GRAFT

GDPR TEMPLATE PACK

Data Protection Impact Assessment (DPIA) — Hybrid Workflow Assessment

Pre-populated DPIA template for clients deploying Graft's hybrid workflow assessment. Covers Jira, Microsoft 365 and Google Workspace integrations.

VERSION	LAST UPDATED	PUBLISHED BY
v1.0	**2026-04-07**	**Graft**

TEMPLATE — REQUIRES LEGAL REVIEW BEFORE USE

This document is a drafting aid. Your DPO, legal counsel or works council should review every section before use.

> **WARNING**
> **Template -- requires legal review.** This document is a starting point. Your DPO, legal counsel, or supervisory authority contact should review and adapt every section before you sign it off. Graft has pre-populated the technical and architectural sections so your reviewer only needs to confirm them, not write them from scratch.

1. Executive summary

This Data Protection Impact Assessment (DPIA) covers `<<CLIENT ORGANISATION NAME>>`'s deployment of **Graft's hybrid workflow assessment**. The processing combines read-only OAuth integrations into existing workplace tools (Jira, Microsoft 365, Google Workspace) with a short employee questionnaire to produce an aggregated workflow map.

The assessment is positioned as a **process improvement** activity, not a productivity monitoring activity. No message content, no keystroke data, no screen capture and no individual performance scoring is collected. Outputs are aggregated to cohorts of `n >= 5` employees by default.

A DPIA is conducted because the processing involves systematic monitoring of employees in a working environment (Art. 35(3)(c) GDPR; ICO list of operations requiring a DPIA).

Conclusion: the residual risk after the controls described in section 8 is assessed as **low**, on the basis that (i) only metadata is processed, (ii) the lawful basis is legitimate interest with a documented balancing test, (iii) employees are notified in advance, and (iv) the n>=5 aggregation rule prevents re-identification of individuals from the published outputs.

2. Description of the processing

2.1 Nature of the processing

Graft connects, through the **Nango** integration platform (self-hosted by Graft), to the following systems using **read-only OAuth scopes**:

System	Source API	Scope	Frequency
Atlassian Jira	Jira Cloud REST API (`/rest/api/3/search`, changelog endpoints)	`read:jira-work`, `read:jira-user`	Hourly delta sync
Microsoft 365	Microsoft Graph Reports API (`/reports/getEmail*`, `/reports/getOffice365Active*`)	`Reports.Read.All`	Daily

Google Workspace	Admin SDK Reports API (`activities/users/all/applications/*`)	`- Daily https://www.googleapis.com/auth/admin.reports.audit.readonly`

Access tokens are issued by the client tenant administrator at the start of the engagement and can be revoked at any time from the source system's admin console. Graft stores no passwords and never has interactive access to user accounts.

2.2 Scope and context

- **Purpose:** to build a baseline workflow map identifying repetitive, automatable, or AI-augmentable tasks across teams.
- **Population:** `<<CLIENT FIELD: list teams / departments in scope>>`. Out of scope by default: contractors, third parties, and any employee who has invoked the opt-out described in section 9.
- **Duration:** initial assessment runs for `<<CLIENT FIELD: typically 14 days>>`, after which the integration can be paused or kept connected for a "continuous improvement" subscription.
- **Geography:** processing takes place in `<<CLIENT FIELD: e.g. EU (Frankfurt) for EU clients on the Business tier>>`. See section 11 for transfers.

2.3 What we collect, what we don't

We collect (metadata)	We do **not** collect
Jira issue IDs, statuses, assignee IDs, transitions, timestamps	Issue descriptions, comments, attachments
M365 Reports API aggregates: emails sent/received counts, Teams active users, Office app usage counts	Email subjects, bodies, attachments, recipients, calendar invitee identities
GWS audit log events: app launched, file accessed (event type only), login events	Drive document content, Gmail content, Meet recordings
Tool-to-tool transition counts (e.g. number of times user moves Jira ’ Slack within 5 minutes)	Browser history, URL paths, screen contents
User identifiers as opaque hashes (`sha256(email + per-org salt)`)	Plain-text email addresses outside the source tenant

We process metadata about _what tools are used and how often_, never _what is being said inside them_.

