IDFLOW / LEGAL

IDFlow App Privacy Notice

Version 2.4 | Last updated 11 November 2026

At a glance

IDFlow connects artists and labels through music discovery, uploads, messages and collaboration tools. We use the information needed to provide those features, keep the service safe, manage subscriptions and answer requests. You choose the material you upload and the information you put in your profile. Profile pictures are optional.

Optional analytics, crash reporting and advertising measurement are off until you choose them. If you enable advertising or campaign measurement, Meta or AppsFlyer receives selected app activity and device identifiers for the advertising purposes explained below. You can use IDFlow without enabling these choices. A public profile, a message to another user and a private full-track file have different audiences.

You can manage your profile, discovery and notification settings in the app, start account deletion in Settings, or contact **privacy@idflowapp.com** about your information. Deleting an IDFlow account does not automatically cancel an Apple or Google Play subscription.

1. Who is responsible for your information

SPIRITBE TECHNOLOGY LTD, company number **16163151**, operates IDFlow and is the controller of the information described in this notice. Our correspondence address is **Flat 84, Meeches House, 3 Pavana Square, Brentford, England, TW8 0YH, United Kingdom**.

Our privacy contact is **Axel Yalmaz, Company Director**, at **privacy@idflowapp.com**. This is our privacy contact, not an appointment as a statutory Data Protection Officer. You can contact us without an account, including after a restriction or deletion.

Our representative under Article 27 of the EU GDPR is **Axel Yalmaz**, **13 Han Asparuh Street, Floor 6, Apartment 18, 4230 Asenovgrad, Bulgaria**; **info@idflowapp.com**; **+44 7780 949156**. You may contact the representative or our privacy contact about your personal information and rights. This appointment is separate from our Digital Services Act representative role.

This notice covers the IDFlow app and its account, invitation, support and connected service operations. It applies to artists, label representatives, applicants, invited users and other people whose information is supplied through those features. It also applies when you read it on our website. It does not cover separate advertising or browsing technologies on our marketing website.

2. The information we process and why

Account creation, sign-in and invitations

We process your email, authentication credentials, account identifier, chosen username, name, role, verification/recovery information and session records to create and secure your account. Email and username availability checks help prevent duplicate registration. Rate limits and connection information, including IP-derived security values, help prevent abuse. Our authentication service handles password hashes; we do not keep a separate profile field containing the password you retype to confirm it.

For invited accounts, we receive the email and invitation details from the person administering the invitation. We keep the account link, onboarding status, inviting staff reference and any applicable offer eligibility. An administrator-created account requires its temporary password to be changed before normal onboarding. An administrative verified-email flag does not mean that you previously completed an email-code check yourself.

Label applications contain the details, contact information and links you submit, together with application decisions and review history. Invited labels may use a pre-approved route. We use this information to assess eligibility, prevent impersonation and manage onboarding. Invitation details are explained when we first contact you or provide access to the account.

IDFlow is for people aged **18 or over**. Registration asks for an age confirmation; this is a self-declaration, not identity-document or date-of-birth verification. Contact us if you believe a child has supplied information through IDFlow so we can investigate and take appropriate action.

Profiles, photographs and contact details

Your profile can include an artist/label name, username, biography or description, country, music genres, image/background and social or website links. These details help other users identify you and discover your music or label. Photographs are optional; an account without one displays an initials avatar. Notifications that conceal a label's identity under the recipient's plan use a generic avatar.

Profile information and content you publish can be visible to other users. Publicly delivered media may also be accessible through its link outside the screen where you first shared it. Do not put information you want kept private in public fields.

Your sign-in email is not a public profile field. An artist can separately provide a business contact email, including by choosing the same address. Eligible labels can receive that business address at an authorised deal stage, subject to the relevant access checks. A recipient can keep its own copy.

The app uses tutorial progress, availability, read and typing information to support onboarding and conversations. Hiding activity status removes the exposed presence information; it does not disable the separate, minimal latest-use timestamp used for account lifecycle management.

Tracks, previews, artwork and released material

We process the audio, video, artwork and files you select, their filenames, size, duration and technical or embedded metadata, and the titles, captions, genres, moods, BPM and editing choices you provide. We use these to validate uploads, store masters, create playable previews and thumbnails, display your material and manage playback/access. Processing also creates file references, status, verification, retry, timing and diagnostic records.

Private masters, voice recordings and report evidence have different access rules from public previews and artwork. Full-track requests, grants, revocations and access events support controlled sharing.

Expiry of a download link or removal from a screen is not the same as erasure of a stored file. A recipient may already have downloaded an authorised copy.

