

MOBILE SECURITY GURU

The Enterprise Mobile Security Gap Report

Why Your Stack Produces Signals — Not Security.
And What a Verifiable Posture Actually Requires.

A field intelligence report for CISOs, IT directors, and security leaders who need board-defensible answers — not vendor dashboards.

3

CONTROL LAYERS

Signal · Enforcement · Verification

\$4.88M

AVG BREACH COST

IBM Cost of a Data Breach 2024

0%

STACK VISIBILITY

into the baseband layer

William Haynes | Mobile Security Guru | 2026 Edition | mobilesecurityguru.com

How to Use This Report

This report is designed to be read in one sitting — approximately 45 minutes — and to produce three outputs by the time you reach the final page:

A documented verification inventory

A clear picture of what your current stack signals, what it enforces, and what it cannot verify — across the five layers of the mobile attack surface.

A board-ready posture statement

Language you can use in your next security committee meeting that is accurate, defensible, and avoids the compliance-as-security failure mode.

A prioritized 30-day action list

Three decisions — executable with existing resources — that close the highest-consequence verification gaps first.

"This report does not evaluate specific vendors. It evaluates whether your current stack — regardless of what you have deployed — produces verifiable mobile security outcomes."

Who this is written for:

For CISOs:

You are accountable for the posture. This report gives you the vocabulary, the framework, and the board-ready summary language.

For IT Directors:

You manage the stack. This report tells you what your current tools signal, what they enforce, and what they cannot verify — and what that means operationally.

For CFOs and COOs:

You own the financial and operational risk. Section 06 translates verification gaps into the numbers that belong in your risk register.

For Security Committee Members:

You ask the questions. Section 12 gives you the questions that separate a verifiable mobile security posture from a well-managed signal dashboard.

For High-Risk Individuals:

Executives, board members, journalists, legal counsel, and staff whose devices concentrate disproportionate risk. Section 10 addresses the asymmetric exposure your fleet controls do not cover.

What's Inside

SECTION 00	The Verification Problem Signal, Enforcement, Verification — and where modern stacks now reach
SECTION 01	The Compliance Illusion Why dashboards produce confidence without producing assurance
SECTION 02	The Pocket Attack Surface The four-factor model for quantifying mobile exposure
SECTION 03	The Five-Layer Attack Surface The mobile threat model the integrated stack was not built for
SECTION 04	Stack Coverage Reality What MDM, MTD, EDR-Mobile, and CA actually verify — and where coverage stops
SECTION 05	The Control-Plane Threat Why attackers target the management infrastructure, not the device
SECTION 06	The Financial Exposure Calculator Translating verification gaps into board-level numbers
SECTION 07	The eSIM Identity Risk Provisioning-layer attacks that bypass device-focused security
SECTION 08	The Baseband Blind Spot The architectural gap no integrated stack currently closes
SECTION 09	The Identity Plane Divergence Where CAE and Token Protection reach — and where they don't
SECTION 10	Executive Device Risk Concentration Why C-suite and high-risk individuals require asymmetric controls
SECTION 11	The Quantum Window HNDL exposure for executive communications, M&A, and legal material
SECTION 12	The Board Readiness Scorecard Six questions. Your answers determine your posture.
SECTION 12B	Verification Metrics That Matter Ten metrics that distinguish policy-only from runtime-verified
SECTION 13	The 30-Day Action Plan Three decisions that close the highest-consequence verification gaps
SECTION 13B	The Board One-Pager What a defensible mobile posture statement looks like

The Verification Problem

"Mobile security programs fail at the board level for one recurring reason: they report control presence as if it were proof of runtime assurance. Enrollment is not verification. Compliance is not compromise status. Conditional access is not total session visibility. And no mainstream enterprise stack fully collapses carrier, eSIM, or baseband risk into a single measurable control plane."

Modern mobile security stacks have become more capable. Integrated deployments now combine MDM policy enforcement, MTD runtime telemetry, EDR-Mobile behavioral signal, app protection policies for both managed and unenrolled access, and Conditional Access tied to risk-based and session-aware controls. That is real progress. It is also the source of the board-level confusion this report addresses.

The mobile security category operates across three distinct control layers. The industry now builds the first two well. It does not yet build the third in a way that survives board, insurer, or regulatory scrutiny. Until that distinction is clear, every gap analysis and risk scorecard continues to mis-diagnose the problem.

The three control layers

Signal

What the stack detects. Device compliance state, threat intelligence matches, risk scores, behavioral anomaly flags. This is where Mobile Device Management (MDM), Mobile Threat Defense (MTD), and Endpoint Detection & Response for mobile (EDR-Mobile) operate. The industry invests heavily here.

Enforcement

What the stack does with the signal. Conditional Access policies, quarantine actions, access revocation, MFA challenges. This is where Conditional Access (CA) and identity providers (IAM) operate. The industry builds this competently.

Verification

What the stack can prove. Continuous, cross-layer confirmation of device, identity, session, and access state — with evidence sufficient to defend the posture to a board, an insurer, or a regulator. Continuous Access Evaluation and Token Protection close part of this gap for supported services and managed device classes, but coverage is uneven across platforms, apps, and session types. No mainstream stack today produces uniform verification across the full mobile attack surface.

"Modern mobile security stacks produce signals and enforce access. What they do not yet produce, uniformly, is verifiable cross-layer posture evidence. The result is partial visibility presented as assurance."

Why this distinction matters at the board level

A CISO who reports a 94% compliance rate and a clean MTD detection feed is reporting Signal. A CISO who reports that Conditional Access blocked 17 high-risk sign-in attempts last quarter is reporting Enforcement. Neither of those is a Verification statement. Verification sounds different:

A verifiable posture statement

"Our mobile-accessible identity surface is X. Our stack produces deterministic signal on Y% of it, reactive enforcement on Z% of it, and has documented blind spots in the following areas — which we have compensated for through the following controls, with the following residual risk."

That statement is defensible. The first two are not — because they describe tool output, not security outcome. The rest of this report maps the distance between those two things across the five layers of the mobile attack surface.

Two frameworks, one matrix

This report uses two frameworks that operate at different axes, and a careful reader will encounter both. The distinction matters, because conflating them produces exactly the kind of imprecise posture statement this report exists to address.

Three control layers — what the stack does

Signal, Enforcement, Verification. These describe the *function* a control performs. Every mobile security capability — whether MDM policy, MTD telemetry, or Conditional Access enforcement — operates as one or more of these three functions.

