Data Protection Impact Assessment

Trellis · v2.8 · October 2026

Document Control

Version	Author(s)	Date	Summary of changes
1.0	Andy Bell, Silva Arapi, Tash Whitaker	17 Jun 2025	Initial draft of DPIA for new Project Plan (e.g. going straight to a transcriber)
1.1	Andy Bell	8 Jul 2025	Tweaks to clarify the models used and the process of optimisation
1.2	Andy Bell, Frank Rankin	30 Aug 2025	Review prior to multi-council workshop. Significant updates to the 'tolerable risk levels' questions
1.3	Harriet Owen, Andy Bell	8 Oct 2025	Clarifications and simplifications
1.4	Andy Bell	20 Jan 2026	Change to 60 days data retention. Clarify EU failover
1.5	Andy Bell, Frank Rankin	18 Feb 2026	Tweaks for ISO 27001 process (e.g. adding Google Workspace as a subprocessor)
2.0	Andy Bell	12 May 2026	Making the DPIA work for English as well as Scottish schools
2.1	Andy Bell	2 Jul 2026	Remove 'Confidential – Internal use only' from footer
2.2	Andy Bell, Frank Rankin	20 Jul 2026	Aligned with Technical Specification v3.1: Azure Speech transcription (replacing Whisper), native iPad app, client-side encryption, shared-schema database with Row-Level Security, PII expiry scrubbing. Added treatment of user (teacher/staff) data as a distinct category. Clarified evaluation and optimisation activities.
2.3	Andy Bell, Frank Rankin	27 Jul 2026	Clarified MFA wording (enforced for Trellis managed accounts; controller policy for federated sign-in); Added 14-day retention of intermediate transcripts and seven-year retention of audit logs; Removed EU failover
2.4	Harriet Owen, Frank Rankin, Andy Bell	10 Aug 2026	Added text relating to the Management Panel; added Notion Labs, Inc. as a sub-processor (internal business systems; no children's data)
2.5	Frank Rankin	19 Aug 2026	Added text responding to feedback from August LA meetings
2.6	Andy Bell	7 Sep 2026	Wording made nation-neutral for English and Scotland: 'council' / 'local authority' replaced with 'controller' where the controller may be a trust; 'Child's Plan' generalised to support plans; Scottish pack reference removed from the front page; scale statement updated
2.7	Andy Bell	20 Sept 2026	Clarified data residency: Trellis content UK (UK South / UK West backups); staff identity data in Microsoft Entra External ID held in the EU under UK EEA adequacy (not a restricted transfer). Per management review 17 Sep 2026 and Frank Rankin advice.

Version	Author(s)	Date	Summary of changes
2.8	Andy Bell/Freddie Paoli	16 Oct 2026	Removal of Notion and Google as sub-processors. Removal of ‘acceptable’ from the risk assessment. Wording clarification at Step 3 question 5.

Contents

Step 1: Describe the project and why a DPIA is required

Step 2: Describe the procedures

Assess the data processors

Step 3: Assess necessity and proportionality

Step 4: AI system risk assessment

Step 5: Consultation process

Step 6: Identify and assess risks

Step 7: Sign off and record outcomes

Step 1: Describe the project and why a DPIA is required

Project Name	Trellis
Project Description	Trellis is designed to reduce the administrative burden on teachers supporting children with additional needs by automating the drafting of plans and minutes. A teacher records the relevant meeting through the Trellis web app or the native Trellis iPad app. Trellis transcribes the recording using Microsoft Azure Speech and generates high-quality, editable draft documents using Azure OpenAI API, based on the transcript and other inputs the teacher provides (e.g. pastoral notes, student and parent views). The teacher reviews and edits every draft before it is used. Trellis makes no decisions about any child.
Brief reason(s) why a DPIA should be carried out	Trellis, as processor, is completing a DPIA because we operate a SaaS solution that uses AI and processes special category data of vulnerable individuals. The controller should determine their own reason for carrying out a DPIA. This document is intended to support each controller's own assessment.
Name of Project Owner	Andy Bell, CEO, Trellis Education Ltd

Step 2: Describe the procedures

Describe the scope of the data being processed

What is the nature of the data, and does it include special category or criminal offence data?	Two distinct categories of personal data are processed: **1. Children's data.** Personal data about children with additional support needs: names, support needs, health-related notes, pastoral notes, and recordings of Team Around the Child (and equivalent) meetings. This includes special category data under Article 9 UK GDPR (e.g. health data). **2. User (teacher/staff) data.** Personal data about the teachers and the school, trust and local authority staff who use Trellis: name, work email address, school and organisation, identity token identifiers (Entra Object ID), and operational records generated by the service: activity and audit logs (actor, action, entity, IP address, user agent, trace identifier), error logs, and survey and feedback responses. This category is ordinary personal data, retained for security, accountability and service improvement, and is treated separately from children's data throughout this DPIA and the DPA.
How will the processing work?	**User journey for creation of a Support Plan** **1. The teacher signs in and starts the recording.** The teacher signs in via Azure Entra External ID (OpenID Connect). Multi-factor authentication is enforced for all accounts hosted in the Trellis tenant. Where sign-in is federated to the controller's own identity provider, the controller's authentication policy (including MFA) applies. The teacher enters attendee names and starts the recording on a school laptop, desktop or iPad, using the Trellis web app or the native Trellis iPad app. **2. The meeting is recorded.** These meetings are wide-ranging needs assessments where the teacher, student, parents and other stakeholders discuss education and emotional support needs. Audio is captured in approximately two-second chunks. Each chunk is encrypted on the device before upload, so audio is never transmitted or stored locally in plaintext, and can only be decrypted server-side. If connectivity drops, encrypted chunks are held on the device (browser IndexedDB, or an encrypted on-disk queue on iPad) and uploaded automatically when the network returns. Nothing is stored on the device once the server has accepted the recording. Data involved: audio file of the meeting recording, attendee names as entered by the teacher, date and time of recording. Storage: Azure Blob Storage, UK South region.

