Security & Privacy Disclosure, Final Report

Chat Me: Talk to Her AI for Android (com.biko.talkme.toyou)

An adult AI companion app starts its advertising and attribution stack at launch, sends chats to a US cloud, tells the store that nothing is encrypted and leaves its age check to one tap.

PUBLIC DISCLOSURE, 21 NOVEMBER 2026, EMBARGO REACHED. EMBARGO 19 NOVEMBER 2026. NO REPLY OF ANY KIND AS OF 21 NOVEMBER 2026.

Target	ZEOSTONE DISTRIBUTING LLC, California (Chat Me)
Versions Audited	2.3.3 / 51 (pulled 2026-08-20, SHA-256 17552e40…7f10). The listing shows an update on 24 October 2026, not analysed.
Notice / Publication	2026-08-21 / 2026-11-19 (90 days)
Findings	8: 6 Medium, 2 Low (reassessed against the notice, see status table)
Status	NO REPLY OF ANY KIND AS OF 21 NOVEMBER 2026. Four messages from RFI-IRFOS in 86 days, no delivery failure.

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

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

Contents

1 Executive Summary 3

2 Finding Status at a Glance 5

3 Subject & Operator Analysis 6

4 Architecture & Data Flow 6

5 Findings 7
- 5.1 H1: Advertising, attribution and mediation SDKs start at launch with a live Adjust event, and the only consent call sits in a Settings row . . . 7
- 5.2 M1: Chat content goes to a US cloud backend and the policy names no recipient, country or transfer safeguard . . . 8
- 5.3 M2: Advertising and mediation recipients are absent from the policy, among them two partners with China-linked infrastructure . . . 9
- 5.4 M3: Data safety label, privacy policy and build disagree on encryption, deletion, messages and location . . . 10
- 5.5 M4: Adult content is on by default behind a self-declared age, no age assurance method is documented or found in the build . . . 11
- 5.6 M5: Operator is a California company with a personal Gmail contact, no EU representative, and a user agreement with an unfilled date . . . 12
- 5.7 L1: Cleartext traffic is permitted for every domain in the network security configuration . . . 13
- 5.8 L2: Firebase key and SDK identifiers ship in the application resources . . . 13

6 Impact Analysis 14

7 The Complete Chain 14

8 Child Protection & Age Assurance 15

9 Data Protection: Article by Article 15

10 Consumer Protection & Deceptive Design 16

11 Platform Accountability & the DSA 16

12 Regulatory & Legal Landscape 16

13 Indicators of Compromise & Reproducible Evidence 17

14 Disclosure Timeline & Outcome 17

15 What the Operator Documents State Clearly 18

16 Residual Risk & Recommendations 18

Executive Summary

Chat Me: Talk to Her AI is an AI companion app for conversations and roleplay with virtual characters, with unlockable image packs and in-app purchases. Google Play shows more than 100,000 downloads, a PEGI 18 rating with the descriptors Users Interact, Shares Location and In-App Purchases, and the marker Contains ads. The legal entity on the developer panel is ZEOSTONE DISTRIBUTING LLC of San Juan Capistrano, California.

RFI-IRFOS analysed version 2.3.3 with root level code analysis and sent the findings under REF CHATME-2026-R1. The notice was delivered on 21 August 2026 to the support address shown on the Google Play page, with the Austrian data protection authority and the European Data Protection Supervisor in copy. A follow-up of 28 August, a reminder of 17 October and a final notice of 15 November 2026 went to the same address. No reply and no delivery failure is recorded as of 21 November 2026, 92 days after delivery.

The start-up code initialises Google Mobile Ads, Adjust and AppLovin MAX and fires a live Adjust event with no consent branch, and the only consent call found sits in a Settings row (H1). The chat backend runs on AWS in the United States, and the privacy policy names no recipient, country or safeguard (M1). Advertising partners, among them ByteDance Pangle and Mintegral, are absent from the policy (M2). The Data safety label says that data is not encrypted and cannot be deleted, while the policy promises both (M3). Adult content is on by default behind a self-declared age (M4). The operator's only contact is a personal Gmail address (M5).

