Coordinated Privacy Disclosure, Final Report

BetterHelp (com.betterhelp)

BetterHelp, Inc., San Jose, California, USA, a Teladoc Health subsidiary

SILENT · CASE CLOSED AND PUBLISHED 24 SEPTEMBER 2026, 94 DAYS AFTER DISCLOSURE, FOUR MESSAGES SENT, ZERO REPLY FROM BETTERHELP AT ANY POINT

Target	BetterHelp, the Android client for BetterHelp's online talk-therapy platform
Package	com.betterhelp
Operator	BetterHelp, Inc., 3155 Olsen Dr, San Jose, CA 95117, USA, a Teladoc Health subsidiary
Initial Disclosure	2026-06-22
Original Embargo	2026-09-19 (90 days from disclosure, as stated in the original notice)
Report Published	2026-09-24, five days after the embargo, an internal publishing gap, not a change in terms
Regulators in CC/BCC	Austrian DSB, CERT.at, EDPS, US Federal Trade Commission (ftc-opa@ftc.gov)
Total Outbound Messages	4, across 94 days, plus one delivery-failure bounce on the first send attempt, corrected the same day
Inbound Replies	0 from BetterHelp. The only automated reply in the thread was a mail-delivery-failure notice, not a company response.
Companion Case	TeenCounseling (com.teencounseling) and Regain (us.regain) were named in the original bundled disclosure as sharing the same codebase and SDK stack. Both are tracked as separate cases on RFI-IRFOS's public ledger and are not independently re-verified in this report; see "Scope Note on TeenCounseling and Regain" below.

rfi.irfos@gmail.com • rfi-irfos.com

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

Contents

Executive Summary

On 22 June 2026, RFI-IRFOS disclosed to BetterHelp, Inc. that its Android app, com.betterhelp, still ships the plumbing for third-party advertising data sharing, a hardcoded Firebase API key, pre-consent SDK auto-initialization, and a full ad-attribution stack (AppsFlyer, Google's Privacy Sandbox attribution permissions, the Play Install Referrer API, and the Iterable marketing-automation SDK bound to the user's real email address), three years after the US Federal Trade Commission fined the company $7.8 million (FTC File No. 2023169, March 2023) for sharing consumers' mental-health-context data with Meta, Snapchat, Pinterest, and Criteo for advertising purposes.

A deep-dive re-audit conducted during this disclosure, re-downloading the identical binary (SHA-256-verified byte-for-byte match to the original pass) and running a full smali decompile instead of a resources-and-strings-only pass, produced two corrections that matter and are credited here plainly. First, no Meta, TikTok, Pinterest, or Snapchat SDK class exists anywhere in the com.betterhelp code tree, confirmed at the strongest evidentiary level this methodology supports, a genuine architectural improvement over the pre-2023 pattern the FTC investigated, though a dangling, non-placeholder Facebook Application ID remains bound in the manifest with no client-side code to consume it (NG1, held as an open question, not an assertion). Second, a real, functioning, fail-closed consent gate was traced in bytecode for both AppsFlyer's SDK bootstrap and Firebase Analytics' collection-enable flag, sourced from a OneTrust cookie set by BetterHelp's own website and relayed through a shared WebView cookie jar, a genuine positive credited in full below.

That consent gate cannot reach the one place it would matter most: Android's ContentProvider initialization, including FirebaseInitProvider's own bootstrap, fires before any Activity exists and therefore before the WebView that carries the OneTrust cookie has ever loaded a page. The pre-consent processing finding (C2) is unaffected by the gate and stands at its full severity. Four messages were sent to security@betterhelp.com, privacy@betterhelp.com, and legal@betterhelp.com across 94 days, cc'ing press@betterhelp.com, the Austrian DSB, CERT.at, and the EDPS, and bcc'ing the FTC's Office of Public Affairs. Not one reply, automated or human, was ever received from BetterHelp, Inc.

Scope: Root-level static analysis of the production BetterHelp Android application (com.betterhelp, version 3.09, versionCode 309), conducted in two passes: an initial apktool resources-and-manifest decode with dex-string grep analysis (6 September 2026), and a full smali baksmali re-decompile of all three classes.dex files against a freshly re-downloaded, SHA-256-verified-identical copy of the same binary (6 September 2026), used to trace the analytics dispatch and consent-gate call graph at the bytecode level. Both APK copies and the full decompiled tree were deleted immediately after the relevant evidence was captured. No dynamic instrumentation, live traffic capture, or interaction with BetterHelp's backend infrastructure was performed.