Marking a track released does not delete its master. The service retains the master, prepared media, track metadata and necessary release/deal and access history while the owning account exists. Extra release-related viewing and processing diagnostics have a six-month limit from first release, subject to specific open cases or justified evidence holds. Minimal publication links and status can remain where needed to preserve the release record. These records preserve the released workflow; they are not permission for unrelated future reuse. Account deletion and applicable copyright or safety removal still apply. Release status does not grant new ownership rights, restore revoked access or guarantee permanent recovery of a file.

Messages, social activity and deals

Messages include participants, text, replies, shared-track references, voice recordings, duration/waveform information and delivery, read, edit or deletion state. We process these to deliver the conversation, play recordings, display previews and retry requested sends. Temporary recording segments or an assembled recording may exist on your device before sending.

Messages are **not end-to-end encrypted**. Authorised staff can access the information needed for support, safety and administration. Other participants can retain their own copies; closing an account cannot recall files or messages independently saved outside IDFlow.

Likes, follows, comments, views, interests, shortlists and blocks support the interactions you choose. Deal and label-inbox records include the track, participants, stage, decisions, permissions, deadlines, reminders and acceptance, decline, cancellation or release history. These help you manage the workflow; they do not certify copyright ownership or an external contract.

Some records, such as blocks and private feed exclusions, are used to apply your choices without being publicly displayed. Where a plan conceals a label's identity in interest/shortlist alerts, the backend still processes the actual label identity to manage the activity.

Notifications and purchases

We use notification events, actor/content references, read state, preferences, device tokens and delivery records to send relevant service updates. Likes on the same video and new followers may be grouped within a fixed ten-minute window; interests and shortlists remain separate. Grouping does not set the retention period.

On supported iPhones, Apple Push Notification service delivers permitted notifications. On supported Android devices, **Google Firebase Cloud Messaging (FCM)** delivers them. For delivery, we send the provider a device registration token, the notification title and body, and references needed to open the relevant account, conversation or content in IDFlow. Google also processes Firebase installation identifiers and the device/network information needed to operate its delivery service. A notification can include an actor name and a short text-message preview; voice messages use a generic description. These push payloads are processed by the delivery provider and are not end-to-end encrypted. They do not include the full music file or voice recording.

Notifications can arrive while the app is in the background or closed, subject to device settings and delivery availability. Content can appear on your lock screen according to device settings. You can change app notification categories and operating-system permissions separately. Service-notification delivery is separate from optional usage analytics, crash reporting and advertising measurement; enabling notifications does not enable those optional purposes.

Apple processes App Store checkout and Google processes Google Play checkout, including payment credentials. We verify purchases with the relevant store to apply a plan, restore access, acknowledge Google Play purchases, reconcile renewals, cancellations, refunds or revocations, prevent duplicate/conflicting claims and provide billing support. This includes the linked IDFlow account, product and transaction/order identifiers, Apple app-account identifiers or Google Play purchase tokens and account-linking identifiers, plan, price/currency and storefront where supplied, purchase environment, subscription state and relevant renewal, expiry, grace, hold and revocation dates. Purchase tokens are restricted verification records, not public profile data. IDFlow does not receive your card number or card security code through either store checkout. These necessary billing operations do not depend on optional analytics consent.

Introductory and other offers use eligibility, purchase/redemption state, store-account linkage and currency/storefront information. The store determines eligibility for an automatic introductory offer. We also retain relevant historical invited-artist offer claims and offer-code records where applicable. An offer's expiry is separate from retention of its claim record.

Safety checks, reports and support

We use **Google Cloud Vision** to check selected video frames and artwork for content safety, labels and visible text. **Google Cloud Natural Language** checks titles, captions and extracted text. The resulting classifications, decision information and file hashes help identify unsafe uploads and route uncertain cases for review. These checks can prevent publication or require staff review. They do not establish a person's guilt or determine copyright ownership. The upload safety pipeline does not transcribe audio or create voiceprints. Audio is not sent to these image/text safety APIs; Google-hosted processing workers separately handle media to prepare your preview.

Before using Google safety checks, the app explains the information sent and asks for your permission. You can choose human review instead. Change your choice in **Settings > Upload safety checks**. Declining or withdrawing permission routes future unapproved uploads to human review; it does not prevent you having an account or submitting an upload. Checks already sent cannot be recalled. We keep your current choice, disclosure version and the time it was saved. Choosing automated checks is not consent to unrelated data uses.

Reports, appeals and support requests can include your explanation, attachments, relevant content snapshots, contact/account details, decisions and staff notes. We can receive information about you from a complainant, another user or a counterparty. We use it to investigate the issue, explain decisions, handle appeals and protect users. Please provide only information relevant to the case and do not email copies of suspected child sexual abuse material; report its location instead.

