Coordinated Privacy Disclosure, Final Report

Eustella Android (com.eustella)

AI Newsrooms Technology GmbH, Vienna, Austria

SILENT · CASE CLOSED AND PUBLISHED 25 SEPTEMBER 2026, TWO DAYS AFTER THE EMBARGO, TWO REAL REMEDIATIONS CONFIRMED, FIVE FINDINGS PERSIST, TWO NEW FINDINGS SURFACED ON RE-VERIFICATION, TWELVE MESSAGES, ZERO REPLY

Target	Eustella, marketed as ”Europe’s privacy-compliant, sovereign AI alternative to ChatGPT and Claude”
Package	com.eustella
Operator	AI Newsrooms Technology GmbH, Karl-Farkas-Gasse 22, 1030 Wien, Austria (CEO Matteo Rosoli)
Initial Disclosure	2026-06-25, the app’s own European launch day
Original Embargo	2026-09-23 (90 days from disclosure, as stated in the original notice)
Report Published	2026-09-25, two days after the embargo, an internal publishing gap, not a change in terms
Regulators	Austrian DSB (lead SA), CERT.at
Total Outbound Messages	12, across 79 days: nine in the original R1 round (25 June to 12 August 2026) and three in an 18 August-12 September 2026 Watchtower follow-up round re-presenting findings verified live against version 1.3.0.
Inbound Replies	0. No reply, automated or human, was ever received on any address.

rfi.irfos@gmail.com • rfi-irfos.com

ZVR 1015608684 • GISA 39261441 • UID ATU83405245 • Steuernummer 68 696/8736 • Graz, Austria

Contents

Executive Summary

On 25 June 2026, the day of Eustella's own European launch, RFI-IRFOS disclosed to AI Newsrooms Technology GmbH that its production Android app, marketed on a claim of EU-sovereign, CLOUD-Act-free hosting, hardcodes its authentication, chat, agent-scheduling, and billing API to da08ctfrofx1b.cloudfront.net, an Amazon Web Services domain, and ships a Firebase API key, hardcoded and unrotated, that remained identical across the version reviewed at launch and the version independently re-verified 54 days later. The same launch-day build also shipped a test-environment configuration, test.eustella.com, an alpha Firebase project, and a local development-server IP address, alongside full React Native developer-menu strings, none of which belongs in a production release.

RFI-IRFOS independently re-verified the live production build (version 1.3.0) on 18 August 2026, instead of relying on the original snapshot. Two findings had genuinely been fixed: pre-consent Firebase auto-initialization was switched off, and the biometric permission pair was removed, both silent, unannounced, and both credited in this report. The AWS-hosted backend, the unrotated Firebase key, and the test-environment leak (with a new but still-internal development IP) all persisted unchanged. The re-verification also surfaced two findings absent from the original review: PostHog's session-replay capability, which records screen content and touch interactions, is compiled into the app with no disclosure in either the app's or the website's privacy documentation that AI chat conversations are subject to screen recording, and the app's entire model-routing, agent-execution, and third-party MCP-server OAuth architecture is proxied through the same AWS backend, meaning no user's interaction with the product, however it is described publicly, avoids passing through US-hosted infrastructure.

Twelve messages were sent across 67 days: nine in the original disclosure round through 12 August 2026, and three more between 18 August and 12 September 2026 presenting the live re-verification findings, all to office@eustella.com, cc'ing the Austrian DSB and CERT.at. Not one reply, automated or human, was ever received on any address at any point. The 90-day embargo, set at first contact on 25 June 2026, elapsed on 23 September 2026. This report and the accompanying closure notice publish two days later, an internal gap on RFI-IRFOS's side, not a change to the terms already disclosed.

