Security & Privacy Disclosure, Final Report

Adobe Scan (AI PDF Scanner, OCR) for Android (com.adobe.scan.android)

Root level code analysis of the document/OCR category, acknowledged by Adobe on 18 September 2026 and answered on 20 November 2026 with operator statements that RFI-IRFOS has not verified.

PUBLIC DISCLOSURE, 21 NOVEMBER 2026, EMBARGO REACHED. PUBLICATION DATE 21 NOVEMBER 2026 AS STATED IN WRITING. ADOBE REPLIED ON 20 NOVEMBER 2026 WITH A SUBSTANTIVE ANSWER TO THE THREE QUESTIONS.

Target (NA/Japan)	Adobe Inc., San Jose, California, USA
Target (EU/EEA and rest of world)	Adobe Systems Software Ireland Limited, 4-6 Riverwalk, Citywest Business Park, Dublin 24, Ireland
Versions Audited	26.08.10 / 262602995 (23 Aug 2026, original notice)
Notice / Publication	2026-08-23 / 2026-11-21 (90 days)
Findings	6: 1 Critical, 3 High, 2 Medium
Status	ADOBE ANSWERED THE THREE QUESTIONS ON 20 NOVEMBER 2026 (63 days after its 18 September message, 89 days after the original notice). The statements are stated by the operator, not verified by RFI-IRFOS.
SHA-256 (base APK)	`0bb8784730e79e167cd3f559026675ed2d072152c0e291727bdf233e176b22c7`

security@rfi-irfos.com • rfi-irfos.com

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

Contents

1 Executive Summary 4

2 Finding Status at a Glance 6

3 Subject & Operator Analysis 7

4 Architecture & Data Flow 7

5 Engagement Summary 7

6 Findings 8
- 6.1 C1: Pre-consent ad/analytics/attribution SDK auto-init (Meta, Firebase, InMobi, Branch); no consent-management platform found . . . 8
- 6.2 H1: Third-party/bystander document data: ID-card and business-card classification, Contacts write path, cloud OCR/GenAI upload . . . 9
- 6.3 H2: Disclosed manual/human review of scanned-document content and AI Assistant prompts, no third-party-specific safeguard . . . 10
- 6.4 H3: Firebase/Google API key and project identifiers hardcoded in distributed binary . . . 11
- 6.5 M1: No certificate pinning on first-party document-upload and GenAI endpoints 11
- 6.6 M2: Play Store Data Safety declaration omits the document/photo-content category the app's own cloud pipeline confirms . . . 12

7 Child Protection & Student Data 12

8 Impact Analysis 13

9 The Complete Chain 13

10 Data Protection: Article by Article 14

11 Regulatory & Legal Landscape 14

12 Indicators of Compromise & Reproducible Evidence 15

13 Disclosure Timeline & Outcome 15

14 What Adobe Got Right 16

15 Residual Risk & Recommendations 16

Executive Summary

Adobe Scan (com.adobe.scan.android) is Adobe's document-scanning and OCR app, distributed on Google Play with 100M+ installs and a PEGI 3 content rating. On 23 August 2026 RFI-IRFOS disclosed six findings from a root level code analysis of version 26.08.10 under REF ADOBESCAN-2026-R1. Adobe already has a genuine EU establishment, Adobe Systems Software Ireland Limited, acting as controller for EU/EEA users with a named Data Protection Officer, so this case does not carry the Art. 27 representative gap common elsewhere in this programme.

The core finding is structural rather than incidental: Adobe Scan's own code contains a purpose-built ID-card capture mode and a business-card classifier that can write a scanned business card's contact details into the device owner's Android address book, and every scan is demonstrably uploaded to Adobe's Document Cloud API and, for AI-Assistant features, into Adobe's GenAI ingestion pipeline. Adobe's own central Privacy Policy affirmatively states that user content "may contain personal information of other individuals" and places the entire compliance burden for that on the account holder. A second, independent pattern recurs across this audit programme: Meta Audience Network, Firebase, InMobi, and Branch.io all initialize via ContentProviders before any consent screen can render, and no consent-management platform was found anywhere in the binary.

