ERGONOMY LIMITED

Privacy Policy

Version v1.2 · Effective 21 September 2026

David is provided by **Ergonomy Limited**, a company registered in England and Wales (company no. 17312959) with its registered office at 102 Crawford Street, London, England, W1H 2HR ("Ergonomy", "we", "us", "our"). This policy explains how we collect, use and protect personal data relating to visitors, individual customers, business contacts and users of the David Service. Contact privacy@meetdavid.ai with privacy questions or requests.

1. Our role and the scope of this policy

We act as controller when we determine why and how personal data is used, including for our account administration, billing, business communications and relevant security operations. This policy explains that processing and provides information about how customer content is handled through the Service.

Where we process content on a business customer's instructions, that customer will normally be the controller and we act as processor under our Data Processing Agreement (DPA). Where the customer processes content for another controller, we may act as its sub-processor. Customer organisations are responsible for explaining their own use of David to their staff and other affected individuals, including administrator reporting, configured automations and any optional monitoring they enable.

We also accept individual customers. Sole traders may act as business controllers, while purely personal or household use is treated according to the applicable legal exemption. We do not require personal users to assume statutory business-controller obligations from which they are exempt. Our role depends on the actual processing, and we remain responsible for our own duties. Individual customers can contact us directly for help with privacy rights.

2. Information we collect and its sources

Account and contact information. We collect your name, email address, account details and information you supply when signing up or contacting us. Business accounts may include company name, job title and administrator or billing contact details. We may receive these details from your organisation or an authorised administrator who invites you.

Sign-in information. For Microsoft or Google sign-in, we receive an account identifier and relevant profile information, such as your verified email address, name and profile picture. For email sign-in, we use your email address and the authentication records needed to send and validate the sign-in link or code. Sign-in is separate from connecting a Microsoft, Google or other service: connecting a service involves a separate authorisation and the permissions granted for that connection, and signing in does not by itself grant access to your email, files or calendar.

Agreement records. We record when an agreement is accepted, the network address used and the document versions accepted, to evidence the agreement and administer the relationship.

Billing and payments. We use Stripe to process subscription payments and purchases of additional credits. We share relevant billing and transaction information with Stripe to administer payments, refunds and disputes. Depending on your payment method, Stripe collects and processes payment details and related information. We receive billing details, payment and customer reference identifiers, transaction amounts and status, and limited payment-method information, such as the last four digits of a card or bank account number and a sort code. We use this information to manage your subscription, reconcile payments and provide billing support.

Stripe acts as our processor for certain payment activities and as an independent controller for

certain purposes, including fraud prevention and compliance with legal obligations. Further information about Stripe's processing, retention, international transfers and your rights is in the Stripe Privacy Policy.

Usage and technical information. We process sign-in records, network addresses, features used, requests made, credits spent and errors for security, reliability, account administration and billing. Per-user usage and adoption statistics provided to an organisation's administrators or its authorised support partner are produced on that organisation's behalf under the DPA.

Website usage. When you visit our public website, we collect limited, non-identifying information about how the site is used, so that we can measure and improve it. This includes the pages and sections viewed, how visitors move through the site, the referring website and any campaign tags in the link you followed, and technical details derived from your request such as approximate location at country level, device type, browser and operating system. We run this analytics ourselves on infrastructure we control and do not share it with a third-party analytics provider; it sets no cookies, stores nothing on your device and does not use persistent identifiers, and we do not use it to identify you or build a profile of you. Section 9 explains this further.

Support and feedback. We collect support correspondence and the diagnostic information needed to investigate problems. Customer content accessed for support remains subject to the DPA where we act as processor. In-product feedback used to improve the Service for a customer is treated as Customer Data. Optional monitoring and separate research uses are explained in section 3.

Customer content. Depending on the permissions and features enabled, David processes communications, documents and records from connected services (for example Microsoft 365, Google Workspace, Microsoft Teams, WhatsApp, Xero and helpdesk systems), uploaded files, generated content, meeting recordings, transcripts and screenshots, dictation audio, and derived knowledge. Content may contain information about customers, colleagues, contacts, correspondents and meeting participants, including sensitive information. The DPA governs processing on customers' behalf.

