

BudgetUp DPIA template and summary

Pre-filled data protection impact assessment for institutions, including learners aged 16 and 17

Template: subject to review by your legal team and your data protection officer. This is BudgetUp's own assessment of the processing it carries out as your processor, pre-filled so you can complete your DPIA faster. Your institution, as controller, owns the DPIA and its sign-off. Sections marked **[Controller]** are for you to complete. It is not legal advice.

Version 0.1 · Prepared 8 October 2026 by Tamsar Global Ltd (trading as BudgetUp), company 17371733 · Based on BudgetUp's internal compliance record (docs/04-compliance.md) · Contact: privacy@budgetup.uk

1. Why a DPIA is needed

A DPIA is likely to be required, or at least good practice, because the processing:

- involves learners who may be **children (aged 16 and 17)** in further education, which the ICO's Age Appropriate Design Code (the Children's Code) treats as needing a DPIA;
- may include **special category data** (optional disability, health and ethnicity information);
- touches on **financial vulnerability**, since learners may reveal debt or hardship;
- includes an optional **AI chat feature** (the AI coach, adults only), which counts as innovative technology on the ICO's list of processing likely to result in high risk.

[Controller] Record your screening decision and the DPIA reference here.

2. Description of the processing

What BudgetUp is. A gamified UK financial education service: short lessons, games and an arcade covering budgeting, student finance, renting, bills, work and tax, benefits, borrowing, saving, pensions, investing basics and scams. It is education, not financial advice. It never names or recommends products or providers.

Who is involved. The institution is the controller for learners it enrolls. Tamsar Global Ltd is the processor under an Article 28 data processing agreement (DPA). Sub-processors are listed in Schedule 3 of the DPA and at <https://budgetup.uk/trust>.

Data subjects. Students, apprentices, FE learners (including 16 and 17 year olds), residents and tenants, employees, and the institution's administrators.

Data collected.

Category	Source	Required?
Name, display name, email, generated pseudonym	Learner, or the institution's SSO / SCIM / LTI	Partly: access-code learners can avoid giving an email
SSO, LTI and directory identifiers	Institution systems	Only if the institution connects them
Age band (not date of birth)	Learner, at onboarding	Asked; "prefer not to say" is allowed and triggers under-18 protections
Learning activity and progress	Generated by use	Yes, it is the service
Onboarding answers and check-in surveys	Learner	Optional
Equality information (disability or health condition, ethnicity, care-experienced, estranged, first in family and similar)	Learner	Optional, explicit consent only
AI coach conversations	Learner (adults only)	Optional, off unless the institution switches it on
Problem reports and feedback	Learner	Optional
Security and audit records	Generated	Yes, minimal (account id only)

Where it is stored. Hetzner Online GmbH, Germany (EU), covered by UK adequacy regulations. Encrypted backups are kept off-site in Cloudflare R2 in the EU jurisdiction.

[Controller] Describe how BudgetUp fits your programme (for example, a money module for first-year students or a tenancy-readiness course), which cohorts will use it, and which access route you will use (SSO, LTI, SCIM or access codes).

3. Necessity and proportionality

- **Lawful basis [Controller].** Typically public task (Article 6(1)(e)) for public bodies, or legitimate interests (Article 6(1)(f)) for others, for core learning data; consent for optional surveys and research; explicit consent (Article 9(2)(a)) for equality information.
- **Minimisation.** An age band is stored instead of a date of birth. Equality information is optional and asked separately with its own consent. Institution reports are aggregates computed live, with any group under 10 learners suppressed. No profiling by default.
- **Transparency.** Learners see BudgetUp's privacy notice; the institution's notice should mention BudgetUp too. The AI coach shows a first-use notice ("AI can be wrong; this is general education, not advice").
- **Data subject rights.** Self-service in the app (see section 5.7), plus processor help within 10 working days.
- **Processor assurance.** An Article 28 DPA, a published sub-processor register with 30 days' notice of changes, and a security overview.