The notice grouped its seven items as three critical, two high, one medium and one informational. This report supersedes that grouping with CVSS v4.0 ratings and lists the change item by item in the status table. The wording on the AppLovin consent flow is narrowed to the fact that no consent flow runs ahead of the initialisation, and the comparison with another target is not carried. The embargo of 19 November 2026 has passed, and this report is published on 21 November 2026 on the terms stated in the notice.

Scope: Root level code analysis of the Google Play release 2.3.3 (versionCode 51, base.apk, pulled 2026-08-20 from a test device): manifest, resources, network security configuration, string tables of all five dex files, and smali of the Application, Settings and launcher classes. A read-only permission dump of the installed package. Passive DNS and IP registry data for the hostnames found. Public pages read by plain HTTP GET on 14 November 2026: the Google Play listing and Data safety page, the privacy policy and the user agreement at chatmeai.top. No runtime tracing, no emulator, no backend call, no account. Findings that need runtime behaviour are marked UNVERIFIED.

ID	Severity	Title
H1	MEDIUM	Advertising, attribution and mediation SDKs start at launch with a live Adjust event, and the only consent call sits in a Settings row
M1	MEDIUM	Chat content goes to a US cloud backend and the policy names no recipient, country or transfer safeguard
M2	MEDIUM	Advertising and mediation recipients are absent from the policy, among them two partners with China-linked infrastructure
M3	MEDIUM	Data safety label, privacy policy and build disagree on encryption, deletion, messages and location
M4	MEDIUM	Adult content is on by default behind a self-declared age, no age assurance method is documented or found in the build
M5	MEDIUM	Operator is a California company with a personal Gmail contact, no EU representative, and a user agreement with an unfilled date
L1	LOW	Cleartext traffic is permitted for every domain in the network security configuration
L2	LOW	Firebase key and SDK identifiers ship in the application resources

ID	Notice item	This report	Status on 21 November 2026

Finding Status at a Glance

ID	Notice item	This report	Status on 21 November 2026
H1	Item 1 (pre-consent SDKs, CRITICAL)	HIGH	Open. Order of calls established, runtime effect UNVERIFIED. Wording on the AppLovin flow narrowed. No reply.
M1	Item 2 (AWS backend, CRITICAL)	MEDIUM	Open. Hosting by passive DNS, content flow inferred from endpoint names. No reply.
M2	Item 2 (Pangle, Mintegral)	MEDIUM	Restated as undisclosed advertising recipients. Ad telemetry only.
M3	New	MEDIUM	Open. Label and policy contradict each other on encryption and deletion. No reply.
M4	Item 5 (age, HIGH)	MEDIUM	Restated on the documents and the build. Server rules unobserved.
M5	Item 3 (operator, HIGH)	MEDIUM	Open. Personal Gmail contact, no representative, unfilled date. No reply.
L1	Item 6 (cleartext, MEDIUM)	LOW	Reassessed. No cleartext request observed.
L2	Item 7 (identifiers, INFORMATIONAL)	LOW	Reassessed. Values printed in full.
-	Item 4 (intimacy mechanic, shared template)	not carried as a finding	A product observation. The comparison with another target is circumstantial and carries no privacy consequence of its own.

Subject & Operator Analysis

The Google Play developer panel names ZEOSTONE DISTRIBUTING LLC, 31896 Plaza Dr Ste E5, San Juan Capistrano, California, with the telephone +44 7915 273738 and the support address mahmoudelbakry474@gmail.com. The developer account name on the listing is Services app. The privacy policy and the user agreement at chatmeai.top list the same Gmail address as their only contact. The domain chatmeai.top was registered on 2025-05-19 through NameSilo with the registrant redacted.

