DPIA

This DPIA follows the process set out in the ICO's DPIA guidance and uses the ICO template.

Submitting controller details

Name of controller	Quinly Ltd
Subject/title of DPO	Data Protection Officer and founder
Name of DPO	Sunita Gordon hello@quinly.ai

Step 1: Identify the need for a DPIA

Explain broadly what project aims to achieve and what type of processing it involves. You may find it helpful to refer or link to other documents, such as a project proposal. Summarise why you identified the need for a DPIA.

Quinly 2.0 is an advanced AI-powered crisis detection and support chatbot designed for British children and young people aged 7-17 who are in distress. This premium version extends Quinly Basic with enhanced features including multilingual support, comprehensive analytics dashboards, sentiment analysis, and crisis trend reporting for school administrators. The system provides immediate empathetic responses using CBT techniques and connects children to trusted UK support services including Childline, Samaritans, and other national support services.

DPIA is required due to high-risk processing involving:

- Children (vulnerable individuals under GDPR)
- Special category data (mental health information under GDPR Article 9)
- Large-scale processing across multiple educational institutions
- Systematic monitoring and profiling through sentiment analysis
- New technology deployment using AI and automated decision-making
- Commercial service model with data retention for analytics purposes

The processing triggers multiple DPIA criteria under UK GDPR Article 35 and must comply with the UK Children's Code (Age Appropriate Design Code) as an information society service likely to be accessed by children.

Step 2: Describe the processing

Describe the nature of the processing: how will you collect, use, store and delete data? What is the source of the data? Will you be sharing data with anyone? You might find it useful to refer to a flow diagram or other way of describing data flows. What types of processing identified as likely high risk are involved?

Data Collection:

- Anonymous conversation messages during chat sessions
- Crisis detection patterns and sentiment indicators
- School-specific usage analytics and interaction trends
- Technical data (timestamps, session duration, crisis categories detected)
- School administrator access logs and dashboard usage
- No user accounts, registration, or personally identifiable information collected
- No tracking cookies or behavioural profiling across sessions

Data Use:

- Real-time processing through Anthropic Claude AI for supportive response generation
- Advanced crisis detection using keyword pattern matching and sentiment analysis
- Generation of anonymised analytics reports for school safeguarding teams
- Trend analysis for institutional mental health monitoring
- Admin dashboard population with school-specific aggregated data
- Automated PDF report generation for safeguarding purposes

Data Storage:

- Anonymised interaction data retained for up to 3 years for analytics purposes
- Crisis pattern data stored for trend analysis and reporting
- Sentiment analysis scores aggregated at school level
- Admin access logs retained for 2 years for security purposes
- All data stored with strong encryption and access controls
- Physical data isolation between different school tenants

Data Sharing:

- Anonymised, aggregated analytics shared with subscribing schools only
- No individual conversation content shared with schools or third parties
- No cross-school data sharing - complete tenant isolation
- Temporary processing by Anthropic AI service for response generation
- No marketing or commercial use beyond contracted services

Describe the scope of the processing: what is the nature of the data, and does it include special category or criminal offence data? How much data will you be collecting and using? How often? How long will you keep it? How many individuals are affected? What geographical area does it cover?
Nature of Data: • Regular personal data: Anonymous conversation content, technical metadata • Special category data: Mental health discussions, crisis indicators, emotional state analysis • No identification data: Names, addresses, contact details actively refused and filtered • Sensitive analytics: Suicide risk indicators, self-harm detection, abuse disclosures • Administrative data: School access patterns, dashboard usage, report generation logs **Data Volume:** • Large-scale processing: Estimated 10,000+ student interactions monthly across subscriber schools • Individual sessions: 5-30 messages per conversation • Analytics processing: Continuous sentiment analysis and trend computation • Long-term retention: Anonymised crisis patterns for institutional reporting **Retention Period:** • Crisis interaction analytics: 3 years after creation for trend analysis • Sentiment analysis data: 1 year in anonymised, aggregated form • Admin access logs: 2 years for security and audit purposes • Raw conversation content: Immediate deletion after session completion • System backups: 30-day rolling retention with cryptographic deletion **Individuals Affected:** • Children and young people aged 7-17 accessing crisis support • Estimated user base: 50,000+ students across participating UK schools annually • Anonymous users maintaining complete anonymity throughout processing • School administrators accessing dashboard analytics (up to 3 per school) **Geographical Coverage:** • United Kingdom exclusively • Multilingual processing: English, Welsh, Urdu, Polish, Arabic, Romanian • UK-specific crisis resources and support service integration