3. Teacher ends the recording. The encrypted chunks are merged server-side. Every uploaded file is scanned by Microsoft Defender for Storage before any further processing; a malicious verdict stops the pipeline.

4. Azure Speech creates a transcript. Microsoft Azure Speech batch transcription transcribes the audio. The transcription service processes the audio but does not retain it. The transcript is stored on Azure servers in the UK.

5. Teacher adds context (optional). The teacher can type in the views of the parent and/or student.

6. Azure OpenAI creates the draft documents. The Azure OpenAI service generates the Plan and Minutes from the transcript and the teacher's inputs, using per-document-type prompts, structured (JSON) output and zero temperature. Azure OpenAI does not store the inputs or outputs, and none of the data is used to train models or available to any other customer. Azure OpenAI processing runs in the Azure UK South region. Output is validated against schemas, normalised to British spelling, and rendered into each authority's own Word templates. High-severity content-filter events alert a Trellis operator; they do not block the pipeline.

7. Teacher reviews, edits and downloads. The teacher reviews the draft in Trellis, can edit it in the app, and downloads the finished documents.

8. Teacher copies documents to the school system. The teacher saves the documents into the school's own file system and continues to edit there as required.

Retention and deletion. By default Trellis retains meeting recordings, generated documents and associated pupil data for 60 days, unless a different retention period is agreed with the controller. Intermediate transcripts are retained for 14 days for quality assurance and accuracy checking. System audit logs, which may contain staff personal data but no meeting content, are retained for seven years for security and accountability purposes. These retention arrangements are documented in the Data Processing Agreement.

Every record carries a PII expiry date. A background scrub process runs continuously: when the expiry passes, audio and documents are deleted from storage and identifying database fields are overwritten, and the deletion is itself written to the audit trail. The application also filters expired data out of responses, so expired data cannot be viewed even before a scrub cycle completes.

Data can also be deleted from Trellis earlier than the defined retention period on instruction from the controller.

Evaluation and optimisation. To evaluate and improve output quality, a small number of Trellis staff (no more than five, named, BPSS-vetted, under confidentiality obligations) review inputs and outputs, including identifiable recordings, transcripts and drafts, strictly under the instructions set out in the Data Processing Agreement. Increasingly, evaluation uses in-app editing data (what teachers change in drafts) rather than direct review of content. We improve quality through prompt engineering and model configuration only. We do not fine-tune models, and no children's data is used to train any AI model; neither ours nor Microsoft's. School-controlled personal data processed through Trellis is not transferred into Trellis's internal business administration systems.

Reporting and metrics. To provide controllers (local authorities and trusts) with management information and metrics, and to allow the analysis of trends over time, controllers will be given access to a Management Panel. This will give visibility of quantitative data on the use of the system within their organisation, with the possibility of reporting at school level.

The values included in the reporting panel are likely to include:

- Aggregated number of successful meetings: cumulatively (graph), week by week (graph), broken down by school (as a table)
- Number (or %) of active users within the organisation

	• Table/database broken down into schools with:
	  – total number of users (count);
	  – number of users that have completed the in-app walkthrough (count);
	  – number of users that have recorded a meeting (count);
	  – active users (recorded in the last 30 days) (count);
	  – churned users (recorded before but not in the last 30 days) (count).
	Future iterations of the Management Panel may include values for:
	• Average number of attendees at each meeting (past 30/60 days, as the attendee list is deleted after 60 days)
	• Percentage of meetings where the young person attended (past 30 days, as the attendee list is deleted after 60 days)
	• Average duration of meetings (past 30 days)
	While the data will be quantitative only, with no identifiers, we recognise that in some cases smaller values will potentially allow a controller to identify specific system users within their organisation.

Where will the data/processing be located (which country/countries)?	**Trellis content (pupil and meeting data): stored at rest and processed in Microsoft Azure UK South. This covers the application, database, recordings storage, speech transcription and document generation. Geo redundant backups remain within the UK geography (UK West).** **Staff identity / authentication data: Microsoft Entra External ID holds staff account data (name, work email, unique ID, and sign in and sign out logs) in the EU. There is no UK Entra External ID tenant option. This does not include pupil data, recordings, transcripts or documents.** **Restricted transfers: Trellis content stays in the UK. Staff identity data in the EU is held under the UK's adequacy arrangements for the EEA and is not a restricted transfer; no additional transfer safeguard is required. (Confirmed with Trellis's data protection adviser, Sep 2026; change of understanding, not of processing.)**
Where are the data subjects based?	All data subjects are in the UK: children in Scottish and English schools, and the school, trust and local authority staff who support them.
How much data do you collect and use?	Trellis is in use by several Scottish local authorities and in many English schools and multi-academy trusts. Each child's record typically contains a small number of structured fields (e.g. name), the meeting recording and transcript, any documents the teacher uploads, and the generated drafts. At national roll-out, Trellis would process data relating to thousands of children per year across participating authorities.
How many individuals are affected?	Currently, the children whose meetings are recorded in the participating schools, their supporting adults, and the staff who use the system. At national roll-out, potentially any child with additional support needs in a participating authority, subject to each controller's decisions about scope.