The app is offered in Austria, so Art. 3(2) GDPR applies to a controller established in California and Art. 27(1) requires a representative in the Union. No establishment or representative in the Union appears in the documents, so every authority where users are affected is competent. The Austrian Datenschutzbehörde and the European Data Protection Supervisor received the first notice in copy.

Architecture & Data Flow

The app is a native Android build, package com.biko.talkme.toyou, version 2.3.3, minimum SDK 24, target SDK 35, with the Java namespace com.android.chatme. The backend api.chatmeai.top serves the chat endpoints /honey/bot/chat_new and /honey/bot/chat_stream_refresh from AWS in the United States. Advertising and measurement run through Google Mobile Ads, AppLovin MAX with Pangle and Mintegral as mediation partners, Adjust for attribution, the Facebook SDK and Firebase Analytics and Crashlytics. Google Play Billing version 7.1.1 handles purchases. A local Room database holds messages and a profile with an intimacy value that gates image unlocks.

The manifest declares no microphone, camera, location or biometric permission, and a read-only dump of the installed package matched the manifest. The app has no voice feature that was found. The Google User Messaging Platform library is present with working code, and its only call site found is the Settings fragment.

Findings

H1: Advertising, attribution and mediation SDKs start at launch with a live Adjust event, and the only consent call sits in a Settings row

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)

RFI-IRFOS blast-radius escalation: raised from MEDIUM/6.3 to HIGH because the launch-time calls tie an advertising identifier to an app whose use alone says something about a person's intimate life, because three advertising and measurement SDKs sit behind the same start-up method, and because the only consent call in the reviewed code sits in a Settings row that onboarding never points to. The Application.onCreate() chain of version 2.3.3 calls MobileAds.initialize, creates the Adjust configuration and fires the Adjust event 52rtvm, then initialises AppLovin MAX, with no consent-status branch in the method.

```
smali_classes5/o6/a.smali              BaseApplication.onCreate(): process-name guard,
     then calls b()
smali_classes1/com/android/chatme/ChatMeApplication.smali   b(): no consent-status
    branch
  MMKV init (local storage)
  MobileAds.initialize(context)
  AdjustConfig(app token 1l82ch5dice8, "production")  Adjust.onCreate
  AdjustEvent(token 52rtvm)                           Adjust.trackEvent   fired at
    launch
  AppLovin MAX init, setMediationProvider("max"), terms and privacy flow setEnabled(
    false)
```

```
smali_classes3/com/android/chatme/ui/fragment/p0$b.smali   click listener of one
    Settings row
  ConsentInformation.requestConsentInfoUpdate (Google User Messaging Platform)
SplashActivity.smali (launcher): no consent API reference
```

The finding rests on the order of calls in the code that was read in full: Application.onCreate(), the b() method, the Settings fragment and SplashActivity. The file classes4.dex was reviewed by string search only, so a consent call in code that was not disassembled is possible and UNVERIFIED. Whether the SDKs transmit before any consent screen needs a runtime trace. The AppLovin setEnabled(false) call switches off AppLovin's own consent flow. Whether that differs from the SDK default was not established. The notice of 21 August 2026 described this call as the app turning off its ad SDK's consent screen. This report states the narrower fact that no consent flow runs ahead of the initialisation.

Section 165(3) TKG 2021 requires consent for storing or reading information on the terminal equipment of a user, unless strictly necessary for a service the user requested. Advertising and attribution identifiers are not strictly necessary for chat.

Impact: Adults who open an app for adult AI companions have no chance to decline before their device is registered with an attribution service and an advertising network. Their use of such an app becomes part of an advertising profile.

Fix: Start no advertising, attribution or mediation SDK before a recorded consent choice, show the consent form in onboarding, and fire attribution events only after consent.