Five attack surface layers — where attacks happen


Device, Application, Network, Identity & Session, Baseband / Carrier. These describe the *architecture* at which attacks operate. Every mobile-originated attack lands at one or more of these five layers.

"The two frameworks are orthogonal, not competing. A defensible posture statement answers, for each of the five attack surface layers: what does the stack signal on, what does it enforce on, and what does it verify? That is the 3 × 5 matrix this report is built around."

Section 04 makes this explicit by populating the matrix directly: five attack surface layers as rows, the stack-capability columns mapping to Signal, Enforcement, and Verification. The result is a single page that names — for any given deployment — exactly where the stack reaches and where it stops.

The evidence map your board actually needs

What follows in this report is essentially one diagram, expanded over thirty pages. Five categories of evidence flow into a single board posture statement. The first four contribute. The fifth — residual blind spots — is the category most posture statements omit, which is the failure this report is written to address.



The first four rows are evidence categories any modern stack can contribute to. The fifth row — residual blind spots — is the category most posture statements omit. A defensible board posture statement names all five.

The Compliance Illusion

There is a question most CISOs cannot answer cleanly when asked by their board: *"How do you know your mobile environment is secure?"*

The answers most are trained to give — enrollment rates, compliance percentages, policy status, threat detection counts from an MTD console — are accurate. They describe Signal and Enforcement. They do not answer the question being asked.

"Compliance documentation describes what you have enforced. It does not describe what remains invisible. The gap between those two is where mobile breaches happen — and where board accountability lives."

What compliance dashboards actually measure

Mobile security platforms — whether standalone MDM, MDM integrated with Mobile Threat Defense, or a full MDM + MTD + EDR-Mobile + Conditional Access stack — were designed to enforce configuration policy and detect a defined set of threat signals at scale. Screen lock, encryption, OS version minimums, approved application lists, known-malicious app detection, unsafe Wi-Fi categorization, web filtering, risk scoring, conditional access enforcement. These are legitimate and necessary controls.

When your dashboard reports 94% compliance and a clean threat detection feed, it means 94% of enrolled devices have those policies applied and the signals your detection stack is watching for have not fired. **That is the complete scope of the claim.**

Five things dashboards do not cover — even with MTD, EDR-Mobile, and CA integrated:

- **Devices that never enrolled in the first place.** Executive-owned phones used for corporate email via OAuth. Contractor devices with SSO access. Former-employee devices with cached credentials. Coverage statistics are calculated on the denominator of what you know about.
- **The baseband and carrier layer of every device in your fleet.** MTD agents run in userspace. They cannot see the cellular modem, the carrier provisioning state, or the eSIM profiles installed on the device. Every device you manage has a second operating system that your stack does not observe.
- **The phone number tied to each device.** SIM swaps, port-outs, and eSIM hijacks transfer the authentication channel — the phone number itself — to an attacker. The device remains compliant. The stack still trusts it. The number no longer belongs to your employee.
- **Threat signals that post-date the vendor's detection model.** MTD and EDR-Mobile detect what their research teams have characterized. Zero-days, novel phishing infrastructure, and new classes of in-app behavior are invisible until vendor telemetry catches up. Your dashboard reports "no threats detected" in exactly the window where you are most exposed.

- **What the dashboard decides not to show you.** Detection thresholds, suppression rules, and risk-scoring weights are set by the vendor. Events below threshold do not surface. The absence of an alert is not the same as the absence of an event.

The test that reveals the gap:

Ask your mobile security stack — MDM, MTD, EDR-Mobile, and CA combined — three questions with documented answers:

1. How many mobile devices accessed corporate data in the last 30 days that were not enrolled in our MDM? Show me the reconciliation against our identity provider's sign-in logs.
2. If an executive's phone number were ported to an attacker-controlled carrier account right now, which alert in our stack would fire, and how long after the port-out would it fire?
3. Show me the last three detection model updates from our MTD vendor, the specific threats they added coverage for, and the date range during which those threats were undetectable in our environment.

If your vendor cannot produce answers with timestamps, the gap is not a product failure. It is a scope boundary your board needs to understand.

The distinction matters at the board level because it changes what questions get asked. A board that believes compliance equals security will not ask about unenrolled device inventory, carrier-layer exposure, executive port-out risk, detection latency, or suppression logic. Your job is to bring those questions into the room — before a breach brings them in for you.

The Pocket Attack Surface

The board metric most enterprise mobile security programs lack is a single, defensible number that describes the size of their mobile-originated risk exposure. Compliance rate does not describe exposure. Enrolled device count does not describe exposure. The right metric — the Pocket Attack Surface, or PAS — is the product of four factors.

$$\text{"PAS = Devices} \times \text{Identities} \times \text{Sessions} \times \text{Accessible Systems"}$$

The four factors

Devices

Every mobile endpoint with active access to corporate resources. Enrolled devices, unmanaged BYOD devices, contractor devices, executive-personal devices, and former-employee devices with cached credentials. The denominator is not what your MDM shows — it is what your identity provider's sign-in logs reveal.

Identities

Every authenticated identity on those devices. A single phone may carry the CEO's corporate identity, a personal identity used for work (shadow IT), and delegated admin accounts. Each identity multiplies the exposure of the device it runs on.

Sessions

Every persistent session token, OAuth grant, refresh token, and saved credential active on those identities. Session tokens survive password rotation. Tokens stolen from a mobile device authenticate the attacker long after the device is reset. Session count is rarely inventoried and almost never governed.

Accessible Systems

The corporate resources each session can reach: email, file storage, CRM, ERP, source control, finance systems, customer databases. Lateral access compounds the blast radius of any single compromised session.

Why the four-factor formula matters

A two-factor model (Devices × Systems) tells you how many endpoints can reach how many resources. It describes infrastructure topology. The four-factor model tells you how many distinct authentication paths exist between an attacker and your most sensitive systems if any one device is compromised. It describes attack surface.

For most mid-market organizations, the four-factor PAS is an order of magnitude larger than the two-factor version. For enterprises with BYOD, contractor access, or executive-personal device usage, the gap is typically two orders of magnitude. This is the number that belongs in the risk register — not the compliance percentage.