Describe the context of the personal data

What is the nature of your relationship with the individuals?	The children whose data is processed are not direct users of Trellis. Their data is entered by teachers and educational professionals acting on behalf of the data controller. Trellis is the data processor. Teachers and staff are direct users; Trellis processes their account and activity data in order to provide, secure and improve the service, as instructed by the controller.

How much control will they have?	It is the responsibility of the controller to decide how much control children or parents have in terms of participation. Trellis will advise on consent flows and parent communications about the system. By default Trellis retains meeting/pupil data for 60 days (configurable per authority). Children and parents do not interact with the system directly, but because processing is carried out by public authorities delivering statutory educational support, they retain all their UK GDPR rights, access, rectification, erasure, objection, exercisable through the data controller.
Would they expect you to use their data in this way?	The tool is used by qualified educators as part of their statutory responsibilities, so the processing aligns with the broader expectations of educational support. This is for the controller to determine. Trellis provides a transparency statement flow at the start of recording: parents and students are told at the outset of a meeting that it is being recorded using Trellis. It is the controller's responsibility to determine lawful basis and to provide privacy notices; Trellis provides technical descriptions of the processing to support this, and can surface the controller's privacy notice in the product if required.
Do they include children or other vulnerable groups?	Yes. The data subjects include children receiving additional support: a vulnerable group requiring heightened protections and careful handling of personal and special category data.
Are there prior concerns over this type of processing or security flaws?	There are no known security breaches involving Trellis. The use of AI in public sector education is a relatively new area, and there may be public concerns around transparency, accuracy and fairness. Microsoft disclosed a security incident involving Azure OpenAI in January 2025 (rooted in credential theft in late 2024), which allowed misuse of AI services. Microsoft responded with legal action, infrastructure seizure and strengthened security processes; no new breaches have been publicly reported since. The use of generative AI with sensitive data requires careful oversight, including clear guardrails on how input data is handled, verification that submitted content is not used to train models, human review of all outputs, and transparent information to controllers, educators and families about how AI is being used to support — not replace — professional judgement. Trellis works with qualified data protection advisors to meet UK GDPR requirements.
Is it a different way of working in any way?	Trellis introduces generative AI to assist teachers in drafting documents, which has not previously been part of standard practice in ASN/SEND education. Teachers already use digital systems to record and manage pupil information, but AI-assisted drafting is novel in this sector. It is designed to reduce admin workload; it raises new considerations around transparency, explainability and oversight, which this DPIA addresses.
What is the current state of technology in this area?	The underlying AI services (Azure Speech, Azure OpenAI) are mature and widely deployed. Trellis has been in live use in Scottish schools since December 2025. Application of large language models in child education planning remains a developing practice requiring close governance and human oversight, which is why every output is reviewed by a qualified professional.
Are you signed up to any approved data protection code of conduct or certification scheme?	There is no approved UK GDPR Article 42 certification scheme applicable to this processing. Trellis is certified to ISO 27001 (certificate: 271956, dated 23rd April 2026), with a scope covering the production service and supporting administrative systems, and undergoes annual third-party penetration testing.

Describe the purposes of the processing

What do you want to achieve?	The goal of Trellis is to reduce the administrative burden on teachers by providing an AI-powered tool that drafts support plans (for example Child's Plans, EHCP contributions

	and One Page Profiles), minutes and related documents efficiently. By automating the first draft from teacher inputs, Trellis frees teachers' time to focus on supporting children directly, and helps ensure plans are consistent, structured and aligned with relevant standards.
What is the intended effect on individuals?	**Positive effects.** For children: faster and more consistent preparation of plans may improve the timeliness and quality of support. For teachers: a significant reduction in documentation time, lowering workload and burnout risk. For schools and authorities: improved efficiency and consistency in documentation. **Potential risks.** For children: inaccurate or poorly worded drafts risk misrepresenting a child's needs. For teachers: possible over-reliance on the tool, or uncertainty over accountability for the final output. For families: AI processing of sensitive information may raise concerns about transparency and trust if not fully explained. All outputs are reviewed and edited by qualified professionals; no decisions are made by AI.
What are the benefits of the processing for you, and more broadly?	For Trellis: it advances our mission of reducing time spent on ASN/SEND admin in schools. For schools, trusts, councils and other public bodies: efficiency and standardisation of statutory documentation. Societally: teacher time reallocated from paperwork to direct child support, and a contribution to responsible innovation in public sector AI.

Assess the data processors

Use this table to assess how well the selected processor(s) meet UK GDPR requirements. Trellis engages Microsoft as a sub-processor to provide the cloud infrastructure, authentication, transcription, AI document generation and security services required to deliver Trellis. The assessment below covers the processing of personal data on behalf of controllers.

Data Processor: Microsoft

Name of Data Processor	Microsoft (contracting entity: Microsoft Ireland Operations, Ltd.)
Headquarters of Data Processor	USA (parent); Dublin, Ireland (contracting entity)
Name and contact details of the Processor DPO	Data subjects may contact the Microsoft DPO via the webform at https://aka.ms/privacyresponse, or by post: Microsoft EU Data Protection Officer, One Microsoft Place, South County Business Park, Leopardstown, Dublin 18, D18 P521, Ireland. Telephone: +353 (1) 706-3117. Source: https://learn.microsoft.com/en-us/compliance/regulatory/gdpr-data-protection-officer#contact
Services provided for this processing	Cloud infrastructure (Azure: compute, PostgreSQL database, Blob Storage, Key Vault, networking), identity (Entra External ID), AI model hosting (Azure Speech transcription, Azure OpenAI), and malware scanning (Defender for Storage).
Has the DPO confirmed that the Data Processing Agreement fulfils all the legal requirements of the ICO/UK GDPR?	Yes. The Microsoft Products and Services Data Protection Addendum has been assessed and the assessment can be shared with the Controller. DPA: https://www.microsoft.com/licensing/docs/view/Microsoft-Products-and-Services-Data-Protection-Addendum-DPA
List of sub-processors for this processing	Per product, published at: https://servicetrust.microsoft.com/DocumentPage/badc200c-02ab-43d9-b092-ed9b93b9b4a8

