Data Protection Impact Assessment Process

Tovie AI UK Limited

Version 2.0

Last updated: February 2026

Contents

---

1. Introduction

Tovie AI UK Limited (the "Company") is fully committed to protecting the personal data of its customers, employees, suppliers and other stakeholders in accordance with the requirements of the European Union General Data Protection Regulation (EU GDPR), the United Kingdom General Data Protection Regulation (UK GDPR), and other applicable data protection legislation.

The Company takes the privacy of personal data very seriously and has implemented a variety of methods and controls to ensure it knows what data it collects and holds and that it protects that data appropriately.

As part of this commitment, the Company ensures that all business activities and projects that involve the use of personal data are subject to a data protection impact assessment (DPIA) where required. The purpose of this assessment is to ensure that the use of personal data is fully understood, that the risks to that data are carefully examined and that all appropriate measures are put in place to protect it throughout its lifecycle.

Context of AI Agent Platform. AI Agent Platform enables Clients to deploy AI agents for automated processing, multi-agent orchestration, and integration with external systems (including LLM providers, CRM, ERP and RAG repositories), that is why DPIAs are particularly critical for assessing risks associated with:

- Large-scale automated decision-making and profiling through AI agents
- Processing of personal data by external LLM providers
- Multi-agent workflows where data flows through multiple processing steps
- Integration with Client systems containing sensitive or special categories of data
- Different deployment models (cloud, on-premises, hybrid) with varying risk profiles

This document sets out the Company's process for carrying out a DPIA and, in conjunction with associated forms and guidance, should be used to ensure that obligations and policies in this area are met.

This document should be read in conjunction with the Company's Privacy Policy and any applicable Data Processing Agreements with Clients.

---

2. Definitions

The following definitions of terms used within this process document are taken from the GDPR and supplemented with Company-specific terminology.

"personal data" means any information relating to an identified or identifiable natural person ("data subject"); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.

"processing" means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction.

"pseudonymisation" means the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person.

"controller" means the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data.

"processor" means a natural or legal person, public authority, agency or other body which processes personal data on behalf of the controller.

"AI Agent" means a software module operating on the Tovie AI Agent Platform that performs tasks such as processing user queries (chat, voice, email), integrating with external systems and automating workflows on behalf of a Client.

"Platform" means the Tovie AI Agent Platform, including all associated services, tools and infrastructure (for example platform.tovie.ai, ds.tovie.ai, cloud.tovie.ai, agent.tovie.ai, help.cloud.tovie.ai).

"Multi-agent system" means coordinated workflows involving multiple specialised AI agents where data flows through multiple processing steps and decision points.

"RAG (Retrieval-Augmented Generation)" means a technique where AI agents search and retrieve information from knowledge bases and documents (for example Jira, Confluence, SharePoint, uploaded files) to enrich responses and decision-making.

"LLM (Large Language Model)" means external AI models (for example OpenAI, Anthropic or others) to which personal data may be transmitted for analysis, generation and reasoning as part of agent workflows.

"Client" means a customer, partner or other entity that uses the Platform to deploy AI agents and process personal data of its own end-users. The Client is the Controller for data processed through its agent scenarios.

"End-user" means an individual who interacts with an AI agent deployed by a Client through the Platform.

"Designated data protection contact" means the person appointed by the Company as the primary contact point for data protection matters (which may be a privacy, legal or security lead, or a Data Protection Officer where one is appointed).

---

3. When to Conduct a DPIA

3.1 Mandatory DPIA Scenarios (GDPR Article 35)

The GDPR specifies that a DPIA shall be required where the proposed processing involves:

1. Systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person.
2. Processing on a large scale of special categories of data referred to in Article 9(1), or of personal data relating to criminal convictions and offences referred to in Article 10.
3. Systematic monitoring of a publicly accessible area on a large scale.

Note. Article 9(1) refers to processing of personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person's sex life or sexual orientation.

3.2 Specific DPIA Triggers

In addition to the mandatory scenarios above, the Company requires DPIAs for projects where one or more of the following applies.

General triggers:

a. Information about living individuals will be collected and processed for the first time.
b. Information about living individuals will be shared with people or organizations that previously did not have access to it.
c. There is a change of use of existing personal data.
d. New technology that collects or uses personal data is introduced (for example biometrics, AI/ML models).
e. Existing personal data will be used to reach decisions as part of an automated process.
f. It might reasonably be expected that an individual may find any aspect of the project intrusive or the data involved private.

