essex.ac.uk

Davide Amato
Sadel, IT
davide.amato@sadel.it

Omar Benjumea
Cylus, Spain
omar@cylus.com

Daniele Bozzolo
ANIE/ASSIFER, Italy
daniele.bozzolo.ext@anie.it

Paolo Volpato
Huawei Technologies, Italy
paolo.volpato@huawei.com

***Abstract*—The Future Railway Mobile Communication System (FRMCS) will replace the Global System for Mobile Communications – Railway (GSM-R) as the communication backbone of the European Rail Traffic Management System (ERTMS). FRMCS inherits strong 5G security mechanisms: concealed subscriber identities, mutual authentication, and secure inter-domain signalling; however, its deployment across national and infrastructure-manager boundaries introduces architectural and governance trade-offs.**

This paper combines architectural analysis with quantitative modelling of (i) domain segmentation as a containment control, (ii) roaming architectures (home-routed vs. visited-routed/local breakout), and (iii) boundary-driven rekeying overhead. Results show that segmentation materially limits lateral compromise propagation, while visited-routed roaming substantially reduces latency compared to home routing. However, excessive fragmentation increases boundary crossings and rekeying events, yielding a measurable resilience–operational-complexity trade-off. These findings are discussed in the context of wider implementation and motivation for a federated Railway Trust Framework for cross-domain assurance.

I. Introduction

The Future Railway Mobile Communication System (FRMCS) is the designated successor to the Global System for Mobile Communications – Railway (GSM-R) within the European railway digitalisation roadmap, and will provide the communication bearer for the European Rail Traffic Management System (ERTMS), including ETCS Level 2 automation and future moving-block signalling. It will support mission-critical voice, data, and video services (MCX) over a 5G-based architecture. Unlike GSM-R, FRMCS adopts a service-based architecture (SBA) with virtualised core functions and cryptographic identity management derived from 5G standards [1].

Although the underlying FRMCS primitives are strong, railway operational reality adds constraints that are less prominent in consumer mobile networks: cross-border operation; multiple infrastructure managers and mobile network operators; long-lived operational technology (OT); safety certification and deterministic performance requirements; and evolving regulatory and standards obligations. In this setting, security design must be expressed as measurable engineering trade-offs rather than solely as *control checklists*.

This paper evaluates three coupled design dimensions:

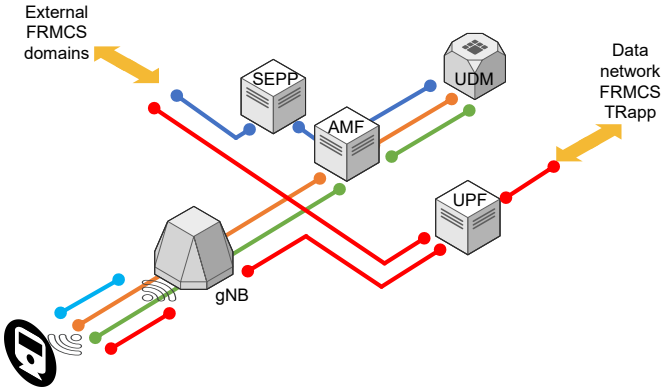

Fig. 1. FRMCS architecture (simplified) based upon 3GPP 5G core functions. The figure highlights security-relevant functions (e.g., UDM, AMF, UPF, SEPP).

- **Containment**: how domain segmentation limits lateral compromise propagation.
- **Performance**: latency impact of roaming architecture choices.
- **Operational overhead**: rekeying and boundary costs under different domain granularities.

Contributions. (i) We formalise a simple, interpretable compromise-propagation model that captures the effect of inter-domain isolation. (ii) We quantify performance differences between home-routed and visited-routed roaming under realistic segmentation schemes. (iii) We model and measure boundary-driven rekeying overhead. (iv) We translate the results into design guidance and governance implications, including requirements for a federated Railway Trust Framework (RTF). The results are made from modelling based on estimates of future network scenarios; as such, the contributions aim to show general trends rather than precise predictions and a methodology that can be used for more precise modelling when real network parameters are known in the future.