3. Data categories and data subjects

Data category	Source	Purpose	Special category?

Workplace identifiers (hashed)	Source tenant	Linking activity across tools per individual for the duration of analysis	No
Activity counts (per user, per day, per app)	M365 / GWS / Jira	Building workflow patterns	No
Workflow transitions and tool-switching events	Derived	Identifying friction and automation candidates	No
Free-text questionnaire responses	Employee	Adding qualitative context to observed patterns	No (employees instructed not to include personal data; PII detection runs on ingest)

Data subjects: employees of `<<CLIENT ORGANISATION NAME>>` whose accounts are within scope. Approximate population: `<<CLIENT FIELD: number of in-scope employees>>`.

No special category data (Art. 9 GDPR) is intentionally collected. Free-text questionnaire fields run through a personal-data detection pass before storage; matches are redacted.

4. Lawful basis

The lawful basis is **legitimate interests** under **Art. 6(1)(f) GDPR**. Consent is **not** relied upon, because the European Data Protection Board has been clear that consent is generally not freely given in the employer-employee context (EDPB Guidelines 05/2020 on consent under Regulation 2016/679, and Recital 43 GDPR on imbalance of power).

4.1 Legitimate interests assessment (LIA)

Test	Result
Purpose test — what is the legitimate interest?	Identifying inefficiencies, friction, and automation opportunities to improve how work is done. This benefits the organisation (efficiency) and the workforce (less repetitive toil).
Necessity test — is processing necessary to achieve it?	Yes. Self-reported surveys alone systematically under-report common tasks ("invisible work") and overestimate exception handling. Tool metadata corrects for this without requiring screen capture or content access.
Balancing test — does the interest override the data subject's rights?	Yes, conditional on the safeguards in section 8: metadata only, ne5 aggregation, advance notification, opt-out, retention limits, and a clear separation from any HR or disciplinary process (see section 9).

The balancing test is documented in the appendix "Legitimate Interests Assessment record" attached to this DPIA on file with `<<CLIENT FIELD: DPO name>>`.

4.2 Why not consent?

- Power imbalance: an employee asked to consent cannot freely refuse without perceived career risk (Recital 43 GDPR).
- Withdrawal: consent must be as easy to withdraw as to give. For an integration that runs across a whole tenant, per-employee withdrawal would be operationally impractical and would create selection bias in the dataset.
- ICO Employment Practices guidance (2023): for monitoring activities, legitimate interests with a documented LIA is the recommended basis, not consent.

5. Necessity and proportionality

Question	Assessment
Is the processing necessary to achieve the purpose?	Yes — see 4.1 necessity test.
Is there a less intrusive way?	Considered alternatives: (a) survey only — rejected, under-reports invisible work; (b) screen capture — rejected, disproportionate, content-level intrusion; (c) keystroke logging — rejected, almost certainly disproportionate per ICO guidance. Metadata-only integration is the least intrusive option that produces useful results.
Is the data minimised?	Yes — only the API endpoints listed in 2.1 are called; only event-level metadata is stored; identifiers are hashed.
Is the retention proportionate?	Yes — see section 10.
Is processing limited to the stated purpose?	Yes — outputs are workflow maps and recommendations, not individual scorecards. The system has no API for "show me employee X's activity" outside of debug mode, which is restricted to the Graft engineering on-call rota and audited.

6. Consultation

This DPIA should be reviewed by:

- `<<CLIENT FIELD: Data Protection Officer or equivalent>>`
- `<<CLIENT FIELD: Workforce representatives -- works council, union, or employee forum>>` (mandatory in Germany under BetrVG §87, in the Netherlands under WOR Art. 27, and recommended elsewhere)
- `<<CLIENT FIELD: IT Security>>`

- `<<CLIENT FIELD: HR -- to confirm separation from performance management>>`

Sign-off recorded in section 14.

7. Risk register

Risks are scored on a 1-5 scale for **likelihood** and **impact on data subjects**. Residual risk is the score after the listed mitigations.