You can challenge a moderation decision and ask for human review at **support@idflowapp.com**, including when signed out or banned. Automated classification is one input to moderation, not a determination that every flag is correct. Restrictions may affect publication, visibility, access or an account; we explain decisions and the available challenge route subject to lawful disclosure limits.

We keep the versions and acceptance date/time of Terms and other legal acknowledgements to establish which information was presented and accepted. Accepting Terms or acknowledging a Privacy Notice is not consent to every use of personal information.

Devices, reliability and security

We and our providers process connection/IP information, request metadata, authentication/security events, device or service identifiers, operating-system/app details where supplied, job status, errors and health/capacity information. We use these to authenticate requests, deliver the service, detect abuse and investigate faults.

The app stores session state, preferences, upload/recording files, retry queues and media caches on your device. Some cleanup occurs when you next open the app or complete a send, discard or sign-out action. Microphone access is used for recordings you choose to make; photo/media access is used for files you choose to upload. You can refuse or revoke those permissions, though the corresponding feature may then be unavailable. These permissions do not authorise unrelated processing.

A country selected in a profile and coarse connection-derived location are different from precise GPS location. The app does not request GPS or address-book access. Social connections and business contact information entered into IDFlow are not a copy of your phone's address book.

We may use necessary copies of service data in separate cloud test services, with access restrictions and deletion requirements. Those copies are stored in cloud services, not only on an individual device. We use a protected deletion record to prevent deleted accounts being restored through covered backup and test workflows. That record is pseudonymous personal data, not anonymous information.

Optional analytics, diagnostics and advertising measurement

In app versions that include these services, **Settings > Privacy preferences** provides separate choices. Accepting Terms, acknowledging this notice, subscribing or granting photo/microphone access does not enable them. Choices are off by default. You can continue without sharing and withdraw a choice in Settings. On iOS, Meta and AppsFlyer also require Apple's tracking permission; declining it keeps those integrations off.

• **Usage analytics - Google Analytics for Firebase:** measures app use and returning installations, including registration, tutorial progress, screens, discovery/playback actions, uploads, social interactions, messages being sent, deal stages and subscription activity. Events use limited categories, counts, timing and status values rather than the content of messages, music, search queries or track titles. Google also processes pseudonymous installation identifiers, app/device/OS information, session information and network-derived approximate geography. We do not set a Firebase user ID to your IDFlow account ID. Our backend temporarily links an opted-in installation to your account to report relevant verified subscription and service events. Firebase advertising-ID collection and advertising consent are disabled in this integration.

• **Crash diagnostics - Firebase Crashlytics:** receives crash/error reports, stack traces, app version, device/OS information and installation identifiers to diagnose faults. This choice is independent of usage analytics. IDFlow's own diagnostic events use fixed error categories; they do not attach your profile, music or messages. Unsent reports are removed when you turn diagnostics off or switch accounts. A report already transmitted cannot be recalled by the switch.

• **Advertising measurement - Meta:** receives app activation, registration, tutorial completion, subscription checkout and eligible verified purchase events, with app/device/network information and permitted advertising identifiers. Meta can match this activity with information from its services to measure and improve IDFlow's Facebook and Instagram advertising. IDFlow does not attach your name, email, music, message contents or Apple transaction ID to these events. Meta's matching means these events are not anonymous. Its processing is governed by its applicable Business Tools terms and privacy policy as well as our instructions where relevant.

• **Campaign measurement - AppsFlyer:** receives installation/session information and selected registration, tutorial, checkout and eligible verified purchase events, with device, network and permitted advertising identifiers. This helps identify which campaigns bring people to IDFlow. With this choice enabled, AppsFlyer and Meta exchange install/device attribution information to match Meta campaigns. This connection does not forward a second copy of IDFlow's in-app conversion events to Meta. IDFlow does not attach your name, email, messages or music. This choice is separate from Meta's direct event reporting.

• **Subscription analytics - RevenueCat (Apple and Google Play subscriptions):** our backend sends a pseudonymous IDFlow account identifier and verified purchase references: Apple receipt/transaction information or Google Play purchase tokens and product identifiers, as appropriate. RevenueCat uses these references and its authorised store connections to retrieve related purchase and subscription information, including products, price/currency, order/transaction identifiers, subscription status and later renewals, expiries and refunds, to help us understand subscription performance. IDFlow does not send your name, email, music, messages or advertising identifiers through this integration. The choice applies across your account. Turning it off stops new IDFlow reporting, removes pending receipt jobs and queues deletion of the RevenueCat customer record. Account deletion queues the same removal. Your Apple or Google Play subscription and IDFlow entitlements are unaffected.