AI Agent Platform-specific triggers:

g. Deployment of AI agents that perform automated decision-making with potential legal or significant effects on data subjects.
h. Integration with external LLM providers where personal data will be transmitted to third-party AI models.
i. Processing special categories of data through AI agents.
j. Any other use of the Platform that, in the assessment of the designated data protection contact, presents novel or elevated risks to the rights and freedoms of data subjects.

These triggers reflect the Company's internal risk management practice and are applied on a risk-based basis; they do not alter the legal criteria for mandatory DPIAs under Article 35 GDPR.

3.3 When in Doubt, Conduct a DPIA

If there is uncertainty regarding whether it is appropriate to carry out a DPIA for a specific project, the project team should err on the side of caution and ensure that one is performed. The designated data protection contact may be consulted for clarification, and guidance issued by relevant supervisory authorities should be taken into account where applicable.

---

4. Data Protection Impact Assessment Process

4.1 Process Diagram

The DPIA process follows these steps:

1. Establish the Need and Context
2. Document the Use of Personal Data
3. Identify the Risks
4. Analyse the Risks
5. Evaluate the Risks
6. Define Risk Treatment Plan
7. Obtain Management Approval for Residual Risks
8. Prior Consultation with Supervisory Authority (if required)
9. Implement Risk Treatment Actions
10. Risk Monitoring and Reporting
11. Regular Review

A visual process diagram may be maintained separately as part of the Company's process documentation.

4.2 Establish the Need and Context

The overall environment in which the DPIA is carried out should be described and the reasons for it explained. This includes a description of the internal and external context of the project and its overall objectives.

For AI Agent Platform projects, this includes:

- Deployment model: cloud-hosted, on-premises or hybrid
- Controller/Processor role: whether the Company acts as Controller (for example platform account data, marketing) or as Processor (for Client agent scenarios)
- Client use case description: business problem addressed by AI agents (for example customer service automation, sales lead qualification, compliance document analysis)
- Agent types involved: text chat agents, voice agents, multi-agent orchestrators, tool agents, RAG agents, action agents
- External integrations: LLM providers, CRM/ERP systems, messaging platforms, document repositories, APIs
- Data sovereignty requirements: geographic restrictions, localisation mandates, regulatory constraints

The scope of the DPIA should be clearly defined. This may be expressed in terms of:

- Geographical location (countries, offices, data centres)
- Organisational units (departments, Client organisations)
- Business processes
- IT services, systems and networks (Platform infrastructure, Client systems, sub-processor systems)
- Customers, products or services
- Agent workflows and data flows (which agents process personal data, where data enters and leaves the system)

Any specific exclusions to the scope should be stated with reasons.

4.3 Identify the Risks

The process of identifying risks to the personal data collected, processed and held will consist of:

- Reviewing processing operations and data flows
- Identifying possible threats and vulnerabilities
- Considering potential impacts on data subjects and the Company

Risk identification should involve:

- Managers responsible for business-critical activities
- Representatives of operational teams
- Providers of inputs and recipients of outputs
- Appropriate third parties with relevant knowledge
- Representatives of supporting services (for example engineering, security)
- Client representatives for Processor scenarios

For AI Agent Platform use cases, common risk categories include:

1. Unauthorised access and disclosure — agent logs containing personal data accessed by unauthorised personnel; personal data exposed through RAG to agents/users who should not see it; data leakage through LLM provider logging or training; misconfigured agent permissions allowing cross-Client data access.
2. Automated decision-making risks — agents making decisions with legal/significant effects without human oversight; biased or discriminatory agent behaviour; lack of transparency/explainability in agent decision logic; inability to challenge or contest agent decisions.
3. Data quality and accuracy — agents acting on outdated or inaccurate personal data; RAG retrieving incorrect or irrelevant personal data leading to wrong decisions; data corruption during agent-to-agent transfers.
4. External integration risks — LLM providers retaining personal data beyond permitted scope; Client CRM/ERP systems with inadequate security receiving data from agents; unauthorised third-party access via compromised API keys or webhooks.
5. Multi-agent complexity risks — loss of data lineage in complex multi-agent workflows; difficulty in identifying which agent caused a data breach or error; inability to fully delete data across all agents and integration points.
6. Consent and rights management — inability to identify and honour data subject rights requests across agent logs and RAG repositories; processing beyond consented scope due to agent misconfiguration; difficulty in implementing "right to be forgotten" in distributed agent architecture.