What countries store the controller's data (including the location of the processors)?	Trellis content (pupil and meeting data) at rest: United Kingdom (Azure UK South; geo redundant backups within the UK geography, UK West). Staff identity and authentication data: European Union (Microsoft Entra External ID). No pupil data, recordings, transcripts or documents are held in Entra.
What International Data Transfer Mechanism are you using?	None is required for UK/EEA processing including Microsoft Entra External ID staff identity data hosted in the EU. (UK adequacy covers the EEA). Microsoft is additionally certified to the UK Extension to the EU-U.S. Data Privacy Framework. Our contract is signed with the Irish entity, Microsoft Ireland Operations Limited, which is the processor for UK and EEA customers. Microsoft Corporation may act as an onward sub-processor to Microsoft Ireland under Standard Contractual Clauses; it is not a sub-processor of Trellis.
What access does the processor or sub-processors have to the controller's data and under what circumstances?	Microsoft can in principle access data stored in its cloud platform, but note that meeting audio is additionally encrypted client-side by Trellis before upload, so an operator with disk-level access to storage cannot read a recording. Microsoft support access is governed by the DPA and logged.
Locations where staff may be based that will have access to the data	Per Microsoft's published sub-processor and staff-location documentation (link above).
Technical and organisational measures in place to protect the Controller's data	Microsoft's measures are set out in Appendix A (Security Measures) of the Microsoft DPA (link above). In summary they cover: security ownership and personnel confidentiality obligations; asset inventory and handling controls; security training; physical and environmental security of data centres; data recovery procedures (multiple copies, separate locations, procedures reviewed at least six-monthly); anti-malware controls; encryption of data over public networks; event logging; access policy, authorisation and least privilege; password and authentication standards; network controls; a documented security-incident response process with notification without undue delay and in any event within 72 hours; and business continuity and redundant storage. The full text was reproduced in earlier versions of this DPIA and is available to controllers on request; the authoritative version is the Microsoft DPA itself.
Security certifications or accreditations (e.g. ISO 27001) in scope for this DPIA	Azure compliance offerings (ISO 27001 and others): https://learn.microsoft.com/en-us/compliance/regulatory/offering-home
Processor's breach notification processes	Per the Microsoft DPA: for each security breach that is a Security Incident, notification will be made without undue delay and, in any event, within 72 hours.
Processes and committed timelines to respond to a controller's request to satisfy an Individual Rights request	Per the Microsoft DPA (Data Subject Rights; Assistance with Requests): Microsoft makes available the ability to fulfil data subject requests, redirects any request it receives to the customer, and complies with reasonable requests for assistance.
Assurance programme over sub-processors and most recent outcomes	Per the Microsoft DPA: sub-processors are engaged under written contracts requiring at least the level of protection of the DPA; Microsoft publishes its sub-processor list, gives at least 6 months' notice of new sub-processors with access to Customer Data, and customers may terminate if they do not approve a new sub-processor.
How the controller's retention schedules will be met	Trellis controls retention within Azure (see Step 3). On subscription termination, per the Microsoft DPA, customer data is retained in a limited-function account for 90 days for extraction and then deleted within a further 90 days.

Step 3: Assess necessity and proportionality

What is your lawful basis for processing?	As the data processor, Trellis does not determine the lawful basis for processing. This is the responsibility of the data controller.
Does the processing actually achieve your purpose?	Yes. Processing personal data is necessary to achieve the core objective: reducing teachers' administrative workload by generating high-quality draft plans. These plans are based on detailed, individualised information: meeting recordings and pastoral notes. Without processing this data, the AI could not generate meaningful, usable drafts.
Is there another way to achieve the same outcome?	Alternatives considered: manual-only drafting (the current standard: time-consuming and a driver of workload stress) and template-based systems without AI (faster for some tasks, but unable to interpret unstructured discussion or personalise output). Neither reduces workload effectively at scale across diverse ASN/SEND contexts.
How will you prevent function creep?	Purpose limitation controls: personal data is processed only for the purposes authorised by the controller, including generating the requested documents, operating and securing the service, and undertaking agreed quality assurance and optimisation activities. Data is not used for advertising, profiling or unauthorised secondary purposes. Contractual safeguards: Trellis operates strictly under a Data Processing Agreement which limits the scope and purpose of processing. New document types are enabled per authority, by the authority. Any significant change in purpose triggers a DPIA review and requires controller sign-off (see Step 4, Ongoing Oversight). All system changes are version-controlled and documented, and the platform's audit trail records every action against a named actor.
How are you ensuring that you are not collecting more data than you need?	Trellis follows a data minimisation approach in line with Article 5(1)(c) UK GDPR. Input fields are limited to the content required to generate the requested documents. By default Trellis retains meeting/pupil data for 60 days; this is configurable to match the controller's retention policy. Trellis does collect operational telemetry and audit data including activity logs (actor, action, entity, IP address, user agent, trace identifiers), error logs and diagnostic events, because these are necessary for security monitoring, accountability and support. They are access-controlled, are retained for seven years as compliance evidence, and are not used for any other purpose. The diagnostics beacon captures client telemetry only; it carries no meeting content. Activity logs identify users by an internal identifier and record IP addresses and user agents. Error logs may contain users' names and email addresses where these appear in error messages. These logs contain staff personal data but no pupil data or meeting content.
How will you ensure data quality?	Trellis does not determine the content or accuracy of data provided by professionals, but ensures data: is not altered during submission or processing; is preserved in its original structure during prompt construction; and is transferred securely (TLS everywhere, plus client-side encryption of audio) to avoid corruption or interception. Generated output is validated against JSON schemas before rendering. Teachers review and edit all AI-generated content before it is finalised or shared, so errors or misinterpretations can be corrected. In-app editing makes this correction step easier and measurable.
What information will you give individuals?	As processor, Trellis does not issue privacy notices directly to data subjects; that is the controller's responsibility. Trellis supports controllers with technical descriptions of the processing, the in-product transparency statement at the start of recording, and assistance with transparency documentation. Work is under way on an accessible explainer for children and families, developed with stakeholder input.