II. Background and Related Work

FRMCS is not simply a telecommunications upgrade to GSM-R; it represents a transition from circuit-switched railway radio to a 5G SBA in which identity, policy and cryptographic trust are first-class design elements [2]. Fig. 1

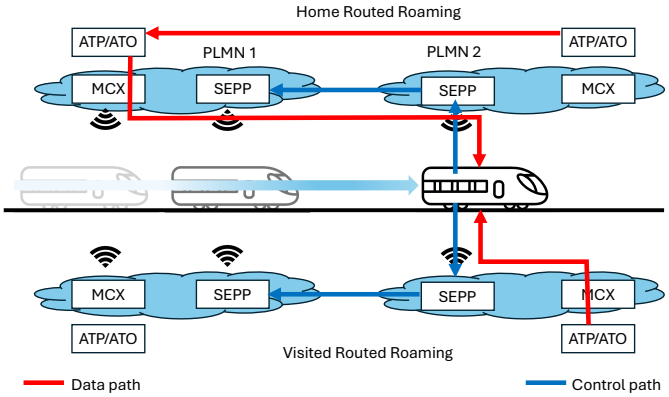

Fig. 2. SEPP-mediated inter-domain roaming in FRMCS (home-routed vs. visited-routed/local breakout).

illustrates the principal 5G-derived core functions relevant to FRMCS security [3], including the User Data Management (UDM) function responsible for subscriber identity and credential storage; the Access and Mobility Management Function (AMF), which handles registration, authentication, and mobility procedures; the User Plane Function (UPF), which forwards user traffic; and the Security Edge Protection Proxy (SEPP), which secures inter-operator signalling. These functions interact with the radio access network via the 5G base station (gNB) and support railway applications such as Mission Critical Services (MCX) communication and European Train Control System (ETCS) data exchange.

Operationally, FRMCS must sustain secure, low-latency connectivity as trains move at high speed across radio cells and, in many cases, across administrative or national boundaries. This mobility requirement makes handover, key continuity, and inter-domain trust central architectural concerns rather than secondary optimisations. Fig. 2 highlights how Security Edge Protection Proxies (SEPPs) mediate signalling across the boundaries of the Public Land Mobile Network (PLMN) and support *home-routed* and *visited-routed* roaming configurations. While consumer mobile deployments commonly use home-routed roaming, FRMCS scenarios may require visited-routed architectures for operational autonomy. This introduces additional inter-PLMN trust and key management constraints as specified in the 5G security architecture [1].

A. Railway security standards and governance context

Railway cyber security is typically engineered alongside safety under lifecycle standards and sector guidance (e.g., CLC/TS 50701 [4]) and is evolving into the international standard IEC 63452 [5], [6]. At the European level, the regulatory environment is increasingly explicit about operational resilience and supply-chain assurance [7], [8]. These pressures motivate architectures that combine technical security primitives with cross-domain governance and verifiable assurance. However, recent analysis has shown that these evolving standards do not take into account recent moves to zero-trust architectures and do not consider the increasing need to manage security across segmented domains [9].

B. Segmentation and zero trust in OT/critical infrastructure

Network segmentation and "zones and conduits" patterns are widely used in OT security to limit lateral movement and to enable enforceable policy boundaries. In parallel, zero-trust principles [10] emphasise continuous verification, explicit trust boundaries, and strong identity as a basis for access decisions. The FRMCS SBA aligns naturally with identity-centric, policy-driven access control, but its deployment across multiple administrative domains raises questions of federation, certificate trust, and operational overhead that we address in this paper.

Beyond the telecommunications layer, FRMCS operates within a railway ecosystem characterised by long asset lifecycles, safety certification constraints, and cross-border governance obligations. Unlike conventional enterprise IT systems, railway OT must satisfy deterministic performance and safety-case requirements under standards such as EN 50126/9 and CLC/TS 50701, with IEC 63452 [11] emerging as the international cyber security reference for railways. The sector has historically adopted "zones and conduits" segmentation patterns [12] aligned with IEC 62443 [13], reflecting the need to control lateral movement between safety-relevant domains.

This broader context is significant for FRMCS because the service-based architecture naturally aligns with identity-centric and policy-driven (zero-trust) security models, yet its deployment must coexist with legacy signalling systems, Radio Block Centres (RBCs), and long-lived onboard assets. Consequently, FRMCS security should be understood not only as a 5G inheritance, but as a component of a layered railway cyber-physical assurance framework that integrates telecommunications security, operational segmentation, and cross-domain governance.