Adobe's Data Protection Officer acknowledged the notice on 18 September 2026, twenty-six days after delivery, confirmed the correspondence was being treated as a coordinated disclosure, and asked seven clarifying questions about the evidence (SDK versions, whether network transmission versus registration-only was observed, the build analysed, and the API key type). RFI-IRFOS answered every one of those seven questions in full the same day, with the exact SHA-256 of the analysed build and a point-by-point response. RFI-IRFOS sent notices on 17 October 2026 (35 days to publication) and on 15 November 2026. Adobe's Data Protection Officer replied on 20 November 2026, one day before publication, with the Adobe Scan group address in copy.

The reply of 20 November 2026 answers the three original questions. On the pre-consent SDKs it states that consent gating has been implemented "going forward" for Meta Audience Network and InMobi, that Branch.io will be moved behind the consent gate, and that no personal data is transmitted by these SDKs before consent in the current build. On third-party document data it places the responsibility on the user, describes ID-card classification as on-device, and concludes that a DPIA "is not required". On the Data Safety declaration it states that Adobe "does not collect or use that content for its own purposes". RFI-IRFOS has verified none of these statements against a build and reports them as stated by the operator, not verified by RFI-IRFOS. The findings below keep their severity and technical evidence.

Scope: Root level code analysis of the Google Play release bundle 26.08.10 (versionCode 262602995, base.apk SHA-256 above): manifest, resources, and dex-string/grep inventory over classes1-4.dex and assets/audience_network.dex (apktool –no-src, memory-safe mode; no baksmali disassembly this pass). Passive review of Adobe's own public Privacy Policy, Generative AI User Guidelines, and Play Store Data Safety declaration, each fetched live and quoted verbatim. No Adobe backend, account, or API key was probed or tested.

ID	Severity	Title
C1	CRITICAL	Pre-consent ad/analytics/attribution SDK auto-init (Meta, Firebase, InMobi, Branch); no consent-management platform found
H1	HIGH	Third-party/bystander document data: ID-card and business-card classification, Contacts write path, cloud OCR/GenAI upload
H2	HIGH	Disclosed manual/human review of scanned-document content and AI Assistant prompts, no third-party-specific safeguard
H3	HIGH	Firebase/Google API key and project identifiers hardcoded in distributed binary
M1	MEDIUM	No certificate pinning on first-party document-upload and GenAI endpoints
M2	MEDIUM	Play Store Data Safety declaration omits the document/photo-content category the app’s own cloud pipeline confirms

ID	Disclosed	This report	Status as of 21 November 2026

Finding Status at a Glance

ID	Disclosed	This report	Status as of 21 November 2026
C1	CRITICAL	CRITICAL	Answered by Adobe on 20 November 2026. Operator states consent gating for Meta Audience Network and InMobi is implemented going forward and Branch.io will follow; no pre-consent transmission in the current build. Not verified by RFI-IRFOS.
H1	HIGH	HIGH	Answered by Adobe on 20 November 2026. Operator places responsibility for third-party data on the user, states ID-card classification runs on-device and that a DPIA is not required. Not verified by RFI-IRFOS.
H2	HIGH	HIGH	Answered in part on 20 November 2026. Operator calls the review text generic Adobe-wide guidelines and states scanned documents are not used for AI model training. The reply names no mechanism that excludes third-party data from human review. Not verified by RFI-IRFOS.
H3	HIGH	HIGH	Adobe's 20 November 2026 reply states that Firebase is used to register crash reports. It does not address the hardcoded key or its server-side restrictions. Open.
M1	MEDIUM	MEDIUM	Not raised in Adobe's clarifying questions or in the reply of 20 November 2026; stands as disclosed.
M2	MEDIUM	MEDIUM	Answered on 20 November 2026. Operator states it does not collect or use the content for its own purposes. Not verified by RFI-IRFOS; the declaration as audited omits the categories.

Subject & Operator Analysis

Adobe operates Adobe Scan for North America and Japan directly as Adobe Inc. (San Jose, California), and for the EU/EEA and the rest of the world through Adobe Systems Software Ireland Limited (Dublin, Citywest Business Park), a genuine EU main establishment with a named Data Protection Officer (DPO@Adobe.com). This is a materially different starting posture than most targets in this programme: Adobe does not have an Art. 27 representative gap, because it already has an EU establishment acting as controller. Adobe's DPO function acknowledged this notice on 18 September 2026, confirmed it was being handled as a co-ordinated disclosure by its security and privacy teams, and sent its written answers on 20 November 2026.