M1: Chat content goes to a US cloud backend and the policy names no recipient, country or transfer safeguard

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:P/PR:N/UI:P/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N, 6.0 (CVSS-B)

The app talks to api.chatmeai.top, with the endpoints /honey/bot/chat_new and /honey/bot/chat_stream_refresh. Public DNS resolves the host to 34.205.204.162, which belongs to Amazon Technologies Inc. (AWS us-east-1, Ashburn, Virginia). The privacy policy lists chat interactions as collected and lists cloud service providers in general terms. It names no provider, no country and no safeguard.

```
api.chatmeai.top  ->  34.205.204.162   OrgName Amazon Technologies Inc., NetName AT
    -88-Z, US
www.chatmeai.top  ->  44.217.127.56    Amazon Data Services Northern Virginia
chatmeai.top registered 2025-05-19 via NameSilo, registrant redacted
Policy, Data Sharing: "Cloud service providers" / "Payment processors (via app stores)
    " /
  "Analytics services (for internal app optimization)" / "Legal authorities (if
    required by law)"
Policy: the words transfer, abroad, SCC and adequacy do not occur
```

Hosting by a US provider can be lawful. Art. 46 or an adequacy decision under Art. 45 then has to carry the transfer, and Art. 13(1)(f) requires the policy to say so. The backend was not called and no request content was observed. That the chat text goes to this host is inferred from the endpoint names. The notice of 21 August 2026 stated it as a fact. The geography rests on passive DNS and IP registry data alone.

Conversations with an AI companion that is marketed for roleplay with adult content can carry sexual preferences, health and relationship details. Art. 9 is not asserted for the content as such.

Impact: People who confide in an AI character, including in sexual roleplay, are not told which company stores the conversation or in which country.

Fix: Name the hosting provider and region in the privacy policy, state the Art. 45 or Art. 46 basis, conclude a processor contract under Art. 28 and state the retention of chat content.

M2: Advertising and mediation recipients are absent from the policy, among them two partners with China-linked infrastructure

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)

The Google Play page shows the marker Contains ads. The build contains Google Mobile Ads (application id ca-app-pub-1985964056515335~1291366486), Adjust, AppLovin MAX, Facebook SDK identifiers, ByteDance Pangle (version 7.5.0.2, endpoint log.byteoversea.com, creatives from sf16-static.i18n-pglstatp.com and p16-sign-sg.tiktokcdn.com) and Mintegral (com.mbridge.msdk, with untranslated Chinese resource strings). The privacy policy mentions analytics services for internal app optimisation and no advertising or attribution recipient.

Component	In the binary (2.3.3)	Named in the policy
Google Mobile Ads / AdMob	yes	no
Adjust (attribution)	yes, live event at launch	no
AppLovin MAX (mediation)	yes	no
ByteDance Pangle	yes, version 7.5.0.2	no
Mintegral	yes	no
Facebook SDK	yes, app id in resources	no
Firebase Analytics and Crashlytics	yes, project chat-me-d96da	”Analytics services” in general terms

The two China-linked partners are ad-mediation SDKs inside the AppLovin waterfall. They carry advertising identifiers and ad interaction data. The audit found no evidence that they carry chat content. Which of them are active at runtime, and where their servers stand for a user in the Union, was not observed and is UNVERIFIED. A sweep for NetEase, Tencent network code, Alibaba, Baidu and other China-based SDKs found none. The one Tencent class, MMKV, is a local storage library.

Impact: A user cannot learn from any document which advertising companies receive their device identifier, and they cannot tell that part of that traffic reaches ByteDance infrastructure.

Fix: List every advertising, attribution and mediation provider with country and purpose in the privacy policy, and state the Art. 46 safeguard for each transfer.