How will you enable a data subject Access right?	If Trellis receives a DSAR or any other rights request, it will promptly notify the controller and await instructions, in line with Article 28(3)(e) UK GDPR. Trellis will assist the controller in identifying and retrieving relevant records within the retention period. All processing activity is logged to support auditability.
How will you enable a data subject Deletion right?	Trellis does not act independently on deletion requests, but actions controller instructions promptly. Records carry structured metadata allowing targeted deletion; deletion is performed by the retention scrub, lifecycle policies or scripted processes as instructed, and every deletion is logged. Data within the system is already subject to automated deletion at the end of the 60 and 14 day retention periods as described above, but Local Authorities can instruct Trellis to override retention and delete data sooner. A signed deletion report is produced.
How will you enable a data subject Rectification right?	Trellis will not modify personal data unless instructed by the controller. Rectification instructions can be fulfilled by direct update of the stored record.
If you are a processor, how will you manage any specific retention periods that the controller may have?	Retention is a property of each record: the authority's retention period, set at onboarding, computes a PII expiry on every record. The continuous scrub process deletes expired audio and documents and overwrites identifying fields, and logs each action. The application additionally filters expired data from responses. Retention settings are agreed with the controller during onboarding and recorded in the DPA. System audit logs sit outside this cycle and are retained for seven years, as recorded in Annex 1 of the DPA.
If you are a processor, will this data be deleted at the end of the contract with the controller and how?	Yes. On termination, and in line with the DPA, Trellis will permanently delete all personal data from live systems, other than the system audit logs retained for seven years under Annex 1 of the DPA.
If you are a processor, will this data be in any churn export at the end of the contract?	No. Trellis does not produce churn reports that include personal data. No personal data is included in Trellis's internal metrics or post-contract reviews.
What measures do you take to ensure processors comply?	Trellis only engages sub-processors with: a Data Processing Agreement in place; verified technical and organisational measures (encryption, access control, ISO 27001); appropriate data residency for Trellis content in the UK (including speech transcription and LLM inference), with staff identity and authentication data held by Microsoft Entra in the EU under UK EEA adequacy; and rights to audit or receive audit documentation. Material changes to the sub-processor list are notified to controllers in advance.
Have you discussed the project with legal/DPO in relation to any international transfers?	Yes. There are no restricted international transfers. Trellis content (including transcription and LLM inference) is stored and processed in the UK. Staff identity and authentication data is held by Microsoft Entra in the EU; under UK adequacy for the EEA this is not a restricted transfer and needs no additional safeguard.

Step 4: AI system risk assessment

This section is required because the project involves an AI system. It assesses the risks specific to AI, including impacts on data subjects, transparency requirements and fairness concerns.

Have the roles and responsibilities of personnel involved in the design, development, deployment, assessment, and monitoring of the AI system been defined and documented?	Yes, the team is as follows: • Andy Bell – CEO and Product Manager • Jasmine Gabbitas – Operations Manager • Thomas Pomfret – CTO • Harriet Owen – Lead AI Engineer • Dani Loader – Software Developer

Will the AI system team provide training for personnel responsible for operating and maintaining the AI system to ensure compliance with data protection principles?	Yes. Trellis provides onboarding training as part of each authority's roll-out, including seminars, recorded video instructions and written guidance. Training emphasises the teacher's responsibility for reviewing and finalising every document, and the known limitations of AI-drafted content (including speaker mis-attribution risk).
Which best describes the AI model: Human-in-the-Loop, Human-on-the-Loop, or Human-in-Command?	Trellis is best described as a **Human-in-Command** system. Trellis creates a draft and the teacher (or other education professional) reviews it before use. The professional is always responsible for reviewing the draft, making any edits, and deciding whether and how the content is used. At no point does the AI make decisions or take autonomous action affecting an individual. Final accountability remains fully with the human user. The AI is a support tool, not a decision-maker.
Do human reviewers have the expertise and authority to override AI decisions and modify outcomes as needed?	Yes. The person using Trellis is the person who would otherwise have written the Plan. Someone with sufficient expertise to write a support plan has the expertise to review an AI-drafted one. We continue to test this assumption through user research and school visits.
Are mechanisms in place to identify and assess the impacts of the AI system on individuals, communities, and society?	Yes, and these are now operating rather than planned: • Machine-learning evaluation reports are produced on a two-monthly cycle, combining automated metrics with manual thematic review. • In-app editing gives inferred feedback: section-level edit rates and edit-type classification show where drafts fall short. • User surveys run every six months. • School visits in each participating authority follow an agreed evaluation format. These mechanisms are designed to surface negative or unintended impacts early and feed directly into prompt and model improvements.
Are potential negative impacts reassessed if there are significant changes to the AI system at any stage of the AI lifecycle?	Yes. Identified potential negative impacts include: inaccurate or misleading draft content; over-reliance on the tool; privacy and data protection risks inherent in processing sensitive data, heightened in the event of a breach; and the risk of bias in how the AI interprets input or generates content, even with human review. Mitigations are built into the system • Human-in-Command review • client-side encryption • role-based access • MFA, (enforced in the Trellis tenant; the controller's policy for federated sign-in, Row-Level Security) • and into process (evaluation reports, expert review, this DPIA's review cycle). Significant system changes trigger reassessment, as recorded in this document's version history.
Are tolerable risk levels defined for the AI system based on business objectives, regulatory requirements, and data sensitivity?	Yes, in line with the sensitivity of pupil data, legal requirements, and the objective of reducing teacher workload without compromising child safety or accuracy: • **Accuracy and quality:** only rare or minor inaccuracies are acceptable, as all drafts remain under Human-in-Command review. Any misrepresentation that could materially affect a child's support needs is intolerable. • **Fairness and bias:** outputs must not contain systematic or discriminatory bias. Minor stylistic variation is tolerable; anything that disadvantages a subgroup of pupils is unacceptable.