Architecture & Data Flow

Every scan in Adobe Scan is locally classified on-device by a BusinessCardClassifier and a generic documentTypeClassifier, and persisted with business_card_classification and document_classification flags in a local Room database (ScanFilePersistentData). A dedicated ID-card capture mode (IDCardFirstPage/IDCardSecondPage) exists with its own front/back UI flow. ContactsContract write-path classes are present alongside the business-card classifier, confirming a path from a scanned business card to the device owner's own Android address book. The raw scan is demonstrably uploaded to Adobe's Document Cloud API (dc-api.adobe.io, dc-asset-transfer.adobe.io) for server-side OCR, and, for the AI Assistant and Generative-summary features, into Adobe's GenAI ingestion pipeline (documentcloud.adobe.com/genai_services/*). Meta Audience Network, Facebook AppEvents/Login/Share, Firebase, InMobi (ads, signals, unifiedId), and Branch.io all initialize via ContentProvider/androidx.startup entries at process creation, before any Activity, and no IAB TCF, OneTrust, Didomi, Usercentrics, or Google UMP consent-management class was found anywhere in the decompiled binary.

Engagement Summary

The first notice went to DPO@adobe.com and Grp-adobescan-android@adobe.com on 23 August 2026, with the Austrian DSB and the EDPS in copy. Neither address bounced; neither replied for twenty days, prompting a follow-up on 12 September 2026. Adobe's Data Protection Officer replied on 18 September 2026, confirming receipt, that the matter had the attention of Adobe's security and privacy teams, and that a substantive response would follow, while asking seven specific clarifying questions about the evidence base (SDK versions and ContentProvider entries observed, whether network transmission or registration-only was seen, the fresh-install state and regulatory configuration tested, the endpoints and payloads for the raw-scan upload, what in the binary evidences GenAI ingestion specifically, the key type extracted, and the exact build/SHA-256 analysed). RFI-IRFOS answered all seven in full the same day, 18 September 2026, including the SHA-256 of the exact build analysed. RFI-IRFOS sent a 35-days-to-publication notice on 17 October 2026 and a further notice on 15 November 2026. Adobe's Data Protection Officer answered the three original questions on 20 November 2026.

Findings

C1: Pre-consent ad/analytics/attribution SDK auto-init (Meta, Firebase, InMobi, Branch); no consent-management platform found

CRITICAL, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:N/SI:N/SA:N, 0.0 (CVSS-B), CVSS scores technical exploitability and gives a pure pre-consent timing defect 0.0

Meta Audience Network, Facebook AppEvents/Login/Share, Firebase, InMobi (ads, signals, unifiedId), and Branch.io all initialize via ContentProvider/androidx.startup entries at Application creation, structurally before any Activity including a first-run consent screen. No OneTrust, Didomi, Usercentrics, consentmanager.net, or Google UMP class was found anywhere in the decompiled binary; the only consent-adjacent surface found is Adobe's own ConsentStatus extension, which gates Adobe's own Experience Cloud analytics, not the third-party ad/attribution stack above.

```
<provider android:authorities="com.adobe.scan.android.AudienceNetworkContentProvider"
    android:exported="false"
          android:name="com.facebook.ads.AudienceNetworkContentProvider"/>
<provider android:authorities="com.adobe.scan.android.firebaseinitprovider" android:
    directBootAware="true"
          android:exported="false" android:initOrder="100" android:name="com.google.
    firebase.provider.FirebaseInitProvider"/>
<provider android:authorities="com.adobe.scan.android.FacebookInitProvider" android:
    exported="false"
          android:name="com.facebook.internal.FacebookInitProvider"/>
```

Declared advertising-identifier permissions: com.google.android.gms.permission.AD_ID, ACCESS_ADSERVICES_TOPICS, ACCESS_ADSERVICES_ATTRIBUTION, ACCESS_ADSERVICES_AD_ID. The ad stack (Meta Audience Network, InMobi) is present and auto-initializing regardless of subscription status: no isPremium/hasSubscription gating string or manifest flag was found on any of these providers.

Adobe's DPO acknowledged this finding on 18 September 2026 and asked which SDK versions and ContentProvider entries were observed and on which build variant, and whether network transmission or registration-only was observed. RFI-IRFOS answered both in full the same day, naming the exact manifest entries above and stating plainly that ContentProvider registration/initialization, not confirmed network transmission, is what the binary demonstrates; Android's own component lifecycle guarantees initialization runs before any Activity regardless.