M3: Data safety label, privacy policy and build disagree on encryption, deletion, messages and location

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:P/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 5.3 (CVSS-B)

The label of 14 November 2026 states that data is not encrypted in transit and that the developer provides no way to request deletion. The privacy policy states that the operator uses encryption and offers deletion on request. The endpoints found in the build use https. The label lists precise location as shared for advertising, while the manifest and the live permission dump of the audited device show no location permission. The label lists no chat content among collected data, while the policy lists AI conversations and user inputs.

```
Shared: crash logs, diagnostics, other performance data (Analytics), approximate
    location (Analytics),
        precise location (Advertising or marketing), device or other IDs, user IDs
Collected: device or other IDs (optional), photos (optional), name (optional), email
    address (optional)
Security: "Data isn't encrypted"  "Data can't be deleted"
```

```
Policy, Data Security      "We use encryption, secure servers, and limited data access
    "
Policy, Your Rights        "Request access or deletion of your data" via
    mahmoudelbakry474@gmail.com
Policy, Information We Collect   "AI conversations and user inputs", IP address,
    device info, purchase history
Manifest                   no ACCESS_FINE_LOCATION, no ACCESS_COARSE_LOCATION
```

Each statement of the label can be wrong in either direction. The label may overstate a risk, as for location, or understate it, as for chat content. The state of the label on 21 August 2026 was not archived. This report rates the label as it stands on 14 November 2026. The manifest also permits cleartext traffic (L1), which fits the label statement on encryption, and every endpoint found in the binary was https. The content rating descriptor Shares Location on the listing falls under the same open question.

Impact: A person who reads the label is told that their data travels unencrypted and cannot be deleted, while the policy promises the opposite. Neither document tells them that their messages are collected.

Fix: Bring the Data safety declaration into line with the policy and the build: declare chat content, remove location entries the app cannot collect, and state the encryption and deletion practice that applies.

M4: Adult content is on by default behind a self-declared age, no age assurance method is documented or found in the build

MEDIUM, CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:P/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 5.3 (CVSS-B)

The listing rates the app PEGI 18 with the descriptors Users Interact, Shares Location and In-App Purchases. The policy collects an "Age declaration (to confirm 18+)" and states that the app is intended strictly for users aged 18 or older. The build string nsfw_switch_change warns that switching to Green Mode will block NSFW content, so adult content is reachable unless the user turns it off. A search of the resources and all five dex files found no date of birth field and no verification step. The only age input, "Age must be 18 - 1000", validates the age of a user-created fictional character.

```
Google Play listing   PEGI 18, Users Interact, Shares Location, In-App Purchases (14
    November 2026)
Privacy policy         "Age declaration (to confirm 18+)", "intended strictly for
    users 18 years or older"
User agreement         "You must be at least 18 years old to use this app", "
    Misrepresent your age" prohibited
Build strings          nsfw_switch_change ("Switching to Green Mode?"),
    report_tag_8_new ("Explicit Or Adult Content")
                       img_age_error "Age must be 18 - 1000."   (character creation,
    not the user)
```

Whether the server applies an age or content rule, and whether the first launch asks for an age, was not observed. Both are UNVERIFIED. The finding rests on the documents and the build. A single self-declaration is the age control described. No minor was observed using the app.

Impact: A thirteen or fourteen year old who taps through a declaration reaches sexual role-play content that the operator itself rates for adults only.

Fix: Switch adult content off by default, add an age assurance step before the first conversation and publish the method.

M5: Operator is a California company with a personal Gmail contact, no EU representative, and a user agreement with an unfilled date

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)

The Google Play developer panel names ZEOSTONE DISTRIBUTING LLC, 31896 Plaza Dr Ste E5, San Juan Capistrano, CA 92675-3724, United States, with the telephone +44 7915 273738. The developer account name on the listing is Services app. The sole contact in the policy, the user agreement and the Play panel is a personal Gmail address. The user agreement opens with "Effective Date: [Insert Date]", and it names no governing law and no dispute route. Neither document names a representative in the Union.

```
Google Play developer panel   ZEOSTONE DISTRIBUTING LLC, San Juan Capistrano,
    California, +44 7915 273738
Google Play developer account  "Services app"
Contact in policy, agreement and Play panel   mahmoudelbakry474@gmail.com
user-policy.html   "Effective Date: [Insert Date]"
EU representative: none   DPO: none   Postal address in the documents: none
```

The app is offered in Austria, so Art. 3(2) GDPR applies to a controller established in California, and Art. 27(1) requires a representative in the Union. None is stated. The app label in the manifest is Chat Me, and the name TalkMe in the case folder belongs to the shorthand only. At least two unrelated apps carry the title TalkMe on Google Play, so every reference must use the package com.biko.talkme.toyou and the title "Chat Me: Talk to Her AI".