Meta and AppsFlyer measurement is used to advertise IDFlow, not to display third-party advertisements inside IDFlow. We do not sell personal information for money. Disclosures to advertising partners can nevertheless fall within legal definitions of sharing for advertising or targeted advertising. Leave the relevant choices off, or withdraw them, to stop that optional sharing by IDFlow. You can also contact **privacy@idflowapp.com** to exercise applicable rights.

Withdrawal stops new optional collection by the relevant integration when the choice takes effect. It does not automatically erase historical provider records or undo information already sent. RevenueCat removal is queued and retried; other requests to delete previously collected information can be made through our privacy contact. We will seek only the identifiers needed to locate the relevant records. Device-scoped analytics identifiers do not always let us locate historical data using an email address alone. You can use the privacy controls in the relevant provider's services as well.

Some choices are stored on the device for the current IDFlow account, so another installation may need its own choice. RevenueCat's choice is stored centrally and applies across the account. A server-side save needs connectivity; if it fails, the app asks you to retry rather than confirming a successful change. These optional services are separate from the necessary authentication, security and billing operations described above.

3. Discovery, search and your choices

We process search terms, filters and page requests to return results. Saved discovery genres remain private preferences used by the feed. The service does not maintain a dedicated permanent search-history feature.

IDFlow's discovery feed uses the genres you choose and private delivery and repeat-prevention records to provide a varied feed, prioritise unseen tracks and avoid repeatedly showing the same artist or track. With active genre filters, only matching genres are eligible; this does not switch off the records used to reduce repetition. Without explicit filters, the feed mixes the genres on your profile with other eligible music. It does not learn a new set of preferred genres from what you listen to. There is no user-facing listening-history page. Tracks can reappear when the eligible catalogue runs low.

We keep the identity of cards delivered/displayed, their delivery and first-display times, and per-track last-display times/counts to continue the feed and reduce repetition. These records are linked to your account. If you enable **Use playback time**, we also collect how long feed previews play while the feed is active, including muted playback. We use that time only as a late tie-breaker between otherwise similar recommendations, giving less-played tracks priority. It is not proof that you heard a track. We do not record audible duration, mute status, loop/replay counts or attribute interests/full-track requests to feed appearances for analytics. Your actual interests and requests still exist as the actions you chose to make.

Settings > Discovery lets you switch personalised discovery off, separately allow or withdraw playback-time measurement, and clear discovery history. Playback-time measurement is optional and off until you explicitly enable it. The account-linked choice records its wording version and the time you enabled it. Turning it off stops acceptance of further playback measurements and erases the saved playback records and totals. Your feed and other account features remain available; you can enable it again later with a fresh choice. The non-personalised feed uses newest-first discovery with your saved genre filters. Switching personalised discovery off also withdraws playback permission, clears the applicable delivery, playback and repeat-prevention records, and stops collecting them. Clearing history alone keeps your discovery choices and starts fresh. Genre filters alone do not turn repeat prevention off. Clearing discovery history does not delete your separate tracks, likes, follows, messages or deals.

A bounded retry queue can temporarily hold unsent display markers and, only with playback permission, playback increments on your device for up to 48 hours. It is tied to your account and cleared by the relevant sign-out/reset action. Devices regularly recheck the saved choice; the server rejects measurements from withdrawn or superseded permission. Server delivery records and playback retry receipts are scheduled for deletion after 30 days; per-track discovery summaries, including optional playback totals, expire 90 days after the track last appeared. Withdrawal or history clearing removes the corresponding records sooner. These schedules run periodically, and provider backups follow section 7.

4. Our legal grounds

For people protected by UK or EU data-protection law, we use the following grounds for the corresponding purpose. We do not treat these as interchangeable reasons for keeping data.

Purpose	Legal ground
Creating and securing the account you request; publishing your chosen profile/content; delivering messages, interactions and collaboration features; purchases and requested support	Performance of our contract with you, or steps you ask us to take before a contract, insofar as that processing is necessary to provide the requested feature.
Administering a label relationship through its employees/representatives	Our legitimate interest in providing and managing that business relationship when the contract is with the organisation rather than the individual.
Proportionate security, abuse prevention, application checks, moderation, service reliability and necessary case/dispute records	Our legitimate interests in protecting users, maintaining a reliable service, preventing misuse and resolving disputes, balanced against your rights. You can object as described in section 10.
Handling statutory privacy requests, applicable legal orders and legally required business records	Compliance with the particular legal obligation that applies. This does not authorise retention of the entire account for every obligation.