#	Risk	Inherent (L × I)	Mitigation	Residual (L × I)
R1	**Function creep** — assessment data is later repurposed for performance management or disciplinary action	3 × 5 = 15	Contractual prohibition in DPA section 4 ("no individual scoring"); ne5 aggregation enforced in product; audit log of every query that returns data; reference to Barclays desk-sensor and Teleperformance webcam-monitoring backlash as case studies of why this fails.	1 × 5 = 5
R2	**Re-identification** — small teams mean an "aggregated" cohort of 5 still effectively identifies individuals	3 × 4 = 12	Default ne5; cohorts smaller than n are suppressed; client admin can raise the threshold but never lower it; differential cell suppression on cross-tabulations.	1 × 4 = 4
R3	**Excessive retention** — raw event data accumulates indefinitely	4 × 3 = 12	90-day raw data retention enforced via scheduled deletion job; 2-year aggregated retention; deletion verifiable via audit log.	1 × 3 = 3

R4	**OAuth token compromise** — leaked credentials grant read access to client tenant	2 × 5 = 10	Tokens stored encrypted at rest (AES-256-GCM via Nango), org-scoped credential isolation, automatic rotation, immediate revocation path documented in section 9.	1 × 5 = 5
R5	**Inadvertent collection of content** — an API change adds content fields to a previously metadata-only endpoint	2 × 4 = 8	Schema validation on ingest rejects unexpected fields; weekly automated drift check against API specs; incident playbook for new field detection.	1 × 4 = 4
R6	**Cross-border transfer concerns** — non-EU sub-processors process EU personal data	3 × 3 = 9	EU data residency option; SCCs in place with Anthropic, OpenAI and Sentry; transfer impact assessment on file. See section 11.	2 × 3 = 6
R7	**Loss of employee trust / surveillance perception**	4 × 3 = 12	Mandatory privacy notice (section 9), works council consultation in DE/NL, public framing as "process improvement, not monitoring", and the opt-out.	2 × 3 = 6
R8	**Free-text PII leak** in questionnaire responses	3 × 3 = 9	PII detection on ingest; redaction of detected entities; instruction to employees not to include personal data.	1 × 3 = 3

R9	**Sub-processor breach** propagates to client data	2 × 4 = 8	Sub-processors selected for ISO 27001 / SOC 2; DPA chain in place; 72-hour breach notification (Art. 33 GDPR).	1 × 4 = 4

The residual risk profile is dominated by R1 (function creep). The single most important control is the contractual and product-level prohibition on individual scoring. If your organisation cannot or will not commit to that, the assessment should not be deployed in its current form.

8. Data subject rights

Employees retain full GDPR rights with respect to data processed by Graft on the client's behalf. The client (as controller) is the first point of contact for any rights request; Graft will action requests via the controller within 72 hours of being instructed.

Right	How it's exercised
Access (Art. 15)	Employee contacts client DPO; Graft returns the hashed-identifier event log for that user within 72 hours of request.
Rectification (Art. 16)	Tool metadata is factual and not normally subject to rectification. Free-text questionnaire responses can be edited or withdrawn by the employee for the duration of the assessment.
Erasure (Art. 17)	Per-employee erasure removes the hashed identifier and all linked events. Aggregated cohort outputs are not affected (and are not personal data).
Restriction (Art. 18)	Employees can be marked as "out of scope", which excludes them from future syncs and from any new aggregations.
Portability (Art. 20)	Event log export available in CSV / JSON.
Objection (Art. 21)	Objection on grounds relating to the employee's particular situation triggers a re-run of the LIA balancing test for that case. Where the balance no longer holds, the employee is added to the opt-out list.
Automated decisions (Art. 22)	Not applicable — no automated decisions with legal or similarly significant effects are made about individuals from this data.

9. Opt-out and transparency

- **Notice:** every employee in scope receives the **plain-language privacy notice** (separate document) **at least 14 days before the integration is enabled**.
- **Opt-out:** employees can opt out by emailing `<<CLIENT FIELD: opt-out mailbox>>`. Opt-out is honoured at the next sync, within 24 hours.
- **Display:** the dashboard shows aggregated counts only; no individual drill-down is exposed to client admins.
- **Separation from HR processes:** the assessment outputs are contractually and technically separated from performance reviews. See R1 above.