Impact: Every installation of a 100M+-install app runs ad-mediation, attribution, and behavioural-analytics SDK initialization before any EU user has seen or made a consent choice, with no consent-management platform gating that initialization anywhere in the binary.

Fix: Answered by the operator on 20 November 2026, stated by the operator, not verified by RFI-IRFOS. Adobe writes: "going forward we have implemented consent gating for the advertising SDKs (Meta Audience Network and InMobi)", and that it will move Branch.io's initialization behind the consent gate. Adobe states that Firebase registers crash reports, that

no personal data is transmitted by these SDKs before consent in the current build, and that Art. 6(1)(a) is the basis once consent is obtained, with Art. 6(1)(b) and (f) for strictly necessary processing. Adobe asked for network-traffic captures. The gated build version was not named, and the Branch.io change is announced, not completed. Severity unchanged.

H1: Third-party/bystander document data: ID-card and business-card classification, Contacts write path, cloud OCR/GenAI upload

---

HIGH, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:H/SI:N/SA:N, 7.5 (CVSS-B), confidentiality impact on a data subject outside the vulnerable system (the scanned third party, not the account holder)

A dedicated ID-card capture mode (IDCardFirstPage/IDCardSecondPage) and an on-device BusinessCardClassifier and documentTypeClassifier exist and persist a business_card_classification and document_classification flag per scan. Classified business cards can be written into the device's own Android Contacts database via confirmed ContactsContract classes. The raw scan is demonstrably uploaded to Adobe's Document Cloud API for server-side OCR and, per the marketed AI Assistant features, into Adobe's GenAI ingestion pipeline.

```
CREATE TABLE ScanFilePersistentData (... business_card_classification INTEGER NOT NULL
    ,
  document_classification INTEGER NOT NULL, captureIntelligenceConversationId TEXT,
  captureIntelligenceGenAiDocId TEXT, ...)
https://dc-api.adobe.io/   https://dc-asset-transfer.adobe.io/schemas/
    authorize_upload_v1.json
https://documentcloud.adobe.com/genai_services/schemas/ingestion_v1.json
```

Adobe's own central Privacy Policy, fetched live 23 August 2026 from adobe.com/privacy/policy.html: "If the content or information that you store in Adobe Services and Software contains personal information of other individuals, you must be legally permitted to disclose the personal information to Adobe."

Adobe's DPO asked on 18 September 2026 for the specific endpoints and payloads observed for the raw-scan upload, and what in the binary evidences GenAI ingestion specifically as opposed to on-device processing. RFI-IRFOS answered both in full the same day, naming the exact dc-api.adobe.io and documentcloud.adobe.com/genai_services endpoints and the DCAssetBlockUploadInitializeBody/DCAssetCreatePdfBody API client classes that call them.

Impact: A business-card owner, an ID-card holder, another signatory on a scanned contract, or a child whose school form is scanned has their personal data captured, classified, OCR'd, and uploaded to Adobe's cloud, without ever having seen a consent screen, a privacy notice, or any communication from Adobe at all.

Fix: Answered by the operator on 20 November 2026, stated by the operator, not verified by RFI-IRFOS. Adobe writes that the user initiates the processing and that responsibility for a third party's data "is on the user", that the ID-card classification "operates entirely on-device", and that a DPIA for Adobe's processing of that content "is not required". Adobe also states it does not systematically extract, index or separately process third-party personal data. The reply does not address the Contacts write path or the Art. 14 notice question. Severity unchanged.

H2: Disclosed manual/human review of scanned-document content and AI Assistant prompts, no third-party-specific safeguard

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:L/SI:N/SA:N, 5.8 (CVSS-B), rated HIGH by blast-radius escalation

RFI-IRFOS blast-radius escalation: raised from MEDIUM/5.8 to HIGH because the reviewed content can include ID cards and business cards by the product's own designed capture modes (H1), meaning human reviewers, not only automated systems, may see a non-consenting third party's identity document on a 100M+-install app with no document-sensitivity routing evidenced anywhere in the binary.

Adobe's central Generative AI User Guidelines, verbatim: "Your prompts and inputs to, and the results generated by, generative AI features in Adobe products may be reviewed through both automated (such as machine learning) and manual methods for abuse prevention and content filtering purposes." The same guidelines tell users not to input sensitive personal information "unless the product is designed to process that type of information", a carve-out Adobe Scan's own ID-card mode falls squarely inside.