Disposition

Five findings were disclosed on 22 June 2026, later re-scored under CVSS v4.0: one MEDIUM (a hardcoded Firebase API key, an availability-only exposure, corrected from an original CRITICAL label), two findings held at HIGH under an explicit blast-radius escalation (a Firebase bootstrap sequence that initializes before any consent screen can appear, at the exact moment a person first considers disclosing mental-health information; and a full ad-attribution stack, AppsFlyer, the Play Install Referrer API, and Google's Privacy Sandbox attribution permissions, present three years after BetterHelp, Inc.'s 2023 FTC settlement for sharing therapy-context data with advertisers), and one finding corrected down to LOW. A fifth, content-rating/age-gate mismatch, and a sixth, a dangling Facebook Application ID binding with no client-side SDK to consume it, are both classified OBSERVATIONAL, real and worth recording but not demonstrated technical attack paths. Four messages were sent across 94 days. Not one reply was ever received from BetterHelp, Inc.

ID	Severity	Title
C1	MEDIUM	Firebase API Key Hardcoded in the Distributed Binary
C2	HIGH	Firebase Initializes Before Any Consent Screen, at the Moment of a Person's First Contact With a Mental-Health Service
H1	HIGH	Full Ad-Attribution Stack Present Three Years After BetterHelp's 2023 FTC Health-Data-Sharing Settlement, Partially Mitigated by a Confirmed Fail-Closed Consent Gate
H2	LOW	Android Auto-Backup Enabled Without a Demonstrated Full Exclusion List for Therapy-App Data
M1	OBSERVATIONAL	Content Rating Lists 'Everyone' With No Client-Side Age-Verification Artifact Found
NG1	OBSERVATIONAL	A Dangling Facebook Application ID Remains Bound in the Manifest, With No Client-Side SDK Left to Consume It

Subject & Operator Analysis

BetterHelp (com.betterhelp) is the Android client for BetterHelp's online talk-therapy platform, operated by BetterHelp, Inc. of San Jose, California, part of Teladoc Health's consumer mental-health portfolio. This is the same company the US Federal Trade Commission fined $7.8 million in March 2023 (FTC File No. 2023169) for sharing consumers' health data, including mental-health-history questionnaire answers, with Meta, Snapchat, Pinterest, and Criteo for advertising purposes, without adequate disclosure or consent.

Several genuine positives are worth naming up front. A full smali sweep of the entire code tree found zero Meta, TikTok, Pinterest, or Snapchat SDK classes anywhere in com.betterhelp, confirmed at the strongest evidentiary level this methodology supports, a real architectural improvement over the pre-2023 pattern the FTC investigated. A real, functioning, fail-closed consent gate was traced in bytecode for both AppsFlyer's SDK bootstrap and Firebase Analytics' collection-enable flag. network_security_config.xml blocks cleartext traffic by default. And a named, reachable EU/UK Data Protection Officer structure exists (Ametros Group, Cork and Hereford), not a dead-end contact form.

Scope Note on TeenCounseling and Regain

The original 22 June 2026 disclosure bundled com.betterhelp together with com.teencounseling (TeenCounseling) and named us.regain (Regain) as a related third product sharing the same Facebook Application ID, on the basis that all three share a common codebase and SDK stack. This closing report covers com.betterhelp only, the app independently re-verified at the bytecode level for this disclosure. TeenCounseling and Regain are tracked as separate cases on RFI-IRFOS's public ledger, disclosed and followed up on their own correspondence threads, and are not re-asserted here beyond what was independently confirmed for com.betterhelp specifically. Claims made about TeenCounseling in the original bundled notice, including an alleged Facebook SDK reference in its manifest, were not re-verified in this pass and are stated here as originally reported, not independently confirmed.