Describe the context of the processing: what is the nature of your relationship with the individuals? How much control will they have? Would they expect you to use their data in this way? Do they include children or other vulnerable groups? Are there prior concerns over this type of processing or security flaws? Is it novel in any way? What is the current state of technology in this area? Are there any current issues of public concern that you should factor in? Are you signed up to any approved code of conduct or certification scheme (once any have been approved)?

Relationship with Individuals:

- Anonymous crisis support service with no direct contact or relationship
- Commercial service relationship with schools (£495 annual subscription)
- Joint controller arrangement with schools for analytics purposes

Individual Control:

- Users maintain complete control over conversation participation
- Immediate conversation termination available at any time
- No forced participation or mandatory usage requirements
- Privacy controls embedded throughout user interface

User Expectations:

- Anonymous, confidential crisis support with immediate deletion of conversations
- School analytics processing clearly communicated through privacy notices
- Compliance with UK Children's Code expectations for child-appropriate privacy
- Transparent AI processing with clear limitations messaging

Vulnerable Groups:

- Primary focus on children and young people in crisis situations
- Heightened protection under UK Children's Code and GDPR Article 8
- Students experiencing mental health difficulties, suicidal ideation, abuse
- Enhanced safeguards for under-13s requiring parental awareness

Technology Context:

- Advanced AI processing using Anthropic Claude for crisis-appropriate responses
- Novel application: AI-powered crisis detection specifically designed for educational settings
- Emerging field: Institutional mental health monitoring through conversational AI
- Public concerns: AI bias, mental health data security, educational surveillance

Compliance Framework:

- UK Children's Code (Age Appropriate Design Code) compliance
- GDPR Article 8 enhanced protection for children
- ICO guidance on AI and automated decision-making
- Educational sector data protection best practices

Describe the purposes of the processing: what do you want to achieve? What is the intended effect on individuals? What are the benefits of the processing – for you, and more broadly?

Primary Objectives:

- Provide immediate, accessible crisis support to children and young people
- Enable early detection and intervention for mental health crises
- Support schools' statutory safeguarding and pastoral care responsibilities
- Deliver institutional analytics for evidence-based wellbeing strategies
- Connect vulnerable young people to appropriate professional resources
- Provide 24/7 support when human services are unavailable

Intended Effects on Individuals:

- Immediate emotional support during crisis situations
- Reduced isolation and improved access to help
- Connection to appropriate professional support services
- Enhanced safeguarding protection through institutional monitoring

Benefits:

- For children: Immediate crisis support, improved safety, connection to help
- For schools: Enhanced safeguarding capabilities, data-driven wellbeing insights, 24/7 availability
- For society: Improved child protection, early intervention, reduced suicide risk
- For Quinly: Sustainable business model enabling continued service development

Step 3: Consultation process

Consider how to consult with relevant stakeholders: describe when and how you will seek individuals' views – or justify why it's not appropriate to do so. Who else do you need to involve within your organisation? Do you need to ask your processors to assist? Do you plan to consult information security experts, or any other experts?

Stakeholder Consultation:

- Technical security review with school IT departments and Multi-Academy Trust technology leads
- Consultation with Designated Safeguarding Leads across subscriber schools
- Review with school Data Protection Officers and business managers
- Engagement with headteachers and senior leadership teams regarding institutional analytics

Individual Consultation:

- Direct consultation not appropriate due to anonymous service design and crisis context
- Child-friendly privacy notices provided with age-appropriate language
- Clear opt-out mechanisms and control information provided
- School-level consultation through parent/carer communication about service availability

Expert Review:

- Child protection specialist review of crisis detection algorithms and response protocols
- Independent information security audit of multi-tenant architecture
- GDPR compliance review by external data protection specialist
- Educational sector consultation on safeguarding analytics appropriateness
- School safeguarding professional review of AI response quality and safety

Step 4: Assess necessity and proportionality

Describe compliance and proportionality measures, in particular: what is your lawful basis for processing? Does the processing actually achieve your purpose? Is there another way to achieve the same outcome? How will you prevent function creep? How will you ensure data quality and data minimisation? What information will you give individuals? How will you help to support their rights? What measures do you take to ensure processors comply? How do you safeguard any international transfers?