10. Retention

Data class	Retention	Deletion mechanism
Raw event logs (per-user, per-event)	**90 days** from collection	Daily scheduled job; cryptographic erasure of encrypted segments
Aggregated cohort outputs (ne5, no identifiers)	**2 years** from collection	Manual deletion at end of retention period
OAuth refresh tokens	Until engagement ends or revoked	Immediate hard-delete on revocation
Audit logs (who queried what)	**7 years** (legal hold)	Standard audit retention

11. Security measures

Control area	Measure
Encryption in transit	TLS 1.2+ (TLS 1.3 preferred) on all API calls and dashboard access
Encryption at rest	AES-256-GCM (Nango-managed credential vault); database encryption at rest via Neon
Credential isolation	Per-organisation OAuth credentials, scoped at the Nango connection level; cross-tenant access prevented at the application layer
Authentication	Better Auth with Google SSO, Microsoft Entra SSO, or SCIM-provisioned accounts; MFA enforced for client admins

Network	Hosting on Vercel (CDN + serverless) with WAF; database on Neon Postgres with private endpoints
Logging and monitoring	Pino structured logging; Sentry error monitoring (PII scrubbing enabled); audit log of every query that returns user-level data
Access control	Least-privilege RBAC inside Graft; production database access limited to two named SREs with break-glass procedure
Vulnerability management	Dependency scanning on every PR; quarterly penetration test (planned)
Backup	Daily automated database snapshots, 30-day retention
Incident response	72-hour breach notification per Art. 33 GDPR; runbook in DPA Schedule 2

12. Cross-border transfers

The default deployment processes and stores data in the **EU region (Frankfurt)** for clients on the Business tier and above. For clients on the Starter tier, data may be processed in the **United States** unless EU residency is specifically requested.

Sub-processors and transfer mechanisms:

Sub-processor	Role	Region	Transfer mechanism
Nango	OAuth + integration sync infrastructure (self-hosted by Graft)	EU (Frankfurt)	N/A — same region as Graft processing
Neon	Database	EU (Frankfurt) for Business+; US for Starter	EU SCCs in place for Starter tier
Vercel	Hosting and CDN	Global edge; primary compute in EU for Business+	EU SCCs
Anthropic	LLM (primary, used for narrative generation in reports)	United States	EU SCCs + Transfer Impact Assessment on file
OpenAI	LLM (fallback)	United States	EU SCCs + Transfer Impact Assessment on file

Sentry (Functional Software, Inc.)	Error monitoring	United States	EU SCCs + Transfer Impact Assessment on file

Transfer Impact Assessments for the three US sub-processors evaluate (i) the legal regime of the recipient country, (ii) the nature of the data transferred (metadata only, no content, no special category data), and (iii) supplementary measures (encryption in transit and at rest, no plain-text personal data sent to LLMs except where explicitly redacted). TIAs are available on request from `legal@graft.bot`.

13. Contacts

Role	Contact
Client Data Protection Officer	`<<CLIENT FIELD: DPO name and email>>`
Client controller representative	`<<CLIENT FIELD: name, role, email>>`
Graft processor contact	`legal@graft.bot`
Graft technical/security contact	`security@graft.bot`
Lead supervisory authority (client)	`<<CLIENT FIELD: e.g. ICO (UK), CNIL (FR), AP (NL)>>`

14. Sign-off

Role	Name	Signature	Date
Data Protection Officer	`<<CLIENT FIELD>>`		
Information Security lead	`<<CLIENT FIELD>>`		
HR lead	`<<CLIENT FIELD>>`		
Workforce representative	`<<CLIENT FIELD>>`		
Business owner	`<<CLIENT FIELD>>`		

Document control

Field	Value

Version	1.0
Last updated	2026-04-07
Next review	2027-04-07, or on any material change to the processing
Owner	Client DPO
Source template	Graft GDPR template pack v1.0

Change log

Version	Date	Author	Change
1.0	2026-04-07	Graft	Initial template release