Public sources and abuse prevention. When verifying a company at signup, we look it up on the Companies House public register and hold the company number, registered name and registered office address returned. Our forms use a self-hosted proof-of-work check to deter abuse rather than a third-party captcha.

Required information. Sign-in and basic account details are needed to provide access; relevant billing details are needed for paid services. If you do not provide necessary information, we may be unable to create the account, take payment or provide the affected feature. Optional profile details, research participation and optional feature permissions are not required for unrelated features.

3. Purposes and lawful bases

Providing individual accounts and paid services. Where you personally contract with us, we process the information necessary to provide your account, administer the subscription, take payment and handle service requests on the basis of performance of that contract, or of steps you request before entering it.

Business account administration. Where our contract is with your organisation, we rely on legitimate interests to administer business-contact details, communicate with authorised users and billing contacts, and support the organisation's account. Our interest is providing and managing the agreed business service. Customer-content processing on the organisation's instructions remains governed by its lawful basis and our processor obligations; we do not rely on a contract with each employee.

Security and reliability. We rely on legitimate interests in protecting accounts and systems, preventing fraud and abuse, investigating errors and maintaining a reliable service. We assess the necessity of this processing and its impact on individuals and apply appropriate safeguards.

Agreement evidence and legal compliance. We rely on legitimate interests in maintaining reliable

agreement records and in establishing or defending legal claims. We rely on legal obligation where applicable law requires particular accounting, tax or other records.

Support and improvement. We rely on contract where necessary to support an individual customer, or on legitimate interests in resolving enquiries and improving service reliability. Customer-specific feedback is handled within the customer's instructions. General service analytics use genuinely anonymous information. Creating anonymous information from personal data must itself be lawful; pseudonymised or merely de-identified information remains personal data where individuals can still be identified.

Optional monitoring for support and improvement. Where an organisation's administrators opt in, our support and engineering personnel may review identifiable usage information and conversation content for that organisation to identify problems, help its users and improve how David performs for it. We do this on the organisation's instruction under the DPA, and the organisation is responsible for informing its users. Access is limited and logged, identifiable material drawn from the review is kept for no more than 90 days unless an incident or dispute requires longer, general learning is applied only in anonymous form, and none of it is used to train AI models. Without the opt-in, personnel access follows the support-access safeguards in section 4.

Optional product research. Any identifiable individual usage research is offered separately with an explanation of the information, purpose, recipients and retention. Where we rely on consent, participation is optional and consent can be withdrawn through the stated control or by contacting privacy@meetdavid.ai. An administrator's instruction is not consent from every affected individual. Research opt-in does not authorise AI model training.

Service communications and marketing. We send necessary account, payment, security and service messages on the lawful basis applicable to the underlying activity. For marketing to individuals, sole traders and other recipients protected by electronic-marketing rules, we obtain consent unless an applicable exception, such as a valid existing-customer soft opt-in, permits the message. Where permitted, we rely on legitimate interests in promoting relevant services. We provide an unsubscribe option and respect objections. We do not sell personal data.

Your right to object to marketing. You can ask us to stop using your personal data for direct marketing at any time, including related profiling, by using an unsubscribe link or contacting privacy@meetdavid.ai. We will stop that marketing use. Necessary service messages may still be sent. We retain minimal suppression information so that your choice is respected.

4. AI processing providers and recipients

We use AI model providers, and providers supporting hosting, search, speech and observability, to deliver the Service. The sub-processor page identifies each provider, its purpose, processing location and engagement basis. Supporting providers are used by default or only when a named feature is enabled. Customers control the enabled models and can request supporting-provider exclusions under the DPA; dependent features may then be unavailable.

We do not use Customer Data to train AI models, and our AI model providers are bound by their terms not to train on data supplied for the Service. Supporting providers are configured or contracted not to train on Customer Data where the provider offers that control; the sub-processor page shows each provider's status. Optional monitoring and research processing is subject to the purposes and safeguards in section 3 and the DPA.

Credentials and integration secrets managed by Ergonomy in David are not supplied to AI models. They are used by the Service to authenticate authorised connections and actions. If users put passwords, secrets or other credentials into messages, documents or other content, that content may be processed by the models used for the task. Users should not include secrets in such content.