A worked example

A 200-person organization with 200 enrolled devices, 50 unmanaged BYOD devices, an average of 1.8 identities per device, 4 active sessions per identity, and access to 30 corporate systems has a PAS of approximately 54,000 authentication paths. The MDM dashboard shows 200 devices. The risk register typically cites the 200. The exposure is 270× larger.

The Five-Layer Attack Surface

The mobile attack surface has five distinct architectural layers — the *where* axis of the matrix introduced in Section 00. These are not the same as the three control layers (Signal, Enforcement, Verification), which describe the *what* axis. Each architectural layer has its own threat model, its own category of vendor tooling, and its own verification requirements. The error most enterprise programs make is assuming that covering one or two of these layers well produces posture across the other three or four. It does not.

LAYER 01

Device Layer

The operating system, configuration state, and installed application inventory. This is where MDM operates and where the industry has the most mature tooling.

What the stack signals:

MDM enforces configuration. MTD detects compromise indicators. EDR-Mobile provides behavioral telemetry.

What remains unverified:

No current stack continuously verifies device integrity between signals. Detection gaps exist between vendor telemetry updates and in the window before a novel threat is characterized.

LAYER 02

Application Layer

The behavior of applications once installed — including legitimate applications used for illegitimate purposes.

What the stack signals:

MTD and EDR-Mobile detect known-malicious apps, unsafe SDKs, and pattern-based anomalies. App reputation services flag risky publishers.

What remains unverified:

Detection is signal-based and retrospective. Novel data exfiltration via legitimate apps — the OAuth token theft pattern, the cloud-storage uploader pattern — is invisible to categorical detection. No stack guarantees benign application behavior.

LAYER 03

Network Layer

The connections the device makes, the traffic those connections carry, and the parties on the other end.

What the stack signals:

MTD detects unsafe Wi-Fi, MITM patterns, and malicious DNS. Web filtering blocks known-malicious domains. Some stacks support VPN enforcement.

What remains unverified:

Network visibility ends at the device's network interface. Encrypted traffic to legitimate destinations is opaque. Split-tunnel VPN configurations and cellular data paths bypass most enforcement. No stack certifies that all traffic is clean.

LAYER 04

Identity and Session Layer

The authentication, authorization, and session state that determines what the device and its user can access.

**What the stack
signals:**

Conditional Access enforces policy on sign-in. Risk-based authentication challenges anomalous attempts. Identity providers log sign-in events.

**What remains
unverified:**

Identity posture and device posture have diverged. A compromised session token authenticates the attacker independent of device state. Phishing-resistant MFA mitigates sign-in attacks but does not address post-authentication token theft. Identity is now a parallel control plane, not governed by device security.

LAYER 05

Baseband and Carrier Layer

The cellular modem, the carrier's provisioning infrastructure, and the phone number itself as an authentication channel.

**What the stack
signals:**

Effectively none. This layer operates below the application processor. MTD agents cannot observe it. MDM cannot configure it.

**What remains
unverified:**

This is a structural blind spot. No enterprise product currently provides visibility here. It is not a coverage gap that will be closed by vendor roadmap — it is a permanent constraint of current mobile architecture.

***"Coverage does not equal control. Detection does not equal prevention.
Signals do not equal assurance."***

Stack Coverage Reality

This section populates the 3 × 5 matrix introduced in Section 00. The five *attack surface layers* from Section 03 are the rows. The columns map to the three *control layers* — Signal, Enforcement, and Verification — across the recognizable stack deployment patterns most enterprises run today. Each cell names what that stack pattern actually does at that architectural layer.

The honest framing is not "how many layers does your stack cover" — it is "which layers does your stack signal on, which does it enforce on, and which does it verify."

LAYER	MDM ONLY	MDM + MTD	FULL STACK	VERIFIED ?
Device	Policy	Policy + Signal	Policy + Signal + Enforce	Partial
Application	None	Partial signal	Partial signal + Enforce	Partial
Network	None	Partial signal	Partial signal + Enforce	Limited
Identity / Session	None	None	CAE + Token Protection*	Partial
Baseband / Carrier	None	None	None	None

What this table means

Policy means the stack can enforce configuration (encryption on, OS version, approved apps). **Signal** means the stack can detect threat indicators at that layer. **Enforce** means the stack can take automated access action on detected signals. **Verified** means the stack can continuously prove posture at that layer with evidence sufficient for board, insurer, or regulatory defense. *CAE = Continuous Access Evaluation. Token Protection is currently in preview on Apple platforms and limited to supported services such as Exchange, Teams, and SharePoint Online.

"Modern integrated stacks have moved the Identity and Session layer from "unverified" to "partially verified" — meaningful progress. Verification across the full mobile attack surface remains uneven. Carrier-side identity events, unmanaged traffic outside protected apps, and baseband execution are the layers no current stack meaningfully covers."

What integrated signal coverage actually produces

A fully integrated stack — MDM + MTD + EDR-Mobile + Conditional Access with Continuous Access Evaluation enabled — produces broad signal coverage, meaningful enforcement at the identity and session layer, and partial verification on supported services. That is real progress over MDM-only deployments. It is not full verification across

the mobile attack surface.

The gap between Enforcement and Verification is where every unresolved mobile incident in 2025 originated: unenrolled devices accessing data via OAuth grants outside protected apps, post-authentication token theft from sessions on services that do not yet support Token Protection, baseband-layer attacks on executive devices, and carrier-layer identity hijacks via eSIM provisioning. None of these produce a signal in the stack that detected nothing.

The Control-Plane Threat

The most consequential mobile security shift of the last 24 months has nothing to do with the devices. It has to do with the infrastructure that manages them.

The strategic pivot attackers have already made

Enterprise mobile management platforms were designed as trusted infrastructure. The assumption was that the management plane itself was outside the threat model — that attackers would target devices, not the systems that managed them. That assumption no longer holds.

CISA's Known Exploited Vulnerabilities catalog now contains repeated additions from the enterprise mobility management and endpoint management category. The exploitation pattern is consistent: attackers compromise the management plane, inherit trust on every device it manages, and operate within the perimeter without ever attacking a single endpoint directly.

Why this changes the threat model

A compromised management plane is fundamentally different from a compromised device. One compromised device exposes one device's worth of data. One compromised management plane exposes every device under management, every policy configuration, every enrollment credential, and every piece of telemetry flowing to the vendor's cloud. The blast radius is the entire fleet.