C. FRMCS security foundations

The initial registration and full 5G-AKA exchange provide strong mutual authentication but are comparatively expensive in signalling terms. While acceptable when a train first attaches to a domain, repeating full authentication at every cell boundary would introduce unacceptable latency and signalling overhead. To address this, FRMCS leverages the hierarchical key-derivation structure defined in 3GPP 5G. Following successful authentication, a master security context is established at the AMF. During intra-domain handover (i.e., when a train moves between gNBs within the same AMF/PLMN), the security context is transferred securely between network functions without repeating the full 5G-AKA procedure. This enables fast handover performance compatible with high-speed railway operation, while ensuring that encryption and integrity protection are never suspended during mobility events. While the exact overhead reduction from this improved handover is difficult to quantify, as it depends upon the hardware implementation, a comparative study shows that it reduces the overhead by several orders of magnitude [14]. This reduction

in the security-context handover allows us to absorb this latency within the general access capture delay we define later.

When a handover coincides with a PLMN or administrative boundary crossing, the control-plane exchange is mediated by SEPP entities (Fig. 2). In this case, authentication context and policy enforcement must be validated across domains. Depending on roaming configuration, either (i) the user plane remains anchored in the home domain (home-routed), or (ii) it is locally terminated in the visited domain (visited-routed/local breakout). The former maximises centralised policy control but introduces additional path length; the latter reduces latency but requires stronger inter-domain trust guarantees. In both cases, the SEPP ensures confidentiality, integrity, and topology hiding for inter-operator signalling. This trade-off between home-routed and visited-routed systems has not previously been studied in the context of FRMCS. Consequently, this, together with other trade-offs in segmentation discussed above, are motivations for this paper.

III. System, Threat Model, and Segmentation Schemes

A. System and roaming model

We consider a pan-European FRMCS deployment consisting of multiple administrative domains (e.g., country or infrastructure-manager networks), each providing radio access (gNBs) and 5G core functions. When trains cross the boundaries of the Public Land Mobile Network (PLMN), inter-domain control-plane signalling is mediated by SEPP. Two roaming modes are considered: (i) *home-routed*, in which user-plane traffic remains anchored in the home network; and (ii) *visited-routed* (local breakout), in which the user plane terminates in the visited domain.

B. Threat model

We focus on the propagation of compromise within the FRMCS-connected railway ecosystem. We take an *attack graph* approach to this compromise propagation [15]. We consider two initial compromise origins in our simulations: **(T)** train onboard network/endpoint compromise, and **(M)** maintenance or back-office compromise. Once an initial node is compromised, the attacker attempts lateral movement to adjacent nodes (logical adjacency) with probability dependent on whether the move is intra-domain or cross-domain (inter-domain). The model does not assume specific exploits; it abstracts the effect of segmentation as increased friction at boundaries once the initial compromise has occurred.

C. Segmentation schemes

We evaluate four segmentation strategies, representative of plausible governance boundaries:

1) **Flat**: a single logical security domain.
2) **Region**: a small number of large domains (e.g., multi-country or macro-regional).
3) **Country**: domain boundaries aligned with national or primary infrastructure-manager boundaries.
4) **Sub-country**: finer granularity (e.g., multiple domains within a country/region).

TABLE I
Segmentation schemes considered and their expected qualitative effects.

Scheme	Boundary count	Expected effect
Flat	None	Minimal overhead; maximal lateral movement potential.
Region	Low	Some containment; may still allow large compromise radius.
Country	Medium	Stronger containment aligned to governance; moderate boundary overhead.
Sub-country	High	Strong containment; increased boundary crossings and rekeying.

Table I summarises the intent and trade-offs of these segmentation schemes.

IV. Model and Simulation Parameters

We implement Monte Carlo simulations over representative network instances for each segmentation scheme. The network was modelled using the EU defined Trans-European Transport Network [16] and using major EU stations for typical EU transit; FRMCS nodes were spaced a mean distance of 8 km as might be assumed from existing modelling [17]. This resulted in 2452 TRACKSIDE elements (each with a RAN_RADIO) and 103 CORE nodes distributed across representative national and regional railway operational data centre locations within the EU model.