Purpose	Legal ground
Requested genre filtering and repeat-aware feed delivery	Performance of the requested feed service for its minimum delivery and repeat-prevention state. This does not authorise optional playback measurement or separate outcome analytics.
Optional playback-time ranking	Your specific consent. It is not inferred from registration, Terms acceptance or the general discovery setting. You can withdraw it in Settings > Discovery without losing other features.
Optional Firebase usage analytics, Crashlytics diagnostics, Meta advertising measurement, AppsFlyer campaign measurement and RevenueCat subscription analytics	Your separate consent for the selected purpose, including optional device storage/access where applicable. You may refuse or withdraw without losing the core service. iOS tracking permission is an additional control, not a substitute for the app's purpose-specific choice.
Device files, service providers and necessary testing	The lawful purpose for the underlying operation, with separate access, minimisation and transfer safeguards; infrastructure is not a separate permission for unrelated use.

If required information is missing, we may be unable to provide its corresponding feature. Optional profile photographs are not a condition of having an account.

Free text, images, private communications and reports may contain sensitive information or allegations about other people. We do not treat ordinary photographs as biometric identification or every content-safety label as information about a person's health or criminal record.

Where case evidence contains specially protected information, we restrict it to what is necessary for the particular issue and use an applicable additional legal condition. For example, **Article 9(2)(f)** of the UK/EU GDPR can apply where sensitive evidence is necessary to establish, exercise or defend a legal claim. In the UK, prevention/investigation of unlawful acts or safeguarding can rely on the relevant Data Protection Act 2018 Schedule 1 condition only where its requirements are met. Criminal-offence allegations require separate authorisation under applicable law; the UK's conditions are not a blanket authorisation in other countries. We do not rely on a complainant's submission as consent from the person complained about. Irrelevant details are excluded or removed, and information without an applicable condition must not be repurposed or kept simply because it was supplied.

Automated upload classifications can hold or reject a particular upload; they do not establish a criminal conviction, determine a legal dispute or themselves impose a permanent account ban. You can request human review, submit your explanation and challenge the outcome. Significant restrictions require an assessment appropriate to their effect and the applicable legal safeguards.

5. Who receives information

We disclose information for the purposes above to the following categories:

• **Other users:** the profile/content you publish, chosen messages, authorised full tracks and relevant business/deal information. Their own copies are subject to their responsibilities and applicable law.

• **Authorised IDFlow staff:** the information needed for administration, support and safety review. Staff access is from **the United Kingdom and Bulgaria**. This does not describe all provider locations.

• **Supabase:** authentication, databases and configured backend, realtime and storage operations.

• **Cloudflare:** R2 file storage, media delivery, caching and relevant security services.

• **Google Cloud:** media-processing infrastructure, the safety-check services named above, monitoring and operational logs. Google Workspace provides our support/invitation email and relevant administrative records.

• **Apple:** App Store purchases and offers, transaction verification and iPhone push delivery. Apple also acts independently for its store, payment and device services under its own privacy notice.

• **Google Play:** Android purchases and offers, purchase verification, acknowledgement and subscription lifecycle updates. Google also acts independently for its store, payment and device services under its own privacy notice.

• **Google Firebase Cloud Messaging:** Android service-notification delivery, including device registration/installation identifiers, notification content and routing references described in section 2. This is separate from optional Firebase analytics and crash reporting.

• **Google Analytics for Firebase and Firebase Crashlytics:** optional usage analytics and diagnostics, respectively, with the information and choices described in section 2.

• **Meta and AppsFlyer:** optional advertising/campaign measurement and the consented partner attribution connection described in section 2. Advertising partners can process and match information under their own applicable terms; they are not all treated as infrastructure processors.

• **RevenueCat:** optional Apple and Google Play subscription reporting and history under its data-processing agreement, as described in section 2.

• **Authorities and professional advisers:** relevant information where a specific legal requirement applies or disclosure is otherwise lawfully necessary to protect rights, address a serious safety matter or establish, exercise or defend a claim. We limit disclosure to that purpose.

We require providers acting on our behalf to protect personal information under their applicable data-processing agreements, restrict its use to authorised purposes and apply protection consistent with this notice and applicable requirements. Independent recipients, including other users, Apple and Google for their respective store services, have their own responsibilities; a provider contract does not make every recipient our processor.

External social/website links open services with their own privacy practices. Remote media and email branding may generate a request to the host when displayed, sometimes through a device's privacy proxy.

6. International processing and safeguards

IDFlow is operated from the UK, with the staff access locations stated above. The primary database is configured in London and Google workers in Belgium. Global content delivery, provider support, email, API processing, logs and backups can involve other countries, including the United States and, for relevant provider operations, Singapore. We do not promise that every copy remains in a single European region.

For transfers requiring protection under UK/EU law, we use an applicable adequacy arrangement or contractual safeguards. Our provider agreements incorporate the European Commission's Standard Contractual Clauses and, for relevant UK transfers, the UK International Data Transfer Addendum or the provider's applicable UK clauses. Where a provider relies on a recognised adequacy framework,

that applies only to the covered recipient and processing. Access restrictions, security measures and transfer assessment supplement the applicable mechanism.