What the attack sequence looks like

- **Step 1.** Attacker exploits a vulnerability in the management platform's administrative interface — typically web-facing, often protected by single-factor authentication.
- **Step 2.** Attacker establishes administrative access, inheriting trust across every enrolled device. From the fleet's perspective, the attacker is the legitimate administrator.
- **Step 3.** Attacker modifies compliance policies, pushes configuration profiles, installs applications, or extracts telemetry — any action the legitimate administrator could take. The fleet obeys, because the management plane tells it to.
- **Step 4.** Attacker exfiltrates data, harvests credentials, or maintains persistence through policy artifacts that survive vendor patches and incident response. The fleet remains compliant throughout.

Five questions your MDM governance program must answer

1. Who has administrative access to the management platform, and is that access protected by phishing-resistant MFA?
2. When was the last time the management platform's web-facing interfaces were tested against the current CISA KEV catalog?

3. If the management platform were compromised today, what is our detection window? What telemetry would surface the compromise, and how long would it take to fire?
4. What is the recovery playbook if the management platform must be rebuilt from scratch — and has it been exercised?
5. Which vendor-side telemetry do we have contractual visibility into, and which do we rely on the vendor to monitor on our behalf?

The strategic framing

The control plane is now a higher-value target than any individual device. Every mobile security program needs a control-plane threat model. Most do not have one.

The Financial Exposure Calculator

Verification gaps translate into numbers that belong in the risk register. This section provides the framework for calculating your organization's mobile exposure in terms the CFO, the board, and your cyber insurance underwriter will recognize.



The honest cost framing

The legacy assumption was that a mobile breach would produce no alert and no forensic record. In 2026, that is rarely true — detection may occur. The problem is where and when:

- **Detection often happens downstream.** Most mobile-originated incidents surface in identity systems, cloud platforms, or customer-facing fraud monitoring — not at the mobile layer where they originated. By the time detection fires, lateral movement has already occurred.
- **Detection rarely happens across all devices.** Unenrolled and partially managed devices produce no telemetry. Detection statistics describe the instrumented fraction of the fleet, not the whole fleet.
- **Detection rarely happens across all attack types.** Baseband attacks, eSIM hijacks, and session token theft post-authentication produce no signal in endpoint-focused tooling.
- **Forensic completeness is often the limiting factor.** Detection without sufficient forensic trail produces an incident that cannot be scoped, communicated to regulators, or defended in litigation.

The three cost categories boards underestimate

Regulatory and contractual exposure

Data breach notification obligations, regulatory fines (HIPAA, GDPR, state-level breach notification laws, SEC cyber disclosure rules), and contractual penalties in enterprise customer agreements. The notification clock starts at discovery — not at containment — and mobile-originated breaches are notoriously hard to scope.

Cyber insurance cost and coverage degradation

Carriers now ask specific questions about mobile security posture at renewal. Inability to produce verification evidence increases premiums, reduces coverage limits, and introduces exclusion clauses for mobile-originated incidents. This is becoming an annual structural cost, not a one-time incident cost.

Customer concentration risk

Enterprise customers increasingly require mobile security attestations as part of vendor onboarding. A security incident — or even a failure to produce current posture documentation — can trigger customer audit rights, contract renegotiation, or churn. For organizations with top-10 customer concentration above 40%, this is the largest single financial exposure.

The CFO-facing framing is simple: the marginal cost of closing verification gaps is bounded and known. The marginal cost of a mobile-originated breach is unbounded and increasing. The insurance market is already pricing that asymmetry.

The eSIM Identity Risk

The phone number tied to a mobile device is not a trivial attribute. It is an authentication channel, a recovery mechanism, a password reset vector, and in many organizations, the de facto second factor for high-privilege accounts. eSIM provisioning has made that channel more portable, more attackable, and less observable than it was five years ago.

The provisioning attack pattern

Traditional SIM swap attacks required physical access to a carrier retail store or social-engineered access to a carrier call center. eSIM provisioning has streamlined both vectors. An attacker who obtains sufficient account information — often harvested from previous data breaches — can provision a new eSIM to an attacker-controlled device and, in some carrier workflows, have it active within minutes.

Once active, the attacker's device receives the victim's calls, SMS messages, and any SMS-based authentication codes. The victim's device shows no signal of the hijack until service stops — which may be hours or days later, depending on the attacker's operational pattern.

Why this bypasses device-focused security

Most controls in your MDM, MTD, and EDR-Mobile stack operate on the device itself. eSIM provisioning attacks operate at the carrier, outside the device. The device remains compliant, trusted, and fully enrolled throughout the attack. Device-focused telemetry produces no signal because the attack never traverses the layer it monitors. Verification at this layer requires controls outside the device-focused vendor stack: carrier-account hardening, port-freeze policies, enterprise-managed eSIM profiles where available, and elimination of SMS as an authentication factor for any privileged account.

The documented legal exposure

In March 2025, a California arbitration panel awarded approximately \$33 million against a major US carrier in a case arising from a SIM-swap account takeover, finding that the carrier's identity verification procedures had been inadequate. The case is not a clean eSIM-specific precedent — the underlying attack predates the most aggressive eSIM-provisioning patterns documented since — but it establishes two things relevant to enterprise programs: carrier liability for inadequate identity verification at the provisioning layer is now a matter of decided arbitration record, and damages from carrier-side identity failures can reach material dollar amounts. The eSIM provisioning workflows documented above operate in the same regulatory and legal envelope.

What enterprise mobile programs must ask

1. Which employees' phone numbers, if hijacked, would compromise the most sensitive corporate accounts? (This is the executive port-out risk concentration question.)

2. What carrier-side controls — port freeze, SIM PIN, account-level MFA on the carrier account — are in place for those specific numbers?
3. What telemetry would alert us to an eSIM provisioning event on an executive device? How long after the event would the alert fire?
4. Have we removed SMS-based MFA from all high-privilege accounts and replaced it with phishing-resistant alternatives?

The strategic framing

eSIM compromise operates at the identity layer without traversing the device layer. It is invisible to every device-focused control you have deployed, and it renders SMS-based MFA worse than useless — because the SMS still arrives, it just arrives on the attacker's device.

The Baseband Blind Spot

Every smartphone is two computers in one case. The application processor — the one running iOS or Android — is the one your MDM configures, your MTD monitors, and your security team thinks about. The baseband processor — the one that handles cellular radio communication — is the one none of your tooling can see.