Architecture & Data Flow

The consent architecture found in com.betterhelp is real but structurally cannot reach the point that matters most. A OneTrust cookie set by BetterHelp's own website, relayed through a shared WebView cookie jar, does gate AppsFlyer's SDK bootstrap and Firebase Analytics' collection-enable flag in a fail-closed manner, confirmed by tracing the call sites in bytecode (H1). What it cannot gate is FirebaseInitProvider itself: Android initializes ContentProviders before any Activity exists, and therefore before any WebView has had the chance to load the page that sets the governing cookie (C2). A hardcoded Firebase key (C1) sits underneath both at the operational-security layer, and a dangling Facebook Application ID with no client-side consumer (NG1) remains an open question for the operator to resolve.

Findings

C1: Firebase API Key Hardcoded in the Distributed Binary

MEDIUM, 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, 6.9.

The Firebase project API key and database URL are present verbatim in every distributed copy of com.betterhelp, extractable with standard tooling in under a minute. Google's own documentation states Firebase API keys are not secrets by design, this finding is about availability, not confidentiality: a party holding the key can exhaust the project's Firebase quota, taking backend services offline for a platform whose users may be mid-crisis and relying on messaging or scheduling continuity with their therapist.

```
res/values/strings.xml
<string name="firebase_database_url">https://betterhelp-141121.firebaseio.com</string>
<string name="google_api_key">AIzaSyAg5YDfsF3fviLetDITUdq2wzvpdVla4mA</string>
<string name="project_id">betterhelp-141121</string>
```

Impact: Quota exhaustion or abuse of the shared Firebase backend can degrade or take offline messaging and scheduling continuity for users in active therapy relationships, a service context where availability failures carry more weight than in an ordinary consumer app.

Fix: Rotate the exposed key, apply Firebase App Check and API key restrictions, and route database access through an authenticated backend layer instead of direct client SDK access. No confirmation that this has occurred was ever received.

C2: Firebase Initializes Before Any Consent Screen, at the Moment of a Person's First Contact With a Mental-Health Service

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. RFI-IRFOS blast-radius escalation: raised from MEDIUM/6.9 to HIGH because this processing occurs at the precise moment a person is deciding, for the first time, whether to trust a mental-health platform with sensitive information, a context CVSS's generic timing metric does not weight.

FirebaseInitProvider is declared with initOrder="100" and directBootAware="true", meaning it initializes before any Activity, including any consent dialog, can render, and before device unlock. No consent-management-platform SDK exists in the app itself; the only consent mechanism found (see H1) is a native cookie parser reading a OneTrust cookie set by BetterHelp's own website, relayed through a shared WebView cookie jar, a mechanism that structurally cannot gate a ContentProvider that has already run by the time any Activity, and therefore any WebView, exists. RFI-IRFOS blast-radius escalation: raised from MEDIUM/6.9 to HIGH because this processing occurs at the precise moment a person is deciding, for the first time, whether to trust a mental-health platform with sensitive information, a context CVSS's generic processing-timing metric does not weight.

```
AndroidManifest.xml
<provider android:authorities="com.betterhelp.firebaseinitprovider"
    android:directBootAware="true" android:exported="false"
    android:initOrder="100" android:name="com.google.firebase.provider.
    FirebaseInitProvider"/>
```

Impact: App-open telemetry, a weak but real signal of mental-health-seeking behavior, is generated and transmitted to Google before the person opening the app has seen a single privacy notice.

Fix: Remove directBootAware and defer FirebaseInitProvider's initOrder until after an app-native consent flow has run; a website cookie relayed through a WebView cannot substitute for a gate at process start. No confirmation that this has occurred was ever received.

H1: Full Ad-Attribution Stack Present Three Years After BetterHelp's 2023 FTC Health-Data-Sharing Settlement, Partially Mitigated by a Confirmed Fail-Closed Consent Gate

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:L/SI:N/SA:N, 6.9. RFI-IRFOS blast-radius escalation: raised from MEDIUM/6.9 to HIGH because BetterHelp, Inc. was already fined $7.8M by the FTC in 2023 for this exact category of conduct, standing regulatory history a bare CVSS vector does not capture.