A. Compromise propagation model

We model compromise as a probabilistic process over a graph with nodes that represent railway-connected assets, including trackside infrastructure and onboard units. The links of the graph represent connectivity that is modelled as either fixed communication systems (e.g. fibre along trackside) or RF connected systems (e.g. the RF connectivity of FRMCS). For an attempted move from node u to adjacent node v:

$$P(u \to v) = \begin{cases} p_{\text{intra}}, & \text{if } u \text{ and } v \text{ in the same domain} \\ p_{\text{inter}} & \text{if } u \text{ and } v \text{ in different domains} \end{cases}$$

where $p_{\text{inter}} \in (0, 1]$ is an isolation factor capturing boundary controls across segmented FRMCS *domains*, e.g., policy enforcement, authentication, monitoring, and administrative separation. Lower p_{inter} corresponds to stronger isolation across FRMCS domains.

This separation of intra-zone likelihood from inter-zone traversal aligns with ENISA's railway zoning-and-conduits guidance and the IEC 62443 zoning/conduit security model, where segmentation limits *blast radius* by constraining movement between zones rather than assuming perfect internal hardening.

Within a domain, lateral movement probabilities, p_{intra}, are parameterised using qualitative bands mapped to numeric ranges across different technology *zones*: High (0.8–0.95) for homogeneous or privileged pathways (e.g.,

MAINT→infrastructure and TRACKSIDE→TRACKSIDE), Medium (0.5–0.8) for operational interconnections (e.g., RAN_RADIO↔CORE↔TRACKSIDE), and Low (0.05–0.2) for constrained or segregated interfaces (e.g., TRAIN_ONBOARD→outward transitions). These values represent conditional success given adjacency and are intentionally pessimistic within-zone assumptions, reflecting plausible exploit reuse and shared technology stacks *once trust is breached*.

Inter-domain movement is assumed to cross controlled conduits with stronger policy enforcement, authentication, monitoring, and administrative separation. We therefore set $p_{\text{inter}} = 0.001$, approximately one to two orders of magnitude below the intra-domain probabilities, to model strong but not perfect boundary isolation. This deliberately preserves a residual probability of cross-domain compromise while ensuring that the simulations primarily test whether segmentation can limit blast radius after an initial compromise. Later we briefly comment on the sensitivity of the model to p_{inter} but note here that it is not a critical value.

Each compromise-propagation run selects an initial compromise origin TRAIN_ONBOARD (T) or MAINT (M) with an arbitrary 100 nodes randomly selected in each case, then iteratively attempts lateral moves until no further moves succeed. We do not model the initial compromise probability; this is because we are most interested to observe the effects of segmentation and compromise location *after a compromise has occurred*. We record the total number of compromised nodes and report empirical cumulative distribution functions (ECDFs).

B. Latency model: home-routed vs. visited-routed

To compare roaming architectures, without committing to a specific vendor timing model, we use a path-based latency model derived from the network graph. For any two nodes u, v we define a one-way latency

$$\ell(u,v) \;=\; \alpha\, d(u,v) \;+\; \beta\, h(u,v) \;+\; \gamma, \tag{1}$$

where $d(u,v)$ is the shortest-path distance, $h(u,v)$ is the corresponding hop count (unweighted), α is the propagation latency, β is the per-hop forwarding/processing latency, and γ is an access/edge term capturing radio scheduling and fixed processing delays that are largely independent of the chosen core routing. This model is intended to support *relative* comparisons between architectures; absolute values depend on deployment-specific constants. We use nominal coefficients e.g. $\alpha = 0.005$ ms/km, $\beta = 0.15$ ms/hop, $\gamma = 5$ ms; chosen to produce plausible magnitudes for 5G transport scenarios [18]. The conclusions are driven by topology and anchoring effects rather than precise coefficient values.

We then model two roaming cases for delivery from a visited-domain control/core anchor:

Visited-routed (local breakout). User-plane delivery remains within the visited domain:

$$L_{\text{local}} \;=\; \ell(c_{\text{vis}}, t), \tag{2}$$

where c_{vis} is a sampled core/control node in the visited domain and t is a sampled trackside attachment node.

Home-routed (anchored). User-plane delivery is anchored via a home-domain core node (a "home agent"), producing a *tromboning* path:

$$L_{\text{home}} \;=\; \ell(c_{\text{vis}}, c_{\text{home}}) \;+\; \ell(c_{\text{home}}, t), \tag{3}$$

where c_{home} is a sampled core/control node in the home domain (chosen to be different from the visited domain when segmentation introduces multiple domains).