What the baseband actually does

The baseband is a separate system-on-chip running its own real-time operating system, independent of the application processor. It manages every cellular network interaction: registration with carrier towers, signaling protocol exchanges, call and SMS handling, data connections, and the increasingly complex semantic messaging of 5G Radio Resource Control (RRC) protocols. It is the first to touch every bit of cellular traffic entering or leaving the device.

What the baseband means for enterprise security

- **The baseband is not covered by MTD.** MTD agents run in userspace on the application processor. They have no access to the baseband's state, memory, or protocol handling.
- **The baseband is not covered by MDM.** MDM configures policies on the application processor. It cannot set baseband configuration, cannot inspect baseband firmware, and cannot observe baseband behavior.
- **The baseband is not covered by EDR-Mobile.** Behavioral telemetry is collected at the operating system layer. Baseband operations occur below that layer and are invisible to it.
- **Baseband vulnerabilities are researched, documented, and exploitable.** Academic research has demonstrated remote exploitation of baseband stacks across multiple chipset vendors. CVE-level disclosures continue to appear. This is not a theoretical threat.
- **Baseband patching is carrier-gated and operator-dependent.** Even when vendors issue baseband firmware updates, deployment depends on carrier approval and device-by-device update cycles. Enterprise security teams have no control over this cadence.

The architectural reality

This is not a gap that vendors will solve in the near term. It is a permanent constraint of current mobile architecture. Every smartphone your organization deploys has a second operating system running cellular protocol stacks that no tool in your stack can observe. The honest posture statement acknowledges this rather than pretending it does not exist.

What can be done

Baseband risk cannot be eliminated with current tooling. It can be managed through three compensating controls:

Carrier-agnostic connectivity strategy

Enterprise eSIM profiles managed by the organization rather than individual consumer carriers. This moves the provisioning and management plane under enterprise control and shortens the patch-cycle dependency.

Device selection with known baseband characteristics

Executive device standards should specify chipset vendor, baseband firmware update cadence, and disclosed vulnerability history. Not all smartphones are equivalent from a baseband risk perspective.

Assumption of baseband compromise in the threat model

High-risk communications should not assume the baseband is trustworthy. End-to-end encrypted applications, out-of-band verification for high-value transactions, and compartmentalization of sensitive workflows to specific devices all mitigate the consequence of baseband compromise.

The Identity Plane Divergence

For most of the last decade, enterprise security treated device posture and identity posture as tightly coupled. A compliant device produced a trusted session. A compromised device revoked that trust. That coupling has broken. Identity is now a parallel control plane, and its posture is not governed by device state.

How the divergence happened

Three shifts, all accelerating: the move from session cookies to long-lived refresh tokens, the proliferation of OAuth grants across SaaS applications, and a transitional period in identity-provider architecture where sign-in assessment has moved ahead of session-state assessment for most workloads. Conditional Access historically evaluated risk at sign-in. That has changed: Microsoft now documents Continuous Access Evaluation for selected services — initially Exchange Online, Teams, and SharePoint Online — and Token Protection for supported native apps and managed Apple devices. These are real advances. They do not yet provide universal post-authentication evaluation across all mobile workloads, all platforms, or all session types.

What this means operationally

- **A stolen session token is valid even after the device is remediated, on services that do not support continuous evaluation.** Wiping the device does not invalidate the token on the attacker's infrastructure. Explicit token revocation is required — and on services without CAE, it is not automatic.
- **Phishing-resistant MFA protects sign-in, not all sessions.** FIDO2 and passkeys prevent credential phishing. Token Protection extends that protection to selected supported scenarios. Outside those scenarios — native mail clients on some platforms, third-party apps, services that have not yet adopted Token Protection — the session token remains a bearer credential.
- **Risk-based authentication operates on sign-in signals.** Impossible travel, device fingerprint changes, and anomalous behavior trigger re-authentication — but the baseline for "normal" is established from prior sessions, which means an attacker who operates within normal parameters avoids triggering it.
- **Token theft from device memory is invisible to device-focused tooling.** MTD and EDR-Mobile detect signals on the device. An attacker who extracts a refresh token from device memory and uses it from attacker-controlled infrastructure produces no signal on the original device — and unless the target service supports CAE or Token Protection, the identity plane has no signal either.

The strategic framing

Identity posture and device posture are increasingly governed by separate control planes. Continuous Access Evaluation and Token Protection are closing important gaps for supported services and managed device classes — but coverage is uneven across platforms, apps, and session types. A verifiable mobile security posture must address both planes and must specifically account for the services and scenarios where post-authentication evaluation does not yet apply.

Three questions this layer introduces

1. What is the maximum age of any refresh token or OAuth grant accessible from a mobile device in our environment? What controls limit that lifetime?
2. If a device is declared compromised at 10:00 AM, by what time are all tokens and sessions issued to that device provably revoked? Which services support automatic revocation via Continuous Access Evaluation, and which require manual action?
3. What percentage of our mobile-accessible identities and supported services are covered by Continuous Access Evaluation and Token Protection — and what is our documented residual risk on the remainder?

Executive Device Risk Concentration

Enterprise mobile security is often framed as a fleet problem. It is not. It is a concentration-of-risk problem. Ten devices in the C-suite, the legal department, the finance function, and the board carry more aggregate risk than ten thousand devices in the general workforce.

"Fleet controls applied uniformly do not address asymmetric risk. The executive device posture problem requires a different answer."

Why executive devices carry asymmetric risk

- **Account access concentration.** Executive identities typically carry administrative rights to financial systems, HR platforms, board portal access, legal privilege repositories, and M&A data rooms. A single compromised identity produces disproportionate blast radius.
- **Communication sensitivity.** Board communications, M&A discussions, litigation hold material, and strategic planning documents disproportionately flow through executive devices. The data at risk is categorically different from fleet data.
- **Phone number targeting value.** Executive phone numbers are the highest-value targets for eSIM provisioning attacks and SIM swap attacks because they authenticate the most sensitive accounts.
- **Social engineering attack surface.** Executive staff — executive assistants, personal security details, family members — are targeted as vectors to the executive device. Fleet awareness training does not reach them.
- **Personal-use overlap.** Executive devices are more likely to be personal devices used for corporate work, or to have personal applications installed on corporate devices. The attack surface expands.

What an executive device standard should require

Hardware selection baseline

Devices with documented baseband vulnerability track records, vendor-supported firmware update cadence of 60 days or better, and hardware-backed key storage.