7. Deployment and migration risks — data exposure during cloud-to-on-premises (or vice versa) migration; inadequate security controls in Client on-premises deployments; data sovereignty violations due to unexpected data transfers (for example LLM calls routing through non-EU regions).
8. Special categories of data — agents processing health/biometric/financial data without appropriate safeguards; non-compliance with Article 9 GDPR requirements.

Identified risks should be recorded with sufficient detail to allow likelihood and impact to be assessed, and each risk should have a clearly assigned owner.

4.4 Analyse the Risks

Risk analysis involves assigning a numerical value to:

- The likelihood of the risk occurring; and
- The impact that the risk could have if it materialises.

These values are multiplied to arrive at a classification level (high, medium or low).

Assess the Likelihood

Likelihood is graded on a numerical scale of 1 (low) to 5 (high).

TABLE 1 - RISK LIKELIHOOD GUIDANCE

Grade	Description	Summary
1	Improbable	Has never happened before and there is no reason to think it is any more likely now
2	Unlikely	There is a possibility that it could happen, but it probably will not
3	Likely	On balance, the risk is more likely to happen than not
4	Very Likely	It would be a surprise if the risk did not occur either based on past frequency or current circumstances

5	Almost certain	Either already happens regularly or there is some reason to believe it is virtually imminent

The rationale for allocating the grade given should be recorded to aid understanding and allow repeatability in future assessments.

Assess the Impact

Impact is graded on a numerical scale of 1 (low) to 5 (high), taking into account effects on:

- Data subjects (privacy harm, discrimination, financial loss, identity theft, reputational damage, physical harm)
- Customers (loss of trust, contract termination, compensation claims)
- Finance (fines, legal costs, remediation, lost revenue)
- Health and Safety
- Reputation
- Knock-on operational impact
- Legal, contractual or organisational obligations.

TABLE 2 - RISK IMPACT GUIDANCE

Grade	**Description**	**Customer impact**	**Financial impact**	**Health and Safety**	**Impact on Reputation**	**Legal impact**
1	Negligible	No effect	Very little or none	Very small additional risk	Negligible	No implications
2	Slight	Some local disturbance to normal business operations	Some	Within acceptable limits	Slight	Small risk of not meeting compliance

3	Moderate	Can still deliver product/service with some difficulty	Unwelcome but could be borne	Elevated risk requiring immediate attention	Moderate	In definite danger of operating illegally
4	High	Business is crippled in key areas	Severe effect on income and/or profit	Significant danger to life	High	Operating illegally in some areas
5	Very High	Out of business; no service to customers	Crippling; the organisation will go out of business	Real or strong potential loss of life	Very High	Severe fines (up to 20M EUR or 4% global turnover) and possible imprisonment of staff

The rationale for allocating the grade given should be recorded to aid understanding and allow repeatability in future assessments.

Risk Classification

The product of likelihood and impact defines the risk score. Each risk is then classified as:

- HIGH - 12 or more
- MEDIUM - 5 to 10 inclusive
- LOW - 1 to 4 inclusive

This classification is used as input to the risk evaluation stage.

4.5 Evaluate the Risks

Risk evaluation determines which risks can be accepted and which must be treated.

- LOW risks may be accepted if they fall within predefined risk appetite.
- MEDIUM and HIGH risks will normally require treatment, especially where they relate to:
  - Special categories of data (Article 9 GDPR)
  - Large-scale processing
  - Automated decision-making with legal or similarly significant effects (Article 22 GDPR)
  - Children's data

Risks are prioritised for treatment based on their score and classification so that the most significant exposures are addressed first.

4.6 Define Risk Treatment Plan

For risks that exceed the Company's acceptance criteria, suitable treatment options must be identified.

Treatment options include:

1. Modify the risk by applying controls to reduce likelihood and/or impact.
2. Avoid the risk by changing or not proceeding with certain processing activities (for example not deploying a high-risk agent workflow, choosing on-premises over cloud for special categories).
3. Share the risk with another party (for example insurer, supplier) subject to appropriate contracts and guarantees.