Lawful Basis for Processing (Article 6):

- Article 6(1)(b) - Contract: Processing necessary for service delivery under £500 annual school subscription
- Article 6(1)(f) - Legitimate Interest: Providing crisis support to children in need
- Necessity test: All processing directly essential for crisis support and institutional safeguarding
- Balancing test: Child welfare and institutional safeguarding benefits clearly outweigh privacy risks

Special Category Data Basis (Article 9):

- Article 9(2)(g) - Substantial Public Interest: Child protection and safeguarding under UK law
- Article 9(2)(h) - Health: Mental health support and crisis intervention

UK Children's Code Compliance:

- Best interests of child primary consideration in all processing decisions
- Age-appropriate privacy notices and controls implemented
- Minimal data collection with strong purpose limitation
- Enhanced security measures for children's data
- Regular impact assessments for child-specific risks

Processing Necessity:

- Crisis response generation requires conversation context understanding
- Sentiment analysis essential for early warning systems
- Analytics processing necessary for institutional safeguarding obligations
- All processing directly linked to child protection and educational purposes

Alternative Methods Assessment:

- Considered human-only support: Insufficient 24/7 availability and scale
- Evaluated external counselling referral only: Misses early intervention opportunities
- Assessed non-AI solutions: Inadequate crisis detection and response quality
- Current approach represents least intrusive method achieving safeguarding objectives

Function Creep Prevention:

- Strict purpose limitation documentation and technical controls
- Regular processing audits against original DPIA purposes
- Technical architecture prevents unauthorised data access or use
- Contractual restrictions with schools on data use scope

Data Minimisation:

- Anonymous processing eliminates unnecessary identification data
- Immediate deletion of conversation content after session completion

- Analytics limited to aggregated, non-identifiable patterns
- No collection of contact details, behavioural tracking, or profiling data

Individual Rights Support:
- Transparency: Child-friendly privacy information prominently displayed
- Control: Immediate conversation termination and clear service limitations
- Erasure: Automatic deletion processes and manual deletion procedures
- Rectification: Analytics correction procedures for school administrators
- Portability: School analytics export functionality where applicable

Limitation of Individual Rights due to Anonymisation:
- **Access, rectification, erasure requests**: Cannot be fulfilled for conversation data due to genuine anonymisation - no mechanism exists to identify which anonymous interactions belong to specific individuals
- **GDPR Recital 26**: When data is genuinely anonymous, GDPR does not apply to that processing
- **ICO guidance**: Rights limitations acceptable when true anonymisation prevents individual identification
- **Mitigation**: Users informed upfront that conversations are immediately anonymised and cannot be retrieved or modified
- **School analytics**: Subject access requests can be fulfilled for school-level aggregated data where applicable"

Processor Compliance:
- Data Processing Agreements with Anthropic AI and hosting providers
- Regular audit rights and security assessment requirements
- Contractual restrictions on data use, retention, and sharing
- Technical and organisational measures verification processes

International Transfers:
- Anthropic AI processing subject to UK adequacy decisions and standard contractual clauses
- No routine international transfers of analytics data
- UK/EU hosting with specific data localisation requirements
- Enhanced safeguards for any necessary international processing

Step 5: Identify and assess risks

Describe source of risk and nature of potential impact on individuals. Include associated compliance and corporate risks as necessary.	**Likelihood of harm**	**Severity of harm**	**Overal l risk**
Unauthorised access to anonymised crisis interaction data during processing or storage	Remote	Medium	Low
Cross-contamination of data between different schools in multi-tenant environment	Remote	High	Medium
Inappropriate AI responses to children in crisis situations leading to harm	Possible	High	Medium
Re-identification of children through analytics patterns or sentiment analysis	Remote	High	Medium