Separate device for sensitive workflows

A dedicated device for board material, M&A, and litigation hold workflows, physically separated from the primary personal-use device. The extra device cost is trivial relative to the risk asymmetry.

Carrier-layer controls

Port freeze enabled, SIM PIN set, account-level MFA on the carrier account, and where possible, enterprise-managed eSIM profiles rather than consumer carrier profiles.

Identity plane isolation

Executive identities governed by stricter Conditional Access policies, shorter token lifetimes, continuous access evaluation on sensitive operations, and phishing-resistant MFA as the only permitted second factor.

Verification cadence

Quarterly structured assessment of executive device posture, not the annual or biannual cadence appropriate for fleet devices. The threat model moves faster at the top.

The framing that closes budget conversations

Enterprise mobile security is not a fleet problem. It is a concentration-of-risk problem. Budget allocated to fleet-uniform controls produces diminishing returns past a certain coverage point. Budget allocated to executive and high-risk individual posture produces compounding returns because it reduces the probability and consequence of the highest-severity incidents.

The Quantum Window

Harvest Now, Decrypt Later (HNDL) is not a present-day decryption capability. It is a present-day harvesting pattern with a future decryption payoff. Nation-state and sophisticated financial adversaries are presumed to be collecting encrypted traffic of long-term strategic value today, on the assumption that cryptographically relevant quantum computers will eventually make that traffic readable.

Why this matters for enterprise mobile programs

Most enterprise data has a short shelf-life. A breach notification decision in 2026 is made from 2026 data. But three categories of mobile-carried data have shelf-life that extends into the decryption window:

Executive communications

Strategic plans, board material, competitive intelligence, and succession planning often remain sensitive for 5–15 years. Communications captured in 2026 may be decrypted in 2031–2036 and remain commercially damaging.

Mergers, acquisitions, and due diligence

M&A material has a narrower sensitivity window — typically 1–3 years — but its concentration value is extreme. A decrypted M&A thread can produce securities-fraud-level consequences or unwind completed transactions.

Legal and litigation material

Attorney-client privileged communications, litigation strategy, and internal investigations material remain sensitive for the duration of the legal matter and beyond. Litigation hold periods commonly extend 7–10 years.

The migration reality

NIST finalized the first three post-quantum cryptography standards — ML-KEM (FIPS 203), ML-DSA (FIPS 204), and SLH-DSA (FIPS 205) — in August 2024. The migration path for enterprise mobile environments is clear: identify cryptographic dependencies, prioritize high-sensitivity long-lived data, and coordinate with vendors on PQC-ready libraries and protocol support. Most enterprise mobile programs have not begun this work.

The pragmatic framing

HNDL is not a current breach risk. It is a long-term exposure multiplier for already-compromised communications. Organizations with concentrated long-lived sensitive content — executive comms, M&A workflows, litigation material — should be tracking PQC migration now. Organizations without that concentration can reasonably defer until vendor support matures.

The Board Readiness Scorecard

The questions below are designed for the security committee conversation. They are not pass/fail. They are diagnostic. Each question maps to a verification gap. Your answers produce the posture statement.

QUESTION 01

Can we produce a count of mobile devices accessing corporate data that are not enrolled in our MDM, reconciled against identity provider sign-in logs?

The denominator problem. If the answer is no, all coverage percentages are calculated on an unknown fraction of the surface.

QUESTION 02

If an executive's phone number were ported to an attacker-controlled carrier account right now, what alert in our stack would fire — and how long after the port-out?

The eSIM and carrier-layer blind spot. The answer should include timestamps and the specific telemetry source.

QUESTION 03

What percentage of our mobile-accessible identities have continuous access evaluation — where identity posture is re-assessed on every sensitive operation, not just at sign-in?

The identity plane divergence. If the answer is close to zero, token-theft attacks succeed invisibly.

QUESTION 04

What is our documented recovery playbook if our mobile management platform is compromised? When was it last exercised?

The control-plane threat. Most organizations do not have this document. Most of those that do have never exercised it.

QUESTION 05

Do our executive and high-risk individual devices have a different posture than our fleet devices — carrier-layer controls, separate devices for sensitive workflows, identity plane isolation?

The asymmetric risk question. Uniform fleet controls systematically underinvest in highest-consequence exposure.

QUESTION 06

Can we produce, on one page, a verification statement describing what our stack signals, what it enforces, and what it cannot verify — with evidence sufficient for a board or cyber insurer?

The verification layer. If the answer is no, you are reporting Signal and Enforcement while being asked about Security.

How to read your answers

If you can answer all six questions with documented evidence and timestamps, your posture is mature and your next quarter's work is optimization. If you can answer three to five, you have a clear remediation path. If you can answer fewer than three, you are reporting Signal as Security — and the first task is scoping the verification inventory. That is what a structured diagnostic produces.

Verification Metrics That Matter

The compliance percentages and detection counts most mobile programs report describe tool output. They do not describe verification quality. The metrics below are the ones that distinguish a policy-only program from a runtime-verified program. A mature mobile security posture reports against these — not as vanity numbers, but as evidence-quality indicators that map directly to board, insurer, and regulator questions.

METRIC	WHY IT MATTERS	DATA SOURCE
Unenrolled mobile access %	Devices accessing corporate data without appearing in MDM enrollment. The denominator problem made visible.	Identity provider sign-in logs, reconciled against MDM inventory
App protection coverage %	Mobile access governed by app-layer protection on enrolled and unenrolled devices.	MAM / app protection policy coverage reports
Runtime threat telemetry %	Devices actively sending behavioral and threat signal to the SOC.	MTD / EDR-Mobile dashboards
Mean time to access restriction	Elapsed time from threat detection to enforced access change. The operational test of risk-based access.	Correlation of MTD events with CA policy actions
Privileged users on phishing-resistant MFA %	Privileged identities using FIDO2, passkeys, or hardware tokens — not SMS.	Identity provider authentication method report
Privileged users with SMS fallback %	Privileged identities where SMS remains an authentication or recovery path. The cleanest board-visible exposure metric.	Identity provider auth method + recovery channel report
Token Protection coverage %	Mobile sessions on supported services and platforms covered by Token Protection.	Conditional Access policy coverage report
Web / network protection coverage %	Mobile devices with active web filtering and unsafe-network detection enabled.	MTD / EDR-Mobile platform-specific feature reports
Carrier / eSIM hardening coverage %	Privileged users with documented carrier-account controls (port freeze, SIM PIN, account MFA).	Telecom operations records, identity-provider recovery policy
Documented residual-risk acceptance %	Mobile users or scenarios where a verification gap is explicitly accepted with board-level documentation.	Risk register, board minutes

How to use this list

These ten metrics are the verification inventory. A mobile program that can produce current values for all ten — with timestamps and data lineage — is a runtime-verified program. A program that produces values for three to five is in transition. A program that cannot produce most of them is reporting policy presence as security, which is the failure mode this report exists to address.