Scope: Root-level static analysis of the production Eustella Android application, performed twice: version 1.2.0 as pulled and decoded on 25 June 2026 (launch day), and an independent live re-verification of version 1.3.0 (versionCode 25) pulled via ADB from a physical device on 18 August 2026, SHA-256 2122aa9b3f648a11c7dc9eae2f57924cb35a3b8294fb6777f4590101e901cff7. Both reviews used static decompilation and manifest/resource analysis only; no dynamic instrumentation, live traffic capture, or interaction with Eustella's, AWS's, or PostHog's backend infrastructure was performed.

Disposition

Seven findings were disclosed on 25 June 2026 and independently re-verified live against production version 1.3.0 on 18 August 2026, then re-scored under CVSS v4.0. Two original findings were genuinely and silently remediated between versions: pre-consent Firebase auto-initialization was disabled, and the biometric permission pair was removed. Five original findings persist unchanged: a payments/chat/agent backend hardcoded to AWS CloudFront, an unrotated hardcoded Firebase key, an internal test-environment leak with a live dev-server IP, a set of undisclosed US processors, and no network security configuration. Live re-verification also surfaced two new findings: undisclosed PostHog screen-recording of AI chat conversations, and an architecture in which all chat, agent-execution, and MCP-server OAuth traffic is proxied through the same AWS backend. Both new findings reach HIGH on their computed CVSS score alone. A structural contradiction between the company's own public "sovereignty" marketing claim and its own published sub-processors page is documented as an observational finding, since the underlying technical facts are already scored under the findings that support it.

ID	Severity	Title
F1	MEDIUM	Payments, Chat, and Agent API Hardcoded to AWS CloudFront, Contradicting the App's Own Sovereignty Claim
F2	MEDIUM	Firebase API Key Hardcoded and Unrotated Across Versions
F3	REMEDIATED	Pre-Consent Firebase Auto-Initialization — Silently Fixed
F4	MEDIUM	Internal Test Environment Shipped in the Production Build, Persisting With a New Internal IP
F5	MEDIUM	US Processors Present in the Binary but Incompletely Named on the Company's Own Sub-Processors Page
F6	MEDIUM	No Network Security Configuration, No Certificate Pinning, on an AI Chat Application
F7	REMEDIATED	Biometric Permissions — Silently Removed
NG1	OBSERVATIONAL	Public "No API Calls Leave Europe" Claim Directly Contradicts the Company's Own Sub-Processors Page and F1's Technical Evidence
NG2	HIGH	PostHog Session Replay Records Screen Content and Touch Interaction of AI Chat Conversations, Undisclosed
NG6	HIGH	Chat, Agent Execution, and Third-Party MCP-Server OAuth All Proxy Through the Same AWS Backend

Subject & Operator Analysis

Eustella is an Austrian AI chat and agent product operated by AI Newsrooms Technology GmbH, marketed specifically on a claim of EU sovereignty, GDPR-native design, and freedom from US CLOUD Act exposure, in explicit contrast to ChatGPT and Claude.

Several genuine positives are worth naming, including two real, credited fixes made between the two versions reviewed (F3, F7). usesCleartextTraffic is correctly set to false in both versions. allowBackup is correctly set to false. No Chucker HTTP interceptor, Charles Proxy debug certificate, or Flipper debug tooling was found in either version. Deep-link handling (app.eustella.com) is correctly bound to the app's own signing key via a verified assetlinks.json, with no host-hijack vector. No unguarded exported component or OTA-takeover vector was found. No PRC-linked SDK or endpoint was found in either version; an earlier apparent hit on umeng/baidu/getui-like substrings was confirmed, on deeper verification, to be unrelated Catalan-language strings and an unrelated method name, not a real third-party integration. Firebase Remote Config is confirmed disabled, closing a potential hidden remote-configuration push vector.

Eleven Messages, Zero Reply, One Independent Live Re-Verification

This case's correspondence record involved a working contact channel from the first message; no address bounced. Nine messages were sent to office@eustella.com in the original disclosure round, 25 June through 12 August 2026. Rather than relying on the original launch-day snapshot for a closing report 92 days later, RFI-IRFOS independently re-pulled and re-verified the live production build on 18 August 2026, confirming which original findings persisted, which had been genuinely fixed, and surfacing two findings absent from the original review. Three further messages presenting this live re-verification were sent between 18 August and 12 September 2026, cc'ing the Austrian DSB and CERT.at. Not one reply, automated or human, was ever received on any address at any point across the full 79 days of outreach.