Low-credit fallback. We use credit-balance and usage information to apply the model-access rules in the Customer Agreement, including running requests on a lower-cost model when an organisation's available balance is at or below its low-credit level. Any automatic switch remains within the organisation's enabled models and providers, its processing-region setting and any

vendor exclusions; if no compliant model can perform a task, that processing pauses. The Customer Agreement explains the level, the charging and the restoration rules. A low balance does not authorise a new processing location or provider.

Other recipients. Personal data is shared with authorised personnel and contractors who need access, Stripe for payments, relevant service providers, and the connected tools and channels selected by the customer. Business administrators and customer-authorised support partners receive information within their configured roles. We may disclose information to professional advisers, courts, regulators or other authorities where necessary and lawful for advice, legal compliance or claims. Each recipient's role depends on the activity; not every recipient is our sub-processor.

Support access to a user's account follows the approval, time-limit, read-only and audit-log safeguards in the Customer Agreement and the DPA. Operational access is restricted to authorised purposes and to personnel bound by confidentiality.

5. Hosting and international transfers

Each organisation's environment, including its database, files and backups, is hosted in the United Kingdom unless its order specifies the European Union or the United States. AI model processing is controlled separately through the UK, EU or full-provider setting and the enabled providers. Selecting a UK or EU AI-model option does not by itself confine every supporting service to that region.

Supporting services, including search, embeddings, transcription and speech features, may process in the United States as disclosed on the sub-processor page and the data-residency page. Relevant overseas support access, logging and onward processing are also taken into account. Customers needing all relevant processing confined to a specified region should agree the restriction under the DPA before the affected processing begins; incompatible features are then disabled.

Where restricted transfers occur, we use an applicable adequacy decision or appropriate contractual safeguards. Depending on the transfer, these may be the EU Standard Contractual Clauses, those clauses with the UK Addendum, or the UK International Data Transfer Agreement. We carry out required assessments and apply necessary supplementary safeguards. You can request information about, or a copy of, the relevant safeguards by emailing privacy@meetdavid.ai, subject to appropriate protection of confidential information. Stripe explains its own international processing in its privacy policy.

6. Retention

We retain personal data only for the relevant purposes, applicable legal requirements and the periods described below. The criteria include whether an account remains active, outstanding payments or disputes, security investigations, legal record-keeping duties and the time in which relevant claims may be brought. Retention is not extended merely because storage is available.

Account and contact records. We keep the necessary records while the account is active. After closure we delete them within 90 days, retaining only what is needed for outstanding matters, legal obligations or claims.

Agreement and billing records. Agreement evidence is retained during the relationship and for six years after it ends. Billing and tax records are retained for six years from the end of the financial year to which they relate, or longer where the law requires.

Security logs, usage records and support. Identifiable security and usage logs are kept for up to twelve months, and traces of AI requests held in our observability tooling for 30 days. Support correspondence and diagnostic traces are kept for up to 24 months after the matter closes. Material drawn from optional monitoring is kept for no more than 90 days after review. Longer retention applies only for a documented incident, dispute or legal requirement. Genuinely anonymous statistics may be retained after the underlying personal data is deleted.

Customer content. After a trial or subscription ends, the environment is suspended and retained for 90 days, allowing restoration, export or earlier deletion, subject to the DPA. Processing during that period is limited to the activities the DPA permits. Defined extensions require a recorded instruction or agreement. Live copies are then deleted within 30 days, environment backups expire after 30 days and previous versions of stored files after 90 days, as set out in the DPA.

Research and marketing choices. Optional research retention is disclosed before participation. We retain only the information needed to administer current marketing choices and honour opt-outs. Withdrawal does not affect the lawfulness of processing already carried out on consent, and records may be retained where another lawful basis requires them.

7. Your rights and requests

Depending on applicable law, the information and the lawful basis, you may have rights to access, correct or erase personal data, restrict processing, receive certain data in a portable format, and object to processing. These rights are subject to legal conditions and exceptions. You may withdraw consent at any time where processing relies on it, without affecting prior lawful processing.

Right to object. In addition to the unconditional right to object to direct marketing described in section 3, you may object on grounds relating to your situation to processing based on legitimate interests. We will consider the objection under the applicable legal requirements.