The 30-Day Action Plan

Three decisions that close the highest-consequence verification gaps first. These actions can often begin with existing logs, policies, and licenses, but some organizations will need new tooling or services to fully operationalize them. All three require documented evidence before they are considered complete.

01

Produce the unenrolled device inventory.

Reconcile your MDM enrollment list against your identity provider's sign-in logs for the last 30 days. Identify every mobile device that accessed corporate data without appearing in MDM enrollment. For each, determine ownership, access justification, and remediation path (enroll, revoke, or accept with documented compensating controls). Target completion: 10 business days.

Board exposure: The denominator problem is the single largest source of false posture assurance. Until it is quantified, every coverage metric is unreliable.

02

Build the executive device baseline.

Identify your top 10–20 highest-risk individuals by account privilege, data access, and communication sensitivity. For each, document: device model and baseband firmware status, carrier-layer controls (port freeze, SIM PIN, carrier account MFA), MFA configuration (phishing-resistant only), and separation of personal vs. sensitive workflows. Where gaps exist, remediate. Target completion: 20 business days.

Board exposure: Asymmetric risk concentration means the highest-severity incidents originate from this population. Uniform fleet controls will not address them.

03

Write the one-page verification statement.

Draft the board-ready posture statement using the format in Section 00: what the stack signals, what it enforces, what it cannot verify, what compensating controls address each verification gap, and what residual risk remains. Review with legal counsel and cyber insurance broker. Target completion: 30 business days.

Board exposure: The one-page verification statement is the artifact that distinguishes a defensible mobile security program from a well-managed signal dashboard. Most programs do not produce it.

What this plan prioritizes

These three actions can often begin with existing logs, policies, and licenses — though some organizations will still need new product or services spend to operationalize them. What the plan deliberately does not lead with is additional MTD, EDR-Mobile, or CA tooling. Most organizations do not have a coverage problem; they have a verification problem. New tooling can add signal. It does not by itself close the verification gap. The 30-day plan addresses verification first; tooling decisions follow from what the verification work reveals, not precede it.

The Board One-Pager

The artifact most mobile security programs do not produce, and that this report exists to help generate. The structure below is the format an MRRE engagement delivers — and the format an internal team can produce themselves with the right inputs. It is what a defensible mobile posture statement looks like.

VERIFIED TODAY

Management can verify policy enforcement on X% of enrolled mobile devices, runtime threat telemetry on Y% of corporate or BYOD mobile access paths, and risk-based access controls on Z% of mobile sessions using protected apps or supported native clients.

VERIFIED WITH LIMITATIONS

Some mobile protections are enabled only for supported apps, supervised devices, or services that support app protection, Continuous Access Evaluation, or Token Protection. Coverage is meaningful but not universal. Specific exclusions are documented.

RESIDUAL BLIND SPOTS

Management cannot currently claim meaningful enterprise visibility into carrier-side number provisioning events, baseband execution, or unmanaged traffic outside protected applications. These risks are either partially controlled or explicitly accepted, with named compensating controls.

EXECUTIVE EXPOSURE

X% of executive or privileged mobile users still rely on SMS fallback, have incomplete phishing-resistant MFA coverage, or lack a hardened mobile recovery path. This is the highest-consequence mobile risk concentration in the environment.

TOP THREE DECISIONS

Retire SMS where feasible for privileged users; extend threat-aware app protection to all supported BYOD access paths; formalize residual-risk acceptance for carrier and baseband exposure and executive-device exceptions in the risk register.

90-DAY OUTCOME

At the end of 90 days, the organization will be able to distinguish policy-only assurance from runtime-validated assurance and present a narrower, more defensible residual mobile-risk statement to the board, the cyber insurance broker, and any regulatory auditor — using the same document.

What makes this defensible

This format names what is verified, what is verified with limitations, what remains unverified, where the highest-consequence concentration sits, what the next moves are, and what success in 90 days looks like. It does not claim universal coverage. It does not pretend the blind spots do not exist. It is the document a board, an insurer, and a regulator can each accept as a current statement of mobile security posture — using the same document for all three audiences.

The Verification Layer

"Mobile Security Guru independently assesses the quality of mobile security evidence across policy, runtime, session, and telecom-adjacent risk — then identifies the residual blind spots that survive even in modern integrated stacks."

Above the stack, not against it

The MRRE is positioned above the vendor stack, not against it. Modern integrated deployments — MDM with policy enforcement, MTD with runtime telemetry, EDR-Mobile with behavioral signal, app protection policies extending to unenrolled access, and Conditional Access with CAE and Token Protection on supported services — produce real and important security outcomes. The verification layer's job is not to replace those investments. It is to assess the quality of the evidence they produce, identify the blind spots that remain even with all of them deployed, and convert that assessment into language a board, an insurer, or a regulator can defend.

Every section of this report has argued the same structural point: even with modern stacks deployed, the gap between control presence and verifiable runtime assurance is where mobile-originated material risk concentrates. The practical question that follows is operational. How does an enterprise program produce evidence-quality verification across the stack it already owns?

Verification is a practice, not a product

Verification cannot be procured as a single SKU. It is the work of reconciling what the stack claims with what the environment actually contains — across devices, identities, sessions, carrier-layer exposure, and control-plane integrity. It requires assessment from a vantage point independent of the vendor whose stack is being evaluated, looking at the evidence rather than the dashboard. Vendors increasingly produce excellent evidence; verification asks how much of that evidence is actually present, current, and complete in this specific environment.

This work can be done internally by a security team with sufficient capacity and independence from the mobile program being assessed. It can be done by qualified external practitioners who specialize in mobile security. It is structurally difficult for the vendors whose stack is being verified to do it themselves, because the scope boundary they report against is the same scope boundary the verification work is examining.