Findings

F1: Payments, Chat, and Agent API Hardcoded to AWS CloudFront, Contradicting the App's Own Sovereignty Claim

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.9. Confirmed unchanged on live re-verification of version 1.3.0, 18 August 2026, 54 days after the original finding.

The app's authentication, chat, agent-scheduling, and billing API is hardcoded to da08ctfrofx1b.cloudfront.net, confirmed live and unchanged in version 1.3.0, with the same production endpoints (/api/friends/invites, /api/agent-schedules, /api/streams, /api/v2/conversations, /api/auth/apple, /api/account/subscription, /api/agent-runs, /api/projects). CloudFront is Amazon Web Services, a US company fully subject to the US CLOUD Act. Every user's every API request transits AWS infrastructure.

```
https://da08ctfrofx1b.cloudfront.net/api/  -- unchanged, v1.2.0 (25 June) and v1.3.0
    (18 August)
/api/friends/invites, /api/agent-schedules, /api/streams, /api/v2/conversations, /api/
    auth/apple, /api/account/subscription
```

Impact: Every API request from every user, including authentication, AI chat, and billing, transits AWS infrastructure subject to US compelled-disclosure law, directly contradicting the company's own public claim that no API calls leave European territory.

Fix: Migrate the API backend off AWS CloudFront to the EU-hosted infrastructure already claimed publicly, or correct the public sovereignty claim to reflect the actual architecture. No confirmation that this has occurred was ever received.

F2: Firebase API Key Hardcoded and Unrotated Across Versions

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:L/VA:L/SC:N/SI:N/SA:N, 6.9. Confirmed unchanged on live re-verification of version 1.3.0, 18 August 2026.

The same Firebase API key, app ID, and storage bucket found in the launch-day build were confirmed present, byte-identical, in version 1.3.0, 54 days and one version bump later, meaning the key was never rotated after disclosure.

```
AIzaSyCa0SQF-yQK6WSSADOIOT0Z8QeBLx1o9_w
app_id: 1:764234159319:android:d7370f68ae7b7e57b30c92
storage bucket: eustella-alpha.firebasestorage.app  -- unchanged across both versions
    reviewed
```

Impact: The key remains trivially extractable and, depending on the security-rule configuration behind it, could enable unauthorized access or denial of service against legitimate users; the fact that it was not rotated after direct disclosure indicates the finding was not acted on.

Fix: Rotate the Firebase API key and restrict it by package name and SHA-1 certificate fingerprint. No confirmation that this has occurred was ever received.

F3: Pre-Consent Firebase Auto-Initialization — Silently Fixed

REMEDIATED between the 25 June and 18 August reviews. Not scored; documented as a confirmed, credited fix.

At launch, firebase_messaging_auto_init_enabled and app_data_collection_default_enabled were both set to true, meaning Firebase began collecting before any consent screen. On live re-verification of version 1.3.0, both flags were confirmed set to false. This is a genuine, unannounced fix and is credited here as such.

```
v1.2.0 (25 June): firebase_messaging_auto_init_enabled=true,
    app_data_collection_default_enabled=true
v1.3.0 (18 August): both confirmed false. FirebaseInitProvider still initOrder=100;
    FCM_PUSH_ENABLED still true.
```

Impact: Pre-consent data collection via this specific mechanism no longer occurs. FirebaseInitProvider itself still initializes early and an FCM push flag remains enabled, so this fix is real but partial.

Fix: No further action required for this specific mechanism; the remaining FirebaseInitProvider initOrder and FCM push flag were not part of this specific finding and are not separately scored here.