4. Consultation

[Controller] Record consultation with your DPO, information security, safeguarding lead, students' union or learner voice, and accessibility team. BudgetUp will join consultation meetings on request.

5. Risks and measures

Risk scoring: likelihood (remote, possible, probable) × severity (minimal, significant, severe). Residual scores are BudgetUp's assessment for you to confirm.

5.1 Learners aged 16 and 17 (Children's Code)

Risk. Children's data processed in ways that are not in their best interests: nudging, profiling, social pressure, marketing, or an AI chat.

Measures.

- **Under-18 mode** switches on when the learner picks the 16–17 band, prefers not to say, has been flagged under 18, or belongs to an institution set to under-18 mode without having given a band. The most protective source always wins, and it applies on the device from the moment the band is tapped, before sign-in.
- In under-18 mode: **no AI coach, no analytics, no browser error reports, no marketing prompts, no advertising tags, no email or push nudges, no individual leaderboards, no friends or social features, no profiling**, and weekly instead of daily streaks.
- High privacy by default (standard 7); nothing optional runs before a choice is made.
- No purchase prompts or prices for under-18s. A learner may choose to "ask a parent or guardian" once, with at most one request per 7 days, no reminders and no countdowns (standards 5 and 13). Only the email the child types is kept, and only until the request is paid or expires.
- No loot boxes, random rewards, paid streak repair or gambling-style mechanics.
- Settings show whether under-18 mode is on, why, and what is different.
- Unit tests cover the age policy and the analytics gate (under-18 × consent × route).

Residual risk: low (remote × significant).

5.2 Special category data is optional and consented

Risk. Disclosure of disability, health or ethnicity information, or discrimination based on it.

Measures.

- Never required. Asked separately, with explicit consent recorded with a version and timestamp; it can be withdrawn at any time in Settings.

- Encrypted at the application layer (AES-256-GCM) with a dedicated key held outside the database.
- Row-level security: no institution staff role can read it. Reports receive only suppressed aggregates (groups under 10 hidden).
- Included decrypted in the learner's own data export; deleted with the account.

Residual risk: low (remote x severe).

5.3 Inference of financial distress

Risk. The service reveals or infers that a learner is in financial difficulty, and this is used against them.

Measures. At-risk flags are off by default. They can be switched on per institution and per learner only, and are rules-based and explainable. There are no individual reports of financial situation to staff; outcomes reports are suppressed aggregates. Every game ends with signposts to free, impartial help (MoneyHelper, StepChange, National Debtline, Citizens Advice, Shelter) and, for crisis, Samaritans and Shout.

Residual risk: low.

5.4 AI coach safeguards (adults only)

Risk. Inaccurate or harmful answers; answers that stray into regulated financial advice; disclosure of sensitive information in chat; children using it.

Measures.

- **Access.** Off by default for every institution, switched on by an institution administrator (audit-logged); always off in under-18 mode; requires sign-in and acceptance of a first-use notice; rate-limited (6 a minute, 40 a day) with a daily spend cap per learner.
- **Safety routing first.** Deterministic rules decide safety routes before any AI call. Self-harm, crisis, scam-in-progress, debt emergency and self-declared under-18 messages get fixed replies with helplines and **no AI-generated answer**. A model-based classifier can only make the route more protective, never less.
- **Grounded answers only.** The coach answers only from reviewed BudgetUp lessons and cites them. If retrieval confidence is low, it says it has no lesson on that and points to MoneyHelper.
- **Output checks on every sentence before display.** Named providers, brands or coins; rankings; personal instructions ("you should take out..."); allocations; guarantees; unapproved phone numbers or links; prompt leaks; and key UK figures that contradict our reviewed figures are removed or corrected. The same content linter used in our editorial process runs on every output sentence.

- **Data.** No names or email addresses are sent to the model; learners are told not to share personal details. Conversations are kept 90 days, visible only to the learner (institution staff have no read path), deletable immediately with "Delete my chats", and never used for training.
- **Testing.** A red-team evaluation suite: on 8 October 2026 it passed 74/74 in fixture mode and 73/74 live (66/66 critical cases).
- **Processing location.** Production processing is planned for Google Cloud Vertex AI in London (europe-west2). **Current position:** the move from the Gemini API to Vertex AI is pending, and we will not switch the coach on for institution learners until it is complete.