C. Boundary rekeying overhead

Finer-grained segmentation increases the frequency of inter-domain boundary events during mobility (e.g., crossing from one operational domain to another), which require additional authentication signalling and security-context establishment. Rather than modelling the full 5G-AKA protocol, we abstract the *added* delay attributable to a boundary event as a fixed number of additional control-plane round trips plus processing time. Indeed, as noted in II-C, the 5G-AKA protocol can be significantly reduced in complexity through hierarchical key-derivation that transfers the security context to visited nodes. Thus, the latency for this process can be absorbed into a fixed processing overhead per-node access (δ below) rather than requiring calculation relative to the originating context.

For each Monte Carlo run, we sample a journey by selecting two trackside attachment nodes and computing the shortest path between them. We then count the number of domain transitions along that path, denoted H. We model added delay as

$$L_{\text{rekey}} \;=\; \sum_{k=1}^{H} \Big(r \cdot 2\ell_k \;+\; \delta \Big), \tag{4}$$

where $r = 2$ is the number of additional round trips required per boundary (e.g., signalling exchanges to establish/refresh a security context), $2\ell_k$ is the control-plane round-trip time for the kth boundary calculated from the latency ℓ as specified in (1); $\delta = 11$ ms captures fixed processing overhead in security functions, based upon reported measurement [19]. This yields a conservative, architecture-driven estimate of the incremental latency cost of crossing more domain boundaries under finer segmentation.

V. Results

A. Containment via segmentation

Figure 3 presents the ECDF of compromised nodes under each segmentation scheme, with compromise originating either from maintenance/back-office systems (M) or onboard train systems (T).

Two clear effects emerge. First, increasing segmentation reduces compromise propagation. Flat and coarse regional topologies exhibit the widest compromise distributions, whereas country and sub-country segmentation shift the ECDF significantly leftwards. Although the ECDF reflects node counts rather than geometric distance, separate analysis (not shown) was carried out to look at finer distribution of trackside

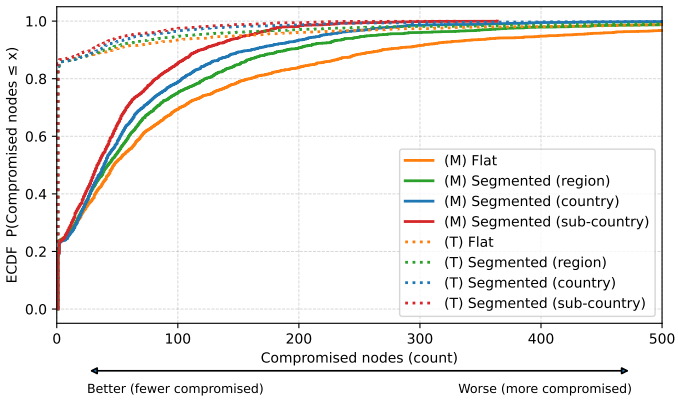

Fig. 3. Compromise size ECDF under varying segmentation schemes and attacks originating from either maintenance/back-office (M) or train onboard (T).

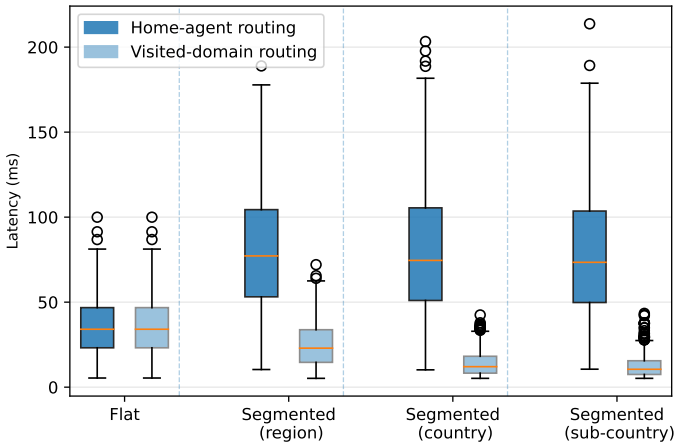

Fig. 4. Latency comparison: home-routed vs. visited-routed (same-domain/local breakout) delivery across segmentation schemes.