You can request information about the safeguards applicable to your data and a copy of the relevant protections at **privacy@idflowapp.com**. We may redact security-sensitive or confidential details while providing meaningful information. Contractual safeguards do not mean that foreign authorities can never lawfully require access.

7. How long we keep information

We retain information for the feature or purpose described, using the periods and triggers below. A shorter period can apply after an eligible removal or account-deletion request. An active dispute, appeal, recovery period, reference to a shared record or specific preservation requirement can delay removal of the relevant subset; it does not justify retaining unrelated information. Scheduled cleanup uses bounded batches and retries, so reaching a period makes a record eligible rather than guaranteeing deletion at that exact second.

Information	Retention period or criterion
Account, current profile, preferences, invitation/onboarding state and current entitlements	While the account remains in existence and the information is needed for its operation, subject to edits, removal requests and account deletion. Incomplete invited/verified accounts follow the account lifecycle; there is no separate automatic invitation-expiry deletion rule.
Unverified, never-used registration attempts	Unverified identities that have never signed in and have no provisioned profile/invitation are eligible for removal after 30 days. Necessary security records have their own shorter/defined periods.
Active tracks, artwork, messages, comments and collaboration records	While needed for the feature/account or until the applicable removal workflow. Removing something from a screen is not necessarily deletion. Message contents marked deleted are eligible for physical cleanup after 30 days, subject to a specific hold.
Released masters and necessary release records	Master, prepared media, track metadata and necessary deal/access history: while the owning account exists, subject to deletion and applicable safety/copyright removal. Extra release-related views and processing diagnostics: six months from first release; ordinary shorter limits still apply outside the release archive. Minimal publication/status links and separately justified open-case evidence have their own purposes.
Declined, expired or cancelled full-track requests	Six months from the later relevant update/expiry, where the request is not linked to a continuing deal and is not held for a specific case.
Feed delivery and optional playback ranking	Server sessions/entries, including display markers and optional playback counters: 30 days from session creation. Optional playback retry receipts: 30 days from receipt. Per-track discovery summaries, including optional playback totals: 90 days after the track last appeared. The device retry queue is bounded to 400 events and 48 hours. Withdrawal deletes playback records and totals; switching discovery off or clearing history removes the corresponding records sooner. The current permission, wording version and saved time remain while enabled and are cleared on withdrawal or account deletion. No outcome attribution or mute/replay analytics.
Ordinary notifications	90 days from creation, including grouped activity.

Information	Retention period or criterion
Push delivery information	Sent/skipped delivery records: seven days after completion; failed records: 30 days. Terminal content payloads are cleared separately. Stale device registrations are eligible after 90 days without relevant device/account activity. Already-delivered device copies follow device controls.
Android push provider information	We set a one-hour delivery lifetime for each message sent to FCM; this is not a promise that all provider records are erased after one hour. Signing out or removing an IDFlow device registration stops our use of that registration; it does not by itself delete the Firebase installation identifier. Google retains Firebase installation identifiers until their deletion is requested through its installation service; after that request, Firebase states that removal from live and backup systems can take up to 180 days. Provider security and operational records follow the applicable Firebase terms and privacy information. Contact privacy@idflowapp.com for help with a deletion request.
Closed support requests and support emails	12 months after case closure. For support email, the first observed Closed label starts the period; reopening or a new reply resets it. An identified retention hold pauses the affected messages. Gmail Trash can retain removed email for a further 30 days.
Routine reports and enforcement records	Resolved no-violation, duplicate or closed reports are eligible after 90 days where not linked to enforcement/held evidence. Ended non-ban enforcement records: 12 months after ending; ended ban records: 90 days. Active restrictions and necessary case evidence follow the separate review criteria below.
Upload moderation details, staging files and appeals	Ordinary result/detail and staging cleanup begins after 30 days from the relevant final decision/review; rejected items without an appeal use 60 days. Closed appeals: 12 months after review. Rejected summary records: 12 months; approved summaries: 90 days after the associated published content is removed. Open cases, holds, live references and released-record preservation can extend these periods for the affected records.
Label review and staff action logs	Label-review events: six months. Administrative action logs: 12 months, subject to specific case preservation.
Recoverable banned accounts	Six calendar months from successful dispatch of the ban notice, plus at least 30 days' pre-erasure notice. Pending appeals, notice-delivery failure or a justified hold can pause erasure. After a rejected appeal, at least 30 days from the decision notice or the original deadline applies, whichever is later.
Continuing registration restrictions and exceptional evidence	Only while the particular restriction, claim or preservation need remains justified. A protected email-matching value may remain to enforce a continuing ban. Such values are personal data. Restriction necessity is reviewed at least annually; exceptional preservation holds require review at least every 30 days and removal when the justified need ends. A review date is not an automatic expiry.
Deletion and restoration safeguards	Detailed completed deletion-job/object records: 90 days. Minimal erasure receipt: 24 months. Separate hashed erasure markers: 90-91 days. Completion-email destinations are cleared after sending or the seven-day failed-delivery window. These records do not contain a retained copy of the deleted account's music/messages.