	• **Confidentiality and security:** zero tolerance for unmitigated risks. Data is encrypted in transit and at rest (audio additionally encrypted client-side); access is restricted by MFA and role-based controls. Any unauthorised disclosure would be intolerable. • **Availability:** medium tolerance. Short outages are acceptable — teachers can revert to manual drafting. Prolonged outages that risk delaying statutory support are not tolerable. • **Transparency:** parents and guardians must not be misled about recording or AI use. Failure to provide adequate notice or opportunity to object is not acceptable. These tolerances are monitored through quantitative quality thresholds in the two-monthly ML evaluation reports and structured qualitative feedback through the Teacher Strategy Group.
Are there processes to address potential malicious use, misuse, or inappropriate use of the AI system?	The authority is expected to be involved in creating acceptable-use statements or policy for Trellis. The narrow use case limits the scope for misuse; in addition, every uploaded file is malware-scanned before processing, the WAF applies UK geo-restriction and additional security controls, with registered-IP exceptions for authorised access from outside the UK , and high-severity content-filter results alert a Trellis operator.
Are red-team exercises or other adversarial testing methods used to evaluate the system under stress and identify failure modes?	An accredited third-party security firm penetration-tests the service annually, with findings tracked to closure and re-validated at the following test. We have not commissioned adversarial red-teaming of the AI pipeline specifically, as Trellis is not a decision-making system.

Transparency, fairness, and accountability

Are there mechanisms for data subjects to challenge AI-generated outcomes?	Trellis does not make decisions or recommendations. The final document is always reviewed and submitted by a human, and data subjects can challenge its content through the controller's normal channels.
Are explanations provided to relevant users and stakeholders about how AI decisions are made?	The system does not make decisions or recommendations. We are developing an accessible explainer, with stakeholder and learner input, so children and families can understand how Trellis works.
Are procedures in place to review the fairness and non-discriminatory nature of AI outputs?	Yes. The AI assists with drafting only; it does not assess pupils, recommend support levels, or make decisions about entitlements. Expert teachers review outputs for accuracy and fairness. ML evaluation reports include review of edits teachers make to drafts.
How will you monitor for unintended bias or discrimination over time?	The two-monthly ML evaluation cycle reviews how the system handles a diverse range of inputs, using automated metrics continuously and manual thematic review periodically. Concerns around fairness or representativeness are addressed before changes are promoted. The risk is structurally limited: the system extracts information from meetings into predefined templates, always reviewed by a qualified professional.
Have applicable legal regulations for privacy been identified and considered before processing personal data to train, develop, or deploy the AI system?	We have taken the advice of qualified data protection experts. Relevant frameworks: the Data Protection Act 2018 and UK GDPR, governing lawful basis, security, minimisation, individuals' rights and special category data. We also aim to abide by the UK government's five AI principles: Safety, security and robustness; Mechanisms for contestability and redress; Accountability and governance; Reasonableness and fairness; Transparency and explainability. To note: no personal data is used to train, fine-tune or develop AI models.
Have measures to achieve privacy by design and default been implemented to mitigate identified privacy risks?	Yes. Privacy by design is implemented structurally: client-side encryption of audio; a single trust-boundary layer for all data access with automatic audit logging; per-user Row-Level Security in the database that fails closed; least-privilege database roles per service; PII expiry on every record with automated scrubbing; and three fully separated environments (production, preview, development) with no developer access to production data. TKM

	Consulting provides expert guidance on privacy by design throughout build and deployment.

Security and incident response for AI systems

Are there specific security controls for AI training data and model parameters?	Not applicable in the training sense: Trellis does not train models or hold model parameters. Prompts and templates (which contain no personal data) are version-controlled and governed through the admin panel.
How will you detect and respond to data poisoning or adversarial attacks against the AI system?	Uploaded files are malware-scanned before any processing; input is confined to the teacher's own recording and documents; the WAF (OWASP Core Rule Set, prevention mode) protects the application surface; and annual penetration tests probe for weaknesses. Content-filter events are alerted to an operator channel.
Are mechanisms in place to assess the security and resilience of the AI system across its lifecycle?	Yes: third-party penetration testing with tracked remediation; dependency vulnerability scans on every change and nightly against deployed images (high-severity findings block release); infrastructure-as-code review before changes apply; comprehensive logging and monitoring; and ISO 27001 management review and audit cycles.
How will you manage the security of AI models if they are shared or reused in different contexts?	Not applicable. Trellis uses managed Azure AI services; no models are shared or reused outside the service.
Are there processes in place to respond to incidents related to AI systems?	Yes. Incident response procedures cover documenting the incident, its impact, and resolution steps, with controller notification obligations set out in the DPA. The audit trail supports investigation, and failure events are recorded even where the originating transaction rolls back.
Are processes in place to modify system security and countermeasures in response to testing and real-world events?	Yes: penetration-test findings are tracked to closure on a register and re-validated; incident-driven improvements, continuous monitoring and alignment with recognised practice strengthen security over time.

