

Data Protection Impact Assessment

1. Purpose

This is Subcontext's assessment of the data protection risks of the platform's core processing: AI agents holding conversations with our customers' customers, and the CRM and insurance features around them. Customers can use it to inform their own impact assessments, which remain their responsibility as controllers.

It follows the structure of the Information Commissioner's Office DPIA template: the need for a DPIA, a description of the processing, consultation, necessity and proportionality, the risks and the measures that reduce them, and sign-off.

2. Need for a DPIA

Subcontext is a processor for the data its customers' customers provide, and a DPIA is formally the controller's duty. We carry out this assessment because the processing has several features that the ICO lists as likely to result in a high risk, and our customers need a sound basis for their own assessments:

Indicator	How it applies
Innovative technology	AI agents hold conversations and take actions using large language models
Special category data	Health information may be collected for insurance applications
Large-scale processing	The platform serves many customers across several channels
Vulnerable individuals	Consumers of financial services may be in vulnerable circumstances
Decisions affecting access to a service	Quotes, applications and payments affect whether a person obtains insurance
Data matching	Leads are combined with data from customers' systems and, where enabled, with address lookup and market quoting

3. Description of the Processing

3.1 Nature

- **Nature:** AI agents converse with individuals on web chat, SMS, WhatsApp, email and voice; capture their details; prepare quotes and applications; take payments; and hand over to the customer's own team
- **Data subjects:** prospective and existing customers of our customers
- **Personal data:** contact details, conversation content, call audio and recordings, application details, payment mandate details
- **Special category data:** health information, where an insurance application requires it, collected only after explicit consent is recorded

- **Recipients:** the customer, and the sub-processors on the published Sub-Processor List
- **Retention:** set by the customer within the limits in the Data Retention & Deletion Policy

Data is collected directly from individuals in conversations and forms, and from customers' own systems through integrations. It is stored in the platform's database and file storage in the AWS Ireland (eu-west-1) region. Text AI processing runs on AWS Bedrock in the UK and EU. Voice calls use speech services in the United States. Card details are never stored by Subcontext: they are handled by Stripe, and Direct Debit bank details are encrypted in the application.

3.2 Scope

Aspect	Detail
Channels	Web chat, SMS, WhatsApp, email and voice
Data	Names, email addresses, phone numbers, postal addresses, dates of birth and other details needed to quote; conversation content; call audio and recordings; application, policy and claim details; Direct Debit mandate details; IP addresses and browser details
Special category data	Health information for insurance applications, only after explicit consent is recorded
Volume	Varies by customer; the platform is designed for continuous use across many customers
Retention	Automatic periods and customer-configurable defaults, set out in the Data Retention & Deletion Policy
Geography	Individuals are mainly in the United Kingdom

3.3 Context

- Customers are UK insurers, brokers and other regulated firms, supervised by the FCA and in some cases the PRA, and subject to the Consumer Duty
- Individuals are members of the public seeking insurance or other financial products, some of whom may be in vulnerable circumstances
- Individuals would reasonably expect a firm they contact to use their details to answer their enquiry and arrange cover, but may not expect to be speaking to an AI unless told
- Customers control which channels, features and integrations are enabled, and configure their agents' knowledge and rules

3.4 Purposes

- To answer enquiries and help individuals find and apply for products, on the customer's behalf
- To prepare quotes and applications, and to take payments and set up Direct Debits
- To hand over to the customer's own staff when needed
- To keep records the customer needs for its regulatory obligations, including a tamper-evident record of compliance decisions

4. Consultation

- **Customers:** customers' requirements, due diligence questions and compliance feedback inform how the platform is designed and how its controls work

- **Sub-processors:** their data protection terms, data locations and model training commitments are assessed under the Supplier Risk Management Policy
- **Individuals:** Subcontext does not consult customers' customers directly. Customers are best placed to seek their views, and complaints and feedback that reach us are taken into account
- **Advisers:** Subcontext will seek independent legal or data protection advice where a change raises a question it cannot resolve internally

5. Necessity and Proportionality

- Processing is performed on the customer's instructions, for the purposes they define
- Only the data needed for the conversation or application is collected
- Customers choose which channels, features and integrations are enabled
- Individuals can opt out of each channel, and every message carries opt-out information where required

Question	Answer
Lawful basis	Determined by the customer as controller, typically steps prior to entering a contract or performance of a contract, legitimate interests, or legal obligation
Special category condition	Explicit consent, recorded before health data is captured, or another condition the customer relies on for insurance purposes
Does the processing achieve the purpose?	Yes. Agents answer enquiries, prepare quotes and applications and hand over to people, which is what customers deploy them for
Is there a less intrusive way?	Customers can limit channels and features to what they need, and the platform collects only what each conversation or application requires
Data quality	Individuals provide their own details; address lookup and phone validation help keep them accurate; customers can correct records in the CRM
Information to individuals	Customers provide privacy information to their customers; an AI disclosure rule can add a disclosure to the first message
Individuals' rights	Supported by the tools described in the Data Subject Rights Procedure
Processors	Bound by data processing terms; listed on the Sub-Processor List
International transfers	Covered by the Transfer Risk Assessment

6. Risks and Mitigations

Each risk is scored for its likelihood and severity of harm to individuals before mitigation, and the residual risk is the assessment after the measures below are applied.