Information	Retention period or criterion
Purchases and offers	Account-linked entitlements/claims generally follow the account lifecycle. Applied verification receipts: 30 days. Unassigned offer codes: 90 days after expiry. Minimal records required for a particular legal/accounting duty remain for that duty's applicable period; Apple and Google independently retain their respective store records.
Optional measurement	IDFlow's Firebase installation/account mappings expire after 35 days without refresh; server event jobs after 48 hours. Withdrawal removes that installation's mapping and pending events. Crashlytics starts removal of crash reports and associated identifiers after 90 days under its published retention policy. Google Analytics event-level retention is two months; user-level retention is 14 months and resets on new activity. These settings do not limit standard aggregate reports. Meta and AppsFlyer retain provider records under their applicable service terms and policies for the measurement purpose, subject to applicable deletion requests and permitted exceptions. Contact us for details and to request deletion. Turning a choice off is not itself a historical-data deletion request.
Optional subscription analytics	RevenueCat customer/subscription history is retained while the account's reporting choice remains enabled and needed for subscription analysis. Withdrawal or account deletion queues remote customer removal after a two-minute allowance for in-flight requests, with retries until confirmed. IDFlow removes delivered receipt payloads and expires undelivered receipt jobs after 30 days. Provider recovery copies and independently required Apple or Google records can remain separately.
Processing jobs, temporary files and abandoned uploads	Eligible finished processing jobs: 30 days, unless still needed by an active operation or justified case. Released-track diagnostics have the six-month limit above; a minimal publication link can remain where required. Records for unfinished uploads: after 24 hours, excluding active processing. Unreferenced raw/voice files and worker scratch files: generally after 24 hours; replaced public artwork/derivatives: after seven days. Upload permissions and active-reference checks can delay safe removal. Incomplete multipart uploads abort after seven days.
Security, presence and operational records	Authentication audit records: 90 days; cleanup/scheduler history: 30 days; health/capacity samples: seven days; health events: 90 days. Expired deletion challenges: a further 24 hours; abuse-rate counters: up to 48 hours; typing signals: 60 seconds; obsolete presence: 30 days.
Provider logs and backups	Configured Google default logs: 30 days; its locked required audit logs: 400 days. Supabase daily backups: seven days; point-in-time recovery is not enabled. Specific provider recovery/deletion processes and relevant holds can operate separately from active-system cleanup.
Device files, test copies and exports	Device files persist until the relevant cleanup, cache eviction, sign-out or removal action can run. Offline files cannot be remotely erased reliably. Covered test/restore workflows replay deletion markers; independently saved downloads, exports or recipient copies need separate removal and are not erased by a marker alone.

Inactivity: there is currently no active automatic two-year dormant-account content deletion rule. If we introduce inactivity deletion, we will explain the rule and provide advance notice and an opportunity to act before it starts.

Preserved evidence: a narrowly applicable duty or justified legal claim can require restricted records to remain after normal deletion. We limit the records, access, purpose and duration, review the need and delete them when it ends. A report or possible future use does not by itself justify keeping

everything indefinitely. Where lawful, we explain a retention exception when responding to your request.

Backups and copies: backup expiry is separate from removal from active systems. Deleted information must remain out of ordinary use and be reconciled against deletion records before a covered restore is reopened. A provider's own lawful records and a recipient's independent downloads are not automatically recalled by IDFlow account closure.

When you cancel an unfinished upload or replace its file, the app requests cleanup of the discarded files rather than retaining them for reuse. Cleanup is queued and retried; active transfers, live references and justified safety or legal holds can delay physical removal. If the app closes or loses connectivity before that request reaches us, local retry records and the abandoned-upload cleanup periods above provide a fallback. Published or released tracks follow their own removal and retention rules.

8. Deleting your account

You can start deletion in **Settings**. The process verifies your request before irreversible deletion. After confirmation, your account becomes unavailable and queued cleanup removes account data and associated content, including retained released masters, subject to the specific exceptions above.

We aim to complete deletion from active systems within 24 hours of confirmation. This is an operational target, not a guarantee that every provider backup or independently saved copy disappears within 24 hours. Technical failures can cause delays; cleanup continues when the app is closed, and we send confirmation when it completes. Contact **privacy@idflowapp.com** if you need help or have not received confirmation. Applicable privacy-rights deadlines still apply.