Known residual limits. The rule-based classifier is English-only and pattern-based, so indirect crisis language can slip past it (the model classifier is the second net). Figure checks cover the key figures, not every number. There is no institution-level monthly budget cap yet.

Residual risk: medium until the Vertex AI move and a red-team re-run on Vertex are complete, then low. You can keep the coach switched off.

5.5 Analytics is consent-gated, and there are no ad pixels in the app

Risk. Tracking learners' behaviour for purposes they would not expect, or leaking it to advertising platforms.

Measures.

- Product analytics (PostHog EU Cloud, Frankfurt) and browser error reports run **only with the learner's opt-in consent**, never in under-18 mode, and never when Global Privacy Control or Do Not Track is set. Rejecting is as easy as accepting.
- Event properties are limited to slugs, numbers and booleans. Identity is a salted one-way hash of the account id. The institution's name is never sent. Session recordings (about 1 in 4 visits) mask all text and typing and never run on settings, admin, coach or parent pages.
- **No advertising pixels inside the learning app.** Advertising tags (Google, LinkedIn) run only on an allow-list of public marketing pages, only with separate marketing consent from adults, and never in under-18 mode. Every app route is denied by default, so a new page cannot leak to an ad platform by accident. Tests enforce the allow-list.

Residual risk: low.

5.6 Retention

Data	Retention
Account and learning data	While the account exists; the learner can delete it at any time
Inactive accounts	Deleted after 24 months with no activity (30 days' warning where we have a verified email)
Data held for the institution	Deleted or returned within 30 days of contract end
AI coach conversations	90 days, or immediately with "Delete my chats"

Data	Retention
Problem reports and feedback	24 months
Records of data-rights requests	3 years after completion
Security and audit logs	Up to 6 years, holding an account id only
Backups	Overwritten within 35 days (30 days off-site, 7 days on the server)

A daily retention job applies the automated rules and logs what it removed. One source of truth in the code drives both the privacy notice and the job, and a test fails if they drift apart.

Residual risk: low.

5.7 Data subject requests (DSAR tooling)

- **Self-service** in Settings, "Your data and account": download my data (JSON or readable HTML), correct my details, delete my account (a two-step confirmation on the page), and "ask about something else" for other rights.
- **Operator queue** for requests by email, post or phone, or passed on by the institution, with the one-month statutory deadline calculated automatically, extensions tracked, and overdue requests flagged.
- Every export, correction and deletion is audit-logged. A test fails if a new database table holding personal data is not covered by export and deletion.
- Where the learner was enrolled by the institution, BudgetUp passes on requests that need the controller's decision and helps within 10 working days.

Residual risk: low.

5.8 Security and tenant isolation

Row-level security in the database separates institutions, tested automatically on every change; encryption in transit; encrypted backups; least-privilege database roles; append-only audit log. Gaps: full-disk encryption on the database server, an independent penetration test and Cyber Essentials for Tamsar Global Ltd are planned, not done. See the security overview.

Residual risk: low to medium (possible × significant) until the planned items are complete.

6. Summary of residual risk

Risk	Residual
Children's data (16 and 17 year olds)	Low
Optional special category data	Low
Financial distress inference	Low
AI coach (adults, if switched on)	Medium until the Vertex AI move; can be avoided by keeping it off
Analytics and advertising	Low

Risk	Residual
Retention	Low
Data subject rights	Low
Security and isolation	Low to medium, pending the planned assurance items

No residual high risk is identified, so prior consultation with the ICO should not be needed. **[Controller]** confirm.

7. Sign-off [Controller]

Item	Name and date	Notes
Measures approved by		
Residual risks approved by		
DPO advice		
DPO advice accepted or overruled		If overruled, give reasons
Learner or students' union views		
Review date		We suggest yearly, or when the AI coach is switched on