The binary declares Google's Privacy Sandbox ACCESS_ADSERVICES_ATTRIBUTION and ACCESS_ADSERVICES_AD_ID permissions, bundles the AppsFlyer attribution SDK, uses the Play Install Referrer API, and binds the Iterable marketing-automation SDK to the user's real email address via setEmail. BetterHelp's own privacy policy names Google Click ID, Google Advertising ID, IDFA, Meta Click ID, and TikTok Click ID as processable 'External Identifiers' for opt-in advertising sharing, the same category of data flow the FTC's 2023 order (File No. 2023169) addressed. A bytecode trace during this disclosure found a genuine, fail-closed mitigation for part of this stack: AppsFlyerLib.init() and .start() fire only when a native cookie parser confirms an explicit 'targeting' consent signal from a OneTrust cookie set on BetterHelp's own website; when that cookie is absent, expired, or never set, the call is skipped entirely, a real, statically-demonstrate d fail-closed gate, credited here in full. This mitigates but does not resolve the finding: the underlying architecture, an attribution and marketing-automation stack wired into a mental-health app's initialization path, remains present in a binary shipped three years after a federal consent order specifically addressed this category of data sharing, and the gate depends on a WebView having loaded a specific web page, not a guaranteed event on every session. RFI-IRFOS blast-radius escalation: raised from MEDIUM/6.9 to HIGH because BetterHelp, Inc. was already fined $7.8M by the FTC in 2023 for this exact category of conduct, standing regulatory history a bare CVSS vector does not capture.

```
AndroidManifest.xml
uses-permission: android.permission.ACCESS_ADSERVICES_ATTRIBUTION
uses-permission: android.permission.ACCESS_ADSERVICES_AD_ID
uses-permission: com.google.android.finsky.permission.
    BIND_GET_INSTALL_REFERRER_SERVICE

smali/f3/a.smali: AppsFlyerLib.init("fvHEhEU8qxEEZZWo5uBe6X", ...) / .start(...)
# fires only when MainActivity.d1("targeting") confirms a OneTrust
# OptanonConsent cookie (C0004:1) read via the shared WebView CookieManager
# absent/expired cookie -> call is skipped (fail-closed, confirmed in bytecode)
```

Impact: An attribution and marketing-automation architecture matching the category the FTC already sanctioned this company for remains wired into the app's core, now partially, not fully, gated behind a website-cookie consent signal a user may never trigger within a given session.

Fix: Replace the website-cookie-relay consent mechanism with a native, app-level consent screen shown before any SDK initialization path, including the currently-ungated FirebaseIn itProvider bootstrap. No confirmation that this has occurred was ever received.

H2: Android Auto-Backup Enabled Without a Demonstrated Full Exclusion List for Therapy-App Data

LOW, CVSS v4.0 AV:P/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 2.4.

allowBackup=”true” is set at the application level. The only backup exclusion rules present are AppsFlyer’s own, excluding AppsFlyer’s attribution data specifically, not the app’s own local caches. This audit did not enumerate the app’s full local storage schema, which would require a live install and backup-extraction pass, so this is stated as an observed configuration risk, not a demonstrated data exposure: allowBackup=true with no app-authored exclusion rules for its own sensitive local state is a materially different posture from a purpose-built exclusion list.

```
AndroidManifest.xml
<application android:allowBackup="true"
    android:dataExtractionRules="@xml/appsflyer_data_extraction_rules"
    android:fullBackupContent="@xml/appsflyer_backup_rules" ... >
# only AppsFlyer's own attribution-data paths are excluded
```

Impact: If the app’s local cache holds session or authentication state, it plausibly rides along with Android’s general-purpose backup mechanism, a real gap, though not demonstrated in this static pass to include therapy content specifically.

Fix: Extend the existing data-extraction-rules exclusion list to cover session/authentica tion state and any locally cached therapist-matching data, following a documented data-classification exercise. No confirmation that this has occurred was ever received.

M1: Content Rating Lists ’Everyone’ With No Client-Side Age-Verification Artifact Found

OBSERVATIONAL — not independently CVSS-scored; a transparency and consumer-protection question, not a demonstrated technical attack path.

Play Store metadata lists BetterHelp under ’Parental guidance recommended’ while the mirrored APKPure listing shows Content Rating ’Everyone.’ No age-gate, date-of-birth field, or COPPA/parental-consent string was found anywhere in the decoded resources or dex string tables. BetterHelp’s own terms of service require adult users, a separate sibling product, TeenCounseling, exists for minors, yet no client-side technical mechanism enforcing that adult-only boundary was found within this app’s binary. A server-side age check at account registration cannot be ruled out from client-side evidence alone; this finding is stated as no client-side age-verification artifact was found, not as an assertion that no verification exists anywhere in the system.