Ongoing oversight

Will the AI system be subject to periodic impact assessments to ensure ongoing compliance with data protection laws?	Yes. This DPIA is updated as the system changes (see version history) and is reviewed with data protection stakeholders each roll-out phase. An EQIA and CRWIA are being updated for sharing with local authorities.
How will you monitor the long-term effects of the AI system on data subjects?	Primarily through user feedback, plus six-monthly surveys of all users and school visits in each authority.
Are there dedicated roles or teams responsible for the ongoing oversight of the AI system?	Yes: the core team (CTO, Lead AI Engineer, Product Manager) supported by expert teachers, with data protection oversight from our external advisors.
Will there be regular reviews of the AI system's impact on fairness, transparency, and accountability?	Yes, through the two-monthly ML evaluation reports and the Teacher Strategy Group.
How will you ensure that the AI system is not used beyond its original purpose without proper re-assessment?	Any significant change in use requires an updated DPIA process to assess risks and define new tolerances, with controller sign-off. Document types are enabled per authority, by the authority.

Are measures in place to address the traceability of the AI system throughout its entire lifecycle?	Yes: versioned prompts and templates; version-pinned schema and libraries; audit and error logs with trace correlation; and lifecycle documentation, although the system is not classified as high risk.

Step 5: Consultation process

What was the feedback from data subjects?	The Controller is responsible for their own consultation with staff and those with parental responsibility if necessary. Employees of controllers are directly involved in the continuous evaluation process: they review AI-generated drafts, provide structured feedback, and their in-app edits inform quality benchmarks. An accessible explainer for children and families is in development, with feedback from at least eight stakeholders including learners planned during 2026.
Who else has been involved, internal or external, and what was their feedback?	Trellis has consulted a number of experts in developing this architecture, notably: • Michael Cochrane – Senior Tech Architect, Education Scotland • Murray Elliot – Security Lead, Education Scotland • Frances Foreman – Senior Education Officer, Education Scotland • Thomas Pomfret – CTO, Trellis • Angela Bradbury – Project Manager, Trellis • Silva Arapi – Data Protection Consultant, Whitaker Consultants • Frank Rankin – Data Protection Consultant, TKM Consulting
What was the feedback from industry experts (e.g. security experts)?	The use of private networking, encryption at rest and in transit, and limited public exposure were validated as aligned with good industry practice. The v3.1 architecture (WAF, private endpoints, managed identities, Row-Level Security) has been reviewed by the Scottish Government system architect as part of the CivTech PCA programme.
What is the feedback from the DPO?	Our data protection advisor has reviewed both this DPIA and the technical architecture. Feedback highlighted: • The importance of clear purpose limitation and avoidance of function creep. • Ensuring that no personal data is used to train AI models. • Making sure temporary storage is properly scoped and time-bound. • Distinguishing children's data from user data throughout the documentation (added in v2.2).

Step 6: Identify and assess risks

Risk scoring. Risk = Likelihood (1–5: Rare, Unlikely, Possible, Likely, Almost certain) × Severity (1–5: Negligible, Minor, Moderate, Major, Catastrophic).

Scores of 1–3 are Low; 4–6 are Medium; 7+ are High. High risks require further mitigation. Where additional controls are identified, the risk is re-estimated. Risks are assessed at current deployment scale (seven authorities), not the original 100-child pilot.

Risk register

What could happen, and effect on data subject	Category	How could it happen?	Existing risk controls	L	S	Risk	Accept? Y/N	New risk controls	Residual (L×S)	Accept? Y/N
The AI generates an inaccurate or poorly worded draft, resulting in misrepresentation of a child's needs. This could lead to inappropriate or delayed support.	Accuracy / Fairness	The AI summarises transcript content inaccurately or omits key context. Teacher may accept output without sufficient review.	Human-in-Command design: a qualified professional always reviews and edits drafts before finalisation. In-app editing lowers the barrier to correction. Output validated against schemas. Section-level edit rates monitored in two-monthly ML evaluation reports. Onboarding training emphasises teacher responsibility for final output.	2	4	8		Continue reviewing sample outputs with expert users; monitor edit-rate metrics for drift after any prompt or model change; authority onboarding materials reinforce review responsibility.	1 × 4 = 4	
A parent objects to the meeting being recorded, feeling they were not properly informed of the purpose of the recording.	Transparency	A child or parent is not informed that the meeting is recorded and processed using AI, and objects after the fact.	Transparency statement flow at the start of recording; controllers can configure the system to pause or cancel recording if an objection is raised; accessible explainer in development for children and families.	1	2	2		—	—	
Service outage delays creation of plans for vulnerable children.	Availability	Transcription or generation temporarily unavailable due to Azure service outages, software bugs or integration failures. The production database and broker are currently single-zone,	Teachers remain in full control of documentation and can revert to the manual drafting process. Zone-redundant gateway and container environment; Azure SLAs; job queue retries and recovery.	2	3	6		High-availability and zone-redundancy hardening of the database and broker is on the delivery roadmap (Tech Spec v3.1, Section 12).	2 × 2 = 4	