Risk	Likelihood (before)	Severity (before)	Mitigation	Residual risk
The AI gives inaccurate or misleading information	Possible	Significant	Agents answer from the customer's own knowledge and product data; conduct rules for FCA conduct and Consumer Duty; a compliance policy gate checks actions; customers test agents before going live	Low
Vulnerable individuals are not identified or are treated unfairly	Possible	Significant	Conversations are assessed for signs of vulnerability; flagged people are held back from sales and outreach and passed to a person	Low
Individuals do not realise they are talking to an AI	Possible	Minimal	An AI disclosure rule customers can switch on, adding a disclosure to the first message of every conversation	Low
Decisions with significant effects are made without a human	Possible	Significant	Agents can request human review; underwriting rules are set by the insurer; escalation to a person is always available	Low
Special category data is over-collected or misused	Possible	Severe	Explicit consent is recorded before health data is captured; access is limited to the customer's organisation	Low
Manipulation of the AI by malicious input	Probable	Significant	Baseline guardrails against prompt injection and jailbreaks apply to every message and cannot be switched off; untrusted content is clearly separated from instructions	Low
Unauthorised access or a data breach	Possible	Severe	Encryption in transit and at rest, access controls with two-factor authentication, an immutable audit log, and the Incident Response Plan	Low
Data is kept too long	Possible	Minimal	Automatic retention sweeps, configurable per customer	Low
International transfers expose data to foreign authorities	Remote	Significant	Text AI processing is kept in the UK and EU; US transfers are covered by the Transfer Risk Assessment	Low
Data is used to train AI models	Remote	Significant	Our AI providers do not train on the data we send them	Low

7. Additional Measures

Beyond the mitigations in the table, the following measures support the assessment:

Measure	Effect
Conduct rules covering FCA conduct, Consumer Duty, medical disclosure, vulnerable customers, complaints handling, financial promotions and data protection	Reduces the risk of unfair or non-compliant agent behaviour

Measure	Effect
Tool approvals and a compliance policy gate on actions	Prevents agents taking actions the customer has not allowed
A hash-chained compliance decision trail, kept 6 years by default	Lets customers evidence how agent decisions were made
Requests to AI models routed outside the UK and EU are refused	Keeps text conversations in the UK and EU
Sensitive values filtered from logs and sensitive AI tool parameters stored redacted	Limits personal data in operational records
Application-level encryption of Direct Debit bank details and integration credentials	Adds protection beyond database encryption
Audited administrator impersonation	Makes any Subcontext access to a customer's account visible
Complaints tracking with acknowledgement and final response deadlines	Helps customers handle complaints about the service fairly and on time

8. Conclusion

With these measures in place, the residual risk to individuals is assessed as low, and the processing can proceed. Prior consultation with the Information Commissioner's Office is not required.

9. Sign-Off and Outcomes

Item	Outcome
Measures approved by	The directors of Subcontext Ltd
Residual risk approved by	The directors of Subcontext Ltd
Advice from the person responsible for data protection	The Security Officer reviewed the assessment and agrees the processing can proceed with the measures above
Integration into the platform	The measures are built into the platform and are checked when the processing changes
Prior consultation with the ICO	Not required, as no high residual risk remains

10. Review

This assessment is reviewed at least annually, and whenever the processing changes significantly: for example a new channel, a new category of data, or a new AI provider.

Changes to the platform are screened during design to decide whether this assessment needs updating, and the outcome of each review is recorded.

11. Related Documents

- Transfer Risk Assessment
- Sub-Processor List
- Data Retention & Deletion Policy
- Data Subject Rights Procedure
- Data Classification Policy
- AI Ethics & Responsible Use Policy
- Incident Response Plan
- Supplier Risk Management Policy

12. Document Control

Item	Detail
Owner	Security Officer
Approved by	The directors of Subcontext Ltd
Review	At least annually, and when the processing changes significantly
Effective	October 2026

Subcontext Ltd, registered in England and Wales, Company No. 17019645. Registered office: 124 City Road, London, EC1V 2NX.
The current version of this document is published at subcontext.com.