


Data Flow

Where customer data comes from, where it is stored and processed, and which providers see it.

1. ROLES

The customer organisation is the **controller** of the personal data in its workspace. Workplacer is its **processor**, under the customer's Data Processing Agreement. The providers named below are Workplacer's subprocessors, listed in Schedule A of that agreement.

2. DIAGRAM



Solid boxes: the customer, and Workplacer's EU-hosted platform. Data at rest stays in the EU.
Dashed box: providers that process limited data in the US, under EU Standard Contractual Clauses (2021).
Numbers match the flows in section 3.

3. FLOWS

Flow	Data	Route and location	Safeguards
1 · Accounts and use	Users' names, work email addresses and roles, sign-in sessions, and an audit log of changes they make.	Users ↔ Workplacer app. Stored in the database (EU).	Every record belongs to one customer workspace, and access follows the customer's own roles. Sessions end 7 days after last activity.

Flow	Data	Route and location	Safeguards
2 · Workforce and occupancy data	Headcount, departments, presence counts and, for commute analysis, home postcodes.	Customer's HR, access-control or sensor exports → Workplacer, by file upload or connector.	Names, employee numbers and card numbers are replaced by keyed pseudonyms as each row is read, and only counts and pseudonyms are kept. UK home postcodes can be cut to the district before storage.
3 · Documents	Leases, contracts, invoices and reports. These can include names and contact details of signatories and counterparties.	Uploaded in the app, forwarded to the workspace's email address, or synced from SharePoint (the customer's own, or a Workplacer drop folder in Microsoft 365, EU). Stored as files (EU).	The customer chooses how much of each document AI may read, or turns AI off (flow 5).
4 · Storage and operations	Everything in flows 1–3, plus rate-limit counters in the cache and error reports.	Database (Neon, Frankfurt), files (Vercel Blob), cache (Upstash, EU) and error monitoring (Sentry, Frankfurt).	TLS 1.3 in transit and AES-256 at rest. Database backups cover a rolling 7-day window.
5 · AI processing	Document text, to the extent the customer's document data mode allows (full document, text only, text with personal data redacted, or none), and aggregated figures.	Workplacer → Anthropic (analysis) and Voyage AI (search embeddings), US.	Personal data is scrubbed before text is embedded. No customer data is used to train models. AI output is advisory. See the AI Governance Policy.
6 · Background jobs	Job events and step results, which carry record IDs and the working data a job needs.	Workplacer → Inngest, US.	Each event is scoped to one customer workspace.
7 · Maps and geocoding	Site addresses and postcodes, never with a person's name.	Workplacer → Mapbox and Google, US. Geoapify (Germany) handles some address lookups in the EU.	Only the place is sent, not who it belongs to.
8 · Notifications	Recipient name and email address, and the notification's content.	Workplacer → Resend (Ireland) → users.	Sent only as part of the customer's own work in the platform. Demo accounts are filtered out.
9 · Single sign-on (optional)	The user's sign-in identity: name and email address.	User's identity provider (Google, Microsoft or Apple) ↔ Workplacer.	Used only when the customer chooses that sign-in method.

4. WHAT DOES NOT FLOW

- Special categories of personal data (GDPR Art. 9). The platform does not ask for them or process them as structured data.
- Customer data to train AI models, by Workplacer or its providers.
- Data between customers. Each workspace is isolated at the query layer.
- Customer data to anyone for their own purposes: no sale, no advertising, no analytics resale.

5. RETENTION

Customer data is kept for the life of the contract. On termination it is returned or deleted at the customer's choice, with a 30-day extraction window. Shorter fixed limits apply to sessions, tokens, notifications and inbound email. The Data Retention Schedule has the full detail.

This policy should be read alongside the Data Processing Agreement (Schedule A is the authoritative subprocessor list), the AI Governance Policy, the Data Retention Schedule and Privacy by Design.