System unavailability during critical safeguarding incidents	Possible	Medium	Low
Analytics data breach exposing school-level mental health trends	Remote	Medium	Low
Admin access token compromise leading to unauthorised dashboard access	Remote	Medium	Low
Retention of data beyond stated periods due to technical failure	Remote	Medium	Low
Children sharing personal information despite anonymisation design	Possible	Low	Low
Third-party AI processor retaining or misusing conversation data	Remote	High	Medium

Step 6: Identify measures to reduce risk

Identify additional measures you could take to reduce or eliminate risks identified as medium or high risk in step 5

Risk	Options to reduce or eliminate risk	Effect on risk	Residual risk	Measure approved
Unauthorised data access	End-to-end encryption, access logging, role-based controls, regular security audits	Reduced	Low	Yes
Cross-school data contamination	Physical database isolation, cryptographic tenant separation, automated validation	Eliminated	Low	Yes
Inappropriate AI responses	Content filtering, crisis-specific prompts, human escalation protocols, response quality monitoring	Reduced	Low	Yes
Re-identification through analytics	K-anonymity techniques, statistical disclosure control, aggregation thresholds	Reduced	Low	Yes
System unavailability	Redundant infrastructure, 24/7 monitoring, automatic failover, clear fallback messaging	Reduced	Low	Yes
Analytics breach	Encryption at rest and in transit, access logging, penetration testing, incident response procedures	Reduced	Low	Yes
Admin token compromise	Cryptographically secure tokens, session management, suspicious activity monitoring	Reduced	Low	Yes
Data retention failures	Automated deletion schedules, audit trails, manual verification processes, backup encryption	Eliminated	Low	Yes
Personal information sharing	Active detection and filtering, privacy education, conversation monitoring, clear guidelines	Reduced	Low	Yes
Third-party processing risks	Data Processing Agreements, audit rights, contractual restrictions, alternative provider assessment	Accepted	Low	Yes

Step 7: Sign off and record outcomes

Item	Name/position/date	Notes
Measures approved by:	Ruth Sparkes, Developer, Quinly Ltd / 23 Aug 25	All technical and procedural safeguards implemented and verified. Medium risk processing requires enhanced controls but provides substantial child protection benefits. Multi-tenant architecture includes strong isolation measures.
Residual risks approved by:	Ruth Sparkes, Developer, Quinly Ltd / 23 Aug 25	All residual risks assessed as low following implementation of comprehensive safeguards. No high risks remaining requiring ICO consultation. Enhanced protection measures address children's data processing requirements.
DPO advice provided:	Sunita Gordon, **Data protection Officer**, Quinly Ltd / 03 Sept 25	Processing complies with GDPR, UK Children's Code, and educational sector requirements. Legitimate interest and contractual bases provide strong legal foundation. Enhanced safeguards appropriate for children's special category data processing. Analytics functionality justified by institutional safeguarding benefits. Recommend annual review given evolving AI technology and regulatory landscape.
Quinly 2.0 processing is lawful, necessary, and proportionate with appropriate safeguards implemented. The service addresses critical child protection needs while maintaining strong privacy protections. Enhanced features provide substantial benefits to educational institutions' safeguarding responsibilities without compromising individual privacy rights. Processing approved for deployment with mandatory annual review and continuous monitoring of risk controls.		
DPO advice accepted or overruled by:	Ruth Sparkes, Developer, Quinly Ltd / 2 September 25	All DPO recommendations accepted and integrated into project plan. Processing approved for educational institution deployment.

Comments: Low risk processing ready for deployment. Stateless design provides enhanced privacy protection compared to typical chatbot services. Enhanced safeguarding capability for schools with minimal privacy impact.		
Consultation responses reviewed by:	Sunita Gordon, DPO, Quinly Ltd / 5 Sept 25	School and expert consultation feedback incorporated into final design. Educational institutions support deployment with implemented safeguards.
Comments: Positive feedback from schools regarding enhanced safeguarding capabilities. Technical review confirms robust privacy protection through stateless architecture.		
This DPIA will kept under review by:	Sunita Gordon, **Data protection Officer**, Quinly Ltd / 5 Sept 25	The DPO should also review ongoing compliance with DPIA