Coordinated Privacy Disclosure, Final Report

Google Gemini Android (com.google.android.apps.bard)

Google LLC, Mountain View, USA

SILENT · CASE CLOSED AND PUBLISHED 25 SEPTEMBER 2026, ONE DAY AFTER THE STATED EMBARGO, FOUR MESSAGES, NO HUMAN REPLY

Target	Google Gemini, an AI assistant application
Package	com.google.android.apps.bard
Operator	Google LLC, Mountain View, USA
Initial Disclosure	2026-06-26
Stated Embargo	2026-09-24, as communicated directly to the target in the initial disclosure
Report Published	2026-09-25, one day after the stated date
Regulators	UK ICO (jurisdiction over Alphabet's UK GDPR obligations), Austrian DSB, CERT.at
Total Outbound Messages	4, across 72 days, to privacy@google.com and security@google.com.
Inbound Replies	0 human replies. One automated bug-bounty triage bot reply directing the report to Google's separate vulnerability-reward program; one bounce from a mistyped regulator address, corrected on the next send.

rfi.irfos@gmail.com • rfi-irfos.com
ZVR 1015608684 • GISA 39261441 • UID ATU83405245 • Steuernummer 68 696/8736 • Graz, Austria

Contents

Executive Summary

On 26 June 2026, RFI-IRFOS disclosed to Google LLC that its Gemini Android app declares no android:networkSecurityConfig. usesCleartextTraffic is correctly set to false, plaintext HTTP is blocked, but without an NSC file there is no certificate pinning for gemini.google.com. Any system-trusted CA, including an enterprise MDM-pushed certificate, can issue a valid certificate for Google domains and read conversation traffic in transit, conversations that routinely contain Art. 9 special-category data such as medical, mental-health, and political content. A second, related gap: because Gemini is architecturally a WebView shell pointing to gemini.google.com, camera and microphone permission prompts identify the browser as the requesting entity, not Gemini itself, leaving the actual data controller unidentified at the point of collection. This computes to HIGH/7.1 under CVSS v4.0, an undershoot correction from the original MEDIUM label; RFI-IRFOS's own real-world assessment at disclosure time predates CVSS v4.0 adoption into the methodology and is not retrofitted here, but the computed band is reported faithfully.

A second finding: Clearcut and usagereporting, two Google telemetry systems, transmit performance metrics, feature-usage signals, and session data to Google infrastructure. A ConsentVerifier Phenotype library is present, showing Google has built a consent-gate mechanism, but in every other app in this campaign, pre-consent auto-init is visible as a ContentProvider initOrder hook; here, no such gate is visible in the Gemini binary; whether this telemetry fires before or after ToS acceptance cannot be verified from the APK alone. This computes to MEDIUM/6.9.

A third finding: Gemini requires a Google account with no anonymous or pseudonymous mode, visible in the binary via a mandatory sign-in path with no guest option. Every conversation is permanently linked to a Google identity that simultaneously holds Gmail, Search history, Maps location history, YouTube history, Calendar, and Drive. Users routinely share Art. 9 special-category data with AI assistants, medical conditions, mental health history, relationship and financial details, and Google's own privacy policy permits human review of conversations for safety and quality purposes. This combination of mandatory identity linkage, cross-product data access, Art. 9 content, and human review capability computes to MEDIUM/6.9. This finding is architectural and policy-level, not a code defect; it is not fully verifiable from the APK alone and is best addressed through privacy policy clarification and an explicit pseudonymous option.

Genuine positives dominate this report in a way no other target in this campaign matched: of 58-plus applications audited, Gemini is the only AI assistant found with zero third-party tracking SDKs, no Firebase Analytics, Braze, Segment, Facebook SDK, Adjust, or AppsFlyer; zero hardcoded credentials; zero ContentProvider auto-init; allowBackup correctly set to false; no advertising-identifier permission; no Privacy Sandbox configuration; and a disabled boot receiver. Four messages were sent across 72 days to privacy@google.com and security@google.com. One automated reply arrived, a bug-bounty triage bot directing the report to Google's separate vulnerability-reward program, not engaging on the data-protection substance. No human reply was ever received. The 90-day embargo, stated directly to the target as 2026-09-24, elapsed one day before this report and its accompanying closure notice publish.

Scope: Root-level static analysis of the production Google Gemini Android application (com.google.android.apps.bard), pulled from Google Play and reviewed on 26 June 2026, as part of a 58-plus application AI-assistant and consumer-app audit series. No account was created and no interaction with Google's backend infrastructure was performed.

Disposition

Three findings were disclosed on 26 June 2026, re-scored under CVSS v4.0. The absence of a network security config, meaning conversation traffic carries no certificate pinning, corrects to HIGH. An unverifiable consent gate for behavioral telemetry, and mandatory linkage of every conversation to a user's full Google identity graph with no pseudonymous option, both correct to MEDIUM. Of the 58-plus applications audited in this campaign, Gemini's own disclosure opened by naming it the cleanest AI assistant reviewed, zero third-party tracking SDKs, zero hardcoded credentials, zero pre-consent auto-init. Four messages were sent across 72 days. No human reply was ever received.