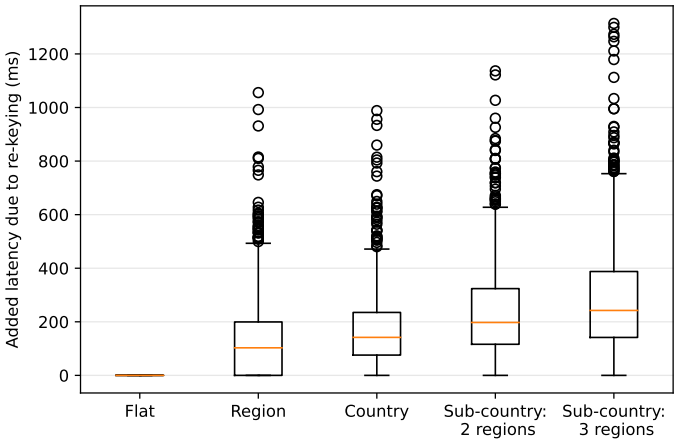

Fig. 5. Additional latency due to inter-domain rekeying under different segmentation granularities.

elements. That analysis confirmed that finer segmentation also reduces effective attack radius. The transition to country-level segmentation is particularly important, making large cascades statistically unlikely. We investigated the sensitivity of the model to values of $p_{\text{inter}} = [1, 0.1, 0.01, 0.001, 0.0001]$. The value $p_{\text{inter}} = 1$ is a degenerate version of the model that effectively has no segmentation and the results confirmed no difference between the scenarios. For smaller values, the spread of the curves shown in Figure 3 separate further for the M or T attack scenarios as p_{inter} decreases but the overall trends are very similar so the results of the sensitivity analysis are not shown here.

Second, compromise originating from maintenance systems is consistently more severe than onboard-origin attacks. Maintenance nodes are directly connected to widely interconnected infrastructure and back-office systems, providing broader lateral reach once compromised, whereas onboard-origin attacks begin from a more constrained position.

These results support segmentation as a primary containment mechanism in FRMCS deployments, particularly where administrative boundaries can be aligned with security domains.

B. Latency: home-routed vs. visited-routed

Figure 4 compares latency distributions under home-routed and visited-routed (local breakout) architectures across segmentation schemes. The dominant effect is architectural: home routing consistently introduces a clear latency increase relative to local breakout. This is the "tromboning" penalty as user-plane traffic is anchored in the home domain, so paths are systematically longer in both physical distance and hop count. Note that in the "flat" case, a single FRMCS spans the whole network so there are no meaningful home- vs. visited-routed scenarios as both are materially the same, but the result is presented for completeness. Domain granularity has limited impact on home-routed performance as latency back to the home domain dominates the latency. For visited-routed, smaller segmentation results in reduced latency as the train will be connected to a more local control centre.

This separation of effects is important: segmentation improves containment (as shown in V-A) without inherently penalising user-plane latency, provided that visited-routing is used. For latency-sensitive safety services (e.g., ETCS data exchange or real-time voice services), these results provide quantitative support for visited-routed architectures.

C. Rekeying overhead under boundary crossings

Figure 5 shows the added latency attributable to boundary-driven rekeying events. As expected, finer-grained segmentation increases the frequency of domain transitions along sampled paths, leading to higher aggregate rekeying delay. The effect is monotonic: sub-country schemes exhibit higher medians and broader variance than country or region schemes.

This comparison reframes the trade-off: the operational cost of segmentation lies primarily in boundary processes, whereas the performance cost of centralisation lies in systematic path inflation. Engineering low-latency boundary authentication therefore offers a path to achieve strong containment without sacrificing user-plane performance.

Taken together, Figs. 3–5 illustrate that segmentation, roaming architecture, and boundary design must be treated as coupled variables rather than independent controls.

VI. Discussion and Governance Implications

The quantitative results reveal that FRMCS security architecture cannot be optimised along a single dimension. Segmentation, roaming design, and boundary operations form a coupled design space in which containment, latency, and operational complexity interact.

Architectural trade-offs

Three primary tensions emerge:

1) Domain scale vs. containment. Large domains reduce boundary frequency and operational overhead, but increase the potential radius of lateral compromise. Flat or overly coarse regional topologies exhibit heavy compromise tails (Fig. 3), indicating systemic risk. Country-aligned (or sub-country) segmentation produces a significant reduction in compromise size, suggesting that governance boundaries are also meaningful security boundaries.

2) Segmentation vs. boundary overhead. Finer-grained segmentation strengthens containment, but increases domain transitions and rekeying events (Fig. 5). However, rekeying overhead remains bounded and can be mitigated by make-before-break mechanisms. This result provides a proxy for increasing overhead implicit in finer segmentation.