F4: Internal Test Environment Shipped in the Production Build, Persisting With a New Internal IP

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.9. Confirmed still present on live re-verification of version 1.3.0, with the internal IP changed but the underlying leak unresolved.

The launch-day build shipped a test API base URL (test.eustella.com), the alpha Firebase project ID, a local development-server IP address, and full React Native developer-menu strings ("Open DevTools", "Enable Fast Refresh", "Toggle Sampling Profiler"). On re-verification, the developer-menu strings had been stripped, a hygiene improvement, but the test host, the alpha project ID, and an internal development IP (changed from .212 to .33, still a private-range address) all persisted in the production binary.

```
API_BASE_URL_ANDROID / API_BASE_URL_IOS: https://test.eustella.com  -- persists in v1
    .3.0
project_id: eustella-alpha  -- persists in v1.3.0
react_native_dev_server_ip: 192.168.31.212 (v1.2.0) -> 192.168.31.33 (v1.3.0, changed
    but still internal)
```

Impact: A production release continues to embed test-environment identifiers and an internal development address, indicating the build pipeline does not cleanly separate development and production configuration, a build-hygiene gap that persisted through at least one full version release.

Fix: Strip all test-environment configuration, the alpha project ID, and development-server

addresses from production release builds. No confirmation that this has occurred was ever received.

F5: US Processors Present in the Binary but Incompletely Named on the Company's Own Sub-Processors Page

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.9. Confirmed unchanged on live re-verification of version 1.3.0; re-verification also cross-checked this finding against the company's own published sub-processors page.

RevenueCat, Amazon IAP, PairIP, and multiple hardcoded Google OAuth client IDs are all confirmed present in the binary, unchanged in version 1.3.0. Cross-checking against the company's own live sub-processors page (fetched 18 August 2026) found Google, Apple, and RevenueCat named, but PairIP, an app-protection and license-check SDK compiled directly into the Application class, was named nowhere on that page.

```
RC_GOOGLE_KEY: goog_FQwQZLOuxvabBIfJbCjjezLlkgP  (RevenueCat, US)
com.amazon.device.iap.ResponseReceiver  (Amazon IAP, US)
com.pairip.application.Application  (PairIP, US, the app's own Application class --
    unnamed on the sub-processors page)
4x hardcoded Google OAuth client IDs
```

Impact: A US-based app-protection processor is compiled directly into the app's core Application class yet is not named on the company's own list of sub-processors, a specific, checkable transparency gap distinct from the app's general US-hosting exposure.

Fix: Add PairIP to the published sub-processors list, or remove it and replace it with an EU-based alternative. No confirmation that this has occurred was ever received.

F6: No Network Security Configuration, No Certificate Pinning, on an AI Chat Application

MEDIUM, CVSS v4.0 AV:A/AC:L/AT:N/PR:N/UI:N/VC:L/VI:L/VA:N/SC:N/SI:N/SA:N, 5.3. Confirmed unchanged on live re-verification of version 1.3.0.

No network_security_config.xml is present in either version reviewed. usesCleartextTraffic is correctly set to false, but with no certificate pinning layer, an attacker on the same network holding a trusted CA certificate could intercept API traffic, which in this app's case carries AI chat content, potentially including personal documents and conversations.

Impact: AI chat content, potentially including personal documents, has no additional transport-security hardening beyond default OS certificate validation, a materially higher-stakes exposure than the same gap in a typical consumer app.

Fix: Add a network security configuration with certificate pinning on all production API endpoints. No confirmation that this has occurred was ever received.

F7: Biometric Permissions — Silently Removed

REMEDIATED between the 25 June and 18 August reviews. Not scored; documented as a confirmed, credited fix.