Contact privacy@meetdavid.ai to exercise rights or to obtain help identifying the relevant controller. We may request proportionate information to verify identity and will respond within the applicable legal timescales, explaining any lawful extension or limitation. Requests about customer-controlled content are forwarded to the relevant controller and supported under the DPA. We handle requests concerning our own controller processing directly. Individual customers are not redirected to themselves as controller.

You may complain to the UK Information Commissioner's Office at ico.org.uk/make-a-complaint. Where the EU GDPR applies, you may also complain to the competent supervisory authority, including in the EU country where you live or work or where an alleged infringement occurred. You may contact us first, but do not have to do so before complaining to a regulator.

8. Security, data isolation and sensitive features

We use technical and organisational controls designed to protect personal data. Customer environments have separate databases, storage and access boundaries, as described on our security page. Access is limited according to authorised roles, customer sharing choices and necessary service operations. These controls are designed to prevent unauthorised access between customers; no system can guarantee that a security incident will never occur.

Meeting and voice features. Recordings, transcripts, screenshots and audio are processed when the relevant features are used. David separates speakers within a transcript; a user may additionally enrol a speaker by name, which creates a voice template held in the organisation's environment and used to recognise that speaker in that user's later recordings. Recognising an individual across recordings from a voice template may involve special-category biometric data; separating speakers within one transcript does not. A user's enrolment is their organisation's instruction under the DPA, and the organisation is responsible for informing the individuals concerned, for a lawful basis and any additional special-category condition, and for any impact assessment. Templates stay in the organisation's environment, are matched only against that user's recordings, are never sent to an AI model provider, and are deleted when the user deletes the enrolled speaker, when the user is deleted, or when the environment is deleted. Speaker recognition is not offered to consumer accounts.

Automated processing. David generates output, performs configured actions and applies usage-based model rules. Customers control the permissions and approvals for their automations. We do not make solely automated decisions that have legal or similarly significant effects on individuals; our account and abuse controls do not decide access on the basis of profiling. If that changes, we

will provide the information and safeguards the law requires, including a way to seek human review, before doing so.

9. Cookies and similar technologies

Our public website does not use advertising or third-party tracking technologies, and we do not load third-party tracking scripts into it. To understand how the site is used and to improve it, we run our own privacy-focused analytics on infrastructure we control; this information is not shared with a third-party analytics provider. This analytics sets no cookies and stores nothing on your device, so it does not require consent; we describe it here and in section 2 for transparency.

It records the pages and sections viewed and how visitors move through the site, together with limited non-identifying technical information derived from your request, such as approximate location at country level, device type, browser and operating system. It does not store your IP address and does not use persistent identifiers, so it does not track you across other websites or over time, and does not build a profile of you. We also record how a visit reached us, including the referring website and any campaign tags contained in the link you followed, for example from an advertisement or email, so that we can see which channels bring people to the site; we remove those campaign tags from the address bar once the visit has been recorded, and we do not store them on your device. We process this information on the basis of our legitimate interest in measuring and improving the website; where it is genuinely anonymous, data-protection law does not apply to it.

The David app and desktop app use strictly necessary cookies and local storage to maintain sign-in and necessary settings; they are not used for advertising or third-party tracking. Our public website also sets a sign-in cookie on internal pages used by our own staff.

Stripe payment pages and other third-party services may use their own cookies or similar technologies under their own notices. Our statement about our own website does not describe those separate services. If we change how we measure the website, or introduce cookies or similar technologies that require consent, we will update this information and implement the consent or other controls the law requires before doing so.

10. Data protection contact

Our data protection contact is Jack Clinton, co-founder, at privacy@meetdavid.ai or the registered office above. Having assessed the nature and scale of our activities, we are not required to appoint a statutory Data Protection Officer under UK GDPR Article 37; we keep that assessment under review.

Ergonomy Limited is registered with the Information Commissioner's Office (registration reference ZC241772). Registration does not constitute ICO approval of the Service.

Where an EU representative under Article 27 of the EU GDPR is required for our activities, its identity and contact details are: [CONFIRM: EU representative].

11. Changes to this policy

We may update this policy to reflect changes in the Service, our practices or the law. Each version is published at a dated, versioned address and remains retrievable. We notify account contacts of material changes and provide any further notice, choices or consent required before introducing relevant new processing. Publishing an updated policy does not by itself authorise a new use of personal data. Contact privacy@meetdavid.ai with questions.