3) Centralised control vs. latency. Home-routed roaming maximises central policy visibility and control, but systematically inflates user-plane path length (Figure 4). The dominant latency penalty arises from anchoring rather than segmentation itself. Visited-routed (local breakout) architectures preserve low latency while remaining compatible with strong containment. It should be noted that home-routed is closer to the methodology used for current mobile roaming and is still suitable for less critical systems such as infotainment systems where latency is not as critical. However, for real-time safety critical systems (e.g. ETCS Level 2 automation) it is clear that the latency of home-routed is too high and a visited-routed architecture is most suitable. Whether both mechanisms can be mixed in a real deployment may depend upon whether multi-tenant solutions are deemed appropriate in a railway scenario.

Taken together, these results suggest that the primary performance risk in FRMCS is not segmentation per se, but excessive architectural centralisation.

Design guidance

For latency- and safety-critical services (e.g., ETCS data exchange), visited-routed roaming should be the default architecture. Segmentation should be aligned with stable operational and governance boundaries, national or infrastructure-manager domains, to maximise containment benefit per boundary cost.

Where finer segmentation is desirable (e.g., sub-country operational zones), boundary processes must be engineered for low-latency operation. This includes efficient security context establishment, edge-assisted authentication, and minimised signalling round-trips. In practical terms, segmentation should be viewed as a risk-reduction mechanism, and boundary engineering as a performance-preservation mechanism.

Toward a European Railway Trust Framework

The architectural preference for visited-routing and finer segmentation increases reliance on cross-domain trust. While SEPP protects inter-operator signalling, it does not by itself define a sector-wide assurance and certificate management model.

The results therefore support a recommendation from Railway Working Group of the Datacom Industry Association that coordinated action is required to formalise a European *Railway Trust Framework* for cross-border identity, certificate, and policy trust in FRMCS deployments [9]. At a high level, such a framework should enable consistent trust semantics across administrative domains, support efficient certificate lifecycle management under mobility, and validate FRMCS operation in zero-trust contexts. Continuous observability must be treated as a core design principle, ensuring that boundary enforcement and anomalous behaviour remain measurable in real time within segmented, federated environments.

Limitations and future work

The models presented intentionally abstract deployment specifics in order to isolate architectural effects. Future work should incorporate richer attacker models, privilege hierarchies, and additional heterogeneous asset classes; empirical measurement of roaming and authentication timing in FRMCS pilot deployments; and formalisation of a Railway Trust Framework governance roles, audit mechanisms, and incident coordination consistent with railway safety and cyber security lifecycles. While many of these aspects were captured in the simulation model graph they would benefit from more detailed contextual values from real scenarios. For example while the model in IV considered a number of asset classes (two trackside nodes, onboard nodes etc.) and the graph had differing link attributes to represent different latencies and attack propagation probabilities, these were set using common random variables obtained from other research rather than real network values. Given that this paper is predicting future networks that do not exist at the scale modelled, this approach is appropriate at this point in time.

Nevertheless, the core architectural findings of this paper are clear: segmentation materially reduces compromise propagation, visited routing preserves latency, and boundary overhead remains manageable when engineered deliberately.

VII. Conclusion

FRMCS provides a robust 5G-derived security foundation, but pan-European deployment introduces measurable trade-offs. Quantitative modelling shows that segmentation significantly reduces compromise propagation, while visited-routed roaming significantly improves latency. Increasing FRMCS domain segmentation increases boundary overhead but within reasonable limits under plausible assumptions. These findings

support a federated, zero-trust, segmentation-aware deployment strategy and motivate a Railway Trust Framework for cross-domain assurance.

ACKNOWLEDGMENT

The authors acknowledge the support and collaborative environment provided by the Railway Working Group of the Datacom Industry Association (DIA) which facilitated sector engagement with the relevant standards bodies and discussion informing this research.

REFERENCES

[1] *3GPP System Architecture for the 5G System (5GS)*, 3GPP Std. 3GPP TS 23.501, 2024.

[2] UIC, "FRMCS System Requirements Specification v2.1.0," Union Internationale des Chemins de fer, Tech. Rep., 2025. [Online]. Available: https://uic.org/IMG/pdf/uic_frmcs_frs_fu-7120_v2.1_0.pdf

[3] *Security Architecture and Procedures for 5G System*, 3GPP Std. 3GPP TS 33.501, 2024.