The choice of treatment option should consider:

- Business strategy and objectives
- Regulatory and legislative constraints
- Technical feasibility
- Commercial and contractual factors
- Client requirements (for Processor scenarios)

A designated privacy or legal lead ensures that all relevant stakeholders, including the risk owner, are consulted when choosing treatment options.

Selection of Controls

Controls may draw on:

- ISO/IEC 27001 and Annex A controls
- ISO/IEC 27002 - Code of practice for information security controls
- ISO/IEC 27017 - Code of practice for information security controls based on ISO/IEC 27002 for cloud services
- ISO/IEC 27018 - Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors

- AI-specific guidance (for example emerging AI governance frameworks, NIST AI Risk Management Framework, EU AI Act when in force, EDPB Guidelines on automated decision-making and profiling, ICO guidance on AI and data protection)
- Vendor security documentation and Client security baselines

For AI Agent Platform risks, common controls include:

- Data masking and redaction for LLM traffic
- Role-based access control and least-privilege principles
- Encryption in transit and at rest
- Audit logging of agent actions and data access
- Human-in-the-loop mechanisms for high-risk automated decisions
- Data Processing Agreements with sub-processors
- On-premises deployments for high-risk or special category data
- Regular testing and validation of agents for bias and errors
- Interfaces to support data subject rights across logs and RAG repositories
- Geographic restrictions on data transfers (for example EU-only LLM endpoints)
- Retention policies and automated deletion routines
- Incident response procedures specific to agent-related data breaches.

The outcome of this stage is a documented risk treatment plan.

Data Protection Impact Assessment Report

The evaluation of the treatment options will result in the production of a DPIA report which will detail:

- A description of the proposed processing operations and the personal data involved
- The purposes of the processing including, where applicable, the legitimate interest of the controller as defined by the GDPR
- An assessment of the necessity and proportionality of the processing
- For AI Agent Platform: description of agent types, workflows, LLM integrations, RAG sources, deployment model, and controller/processor roles
- The results of the assessment of the risks to the rights and freedoms of the data subjects
- Whether each risk is recommended for acceptance or treatment
- Priority of risks for treatment
- Risk owners
- Recommended treatment option
- Control(s) to be implemented
- Responsibility for the identified actions
- Timescales for actions
- Residual risk levels after the controls have been implemented.

4.7 Obtain Management Approval for Residual Risks

Once proposed controls are identified and applied, residual risks must be evaluated.

- Top management is accountable for approving residual risks and deciding whether processing may proceed.
- Risk owners confirm acceptance or request further treatment.

In Processor scenarios, the Client (Controller) must approve residual risks for its own use cases.

Sign-off is documented in accordance with Company documentation standards.

4.8 Prior Consultation with Supervisory Authority

If a DPIA indicates that the processing would result in a high risk to data subjects in the absence of measures taken by the Controller to mitigate the risk, and if suitable measures cannot sufficiently reduce the risk, the GDPR (Article 36) requires prior consultation with the competent supervisory authority before processing begins.

Information provided to the supervisory authority includes:

- Responsibilities of the controller, joint controllers and processors, where applicable
- Purposes and means of processing
- Controls implemented to protect personal data
- Contact details for the designated data protection contact (or Data Protection Officer where one is appointed)
- A copy of the DPIA report

The supervisory authority has eight weeks (extendable by a further six weeks) to provide a judgment on the proposed processing and, if appropriate, give details of what must be done to make the processing acceptable under the GDPR.

For AI Agent Platform use cases, high residual risk scenarios likely to trigger prior consultation include:

- Large-scale automated decision-making with legal/significant effects on data subjects
- Processing special categories of data through AI agents without adequate safeguards
- Cross-border transfers to countries without adequacy decisions and high risk of government access
- Novel or untested AI technologies with unclear privacy implications
- Situations where identified risks cannot be sufficiently mitigated within technical or commercial constraints

---

4.9 Implement Risk Treatment Actions

Once the risk treatment plan has been approved, actions should be tracked and completed as part of normal project and operational management.