Impact: A user in Austria who wants to see or delete their conversations has a personal Gmail address and no local representative, and the agreement they accepted carries no date.

Fix: Publish the controller's postal address and a business contact, appoint an Art. 27 representative, date the user agreement and name the governing law.

L1: Cleartext traffic is permitted for every domain in the network security configuration

LOW, CVSS v4.0 AV:A/AC:H/AT:P/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 2.3 (CVSS-B)

res/xml/network_security_config.xml sets base-config cleartextTrafficPermitted= ”true” with no domain scoping and no pinning, and the manifest sets android: usesCleartextTraffic=”true” on the application element. Both settings re-enable what target SDK 28 and above block by default. Every endpoint found in the build uses https, and no cleartext request was observed.

```
res/xml/network_security_config.xml:  <base-config cleartextTrafficPermitted="true" />
AndroidManifest.xml:  android:usesCleartextTraffic="true"
```

The permission opens a path for an attacker on the same network if any request is ever made over http, for example by an advertising SDK. The finding carries a high attack complexity for that reason. Rated LOW.

Impact: Users on shared Wi-Fi are exposed only if a cleartext request occurs, which was not observed.

Fix: Remove both cleartext settings and add a network security configuration that allows https only.

L2: Firebase key and SDK identifiers ship in the application resources

LOW, CVSS v4.0 AV:N/AC:L/AT:P/PR:N/UI:P/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 2.3 (CVSS-B)

The resources carry the Firebase API key AIzaSyBQnwQjYCRS8KCn5N2kSxiSEkYw4b_3ebg, project chat-me-d96da, sender 590751886796, the storage bucket chat-me-d96da.firebasestorage.app, Facebook app id 1538964724548443 with client token 5f96bf6dae176d622a14ea5b032b0b03, the AdMob application id ca-app-pub-1985964056515335~1291366486, the Adjust app token 1l82ch5dice8 and the AppLovin MAX SDK key SpnWfLTbXp8sdan9TB8hhKsNN4ZOcQuphV6WQEOiaWyVO1uCui8v1Qln9_mbSapfSrJzNbXJMDDJQcCrC5L92J. These values ship in every build by design. Restrictions on the keys and the security rules of the Firebase project were not tested, because no request was sent to any vendor endpoint.

The notice of 21 August 2026 states that these are publicly shipped identifiers and not secrets. The finding is limited to the open question of key restrictions and project rules. Rated LOW.

Impact: The risk reaches users only if the project rules are open, which was not tested.

Fix: Restrict the Firebase key to the application package and signing certificate, and review the security rules of the project.

Impact Analysis

The people affected are adults who use an AI companion for conversation and roleplay. The product is built for emotional closeness and sexual content, and it sells image packs and tokens. Users therefore bring private details into the conversation and spend money on the relationship.

Three points carry the weight. The device is registered with advertising and attribution services before any consent choice (H1). The conversation is stored in the United States by a host that no document names (M1). The age control is one tap on a declaration, while adult content is on by default (M4).

The remaining findings concern the completeness and consistency of the documents (M2, M3, M5) and hygiene items (L1, L2).

The Complete Chain

Step 1, the store page: a user sees PEGI 18, Contains ads and a Data safety section that says data is not encrypted and cannot be deleted (M3).