What could happen, and effect on data subject	Category	How could it happen?	Existing risk controls	L	S	Risk	Accept? Y/N	New risk controls	Residual (L×S)	Accept? Y/N
		which gates effective availability.								
Unauthorised party gains access to child data via compromised login.	Confidentiality	Misconfigured role-based access; credentials reused or leaked; access from an unmanaged network.	Entra ID sign-in with MFA (enforced in the Trellis tenant; controller policy for federated sign-in); WAF with UK geo-restriction, traffic filtering and additional security checks; per-user Row-Level Security failing closed; least-privilege database roles; audit logging of every access.	1	4	4		—	—	
Personal data kept longer than the controller allows, increasing breach exposure.	Storage limitation	Automated deletion mechanisms fail or are misconfigured; data persists beyond the agreed retention period.	60-day default retention, configurable per authority. PII expiry stamped on every record; continuous scrub deletes expired blobs and nullifies fields; API filters expired data from responses; every deletion logged; signed deletion reports produced.	1	3	3		Quarterly review of scrub operation; automated alert for records past expiry without deletion.	1 × 3 = 3	
Function creep: Trellis used beyond original scope without a new DPIA.	Purpose limitation	Product expansion, local configuration by schools or authorities, or integration with other systems without proper governance, DPIA update or controller approval.	Processing scope documented in DPA and DPIA; document types enabled per authority by the authority; no integrations without controller request and contract amendment.	3	3	9		Mandatory DPIA review triggered by any change of purpose; controller sign-off for scope changes.	2 × 3 = 6	
Speaker mis-attribution in a transcript discloses confidential remarks to the	Accuracy	Diarisation fails due to overlapping speech, background noise, or	Manual review of transcript and draft (Human-in-Command); no	3	4	12		Guidance to teachers on reviewing transcripts with care where multiple	1 × 4 = 4	

What could happen, and effect on data subject	Category	How could it happen?	Existing risk controls	L	S	Risk	Accept? Y/N	New risk controls	Residual (L×S)	Accept? Y/N
wrong family or misattributes a statement in the record.		speakers not being clearly distinguishable. A comment from one speaker is attributed to another and carried into a document.	automatic sending of documents to families; Azure Speech native diarisation.					speakers are present; example cases in training materials; diarisation quality tracked in ML evaluation reports.		
Bias or discriminatory wording appears in AI output, meaning certain categories of student get systematically different treatment.	Fairness / Accuracy	Model inherits societal bias; poorly designed prompts; unrepresentative evaluation samples.	Systematic evaluation against predefined metrics on a two-monthly cycle; Teacher Strategy Group review; diverse sample plans in evaluation sets; training makes teachers aware of LLM limitations; no fine-tuning.	2	4	8		Continue diverse-input evaluation each cycle; escalate any systematic pattern to prompt/model change before promotion.	1 × 4 = 4	
Sub-processor (e.g. Microsoft) breach exposes data.	Confidentiality / Integrity	Cloud hypervisor exploit or insider threat at vendor.	Microsoft Azure compliance offerings and certifications; client-side encryption of audio means storage-level access cannot read recordings; private endpoints; 72-hour breach notification in Microsoft DPA.	1	5	5		—	—	
Social-engineering attack leads to credential compromise on the controller side.	Confidentiality	Phishing email or phone scam posing as IT support.	MFA for all users (enforced in the Trellis tenant; the controller's policy for federated sign-in); least-privilege RBAC; WAF security controls and UK geo-restriction provide additional protection against unauthorised access; SSO federation (as it rolls out)	2	4	8		Trellis recommends controllers implement phishing-awareness training and endpoint protection policies if not already in place.	1 × 4 = 4	

What could happen, and effect on data subject	Category	How could it happen?	Existing risk controls	L	S	Risk	Accept? Y/N	New risk controls	Residual (L×S)	Accept? Y/N
			keeps credential hygiene in the controller's own directory.							
A Trellis staff member recognises someone in a recording during evaluation.	Confidentiality	For evaluation of outputs, a small number of Trellis staff access recordings and drafts.	Access limited to five or fewer named staff on a need-to-see basis; BPSS vetting as a minimum; MFA and audited access; staff geographically distant from data subjects; recusal policy - if staff realise they know an individual they stop reviewing that data.	1	2	2		—	—	
Staff (teacher/user) personal data in audit logs is over-retained or accessed inappropriately.	Storage limitation / Confidentiality	Audit logs (actor, IP, user agent, actions) are retained for seven years as compliance evidence and could be repurposed or exposed.	Audit data access-controlled and written through a dedicated role; used only for security, accountability and support; contains no meeting content; retention as compliance evidence documented in the DPA.	1	2	2		•	-	
Controller uses small-value data from the Management Panel to identify specific users.	Purpose limitation / Fairness	A manager at the controller (local authority or trust) accesses low-value data in the Management Panel and uses their own knowledge to identify specific users.	Access control to the Management Panel. The controller's own data protection policy and training. Trellis user documentation to make clear that Management Panel data should only be used for system management and metrics purposes – not for performance management of individuals.	3	2	6		• -	-	

Step 7: Sign off and record outcomes

Item	Name / position / date	Notes
DPIA signed off by:	Name: Andy Bell Position: CEO, Trellis Date: to be dated on signature	Integrate actions (new risk controls) back into the project plan, with date and responsibility for completion.
Residual risks accepted by:	To be completed on signature	All residual risks scored 6 or below after mitigation.
Consultation responses reviewed by:	Name: Angela Bradbury Position: Project Manager, Trellis Date: to be dated on signature	16 July 2026 local authority feedback incorporated; DPA revisions progressing in parallel with Frank Rankin (TKM).
This DPIA will be kept under review by:	Name: Andy Bell Position: CEO, Trellis	The DPO should also review ongoing compliance with the DPIA. Next scheduled review: alongside Data Protection Information Pack v1 (M15).