The launch-day build declared USE_BIOMETRIC and USE_FINGERPRINT alongside RECORD_AUDIO, with no explicit disclosure of biometric processing under Art. 9. On live re-verification of version 1.3.0, both biometric permissions were confirmed absent (grep count zero); only RECORD_AUDIO, plausible for voice input to an AI chat product, remains among dangerous permissions.

```
v1.2.0 (25 June): RECORD_AUDIO, USE_BIOMETRIC, USE_FINGERPRINT declared
v1.3.0 (18 August): USE_BIOMETRIC and USE_FINGERPRINT confirmed removed. RECORD_AUDIO
    remains, dangerous-permission count reduced to one.
```

Impact: The Art. 9-adjacent biometric permission vector is closed; the app's permission model at re-verification was confirmed genuinely lean.

Fix: No further action required for this finding.

NG1: A Public "No API Calls Leave Europe" Claim Directly Contradicts the Company's Own Sub-Processors Page and F1's Technical Evidence

OBSERVATIONAL — not independently CVSS-scored; the underlying technical facts are already scored under F1 and F5. This finding documents a structural contradiction between public claims and those facts, not a separate technical exposure.

Eustella's own live /sovereignty page (fetched 18 August 2026) states "No API calls leave European territory. No US or Chinese corporation can revoke access, change model behaviour, or harvest user data." The company's own live sub-processors page, fetched the same day, lists Google LLC, Apple Inc., RevenueCat Inc., and PostHog Inc. as processors and states plainly, "for a small number of services, no European company offers what we need." F1 independently confirms the production API backend runs on AWS CloudFront, a US company. The sovereignty claim, the sub-processors page, and the binary's own architecture do not agree with each other.

```
/sovereignty (live, 2026-08-18): "No API calls leave European territory. No US or
    Chinese corporation can revoke access..."
/services/privacy/sub-processors (live, 2026-08-18): lists Google, Apple, RevenueCat,
    PostHog as processors; "no European company offers what we need"
```

Impact: A specific, checkable legal assurance made to users and prospective customers is contradicted by the company's own published sub-processors disclosure and by the binary's own confirmed architecture, independent of the standing technical findings F1 and F5 that document the underlying facts.

Fix: Bring the sovereignty marketing claim into agreement with the sub-processors page and the actual production architecture. No confirmation that this has occurred was ever received.

NG2: PostHog Session Replay Records Screen Content and Touch Interaction of AI Chat Conversations, Undisclosed

HIGH, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N, 8.7. This finding reaches HIGH on the computed CVSS v4.0 score alone; no editorial escalation was applied or needed. First identified on live re-verification of version 1.3.0, 18 August 2026.

PostHog's React Native session-replay SDK is compiled into the app (posthog-react-native-session-replay, sessionReplayConfig.sampleRate, SessionRecording classes), confirmed via bundle and DEX string analysis. This capability records screen content and touch interactions, meaning, in an AI chat application, the actual content of a user's conversations and any documents shown on screen. Neither the app's nor the website's privacy documentation, as reviewed, names session replay or screen recording as a data-processing activity; PostHog is described only as pseudonymized usage analytics. The SDK references both an EU asset host (eu.i.posthog.com) and a US asset host (us-assets.i.posthog.com), and which region actually serves a given user's replay data is not resolvable from static analysis alone.

```
posthog-react-native-session-replay (RN SDK name, bundle-confirmed)
sessionReplayConfig.sampleRate, _sessionReplayOptions, sessionReplayNativeInitialized
SessionRecording, SessionRecordingComposerializable...
eu.i.posthog.com AND us-assets.i.posthog.com both referenced
```

Impact: Screen and touch-level recording of a user's actual AI chat conversations and any documents shown on screen occurs with no disclosure in either privacy document reviewed, and the geographic routing of that recorded content cannot be confirmed as EU-only from the binary alone.

Fix: Disclose session replay explicitly as screen/touch recording of chat content in both privacy documents, confirm and document the replay data's actual processing region, and consider whether AI chat content constitutes special-category data requiring an Art. 9 basis. No confirmation that this has occurred was ever received.