Step 2, first launch: the start-up code initialises Google Mobile Ads, Adjust and AppLovin MAX and fires an Adjust event (H1). The consent form sits behind a Settings row.

Step 3, age: one declaration confirms 18+, and adult content stays on until the user switches to Green Mode (M4).

Step 4, conversation: chat requests go to api.chatmeai.top at AWS in the United States (M1). The policy names a cloud provider in general terms only.

Step 5, advertising: Pangle, Mintegral, AppLovin, Adjust and Google receive identifiers and ad interaction data (M2).

Step 6, rights: a user who wants access or deletion finds a personal Gmail address, a California address on the store panel and no representative in the Union (M5).

The operator has not answered any of these points.

Child Protection & Age Assurance

The operator sets the age at 18 in the policy and in the user agreement, and the store rates the app PEGI 18. The only control described is the declaration ”to confirm 18+”. The build has no date of birth field and no verification step. The one age input validates the age of a fictional character that the user creates, and the same screen asks creators to ensure that the character is 18 or older.

A string in the build shows that adult content is reachable by default and has to be switched off. A minor who accepts a declaration can reach that content. Art. 8 GDPR sets the age for valid consent to information society services in Austria at 14 under section 4(4) DSG. A user under 14 would need the consent of a parent, and no document describes such a step. Whether the server applies any rule was not observed and is UNVERIFIED.

Data Protection: Article by Article

Article 5(1)(a), transparency: engaged by M1, M2, M3 and M4. The label, the policy and the build differ from each other.

Article 6(1)(a), Article 7 and Article 4(11), consent: engaged by H1. The order of calls in the start-up code shows no consent gate.

Article 8 and Article 25: engaged by M4 and H1.

Article 9: engaged only as a contingent point in M1. Content is not known to be special category data.

Article 12(1), Article 13(1)(a), (b), (c), (e) and (f): engaged by M1, M2, M3 and M5.

Article 17: engaged by M3. The label says deletion is not possible, and the policy says it is.

Article 27(1) and Article 3(2): engaged by M5. No representative is named.

Article 28(1): engaged by M1 for the hosting provider.

Article 32(1)(a) and (b): engaged by L1 and L2 as hygiene points.

Articles 44 to 46: engaged by M1 and M2. The documents name no third country and no safeguard.

Directive 2002/58/EC Article 5(3), implemented in section 165(3) TKG 2021: engaged by H1.

Consumer Protection & Deceptive Design

The build contains a monetised intimacy mechanic. The local Room tables CacheMessage (columns unlock, point, intimacy_value, surplus) and Profile (intimacy_value, can_get_image) and the pricing function NewPrice(intimacy) tie image unlocks to the intimacy value that grows with chat engagement. The user agreement lists unlockable image packs and AI tokens, states that purchases are non-refundable unless the law requires otherwise and excludes liability for emotional distress arising from roleplay. The design is common for the product category. It was not scored, and its effect on users was not measured. The Unfair Commercial Practices Directive 2005/29/EC applies to the commercial use of such mechanics.

The policy states that AI responses are fictional, offers reporting by e-mail and says that reports are reviewed within 48 hours. Those statements are specific. Their accuracy was not tested.

Platform Accountability & the DSA

Google Play presents the Data safety section with the words ”The developer says”. The declaration remains the developer’s statement, and the store page of 14 November 2026 shows the contradiction with the developer’s own policy. The Digital Services Act (Regulation (EU) 2022/2065) applies to the distribution platform. A report of an inaccurate label to the platform is a separate step from this disclosure.

Regulatory & Legal Landscape

ZEOSTONE DISTRIBUTING LLC has no establishment in the Union in the documents read. Every supervisory authority where users are affected is competent, and the Austrian authority and the European Data Protection Supervisor hold the first notice.

Chat Me belongs to a group of AI companion apps that RFI-IRFOS examines with the same method. The method reads the binary, the published documents and the store pages against each other.