[4] *Railway Applications – Cybersecurity*, CENELEC Std., 2023, CLC/TS 50701.

[5] S. Benoliel, "IEC 63452 – Toward the First International Standard on Railway Cybersecurity," in *Proceedings of the 8th ICT for Railways Conference*, Madrid, Spain, October 2025, 2–3 October 2025. [Online]. Available: https://www.ict-for-railways.eu/download.aspx?id=d3e9122c-98ec-422a-94de-86fa157f6f2c

[6] International Electrotechnical Commission (IEC), "IEC 63452 – Railway applications – Cybersecurity (in development)," IEC TC 9 Project Team 63452, expected publishing date September 2026, Committee Draft under development; successor to CLC/TS 50701. [Online]. Available: https://www.iec.ch/dyn/www/f?p=103:23

[7] European Union, "Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity Across the Union (NIS2)," 2022. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[8] European Parliament and Council of the European Union, "Regulation (EU) 2024/2847 of the European Parliament and of the Council on Horizontal Cybersecurity Requirements for Products with Digital Elements (Cyber Resilience Act)," Nov 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/2847/oj

[9] O. Benjumea, D. Bozzolo, M. Reed, and P. Volpato, "Cyber Security for the Evolving FRMCS Deployment in the European Railway Sector," Datacom Industry Association (DIA) Railway Working Group, Brussels, Belgium, White Paper, Dec 2025. [Online]. Available: https://datacom-ia.eu/wp-content/uploads/2026/01/DIA-RailVert-CS-FRMCS-final.pdf

[10] S. W. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," National Institute of Standards and Technology (NIST), Gaithersburg, MD, USA, Special Publication (SP) 800-207, Aug 2020. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-207

[11] S. Benoliel, "Session 3-1: IEC PT 63452 Update," in *Proc. 4th ERA–ENISA Conference on Cybersecurity in Railways*, Lille, France, Oct 2024, presentation on the development and timeline of IEC 63452. [Online]. Available: https://www.era.europa.eu/system/files/2024-10/session%203-1%20-%20iec%20pt%2063452%20update%20-%20serge%20benoliel.pdf

[12] European Union Agency for Cybersecurity (ENISA), "Zoning and Conduits for Railways: Security Architecture," ENISA, Tech. Rep., 2021. [Online]. Available: https://www.enisa.europa.eu/publications/zoning-and-conduits-for-railways

[13] International Electrotechnical Commission (IEC), *IEC 62443-3-3: System security requirements and security levels*, IEC Std., 2013.

[14] J. Huang and Y. Qian, "A secure and efficient handover authentication and key management protocol for 5G networks," *Journal of Communications and Information Networks*, vol. 5, no. 1, pp. 40–49, 2020.

[15] M. Alhomidi and M. Reed, "Risk assessment and analysis through population-based attack graph modelling," in *World Congress on Internet Security (WorldCIS-2013)*, 2013, pp. 19–24.

[16] European Commission, Directorate-General for Mobility and Transport, "Trans-European Transport Network (TEN-T): Overview and European Transport Corridors," 2024, trans-European Transport Network policy overview, including core and extended transport corridors initiatives. [Online]. Available: https://transport.ec.europa.eu/transport-themes/infrastructure-and-investment/trans-european-transport-network-ten-t

[17] R. He, B. Ai, G. Wang, K. Guan, Z. Zhong, A. F. Molisch, C. Briso-Rodriguez, and C. P. Oestges, "High-speed railway communications: From GSM-R to LTE-R," *IEEE Vehicular Technology Magazine*, vol. 11, no. 3, pp. 49–58, 2016.

[18] B. Coll-Perales, M. C. Lucas-Estañ, T. Shimizu, J. Gozalvez, T. Higuchi, S. Avedisov, O. Altintas, and M. Sepulcre, "End-to-end V2X latency modeling and analysis in 5G networks," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 4, pp. 5094–5109, 2023.

[19] H.-T. Nguyen, N.-Q. Nguyen, V. H. Le, S. D. Tran, and L. D. Hoang, "Proposed data roaming network simulation environment for 5G network," in *Proceedings of the 2025 10th International Conference on Cloud Computing and Internet of Things*, ser. CCIOT '25. Association for Computing Machinery, 2025, p. 138–147.