Adobe's DPO asked on 18 September 2026 for the mechanism, if any, that excludes a third party's data from this review. RFI-IRFOS's answer of the same date stated that no such mechanism was found evidenced in the binary and asked Adobe to confirm whether one exists server-side.

Impact: Documents containing a non-consenting third party's identity or signature data can be reviewed by a human reviewer under a generic abuse-prevention policy, with no disclosed routing that would exclude or redact that content first.

Fix: Answered in part on 20 November 2026, stated by the operator, not verified by RFI-IRFOS. Adobe calls the quoted text "generic Adobe-wide AI usage guidelines", states that cloud features run only when the user invokes them and that "scanned documents are not used for AI model training". The reply names no mechanism that excludes or redacts a third party's data before a human reviewer sees it. Severity unchanged.

H3: Firebase/Google API key and project identifiers hardcoded in distributed binary

LOW, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:L/SC:N/SI:N/SA:N, 2.1 (CVSS-B), rated HIGH by blast-radius escalation

RFI-IRFOS blast-radius escalation: raised from LOW/2.1 to HIGH because this key sits in the core document-upload pipeline binary of a 100M+-install app, and no App Check or Firebase Security Rules confirmation has been given for the project it unlocks (adobe-scan-4df73), which also receives the scan/OCR traffic described in H1.

```
<string name="google_api_key">AIzaSyBUUctoBrbNnn4ql4WcfftiLm3yRykH8M0</string>
<string name="google_app_id">1:827597488311:android:b2137c31aea0d481</string>
<string name="project_id">adobe-scan-4df73</string>
<string name="firebase_database_url">https://adobe-scan-4df73.firebaseio.com</string>
<string name="google_storage_bucket">adobe-scan-4df73.appspot.com</string>
<string name="gcm_defaultSenderId">827597488311</string>
```

Adobe's DPO asked on 18 September 2026 to confirm the specific key type and scope extracted. RFI-IRFOS confirmed in full the same day: a Firebase project API key, app ID, sender ID, and Cloud Storage bucket, all extractable in seconds from any copy of the distributed APK.

Whether Firebase Security Rules and Cloud Storage bucket permissions are correctly locked down server-side on the receiving end was not tested and is UNVERIFIED from the client binary alone; no live request was sent against adobe-scan-4df73.

Impact: The key can be used to probe or exhaust project quotas for the adobe-scan-4df73 project, which also receives the scan/OCR traffic described in H1, and to attempt reads against the project's Cloud Storage bucket and Firebase database.

Fix: Open. Adobe's reply of 20 November 2026 mentions Firebase only as used for crash reports and does not address the key restrictions or the Security Rules.

M1: No certificate pinning on first-party document-upload and GenAI endpoints

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:P/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.3 (CVSS-B), no pinning found on dc-api.adobe.io, dc-asset-transfer.adobe.io, or documentcloud.adobe.com

android:usesCleartextTraffic="false" is explicitly set (a genuine positive, see What Adobe Got Right), but no base-config or per-domain pin-set exists in network_security_config.xml, so there is no certificate pinning for any first-party endpoint carrying scanned-document content.

Impact: For an app whose core function is transmitting document images that may contain government ID numbers and signatures, the TLS channel depends solely on the OS trust store, which is interceptable on a managed device or by an attacker able to install a trusted CA.

Fix: Open; not raised in Adobe's clarifying-questions message or in its reply of 20 November 2026.

M2: Play Store Data Safety declaration omits the document/photo-content category the app's own cloud pipeline confirms

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:N/SI:N/SA:N, 0.0 (CVSS-B), transparency finding, not a technical defect

Adobe's self-certified Play Store Data Safety section lists collected data as App activity and Personal info (name, email, user IDs, phone number) only, and shared data as App info and performance only. No Files and docs or Photos and videos category appears, despite the Document Cloud upload and GenAI ingestion pipeline confirmed in H1.

Impact: A prospective user reviewing the store listing's own privacy summary before installing would not learn that scanned document content, the app's entire reason to exist, leaves the device for Adobe's cloud.

Fix: Answered by the operator on 20 November 2026, stated by the operator, not verified by RFI-IRFOS. Adobe writes that the declaration omits "Files and docs" and "Photos and videos" because "Adobe does not collect or use that content for its own purposes". The declaration as audited is unchanged in this report.