Indicators of Compromise & Reproducible Evidence

Every point below is reproducible from the public Google Play build and the public pages named.

```
Package       com.biko.talkme.toyou
Version       2.3.3 / versionCode 51   minSdk 24, targetSdk 35   base.apk 53.3 MB
SHA-256       17552e402e000b3ebd93c4f7814fe708fb23cd83642321583b90c62730267f10
Firebase      project chat-me-d96da   sender 590751886796   bucket chat-me-d96da.
    firebasestorage.app
Backend       api.chatmeai.top (34.205.204.162)   www.chatmeai.top (44.217.127.56)
Adjust        app token 1l82ch5dice8   launch event 52rtvm
AdMob         ca-app-pub-1985964056515335~1291366486
Pangle        7.5.0.2   log.byteoversea.com
Start-up      smali_classes5/o6/a.smali -> smali_classes1/com/android/chatme/
    ChatMeApplication.smali b()
```

```
Google Play listing and Data safety   playstore_snapshot_2026-11-14.txt
chatmeai.top/privacy-policy.html      chatme_privacy_policy_2026-11-14.txt
chatmeai.top/user-policy.html         chatme_user_agreement_2026-11-14.txt
Permission dump of the audited device  dumpsys_package.txt
```

Reproduction for H1: open the ChatMeApplication class and read method b(). Reproduction for M3: open the Data safety page and the privacy policy and compare the statements on encryption and deletion. Reproduction for M4: search the string resources for nsfw_switch_change.

Disclosure Timeline & Outcome

2026-08-21 17:32 UTC Delivery of the notice REF CHATME-2026-R1 to mahmoudelbakry474@gmail.com, the support address on the Google Play page, copy to dsb@dsb.gv.at and edps@edps.europa.eu. The 90 day embargo runs from this delivery.

2026-08-28 First follow-up to the same address. No delivery failure was returned.

2026-10-17 Reminder from security@rfi-irfos.com to the same address with the three questions repeated and the embargo date of 19 November 2026.

2026-11-14 Listing, Data safety page, privacy policy and user agreement read on this date.

2026-11-15 12:58 UTC Final notice from security@rfi-irfos.com to the same address, four days before the embargo date, with the revised ratings and the three questions. No delivery failure was returned.

2026-11-19 Embargo date, 90 days after delivery of the notice of 21 August 2026.

2026-11-21 Publication of this report. No reply of any kind has arrived in 92 days, and no delivery failure was recorded.

The item grouping of the notices (three critical, two high, one medium, one informational) is superseded by the ratings in this report.

What the Operator Documents State Clearly

The policy and the agreement state the age limit of 18 in plain words. The policy offers access, correction and deletion on request, promises that personal data is not sold, discloses that chat interactions are collected for quality improvement and safety, and offers a reporting route for inappropriate AI responses. The build asks for a small set of permissions, with no microphone, camera, location or biometric access, and the permission dump of the device matched the manifest. The Google User Messaging Platform library is present with working code. The listing carries the PEGI 18 rating publicly.

Residual Risk & Recommendations

Open on 21 November 2026: all eight findings, because the operator has not replied. The runtime questions stay open. They are whether the SDKs transmit before any consent screen, whether code in classes4.dex calls a consent flow, which regions serve the Pangle and Mintegral traffic, what the server enforces on age, and whether messages leave the device over any cleartext connection. Builds since 2.3.3 were not analysed.

Recommendations: gate every advertising and attribution SDK behind consent shown at onboarding. Name the hosting provider, region and safeguard. List all advertising partners. Correct the Data safety declaration. Switch adult content off by default and add an age assurance step. Name an Art. 27 representative and a business contact.

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 application

plus public pages read by plain GET. No operator servers or accounts were probed. Final report, published 21 November 2026. ISO/IEC 29147, Art. 17 StGG, GDPR Art. 89. RFI-IRFOS, ZVR 1015608684, GISA 39261441. REF CHATME-2026-R1.