ID	Severity	Title
M1	HIGH	No Network Security Config: Conversation Traffic Carries No Certificate Pinning
M2	MEDIUM	Behavioral Telemetry Consent Gate Not Verifiable in the Binary
M3	MEDIUM	Every Conversation Is Mandatorily Linked to the Full Google Identity Graph

What Gemini Got Right

Of the 58-plus applications audited in this campaign, Gemini is the only AI assistant found with zero third-party tracking SDKs, no Firebase Analytics, Braze, Segment, Facebook SDK, Adjust, or AppsFlyer. Zero hardcoded credentials, no Firebase keys, OAuth secrets, or database URLs. Zero ContentProvider auto-init. allowBackup is correctly set to false, conversation history is not ADB-extractable. usesCleartextTraffic is correctly set to false. The advertising-ID permission is absent, and no Privacy Sandbox configuration is present. The boot receiver is present but explicitly disabled. Three declared permissions total: WAKE_LOCK, ACCESS_NETWORK_STATE, GET_PACKAGE_SIZE. This is not a coincidence; it is engineering discipline applied consistently, and it sets the standard against which every other AI assistant in this series was measured.

Four Messages, No Human Reply

Four messages were sent across 72 days to privacy@google.com and security@google.com. The Austrian DSB, CERT.at, and the UK ICO, which holds jurisdiction over Alphabet's UK GDPR obligations, were copied throughout, with one initial regulator address correction after a bounce. One automated reply arrived, a bug-bounty triage bot directing the report toward Google's separate vulnerability-reward program instead of engaging with the data-protection questions the disclosure actually raised. No human reply of any kind was ever received.

Findings

M1: No Network Security Config: Conversation Traffic Carries No Certificate Pinning

HIGH, CVSS v4.0 AV:A/AC:L/AT:N/PR:N/UI:N/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N, 7.1.

The app declares no android:networkSecurityConfig. usesCleartextTraffic is correctly set to false, but without an NSC file there is no certificate pinning for gemini.google.com. Because Gemini is architecturally a WebView shell, camera and microphone permission prompts identify the browser as the requesting entity, not Gemini itself.

```
android:networkSecurityConfig -> NOT DECLARED
android:usesCleartextTraffic="false" -> present (correct)
certificate pinning for gemini.google.com -> ABSENT
```

Impact: Any system-trusted CA, including an enterprise MDM-pushed certificate on a managed device, can issue a valid certificate for Google domains and read conversation traffic in transit, conversations that routinely contain Art. 9 special-category data.

Fix: Declare a network security config with certificate pinning for gemini.google.com, and attribute WebView camera/microphone requests to Gemini specifically at the point of collection. No confirmation that this has occurred was ever received.

M2: Behavioral Telemetry Consent Gate Not Verifiable in the Binary

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.

Clearcut (ClearcutMetricSnapshotTransmitter) and com.google.android.gms.usagereporting transmit session data, feature-usage signals, and performance metrics. A ConsentVerifier Phenotype library is present, but no ContentProvider initOrder hook is visible, meaning consent-gate timing cannot be verified from the APK alone.

Impact: Behavioral data about how users engage with an AI assistant, which conversation modes are used, how often document and image analysis is triggered, session frequency, is transmitted with no verifiable proof that this occurs only after explicit consent.

Fix: Make the Clearcut and usagereporting consent gate mechanically verifiable from the production binary, or document the consent-gate timing for regulator review. No confirmation that this has occurred was ever received.

M3: Every Conversation Is Mandatorily Linked to the Full Google Identity Graph

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.

Gemini requires a Google account with no anonymous or pseudonymous mode, confirmed via a mandatory sign-in path with no guest option. Every conversation is permanently linked to a Google identity that simultaneously holds Gmail, Search history, Maps location history, YouTube history, Calendar, and Drive.

Impact: Art. 9 special-category data shared routinely in AI conversations, medical, mental-health, financial, and political content, combines with cross-product identity linkage and Google's own stated human-review capability in a way most users do not understand when starting what they believe is a private conversation.

Fix: Offer an explicit pseudonymous or guest conversation mode, and document the specific consent basis for using conversation data beyond the chat service itself. 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. 32(1)(a), Art. 13(1)(a)	M1	No network security config, conversation traffic carries no certificate pinning.
Art. 6(1)(a), Art. 13(1)(e)	M2	Behavioral telemetry consent gate not verifiable in the binary.
Art. 5(1)(b), Art. 5(1)(c), Art. 9(2)(a), Art. 22	M3	Every conversation mandatorily linked to the full Google identity graph.

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. In this report, M1 corrects to HIGH, an undershoot correction from the original real-world MEDIUM label assigned at disclosure time, before CVSS v4.0 was adopted into the methodology; M2 and M3 correct to MEDIUM, matching the original label. Disclosure conducted under ISO/IEC 29147 and the Austrian Forschungsfreiheitsgesetz, Art. 17 StGG. REF REF: GEMINI-2026.