The Mobile Risk Reduction Exercise (MRRE)

The MRRE is Mobile Security Guru's structured diagnostic engagement — the verification work operationalized as a defined scope, defined timeline, and defined deliverables. It is one way to produce verification. Other qualified practitioners can do this work, and internal teams with the right capacity can do it themselves. The MRRE exists because most organizations do not have that internal capacity, and because the engagement needed to be structured around a specific output: a board-ready verification statement.

What the MRRE produces

The verification inventory

A documented reconciliation of enrolled devices vs. identity-authenticated devices, active sessions vs. governed sessions, and accessible systems vs. privilege-justified access. This is the Pocket Attack Surface calculation with real numbers from your environment.

The stack capability map

For each of the five attack layers, what your current stack signals, enforces, and cannot verify — with specific evidence from your deployment rather than vendor marketing. This produces the Section 04 table populated with your actual posture.

The verification-gap register

The specific verification gaps in your current posture, ranked by board-level consequence. For each gap: the compensating control options, the residual risk if no action is taken, and the estimated cost and timeline of closing it.

The one-page board statement

The artifact described in Section 13, Finding 03 — drafted as a deliverable, not as homework. Reviewed with you, revised until it is something you can present to a security committee, a cyber insurance broker, and a regulatory auditor with the same document.

How the engagement runs — four weeks, four outputs

The MRRE is structured as a four-week sequence that aligns with how mobile security evidence actually accumulates inside an organization. Each week produces a defined intermediate output. The final week produces the board-ready artifact.

WEEK 1

Inventory mobile access paths

Reconcile MDM enrollment against identity provider sign-in logs. Separate enrolled, unenrolled, and protected-app-only traffic. Produce the Pocket Attack Surface calculation with real numbers from this environment.

WEEK 2	Validate runtime signals Confirm MTD, EDR-Mobile, and partner telemetry coverage. Map enforcement paths from detection event to access action. Identify supported services covered by Continuous Access Evaluation and Token Protection. Document what each layer actually verifies.
WEEK 3	Review identity and executive posture Phishing-resistant MFA coverage, SMS fallback exposure, carrier-account controls for privileged users, eSIM provisioning controls, session protections. The concentration-of-risk inventory.
WEEK 4	Publish the board one-pager Verified evidence, residual blind spots, executive exception register, top three decisions, 90-day sequence. The artifact described in Section 13B, delivered as a finished document — not as homework.

What the MRRE explicitly does not include

- Vendor product recommendations. MSG does not sell or resell mobile security products and does not have affiliate relationships with any vendor in the category.
- A penetration test or red team exercise. The MRRE is a posture and verification diagnostic, not an offensive security engagement.
- Implementation services. The MRRE produces the verification-gap register and the remediation plan. Implementation is owned by the client organization or referred to qualified partners.

The strategic framing

The MRRE is the verification layer the category has not built, delivered as a structured engagement. It is not the only way to produce verification — but it is a way to produce it in 15 business days, with a defined deliverable, and without adding tooling to a stack that may not need more tooling in the first place.

Validate Your Mobile Security Exposure in 15 Minutes

No pitch. No vendor agenda. This is a controlled assessment to isolate your highest-risk gap and quantify its operational impact.

Book Your 15-Minute Mobile Security Diagnostic

We will:

- Identify your highest-risk mobile exposure
- Map where your current stack is blind
- Show what a structured remediation would produce

If there's no material risk, we'll tell you directly.

Best suited for:

- Security leaders and decision-makers
- Organizations with BYOD or distributed teams
- High-risk individuals managing sensitive data

Limited diagnostic capacity each week.

IDENTIFY MY RISK EXPOSURE →

mobilesecurityguru.com/report · william@mobilesecurityguru.com

Methodology, Citations & About

Methodology

This report synthesizes published industry research, regulatory and standards guidance, publicly available vendor documentation, and field observations from enterprise mobile security engagements. No client data is referenced. All statistics are drawn from publicly available sources cited below. The Signal / Enforcement / Verification framework, the Pocket Attack Surface (PAS) formula, and the five-layer attack surface model are Mobile Security Guru's proprietary analytical frameworks. Where modern stack capabilities are referenced — including Continuous Access Evaluation and Token Protection — the descriptions are drawn from current vendor documentation as of April 2026. Coverage statements are scoped to the conditions documented by those vendors and are subject to change as platform capabilities evolve.

Primary Source Citations

IBM Cost of a Data Breach Report 2024. Global average breach cost of \$4.88M. Published July 2024. The figure is a global all-cause average and is not attributed to any specific initial attack vector.

Verizon Mobile Security Index 2024. Mobile compromise prevalence data and trends in BYOD and unmanaged device exposure. Published May 2024.

NIST SP 800-63B-4 (Digital Identity Guidelines). Authentication and lifecycle management guidance. Classifies SMS OTP as a restricted authenticator. Published July 2025.

NIST SP 800-124 Rev 2. Guidelines for Managing the Security of Mobile Devices in the Enterprise. Published May 2023. Includes centralized device management, endpoint protection, and MTD within scope.

NIST FIPS 203, 204, 205 (Post-Quantum Cryptography Standards). ML-KEM, ML-DSA, SLH-DSA. Finalized August 2024.

CISA Mobile Communications Best Practice Guidance. Mobile communications security guidance for high-risk users following Salt Typhoon intrusions. Published December 2024.

CISA Known Exploited Vulnerabilities Catalog. Ongoing reference for actively exploited vulnerabilities, including management-platform CVEs cited in Section 05.

Microsoft Entra documentation. Continuous Access Evaluation and Token Protection capability and scope, including current preview status and supported services. Referenced as of April 2026.

Huang et al., CONSET — USENIX Security 2026 (forthcoming). Pre-authentication 5G baseband vulnerabilities. Texas A&M; and University of Buffalo. Announced December 2025.

About the author

William Haynes is the founder of Mobile Security Guru and brings nearly two decades of mobile security expertise, most recently as a Senior Security Consultant at T-Mobile for Business. He advises CISOs, COOs, IT directors, and CFOs at enterprise organizations in healthcare, financial services, and defense contracting on building mobile security postures that are verifiable, board-defensible, and built for the actual threat environment — not the compliance checklist.

© 2026 Mobile Security Guru. All rights reserved. This report is provided for informational purposes and does not constitute legal, financial, or regulatory advice. Statistics cited reflect published industry research as of the date of publication. No vendor names appear in this report by design — Mobile Security Guru's advisory is vendor-agnostic and commercially independent.