Child Protection & Student Data

Adobe Scan carries a PEGI 3 content rating with no date-of-birth, age-gate, or COPPA-block string or flow found in the app's own code. The app ships a Student-facing "Student Space" feature (IDS_ARSS_ADD_TO_STUDENT_SPACE, "Upload files, links, or text to a Student Space, they'll be processed in the cloud automatically"), and Adobe's own Privacy Policy states that Adobe's "Document Cloud apps" are "available to students, including children, for educational purposes under the supervision of school personnel" under a separate COPPA-covered Student Services notice. The only COPPA-related string found in the binary (coppa_store) belongs to the InMobi ad SDK's own child-directed-treatment flag, a developer-set value for ad requests, not a user-facing age gate; whether it is set true is unverified from strings alone. A PEGI-3, no-age-gate, ad-monetised (Meta Audience Network, InMobi) document scanner sits alongside this student/child-facing surface without any reconciliation evidenced in the binary between the all-ages rating and the ad SDK stack.

Impact Analysis

The population most exposed by H1 and H2 is not Adobe Scan's own 100M+ account holders but the people pictured, named, or signed on the documents those account holders scan: a business-card owner, a person on the other side of a scanned contract, a patient whose medical form is scanned, or a child whose report card is scanned. None of them installed Adobe Scan, saw a consent screen, or received a privacy notice from Adobe, and the entire compliance burden for having disclosed their data to Adobe is, per Adobe's own Privacy Policy, placed on the account holder, not discharged by Adobe. C1's pre-consent ad/analytics initialization affects every one of the 100M+ installs on every cold start, regardless of jurisdiction. H3's key sits in the same binary that handles the H1 upload pipeline, so an availability failure on that key degrades a document-upload tool's core function.

The Complete Chain

Point of contact: a user installs Adobe Scan, ad-supported and an optional Premium subscription, and scans a document, which may be their own or may be someone else's, a business card handed to them, an ID they are asked to copy, a form a third party filled in. Processing: the scan is locally classified by the BusinessCardClassifier/documentTypeClassifier before the user makes any further choice, and, simultaneously, Meta, Firebase, InMobi, and Branch.io SDKs have already initialized via ContentProvider at process start, before any consent screen. Transport: the raw scan uploads to Adobe's Document Cloud API for OCR, and for AI Assistant use, into Adobe's GenAI ingestion pipeline, both demonstrated in H1. Review: per H2, that content may be reviewed by a human reviewer under a generic abuse-prevention policy with no third-party-data exclusion evidenced. Consequence for the pictured third party: their name, face, signature, or ID number now exists inside Adobe's cloud, classified, potentially analysed by a GenAI pipeline, and potentially seen by a human reviewer, with no notice ever sent to them and no channel of their own to request its deletion, because Adobe's own Privacy Policy frames the account holder, not Adobe, as the party responsible for having put it there.

Data Protection: Article by Article

- Art. 5(1)(a)/(b)/(c): lawfulness, purpose limitation, and minimisation are all engaged by processing a non-account-holder's document content with no lawful basis of its own (H1).
- Art. 6(1): Adobe needs its own lawful basis for processing the personal data of a person who is not its account holder and never interacted with Adobe (H1).
- Art. 7(1): the pre-consent ad/analytics SDK initialization in C1 processes before any consent decision exists.
- Art. 9: where a scanned document is a special-category document (a medical form, for instance), Adobe's processing obligations are correspondingly higher (H1).
- Art. 14: the business-card owner or ID-holder was never given a privacy notice by Adobe, because personal data was not obtained from them directly (H1).
- Art. 32(1)(a)/(b): no certificate pinning on first-party document-upload endpoints, and an unconfirmed server-side restriction on the hardcoded Firebase key (M1, H3).
- Art. 35: a Data Protection Impact Assessment question stands open for the ID-card capture mode specifically, given its designed, marketed third-party-data exposure (H1, H2).

Regulatory & Legal Landscape

Adobe Systems Software Ireland Limited is a genuine EU main establishment acting as controller for EU/EEA users, so under the GDPR one-stop-shop mechanism the lead supervisory authority for a cross-border complaint is the Irish Data Protection Commission, not the Austrian DSB. As an Austrian-initiated matter, RFI-IRFOS still addresses and copies the Austrian DSB and the EDPS per standing practice; the DPC is named here so this report does not misstate the lead authority. Adobe's own Generative AI User Guidelines and central Privacy Policy, both quoted verbatim above, are first-party sources, not inferred or reconstructed.