NG6: Chat, Agent Execution, and Third-Party MCP-Server OAuth All Proxy Through the Same AWS Backend

HIGH, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N, 8.7. This finding reaches HIGH on the computed CVSS v4.0 score alone; no editorial escalation was applied or needed. First identified on live re-verification of version 1.3.0, 18 August 2026.

Beyond the general API hosting documented in F1, the app's bundle confirms a curated MCP (Model Context Protocol) server marketplace and an OAuth broker flow (/api/mcp/marketplace/categories, /api/mcp/servers/oauth/start, /api/mcp/servers/oauth/complete, with a matching deep-link callback), all relative to the same AWS CloudFront backend, alongside model-routing tokens for multiple third-party AI providers. This means a user connecting any third-party MCP server to their account has that connection's OAuth handshake brokered through AWS-hosted infrastructure, not a direct, user-controlled connection.

```
/api/mcp/marketplace/categories_* -- curated MCP server directory on Eustella's
    backend
/api/mcp/servers/oauth/start + /api/mcp/servers/oauth/complete -- OAuth brokered via
    CloudFront
/settings/mcp-servers/oauth-callback -- matching deep link
model-routing tokens: verda, xai, gpt-user, gemini, perplexity, together
```

Impact: Every layer of the product, chat content, agent task execution, and any third-party service a user connects via MCP, is proxied through the same US-hosted broker documented in F1, meaning the sovereignty claim in NG1 fails not only for chat but for the entire connected-services layer of the product.

Fix: Document the MCP-OAuth broker's actual hosting location and processor status under Art. 28, and clarify whether user-added MCP servers can be reached directly instead of through the AWS broker. No confirmation that this has occurred was ever received.

Data Protection, Article by Article

The findings map onto a small, specific set of GDPR provisions, set out below in one place instead of scattered across the findings themselves.

Article	Finding	Basis
Art. 44-49, Art. 5(1)(a)	F1, NG1	US-hosted backend directly contradicting a public no-third-country-transfer claim.
Art. 32	F2	Hardcoded, unrotated Firebase API key across two production versions.
Art. 32, Art. 5(1)(f)	F4	Test-environment configuration and an internal development address shipped in production.
Art. 13(1)(e)	F5	A US app-protection processor compiled into the app but unnamed on the company's own sub-processors page.
Art. 32	F6	No certificate pinning on an application transmitting AI chat content.
Art. 13(1)(e), Art. 9, Art. 44	NG2	Undisclosed screen/touch recording of AI chat conversations with an unresolved processing region.
Art. 28, Art. 13(1)(e), Art. 44	NG6	Chat, agent, and third-party MCP-OAuth traffic all proxied through the same US-hosted broker.

RFI-IRFOS rfi.irfos@gmail.com • rfi-irfos.com

ZVR 1015608684 • GISA 39261441 • UID ATU83405245 • Steuernummer 68 696/8736 • Elisabethinergasse 25/10, 8020 Graz, Austria CVSS v4.0 is RFI-IRFOS's scoring guideline, not a ceiling or a floor. It gives a reproducible, third-party-legible starting point for severity, and it does not capture real-world blast radius on its own, factors such as user population size, a vulnerable or protected class of users, or a standing corporate governance exposure the technical vector alone cannot express. RFI-IRFOS reserves the right to weigh that blast radius above the bare CVSS score, and, more rarely, to hold a score below what CVSS alone would suggest when the underlying claim is unconfirmed or disputed with evidence. In this report, NG2 and NG6 reach HIGH on their computed CVSS v4.0 scores alone, no editorial escalation applied; F1, F2, F4, F5, and F6 correct to MEDIUM to match their own computed bands; F3 and F7 are documented as confirmed, credited remediations

instead of scored findings; NG1 is classified OBSERVATIONAL because its underlying technical facts are already scored under F1 and F5. Disclosure conducted under ISO/IEC 29147 and the Austrian Forschungsfreiheitsgesetz, Art. 17 StGG. REF REF: EUSTELLA-2026-R1.