Deleting the account does not automatically cancel Apple or Google Play billing. Manage renewal through Apple subscription settings or Google Play subscriptions, depending on where you subscribed. You do not have to wait until a subscription ends to request account deletion. Subscription cancellation and any refund rights are explained in the Terms.

You can also request account deletion without signing in or reinstalling the app using our account-deletion page, or by emailing privacy@idflowapp.com. We verify requests before deleting an account.

9. Waitlist invitations and related enquiries

Where you join the waitlist or send an enquiry, we use the role, name/representative name, email, country, phone if supplied, message and relevant referral/submission details to handle the request and arrange invitations. Anti-abuse checks protect submission. Relevant records pass through our website service, email and Google Sheets; the website host uses Amazon Web Services Lightsail. This invitation processing does not itself enable the optional app measurement described in section 2.

The waitlist remains active until the app's actual public release. The identified signup/outbox, Sheet row and waitlist-mail copies become eligible for removal **90 days after public release, or 90 days after a later signup**. An account created from an invitation has its own account-retention period. A support case follows its closure period. A specific hold may delay the affected copy, and email Trash adds up to 30 days. Sheet revision history, provider recovery copies and separately downloaded exports have separate removal processes; clearing a row does not certify immediate erasure of every historical copy.

A request to hear about an invitation is not agreement to unrelated marketing. You can ask us to stop contacting you about the waitlist at **privacy@idflowapp.com**. Necessary security, billing and account messages are distinct from promotional communications.

10. Your privacy rights

You can edit your profile, manage discovery and notification choices, block/report users and start deletion through the app. Device settings control permissions. You can also email **privacy@idflowapp.com** to request access/a copy, correction, deletion, restriction or portability of eligible information, or to make a complaint.

You can object to processing based on legitimate interests, including associated profiling. Where we rely on consent, you can withdraw it without affecting processing that was lawful before withdrawal. Direct-marketing objections do not require you to give a reason. You can also request human involvement and challenge a decision where applicable automated-decision rights apply.

We may ask for proportionate information to verify a request and protect other people's data. Do not send us your password or sign-in code. Rights have lawful limits; if we cannot fulfil a request, we explain the reason and applicable complaint route. Requests are normally free, subject only to permitted legal exceptions.

For UK/EU requests, we normally respond within one month. If the law permits an extension for complexity or volume, we explain it within the initial period; an extension can be up to two further months. Applicable verification/clarification rules may affect the timetable. We acknowledge UK privacy complaints within 30 days and investigate/respond without undue delay.

You may complain to the UK Information Commissioner's Office or the relevant EU supervisory authority, including where you live, work or believe an infringement occurred. Contacting us first does not remove that right.

11. Additional information for US residents

The categories in this notice include identifiers/contact details; professional and profile information; account credentials; purchases and collaboration records; internet/service activity; photos, audio, video and communications; support/safety information; and saved genre preferences, feed delivery/repeat-prevention records and optional playback-time records. Sources are you, invited-account administrators, other users/counterparties, store/service providers and operation of the app. Sections 2-7 explain purposes, recipients and retention by category.

We do not sell app personal information for money. If you enable Meta advertising measurement or AppsFlyer campaign measurement, the advertising disclosures described in section 2 can constitute sharing for cross-context behavioural advertising or targeted advertising under applicable state law. Both choices are off by default. You can stop this optional sharing in Settings > Privacy preferences without losing access to the app. This statement covers the app; marketing-website technologies have separate controls.

Depending on the state law applicable to you, you may have rights to confirm processing, access, correct, delete and obtain a portable copy, receive information about disclosures, and opt out of sale, targeted advertising or certain significant-effect profiling. Additional rights can apply to sensitive information, including account access details and private communications. We do not unlawfully discriminate for exercising privacy rights.

Submit a request to **privacy@idflowapp.com**; account deletion is also available in Settings. An authorised agent can act where permitted, subject to appropriate verification. We apply the deadline and exceptions required by your state's law. For covered California requests to know, correct or delete, we acknowledge within ten business days and ordinarily respond within 45 calendar days; a permitted extension is explained. A denial explains available review or complaint options. To appeal, email the same address with **"Privacy appeal"** so the decision can be reviewed and the applicable state deadline/regulator route applied.

The app does not provide a web-browser Do Not Track setting. A browser signal does not by itself change your in-app discovery preference. You can use Settings > Privacy preferences for optional measurement, Settings > Discovery for discovery preferences, and the contact route above. These app controls do not govern cookies or tracking on our separate marketing website.

12. Changes to this notice

We identify the version and effective date when publishing this notice. Material changes will be explained through an appropriate in-app or other direct notice. Where a new purpose requires permission or consent, publishing a changed notice or accepting changed Terms does not replace that choice.