Indicators of Compromise & Reproducible Evidence

Artifact	Value
Package	com.adobe.scan.android
Base APK SHA-256 (26.08.10)	0bb8784730e79e167cd3f559026675ed2d072152c0e291727bdf233e176b22c7
Bundle SHA-256 (adobescan.apkm)	98ebf766e927ecb061174dc8aef58fc540c4ec8e39b1d51b7783320888b36053
Firebase API key	AIzaSyBUUctoBrbNnn4ql4WcfftiLm3yRykH8M0
Firebase project	adobe-scan-4df73 (app ID 1:827597488311:android:b2137c31aea0d481, sender ID 827597488311)
Document Cloud upload endpoints	dc-api.adobe.io, dc-asset-transfer.adobe.io
GenAI ingestion endpoints	documentcloud.adobe.com/genai_services/schemas/ingestion_v1.json and related

Disclosure Timeline & Outcome

23 August 2026: original notice to DPO@adobe.com and Grp-adobescan-android@adobe.com, four findings and three questions, with the Austrian DSB and EDPS in copy. 12 September 2026: twenty days, no reply of any kind, RFI-IRFOS restates the findings and questions unchanged. 18 September 2026: Adobe's Data Protection Officer replies, confirms receipt and that the matter is being handled as a coordinated disclosure by its security and privacy teams, and asks seven specific clarifying questions about the evidence base. RFI-IRFOS answers all seven in full the same day, naming the exact SHA-256 of the build analysed and the specific manifest entries, endpoints, and API client classes behind each finding. 17 October 2026: RFI-IRFOS sends a 35-days-to-publication notice. 15 November 2026: RFI-IRFOS sends a further notice confirming publication on 21 November 2026. 20 November 2026, 14:22 UTC: Adobe's Data Protection Officer answers the three original questions in writing, with the Adobe Scan group address in copy, and invites network-traffic captures. 21 November 2026: embargo reached, this report published. Adobe's statements in the reply are stated by the operator, not verified by RFI-IRFOS.

What Adobe Got Right

- A genuine EU establishment and a named DPO: Adobe Systems Software Ireland Limited (Dublin) is the controller for EU/EEA users, with a real postal address and DPO@Adobe.com, so there is no Art. 27 representative gap, unlike most targets in this programme.
- android:usesCleartextTraffic="false" is explicitly set, and android:allowBackup="false" blocks adb backup exfiltration of local app data.
- A real, functioning opt-out exists for Adobe's own Experience Cloud SDK: sendOptOutHit/handleOptOut demonstrably clears queued Analytics hits and stored state data on opt-out.
- No Chinese/PRC or other sanctioned-jurisdiction SDK was found anywhere in the binary, a meaningful and explicitly tested-for absence.
- A per-file protection_status field exists in the local scan database, a genuine password/encryption option per scanned document.
- The in-app GDPR/cloud-upload disclosure is honest about the fact of cloud upload, even where it is not consistently reflected in the Play Store's own Data Safety self-declaration.
- The AI Assistant carries a visible "AI responses may be inaccurate" disclaimer, a reasonable EU AI Act Art. 50-adjacent transparency practice.
- A prompt, substantive acknowledgement from Adobe's Data Protection Officer within four weeks, with specific, well-targeted clarifying questions rather than a form response.

Residual Risk & Recommendations

The technical findings in this report stand as disclosed on 23 August 2026. Adobe's Data Protection Officer engaged on 18 September 2026 with seven clarifying questions, which RFI-IRFOS answered in full that day, and answered the three original questions on 20 November 2026, 63 days later and 89 days after the original notice. The announced consent gating for Meta Audience Network and InMobi, and the planned change for Branch.io, would address C1 if confirmed in a shipped build; RFI-IRFOS has not verified them. RFI-IRFOS continues weekly APK diffing against the archived 26.08.10 build; any change to the ID-card/business-card pipeline, the consent-timing architecture, the hardcoded key, or the Play Store Data Safety declaration will be recorded against this correspondence record.

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

ZVR 1015608684 • GISA 39261441 • UID ATU83405245 • Steuernummer 68 696/8736 • Elisabethinergasse 25/10, 8020 Graz, Austria Root level code analysis of the published Google Play release bundle plus the operator's own public legal documents and written correspondence. No Adobe servers, accounts, or API keys were probed or tested. Research conducted under ISO/IEC 29147 and the freedom of scientific research (Art. 17 StGG). REF ADOBESCAN-2026-R1.