If actions are delayed or cannot be completed, the impact on personal data protection must be reassessed, and management must decide whether to proceed, adjust or halt the project. If the untreated risk is sufficiently serious, this may have a significant impact on the viability of the project from a compliance viewpoint and advice should be sought from the designated data protection contact and/or the supervisory authority in the country or countries affected.

For AI Agent Platform use cases, implementation may involve:

1. Testing controls in non-production / sandbox environments
2. Pilot deployments with enhanced monitoring
3. Phased roll-out to production
4. Continuous monitoring of key indicators.

4.10 Risk Monitoring and Reporting

Key performance indicators (KPIs) should be defined to measure effectiveness of controls and identify emerging risks.

These indicators will be reported on a regular basis and trend information produced so that exception situations can be identified and dealt with as part of the management review process.

AI Agent Platform-specific KPIs may include:

- Number of unauthorised access attempts to agent-processed data
- Volume of personal data sent to LLM providers (with/without masking)
- Timeliness of data subject rights request fulfilment for agent-related data
- Number and severity of agent errors affecting data subjects
- Incidents of biased or discriminatory agent behaviour detected
- Time to detect and remediate agent-related data breaches
- Compliance with data retention and deletion policies.

4.11 Regular Review

DPIAs are not static. They are reviewed:

- At least annually; and
- Whenever there is a material change, such as:
  - New or changed AI agents or workflows
  - New Client use cases with different risk profiles
  - New LLM providers or sub-processors
  - Changes to deployment models (cloud to on-premises or vice versa)
  - Regulatory changes affecting AI or data protection (for example entry into force of EU AI Act)
  - Incidents or breaches involving the processing in scope
  - Office moves, mergers and acquisitions
  - Introduction of new or changed IT services

Updates to DPIAs and any resulting changes to controls or residual risk decisions are documented.

---

5. Roles and Responsibilities

Effective DPIA execution depends on clear roles and responsibilities.

5.1 RACI Chart

The table below clarifies responsibilities at each step using the RACI model:

- R = Responsible
- A = Accountable

- C = Consulted
- I = Informed

Step	**Assessment Lead**	**Risk Owners**	**Top Management**	**Privacy Lead**
Establish the need and context	R	C	A	C
Document the use of personal data	R	C	A	C
Identify the risks	C	R	A	C
Analyse the risks	C	R	A	C
Evaluate the risks	C	R	A	C
Define risk treatment plan	R	C	A	C
Management approval for residual risks	C	C	A/R	C
Prior consultation with supervisory authority	R	C	A	C

Implement risk treatment plan	R	R	A/R	C
Monitor and report	R	I	A	C
Regular review	R	C	A	C

5.2 Additional Roles

- Client (Controller). In Processor scenarios, the Client is accountable for initiating and approving DPIAs for its own use cases. The Company provides information about the Platform, sub-processors and controls and cooperates with the Client's DPIA process.
- Platform Engineering Lead. Responsible for assessing technical feasibility of controls and implementing them in the Platform.
- Security / Information Security Lead. Consulted on security controls, threat modelling and incident response.

Where the Company appoints a Data Protection Officer, that person's views will be sought on high-risk processing and DPIAs in line with GDPR requirements. If no DPO is appointed, DPIAs are coordinated by the Assessment Lead with support from the Privacy Lead and relevant technical and security experts.

---

6. Conclusion

The DPIA process is fundamental to the implementation of successful projects that handle personal data and is a significant requirement under the GDPR. Only by fully understanding risks to personal data can an organisation ensure that the controls it has in place are sufficient to provide an appropriate level of privacy and meet the standard expected of it.

For a cloud and AI service provider operating an AI Agent Platform, regular assessment of risks to personal data and the application of comprehensive controls is vital to maintaining the confidence of Clients and meeting obligations to protect personal data.

As the Platform evolves to support new AI agents, multi-agent workflows, external LLM integrations and RAG-based processing, the Company recognizes that new and complex privacy risks will emerge. DPIAs are therefore an ongoing governance practice rather than a one-off exercise.

By following this process, the Company helps ensure that the risks it faces in the day-to-day operation of its AI Agent Platform are effectively identified, assessed, managed and controlled, in compliance with GDPR Article 35 and in alignment with the principles of privacy by design and by default.

For questions about the DPIA process, please contact privacy@tovie.ai.