Impact: A transparency mismatch between the app’s own adult-only terms of service and its public content rating, worth the company clarifying, short of a demonstrated technical failure.

Fix: Align the public content rating with the app’s actual adult-only terms of service, or document the server-side mechanism that enforces this boundary. No confirmation that this has occurred was ever received.

NG1: A Dangling Facebook Application ID Remains Bound in the Manifest, With No Client-Side SDK Left to Consume It

OBSERVATIONAL — not independently CVSS-scored; an open architectural question, not a demonstrated data flow.

The manifest binds a real, non-placeholder, numeric Meta/Facebook Application ID (740224816069682) alongside Facebook and Instagram package-visibility queries, yet a full smali sweep of the entire code tree, all three classes.dex files, found zero Facebook SDK classes anywhere to consume that ID. This is an open question, not an assertion: it may be a stale manifest-merger leftover from a previously integrated, since-removed Facebook SDK dependency, or a live identifier referenced exclusively by server-side Meta Conversions API postback matching that a client binary cannot show either way. Both explanations are consistent with H1's already-stated architecture. What is a demonstrated fact, not an inference, is that a real Meta Application ID remains wired into the current production manifest, three years after the FTC's 2023 settlement with this same company over Meta data sharing.

```
AndroidManifest.xml
<meta-data android:name="com.facebook.sdk.ApplicationId" android:value="@string/
    facebook_app_id"/>
res/values/strings.xml: <string name="facebook_app_id">740224816069682</string>
# full smali sweep, all 3 classes.dex: zero Lcom/facebook/ classes found
```

Impact: Whether this ID is inert or actively used server-side cannot be determined from the client binary alone, and is named explicitly so the operator, who can see the backend the client binary cannot, can resolve it directly.

Fix: Confirm and document whether the Facebook Application ID binding is a manifest-merger leftover or an active server-side reference, and remove it if unused. No confirmation that this has occurred was ever received.

Impact Analysis

C2 and H1 compound, they do not sit side by side: C2 establishes that Firebase's bootstrap fires structurally ahead of any consent mechanism the app has, and H1 shows that the one consent gate that does exist and does work, confirmed in bytecode, sits downstream of that same structural gap and depends on a website page having loaded in the current session. Both findings carry the same standing context: BetterHelp, Inc. already paid $7.8 million to US regulators for this category of data handling.

Four Messages, 94 Days, Zero Reply From BetterHelp

This case's correspondence record is silence, not deflection. Four messages were sent to security@betterhelp.com, privacy@betterhelp.com, and legal@betterhelp.com across 94 days, cc'ing press@betterhelp.com, the Austrian DSB, CERT.at, and the EDPS, and bcc'ing the FTC's Office of Public Affairs. The only automated reply in the thread was a mail-delive ry-failure notice on the first send attempt, corrected the same day, not a response from the company. No human or automated reply from BetterHelp, Inc. itself 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)(b), Art. 25	C1	Hardcoded Firebase credential, availability risk on a crisis-adjacent service's shared backend.
Art. 7(1), Art. 5(1)(a), Art. 9	C2	Consent-gated SDK bootstrap fires before any consent mechanism the app has can ever apply.
Art. 9, Art. 5(1)(c), Art. 44/46	H1	Ad-attribution architecture matching a category of conduct already subject to US federal enforcement, partially but not fully consent-gated.
Art. 32(1)(a)/(b), Art. 25	H2	Backup scope broader than demonstrated necessary for app-local sensitive state.
Art. 8 (conditional)	M1	Content-rating transparency mismatch against the app's own adult-only terms of service.

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 and regulatory history 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. Any such escalation is always stated explicitly in the finding itself, never left implicit, as C2 and H1 are escalated and C1 and H2 are corrected down in this report; M1 and NG1 are classified OBSERVATIONAL because they are real, transparency-relevant findings without a demonstrated technical attack path to score. Disclosure conducted under ISO/IEC 29147 and the Austrian Forschungsfreiheitsgesetz, Art. 17 StGG. REF REF: BETTERHELP-C1-R1.